

FLoRA: Enhancing Vision-Language Models with Parameter-Efficient Federated Learning

Duy Phuong Nguyen¹, J. Pablo Muñoz², and Ali Jannesari¹

¹Department of Computer Science, Iowa State University, Ames, IA, USA

²Intel Corporation, Santa Clara, CA, USA

¹{dphuong, jannesari}@iastate.edu

²pablo.munoz@intel.com

Abstract—In the rapidly evolving field of artificial intelligence, multimodal models, e.g., integrating vision and language into visual-language models (VLMs), have become pivotal for many applications, ranging from image captioning to multimodal search engines. Among these models, the Contrastive Language–Image Pre-training (CLIP) model has demonstrated remarkable performance in understanding and generating nuanced relationships between text and images. However, the conventional training of such models often requires centralized aggregation of vast datasets, posing significant privacy and data governance challenges. To address these concerns, this paper proposes a novel approach that leverages Federated Learning and parameter-efficient adapters, i.e., Low-Rank Adaptation (LoRA), to train VLMs. This methodology preserves data privacy by training models across decentralized data sources and ensures model adaptability and efficiency through LoRA’s parameter-efficient fine-tuning. Our approach accelerates training time by up to $34.72\times$, and requires $2.47\times$ less memory usage than full fine-tuning.

I. INTRODUCTION

In recent years, the fields of computer vision and natural language processing have witnessed remarkable advancements fueled by the power of deep learning models [1]–[3]. Vision-language models (VLMs) [4], in particular, have gained significant attention due to their ability to bridge the gap between visual and textual information, enabling a wide range of applications such as image captioning, visual question answering, and text-to-image synthesis. However, the success of VLMs heavily relies on the availability of vast amounts of labeled data for training. Collecting and annotating large-scale datasets can be time-consuming, resource-intensive, and potentially privacy-sensitive. Additionally, centralizing all data in a single location for training poses challenges such as data privacy, communication costs, and scalability. Federated learning (FL) has emerged as a promising approach to address these limitations. Federated learning allows training models across a distributed network of devices, where each device holds its data and collaboratively learns a shared model while keeping the data locally. This decentralized approach respects data privacy and reduces the communication overhead in transmitting raw data to a centralized server.

This paper delves into the domain of FL with a particular focus on fine-tuning VLMs. By adapting the fine-tuning techniques to a federated setting, we seek to harness the dual

strengths of vision-language pre-training and distributed learning. Our method utilizes the LoRA (Low-Rank Adaptation) [5] technique to fine-tune the CLIP model, aiming to update only a small subset of the model’s parameters while keeping the bulk of the pre-trained weights intact.

Figure 1 illustrates the overarching process. We investigate efficient strategies for federated optimization, employing LoRA as a mechanism to perform fine-tuning that respects the limitations of data privacy and the necessity for reduced communication overhead. Our approach relies on the inherent alignment capabilities of the CLIP model, fine-tuning it to enhance its performance on tasks involving federated datasets. We investigate novel techniques for collaborative model training and federated optimization strategies that enable VLMs to learn from distributed data sources while preserving privacy and minimizing communication costs.

Through rigorous experimental evaluations, we strive to exhibit the effectiveness of federated fine-tuning, positing that our methodology achieves competitive performance on several benchmarks. The implications of our work extend to the advancement of FL techniques and open new pathways for the development of collaborative, privacy-aware VLMs.

This paper embarks on an in-depth examination of federated fine-tuning learning’s potential to revolutionize VLMs, addressing critical concerns such as data privacy, communication efficacy, and scalability in decentralized environments.

Our contributions are outlined as follows:

- We pioneer the application of LoRA to VLMs within a FL framework, particularly focusing on the fine-tuning of the CLIP model’s text encoder.
- Our adaptation of the LoRA adapter into the CLIP model’s text encoder projection layers not only refines the model’s specificity to textual nuances but also results in up to a $4766\times$ reduction in communication overhead during federated training—a significant stride towards practical FL deployment.
- Our methodology, FLoRA, has been rigorously evaluated across datasets with varying characteristics, leading to performance gains of up to 30% in accuracy metrics compared to traditional federated learning baselines, underscoring the robustness and versatility of our approach.




Fig. 1: This schematic illustrates the critical steps in federated learning. In each round of federated learning, clients first download the transferred model from the server and independently train models on their own data (step ①). Then, locally-trained models are then sent to a central server (step ②). After that, the server aggregates these models to form a average global model (step ③), which is subsequently redistributed to the clients for further training or inference (step ④).

Algorithm 1: FLORA Client Update

Input : N : number of clients, ρ : client sample ratio, \mathcal{L} : Loss function, δ : local learning rate, Θ^t previous global model
Output: local model $\Theta_1^{t+1}, \dots, \Theta_N^{t+1}$
 $\Theta^t \rightarrow$ local models
 $\mathcal{D}_i \rightarrow$ local data of client $i \in [1, N]$
 Clients incorporate Θ^t into Θ_{CLIP}
for CLIENT $i \in [1, \rho \times N]$ **do**
 \triangleright **Local training**
 $\Theta_i^{t+1} \leftarrow \Theta_i^t - \delta \nabla [\mathcal{L}(\Theta_i^t; \mathcal{D}_i)]$
 \triangleright **Uploading**
 Client i sends Θ_i^{t+1} to the server

- Our research includes an extensive ablation study, meticulously investigating the optimal integration points and methods for the LoRA adapter within the CLIP model. This investigation is pivotal in optimizing performance within the FL context, ensuring that our approach enhances efficiency and maximizes the efficacy of the model’s learning capabilities.

In subsequent sections, we outline related work, delineate our methodology, present experimental results, and discuss the broader impact and future trajectory of federated fine-tuning for VLMs.

II. RELATED WORK

The synthesis of FL and VLMs has garnered substantial attention in recent times. This section encapsulates seminal contributions to these fields, laying the groundwork for our investigation into federated fine-tuning of VLMs.

Algorithm 2: FLORA Server Update

Input : T : number of communication rounds, N : number of clients, ρ : client sample ratio
Output: Global model Θ^t
if $t == 0$ **then**
 \triangleright **Initial stage**
 Initialize Θ^0
 Server receives Θ_i^t from client $i \in [1, N \times \rho]$
for ITERATION $t = 1, \dots, T$ **do**
 \triangleright **Averaging weights**
 $\Theta^t \leftarrow \sum_{i=1}^{[N \times \rho]} \frac{|\mathcal{D}_i|}{\sum_{j=1}^{[N \times \rho]} |\mathcal{D}_j|} \Theta_i^t$
 \triangleright **Downloading**
 Server sends Θ^t to all clients $i \in [1, N]$

A. Federated Learning

Federated learning, introduced by McMahan et al. [6], has emerged as a powerful approach for training models collaboratively across distributed devices while preserving data privacy. Various studies have explored the application of FL in different domains, including computer vision and natural language processing. McMahan et al. [6] proposed federated averaging (FedAvg), a communication-efficient algorithm for training deep learning models in a decentralized manner. Li et al. [7] extended FL to handle large-scale datasets through hierarchical aggregation and compression techniques.

B. Vision-Language Models

Vision-language models serve as a conduit between the visual and linguistic realms, paving the way for multifaceted applications. Early innovations were predominantly in image captioning, as exemplified by the work of Vinyals et al. [8] and Xu et al. [9]. Subsequent models like ViLBERT [10] and LXMERT [11] capitalized on extensive pre-training to set new benchmarks across various VLMs.

C. Contrastive Language Image Pre-training (CLIP)

The CLIP model, unveiled by OpenAI [4], marked a paradigm shift by leveraging contrastive pre-training over a vast corpus of image-text pairs. It encodes images and text into a shared vector space, facilitating versatile zero-shot capabilities. This property has allowed CLIP to deliver impressive results across diverse datasets, even outperforming models fine-tuned on specific benchmarks such as ImageNet and its variants [12].

D. Federated Learning for Vision-Language Models

Bridging FL with VLMs remains an emergent research domain. Notable are the work by Guo et al. [13], [14], which proposes a privacy-centric federated framework for image classification. Yet, the fusion of FL with the complexity of vision-language tasks necessitates further exploration. Qiu et al [15] proposed FedTPG, a method that tailors pre-trained VLMs to FL by employing task-specific prompt generation. This technique enhances the relevance of the generated prompts to

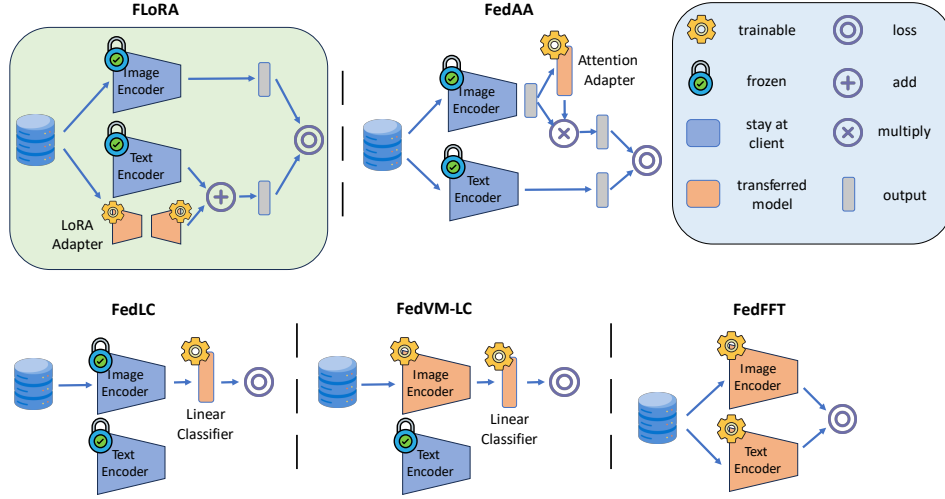


Fig. 2: Different methods for local model update. The clients download the transferred model from the server, train locally with the local data, and upload it to the server. **FLORA**: only transfer LoRA Adapter; **FedAA**: only transfer Attention Adapter; **FedLC**: only transfer Linear Classifier; **FedVM-LC**: only transfer Vision Model and Linear Classifier; **FedFFT**: transfer the whole CLIP model (full fine-tuning). Our method (**FLORA**) is highlighted in the green box. The figure is best viewed in color.

each client’s data, thereby improving the model’s performance on localized tasks. However, FedTPG’s prompt generation process still relies on the exchange of prompt-related parameters, which could be optimized to reduce communication costs further. Another one by Lu et al. [16] demonstrates the feasibility of adapting CLIP to federated settings, yet it leaves room for efficiency improvements, particularly in communication overhead and model adaptability. Our work extends these efforts, integrating FL with the fine-tuning of VLMs, eschewing traditional prompt learning for a focused adaptation using LoRA. The process is visualized in Figure 1, marking a significant stride toward privacy-conscious, efficient, and scalable vision-language model training, especially in terms of reducing the bandwidth required for transmitting updates and ensuring the model remains flexible to the nuances of distributed datasets.

1) *Client-Side Training*: In our framework, client-side training involves local updates using a LoRA-adapted CLIP model, where Adam optimizer [17] aids in gradient-based optimization. This process, outlined in Algorithm 2, ensures that updates are focused and computationally frugal.

2) *Server Aggregation*: Our approach employs a server that orchestrates the aggregation of updates from clients, utilizing a weighted averaging mechanism akin to FedAvg [6]. This strategy balances the contributions across diverse datasets, promoting an egalitarian learning process. Details of this implementation are provided in Algorithm 1.

E. Low-Rank Adaptation for Vision-Language Models

Low-Rank Adaptation (LoRA) [5] presents a novel and efficient approach to fine-tuning deep learning models, particularly in the context of VLMs. This technique, by focus-




Fig. 3: Schematic diagram of LoRA finetuning Mechanism. This figure demonstrates the LoRA fine-tuning process applied to a pre-trained model. Input x is processed through low-rank matrices A and B , which are the core components of the LoRA approach. These matrices modify the pre-trained weights W by a scaling factor α , allowing for precise adjustments to the model without the need to retrain the entire network. This targeted fine-tuning strategy leads to efficient adaptation and output generation.

ing on modifying only a small subset of model parameters, significantly reduces the computational overhead and memory requirements traditionally associated with training large neural networks. In the realm of vision-language tasks, where models must capture complex, multimodal relationships, LoRA offers a path to enhance model adaptability and performance without necessitating extensive computational resources.

LoRA operates by introducing trainable low-rank matrices into specific layers of a pre-trained model. These matrices serve as adapters, adjusting the model’s behavior to better suit the task at hand with minimal interference to the pre-trained parameters. This approach is particularly beneficial in FL environments, where computational resources are dis-

persed and data privacy is of paramount importance. Figure 3 demonstrates the mechanism of LoRA. This methodology is premised on the addition of low-rank matrices, denoted as \mathbf{A} and \mathbf{B} , which are applied to the input x in conjunction with the original weight matrix \mathbf{W} . The key operation can be succinctly represented as

$$\mathbf{W}' = (\mathbf{W} + \alpha \cdot \mathbf{A} \cdot \mathbf{B}) \quad (1)$$

where \mathbf{W}' is the new updated weight matrix, while α , the scaling factor and the rank r are our hyperparameters. The rank r is a hyperparameter that controls the inner dimension of the matrices \mathbf{A} and \mathbf{B} . Then the output can be obtained as: $x \cdot \mathbf{W}'$ during inference. This formulation effectively separates the low-rank matrices from the original weights, facilitating a unique adaptation mechanism that is both memory-efficient and flexible. Note that if $\mathbf{W} \in \mathbb{R}^{d \times k}$, then $\mathbf{A} \in \mathbb{R}^{d \times r}$ and $\mathbf{B} \in \mathbb{R}^{r \times k}$.

The adoption of LoRA within FL frameworks endows clients with the capacity for localized, efficient model updates, which is a cornerstone for collaborative learning that upholds the sanctity of data privacy. Recent studies have ventured into integrating LoRA specifically within the fully connected layers of deep neural networks in a FL context. For instance, Wu et al. [18] examined the implementation of LoRA adapters to fine-tune fully connected layers, while Yi et al. [19] explored the prospects of personalized FL through parameter-efficient LoRA updates. These approaches underscore the growing interest and potential in enhancing model adaptability and efficiency, crucial for FL's scalability and practicality. Of late, parameter-efficient low-rank adapters have benefited from Neural Architecture Search (NAS) techniques that can discover high-performing configurations in a search space of elastic LoRA adapters [20], [21].

The application of LoRA to VLMs, such as the CLIP model, allows for nuanced and efficient adaptation to diverse datasets encountered in federated networks. By selectively updating parameters through low-rank matrices, the model can quickly adjust to new visual and textual information, enhancing its ability to generate accurate and contextually relevant outputs. This capability is especially valuable in scenarios where data distributions vary significantly across clients, a common challenge in FL settings.

Furthermore, LoRA's parameter-efficient fine-tuning aligns with the core principles of FL by enabling clients to contribute to a global model without the need for extensive computational power or bandwidth. This approach not only mitigates the risk of data privacy breaches but also democratizes access to advanced AI technologies, allowing entities with limited resources to participate in the collaborative development of robust, multimodal AI systems.

III. METHODOLOGY

In this section, we present the methodology employed in our study on federated fine-tuning learning for VLMs. We outline the framework and techniques used to enable collaborative

TABLE I: The test accuracy (%) of the image classification tasks in IID setting with $N = 10$.

Dataset	FedFFT	FedLC	FedVM-LC	FedAA	Ours
F-MNIST	79.69	70.13	82.39	81.63	89.89
CIFAR-10	87.26	89.69	86.31	89.57	94.49
CIFAR-100	64.88	60.61	61.52	66.68	79.98
TINY	59.77	60.30	59.77	64.46	75.99
OxfordPets	80.66	81.62	81.41	84.13	91.21
Flowers102	85.5	53.5	53.5	67.0	98.0
Aircraft	16.66	16.66	19.29	21.49	56.14
Cars	44.85	46.81	44.85	50.0	85.53
DTD	55.30	44.69	43.18	46.21	85.60
EuroSAT	86.11	50.15	89.04	78.08	95.67
FER2013	41.14	37.40	52.61	66.33	74.06
Caltech101	84.65	79.36	78.83	83.06	97.88
Food101	77.30	77.30	77.30	76.99	82.62
Country211	18.56	18.69	18.56	18.06	27.94
SUN397	65.97	66.58	65.64	68.67	81.22
R-SST2	97.42	75.53	80.68	91.41	82.40

training across distributed devices while incorporating fine-tuning for enhanced performance and privacy preservation.

A. Baseline fine-tuning methods

1) *Full fine-tuning (FFT)*: In full fine-tuning, we update the whole CLIP model, i.e., both image and text encoder, by minimizing the cross-entropy loss on labeled downstream data.

2) *Linear classifier (LC)*: We learn a linear classifier $h_{\text{class}} \in \mathbb{R}^{d \times k}$ on top of frozen image embedding $\bar{f}(I)$ by minimizing the cross-entropy on downstream data [4]. The text encoder is not involved during this process. As Wortsman et al. [22] have suggested, instead of randomly initializing the linear classifier, the zero-shot weights of the CLIP's text encoder are applied to it.

3) *Fine-tune vision model with linear classifier (VM-LC)*: Inspired by Goyal et al. [23], in VM-LC, we update both a linear head and the parameters of the image encoder Θ_{img} (initialized at the pre-trained value) by minimizing the cross-entropy loss on labeled downstream data. The text encoder is not involved during this process. Similar to the previous method, we initialize the linear head the same as the pre-trained weights of the text encoder.

4) *Attention Adapter (AA)* [16]: Here, we integrate an attention adapter after the image encoder to perform a fine-tuning process. The whole CLIP model is kept frozen during training.

B. Approach

Our approach innovatively adapts the CLIP model, leveraging its powerful pre-trained representation for image and text embeddings within a FL framework. At the core of our strategy lies the integration of LoRA fine-tuning, which allows for efficient model adaptation across distributed clients while preserving the intrinsic knowledge captured during CLIP's extensive pre-training.

Given a pre-trained CLIP backbone, the input to the text encoder is crafted in the form "a photo of a [class]", where [class] denotes the category of the image within the dataset [4]. Leveraging CLIP's ability to predict image-text matchings, the

TABLE II: The test accuracy (%) of the image classification tasks in practical non-IID setting with $N = 10$ (method names omitted the word "Fed" for saving space).

Dataset	β	FFT	LC	VM-LC	AA	Ours
F-MNIST	0.1	63.88	63.88	63.88	59.96	84.30
	1.0	75.36	63.88	72.12	77.82	89.38
CIFAR-10	0.1	79.73	85.42	79.73	81.10	93.06
	1.0	88.27	88.27	88.27	91.06	94.39
CIFAR-100	0.1	60.61	60.74	60.61	62.10	74.04
	1.0	62.31	62.31	62.31	66.45	77.09
TINY	0.1	59.77	59.77	59.77	61.52	71.67
	1.0	59.59	59.59	59.59	61.90	71.29
OxfordPets	0.1	80.66	81.21	80.66	77.46	87.13
	1.0	80.66	81.00	80.66	81.82	90.74
Flowers102	0.1	53.5	54.0	53.5	47.5	91.99
	1.0	65.11	66.46	65.11	67.86	95.11
Aircraft	0.1	16.66	16.66	16.66	18.42	44.29
	1.0	17.84	17.84	19.25	21.05	44.75
Cars	0.1	44.85	46.81	44.85	50.0	86.76
	1.0	47.51	47.85	47.51	52.48	75.84
DTD	0.1	40.15	40.90	40.15	40.90	65.15
	1.0	48.84	41.13	43.26	43.61	70.92
EuroSAT	0.1	43.20	49.38	43.20	37.96	92.43
	1.0	65.38	45.18	68.59	68.96	94.44
FER2013	0.1	39.40	37.15	37.15	40.39	62.84
	1.0	40.41	44.88	40.67	44.78	64.79
Caltech101	0.1	78.30	78.30	78.30	79.89	94.58
	1.0	88.18	88.24	88.18	87.49	95.96
Food101	0.1	77.30	77.30	77.30	74.79	78.02
	1.0	79.59	79.59	79.59	79.62	82.79
Country211	0.1	18.56	18.56	18.56	17.68	25.73
	1.0	15.13	15.13	15.13	15.20	23.05
SUN397	0.1	65.64	65.64	65.64	66.78	77.58
	1.0	59.54	61.56	59.54	61.12	73.39
R-SST2	0.1	97.42	97.42	97.85	92.27	88.84
	1.0	60.19	60.19	60.19	54.31	68.00




Fig. 4: The data distribution of all clients on CIFAR-10 in IID and non-IID setting with varying β . The size of a circle means the number of data samples.

classification loss and logits can be calculated by aligning the embedding spaces for images and textual descriptions corresponding to "[class]".

Formally, let $\mathcal{I}(\cdot)$ denote the image encoder, and $\mathcal{T}(\cdot)$ the text encoder, which remains fixed except for the low-rank adaptations through LoRA. For a given class i , let \mathbf{t}_i be the word embedding vector of the class label $[\text{class}]_i$, with $i \in [1, K]$ and K representing the number of classes.

We utilize the cosine similarity function $\cos[\cdot, \cdot]$, as em-

ployed by CLIP, to measure the similarity between the image and text features. The classification prediction probabilities and logits for an image \mathbf{x} against a class i are calculated as follows:

$$p(\mathbf{y} = i|\mathbf{x}) = \frac{\exp(\cos[\mathcal{I}(\mathbf{x}), \mathcal{T}(\mathbf{t}_i)]/\tau)}{\sum_{j=1}^K \exp(\cos[\mathcal{I}(\mathbf{x}), \mathcal{T}(\mathbf{t}_j)]/\tau)} \quad (2)$$

where τ is the temperature coefficient learned by CLIP. The text encoder's LoRA parameters are the only variables updated during local backpropagation and subsequently aggregated by the federated server. The image encoder, the main body of the text encoder, and the word embedding \mathbf{t}_i is kept frozen to preserve the foundational knowledge embedded within the pre-trained CLIP model.

In this framework, the loss function for a batch of images and their associated classes is the cross-entropy loss between the predicted probabilities $p(\mathbf{y} = i|\mathbf{x})$ and the ground truth labels. By minimizing this loss, we perform selective fine-tuning of the model's textual understanding to better align with the domain-specific data encountered in each client's local dataset while capitalizing on the pre-trained multimodal embeddings of the CLIP model.




Fig. 5: Learning curve for DTD dataset in different settings.

IV. EXPERIMENTS


TABLE III: The number of parameters and size of transferred models for each method. For **FedLC** and **FedVM-LC**, there are ranges of trainable parameters and size depending on the dataset.

	Number of trainable parameters	Size (MB)
FedFFT	151,277,313	577.078
FedLC	1026 – 203661	0.004 – 0.777
FedVM-LC	87,457,026 – 87,659,661	333.622 – 334.395
FedAA	525,312	2.004
Ours	24,576	0.094


We conduct comprehensive experiments to evaluate the effectiveness of our federated fine-tuning approach.

A. Datasets and Experimental Setting


1) *Datasets*: We use publicly available vision datasets: Fashion-MNIST (F-MNIST) [24], CIFAR-10, CIFAR-100 [25], Tiny-ImageNet (TINY) [26], Oxford-IIIT



(a) Comparison of the best two approaches in CIFAR-10 that reach 89% and 85% target accuracy for IID and Non-IID respectively.



(b) Comparison of the best two approaches in Flowers102 that reach 85% and 67% target accuracy for IID (FedFFT and FLoRA) and Non-IID (FedAA and FLoRA) respectively.



(c) Comparison of the best two approaches in EuroSAT that reach 89% and 68% target accuracy for IID and Non-IID respectively.

Fig. 6: Comparison of computation and communication overhead in selected datasets to reach target accuracies. FLORA requires fewer communication rounds while having significantly less communication cost and training time compared to the second best method.

Pet (OxfordPets) [27], Oxford 102 Flower (Flowers102) [28], FGVC-Aircraft (Aircraft) [29], Stanford Cars (Cars) [30], Describable Textures Dataset (DTD) [31], EuroSAT [32], [33], FER2013 [34], Caltech101 [35], Food101 [36], Country211 [4], SUN397 [37], [38], and Rendered SST2 (R-SST2) [4].

To simulate the practical pFL setting, we evaluate the model on the same testing data, 20% of the total dataset. The remaining 80% data is used for training which is partitioned for all clients in different settings.




Fig. 7: Memory usage of all methods.

B. Experimental setting

We train all methods for 50 rounds with one local epoch. Total number of client is $N = 10$, the sample rate $\rho = 1$ with batch size equal to 128. For Adam optimizer, learning rate is 5×10^{-5} , $\epsilon = 10^{-6}$, and 0.2 weight decay. All the experiments are done using 1 GPU A6000 with 48 GB. The pre-trained CLIP model is utilized with Vit-B/32 [39] as the base image encoder. For LoRA paramters, we set the rank $r = 2$ and scaling factor $\alpha = 32$. We create two popular statistically heterogeneous settings to simulate the FL environment: balanced IID [6] and label skew practical non-IID [40]–[42].

1) *Homogeneous setting*: Data on each client are separated in IID fashion [6].

2) *Practical heterogeneous setting*: The second scenario is the practical heterogeneous setting [40], [41], which is controlled by the Dirichlet distribution, denoted as $Dir(\beta)$. The smaller the β is, the more heterogeneous the setting is. We set $\beta = 0.1$ and $\beta = 1.0$ for the heterogeneous setting [40], [43]. Figure 4 illustrates the distribution of training data on CIFAR-10 dataset.

3) *Evaluation Metrics*: Our evaluation focuses on task-specific performance metrics such as top-1 accuracy. Additionally, we analyze communication efficiency, memory consumption and few-shot evaluation.

C. Learning Analysis

In our comprehensive assessment of federated learning methods for image classification, FLORA —our proposed method—stands out for its efficient parameter updates and effective learning from decentralized data. Empirical evaluations across diverse datasets have demonstrated FLORA’s superior performance, particularly in IID settings with $N = 10$ clients, and its robust adaptability to varied classification tasks. Notably, except for the binary classification dataset R-SST2, FLORA outshines other baseline methods, as depicted in Table I.

In Table II, we delve into the performance of various FL methods under non-IID data distribution, which presents a more practical and challenging scenario. The table reflects test accuracy percentages for different values of β , where β represents the degree of data distribution skewness among the clients. We aim to evaluate the resilience of each method to data heterogeneity, an aspect critical to FL applications.

The results across datasets highlight the varying ability of each FL method to cope with non-IID data. Our method demonstrates consistent performance superiority across almost all datasets and β settings, emphasizing its robustness and suitability for practical FL deployments.

The learning curves for the DTD dataset, as visualized in Figure 5, corroborate these findings, where FLORA achieves faster convergence and maintains high accuracy with significantly fewer communication rounds compared to other methods. This rapid and robust performance, particularly in non-IID conditions, highlights FLORA’s potential for practical federated deployments, offering a well-suited solution for diverse and distributed data environments. Notice that all of the methods start at almost the same accuracy (as in zero-shot classification thanks to the strong generalization of the pre-trained CLIP model), but other time only FLORA show significant improvement and maintain it’s performance through the training.

TABLE IV: The distribution of classes across clients in few short learning pathological experiment.

Dataset	Σ class	N	# classes of each client	
			$1^{th} \rightarrow (N-1)^{th}$	N^{th}
F-MNIST	10	5	2	2
CIFAR-10	10	5	2	2
CIFAR-100	100	10	10	10
TINY	200	10	20	20
OxfordPets	37	6	6	7
Flowers102	102	6	17	17
Aircraft	100	10	10	10
Cars	196	7	28	28
DTD	47	7	6	11
EuroSAT	10	5	2	2
FER2013	7	3	2	3
Caltech101	101	10	10	11
Food101	101	10	10	11
Country211	211	10	21	22
SUN397	397	10	39	46
SST2	2	2	1	1




Fig. 8: Average results of few-shot learning on various datasets.

D. Efficiency Analysis

In FLORA, a key attribute is that we incorporate a LoRA adapter independent of the datasets. Hence, it only communicates the tiny size network through training and inference, resulting in lower communication costs than FL algorithms

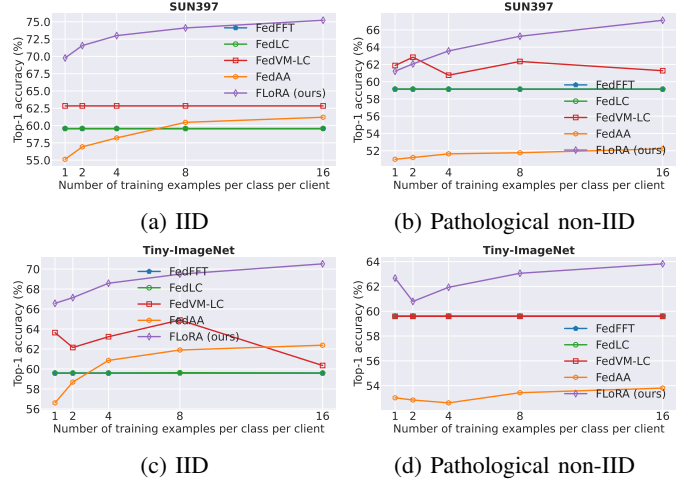


Fig. 9: Result of few-shot learning on selected datasets.

that use whole models for aggregation. We evaluate the communication cost in way such that we train all the models to a target accuracy and calculate the communication cost. The communication cost is represented by:

$$N \times \rho \times \text{cost per round}, \quad (3)$$

where the cost per round is $2 \times$ the size of the exchanged model (for downloading and uploading the transferred model). The transferred model size details are given in Table III.

As demonstrated in Figure 6, our method, FLORA, achieves the target accuracies in fewer communication rounds while substantially reducing both the communication and training costs, as analyzed for the CIFAR-10, Flowers102, and FER2013 datasets. Notably, in the IID setting for CIFAR-10, FLORA requires just one communication round—significantly fewer than the 15 rounds required by competing methods like FedLC. This efficiency carries over to non-IID settings, where FLORA also outperforms in communication costs and training time, suggesting rapid convergence capabilities and suitability for practical federated learning deployments.

Figure 7 further reveals that FLORA maintains a lower GPU memory usage of 11.84 GB, offering an optimal balance between resource economy and model performance. This is attributed to the strategic integration of LoRA adapters, which refine a small yet impactful subset of model parameters, enhancing computational efficiency—a crucial advantage for deploying advanced machine learning on devices with constrained resources.

The efficiency gains from FLORA suggest the model’s potential to facilitate federated learning in environments where minimizing resource use is as important as maintaining data privacy and reducing bandwidth. These findings underscore the viability of FLORA as an effective solution to the challenges of scalability and communication overhead in distributed learning contexts.

E. Few-Shot Evaluation

Few-shot learning is a challenging machine learning paradigm where models must learn from a very limited amount of data. This is particularly relevant in FL, where each client may only have a small number of examples.

In addition to IID scenario, we create pathological heterogeneous setting to simulate the FL environment [6], [44]. For the pathological label skew setting, we sample data with a fraction of total label amount for each client on each dataset from their total of categories, with disjoint data and different numbers of data samples. The details of the number of classes per client in the pathological setting are presented in Table IV.

1) *Homogeneous setting*: Each client will have n -shot training samples of each class (each client has every class). There is a common testing set by randomly selecting 20% of the total number of samples. The data for testing is selected from the remaining samples after allocating the training data, ensuring there is no overlap between training and testing sets. The experiment will be carried up to 16-shot. Due to the limited total number of samples of several datasets, we are only able to do as few as 4-shot learning.

2) *Pathological heterogeneous setting*: Here, each client owns k disjoint classes with n -shot training samples of each of these k classes if the number of classes is divisible by the number of clients (N) to evenly distribute k classes per client. For the cases where the number of classes is not perfectly divisible by the number of clients, we distribute the remainder of the classes to the last client. The experiment will be carried up to 16-shot.

In the analysis of few-shot learning, FLORA emerges as a strong method in federated learning contexts. It outshines established baselines across 16 datasets, exhibiting a robust capability to generalize well and optimize performance even with limited data availability, as depicted in Figure 8. The empirical data show that FLORA not only thrives in IID settings, consistently delivering high top-1 accuracy as the number of per-client training examples increases, but it also maintains a competitive edge in pathological settings that mimic real-world non-IID conditions.

Despite the inherent challenges presented by few-shot learning, FLORA’s performance suggests a particular proficiency in managing data scarcity and distribution variability. While other methods like FedFFT and FedVM-LC plateau or falter, especially when faced with highly skewed data, FLORA’s consistent accuracy—even with minimal training examples—underscores its potential in practical applications that must contend with diverse and imbalanced datasets. Figure 9 shows the results of two selected datasets with relatively big number of classes, SUN397 and TINY. These insights reinforce the relevance of FLORA for few-shot federated learning and underscore its potential impact on the broader deployment of machine learning solutions in data-constrained environments.

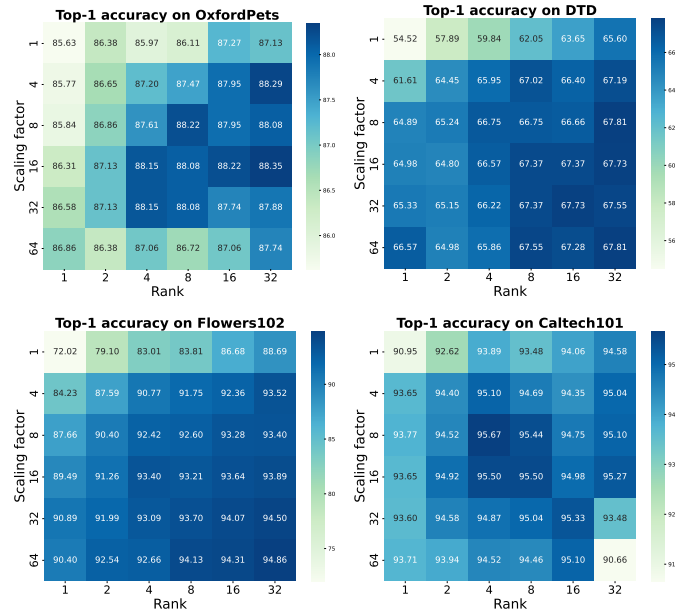


Fig. 10: Different configurations of LoRA adapters placed on projection layers on CLIP’s text encoder in non-IID setting ($\beta = 0.1$).

V. ABLATIONS

The ablation study we conducted rigorously evaluates the impact of varying LoRA hyperparameters—scaling factors and ranks—on the performance of the CLIP model’s text and image encoders within a non-IID federated learning setting. As delineated in Figures 10 and 11, the study dissects the individual and synergistic effects of these hyperparameters on the top-1 accuracy across four distinct datasets, with a particular focus on a non-IID setting characterized by $\beta = 0.1$.

The study explores a range of scaling factors (1, 4, 8, 16, 32, 64) and ranks (1, 2, 4, 8, 16, 32) to understand their individual and combined impact on model performance. The scaling factor likely pertains to the amplification of the projection layer’s adjustments, while rank could represent the number of modifications within the adapter that are allowed to vary.

Our findings, summarized in Table V, reveal that a LoRA adapter with a rank of 2 strikes a strategic balance, significantly reducing the number of trainable parameters to 24, 576 for the text encoder and to 36, 864 for the image encoder, while preserving competitive accuracies. This balance showcases the potential for achieving considerable model performance with a reduced computational footprint—a key factor in the resource-constrained environments typical of federated learning.

Notably, the integration of the LoRA adapter into the text encoder consistently outperforms the image encoder integration while simultaneously maintaining a lower parameter count. This indicates a clear preference for text encoder adaptation in scenarios where both performance optimization and resource efficiency are desired.

TABLE V: Different configurations of LoRA within CLIP. We have selected $r = 2$ because it offers a good tradeoff between accuracy and number of trainable parameters. The scaling factor α does not affect the number of trainable parameters.

	Text Encoder						Image Encoder					
Rank	r = 1	r = 2	r = 4	r = 8	r = 16	r = 32	r = 1	r = 2	r = 4	r = 8	r = 16	r = 32
LoRA size	12,288	24,576	49,152	98,304	196,608	393,216	18,423	36,864	73,728	147,456	294,912	589,824

The conclusive evidence from our study attests to the rank 2 LoRA adapter’s role in bolstering federated learning efficiency. By adopting this configuration, FedLoRA presents as an optimized framework that promises not only swift model convergence and communication efficacy but also practical applicability across various federated learning contexts, especially those with stringent resource limitations.

Our ablation study examined a series of scaling factors, finding that a judicious selection of α is essential for achieving the best trade-off between model accuracy and computational efficiency. While a larger α may potentially lead to greater improvements in accuracy by allowing more pronounced updates to the model, it also risks overfitting and straying too far from the valuable features learned during pre-training. Conversely, a smaller α ensures that the updates remain subtle, preserving the pre-trained model’s generalizability but may not capture the nuances of the new task as effectively.

The empirical results from our study suggest that a moderate scaling factor offers the most effective balance, allowing for sufficient adaptation to the target task without overwhelming the pre-existing strengths of the model. By carefully calibrating α , we enable FedLoRA to adapt to new data while maintaining the integrity of the original CLIP architecture, thus ensuring that our approach is not only resource-conscious but also retains the rich representational capabilities of the foundational model.

Incorporating LoRA within the CLIP model presents a wealth of potential beyond the scope of this paper, such as adapting LoRA to other components like the query, value, key, and multi-layer perceptron elements, as well as across different heads within the model’s architecture. Each of these elements offers a unique set of possibilities for parameter-efficient adaptation with varying ranks and scaling factors. While our current study does not extend to these configurations, they represent intriguing avenues for future investigation. Exploring these alternatives could further illuminate the ways in which federated learning can benefit from LoRA’s efficient fine-tuning, potentially leading to even more sophisticated and resource-aware federated learning models.

VI. CONCLUSION

In this paper, we have rigorously evaluated the efficacy of federated learning (FL) techniques for fine-tuning the CLIP model via LoRA adapters across a spectrum of IID, non-IID, and few-shot learning settings. Our proposed methodology, FLORA, consistently surpasses conventional FL benchmarks, demonstrating superior accuracy and adaptability to diverse datasets and learning environments.

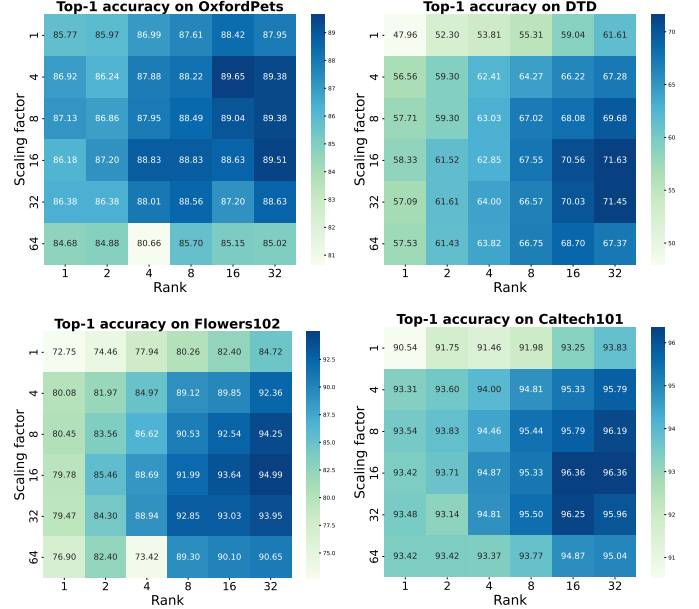


Fig. 11: Different configurations of LoRA adapters placed on projection layers on CLIP’s Image Encoder in non-IID setting ($\beta = 0.1$).

The significant performance gains of FLORA are highlighted in our extensive evaluations, which reveal its remarkable competence in handling both data abundance and scarcity. Ablation studies further distill the critical hyperparameters that underpin the success of multimodal FL models, ensuring both adaptability and computational efficiency.

Ultimately, our investigation into FLORA advances the field of FL by showcasing a method that not only enhances model performance in distributed settings but also offers a strategic blueprint for efficient multimodal model fine-tuning. The results of our work suggest promising directions for future FL systems poised to navigate the complexities of practical data challenges.

In conclusion, this paper presents a step forward in FL, proposing a method that improves the accuracy of models trained across decentralized data and introduces a framework for efficiently fine-tuning multimodal models in a federated context. Our findings open up new avenues for research into adaptive FL systems, which can handle real-world data intricacies while maintaining user privacy and minimizing communication costs.

REFERENCES

- [1] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, “BERT: Pre-training of deep bidirectional transformers for language understanding,” in *Association for Computational Linguistics (ACL)*, 2019. 1
- [2] A. Radford, J. Wu, R. Child, D. Luan, D. Amodei, I. Sutskever *et al.*, “Language models are unsupervised multitask learners,” *OpenAI blog*, vol. 1, no. 8, p. 9, 2019. 1
- [3] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan, R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler, M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever, and D. Amodei, “Language models are few-shot learners,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 1
- [4] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning transferable visual models from natural language supervision,” in *International Conference on Machine Learning (ICML)*, 2021. 1, 2, 4, 6
- [5] E. J. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, S. Wang, L. Wang, and W. Chen, “LoRA: Low-rank adaptation of large language models,” in *International Conference on Learning Representations (ICLR)*, 2022. 1, 3
- [6] H. B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, “Communication-efficient learning of deep networks from decentralized data,” in *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2017. 2, 3, 6, 8
- [7] X. Li, K. Huang, W. Yang, S. Wang, and Z. Zhang, “On the convergence of FedAvg on Non-IID data,” in *International Conference on Learning Representations (ICML)*, 2020. 2
- [8] O. Vinyals, M. Fortunato, and N. Jaitly, “Pointer networks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2015. 2
- [9] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhutdinov, R. Zemel, and Y. Bengio, “Show, attend and tell: Neural image caption generation with visual attention,” in *International Conference on Machine Learning (ICML)*, 2015. 2
- [10] J. Lu, D. Batra, D. Parikh, and S. Lee, “Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 2
- [11] H. Tan and M. Bansal, “LXMERT: Learning cross-modality encoder representations from transformers,” in *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2019. 2
- [12] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, “ImageNet Large Scale Visual Recognition Challenge,” in *International Journal of Computer Vision (IJCV)*, 2015. 2
- [13] T. Guo, S. Guo, J. Wang, X. Tang, and W. Xu, “Promptfl: Let federated participants cooperatively learn prompts instead of models - federated learning in age of foundation model,” in *IEEE Transactions on Mobile Computing (TMC)*, 2023. 2
- [14] T. Guo, S. Guo, and J. Wang, “pfedprompt: Learning personalized prompt for vision-language models in federated learning,” in *ACM Web Conference (WWW)*, 2023. 2
- [15] C. Qiu, X. Li, C. K. Mummadi, M. R. Ganesh, Z. Li, L. Peng, and W.-Y. Lin, “Text-driven prompt generation for vision-language models in federated learning,” in *International Conference on Learning Representations (ICLR)*, 2024. 2
- [16] W. Lu, X. Hu, J. Wang, and X. Xie, “Fedclip: Fast generalization and personalization for CLIP in federated learning,” *IEEE Data Eng. Bull.*, vol. 46, no. 1, pp. 52–66, 2023. [Online]. Available: <https://sites.computer.org/debull/A23mar/p52.pdf> 3, 4
- [17] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *International Conference on Learning Representations (ICLR)*, 2017. 3
- [18] X. Wu, X. Liu, J. Niu, H. Wang, S. Tang, and G. Zhu, “FedLoRA: When personalized federated learning meets low-rank adaptation,” 2024. [Online]. Available: <https://openreview.net/forum?id=bZh06ptG9r> 4
- [19] L. Yi, H. Yu, G. Wang, X. Liu, and X. Li, “pfedlora: Model-heterogeneous personalized federated learning with lora tuning,” *arXiv preprint arXiv:2310.13283*, 2024. 4
- [20] J. P. Muñoz, J. Yuan, Y. Zheng, and N. Jain, “Lonas: Elastic low-rank adapters for efficient large language models,” in *The 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation*, 2024. 4
- [21] J. P. Muñoz, J. Yuan, and N. Jain, “Shears: Unstructured sparsity with neural low-rank adapter search,” in *2024 Annual Conference of the North American Chapter of the Association for Computational Linguistics - Industry Track*, 2024. 4
- [22] M. Wortsman, G. Ilharco, J. W. Kim, M. Li, S. Kornblith, R. Roelofs, R. Gontijo-Lopes, H. Hajishirzi, A. Farhadi, H. Namkoong, and L. Schmidt, “Robust fine-tuning of zero-shot models,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 4
- [23] S. Goyal, A. Kumar, S. Garg, Z. Kolter, and A. Raghunathan, “Finetune like you pretrain: Improved finetuning of zero-shot vision models,” in *Computer Vision and Pattern Recognition Conference (CVPR)*, 2022. 4
- [24] H. Xiao, K. Rasul, and R. Vollgraf, “Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms,” *arXiv preprint arXiv:1708.07747*, 2017. 5
- [25] A. Krizhevsky and G. Hinton, “Learning multiple layers of features from tiny images,” Master’s thesis, Department of Computer Science, University of Toronto, 2009. 5
- [26] P. Chrabaszcz, I. Loshchilov, and F. Hutter, “A downsampled variant of imagenet as an alternative to the cifar datasets,” *arXiv preprint arXiv:1707.08819*, 2017. 5
- [27] O. M. Parkhi, A. Vedaldi, A. Zisserman, and C. V. Jawahar, “Cats and dogs,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012. 6
- [28] M.-E. Nilsback and A. Zisserman, “Automated flower classification over a large number of classes,” in *Indian Conference on Computer Vision, Graphics and Image Processing*, 2008. 6
- [29] S. Maji, E. Rahtu, J. Kannala, M. Blaschko, and A. Vedaldi, “Fine-grained visual classification of aircraft,” *arXiv preprint arXiv:1306.5151*, 2013. 6
- [30] J. Krause, M. Stark, J. Deng, and L. Fei-Fei, “3d object representations for fine-grained categorization,” in *International IEEE Workshop on 3D Representation and Recognition (3dRR)*, 2013. 6
- [31] M. Cimpoi, S. Maji, I. Kokkinos, S. Mohamed, , and A. Vedaldi, “Describing textures in the wild,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2014. 6
- [32] P. Helber, B. Bischke, A. Dengel, and D. Borth, “EuroSAT: A novel dataset and deep learning benchmark for land use and land cover classification,” *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 2019. 6
- [33] —, “Introducing eurosat: A novel dataset and deep learning benchmark for land use and land cover classification,” in *IGARSS IEEE International Geoscience and Remote Sensing Symposium*, 2018. 6
- [34] Dumitru, I. Goodfellow, W. Cukierski, and Y. Bengio, “Challenges in representation learning: Facial expression recognition challenge,” 2013. 6
- [35] L. Fei-Fei, R. Fergus, and P. Perona, “Learning generative visual models from few training examples: An incremental bayesian approach tested on 101 object categories,” in *Computer Vision and Pattern Recognition (CVPR) Workshop*, 2004. 6
- [36] L. Bossard, M. Guillaumin, and L. Van Gool, “Food-101 – mining discriminative components with random forests,” in *European Conference on Computer Vision (ECCV)*, 2014. 6
- [37] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba, “Sun database: Large-scale scene recognition from abbey to zoo,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010. 6
- [38] J. Xiao, K. A. Ehinger, J. Hays, A. Torralba, and A. Oliva, “Sun database: Exploring a large collection of scene categories,” in *International Journal of Computer Vision (IJCV)*, 2014. 6
- [39] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, J. Uszkoreit, and N. Houlsby, “An image is worth 16x16 words: Transformers for image recognition at scale,” in *International Conference on Learning Representations (ICLR)*, 2021. 6
- [40] T. Lin, L. Kong, S. U. Stich, and M. Jaggi, “Ensemble distillation for robust model fusion in federated learning,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 6
- [41] Q. Li, B. He, and D. Song, “Model-contrastive federated learning,” in *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 6

- [42] Q. Li, Y. Diao, Q. Chen, and B. He, “Federated learning on non-iid data silos: An experimental study,” in *International Conference on Data Engineering (ICDE)*, 2022. 6
- [43] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, “Tackling the objective inconsistency problem in heterogeneous federated optimization,” in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020. 6
- [44] A. Shamsian, A. Navon, E. Fetaya, and G. Chechik, “Personalized federated learning using hypernetworks,” in *International Conference on Machine Learning (ICML)*, 2021. 8