

# PRIVACY-PRESERVING PERSONALIZED FEDERATED PROMPT LEARNING FOR MULTIMODAL LARGE LANGUAGE MODELS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Multimodal Large Language Models (LLMs) are pivotal in revolutionizing customer support and operations by integrating multiple modalities such as text, images, and audio. Federated Prompt Learning (FPL) is a recently proposed approach that combines pre-trained multimodal LLMs such as vision-language models with federated learning to create personalized, privacy-preserving AI systems. However, balancing the competing goals of personalization, generalization, and privacy remains a significant challenge. Over-personalization can lead to overfitting, reducing generalizability, while stringent privacy measures, such as differential privacy, can hinder both personalization and generalization. In this paper, we propose a Differentially Private Federated Prompt Learning (DP-FPL) approach to tackle this challenge by leveraging a low-rank adaptation scheme to capture generalization while maintaining a residual term that preserves expressiveness for personalization. To ensure privacy, we introduce a novel method where we apply local differential privacy to the two low-rank components of the local prompt, and global differential privacy to the global prompt. Our approach mitigates the impact of privacy noise on the model performance while balancing the tradeoff between personalization and generalization. Extensive experiments demonstrate the effectiveness of our approach over other benchmarks.

## 1 INTRODUCTION

In recent years, there has been rapid advancement in multimodal large language models (LLMs) that integrate multiple modality information, including text, images, audio, and video, to enhance the comprehension and generation capabilities. Vision-Language Models (VLMs) such as CLIP (Radford et al., 2021) are a variant of multimodal LLMs that learn transferable image and text representations, making them highly effective in applications such as image captioning and visual search. One proposed setting for deploying VLMs is a Federated Learning framework that allows multiple organizations or clients to collaboratively train a global model without directly sharing their local training data. However, fine-tuning pre-trained VLMs in a FL system is time-consuming and resource-intensive given the massive number of parameters each VLM has. This gives rise to Federated Prompt Learning (FPL) which only fine-tunes the soft prompt embedding while freezing the rest of the VLM model parameters in the FL system (Guo et al., 2023b; Cui et al., 2024). In FPL, each client fine-tunes their customized prompt using their local data and shares the prompt with a central server for generalization purposes. The clients can distribute their fine-tuned prompts as prompt providers to public users who wish to perform downstream inference tasks, also known as the prompt as a service (PaaS) paradigm (Wu et al., 2024; Yao et al., 2024; Huang et al., 2023).

One significant challenge in such distributed systems is the presence of data heterogeneity, i.e., organizations often have non-identical and non-independent (non-IID) data distributions, which can vary widely due to factors such as demographics, usage patterns, or device capabilities. To address this, personalized FL has emerged to tailor models to the unique data characteristics of each client rather than solely improving a global model. Personalized FL focuses on learning customized models for each client, reflecting the heterogeneity of their data (He et al., 2020; Dinh et al., 2020). In the context of personalized FPL, the goal is for each client to learn and utilize personalized prompts that better align with their specific data and application needs. Nevertheless, over-personalization can

054 lead to local data overfitting, preventing the model to generalize well on non-training data. Clients  
 055 in an FPL framework may opt to distribute their customized prompts post training to public users for  
 056 downstream tasks. However, these users may have different types of inputs that the client’s prompt  
 057 is not well generalized to, resulting in suboptimal performance. Consequently, it is crucial to achieve  
 058 a nuanced balance between personalization and generalization in a heterogeneous FPL system.

059 In addition to balancing the tradeoff between personalization and generalization, privacy poses an-  
 060 other critical concern in FPL, especially in sensitive domains such as finance, law, and healthcare. In  
 061 the PaaS framework, the distributed trained prompts are shown to be susceptible to Membership In-  
 062 ference Attack (MIA), potentially exposing details about individual clients’ training data (Wu et al.,  
 063 2024). To address this issue, one may consider Differential Privacy (DP) (Dwork et al., 2014) which  
 064 ensures an adversary cannot reliably detect the presence or absence of a data sample based on the  
 065 output information. However, balancing personalization and privacy under data heterogeneity is a  
 066 challenging task. The non-IID nature of the data allows clients to better learn their personalized  
 067 prompt, but it amplifies the performance degradation caused by DP due to the high data sensitivity,  
 068 impairing both personalization and generalization capabilities. Thus, the key question we aim to  
 069 address is: *How can we effectively balance personalization, generalization, and privacy in a data*  
 070 *heterogeneous FPL system?*

071 To tackle the above question, we proposed a Differentially Private Federated Prompt Learning (DP-  
 072 FPL) approach that leverages Low-Rank Adaption (LoRA) scheme and DP as part of the prompt  
 073 learning process. In our framework, each client simultaneously learns a global prompt and a local  
 074 prompt. The global prompt is shared in a FL manner for generalized knowledge transfer, while the  
 075 local prompt is retained at each client site for personalization. Our contributions are threefold.

- 076 • We propose a privacy-preserving personalized federated prompt learning approach with  
 077 Differential Privacy for multimodal LLMs. We adopt LoRA to decompose the local prompt  
 078 into two lower rank components with an additional residual term. The factorized low-rank  
 079 components allow the model to capture broader patterns that are beneficial across different  
 080 data distributions, aiding the generalization capability of each client. The residual term  
 081 is crucial for retaining the expressiveness lost during the decomposition process, thereby  
 082 preserving the client-specific learning and improving personalization.
- 083 • We preserve privacy by utilizing both Global Differential Privacy (GDP) and Local Differ-  
 084 ential Privacy (LDP). **Unlike conventional methods that apply noise uniformly to the entire**  
 085 **prompt, we judiciously apply LDP to the two low-rank components of the local prompt,**  
 086 **and GDP to the global prompt.** Our privacy mechanism mitigates the effect of DP noise on  
 087 model performance while preserving the privacy guarantee post training.
- 088 • We conduct extensive experiments on widely adopted datasets to evaluate our proposed  
 089 method against other benchmarks. The experimental results demonstrate superior per-  
 090 formance of our proposed method in balancing personalization and generalization while  
 091 mitigating the model degradation caused by DP noise.

## 093 2 RELATED WORK

094 **Personalized Federated Learning.** There are several existing approaches that aim to learn per-  
 095 sonalized models for clients in FL settings, including clustering (Ghosh et al., 2020; Berlo et al.,  
 096 2020; Shahid et al., 2021), regularization (Shoham et al., 2019; Dinh et al., 2020; Li et al., 2020)  
 097 and knowledge distillation (Li & Wang, 2019; He et al., 2020; Fang & Ye, 2022). Personalized FL is  
 098 most commonly approached as a multi-task learning problem that simultaneously learns two models  
 099 for each client: a global model for generalized knowledge and a local model for personalized data.  
 100 Existing methods accomplish this by decoupling the model parameters or layers into global and lo-  
 101 cal learning components (Arivazhagan et al., 2019; Deng et al., 2020; Zhang et al., 2020; Collins  
 102 et al., 2021; Jeong & Hwang, 2022; Zhang et al., 2023). In the existing literature on personalized FL,  
 103 private multi-task learning approaches aim to protect training data by retaining personalized param-  
 104 eters and sharing differentially private generalized parameters. Examples include Jain et al. (2021),  
 105 Hu et al. (2021b), Bietti et al. (2022), Yang et al. (2023b) and Xu et al. (2024). However, these  
 106 methods are designed for full model training and cannot be directly applied to prompt tuning due to  
 107 the difference in the parameter space. Sun et al. (2024) proposed a similar approach to ours where

they incorporate LoRA with DP in a standard FL setting, but they do not consider personalization and prompt learning, making their method not applicable to our setting.

**Federated Prompt Learning.** Recent advances in personalized FPL have garnered significant attention (Guo et al., 2023a;b; Li et al., 2023; Yang et al., 2023a; Sun et al., 2023; Deng et al., 2024; Li et al., 2024; Cui et al., 2024). See Table 1 for comparisons. With the exception of Zhao et al. (2023) which introduces a privacy-preserving FPL method that leverages DP to protect the underlying private data, none of the prior literature considers the privacy issue. Many of these works require modification to the backbone model, which is not relevant to our approach as we want to protect the personalized prompt, not the model. Zhao et al. (2023) does not account for the crucial aspects of personalization. Similar to our work, Cui et al. (2024) also factorizes the local prompt into two learnable low-rank components for balancing personalization and generalization. We instead have the learnable full-rank local prompt, and only keep the low-rank terms non-permanent for generalization with an additional residual to retain the expressiveness for personalization.

Table 1: Recent Federated Prompt Learning algorithms

FPL Algorithm	Consider personalization	No model modification	Adopt low-rank factorization	Provide privacy guarantee
pFedPrompt (Guo et al., 2023a)	✓	✗	✗	✗
PromptFL(Guo et al., 2023b)	✓	✓	✗	✗
pFedPT (Li et al., 2023)	✓	✗	✗	✗
pFedPG (Yang et al., 2023a)	✓	✗	✗	✗
Fedperfix (Sun et al., 2023)	✓	✗	✗	✗
SGPT (Deng et al., 2024)	✓	✗	✗	✗
FedOTP (Li et al., 2024)	✓	✓	✗	✗
FedPGP (Cui et al., 2024)	✓	✓	✓	✗
Fedprompt (Zhao et al., 2023)	✗	✓	✗	✓
DP-FPL (ours)	✓	✓	✓	✓

### 3 PROPOSED METHOD

We introduce our proposed method, Differentially Private Federated Prompt Learning (DP-FPL), shown in Figure 1. Our approach leverages LoRA with an additional residual term to balance personalization and generalization in a differentially private FPL system.

#### 3.1 PRELIMINARIES ON PERSONALIZED FEDERATED PROMPT LEARNING

Our system follows a standard FPL setting that consists of a set of  $N$  clients and a central server. Let the global dataset be  $D = [D_1, D_2, \dots, D_N]$ , each client  $i$  holds a local subset  $D_i$  of  $n_i$  samples. Each client local model involves a frozen pre-trained VLM such as a CLIP model and a prompt learner, and their goal is to learn the representation between the visual and prompt information to improve multimodal classification tasks. Specifically, the frozen CLIP model involves a text encoder  $f(\cdot)$  and an image encoder  $g(\cdot)$  that respectively transform the prompt and an image  $x$  into text and image features. The prompt learner trains a soft prompt  $p_i$  for client  $i$  that is optimized to align with the visual features. Using  $\cos[\cdot, \cdot]$  to denote the cosine similarity used by CLIP model, the classification prediction probability for each client  $i$  is computed as:

$$\mathbf{p}(\hat{y} = j|x) = \frac{\exp(\cos[f(p_i, c_j), g(x)]/\tau)}{\sum_{k=1}^C \exp(\cos[f(p_i, c_k), g(x)]/\tau)}, \quad (1)$$

where  $\hat{y}$  denotes the predicted label,  $c_j$  denotes label  $j$  out of  $C$  number of classes, and  $\tau$  denotes the temperature parameter of CLIP. The client personalized prompt  $p_i$  is optimized with cross-entropy loss:

$$\mathcal{L} = -\frac{1}{|D_i|} \sum_{(x,y) \in D_i} \sum_{j=1}^C y \log \mathbf{p}(\hat{y} = j|x). \quad (2)$$

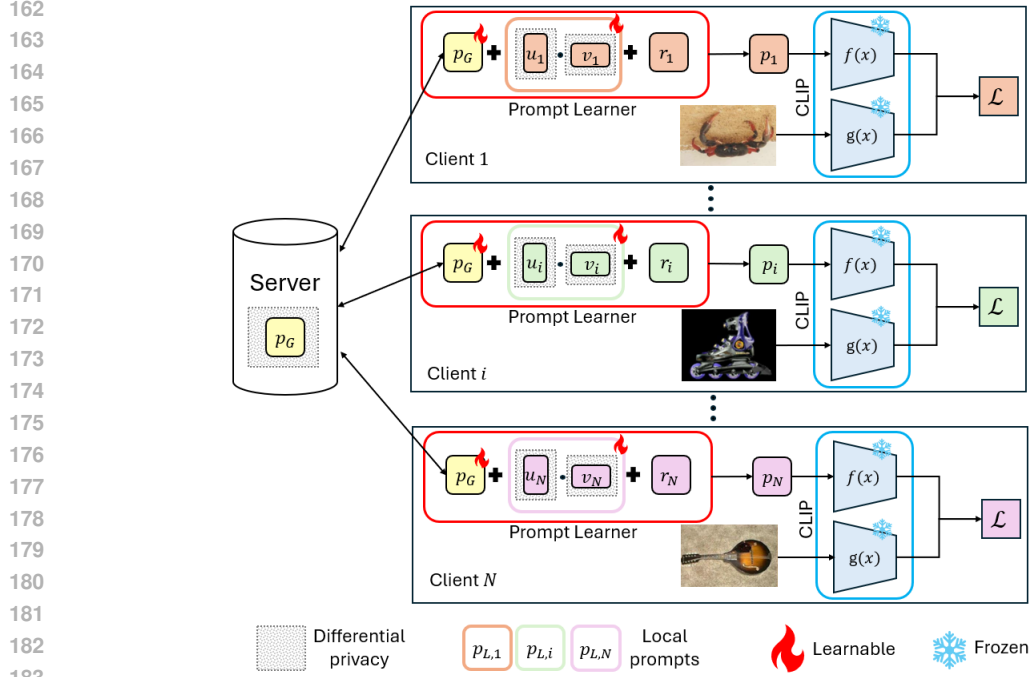


Figure 1: Architecture of DP-FPL with frozen CLIP models. Each client  $i$  trains global prompt  $p_{G,i}$  and local prompt  $p_{L,i}$ . The local prompt is factorized at each training iteration as  $p_{L,i} = u_i v_i + r_i$ .

In a data heterogeneity setting, each client’s local dataset  $D_i$  is drawn from a distinct data distributions. The difference among clients’ data can lead to the drift problem, where the local model converges toward local solutions optimal for their specific data but fails to align with the global model objective. To address this challenge, various personalized FPL solutions such as clustering, local fine-tuning and knowledge distillation have been proposed. Our work focuses on the multi-task learning approach that aims to learn two models for each client: one for generalization and one for personalization. In particular, we separate each client’s prompt  $p_i$  into a global prompt  $p_{G,i}$  and the local prompt  $p_{L,i}$ . The global prompt  $p_{G,i}$  is shared and aggregated to improve the global learning, while the local prompt  $p_{L,i}$  is retained at the client level for personalized learning.

An example of the personalized prompt is illustrated in Figure 1, in which each client has a collection of images captured from different angles, reflecting the heterogeneity in their local data. Consequently, client 1 might use the prompt “an upside-down photo of [class]”, while client  $i$  could have the prompt “a upright photo of [class]”, and client  $N$  might use “a rotated photo of [class]”. In this instance, the template “a photo of [class]” is the generalized global prompt shared across clients, and the terms “upside-down”, “upright” and “rotated” represent the personalized characteristics of each client’s prompt. These variations in prompts allow the client model to better adapt to the specific characteristics of their data, improving personalization.

The training process of FPL over  $T$  iterations is structured as follow. For each global training round  $t$ , each client  $i$  initializes the global prompt  $p_{G,i}^t \leftarrow p_{G,i}^{t-1}$  and the local prompt  $p_{L,i}^t \leftarrow p_{L,i}^{t-1}$ . Client  $i$  then trains their personalized prompt  $p_i$  using their local private data and obtains the cross-entropy loss  $\mathcal{L}$ . At the end of the local training round, client  $i$  updates their local prompt  $p_{L,i}^t \leftarrow p_{L,i}^t - \eta_{L,i} \nabla_{L,i} \mathcal{L}$  and sends the gradient w.r.t the global prompt  $\nabla_{G,i} \mathcal{L}$  to the server for aggregation. The server computes the average gradient  $\nabla_G \leftarrow \frac{1}{N} \sum_{i=1}^N \nabla_{G,i} \mathcal{L}$  and updates the new global prompt to be  $p_G^t \leftarrow p_G^{t-1} - \eta_G \nabla_G$ . The learning objective function of FPL system is:

$$\min_{p_G} \sum_{i=1}^N \frac{n_i}{\sum_{j=1}^N n_j} \mathcal{L}_{D_i}, \quad (3)$$

where  $\mathcal{L}_{D_i}$  is the loss computed on dataset  $D_i$  of client  $i$ .

### 3.2 BALANCING PERSONALIZATION AND GENERALIZATION

Prior research (Aghajanyan et al., 2020; Cui et al., 2024) has demonstrated that fine tuning pre-trained LLMs with lower dimension reparameterization promotes generalization capability across various tasks. Additionally, lower intrinsic dimension improves the model utility under the effect of DP noise (Yu et al., 2021; Xu et al., 2024). Inspired by this, we adopt Low-Rank Adaptation (LoRA) (Hu et al., 2021a) in our FPL framework to balance the personalization and generalization learning under the influence of DP noise.

LoRA has been used to fine-tune pre-trained models more efficiently by introducing rank decomposition into the weight matrices of neural networks. However, LoRA has difficulty matching the performance of full-rank training in many difficult tasks (Liu et al., 2024; Biderman et al., 2024; Ivson et al., 2023; Zhuo et al., 2024). This is because low-rank decomposition methods restrict the parameter space, throwing away some of the information of the full-rank space (Konečný, 2016). Recent work Cui et al. (2024) also used the low-rank approximation in non-private FPL systems, however, they factorize the local prompt only once at the beginning and train iteratively with the low-rank components. This approach can reduce the overall expressive power over the training phase. As we will show in Section 4, the loss of expressiveness can amplify the adverse effects of the DP noise, further diminishing the model’s performance. To overcome this issue, we perform the factorization process in every training round rather than only at the beginning like Cui et al. (2024), and we incorporate a residual term to compensate for the lost expressiveness. We analyze in detail the benefit of the residual term in Section 3.3. The overall personalized prompt of client  $i$  can then be expressed as  $p_i = p_{G,i} + u_i v_i + r_i$ , where  $u_i$  and  $v_i$  are the factorized low-rank components and  $r_i$  is the additional residual term.

We utilize a version of the Reparametrized Gradient Perturbation (RGP) method introduced in Yu et al. (2021) as our LoRA decomposition scheme. We perform the factorization process for each client local prompt  $p_{L,i}$  in every training iteration to get the temporary low-rank prompt components  $u_i$  and  $v_i$ . We also compute a residual term  $r_i$  which is the remainder of the factorization process, i.e.  $r_i = p_{L,i} - u_i v_i$ . As shown by Yu et al. (2021), each client can reconstruct the gradient of  $p_{L,i}$  using the gradient of the low-rank terms in back propagation as follows:

$$\nabla_{L,i} \mathcal{L} = (\nabla_u \mathcal{L}) v_i + u_i (\nabla_v \mathcal{L}) - u_i u_i^T (\nabla_u \mathcal{L}) v_i \quad (4)$$

We note that the residual  $r_i$  is used as part of the forward process, but it is not involved in the full-rank local prompt recomputation.

### 3.3 PRESERVING PRIVACY

In the personalized FPL framework, each client as a prompt provider may distribute the trained customized prompt to public users for downstream inference tasks. These fully trained prompts, if not properly protected with privacy mechanism, are vulnerable to MIA which aims to infer if an image sample was used for training or not (Wu et al., 2024). Moreover, the shared global prompt may leak information about the private data during the training process as it contains the gradient information of the loss function.

We assume that clients and server are honest. We consider adversary to be potential user with access to a trained customized prompt obtained from a FPL client. The adversary also has access to the publicly available pre-trained CLIP model for downstream inference purposes. The goal of the adversary is to infer whether an image data was part of the target FPL client’s training data utilizing the MIA. In this setting, the target client’s trained prompt  $p_i$  is shared with a potential adversarial user and is susceptible to privacy breach, so it is necessary to protect  $p_i$  with the DP mechanism.

One may straightforwardly apply privacy noise to the client’s trained  $p_i$  before distributing it to public users. However, this requires a large DP noise to effectively prevent MIA (Wu et al., 2024). On the other hand, injecting noise gradually over the training phase allows for more control over the influence of noise on the model, leading to better utility (Abadi et al., 2016). In our setting, we perform gradient updates on the global prompt  $p_{G,i}$  and local prompt  $p_{L,i}$  of each client, so we add DP noise to the gradients w.r.t these two terms in each training step.

As part of the FPL procedure, the global prompt of each client  $p_{G,i}$  is aggregated and averaged by the server, while the client’s local prompt  $p_{L,i}$  stays locally. The final trained prompt  $p_i$  is composed

of the synchronized global prompt  $p_G$  and the local prompt  $p_{L,i}$ . To provide privacy guarantee, we use Global Differential Privacy (GDP) for the global prompt and Local Differential Privacy (LDP) for the local prompt. This is an unconventional way of introducing DP noise to selective parts of the prompt, unlike a vanilla method that directly adds noise to the whole prompt before publishing it. We define these two DP notions as follow.

**Definition 3.1. Global Differential Privacy (GDP)** A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  with domain  $\mathcal{D}$  and range  $\mathcal{R}$  satisfies  $(\epsilon, \delta)$ -GDP if for any two adjacent datasets  $D, D' \in \mathcal{D}$  (i.e., datasets that differ in exactly one sample) and for any subset of outputs  $S \in \mathcal{R}$  it holds that

$$\Pr[\mathcal{M}(D) \in S] \leq e^\epsilon \Pr[\mathcal{M}(D') \in S] + \delta$$

**Definition 3.2. Local Differential Privacy (LDP)** A randomized mechanism  $\mathcal{M} : \mathcal{D} \rightarrow \mathcal{R}$  with domain  $\mathcal{D}$  and range  $\mathcal{R}$  satisfies  $(\epsilon, \delta)$ -LDP if for any two adjacent samples  $x, x' \in \mathcal{D}$  where  $D \in \mathcal{D}$  and for any subset of outputs  $S \in \mathcal{R}$  it holds that

$$\Pr[\mathcal{M}(x) \in S] \leq e^\epsilon \Pr[\mathcal{M}(x') \in S] + \delta$$

We apply GDP to the global prompt  $p_{G,i}$  at the server because the impact of the GDP noise on the model utility is much smaller compared to LDP (Arachchige et al., 2019). To provide privacy guarantee for the local prompt  $p_{L,i}$  which is not shared with the server for aggregation, each client needs to obfuscate  $p_{L,i}$  locally with LDP noise. However, directly applying LDP noise to the full-rank local prompt  $p_{L,i}$  can heavily impair the model performance (Yu et al., 2021; Xu et al., 2024). In addition, the local prompt information is predominantly captured within the low-rank components as a result of the factorization process. Therefore, we inject LDP noise only to the low-rank component  $u_i$  and  $v_i$  to mitigate the negative effect of privacy noise, while still effectively protecting the local prompt during the recomputation process. We note that we do not add noise to the residual component  $r_i$  because it is not used for the local prompt recomputation.

We achieve GDP and LDP using the Gaussian noise mechanism. Given a function  $f : \mathcal{D} \rightarrow \mathcal{R}$ , the Gaussian noise mechanism  $\mathcal{M}$  is defined as

$$\mathcal{M}(d) \triangleq f(d) + \mathcal{N}(0, \sigma^2) \quad (5)$$

where  $\mathcal{N}(0, \sigma^2)$  is the normal distribution with mean 0 and variance  $\sigma^2$ . The standard deviation  $\sigma$  is typically chosen based on the function  $f$ 's sensitivity  $S_f$  to satisfy an  $(\epsilon, \delta)$ -DP guarantee.

The two main building blocks in our approach, i.e., low-rank decomposition and DP, naturally introduce error to the training process. We conjecture that this error acts as a regularization term that prevents clients from overfitting to local data, reducing personalization and improving generalization. However, under strictly private conditions (lower rank and higher DP noise), the accumulated error may become too large and potentially destroy the personalization capability. In this case, the added residual term compensates for the regularization-like error and helps improve local learning, balancing personalization and generalization. We demonstrate the benefit of the residual term in the ablation study in Section 4.3, supporting our hypothesis. Further theoretical analysis of the residual term can be a potential future work direction.

### 3.4 ALGORITHM

We are now ready to describe our proposed method in detail, as shown in Algorithm 1.

At the initial stage, the server randomizes a starting global prompt  $p_G^0$  and each client  $i$  sets up their starting local prompt  $p_{L,i}^0$  (line 1). The variances  $\sigma_G$  and  $\sigma_L$  are chosen to satisfy a certain  $(\epsilon, \delta)$ -LDP and  $(\epsilon, \delta)$ -GDP guarantee. The algorithm runs for  $T$  iterations. In each iteration  $t$ , each client updates  $p_{G,i}^t$  to be the previously aggregated  $p_G^{t-1}$  and  $p_{L,i}^t$  to be the previous  $p_{L,i}^{t-1}$  (line 4). A minibatch  $\mathcal{B}^t$  is sampled from the local dataset  $D_i$  for training (line 5).

Each client performs parameter factorization using the power method with rank  $k$ , and then recomputes the local prompt  $p_{L,i} \leftarrow u_i v_i + r_i$  (lines 6 – 7). Lines 8 – 9 describe the forward pass where each client runs their local frozen CLIP model to get the text features and image feature, and uses them to calculate the loss  $\mathcal{L}$  using Equation 2. Each client then computes and clips the gradient w.r.t the two low-rank prompt components  $\nabla_u \mathcal{L}$  and  $\nabla_v \mathcal{L}$  with threshold  $C_{th}$  and add local DP noise with standard deviation  $\sigma_L$  according to 5 (lines 10 – 11).

**Algorithm 1** DP-FPL

---

```

1: Initialize:  $p_G^0, p_{L,i}^0$  for  $i = 1 \dots N$ , variances  $\sigma_L$  and  $\sigma_G$ 
2: for  $t \leftarrow 1 \dots T$  do
3:   for  $i \leftarrow 1 \dots N$  in parallel do
4:     Initialize  $p_{G,i}^t \leftarrow p_{G,i}^{t-1}$  and  $p_{L,i}^t \leftarrow p_{L,i}^{t-1}$ 
5:     Sample  $\mathcal{B}^t$  from  $D_i$ 
6:     Compute low-rank components  $u_i, v_i, r_i \leftarrow \mathbf{Factorize}(p_{L,i}, k)$ 
7:     Recompute  $p_{L,i} \leftarrow u_i v_i + r_i$  for forward process
8:     Obtain global text feature  $f(p_{G,i}^t)$ , local text feature  $f(p_{L,i}^t)$ , image feature  $g(x)$  ( $x \in \mathcal{B}^t$ )
9:     Calculate loss  $\mathcal{L}$  according to Equation 2 and compute gradients  $\nabla_u \mathcal{L}$  and  $\nabla_v \mathcal{L}$ 
10:    Clip gradients  $\nabla_{G,i} \mathcal{L}, \nabla_u \mathcal{L}$  and  $\nabla_v \mathcal{L}$  with threshold  $C_{th}$ 
11:    Add local DP noise:  $\tilde{\nabla}_u \mathcal{L} \leftarrow \nabla_u \mathcal{L} + \mathcal{N}(0, \sigma_L^2)$  and  $\tilde{\nabla}_v \mathcal{L} \leftarrow \nabla_v \mathcal{L} + \mathcal{N}(0, \sigma_L^2)$ 
12:    Recompute noisy gradient w.r.t  $\tilde{\nabla}_{L,i} \mathcal{L}$  using  $\tilde{\nabla}_u \mathcal{L}$  and  $\tilde{\nabla}_v \mathcal{L}$  according to equation 4
13:    Update local prompt  $p_{L,i}^t \leftarrow p_{L,i}^t - \eta_{L,i} \tilde{\nabla}_{L,i} \mathcal{L}$ 
14:    Send  $\nabla_{G,i} \mathcal{L}$  to the server
15:  end for
16:  Server computes average gradient  $\nabla_G \leftarrow \frac{1}{N} \sum_{i=1}^N \nabla_{G,i} \mathcal{L}$ 
17:  Server adds global DP noise to  $\tilde{\nabla}_G \leftarrow \nabla_G + \mathcal{N}(0, \sigma_G^2)$ 
18:  Server updates global prompt  $p_G^t \leftarrow p_G^{t-1} - \eta_G \tilde{\nabla}_G$ 
19: end for

```

---

The noisy gradient w.r.t the local prompt  $\tilde{\nabla}_{L,i} \mathcal{L}$  can be reconstructed from the noisy gradients  $\tilde{\nabla}_u \mathcal{L}$  and  $\tilde{\nabla}_v \mathcal{L}$  using equation 4, and then is updated accordingly (lines 12–13). Each client also computes the gradient w.r.t the global prompt  $\nabla_{G,i} \mathcal{L}$  and sends it to the server for aggregation (line 14). Upon receiving the locally computed gradients, the server computes the average gradient and adds global DP noise with standard deviation  $\sigma_G$  to perturb the gradient (lines 16 – 17). The server then updates the new global prompt for the next training round (line 18).

Low-rank decomposition via traditional SVD method requires significant runtime. Instead, we use the power method with one iteration, significantly reducing the computational cost. Given the full-rank matrix of size  $m \times n$  (assuming  $m \leq n$ ), the computational cost of SVD scales with  $\mathcal{O}(m^2 n)$ , while the power iteration only scales with  $\mathcal{O}(kmn)$  where  $k$  is the reduced rank and  $k \ll m$ .

### 3.5 PRIVACY ANALYSIS

DP has several properties and compositions that make it easier to analyze the privacy budget in repetitive algorithms such as machine learning, where the privacy loss accumulates across multiple training iterations. When DP mechanisms are applied repeatedly to the same dataset, the overall privacy budget accumulates sequentially using the advanced composition theorem (Dwork et al., 2014). Conversely, when DP mechanisms are applied independently to disjoint subsets of a dataset, the overall privacy loss does not accumulate. In this case, the privacy guarantee remains bounded by the maximum privacy loss of any subset using parallel composition (Dwork et al., 2014). The privacy budget in term of LDP and GDP of Algorithm 1 is given by the following theorem.

**Theorem 3.3.** *There exist constants  $c_1, c_2$  so that given the number of global rounds  $T$ , for any  $\delta > 0$ , DP-FPL satisfies  $(\epsilon, \delta)$ -LDP and  $(\epsilon, \delta)$ -GDP if we choose  $\sigma_L$  and  $\sigma_G$  as following:*

$$\sigma_L \geq c_1 \frac{S_L \sqrt{T \log(1/\delta)}}{\epsilon} \quad \sigma_G \geq c_2 \frac{S_G \sqrt{T \log(1/\delta)}}{\epsilon}$$

where  $S_L = \frac{C_{th}}{|\mathcal{B}|}$  and  $S_G = \frac{C_{th}}{N|\mathcal{B}|}$  are the local and global sensitivity respectively.

*Proof* By definition, a single application of the Gaussian noise mechanism satisfies  $(\epsilon, \delta)$ -DP if we choose  $\sigma \geq \frac{S \sqrt{2 \log(1.25/\delta)}}{\epsilon}$  where  $S$  is the sensitivity. Under the advanced composition theorem of DP, the Gaussian noise mechanism after  $T$  training steps results in an accumulated privacy loss of

( $\mathcal{O}(\epsilon\sqrt{T}), \delta$ )-DP. Thus, to achieve  $(\epsilon, \delta)$ -DP, one would need to choose

$$\sigma \geq c \frac{S\sqrt{T \log(1/\delta)}}{\epsilon} \quad (6)$$

for the Gaussian noise mechanism where  $c'$  is a constant.

Applying Equation 6, we can choose  $\sigma_L \geq c_1 \frac{S_L\sqrt{T_L \log(1/\delta)}}{\epsilon}$  where  $c_1$  is a constant to make Algorithm 1 satisfy  $(\epsilon, \delta)$ -LDP with respect to each client. Since each client  $i$  operates the Gaussian noise mechanism independently on disjoint local subset  $D_i$  of the global dataset  $D$ , the release of all clients' noisy mechanism output still satisfies  $(\epsilon, \delta)$ -LDP by the parallel composition of DP. Similarly, by choosing  $\sigma_G \geq c_2 \frac{S_G\sqrt{T \log(1/\delta)}}{\epsilon}$  according to Equation 6, DP-FPL satisfies  $(\epsilon, \delta)$ -GDP.

We proved in the theorem above that Algorithm 1 satisfies  $(\epsilon, \delta)$ -LDP and  $(\epsilon, \delta)$ -GDP when publishing the customized prompt to potential users. Since the GDP noise is added to the aggregated gradient, the distribution of the aggregated gradient to all clients also satisfies  $(\epsilon, \delta)$ -GDP by the post-processing property of DP. According to Theorem 3.3, we can calculate the standard deviations  $\sigma_L$  and  $\sigma_G$  to achieve a certain  $(\epsilon, \delta)$ -LDP and  $(\epsilon, \delta)$ -GDP. The pair  $\epsilon, \delta$  can be chosen to match a desired MIA accuracy rate, preferably lower than random guess (50%) (Thudi et al., 2022).

## 4 EXPERIMENTS

### 4.1 SETUP

**Datasets.** We select four visual classification datasets to investigate the task of balancing personalization, generalization and privacy: Caltech101 (Fei-Fei et al., 2004), OxfordPets (Parkhi et al., 2012), OxfordFlowers (Nilsback & Zisserman, 2008) and Food101 (Bossard et al., 2014). We utilize the pathological data split among 10 clients. Each client model is trained on its local classes, and evaluated on both its local classes for personalization capability and neighbor classes (classes owned by other clients) for generalization capability. We also evaluate the large-scale dataset CIFAR-100 (Krizhevsky et al., 2009) with the Dirichlet data split among 25 and 50 clients. The final test accuracy is obtained by averaging the performance across all clients. Details of our implementation and hyperparameters are provided in the appendix. Our code is available at <https://anonymous.4open.science/r/DP-FPL-84F8>.

Table 2: Mean test accuracy on local classes averaged across 10 clients. The baseline FedPGP and our method DP-FPL have factorization on the local prompt with rank 8.

Dataset	Noise $\epsilon$	PromptFL	FedOTP	FedPGP	DP-FPL
Caltech 101	None	94.45±0.30	<b>97.06±0.48</b>	95.21±0.19	96.06±0.18
	0.4	82.53±0.52	95.09±0.50	95.12±0.46	<b>95.74±0.63</b>
	0.2	81.52±0.45	87.61±0.58	90.98±0.40	<b>95.28±0.46</b>
	0.1	80.42±0.65	84.98±0.39	80.01±0.39	<b>92.71±0.38</b>
	<b>0.05</b>	78.61±1.39	83.61±0.38	77.24±0.38	<b>87.64±0.79</b>
	<b>0.01</b>	78.52±1.68	78.86±0.38	77.23±0.37	<b>85.21±0.85</b>
Oxford Pets	None	76.85±0.96	<b>99.63±0.2</b>	94.66±0.31	96.91±0.76
	0.4	74.36±0.26	80.03±0.93	86.56±0.69	<b>95.13±0.52</b>
	0.2	73.56±0.16	65.97±0.89	67.11±0.44	<b>93.09±0.43</b>
	0.1	72.77±0.16	59.54±0.58	63.21±0.62	<b>85.25±0.18</b>
	<b>0.05</b>	52.39±0.53	58.97±1.02	57.98±0.97	<b>81.26±1.10</b>
	<b>0.01</b>	43.68±0.67	54.08±1.04	45.49±1.33	<b>73.71±0.40</b>
Oxford Flowers	None	84.04±0.32	<b>97.84±1.16</b>	79.11±0.45	85.75±0.62
	0.4	60.31±1.28	79.89±0.80	77.13±0.52	<b>80.09±1.41</b>
	0.2	40.33±0.83	65.96±0.96	70.77±0.61	<b>76.75±1.05</b>
	0.1	38.25±1.37	42.31±0.71	52.42±1.58	<b>72.11±1.37</b>
	<b>0.05</b>	37.18±0.92	38.89±0.66	39.52±0.77	<b>69.80±1.34</b>
	<b>0.01</b>	36.11±0.60	33.98±0.63	35.23±0.64	<b>51.55±1.07</b>
Food 101	None	86.50±0.26	<b>86.65±0.23</b>	84.40±0.09	86.08±0.12
	0.4	78.70±0.39	79.45±0.23	80.58±1.52	<b>81.45±0.21</b>
	0.2	71.84±0.91	77.36±0.41	77.72±1.50	<b>81.25±0.18</b>
	0.1	69.00±0.40	70.48±1.39	75.18±0.22	<b>80.57±0.46</b>
	<b>0.05</b>	68.36±1.23	62.98±1.25	73.72±1.04	<b>78.23±0.43</b>
	<b>0.01</b>	67.47±1.20	54.70±1.12	71.82±1.19	<b>77.45±0.40</b>



**Baselines.** To demonstrate the effectiveness of our proposed method in balancing personalization, generalization and privacy, we compare with three baseline cases: (1) PromptFL (Guo et al., 2023b), (2) FedOTP (Li et al., 2024) and (3) FedPGP (Cui et al., 2024). More details of the baselines are listed in the appendix.

**Privacy levels.** We consider different noise levels for LDP and GDP:  $\epsilon = \{0.01, 0.05, 0.1, 0.2, 0.4\}$ . We pick  $\delta = 10^{-5}$  and the clipping threshold  $C_{th} = 10$ . We provide details of how the noise is added to each baseline in the appendix.

## 4.2 PERFORMANCE RESULTS

**Improving personalization in private setting.** Table 2 shows the average test accuracy over the last 10 epochs on local classes. In the non-private scenario (i.e., Noise = None), FedOTP shows the highest utility, however, as we add more noise, DP-FPL comes in the first place with the highest local classes accuracy. Under strict privacy levels ( $\epsilon = 0.01$ ), we see a noticeable decrease in test accuracy. Nevertheless, DP-FPL still consistently outperforms other baselines across all datasets, showing the robustness of our method even under strictly private conditions. Table 4 shows the overall accuracy utilizing a Dirichlet data split and ResNet50 as the backbone model. As shown in Table 4, DP-FPL demonstrates superior performance compared to baseline methods across all datasets. This shows the effectiveness and applicability of our approach in various complex settings, including different data distributions, models, and number of clients.

**Privacy improves generalization.** Table 3 shows the average test accuracy over the last 10 epochs on neighbor classes. Similar to local classes, DP-FPL exhibits the highest utility in neighbor classes under the presence of DP noise. As expected, the accuracy degrade for higher noise levels, however, we see an improvement in Caltech101 utility when we increase privacy level from 0.4 to 0.1. This is because privacy noise act as a form of regularization that prevents overfitting, and hence improving generalization. Nevertheless, when the noise level is large enough ( $\epsilon = \{0.01, 0.05\}$ ), the overall utility is degraded and the neighbor accuracy no longer improves.

Table 3: Mean test accuracy on neighbor classes averaged across 10 clients. The baseline FedPGP and our method DP-FPL have factorization on the local prompt with rank 8.

Dataset	Noise $\epsilon$	PromptFL	FedOTP	FedPGP	DP-FPL
Caltech 101	None	92.88±0.19	74.91±0.30	<b>93.44±0.75</b>	91.54±0.15
	0.4	82.66±1.22	84.26±0.84	88.24±0.78	<b>88.58±0.38</b>
	0.2	81.93±1.47	84.08±0.98	86.05±0.79	<b>89.72±0.98</b>
	0.1	80.83±0.72	78.91±1.02	79.31±0.82	<b>90.02±1.47</b>
	0.05	79.96±0.40	73.37±1.03	75.62±0.82	<b>82.76±1.12</b>
	0.01	78.89±0.33	73.54±1.00	75.60±0.82	<b>80.60±0.57</b>
Oxford Pets	None	76.34±0.37	65.63±0.14	<b>89.71±0.49</b>	80.19±0.85
	0.4	74.43±0.87	60.95±0.81	72.54±0.54	<b>81.82±0.42</b>
	0.2	73.74±0.82	59.74±0.84	61.68±0.46	<b>80.67±0.28</b>
	0.1	73.12±0.39	58.83±0.93	59.79±0.79	<b>77.12±0.52</b>
	0.05	52.44±0.89	54.08±0.89	51.63±1.01	<b>74.13±0.48</b>
	0.01	38.27±0.86	53.63±0.91	40.20±0.42	<b>71.89±0.48</b>
Oxford Flowers	None	69.44±0.61	38.29±1.09	<b>75.79±0.87</b>	69.51±0.45
	0.4	48.03±0.69	56.65±0.77	65.93±0.86	<b>67.67±0.45</b>
	0.2	38.19±0.97	55.44±0.90	63.49±0.65	<b>67.51±0.58</b>
	0.1	37.76±1.23	37.53±1.09	46.24±1.86	<b>66.44±1.51</b>
	0.05	38.78±1.18	33.48±1.00	35.27±1.28	<b>56.75±1.85</b>
	0.01	34.81±1.19	31.26±0.96	34.72±1.60	<b>43.21±1.72</b>
Food 101	None	86.19±0.13	84.03±0.33	<b>86.23±0.06</b>	86.08±0.11
	0.4	76.88±0.23	80.11±0.47	78.94±1.07	<b>81.00±0.25</b>
	0.2	70.99±0.83	76.44±0.62	77.21±0.90	<b>80.79±0.22</b>
	0.1	67.80±1.59	73.12±1.65	76.92±0.83	<b>78.14±0.53</b>
	0.05	66.76±1.27	71.82±0.79	73.61±1.35	<b>77.18±0.50</b>
	0.01	61.42±1.49	67.08±1.11	72.99±1.53	<b>76.87±0.62</b>

## 4.3 ABLATION STUDY

In this subsection, we investigate the efficacy of the key parameters that directly affect the tradeoff between personalization, generalization, and privacy: residual term, noise level  $\epsilon$  and rank value. The results for Caltech101 are presented in Figure 2; we include other dataset results in the appendix.

Table 4: Mean test accuracy averaged across all clients under Dirichlet data distribution. The baseline FedPGP and our method DP-FPL have factorization on the local prompt with rank 8.

Noise $\epsilon$	CIFAR-100 with 25 clients				CIFAR-100 with 50 clients			
	PromptFL	FedOTP	FedPGP	DP-FPL	PromptFL	FedOTP	FedPGP	DP-FPL
None	71.20±0.18	68.13±0.33	<b>71.54±0.17</b>	69.84±0.19	<b>71.30±0.10</b>	68.57±0.14	70.82±0.12	69.48±0.11
0.4	53.69±0.53	47.22±0.97	58.87±0.35	<b>66.23±0.23</b>	55.44±0.31	58.23±0.42	58.43±0.51	<b>66.40±0.18</b>
0.2	53.02±0.39	47.02±1.13	56.75±0.31	<b>62.92±0.30</b>	54.65±0.29	57.09±0.40	56.34±0.44	<b>64.39±0.29</b>
0.1	50.86±0.68	46.95±1.06	55.76±0.33	<b>59.53±0.25</b>	53.17±1.06	57.05±0.22	53.39±0.48	<b>60.49±0.26</b>
0.05	50.20±0.40	44.71±1.16	53.80±0.30	<b>57.97±0.16</b>	53.23±0.83	55.00±0.37	52.38±0.49	<b>58.35±0.30</b>
0.01	50.65±0.41	44.27±1.17	51.75±0.34	<b>53.31±0.35</b>	51.92±1.56	54.53±0.29	52.32±0.46	<b>56.09±0.56</b>

**Effect of residual term.** We investigate the effectiveness of the residual term by separately testing the model without the residual component in Figure 2. We note that this setting is different from FedPGP because the factorization process is performed every training round instead of at the beginning. In Figure 2, we observe better performance in both local and neighbor classes when incorporating the residual term. The difference in accuracy is more prominent under strict privacy conditions (rank 1 and  $\epsilon = 0.01$ ), confirming our conjecture about the benefit of the residual described in Section 3.3. Lower rank and higher noise significantly reduce the overall utility due to the large accumulated error, and the residual term plays a crucial role in improving local learning, balancing personalization and generalization.




Figure 2: Test accuracy of ablation study on noise level, rank and residual term for Caltech101

**Effect of privacy noise  $\epsilon$ .** We study how different privacy level affects the model performance on both local and neighbor classes in Figure 2 (b) and (c). As expected, the local classes accuracy gradually decreases when noise level increases. However, we see some unexpected improvement in neighbor utility when we increase DP noise from 0.4 to 0.1. As explained in Section 4.2, certain privacy noise range acts like regularization error that prevents overfitting, resulting in better generalization for higher noise level.

**Effect of factorization rank.** We explore the impact of the factorization rank by comparing different rank values with full-rank setting in Figure 2 (b) and (c). Intuitively, higher rank values lead to better performance, however, this is not always the case. For local classes, rank 8 generally performs better than full-rank under higher noise levels ( $\epsilon = \{0.01, 0.1\}$ ). Lower-rank has fewer entries, hence the amount of noise added is less than full-rank for the same level of privacy guarantee, which leads to better accuracy. For neighbor classes, lower rank value tends to have higher utility because lower rank introduces more regularization error that prevents overfitting and improves generalization.

## 5 CONCLUSION

In this paper, we presented a novel approach to address the critical challenges of personalization, generalization, and privacy in FPL for multimodal LLMs. Our proposed framework leverages low-rank decomposition to balance the tradeoffs between these competing objectives. By decomposing the local prompts into low-rank components iteratively while incorporating a residual term, our method effectively preserves both generalization and personalization. Moreover, we introduced a privacy-preserving mechanism that applies both global and local DP to safeguard sensitive client data. Unlike conventional methods, we selectively applied DP noise to the low-rank components, allowing us to maintain privacy without significantly degrading model performance. The critical role of the residual term in mitigating the effects of DP noise was also demonstrated, highlighting its importance for maintaining model expressiveness and personalization.

## REFERENCES

- 540  
541  
542 Martin Abadi, Andy Chu, Ian Goodfellow, H. Brendan McMahan, Ilya Mironov, Kunal Talwar, and  
543 Li Zhang. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC*  
544 *Conference on Computer and Communications Security, CCS '16*, pp. 308–318. Association for  
545 Computing Machinery, 2016.
- 546 Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. Intrinsic dimensionality explains the ef-  
547 fectiveness of language model fine-tuning. *arXiv preprint arXiv:2012.13255*, 2020.
- 548 Pathum Chamikara Mahawaga Arachchige, Peter Bertok, Ibrahim Khalil, Dongxi Liu, Seyit  
549 Camtepe, and Mohammed Atiquzzaman. Local differential privacy for deep learning. *IEEE*  
550 *Internet of Things Journal*, 7(7):5827–5842, 2019.
- 551 Manoj Ghuhan Arivazhagan, Vinay Aggarwal, Aaditya Kumar Singh, and Sunav Choudhary. Fed-  
552 erated learning with personalization layers, 2019. URL [https://arxiv.org/abs/1912.](https://arxiv.org/abs/1912.00818)  
553 [00818](https://arxiv.org/abs/1912.00818).
- 554 Bram Berlo, Aaqib Saeed, and Tanir Ozcelebi. Towards federated unsupervised representation learn-  
555 ing. In *Proceedings of the Third ACM International Workshop on Edge Systems, Analytics and*  
556 *Networking, EdgeSys '20*, pp. 31–36, New York, NY, USA, 2020. Association for Computing  
557 Machinery.
- 558 Dan Biderman, Jacob Portes, Jose Javier Gonzalez Ortiz, Mansheej Paul, Philip Greengard, Con-  
559 nor Jennings, Daniel King, Sam Havens, Vitaliy Chiley, Jonathan Frankle, Cody Blakeney, and  
560 John Patrick Cunningham. LoRA learns less and forgets less. *Transactions on Machine Learn-*  
561 *ing Research*, 2024. ISSN 2835-8856. URL [https://openreview.net/forum?id=](https://openreview.net/forum?id=al0Eru2qCG)  
562 [al0Eru2qCG](https://openreview.net/forum?id=al0Eru2qCG). Featured Certification.
- 563 Alberto Bietti, Chen-Yu Wei, Miroslav Dudik, John Langford, and Steven Wu. Personalization im-  
564 proves privacy-accuracy tradeoffs in federated learning. In *International Conference on Machine*  
565 *Learning*, pp. 1945–1962. PMLR, 2022.
- 566 Lukas Bossard, Matthieu Guillaumin, and Luc Van Gool. Food-101 – mining discriminative com-  
567 ponents with random forests. In *European Conference on Computer Vision*, 2014.
- 568 Liam Collins, Hamed Hassani, Aryan Mokhtari, and Sanjay Shakkottai. Exploiting shared represen-  
569 tations for personalized federated learning. In Marina Meila and Tong Zhang (eds.), *Proceedings*  
570 *of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Ma-*  
571 *chine Learning Research*, pp. 2089–2099. PMLR, 18–24 Jul 2021.
- 572 Tianyu Cui, Hongxia Li, Jingya Wang, and Ye Shi. Harmonizing generalization and personalization  
573 in federated prompt learning. In Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian  
574 Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (eds.), *Proceedings of the 41st*  
575 *International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning*  
576 *Research*, pp. 9646–9661. PMLR, 21–27 Jul 2024.
- 577 Wenlong Deng, Christos Thrampoulidis, and Xiaoxiao Li. Unlocking the potential of prompt-tuning  
578 in bridging generalized and personalized federated learning. In *Proceedings of the IEEE/CVF*  
579 *Conference on Computer Vision and Pattern Recognition*, pp. 6087–6097, 2024.
- 580 Yuyang Deng, Mohammad Mahdi Kamani, and Mehrdad Mahdavi. Adaptive personalized federated  
581 learning, 2020. URL <https://arxiv.org/abs/2003.13461>.
- 582 Canh Dinh, Nguyen Tran, and Josh Nguyen. Personalized federated learning with moreau envelopes.  
583 In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural*  
584 *Information Processing Systems*, volume 33, pp. 21394–21405. Curran Associates, Inc., 2020.
- 585 Alexey Dosovitskiy. An image is worth 16x16 words: Transformers for image recognition at scale.  
586 *arXiv preprint arXiv:2010.11929*, 2020.
- 587 Cynthia Dwork, Aaron Roth, et al. The algorithmic foundations of differential privacy. *Foundations*  
588 *and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.
- 589  
590  
591  
592  
593

- 594 Xiuwen Fang and Mang Ye. Robust federated learning with noisy and heterogeneous clients. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10072–  
595 10081, 2022.
- 596  
597 Li Fei-Fei, R. Fergus, and P. Perona. Learning generative visual models from few training exam-  
598 ples: An incremental bayesian approach tested on 101 object categories. In *2004 Conference on*  
599 *Computer Vision and Pattern Recognition Workshop*, pp. 178–178, 2004.
- 600  
601 Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for  
602 clustered federated learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin  
603 (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 19586–19597. Curran  
604 Associates, Inc., 2020.
- 605  
606 Tao Guo, Song Guo, and Junxiao Wang. Pfedprompt: Learning personalized prompt for vision-  
607 language models in federated learning. In *Proceedings of the ACM Web Conference 2023*, pp.  
608 1364–1374, 2023a.
- 609  
610 Tao Guo, Song Guo, Junxiao Wang, Xueyang Tang, and Wenchao Xu. Promptfl: Let federated  
611 participants cooperatively learn prompts instead of models-federated learning in age of foundation  
612 model. *IEEE Transactions on Mobile Computing*, 2023b.
- 613  
614 Chaoyang He, Murali Annavaram, and Salman Avestimehr. Group knowledge transfer: Federated  
615 learning of large cnns at the edge. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and  
616 H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 14068–14080.  
Curran Associates, Inc., 2020.
- 617  
618 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recog-  
619 nition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp.  
620 770–778, 2016.
- 621  
622 Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang,  
623 and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint*  
*arXiv:2106.09685*, 2021a.
- 624  
625 Shengyuan Hu, Zhiwei Steven Wu, and Virginia Smith. Private multi-task learning: Formulation  
626 and applications to federated learning. *arXiv preprint arXiv:2108.12978*, 2021b.
- 627  
628 Hai Huang, Zhengyu Zhao, Michael Backes, Yun Shen, and Yang Zhang. Prompt backdoors in  
629 visual prompt learning. *arXiv preprint arXiv:2310.07632*, 2023.
- 630  
631 Hamish Ivison, Yizhong Wang, Valentina Pyatkin, Nathan Lambert, Matthew Peters, Pradeep  
632 Dasigi, Joel Jang, David Wadden, Noah A Smith, Iz Beltagy, et al. Camels in a changing cli-  
633 mate: Enhancing lm adaptation with tulu 2. *arXiv preprint arXiv:2311.10702*, 2023.
- 634  
635 Prateek Jain, John Rush, Adam Smith, Shuang Song, and Abhradeep Guha Thakurta. Differentially  
636 private model personalization. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and  
637 J. Wortman Vaughan (eds.), *Advances in Neural Information Processing Systems*, volume 34, pp.  
638 29723–29735. Curran Associates, Inc., 2021.
- 639  
640 Wonyong Jeong and Sung Ju Hwang. Factorized-fl: Personalized federated learning with parameter  
641 factorization & similarity matching. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave,  
642 K. Cho, and A. Oh (eds.), *Advances in Neural Information Processing Systems*, volume 35, pp.  
643 35684–35695. Curran Associates, Inc., 2022.
- 644  
645 Jakub Konečný. Federated learning: Strategies for improving communication efficiency. *arXiv*  
646 *preprint arXiv:1610.05492*, 2016.
- 647  
648 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images.  
2009.
- 649  
650 Daliang Li and Junpu Wang. Fedmd: Heterogenous federated learning via model distillation. *CoRR*,  
abs/1910.03581, 2019. URL <http://arxiv.org/abs/1910.03581>.

- 648 Guanhao Li, Wansen Wu, Yan Sun, Li Shen, Baoyuan Wu, and Dacheng Tao. Visual prompt based  
649 personalized federated learning. *arXiv preprint arXiv:2303.08678*, 2023.
- 650  
651 Hongxia Li, Wei Huang, Jingya Wang, and Ye Shi. Global and local prompts cooperation via optimal  
652 transport for federated learning. *arXiv preprint arXiv:2403.00041*, 2024.
- 653  
654 Tian Li, Anit Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated  
655 optimization in heterogeneous networks. In I. Dhillon, D. Papailiopoulos, and V. Sze (eds.),  
656 *Proceedings of Machine Learning and Systems*, volume 2, pp. 429–450, 2020.
- 657  
658 Shih-Yang Liu, Chien-Yi Wang, Hongxu Yin, Pavlo Molchanov, Yu-Chiang Frank Wang, Kwang-  
659 Ting Cheng, and Min-Hung Chen. Dora: Weight-decomposed low-rank adaptation. *arXiv  
preprint arXiv:2402.09353*, 2024.
- 660  
661 Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number  
662 of classes. In *2008 Sixth Indian Conference on Computer Vision, Graphics and Image Processing*,  
663 pp. 722–729, 2008.
- 664  
665 Omkar M Parkhi, Andrea Vedaldi, Andrew Zisserman, and C. V. Jawahar. Cats and dogs. In *2012  
IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3498–3505, 2012.
- 666  
667 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,  
668 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual  
669 models from natural language supervision. In *International conference on machine learning*, pp.  
670 8748–8763. PMLR, 2021.
- 671  
672 Osama Shahid, Seyedamin Pouriyeh, Reza M. Parizi, Quan Z. Sheng, Gautam Srivastava, and Liang  
673 Zhao. Communication efficiency in federated learning: Achievements and challenges. *CoRR*,  
abs/2107.10996, 2021. URL <https://arxiv.org/abs/2107.10996>.
- 674  
675 Neta Shoham, Tomer Avidor, Aviv Keren, Nadav Israel, Daniel Benditkis, Liron Mor-Yosef, and  
676 Itai Zeitak. Overcoming forgetting in federated learning on non-iid data. *CoRR*, abs/1910.07796,  
2019. URL <http://arxiv.org/abs/1910.07796>.
- 677  
678 Reza Shokri, Marco Stronati, Congzheng Song, and Vitaly Shmatikov. Membership inference at-  
679 tacks against machine learning models. In *2017 IEEE symposium on security and privacy (SP)*,  
680 pp. 3–18. IEEE, 2017.
- 681  
682 Guangyu Sun, Matias Mendieta, Jun Luo, Shandong Wu, and Chen Chen. Fedperfix: Towards  
683 partial model personalization of vision transformers in federated learning. In *Proceedings of the  
IEEE/CVF International Conference on Computer Vision*, pp. 4988–4998, 2023.
- 684  
685 Youbang Sun, Zitao Li, Yaliang Li, and Bolin Ding. Improving lora in privacy-preserving federated  
686 learning. *arXiv preprint arXiv:2403.12313*, 2024.
- 687  
688 Anvith Thudi, Ilia Shumailov, Franziska Boenisch, and Nicolas Papernot. Bounding membership  
689 inference. *arXiv preprint arXiv:2202.12232*, 2022.
- 690  
691 Yixin Wu, Rui Wen, Michael Backes, Pascal Berrang, Mathias Humbert, Yun Shen, and Yang Zhang.  
692 Quantifying Privacy Risks of Prompts in Visual Prompt Learning. In *USENIX Security Sym-  
posium (USENIX Security)*. USENIX, 2024.
- 693  
694 Jie Xu, Karthikeyan Saravanan, Rogier van Dalen, Haaris Mehmood, David Tuckey, and Mete Ozay.  
695 Dp-dylora: Fine-tuning transformer-based models on-device under differentially private federated  
696 learning using dynamic low-rank adaptation. *arXiv preprint arXiv:2405.06368*, 2024.
- 697  
698 Fu-En Yang, Chien-Yi Wang, and Yu-Chiang Frank Wang. Efficient model personalization in feder-  
699 ated learning via client-specific prompt generation. In *Proceedings of the IEEE/CVF International  
Conference on Computer Vision*, pp. 19159–19168, 2023a.
- 700  
701 Xiyuan Yang, Wenke Huang, and Mang Ye. Dynamic personalized federated learning with adap-  
tive differential privacy. *Advances in Neural Information Processing Systems*, 36:72181–72192,  
2023b.

Hongwei Yao, Jian Lou, Zhan Qin, and Kui Ren. Promptcare: Prompt copyright protection by watermark injection and verification. In *2024 IEEE Symposium on Security and Privacy (SP)*, pp. 845–861, 2024.

Da Yu, Huishuai Zhang, Wei Chen, Jian Yin, and Tie-Yan Liu. Large scale private learning via low-rank reparametrization. In *International Conference on Machine Learning*, pp. 12208–12218. PMLR, 2021.

Jianqing Zhang, Yang Hua, Hao Wang, Tao Song, Zhengui Xue, Ruhui Ma, and Haibing Guan. Fedcp: Separating feature information for personalized federated learning via conditional policy. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, KDD '23*, pp. 3249–3261, New York, NY, USA, 2023. Association for Computing Machinery.

Michael Zhang, Karan Sapra, Sanja Fidler, Serena Yeung, and Jose M Alvarez. Personalized federated learning with first order model optimization. *arXiv preprint arXiv:2012.08565*, 2020.

Haodong Zhao, Wei Du, Fangqi Li, Peixuan Li, and Gongshen Liu. Fedprompt: Communication-efficient and privacy-preserving prompt tuning in federated learning. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2023.

Terry Yue Zhuo, Armel Zebaze, Nitchakarn Supattarachai, Leandro von Werra, Harm de Vries, Qian Liu, and Niklas Muennighoff. Astraios: Parameter-efficient instruction tuning code large language models. *arXiv preprint arXiv:2401.00788*, 2024.

## A EXPERIMENTAL DETAILS

### A.1 DATASETS

**Caltech101** (Fei-Fei et al., 2004) is an image dataset that contains images from 101 object categories (e.g., “helicopter”, “elephant” and “chair” etc.). There are about 40 to 800 images for each object category, but most categories have about 50 images. The dataset is available for download on [http://www.vision.caltech.edu/Image\\_Datasets/Caltech101/101\\_ObjectCategories.tar.gz](http://www.vision.caltech.edu/Image_Datasets/Caltech101/101_ObjectCategories.tar.gz). The dataset contains 6,593 samples, including 4,128 training samples and 2,465 testing samples.

**Oxford Pets** (Parkhi et al., 2012) is an image dataset with 37 object classes that can be downloaded at <https://www.robots.ox.ac.uk/~vgg/data/pets/data/images.tar.gz>. The dataset consists of 6,613 pet images with roughly 200 images for each class. The dataset is divided into training set of 2,944 images and test set of 3,669 images.

**Oxford Flowers** (Nilsback & Zisserman, 2008) is an image classification dataset consisting of 102 flower categories, each class has between 40 and 258 images. The dataset is can be retrieved at <https://www.robots.ox.ac.uk/~vgg/data/flowers/102/102flowers.tgz>. There are 6,556 total images, including 4,093 training images and 2,463 testing images.

**Food101** (Bossard et al., 2014) is a large-scale dataset containing images of 101 different types of food. The dataset is are available for download at [https://data.vision.ee.ethz.ch/cvl/datasets\\_extra/food-101/](https://data.vision.ee.ethz.ch/cvl/datasets_extra/food-101/). There are 80,800 total images, and we split the dataset into train set of size 50,500 and test set of size 30,300.

**CIFAR-100** (Krizhevsky et al., 2009) is a another large-scale dataset containing images of 100 different object classes. The dataset is are available for download via `torchvision.datasets.CIFAR100`. The dataset consists of 60,000 32x32 images, with 6,000 images per class. We divide the dataset into training set of 50,000 images and test set of 10,000 images.

### A.2 IMPLEMENTATION DETAILS

For the first four datasets Caltech101, OxfordPets, OxfordFlowers and Food101, we use the Vision Transformer ViT-B16 (Dosovitskiy, 2020) as the backbone for the frozen CLIP model for Caltech101, OxfordPets, OxfordFlowers and Food101. For each dataset, we run experiments with

$N = 10$  clients for  $T = 100$  global training rounds. We use batch size  $|\mathcal{B}| = 32$  for training and  $|\mathcal{B}| = 100$  for testing. We set the global learning rate  $\eta_G = 0.0001$  and local learning rate  $\eta_L = 0.0001$  with SGD optimizer.

We adopt the Pathological setting for data heterogeneity among clients as implemented in <https://github.com/KaiyangZhou/CoOp/blob/main/DATASETS.md>. The class labels are splitted randomly among 10 clients without overlapping, and each client owns disjoint set of local classes. For each client, we train the local model on data associated with their assigned local classes. We then evaluate each client’s local model on two test sets: local class test set for personalization and neighbor class test set for generalization. The local class test set involve all the test image associated with the client’s local class labels. The neighbor class test set is the set of all test images whose labels are owned by other clients.

For CIFAR-100, we adopt ResNet50 (He et al., 2016) as the backbone model and run experiments with  $N = 25$  and  $N = 50$  clients for  $T = 200$  global training rounds. We use batch size  $|\mathcal{B}| = 32$  for training and  $|\mathcal{B}| = 100$  for testing. We set the global learning rate  $\eta_G = 0.0001$  and local learning rate  $\eta_L = 0.0001$  with SGD optimizer. We use Dirichlet data distribution to simulate the real-world non-IID setting with parameter  $\alpha = 0.3$ . The test accuracy is averaged across all clients.

We provide a summary of the experiment set up in Table 5 below.

Table 5: Details of datasets and experimental set up

Dataset	Data distribution	Model	Number of clients	Number of training rounds	Training batch size	Testing batch size
Caltech101	Pathological	ViT-B16	10	100	32	100
OxfordPets	Pathological	ViT-B16	10	100	32	100
OxfordFlowers	Pathological	ViT-B16	10	100	32	100
Food101	Pathological	ViT-B16	10	100	32	100
CIFAR-100	Dirichlet	ResNet50	25, 50	200	32	100

For the prompt learner, the length of prompt vectors is  $b = 16$  with a dimension of  $d = 512$ , and the token position is “end” with “random” initialization. For the factorization process, we experiment with four different decomposition rank 1, 2, 4, 8. We consider three different DP noises with privacy level from low to high:  $\epsilon \in \{0.4, 0.2, 0.1\}$ . The clipping threshold is chosen to be  $C_{th} = 10$  for both GDP and LDP applications.

For the baseline methods, we consider three settings: PromptFL (Guo et al., 2023b), FedOTP (Li et al., 2024) and FedPGP (Cui et al., 2024) with the main difference lies in the prompt learner structure. We describe each baseline in detail below.

- PromptFL:** PromptFL follows the traditional federated learning framework where each client has one single prompt  $p_i$  and the aggregated prompt is the average of all clients’  $p_i$ . In this setting, the privacy noise is added directly to each client’s  $p_i$  before sharing with the server for aggregation.
- FedOTP:** In FedOTP, each client’s customized prompt  $p_i$  involves two full-rank global prompt  $p_{G,i}$  and local prompt  $p_{L,i}$ . To incorporate privacy, the GDP noise is added to the averaged  $p_G$  by the server and the LDP noise is added to each client’s local prompt  $p_{L,i}$ .
- FedPGP:** In FedPGP, each client’s customized prompt  $p_i$  includes a full-rank global prompt  $p_{G,i}$  and two low-rank local components  $u_i, v_i$ . Each client then train three parameters  $p_{G,i}, u_i$  and  $v_i$  across the training process. In this baseline, the GDP noise is added to the averaged  $p_G$  by the server and the LDP noise is added to the two low-rank terms  $u_i, v_i$  of each client.

We run our experiment on a computer cluster, each node has a 6x NVIDIA Tesla V100 GPUs with 32 GiB of memory and 512 GiB RAM and 2x IBM Power 9 processors. The final result is the mean accuracy over the last 10 training epochs, averaged across all clients.

### A.3 MEMBERSHIP INFERENCE ATTACK

In this section, we evaluate the privacy-preserving performance of our method DP-FPL against Membership Inference Attack (MIA). We implement MIA as described in Shokri et al. (2017), where the goal of the attack is to infer the appearance of data samples in one client’s training dataset. We first train a set of 50 shadow models with the same model architecture as the target client’s model. The shadow models generate synthetic training data that is used to train the attack model. The attack model is a two-layer MLP and a classification head that predicts if a given sample is part of the target client’s training data or not. We run the attack on Caltech101, Oxford Pets and Oxford Flowers with different privacy levels. For each dataset, we perform separate attack on each class and compute the average success rate, i.e. percentage of correct guesses.

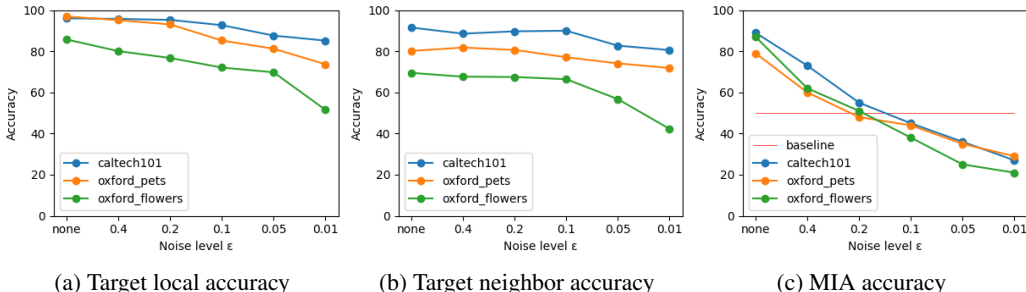


Figure 3: Target model performance (a, b) and MIA performance (c) with rank 8. The baseline of MIA accuracy is set to 50% (random guessing).

Figure 3 demonstrates the performance of the target model (local and neighbor classes) and the MIA. We set the MIA baseline accuracy to be 50%, representing the expected success rate of random guessing. Looking at Figure 3c, the MIA accuracy is low ( $< 50\%$ ) when  $\epsilon = 0.2$  for Oxford Pets and  $\epsilon = 0.1$  for Caltech101 and Oxford Flowers. In addition,  $\epsilon = 0.1$  causes minimal loss in the target model accuracy for both local and neighbor classes (Figures 3a and 3b). Therefore, one can balance the utility-privacy tradeoff by setting  $\epsilon = 0.1$  for any dataset. This shows that our approach effectively protects the training data from MIA while still maintaining good model performance.

### A.4 ADDITIONAL EXPERIMENTAL RESULTS

In this section, we include additional results from the experiments introduced in Section 4 to further investigate how the DP parameter  $\epsilon$ , factorization rank and the residual term affect the tradeoff between personalization, generalization and privacy guarantee. Figures 4, 5 and 6 continue the ablation study on the effect of the key parameters that directly affect the tradeoff: residual term, noise level and rank. We summarize the results for each dataset below.

Figure 4 shows the test accuracy with different parameter settings for Oxford Pets. In general, there is an increase in both local and neighbor classes when incorporating the residual term, highlighting the benefit of the residual in model utility in different datasets. The difference in accuracy is more consistent across different noise and rank values compared to Figure 2. In addition, there is minimal growth in accuracy when we increase the rank value. Therefore, it is beneficial to set any rank value for this particular dataset.




Figure 4: Test accuracy of ablation study on noise level, rank and residual term for Oxford Pets



Figure 5 shows the test accuracy with different parameter settings for Oxford Flowers. The addition of the residual term still improves the overall utility, though the benefit is modest for this dataset. Similar to Oxford Pets, the rank value does not significantly affect the accuracy for both local and neighbor classes. We also observe more drastic drop in accuracy under strong privacy level ( $\epsilon = 0.01$ ). Overall, it is more difficult to balance the tradeoff for Oxford Flowers. One needs to sacrifice strong data protection and set  $\epsilon \geq 0.1$  to achieve good personalization and generalization utility.




Figure 5: Test accuracy of ablation study on noise level, rank and residual term for Oxford Flowers

Figure 6 shows the test accuracy with different parameter settings for Food101. We see an overall increase in local and neighbor accuracy with the introduction of the residual term, and the difference is more significant under higher noise level. In addition, there is minimal reduction in model utility when we increase privacy noise. This behavior is consistent across all rank values, indicating the robustness of our method under strict privacy constraints.




Figure 6: Test accuracy of ablation study on noise level, rank and residual term for Food101

Overall, the affect of the key parameters (noise level, rank and residual term) on the tradeoff between personalization, generalization and privacy varies widely among different datasets. Nonetheless, the results across all datasets show consistent trends, indicating the effectiveness and applicability of our method.

For completeness, we include Tables 6, 7, 8 and 9 which detail the test accuracy with standard deviation of all the ablation experiments demonstrated in Figures 2, 4, 5 and 6.

Table 6: Mean test accuracy of our method averaged across 10 clients under non-private setting (without any DP noise).

Dataset	Rank	Local classes		Neighbor classes	
		without residual	with residual	without residual	with residual
Caletech101	1	86.64±0.78	<b>93.55±0.18</b>	89.70±0.81	<b>93.40±0.28</b>
	2	87.54±0.83	<b>94.51±0.17</b>	87.95±0.42	<b>92.79±0.72</b>
	4	89.57±0.47	<b>95.69±0.57</b>	86.86±0.72	<b>92.32±0.71</b>
	8	92.12±0.47	<b>96.06±0.39</b>	87.58±0.17	<b>91.54±0.73</b>
Oxford Pets	1	89.78±0.71	<b>93.46±0.63</b>	89.54±0.71	<b>92.17±0.27</b>
	2	88.24±0.73	<b>94.49±0.91</b>	87.37±0.38	<b>91.96±0.71</b>
	4	91.28±0.64	<b>96.36±0.28</b>	88.70±0.83	<b>87.61±0.53</b>
	8	92.42±0.81	<b>96.91±0.58</b>	90.71±1.09	<b>80.19±0.64</b>
Oxford Flowers	1	67.41±0.19	<b>75.45±0.20</b>	67.22±0.63	<b>71.49±0.71</b>
	2	66.18±0.89	<b>78.56±0.90</b>	66.16±0.10	<b>69.76±1.11</b>
	4	67.04±0.61	<b>81.53±0.49</b>	67.02±0.39	<b>68.33±0.32</b>
	8	68.35±0.61	<b>85.75±0.91</b>	68.75±0.47	<b>69.51±0.61</b>
Food101	1	74.66±1.14	<b>86.12±0.38</b>	74.58±0.51	<b>85.97±0.72</b>
	2	74.91±0.38	<b>86.06±0.27</b>	74.96±0.49	<b>85.89±0.52</b>
	4	74.95±0.91	<b>86.18±0.63</b>	74.84±0.99	<b>86.16±1.06</b>
	8	76.95±1.02	<b>86.08±0.95</b>	76.02±0.77	<b>86.08±0.69</b>

Table 7: Mean test accuracy of DP-FPL averaged across 10 clients. The DP noise is set to  $\epsilon = 0.4$ .

Dataset	Rank	Local classes		Neighbor classes	
		without residual	with residual	without residual	with residual
Caltech101	1	87.66±0.79	<b>92.93±1.18</b>	90.19±1.11	<b>92.94±0.68</b>
	2	88.98±1.02	<b>93.03±0.90</b>	86.99±1.34	<b>88.82±0.79</b>
	4	88.01±1.21	<b>93.58±0.87</b>	85.20±1.03	<b>89.15±0.90</b>
	8	90.13±1.19	<b>95.74±0.63</b>	85.46±1.08	<b>88.58±0.38</b>
Oxford Pets	1	74.87±0.69	<b>90.43±0.79</b>	72.41±0.49	<b>87.83±0.95</b>
	2	74.66±0.88	<b>91.67±0.62</b>	72.69±0.33	<b>86.26±0.42</b>
	4	74.24±0.74	<b>93.22±0.40</b>	71.83±0.39	<b>83.53±0.60</b>
	8	77.22±0.82	<b>95.13±0.52</b>	73.86±0.42	<b>81.82±0.42</b>
Oxford Flowers	1	62.48±0.83	<b>75.45±0.55</b>	64.41±0.58	<b>71.49±0.55</b>
	2	65.67±0.96	<b>77.39±0.69</b>	64.53±0.75	<b>70.97±0.58</b>
	4	66.23±0.79	<b>78.62±1.04</b>	63.97±0.77	<b>68.10±0.67</b>
	8	68.03±1.41	<b>80.09±0.51</b>	63.22±0.88	<b>67.67±0.45</b>
Food101	1	50.55±1.05	<b>69.76±0.43</b>	50.29±1.14	<b>68.11±0.94</b>
	2	56.07±0.85	<b>72.05±0.66</b>	57.10±1.13	<b>71.15±0.65</b>
	4	61.02±1.09	<b>77.76±0.98</b>	61.25±0.58	<b>76.48±0.57</b>
	8	69.53±0.20	<b>81.45±1.01</b>	67.52±0.40	<b>81.00±0.90</b>

Table 8: Mean test accuracy of DP-FPL averaged across 10 clients. The DP noise is set to  $\epsilon = 0.2$ .

Dataset	Rank	Local classes		Neighbor classes	
		without residual	with residual	without residual	with residual
Caltech101	1	87.79±0.81	<b>92.16±0.38</b>	90.43±0.91	<b>91.61±1.02</b>
	2	88.99±0.63	<b>93.38±0.82</b>	87.22±0.47	<b>90.55±0.49</b>
	4	89.25±0.61	<b>94.22±0.54</b>	85.31±0.28	<b>90.73±0.10</b>
	8	91.75±0.19	<b>95.28±0.93</b>	86.56±0.71	<b>89.72±0.83</b>
Oxford Pets	1	70.06±0.73	<b>89.79±0.59</b>	73.27±0.51	<b>87.69±1.08</b>
	2	72.71±0.42	<b>90.12±0.49</b>	72.04±0.39	<b>86.74±0.18</b>
	4	72.27±0.84	<b>91.45±0.65</b>	70.16±0.77	<b>83.62±0.94</b>
	8	73.73±0.94	<b>93.09±0.30</b>	71.04±0.55	<b>80.67±0.69</b>
Oxford Flowers	1	55.18±0.92	<b>72.19±0.17</b>	64.53±0.81	<b>70.35±0.97</b>
	2	59.52±0.53	<b>73.58±0.57</b>	61.41±0.31	<b>70.50±0.29</b>
	4	63.89±0.83	<b>75.18±0.57</b>	61.50±0.81	<b>68.65±0.46</b>
	8	63.86±0.86	<b>76.75±0.97</b>	62.65±0.65	<b>67.51±0.73</b>
Food101	1	57.38±0.35	<b>69.94±0.82</b>	58.11±0.37	<b>63.63±0.92</b>
	2	59.22±0.16	<b>70.52±0.75</b>	58.17±0.57	<b>67.64±0.38</b>
	4	65.12±0.52	<b>75.11±0.46</b>	64.74±0.44	<b>70.21±0.80</b>
	8	70.16±0.32	<b>81.25±0.51</b>	70.21±0.75	<b>80.79±0.17</b>

Table 9: Mean test accuracy of DP-FPL averaged across 10 clients. The DP noise is set to  $\epsilon = 0.1$ .

Dataset	Rank	Local classes		Neighbor classes	
		without residual	with residual	without residual	with residual
Caltech101	1	84.73±0.32	<b>91.39±0.58</b>	89.33±0.85	<b>90.81±0.59</b>
	2	85.32±0.38	<b>91.11±0.86</b>	86.32±0.28	<b>89.18±0.67</b>
	4	87.82±0.81	<b>92.26±1.28</b>	84.73±0.97	<b>87.92±0.80</b>
	8	88.82±0.84	<b>92.71±0.61</b>	85.32±0.86	<b>90.02±0.48</b>
Oxford Pets	1	67.59±0.19	<b>80.41±0.38</b>	71.47±0.32	<b>79.24±0.61</b>
	2	69.22±0.48	<b>81.65±1.06</b>	67.38±0.27	<b>79.65±0.59</b>
	4	69.35±0.89	<b>82.73±0.49</b>	66.11±0.66	<b>78.73±0.87</b>
	8	74.48±0.91	<b>85.25±0.30</b>	66.10±0.49	<b>77.12±0.63</b>
Oxford Flowers	1	54.18±0.65	<b>66.78±1.02</b>	59.51±0.75	<b>66.93±0.49</b>
	2	57.16±0.39	<b>66.21±0.29</b>	65.35±0.64	<b>65.21±0.94</b>
	4	59.71±0.83	<b>67.32±0.27</b>	67.71±0.41	<b>67.26±0.32</b>
	8	59.81±0.97	<b>72.11±1.05</b>	66.98±0.38	<b>66.44±0.17</b>
Food101	1	62.75±0.37	<b>68.65±0.81</b>	58.05±0.42	<b>68.60±0.28</b>
	2	62.43±0.76	<b>71.84±0.95</b>	58.86±0.26	<b>69.66±0.93</b>
	4	64.66±0.75	<b>75.64±0.41</b>	65.22±0.28	<b>72.61±0.30</b>
	8	67.20±0.48	<b>80.57±0.83</b>	67.06±0.99	<b>78.14±0.37</b>

972  
973  
974  
975  
976  
977  
978  
979  
980  
981  
982  
983  
984  
985  
986  
987  
988  
989  
990  
991  
992  
993  
994  
995  
996  
997  
998  
999  
1000  
1001  
1002  
1003  
1004  
1005  
1006  
1007  
1008  
1009  
1010  
1011  
1012  
1013  
1014  
1015  
1016  
1017  
1018  
1019  
1020  
1021  
1022  
1023  
1024  
1025

Table 10: Mean test accuracy of DP-FPL averaged across 10 clients. The DP noise is set to  $\epsilon = 0.05$ .

Dataset	Rank	Local classes		Neighbor classes	
		without residual	with residual	without residual	with residual
Caltech101	1	78.62±0.24	<b>81.59±1.55</b>	78.75±0.23	<b>79.07±1.59</b>
	2	78.63±0.18	<b>85.42±0.93</b>	77.67±0.34	<b>77.42±0.92</b>
	4	78.57±0.48	<b>85.05±1.06</b>	77.51±0.48	<b>82.28±1.07</b>
	8	78.72±0.49	<b>87.64±0.98</b>	76.83±0.42	<b>82.76±1.69</b>
Oxford Pets	1	63.28±0.57	<b>70.60±1.30</b>	61.75±0.58	<b>72.18±1.32</b>
	2	63.32±0.68	<b>73.66±0.94</b>	61.76±0.70	<b>72.62±1.39</b>
	4	64.26±0.91	<b>74.29±1.35</b>	61.71±0.91	<b>68.73±1.31</b>
	8	65.22±0.89	<b>81.26±1.00</b>	61.69±0.62	<b>74.13±0.99</b>
Oxford Flowers	1	32.98±1.13	<b>59.62±0.99</b>	29.91±0.61	<b>46.80±0.91</b>
	2	32.84±1.80	<b>59.74±1.05</b>	30.02±1.14	<b>49.19±1.47</b>
	4	32.45±1.15	<b>65.72±1.17</b>	30.40±0.75	<b>50.32±1.18</b>
	8	33.09±1.62	<b>69.80±0.94</b>	30.21±1.28	<b>56.75±1.23</b>
Food101	1	58.72±0.36	<b>67.19±0.92</b>	56.37±0.25	<b>67.09±0.46</b>
	2	57.01±0.41	<b>67.71±1.42</b>	55.59±0.77	<b>67.27±0.57</b>
	4	59.03±0.45	<b>68.13±1.06</b>	55.45±0.19	<b>69.44±0.55</b>
	8	64.88±0.37	<b>78.23±1.36</b>	63.15±0.55	<b>77.18±0.77</b>

Table 11: Mean test accuracy of DP-FPL averaged across 10 clients. The DP noise is set to  $\epsilon = 0.01$ .

Dataset	Rank	Local classes		Neighbor classes	
		without residual	with residual	without residual	with residual
Caltech101	1	73.19±0.25	<b>80.89±1.62</b>	82.59±0.22	<b>87.98±1.59</b>
	2	75.18±0.50	<b>83.08±1.16</b>	82.60±0.85	<b>85.70±0.79</b>
	4	78.19±0.48	<b>84.35±1.07</b>	80.61±0.48	<b>82.81±1.07</b>
	8	83.18±0.60	<b>85.21±2.27</b>	77.59±0.57	<b>80.60±1.12</b>
Oxford Pets	1	53.05±0.59	<b>68.73±1.29</b>	48.54±1.15	<b>65.35±1.36</b>
	2	53.04±1.37	<b>71.62±1.10</b>	49.54±1.40	<b>67.58±1.10</b>
	4	53.06±0.91	<b>73.14±1.33</b>	50.52±1.35	<b>70.39±1.36</b>
	8	54.29±1.01	<b>73.71±1.19</b>	51.48±1.48	<b>71.89±1.48</b>
Oxford Flowers	1	32.44±1.07	<b>39.54±0.97</b>	30.40±0.90	<b>31.24±1.08</b>
	2	33.11±1.84	<b>42.53±1.78</b>	30.94±1.53	<b>38.32±1.34</b>
	4	33.08±1.13	<b>45.27±1.23</b>	30.91±1.03	<b>41.24±1.16</b>
	8	33.15±1.72	<b>51.55±1.21</b>	30.95±1.80	<b>42.31±1.85</b>
Food101	1	57.63±0.44	<b>67.16±0.88</b>	56.57±0.31	<b>68.79±0.66</b>
	2	57.40±0.43	<b>68.28±1.40</b>	56.70±1.02	<b>68.48±0.60</b>
	4	58.98±0.55	<b>69.13±1.05</b>	59.39±0.20	<b>69.88±0.77</b>
	8	59.63±0.54	<b>77.45±3.65</b>	59.58±0.73	<b>76.87±0.55</b>