

EXPLORING PARAMETER-EFFICIENT FINE-TUNING FOR IMPROVING COMMUNICATION EFFICIENCY IN FEDERATED LEARNING

Anonymous authors

Paper under double-blind review

ABSTRACT

Federated learning (FL) has emerged as a promising paradigm for enabling the collaborative training of models without centralized access to the raw data on local devices. In the typical FL paradigm (e.g., FedAvg), model weights are sent to and from the server each round to participating clients. However, this can quickly put a massive communication burden on the system, especially if more capable models beyond very small MLPs are employed. Recently, the use of pre-trained models has been shown effective in federated learning optimization and improving convergence. This opens the door for new research questions. Can we adjust the weight-sharing paradigm in federated learning, leveraging strong and readily-available pre-trained models, to significantly reduce the communication burden while simultaneously achieving excellent performance? To this end, we investigate the use of parameter-efficient fine-tuning in federated learning. Specifically, we systemically evaluate the performance of several parameter-efficient fine-tuning methods across a variety of client stability, data distribution, and differential privacy settings. By only locally tuning and globally sharing a small portion of the model weights, significant reductions in the total communication overhead can be achieved while maintaining competitive performance in a wide range of federated learning scenarios, providing insight into a new paradigm for practical and effective federated systems.

1 INTRODUCTION

Federated learning (FL) has become increasingly prevalent in the research community, having the goal of enabling collaborative training with a network of clients without needing to share any private data. One key challenge for this training paradigm is overcoming data heterogeneity. The participating devices in a federated system are often deployed across a variety of users and environments, resulting in a non-IID data distribution. As the level of heterogeneity intensifies, optimization becomes increasingly difficult. Various techniques have been proposed for alleviating this issue. These primarily consist of modifications to the local or global objectives through proximal terms, regularization, and improved aggregation operations (Li et al., 2020; Karimireddy et al., 2020; Mendieta et al., 2022; Acar et al., 2021; Wang et al., 2020a). More recently, some works have investigated the role of model initialization in mitigating such effects (Nguyen et al., 2022; Chen et al., 2022b). Inspired by the common usage of pre-trained models for facilitating strong transfer learning in centralized training, researchers employed widely-available pre-trained weights for initialization in FL and were able to close much of the gap between federated and centralized performance.

Still, while pre-trained initializations are effective for alleviating heterogeneity effects in FL, another key challenge is left unaddressed; that is, communication constraints. This is often the primary bottleneck for real-world federated systems (Kairouz et al., 2021). In the standard FL framework (McMahan et al., 2017), updates for all model parameters are sent back and forth between the server and participating clients each round. This can quickly put a massive communication burden on the system, especially if more capable models beyond very small MLPs are used. When employing strong pre-trained models, the number of parameters can be large, such as for current state-of-the-art transformers. For example, ViT-Base (Dosovitskiy et al., 2021) has 84 million parameters, which would simply exacerbate the communication overhead to insurmountable levels.

However, we note that pre-trained models have strong representations, and updating all the weights during fine-tuning is often not necessary. Various parameter-efficient fine-tuning methods (e.g., fine-tuning only a subset of the parameters or the bias terms) for centralized training have been proposed in the literature and show that successful and efficient adaptation is possible, even under domain shift (Jia et al., 2022; Cai et al., 2020; Pfeiffer et al., 2020). We reason that such insight is applicable to FL, where each client can be thought of as a shifted domain on which we are fine-tuning. By leveraging pre-trained weights, it may be possible to simply update a small portion of the weights for each client. This will significantly reduce the communication burden on the system, as the updates communicated with the server will consist of just a fraction of the total model parameters.

Can we reap these potential communication benefits while still achieving strong performance in FL? Unfortunately, operating conditions in FL are difficult, requiring successful convergence under varying data heterogeneity levels, random client availability, and differential privacy procedures. Therefore, we are unable to properly assess this possibility of benefit based on existing literature, as the effectiveness of parameter-efficient fine-tuning methods has not been explored in such situations. To fill this gap, we explore the viability of a Federated Parameter-Efficient Fine-Tuning (FedPEFT) framework, illustrated in Figure 1. We deploy parameter-efficient fine-tuning methods to adapt pre-trained models to enable massive reductions in communication overheads. The contribution of this paper can be summarized as follows.

- We introduce FedPEFT, a new federated learning framework that simultaneously addresses data heterogeneity and communication challenges. FedPEFT is the first federated learning framework that enables the leveraging of strong pre-trained models in FL while maintaining an extremely low communication cost.
- We present a systematic study of the FedPEFT framework with various fine-tuning methods under heterogeneous data distributions, client availability ratios, and increasing degrees of domain gap relative to the pre-trained representations. (Sections 5 and 6)
- To ensure FedPEFT is practical for the complex environments of FL, we further analyze the robustness of FedPEFT among different backbone sizes, low-data regimes, and differential privacy operations. (Sections 7 and 8)

2 RELATED WORK

Federated Learning. Federated learning is a decentralized training paradigm composed of two procedures: local training and global aggregation. Therefore, most existing work focuses on either local training (Mendieta et al., 2022; Li et al., 2021; 2020) or global aggregation (Wang et al., 2020b; Yurochkin et al., 2019) to learn a better global model. Another line of work cuts into this problem by applying different initialization to help both procedures. (Chen et al., 2022b) show that initializing the model with pre-trained weights can make the global aggregation of FedAvg more stable, even when pre-trained with synthetic data. Furthermore, Nguyen et al. (2022) present the effectiveness of pre-training with different local and global operations. However, these works focus purely on the effect of initialization in a standard FedAvg framework and do not consider the communication constraints of the system. Our work pushes the envelope further by leveraging strong pre-trained models (even large, capable transformers) in federated learning while effectively handling the communication issue via parameter-efficient fine-tuning.

Communication in Federated Learning. Communication constraints are a primary bottleneck in federated learning. To reduce the communication cost, several previous work leverage model compression techniques (Konečný et al., 2017; Suresh et al., 2017). Such works do not change the training paradigm but rather post-processes the local model to reduce communication costs. For instance, Konečný et al. (2017) propose approaches that parameterize the model with fewer variables and compress the model in an encoding-decoding fashion. However, the minimal requirement to maintain all the information is still high when facing today’s large models. Meanwhile, another line of work changes the training paradigm by learning federated ensembles based on several pre-trained base models (Hamer et al., 2020). In this way, only the mixing weights of the base models will be communicated in each round. This approach aims to reduce the burden of downloading and uploading the entire model in each round. However, the base models are not directly trained, and the final performance is highly related to the base models. Meanwhile, model ensembles will take more time and space, which is often limited on the client side. Our framework follows the strategy of this

line of work that does not transmit the entire model, but we use only one pre-trained model instead of several base models and only transmit a subset of the parameters instead of the model ensembles. Therefore, no additional time or space is required.

Parameter-Efficient Fine-tuning. Fine-tuning is a prevalent topic in centralized transfer learning, especially in this era of the ‘‘Foundation Model’’ (Bommasani et al., 2022). A significant line of work is to reduce the trainable parameter number, i.e., parameter-efficient fine-tuning (PEFT) (Chen et al., 2022a; Pan et al., 2022; Liu et al., 2022). This enables easier access and usage of pre-trained models by reducing the memory cost needed to conduct fine-tuning due to fewer computed gradients. In federated learning, PEFT has an additional benefit that is not salient in centralized training: reduction of the communication cost. By introducing PEFT to federated learning, our work can take advantage of a strong (and even large) pre-trained model meanwhile significantly reducing communication costs.

3 FEDERATED PARAMETER-EFFICIENT FINE-TUNING

3.1 RECAP: CONVENTIONAL FEDERATED LEARNING WITH FEDAVG

In this section, we formally describe the federated learning objective and federated parameter-efficient fine-tuning. Before we dive into FedPEFT, we recap the formulation of conventional federated learning. Using a classification task as an example, K samples in a dataset $\mathbb{D} = \{(\mathbf{x}_k, y_k)_{k=1}^K\} = \cup_{n=1}^N \mathbb{D}_n$, where \mathbf{x} is the input and $y \in \{0, 1, \dots, C - 1\}$ is the label, are distributed among N clients. Each client has a local model $\{f_n = c_n \circ \phi_n\}_{n=1}^N$ parameterized by $\{\theta_n\}_{n=1}^N$, where ϕ is the feature extractor and c is the classification head. The goal of federated learning is to learn a global model $F = c \circ \phi$ parameterized by θ on the server from M sampled client models in T communication rounds.

At the beginning of training, the global model $F^{(0)}$ is randomly initialized or initialized with pre-trained weights, where the superscript t indicates the model at round t . In each round t , $f_n^{(t)}$ will be initialized by $F^{(t)}$ and updated by

$$\min_{\theta_n} \mathcal{L}(f_n^{(t)}) = \frac{1}{|\mathbb{D}_n|} \sum_{k=1}^{|\mathbb{D}_n|} \ell(y_k, f_n^{(t)}(\mathbf{x}_k)),$$

where ℓ is the loss function. After the local updates, the server will sample M clients from all N clients and aggregate $\{f_m^{(t)}\}_{m=1}^M$ with the FedAvg algorithm to a new global model

$$F^{(t+1)} = \sum_{m=1}^M \frac{|\mathbb{D}_m|}{\sum_{i=1}^M |\mathbb{D}_i|} f_m^{(t)}.$$

This procedure is repeated from $t = 0$ to $t = T - 1$. During the client-server communication, we only take the communication cost for the model into consideration, assuming the remaining communication costs are fixed. Therefore, the communication cost C is proportional to the transmission parameters number, thus can be formulated as

$$C \propto |\mathbb{P}| \cdot M,$$

where $\mathbb{P} \subseteq (F \text{ or } f)$ is the set of parameters to transmit in F for global model or f for client model. The final goal of this problem is to minimize the C while maintaining the server accuracy.

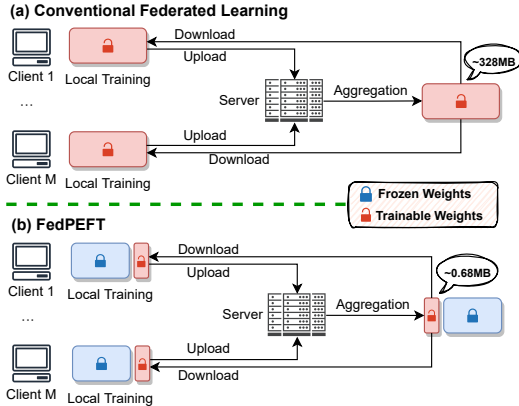


Figure 1: Process in a federated learning communication round with M participating clients. We use ViT-Base as an instance to analyze the communication costs. (a) Conventional federated learning framework, where the entire model will be sent during the communication. (b) FedPEFT, which is our proposed parameter-efficient framework for federated learning. Instead of uploading and downloading the entire model during the communication, only a small subset of the model will be updated and communicated, leading to a huge relief in communication, specifically from about 328MB per round to about 0.68MB per round.

3.2 FEDPEFT

In conventional federated learning, updates for the entire model need to be repeatedly sent to and from the server, resulting in significant communication costs, especially when larger, more capable modern neural network architectures are employed. To reduce this heavy burden, we deploy parameter-efficient fine-tuning methods to adapt pre-trained models to the local clients rather than fully fine-tuning all parameters, which is described in Algorithm 1. In the FedPEFT framework, illustrated in Figure 1, only a small amount of parameters in the local model will be downloaded, trained, and uploaded in each communication round. For instance, FedPEFT reduces the size of communication each round from 328MB/Client to 0.68MB/Client when using a pre-trained ViT-Base (Dosovitskiy et al., 2021) as the backbone.

To implement FedPEFT, we provide a canonical baseline approach (head-tuning) and three prototypes leveraging different parameter-efficient fine-tuning methods (Bias, Adapter, and Prompt), which are detailed in the following.

To reduce the number of trainable parameters, one intuitive method is to freeze the backbone ϕ and only train the head c . This method is historically the most common fine-tuning procedure, and therefore we use it as a baseline for FedPEFT. However, the adaptation ability of this method is limited, as no adjustment is made to the network representation prior to the final output head. This can be problematic in the presence of a more intense domain shift. Therefore, we consider the following approaches as primary prototypes for FedPEFT:

FedPEFT-Bias. Bias-tuning (Cai et al., 2020) aims to adapt the pre-trained backbone ϕ with only fine-tuning a specific group of parameters, the bias term. In this way, the backbone can be trained with moderate adjustments to prevent damaging the upstream representation.

FedPEFT-Adapter. Instead of directly tuning existing parameters in the backbone like Bias-tuning, Adapter-tuning (Pfeiffer et al., 2020) adds a few parameters called adapters inside the backbone ϕ instead. Usually, adapters will be deployed in each layer of the backbone to perform transformations on different levels of the pre-trained feature while the backbone stays frozen.

FedPEFT-Prompt. Prompt-tuning (Jia et al., 2022) takes a slightly different approach from the other fine-tuning methods. Specifically, it concatenates trainable parameters, called prompt embeddings, to the input embedding and hidden states in each layer.

We illustrate the differences between all baseline and prototype methods in Figure 2. We also provide the algorithm pseudo-code for each prototype in Appendix A.1.

4 EXPERIMENTS

To verify the performance of FedPEFT, we evaluate the server accuracy with each method from three perspectives and aim to answer the following questions:

Communication Analysis: When faced with a limited communication budget, there are several conventional solutions to reduce costs, e.g., sampling fewer clients each round or using a lightweight model. Can FedPEFT outperform other solutions in terms of *communication cost and accuracy*? **(RQ1)**

Algorithm 1 Algorithm of FedPEFT framework

Input: N clients indexed by n , participating-client number M , communication rounds T , transmission parameter set \mathbb{P} of the model where $|\mathbb{P}| \ll |\text{Parameters in } F|$, pre-trained model weights θ .

Server executes:

```

initialize  $F^{(0)}$  and  $\{f_n\}_{n=1}^N$  with  $\theta$ 
for each round  $t = 0, 1, 2, \dots, T - 1$  do
   $\mathbb{S}_t \leftarrow$  (random set of  $M$  clients)
  for each client  $m \in \mathbb{S}_t$  in parallel do
     $\theta^{(t)} \leftarrow$  weights for parameters in  $\mathbb{P}$  of  $F^{(t)}$ 
     $\theta_m^{(t)} \leftarrow \text{ClientUpdate}(m, \theta^{(t)}, \mathbb{P})$ 
     $\theta^{(t)} = \sum_{m=1}^M \frac{|\mathbb{D}_m|}{\sum_{i=1}^M |\mathbb{D}_i|} \theta_m^{(t)}$   $\triangleright \mathbb{D}_m$  is the set of
    the local data on the client with index  $m$ .
   $F^{(t+1)} \leftarrow$  load  $\mathbb{P}$  with  $\theta^{(t)}$ 

```

ClientUpdate (m, θ, \mathbb{P}):

```

 $f_m \leftarrow$  load  $\mathbb{P}$  with  $\theta$ 
perform local training for  $f_m$  with  $\mathbb{D}_m$ 
 $\theta \leftarrow$  weights for parameters in  $\mathbb{P}$  of  $f_m$ 
return  $\theta$  to server

```




Figure 2: Methods to fine-tune the pre-trained model, where x means the input, ϕ means the backbone, ϕ_i means the i -th layer of ϕ with l layers in total, c means the classification head, and h means the hidden states. (a) Fully fine-tune the entire model. (b) Only tune the model head. (c) Only tune the bias term and the model head. (d) Concatenate prompts to the input and hidden states in each layer, and only tune the prompts and the model head. The purple rectangles denote the tunable prompts. (e) Add an adapter to each layer of the backbone, and only tune the adapters and the model head. The orange rectangles denote the tunable adapters.

Capability Analysis: When the communication budget is amply sufficient for all approaches, we want to evaluate the trade-off of training fewer parameters with FedPEFT. Can FedPEFT outperform full fine-tuning and training from scratch within *various federated learning settings and increasing levels of downstream domain gap?* (RQ2)

Robustness Analysis: In a lot of application scenarios, there will be additional challenges for FL, such as privacy-preserving requirements (i.e., differential privacy) and data scarcity (i.e., very small amount of data on each client). We want to evaluate the robustness of each method under such scenarios. Can FedPEFT outperform full fine-tuning in terms of *robustness?* (RQ3)

4.1 EXPERIMENTS DETAILS

Dataset. For our study, we focus on computer vision (CV) applications as our testbed. Therefore, we employ ImageNet-21K (Ridnik et al., 2021) as the pre-training dataset, which is a primary dataset for pre-training models in CV. Then we select two datasets for the downstream tasks that have *increasing degrees of domain gap* compared to ImageNet-21k: CIFAR-100 (Krizhevsky, 2012) and PCam (Veeling et al., 2018). CIFAR-100 contains 50,000 training samples and 10,000 testing samples for 100 classes of nature objects (e.g., sunflowers, buses), which has a small domain gap with ImageNet-21K. PCam is a medical dataset containing 32,768 colored images extracted from histopathologic scans of lymph node sections, which has a much larger domain gap. Each image is annotated with a binary label indicating the presence of metastatic tissue. We sample 20,000 images from the original train set for training and use the entire test set for evaluation.

Experimental Setting. Our default experimental setting is to split the dataset across $N = 64$ clients and sample $M = 8$ clients each round. The global aggregation will be performed after $E = 10$ local epochs. A total of $T = 50$ rounds of communication will be performed. To simulate heterogeneous data, we partition samples in each class to all clients following a Dirichlet distribution, as common in the literature (Mendieta et al., 2022; Acar et al., 2021; Li et al., 2021), with $\alpha = 0.1$ for CIFAR-100 and $\alpha = 0.5$ for PCam based on the class number. Any modifications to this setting in subsequent experiments will be clearly noted.

Implementation Detail. We choose ViT-B (Dosovitskiy et al., 2021) with image size 224 and patch size 16 as our backbone. The backbone is pre-trained on ImageNet-21K (Ridnik et al., 2021), as available in the timm library (Wightman, 2019). The images for the downstream datasets are resized to 224×224 . Images from CIFAR-100 are augmented by random cropping with a padding of 4 and random horizontal flipping, and PCam is augmented only with random horizontal flipping. We perform the experiments on 8 Nvidia RTX A5000 GPUs with a batch size of 64. All reported numbers are run multiple times and averaged. More details about hyperparameter searching can be found in Appendix A. The base hyperparameters for each method are described below.

- Full fine-tuning: We use SGD (Ruder, 2017) optimizer with learning rate 0.001 and weight decay 0.0001.
- Head-tuning: We use one linear layer as the classification head and use the SGD optimizer with learning rate 0.005 and weight decay 0.0001.

Table 1: **Communication analysis.** The communication cost is computed with 4B/parameter, and the accuracy is the final accuracy, i.e., $t = T - 1 = 49$. The first section shows the change of accuracy when decreasing the participating-client number. * indicates the baseline performance with no decrease of participating-client number. The second section shows the change of accuracy when we reward the low communication cost of head-tuning by increasing the participating clients. The third section shows the accuracy when we fully fine-tune a lightweight model, ShuffleNet V2 $0.5\times$ (Ma et al., 2018). The fourth section shows the accuracy of each prototype of FedPEFT.

Model	Method	# Tuned Params \times Clients	Comm. Cost	CIFAR-100	PCam
ViT-B	Full Fine-tuning	$85.88\text{M} \times 8$	2.56GB	92.09*	84.82
ViT-B	Full Fine-tuning	$85.88\text{M} \times 4$	1.28GB	89.84	82.64
ViT-B	Full Fine-tuning	$85.88\text{M} \times 2$	656MB	86.96	81.31
ViT-B	Full Fine-tuning	$85.88\text{M} \times 1$	328MB	73.10	78.29
ViT-B	Head-tuning	$0.08\text{M} \times 8$	2.44MB	72.55	72.38
ViT-B	Head-tuning	$0.08\text{M} \times 64$	19.53MB	75.47	78.24
ShuffleNet	Full Fine-tuning	$0.44\text{M} \times 8$	13.43MB	51.44	72.91
ViT-B	FedPEFT-Bias	$0.18\text{M} \times 8$	5.49MB	91.02	85.29
ViT-B	FedPEFT-Adapter	$0.23\text{M} \times 8$	7.02MB	88.05	79.63
ViT-B	FedPEFT-Prompt	$0.17\text{M} \times 8$	5.19MB	89.90	87.25

- FedPEFT-Bias: There is no additional hyperparameter for Bias. We use SGD optimizer with learning rate 0.01 and weight decay 0.0001.
- FedPEFT-Adapter: We use a Bottleneck Adapter with residual connections as the adapter. We insert the adapter to each layer of the backbone after the feed-forward block following Pfeiffer et al. (2020) with a reduction factor of 8 and the GELU as the activation function. We use SGD optimizer with learning rate 0.005 and weight decay 0.0001.
- FedPEFT-Prompt: We follow the design of VPT-Deep (Jia et al., 2022), and use a prompt length as 10 to prepend prompt tokens to the input and the hidden embedding in each layer. We use SGD optimizer with learning rate 0.01 and weight decay 0.0001.

Table 2: **Total communication cost for reaching a target accuracy.** The number of communication rounds is not fixed in this experiment. The number in the bracket next to the dataset indicates the target accuracy %. FFT(m) means Full Fine-tuning with m participating clients. Bias, Adapter, and Prompt, respectively, denotes the corresponding FedPEFT prototype. Head-tuning and ShuffleNet fail to achieve such targets in most cases, which are not listed here. – indicates that the target accuracy was not achieved at convergence.

Dataset	FFT(8)	FFT(4)	FFT(2)	FFT(1)	Bias	Adapter	Prompt
CIFAR-100 (85)	10.24GB	16.64GB	19.22GB	–	38.34MB	42.12MB	31.14MB
CIFAR-100 (90)	25.6GB	51.2GB	–	–	115.29MB	–	88.23MB
PCam (75)	7.68GB	3.84GB	1.92GB	1.28GB	16.47MB	63.18MB	15.57MB
PCam (80)	30.72GB	14.08GB	12.17GB	–	54.9MB	287.82MB	15.57MB

5 COMMUNICATION ANALYSIS

To verify the effectiveness of FedPEFT and answer the first research question (RQ1, 4), we compare it with three baselines while monitoring the communication budget: a) Full fine-tuning of our default model (ViT-B). We vary the number of participating clients to show different levels of communication requirements. b) Head-tuning. The communication cost of head-tuning is naturally lower than other methods, so we increase the participating clients to make it a stronger baseline. c) Fully fine-tune a light-weighted model (ShuffleNet V2 $0.5\times$ (Ma et al., 2018)) with a similar parameter scale.

As demonstrated in Table 1, all FedPEFT methods achieve better results in many cases compared with other approaches, even with significantly fewer communicated parameters. We find that full fine-tuning needs at least $187\times$ and $477\times$ more parameters to reach and outperform FedPEFT on CIFAR-100, head-tuning lags behind most other approaches, and the ShuffleNet model only

Table 3: **Capability analysis among different federated learning settings on CIFAR-100**. Bold-style shows the best performance among all methods or among prototypes in FedFEFT.

Client	Method	# Tuned Params/Client	Homogeneous	Heterogeneous
$N = 16, M = 16$	Scratch	85.88M	38.44	35.72
	Full Fine-tuning	85.88M	93.70	93.50
	Head-tuning	0.08M	78.11	77.59
	FedPEFT-Bias	0.18M	91.89	90.25
	FedPEFT-Adapter	0.23M	90.21	88.77
	FedPEFT-Prompt	0.17M	92.09	90.37
$N = 16, M = 2$	Full Fine-tuning	85.88M	93.32	87.01
	Head-tuning	0.08M	76.65	62.80
	FedPEFT-Bias	0.18M	91.35	86.18
	FedPEFT-Adapter	0.23M	89.48	80.08
	FedPEFT-Prompt	0.17M	91.60	85.54
$N = 64, M = 64$	Full Fine-tuning	85.88M	93.66	92.81
	Head-tuning	0.08M	78.45	75.51
	FedPEFT-Bias	0.18M	92.71	91.71
	FedPEFT-Adapter	0.23M	90.50	89.26
	FedPEFT-Prompt	0.17M	91.87	90.96
$N = 64, M = 8$	Full Fine-tuning	85.88M	93.50	92.09
	Head-tuning	0.08M	77.59	72.55
	FedPEFT-Bias	0.18M	92.49	91.02
	FedPEFT-Adapter	0.23M	90.39	88.05
	FedPEFT-Prompt	0.17M	92.00	89.90

achieves 57% and 84% of accuracy on CIFAR-100 and PCAM with **2.4** \times the communication cost compared with FedPEFT-Bias.

In Table 2, we also report the total communication cost to reach reasonable target accuracies in each dataset. All FedPEFT prototypes only require megabytes level communication, while full fine-tuning requires gigabytes level communication, indicating the efficiency of FedPEFT. For the inter-prototype comparison, FedPEFT-Prompt stands out for its fast convergence speed. We provide further discussions on the performance of each prototype in Section 6.

6 CAPABILITY ANALYSIS

To study and understand our second research question (RQ2, 4), we systemically perform experiments on CIFAR-100 across different federated learning scenarios by varying client status and data distribution (Section 6.1) and analyze the impact of the domain gap between the model pre-training dataset and the dataset for FL (Section 6.2).

6.1 CAPABILITY WITH DIFFERENT FEDERATED LEARNING SETTINGS

In application scenarios, the setting of federated learning can vary substantially. It is important to show the capability to maintain high performance in diverse settings. In Table 3, we present results for all approaches with different client availability ratios and data distributions and draw the following conclusions from the experiments:

First, we see that fine-tuning on the pre-trained model shows a significant improvement over training from scratch, especially in heterogeneous scenarios. This finding is in agreement with other very recent works (Chen et al., 2022b; Nguyen et al., 2022), which note the stabilization effect of pre-trained initialization in federated optimization. When only fine-tuning the head, the performance is still much better than training the entire model from scratch (Scratch) but remains low in comparison to other methods across all settings. We again find that head-tuning simply lacks adaptation ability, holding too closely to the upstream representation.

On the other hand, we find that FedPEFT achieves comparable results to full fine-tuning with less than 0.3% of the trainable parameters. This ability to maintain accuracy performance in various scenarios is crucial for FL, as oftentimes, the exact setting and distributions are not known ahead of time. With this in mind, we further investigate the viability of FedPEFT as a practical FL framework in some more extreme situations in Section 7.

6.2 CAPABILITY WITH DOMAIN GAP

In Table 4, we present the performance of all approaches with an increasing degree of domain gap compared to the ImageNet-21k pre-training dataset. Interestingly, full fine-tuning falls further behind as the data domain gap widens in the PCam scenario, largely unable to keep up with FedPEFT despite requiring a massive communication budget. This phenomenon when a pre-trained model meets out-of-domain data has been studied under centralized settings (Shen et al., 2021). It was found that the pre-trained upstream representations are still meaningful even with a domain gap. Therefore, fully fine-tuning the backbone with out-of-domain data can damage the high-level semantics inside these upstream representations due to overfitting, especially when the data size is small. This is particularly relevant in FL, where overfitting and subsequent client drift (Varno et al., 2022; Mendieta et al., 2022; Karimireddy et al., 2020) are prone to occur.

On the opposite end of the spectrum, by not fine-tuning the backbone at all, head-tuning maintains similar accuracy despite the domain gap. This shows the robustness of the pre-trained high-level semantics across domains, supporting the conclusion that *there is meaningful high-level semantics inside of the upstream representations*.

Still, the tight restriction on head-tuning is perhaps a bit too far, as the accuracy on both datasets is still low overall. Between the two extremes of head-tuning and full fine-tuning, FedPEFT approaches may be able to suitably adapt the upstream representations without excessively damaging them. Specifically, FedPEFT-Bias operates with parameter-level control for each parameter pair containing weight and bias terms. The representation can then preserve the high-level semantics by freezing the weight term (maintaining the direction in feature space) and still adapting via the bias term (shifting in the feature space). FedPEFT-Adapter and FedPEFT-Prompt have slightly different mechanisms (layer-level), controlling of the backbone by transforming the intermediate hidden representations via adapters and prompts. Of these approaches, FedPEFT-Prompt is the most stable under domain gap, surpassing full fine-tuning by 2.4% on PCam. Overall, we hypothesize that more fine-tuning freedom will be better when the domain gap is minor, but moderate fine-tuning is needed to maintain, as well as control, the high-level semantics when the domain gap is large.

7 ROBUSTNESS ANALYSIS

In this section, we further investigate our third research question (**RQ3**) in two critical FL scenarios.

Differential Privacy. A fundamental property of federated learning is privacy protection. However, various works (Huang et al., 2021; Hatamizadeh et al., 2022) have demonstrated how the client data can be reconstructed from the raw gradient updates received by the server in some scenarios. To protect client data privacy from such attacks, differential privacy (DP) (Kairouz et al., 2021; Chamikara et al., 2020; Dwork, 2008; Dwork et al., 2014) has become standard practice. Therefore, we first study FedPEFT and other baselines under DP.

To integrate DP, we apply a Gaussian mechanism within the local optimization of each iteration (Dwork et al., 2014) with $\epsilon = 5$ and $\delta = 0.001$. We maintain the remaining FL settings as described in Section 4.1, and show the results in Table 5. Interestingly, when comparing all methods, full fine-tuning experiences the sharpest drop with DP. This causes its accuracy to fall lower than all the FedPEFT prototypes. To understand this effect, we note that DP applies noise to all trainable parameter gradients. Full fine-tuning, therefore, requires such noise on all model parameters, resulting in a more pronounced negative effect on final performance. On the other hand, the others fine-tuning methods maintain some part of the backbone frozen and have significantly fewer trainable parameters on which adding noise is necessary, limiting the performance drop. Overall, FedPEFT allows for

Table 4: **Capability analysis when there is a domain gap between the pre-training dataset and the dataset for FL.** We use full fine-tuning as a baseline, and the number in brackets shows the performance difference compared with the baseline.

Method	CIFAR-100	PCam
Full Fine-tuning	92.09	84.82
Head-tuning	72.55 (-19.54)	72.38 (-12.44)
FedPEFT-Bias	91.02 (-1.07)	84.98 (+0.47)
FedPEFT-Adapter	88.05 (-4.04)	79.63 (-5.19)
FedPEFT-Prompt	89.90 (-2.19)	87.25 (+2.43)

Table 5: **Robustness analysis for privacy-preserving.** The red number indicates the performance difference when DP is applied.

Method	w/o DP	w/ DP
Full Fine-tuning	92.09	77.61 (-14.48)
Head-tuning	72.55	62.20 (-10.35)
FedPEFT-Bias	91.02	84.98 (-6.04)
FedPEFT-Adapter	88.05	79.05 (-9.00)
FedPEFT-Prompt	89.90	78.35 (-11.55)

stronger accuracy in DP-enabled federated systems than even full fine-tuning while still maintaining extremely low communication needs.

Data Scarcity. We explore another common yet challenging robustness condition in FL; that is when very little data is available on individual clients. Such data scarcity scenarios are even a tricky problem in centralized training. For example, (Wang et al., 2019) show that fewer training data will incur damage to the pre-trained representation due to overfitting. In our evaluation for FL, we reduce the total sample number K to 1000, 1500, and 2000. As shown in Table 6, we find that FedPEFT outperforms full fine-tuning and head-tuning under such low-data scenarios, further revealing its capability to appropriately adapt pre-trained representations to the FL task at hand.

8 PRE-TRAINING ABLATIONS

In the experiments above, we maintain a consistent backbone (ViT-B). However, since model size requirements of different federated systems may vary, we further investigate the use of ViT-S with FedPEFT and all baselines. Additionally, we study the effect of pre-training dataset size on the final performance of all approaches. CIFAR-100 is used for the experiments.

Impact of Model Size. We replace ViT-B with a smaller version, ViT-S. As expected, most methods experience a slight performance drop. Nonetheless, FedPEFT methods maintain strong performance even with smaller models. Surprisingly, the performance of head-tuning is not damaged by reducing the model size, but actually increases slightly. We interpret it as due to the fixed representation of head-tuning. Representations from a pre-trained model can have similar semantics whatever the model size, thus the performance of head-tuning is stable.

Impact of Pre-training Dataset. We also evaluate the performance when a backbone is trained from ImageNet-21K and ImageNet-1K (Russakovsky et al., 2015). Interestingly, the performance stays relatively consistent, with Head-tuning and FedPEFT-Adapter achieving higher accuracy with the ImageNet-1K models. This phenomenon indicates that the performance of fine-tuning is not always proportional to the data size, but overall FedPEFT performance is consistent with different scales of the pre-training dataset.

9 CONCLUSION

In this paper, we introduce FedPEFT, a new federated learning framework leveraging strong pre-trained models and massively reducing communication costs. We integrate three effective prototypes within the FedPEFT framework: Bias, Adapter, and Prompt. With a thorough empirical study, we then evaluate FedPEFT and other baselines in three key areas: communication, capability, and robustness. We find FedPEFT to be a promising approach for practical FL systems, capable of handling many of the harsh conditions in FL while alleviating the critical communication bottleneck. As a general framework, FedPEFT can also be leveraged in application domains other than computer vision, which we leave for future work. We hope this work can inspire new perspectives for improving the communication efficiency in federated learning through the combined innovation of strong pre-trained models and parameter-efficient fine-tuning methodologies.

Table 6: **Robustness analysis for data scarcity.** K indicates the total sample number of all clients.

Data	Method	Server Acc (%) \uparrow
$K = 1000$	Full Fine-tuning	66.52
	Head-tuning	52.13
	FedPEFT-Bias	76.40
	FedPEFT-Adapter	71.34
$K = 1500$	FedPEFT-Prompt	63.77
	Full Fine-tuning	67.47
	Head-tuning	56.52
	FedPEFT-Bias	81.14
$K = 2000$	FedPEFT-Adapter	76.91
	FedPEFT-Prompt	71.94
	Full Fine-tuning	77.67
	Head-tuning	60.15
$K = 2000$	FedPEFT-Bias	83.83
	FedPEFT-Adapter	79.22
	FedPEFT-Prompt	76.89

Table 7: Ablation study on network backbone and pre-training dataset (**PD**).

Method	Backbone	PD	Server Acc (%)
Full Fine-tuning	ViT-B	ImageNet-21K	92.09
	ViT-B	ImageNet-1K	92.09
	ViT-S	ImageNet-21K	89.33
Head-tuning	ViT-B	ImageNet-21K	72.55
	ViT-B	ImageNet-1K	80.78
	ViT-S	ImageNet-21K	74.40
FedPEFT-Bias	ViT-B	ImageNet-21K	91.02
	ViT-B	ImageNet-1K	90.73
	ViT-S	ImageNet-21K	88.40
FedPEFT-Adapter	ViT-B	ImageNet-21K	87.99
	ViT-B	ImageNet-1K	88.69
	ViT-S	ImageNet-21K	84.15
FedPEFT-Prompt	ViT-B	ImageNet-21K	89.63
	ViT-B	ImageNet-1K	89.29
	ViT-S	ImageNet-21K	87.39

REFERENCES

- Durmus Alp Emre Acar, Yue Zhao, Ramon Matas Navarro, Matthew Mattina, Paul N Whatmough, and Venkatesh Saligrama. Federated learning based on dynamic regularization. *arXiv preprint arXiv:2111.04263*, 2021.
- Rishi Bommasani et al. On the Opportunities and Risks of Foundation Models, July 2022. arXiv:2108.07258 [cs].
- Han Cai, Chuang Gan, Ligeng Zhu, and Song Han. TinyTL: Reduce Memory, Not Parameters for Efficient On-Device Learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 11285–11297. Curran Associates, Inc., 2020.
- M. A. P. Chamikara, P. Bertok, I. Khalil, D. Liu, S. Camtepe, and M. Atiquzzaman. Local Differential Privacy for Deep Learning. *IEEE Internet of Things Journal*, 7(7):5827–5842, July 2020. ISSN 2327-4662, 2372-2541. doi: 10.1109/JIOT.2019.2952146. arXiv:1908.02997 [cs].
- Hao Chen, Ran Tao, Han Zhang, Yidong Wang, Wei Ye, Jindong Wang, Guosheng Hu, and Marios Savvides. Conv-Adapter: Exploring Parameter Efficient Transfer Learning for ConvNets, August 2022a. arXiv:2208.07463 [cs].
- Hong-You Chen, Cheng-Hao Tu, Ziwei Li, Han-Wei Shen, and Wei-Lun Chao. On Pre-Training for Federated Learning, June 2022b. arXiv:2206.11488 [cs].
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale, June 2021. arXiv:2010.11929 [cs].
- Cynthia Dwork. Differential Privacy: A Survey of Results. In Manindra Agrawal, Dingzhu Du, Zhenhua Duan, and Angsheng Li (eds.), *Theory and Applications of Models of Computation*, Lecture Notes in Computer Science, pp. 1–19, Berlin, Heidelberg, 2008. Springer. ISBN 978-3-540-79228-4. doi: 10.1007/978-3-540-79228-4_1.
- Cynthia Dwork, Aaron Roth, and others. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014. Publisher: Now Publishers, Inc.
- Jenny Hamer, Mehryar Mohri, and Ananda Theertha Suresh. FedBoost: A Communication-Efficient Algorithm for Federated Learning. In *Proceedings of the 37th International Conference on Machine Learning*, pp. 3973–3983. PMLR, November 2020. ISSN: 2640-3498.
- Ali Hatamizadeh, Hongxu Yin, Holger Roth, Wenqi Li, Jan Kautz, Daguang Xu, and Pavlo Molchanov. GradViT: Gradient Inversion of Vision Transformers, March 2022. arXiv:2203.11894 [cs].
- Yangsibo Huang, Samyak Gupta, Zhao Song, Kai Li, and Sanjeev Arora. Evaluating Gradient Inversion Attacks and Defenses in Federated Learning, November 2021. arXiv:2112.00059 [cs] version: 1.
- Menglin Jia, Luming Tang, Bor-Chun Chen, Claire Cardie, Serge Belongie, Bharath Hariharan, and Ser-Nam Lim. Visual Prompt Tuning, July 2022. arXiv:2203.12119 [cs].
- Peter Kairouz et al. Advances and Open Problems in Federated Learning. *arXiv:1912.04977 [cs, stat]*, March 2021. arXiv: 1912.04977.
- Sai Praneeth Karimireddy, Satyen Kale, Mehryar Mohri, Sashank Reddi, Sebastian Stich, and Ananda Theertha Suresh. Scaffold: Stochastic controlled averaging for federated learning. In *International Conference on Machine Learning*, pp. 5132–5143. PMLR, 2020.
- Jakub Konečný, H. Brendan McMahan, Felix X. Yu, Peter Richtárik, Ananda Theertha Suresh, and Dave Bacon. Federated Learning: Strategies for Improving Communication Efficiency, October 2017. arXiv:1610.05492 [cs].

- Alex Krizhevsky. Learning Multiple Layers of Features from Tiny Images. *University of Toronto*, 2012.
- Qinbin Li, Bingsheng He, and Dawn Song. Model-Contrastive Federated Learning, March 2021. arXiv:2103.16257 [cs].
- Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated Optimization in Heterogeneous Networks, April 2020. arXiv:1812.06127 [cs, stat].
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin Raffel. Few-Shot Parameter-Efficient Fine-Tuning is Better and Cheaper than In-Context Learning, August 2022. arXiv:2205.05638 [cs].
- Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 116–131, 2018.
- H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-Efficient Learning of Deep Networks from Decentralized Data, February 2017. arXiv:1602.05629 [cs].
- Matias Mendieta, Taojiannan Yang, Pu Wang, Minwoo Lee, Zhengming Ding, and Chen Chen. Local Learning Matters: Rethinking Data Heterogeneity in Federated Learning. *arXiv:2111.14213 [cs]*, March 2022. arXiv: 2111.14213.
- John Nguyen, Kshitiz Malik, Maziar Sanjabi, and Michael Rabbat. Where to Begin? Exploring the Impact of Pre-Training and Initialization in Federated Learning, June 2022. arXiv:2206.15387 [cs].
- Junting Pan, Ziyi Lin, Xiatian Zhu, Jing Shao, and Hongsheng Li. Parameter-Efficient Image-to-Video Transfer Learning, June 2022. arXiv:2206.13559 [cs].
- Jonas Pfeiffer, Andreas Rücklé, Clifton Poth, Aishwarya Kamath, Ivan Vulić, Sebastian Ruder, Kyunghyun Cho, and Iryna Gurevych. AdapterHub: A Framework for Adapting Transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 46–54, Online, 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-demos.7.
- Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. ImageNet-21K Pretraining for the Masses, August 2021. arXiv:2104.10972 [cs].
- Sebastian Ruder. An overview of gradient descent optimization algorithms, June 2017. arXiv:1609.04747 [cs].
- Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Zhiqiang Shen, Zechun Liu, Jie Qin, Marios Savvides, and Kwang-Ting Cheng. Partial Is Better Than All: Revisiting Fine-tuning Strategy for Few-shot Learning, February 2021. arXiv:2102.03983 [cs].
- Ananda Theertha Suresh, Felix X. Yu, Sanjiv Kumar, and H. Brendan McMahan. Distributed Mean Estimation with Limited Communication, September 2017. arXiv:1611.00429 [cs].
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Farshid Varno, Marzie Saghay, Laya Rafiee, Sharut Gupta, Stan Matwin, and Mohammad Havaei. Minimizing Client Drift in Federated Learning via Adaptive Bias Estimation. *arXiv:2204.13170 [cs]*, April 2022. arXiv: 2204.13170.

Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation Equivariant CNNs for Digital Pathology. June 2018. [_eprint: 1806.03962](#).

Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020a.

Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H. Vincent Poor. Tackling the Objective Inconsistency Problem in Heterogeneous Federated Optimization, July 2020b. [arXiv:2007.07481](#) [cs, stat].

Yaqing Wang, Quanming Yao, James Kwok, and Lionel M. Ni. Generalizing from a Few Examples: A Survey on Few-Shot Learning. April 2019. [doi: 10.48550/arXiv.1904.05046](#).

Ross Wightman. PyTorch Image Models, 2019. Publication Title: GitHub repository.

Mikhail Yurochkin, Mayank Agarwal, Soumya Ghosh, Kristjan Greenewald, Trong Nghia Hoang, and Yasaman Khazaeni. Bayesian Nonparametric Federated Learning of Neural Networks, May 2019. [arXiv:1905.12022](#) [cs, stat].

APPENDIX

A IMPLEMENT DETAILS

A.1 PSEUDO-CODE FOR EACH PROTOTYPE

We provide a detailed pseudo-code for each prototype in a Pytorch style. The key code initialization and forward pass are given, which can be found in Algorithm 2, 3, 4.

Table 8: **Optimal learning rate for each setting.** FFT: Full Fine-tuning, HT: Head-tuning. Homo: homogeneous data, α : parameter for the Dirichlet distribution.

Setting	FFT	HT	Bias	Adapter	Prompt
$N = 64, M = 8, \alpha = 0.1$	5e-4	1e-2	1e-2	5e-3	1e-2
$N = 64, M = 8, \text{homo}$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 64, M = 64, \alpha = 0.1$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 64, M = 64, \text{homo}$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 16, M = 16, \alpha = 0.1$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 16, M = 16, \text{homo}$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 16, M = 2, \alpha = 0.1$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 16, M = 2, \text{homo}$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 64, M = 8, \alpha = 0.1, \text{DP}$	1e-4	5e-4	1e-3	5e-4	3e-4
$N = 64, M = 8, \alpha = 0.1, K = 1000$	1e-3	5e-3	1e-2	1e-3	1e-2
$N = 64, M = 8, \alpha = 0.1, K = 1500$	1e-3	5e-3	1e-2	5e-3	1e-2
$N = 64, M = 8, \alpha = 0.1, K = 2000$	1e-3	5e-3	1e-2	5e-3	1e-2

Algorithm 2 Pseudo-code for FedPEFT-Bias

```

class FedPEFT_Bias:
    def __init__(basic_model):
        self.basic_model = basic_model
        # freeze backbone except for the bias terms
        for p in self.basic_model.backbone.parameters():
            if p is not a bias term:
                p.requires_grad = False
    def forward(x):
        return self.basic_model(x)

```

A.2 HYPERPARAMETER SEARCHING

The most crucial hyperparameter in our experiments is the learning rate. We search the hyperparameter with several options in a grid-search-like fashion. The initial options for learning rate are $\{1e-4, 5e-4, 1e-3, 5e-3, 1e-2, 5e-2\}$. The specific optimal learning rate for each setting is shown in Table 8. Besides, some compared method’s optimal learning rate is listed here:

- ShuffleNet in CIFAR-100: 5e-2
- ShuffleNet in PCam: 1e-1
- Training from Scratch in CIFAR-100 with homogeneous data: 1e-2
- Training from Scratch in CIFAR-100 with heterogeneous data: 1e-2

Algorithm 3 Pseudo-code for FedPEFT-Adapter

```

class FedPEFT_Adapter:
    def __init__(basic_model, reduction_factor):
        self.basic_model = basic_model
        num_layers = layer number of basic_model.backbone
        # build adapters
        self.adapter_downsample = nn.Linear(
            embed_dim,
            embed_dim // reduction_factor
        )
        self.adapter_upsample = nn.Linear(
            embed_dim // reduction_factor,
            embed_dim
        )
        self.adapter_act_fn = nn.functional.gelu
        # freeze backbone
        for p in self.basic_model.backbone.parameters():
            p.requires_grad = False
    def forward(x):
        # get the embedding of x
        x = embedding(x)
        for i in range(num_layers):
            # forward normal blocks
            block = self.basic_model.blocks[i]
            x = x + block.drop_path1(block.ls1(block.attn(block.norm1(x))))
            h = x
            x = block.drop_path2(block.ls2(block.mlp(block.norm2(x))))
            # adapter
            adpt = self.adapter_downsample(x)
            adpt = self.adapter_act_fn(adpt)
            adpt = self.adapter_upsample(adpt)
            x = adpt + x

            x = x + h

        x = self.basic_model.norm(x)
        x = self.basic_model.head(x)
        return x

```

Algorithm 4 Pseudo-code for FedPEFT-Prompt

```

class FedPEFT_Prompt:
    def __init__(basic_model, prompt_num):
        self.basic_model = basic_model
        num_layers = layer number of basic_model.backbone
        # build prompts
        prompt_token = nn.Parameter(
            torch.zeros(num_layers, prompt_num, prompt_dim)
        )
        # freeze backbone
        for p in self.basic_model.backbone.parameters():
            p.requires_grad = False
    def forward(x):
        for i in range(len(self.basic_model.blocks)):
            torch.cat(prompt_token, x)
            x = self.basic_model.blocks[i](x)
            x = remove_prompt_tokens
        x = self.basic_model.norm(x)
        x = self.basic_model.head(x)

```

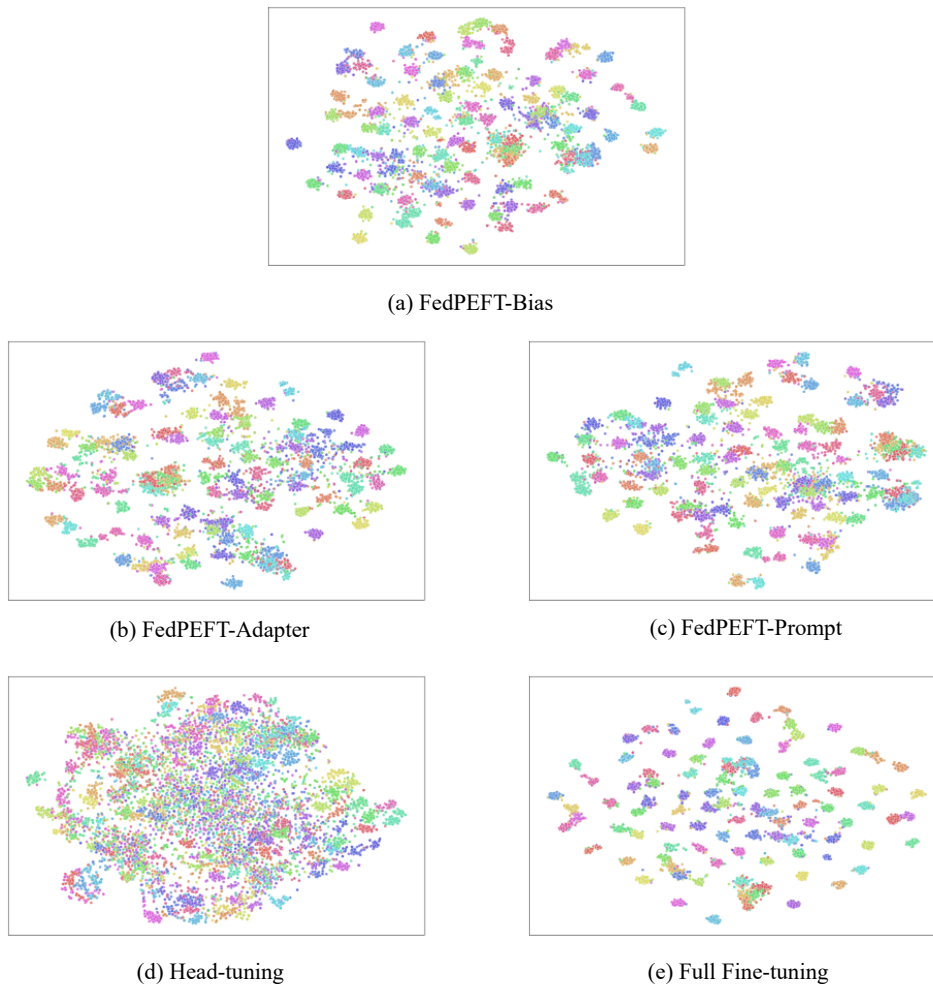


Figure 3: t-SNE Plots on CIFAR-100.

B VISUALIZATION

We visualize the t-SNE (Van der Maaten & Hinton, 2008) plots of the extracted feature from the server model trained with Full fine-tuning, Head-tuning, FedPEFT-Bias, FedPEFT-Adapter, and FedPEFT-Prompt under the default setting described in Section 4.

We make a few observations here. First, we note that FedPEFT-Bias, FedPEFT-Adapter, and FedPEFT-Prompt are able to reasonably adapt the backbone representation to the target data. With head-tuning, the backbone is kept frozen, and therefore we do not see the same clear separation of each class as with the FedPEFT prototypes. Nonetheless, the frozen backbone does seem to contain relevant semantics, with visible clusters throughout, particularly around the perimeter regions. These observations are consistent with our results and analysis in Section 6.