

Communication-Efficient LLM Training for Federated Learning

Arian Raje

CMU-CS-24-123

May 2024

Computer Science Department
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Virginia Smith, Chair
Zihao Jia
Gauri Joshi

*Submitted in partial fulfillment of the requirements
for the Master's degree in Computer Science.*

Copyright © 2024 **Arian Raje**

Keywords: Federated Learning, Sparsity, Efficiency, LLMs

Abstract

Federated learning (FL) is a recent model training paradigm in which client devices collaboratively train a model without ever aggregating their data. Crucially, this scheme offers potential privacy and security benefits for users by only ever communicating updates to the model weights to a central server as opposed to traditional machine learning (ML) training which directly communicates and aggregates data. However, FL training suffers from statistical heterogeneity as clients may have differing distributions of local data. Large language models (LLMs) offer a potential solution to this issue of heterogeneity given that they have consistently been shown to be able to learn on vast amounts of noisy data. While LLMs are a promising development for resolving the consistent issue of non-I.I.D. clients in federated settings, they exacerbate two other bottlenecks in FL: limited local compute and expensive communication. This thesis aims to develop efficient training methods for LLMs in FL. To this end, we employ two critical techniques in enabling efficient training. First, we use low-rank adaptation (LoRA) to reduce the computational load of local model training. Second, we communicate sparse updates throughout training to significantly cut down on communication costs. Taken together, our method reduces communication costs by up to 10x over vanilla LoRA and up to 5x over more complex sparse LoRA baselines while achieving greater utility. We emphasize the importance of carefully applying sparsity and picking effective rank and sparsity configurations for federated LLM training.

Acknowledgments

First and foremost, I would like to thank my advisor, Professor Virginia Smith, for all her guidance and support. I could not have had a more positive experience learning to conduct research, and I am extremely grateful that I had the opportunity to join her lab as an undergraduate. Next, I would like to thank my other committee members, Professor Gauri Joshi and Professor Zhihao Jia. The opportunity to share some of my research and get valuable feedback is extremely meaningful. I am additionally grateful to my mentor Kevin Kuo, who has supervised my research and provided consistent help since the very beginning. Finally, I owe so much to the friends and family who have supported me throughout my education and over this past year. You all mean so much to me, and I would not be where I am without you. Aaju, you are the best.

Contents

1	Introduction	1
2	Background	3
2.1	Federated Learning Methods	3
2.1.1	Problem Statement	3
2.1.2	Optimization Techniques	4
2.2	Large Language Model Fine-Tuning	5
2.3	Sparsity and Model Compression	7
3	Related Works	11
3.1	Communication Efficiency in Federated Learning	11
3.2	Fine-Tuning in Federated Learning	12
3.3	Pruning Adapters	12
4	Methods	15
4.1	Federated LoRA	15
4.2	Sparsity for LoRA	16
4.3	FLoSS Algorithm	16
4.4	Benchmarks	17
5	Results	19
5.1	Model Performance with Fixed Communication	19
5.2	Robustness to Statistical Heterogeneity	20
5.3	Improvements to Communication Efficiency	21
5.4	Privacy-Preserving Variations	22
6	Discussion and Future Work	25
	Bibliography	27

List of Figures

- 2.1 FL procedures involve 4 steps in each communication round. (1) Clients download model weights from the central server. (2) Clients train the model on local data. (3) Clients upload local model weights to the central server. (4) The central server aggregates client updates into a new global model. 4
- 2.2 LLM fine-tuning involves taking an open-source model trained on generic data and adapting the weights on task-specific data. 5
- 2.3 LoRA inserts small trainable matrices into the model architecture and freezes the remaining model weights. 6
- 2.4 Structured pruning sets entire structures to 0 while unstructured pruning sets individual weights to 0. 7
- 2.5 Pruning at initialization prunes once before the start of model training while iterative pruning prunes repeatedly throughout the training procedure. 9

- 4.1 FLoSS procedure involves (1) sparsifying adapters prior to download (2) training dense adapters (3) sparsifying adapters prior to upload (4) aggregating using FedAdam. 17

- 5.1 Comparison of communication-efficient LoRA methods in FL. 20
- 5.2 Comparison of communication-efficient LoRA methods in FL with non-I.I.D clients. 21
- 5.3 Performance of FLoSS on CIFAR10 with varying rank and sparsity configurations. 22
- 5.4 Performance of FLoSS with varying degrees of DP. 23

- 6.1 Comparison of communication time vs. performance for FLoSS on CIFAR10 . . 26

List of Tables

4.1	Statistics of the datasets used in the experiments.	18
-----	-------------------------------------------------------------	----

Chapter 1

Introduction

The increased ubiquity of edge devices, such as IoT devices and smartphones, has made federated learning (FL) paradigms a feasible way to train machine learning (ML) models [35]. In traditional ML applications, a central server would first aggregate the data from disparate sources. A model would then be trained at the server level using this aggregated data. While FL applications similarly involve a central server and disparate data sources, which we refer to as “clients”, FL offers an alternative to traditional training procedures. FL applications will instead first distribute the model parameters from the central server to the clients (download phase). The clients then independently train the model parameters they receive from the central server on their local data. Once local training is complete, the clients send the model parameters back to the central server (upload phase) where the central server aggregates the clients’ models into a new global model. In total, the above steps constitute a single “communication round” for the training procedure. This process is repeated for multiple communication rounds, with each round involving sampling clients, training local models, and aggregating the models at the central server. In this training scheme, data never leaves the clients’ devices and only the model parameters are communicated between the central server and the clients. As a consequence, FL offers potential privacy benefits over traditional ML [5]. FL training offers other benefits include improved personalization [45] and scalability [59]. FL-based training schemes are particularly useful in large networks of similar devices and have already been productionized at scale [6, 24].

While FL continues to become more valuable in practice, practical concerns about its utility remain. In particular, the use of Large Language Models (LLMs) has become standard for a vast number of ML problems [47]. In many settings, the LLM being used may have billions of trainable parameters. The scale of these LLMs presents serious compute and communication bottlenecks for distributed training schemes. Since the advent of LLMs, a common strategy to use LLMs has been the pretrain-then-fine-tune framework. In essence, LLMs like GPT [8] or BERT [29], which have already been pretrained on a large corpus of data, can be fine-tuned on a downstream task. A consequence of this framework is that the updates to the pretrained model weights take on low-rank structures, eliminating the need for full fine-tuning of all the model weights. In recent years, adapter methods for LLMs have been developed to incorporate these ideas by injecting a small set of trainable parameters into each transformer block and freezing the remaining parameters of the model [21, 23]. Therefore, when a pretrained LLM is being

fine-tuned on a downstream task, only a smaller portion of parameters must be trained.

Adapter methods offer a concrete way to reduce the compute and communication loads of LLM training in a distributed or federated setting. However, since FL training can still require a large number of communication rounds and clients, the communication costs of adapter methods in FL can still be prohibitively high. While adapter methods reduce some of the communication bottlenecks associated with distributed LLM training, in practice they can still lead to slow communication while seeing more significant drops in model utility. They may additionally increase storage and inference costs by increasing the number of total parameters included in the model [51].

LLMs broadly have achieved state-of-the-art performance in multiple domains and in many ways have become the standard approach for modeling with large amounts of data. Nonetheless, they remain difficult to implement in FL. The fact that FL applications involve communication over a wireless network means that coordinating large model updates from a large array of clients is specifically difficult and has been a primary concern in utilizing LLMs for FL. **The goal of this thesis is to introduce a method to perform communication-efficient LLM training for FL.** To this end, we introduce **Federated LoRA with Simple Sparsity (FLoSS)**. We specifically employ sparsity to low-rank adaptation (LoRA) during only the download and upload phases of FL training in order to retain the model’s utility while restricting communication. We additionally suggest heuristics to accurately select a rank and download/upload sparsity ratio to perform accurate model training given an arbitrary communication budget. We summarize our contributions as follows -

1. To the best of our knowledge, we are the first to apply unstructured sparsity to LoRA for efficient federated fine-tuning. We focus on unstructured (weight-level) sparsity because it has been shown to outperform structured (block-level) sparsity in centralized settings.
2. We propose FLoSS, a simple baseline that applies a constant top- k sparsity only to communication. This method can reduce communication costs up to 10× while matching the performance of dense LoRA on several FL image and text tasks.
3. We simulate an FL training procedure on a network with realistic download and upload communication speeds. Given a communication budget, we recommend heuristics to accurately select a LoRA rank and download/upload sparsity ratios that maximize the model’s utility within that budget.

FLoSS aims to make LLM training more feasible in a federated setting. LLM training in resource-constrained environments remains an open research problem and an important field of study as LLMs grow in size. We hope to contribute to this growing field of work by proposing methods to improve efficiency while retaining utility in real-world settings.

Chapter 2

Background

2.1 Federated Learning Methods

2.1.1 Problem Statement

FL applications train an ML model across a distributed network of clients [35]. Generally, FL applications involve a central server and a set of k client devices. At each communication round, a subset of the k devices each download a copy of the global model weights, train the model on a local objective, and upload the locally trained model weights back to the central server. The communication round ends with the central server aggregating the weights into a new global model. The process is depicted below in Figure 2.1. The objective of an FL application is traditionally described as follows where the model is parameterized by weights w -

$$\min_w F(w), \text{ where } F(w) := \sum_{i=1}^k p_i F_i(w) \quad (2.1)$$

F_i describes the local objective for client i . In our setup, we treat F_i as the loss of the model parameterized by w with respect to the local data on client device i -

$$F_i(w) = \frac{1}{m_i} \sum_{j=1}^{m_i} l_i(x^{(j)}, y^{(j)}; w) \quad (2.2)$$

Each client has m_i local examples and local loss function l_i . Our global objective F is weighted by parameters p_i where each $p_i \geq 0$ and $\sum_{i=0}^k p_i = 1$. While multiple weighting schemes for clients exist, we define $p_i = \frac{1}{k} \forall i \in [1 \dots k]$. With this weighting, the global objective function treats each client equally regardless of the distribution of local data. Therefore, an FL training procedure aims to find parameters w that minimize the average of the loss across the k clients. We note that there exist certain FL settings that train multiple models (optimizing multiple w_i 's) [4] or have different tasks for different clients (using different l_i 's) [15]. Our focus is single-task, single-model FL where each client uses the same loss function and optimizes a single set of parameters w . The next section described two common procedures used to achieve this objective, namely FedAvg and FedAdam.




Figure 2.1: FL procedures involve 4 steps in each communication round. (1) Clients download model weights from the central server. (2) Clients train the model on local data. (3) Clients upload local model weights to the central server. (4) The central server aggregates client updates into a new global model.

2.1.2 Optimization Techniques

We present two algorithms that aim to find model parameters w that minimize the objective F . These methods draw from historical works in distributed optimization. The first method is Federated Averaging (FedAvg) [35]. Consider a total of k client devices. At communication round t , n clients are sampled from the k total clients. These n clients each download a copy of the global model with parameters w_t from the central server and train the model for e epochs on local client data. This results in local models $w_t^{(1)}, w_t^{(2)}, \dots, w_t^{(n)}$. These models are uploaded back to the central server. The central server then aggregates by averaging these models to define the new global model. This update is formulated as follows -

$$w_{t+1} \leftarrow \frac{1}{n} \sum_{i=1}^n w_t^{(i)} \quad (2.3)$$

This process is repeated for a total T communication rounds. In this straightforward method, successive averaging of client updates aims to make the final global model w_T accurate for all clients k despite sampling a fraction of the clients in each communication round. In settings with large k , it is possible for $n \ll k$ to still result in an accurate final model as has been shown empirically in many studies. Federated Adam (FedAdam) [40] proceeds similarly to FedAvg with a slight adjustment made to the aggregation mechanism. In FedAdam, once the clients performing local training to calculate $w_t^{(1)}, w_t^{(2)}, \dots, w_t^{(n)}$, they each calculate $\Delta_t^{(i)} = w_t^{(i)} - w_t$ and upload $\Delta_t^{(i)}$ back to the central server. The central server calculates the averages of these differences $\Delta_t = \frac{1}{n} \sum_{i=1}^n \Delta_t^{(i)}$ and uses this Δ_t as a “pseudo-gradient” for an Adam optimizer.

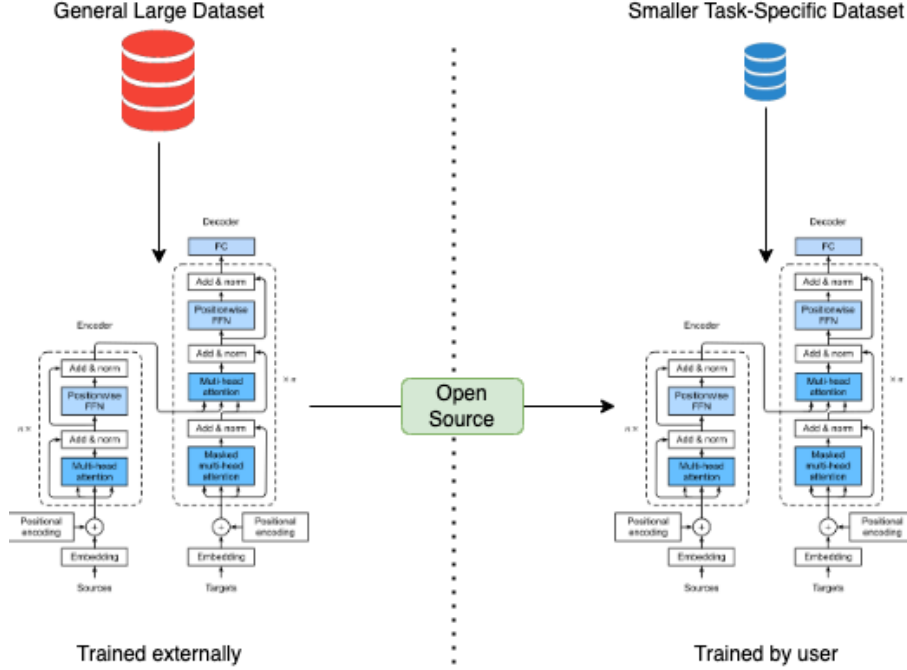


Figure 2.2: LLM fine-tuning involves taking an open-source model trained on generic data and adapting the weights on task-specific data.

In other words the update can be formulated as -

$$\begin{aligned}
 m_{t+1} &\leftarrow \beta_1 m_t + (1 - \beta_1) \Delta_t \\
 v_{t+1} &\leftarrow \beta_2 v_t + (1 - \beta_2) \Delta_t^2 \\
 w_{t+1} &\leftarrow w_t + \gamma \frac{m_t}{\sqrt{v_t + \epsilon}}
 \end{aligned} \tag{2.4}$$

FedAdam has been shown to outperform FedAvg in a variety of settings. For this reason, we utilize the FedAdam optimization method for the remainder of our experiments. We aim to benchmark FLOSS against the best possible baseline. Therefore, using FedAdam was the natural choice to compare against for our communication-efficient method for federated training.

2.2 Large Language Model Fine-Tuning

LLMs have become a standard approach for a variety of ML problems due to their state-of-the-art performance in a broad host of domains (e.g. language translation [39], question answering [61], next-word prediction [56], large-scale vision [13], etc.). An important finding is that the performance of LLMs scales well with respect to the number of parameters in the model [28]. This runs contrary to other architectures like CNNs, RNNs, and LSTMs where performance on a task reaches a saturation point despite an increase in the number of parameters. However, as LLMs continue to grow in size, they become increasingly hard to train efficiently. Coupled with the fact that LLMs require inordinate amounts of data in order to achieve their state-of-the-art




Figure 2.3: LoRA inserts small trainable matrices into the model architecture and freezes the remaining model weights.

accuracy, training LLMs presents serious logistical and computational challenges. To address some of these challenges, recent literature has focused on a pretrain-then-fine-tune setup for LLM training. The core idea is to use open-access LLMs that have been pretrained on a large corpus of public data and then fine-tune the model weights on a domain-specific task [18]. Models like LLaMA [46], GPT [8], BERT [29], and ViT [13] have become standard open-source models to use in this training paradigm for tasks including representation learning, chat, and image classification. While this pretrain-then-fine-tune setup helps resolve some of the massive data requirements for LLM training, it still suffers from the computational, memory, and storage requirements for training all the weights of an LLM.

Adapter methods improve the computational efficiency of LLM training by reparameterizing the updates to the model [18]. Instead of training all the weights of the LLM, adapter methods inject a small set of trainable weights into the model architecture and freeze the original model weights at their pretrained value. These methods are inspired by the idea that the change in weights from a pretrained model to a fine-tuned model exist in a low-rank space. The most frequently used adapter is Low-Rank Adaptation (LoRA) [23]. LoRA reparameterizes weight updates as follows. Consider an initial weight matrix $W_0 \in \mathbb{R}^{d \times d}$. The update to W_0 , which we call $\Delta W \in \mathbb{R}^{d \times d}$ can be defined as a product BA where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times d}$. In this case r is a hyperparameter and is generally defined in a way such that $r \ll d$. In order to make training using LoRA more efficient, W_0 is frozen at its pretrained value and only B and A receive gradient updates. The forward pass of the model with input x can simply be defined as follows -

$$Wx = (W_0 + BA)x = W_0x + BAx \quad (2.5)$$

B is initialized as the 0 matrix while A is initialized as $N(0, \sigma^2)$ so at the initialization of the LoRA parameters, $Wx = (W_0 + 0A)x = W_0x$. Ultimately, by training the LoRA parameters instead of W_0 directly in this context, we train only $2dr$ parameters as opposed to d^2 parameters. Choosing an adequately small r can greatly improve the efficiency of fine-tuning.

Inserting LoRA parameters into an LLM is straightforward. In each transformer block of an LLM

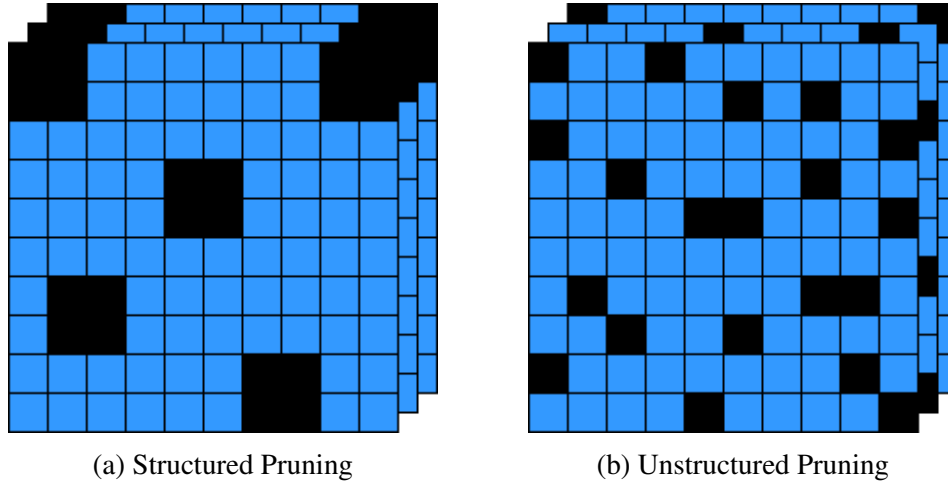


Figure 2.4: Structured pruning sets entire structures to 0 while unstructured pruning sets individual weights to 0.

there is a Multi-head Self-Attention (MSA) mechanism with weight matrices $W_Q, W_K, W_V \in \mathbb{R}^{d \times d}$ where d is the embedding dimension [47]. With input X , we define $Q = XW_Q$, $K = XW_K$, and $V = XW_V$. MSA is then calculated as -

$$\text{MSA}(X) = \text{softmax}\left(\frac{QK^T}{\sqrt{d}}\right)V \quad (2.6)$$

In this calculation, W_Q, W_K , and W_V are all trainable parameters. With LoRA, we insert B_Q, A_Q, B_K, A_K , and B_V, A_V and freeze W_Q, W_K , and W_V at their pretrained weights. We use LoRA for our experiments for the following reasons -

- LoRA parameters can be easily merged back into the model by calculating $W = W_0 + BA$ to reduce inference and storage costs.
- LoRA can be easily integrated with any transformer-based architecture as these architectures all use some form of MSA and LoRA parameters can be defined for the weight matrices used in MSA.
- Using LoRA significantly reduces VRAM consumption making local LLM training feasible for client devices in an FL setting.

2.3 Sparsity and Model Compression

Model compression techniques aim to reduce the size of an ML model by altering the weights or structure of the model [60]. These methods have been used for various purposes including improving the training efficiency, decreasing the storage requirements, and reducing the inference latency of ML models [50]. While there exist multiple methods for model compression, we focus on pruning/sparsity. Pruning methods set a large fraction of model weights to zero and compactly


represent the model in a sparse matrix format [10]. These methods rank the model weights according to a scoring function and prune a fraction of the weights with the lowest scores. While there are numerous approaches to pruning and sparsification, we broadly divide current pruning literature along the following axes: structured vs. unstructured and pruning at initialization vs. iterative pruning. We describe these categories and their consequences. First, we compare structured and unstructured pruning -

- Structured pruning - In structured pruning, the scoring function ranks existing structures within the model architecture (filters, channels, layers, etc.) [52]. Structures are then pruned in their entirety by setting every parameter that exists in that structure to 0. Structured pruning is useful in accelerating training because block level sparsity can speed up matrix multiplication on GPUs. However, the downside to structured pruning is that there is less flexibility in pruning patterns, often resulting in a model that is significantly less accurate than its dense counterpart.
- Unstructured pruning - In unstructured pruning, the scoring function ranks individual weights within the model architecture [57]. Weights with the lowest scores are set to 0 independently of other weights in the model. While unstructured pruning offers more flexibility in sparsity patterns and usually results in more accurate models, they offer few benefits in terms of training efficiency and inference latency. This is because random sparsity patterns require custom hardware to see noticeable speedups for operations like matrix multiplication. For general consumer hardware, there are few efficiency advantages to applying unstructured pruning [49].


The next dimension on which pruning techniques differ is pruning at initialization and iterative pruning -

- Pruning at initialization - In pruning at initialization, model weights or structures are ranked according to the scoring function prior to any model training [14]. Weights or structures with the lowest scores are immediately set to 0 and masked for the remainder of model training. This method allows you to pre-define a sparsity ratio and retain that sparsity ratio for the entire training procedure. However, the scores assigned to weights or structures at initialization may be inaccurate and fail to reflect their actual importances as a consequence of model training.
- Iterative pruning - In iterative pruning, at certain points during model training, the scoring function ranks the current weights or structures in the model [37]. A smaller fraction of the model's weights are set to 0 and masked for the remainder of model training. In this way, the model becomes progressively more sparse throughout the training procedure. Iterative pruning tends to result in more accurate models because weights or structures are pruned based on their current value during model training and fewer weights are pruned at once. However, in initial rounds, the model is almost entirely dense resulting in little advantage over dense training until the model becomes sparser later in training.

We choose to focus on sparsity because recent literature has demonstrated that models can be



(a) Pruning at Initialization



(b) Iterative Pruning

Figure 2.5: Pruning at initialization prunes once before the start of model training while iterative pruning prunes repeatedly throughout the training procedure.

significantly compressed using these techniques while retaining performance close to the original model. In contrast, techniques like quantization, which aim to represent weights in lower-bit precision formats, seem to degrade in accuracy more significantly with fewer memory savings. Even still, FLoSS must balance the utility-memory trade-offs described in the above definitions. Without a careful consideration of system design, it remains difficult to train a model using high levels of sparsity that is as accurate as its dense counterpart. We describe the procedure to achieve considerable performance and efficiency in our methods section.

Chapter 3

Related Works

3.1 Communication Efficiency in Federated Learning

A bottleneck in FL systems is the communication cost associated with training [35, 55]. In FL applications, clients communicate wirelessly with the central server. Additionally, in large-scale FL systems that utilize millions of clients, client devices can exist anywhere in the world. Taken together, this means that communication for FL training can be significantly slower than distributed data center training where devices are connected and local. Slow communication can make FL training impractical as long-lasting communication rounds could make it difficult for the model to converge in a timely fashion. For this reason, many FL applications consider communication efficiency an important principle when considering system design.

Modern approaches consider two different ways to tackle the issue of communication efficiency in FL, namely either reducing the number of communication rounds or reducing the size of communicated messages. Some approaches that fall into these categories are described below -

- Reducing the number of communication rounds - Approaches that aim to reduce the number of communication rounds focus on systematically improving model convergence time. For example, CMFL [48] does not upload outlier updates from clients to keep updates in each round relevant to global model convergence. CA-FL [1] and FL+HC [7] use clustering to determine the most representative update from a set of clients and only sends that update to the central server. FedBoost [17] uses an ensemble of models that converge faster than a single larger model to reduce the number of updates required for each of the ensemble models. The method described in [9] uses a probabilistic model to select devices most likely to contribute to faster model convergence. Finally, methods like One-Shot FL [16] and k -Fed [11] propose performing all FL training in a single communication round as opposed to iteratively updating the global model over multiple communication rounds. While some of these methods have empirically demonstrated improvements over vanilla FL, many have concrete failure modes that make it difficult to demonstrate their effectiveness in a broad array of real-world settings. Additionally, many of these methods require changes to the FL training procedure that are difficult to implement at scale.

- Reducing the size of communicated messages - Methods that reduce the size of communicated messages focus on sending or receiving partial updates or compressing the size of updates. LFL [2] and FedPq [41] use quantization to represent model updates at a lower bit precision. FedMP [27], PruneFL [26], and model pruning for HFL [34] all use various types of pruning during the communication phases of FL. Finally, FedKD [53] and DS-FL [25] use knowledge distillation to compress the model parameters and improve communication efficiency. We focus on these methods as they demonstrably reduce the communication cost at each communication round in training. However, we note that a significant drawback of these methods is that reducing the amount of information communicated at each round may degrade the model’s utility. We design FLoSS in a way that balances communication efficiency and model performance across clients.

3.2 Fine-Tuning in Federated Learning

Training LLMs in FL is a difficult task due to computational constraints at the client-level and communication constraints across the network. These constraints necessitate the use of parameter-efficient fine-tuning (PEFT) methods like adapters [21, 23], prompt-tuning [31], Bit-Fit [3] etc. By using these PEFT methods for local computation, resource-constrained clients can more effectively train models on local data. A few recent works have examined the idea of fine-tuning open source LLMs in a federated setting. FedPETuning [58] uses methods like the Housby adapter [21] and prefix tuning to demonstrate efficiently fine-tuned LLMs in FL are robust to privacy attacks. The method described in [20] proposes fine-tuning all the weights of model architectures like BERT and DistilBERT in FL using a fixed compute budget per client. FedPEFT [43] benchmarks methods like bias-tuning and prompt-tuning in FL across a variety of dataset/model pairings. Each of these methods look to reduce the computational load of local client training. In our work, we consider dense PEFT as a naive baseline and study how to further reduce its message size using sparsity.

3.3 Pruning Adapters

Some recent literature has used pruning to enhance the efficiency of adapters. Specifically, SparseAdapter [19] proposes pruning adapters once at initialization. The pruned weights are set to 0 for the entire training procedure. Thus, sparsity is only applied once prior to the beginning of training and the sparsity pattern never changes throughout the training process. AdapterLTH (Lottery Ticket Hypothesis) [54] performs iterative pruning by alternating between pruning away a small fraction of the lowest scoring weights and retraining the remaining weights of an adapter module such as LoRA. While these works improve the storage efficiency of LoRA, they otherwise have marginal practical benefits in centralized settings. First, LoRA parameters can be merged back into the model backbone to eliminate storage costs after training. Second, Unstructured sparsity often requires specialized hardware and software to accelerate computation. Otherwise, training and inference are no more efficient than that of a dense counterpart. We argue that combining unstructured sparsity with LoRA is particularly effective for handling issues

of communication in FL. There are important design considerations when utilizing sparse LoRA methods in FL. For example, depending on how sparsity is applied, communication rounds may have varying levels of sparse communication. The result is that certain rounds may be slower than others during the training process leading to lags and inefficient training. More importantly, sparsity, as a form of lossy model compression, can result in inaccurate models that fail to perform at the level of their dense counterpart. Therefore, applying extreme levels of sparsity, while helpful for efficient communication, may diminish the model's utility. For these reasons, there are critical design elements in a sparse LoRA implementation for FL that do not exist in centralized settings.

Chapter 4

Methods

This section details the combination of adapter methods and sparsity techniques we employ in FLoSS. We additionally highlight the importance of each step in resolving real-world bottlenecks, such as on-device computation and communication constraints, for FL training. Finally, we describe a set of benchmarks we compare against to demonstrate the effectiveness of our training procedure.

4.1 Federated LoRA

FLoSS utilizes the LoRA adapter in a federated setting. Federated LoRA training functions as follows. The training procedure starts with a pretrained LLM that exists at the central server. Each client then downloads a copy of the weights of the pretrained LLM. While this is an expensive operation, we note that in this training procedure the full weights of the LLM are only ever communicated once. Additionally, since the full LLM weights are only ever communicated during a download phase, and download is typically much faster than upload, we can afford to incur this one-time communication cost during this training procedure. The central server then initializes LoRA parameters in each transformer block of the pretrained LLM. FL training proceeds similarly to the description provided in section 2.1 with one critical difference - only the LoRA parameters are communicated between the central server and the client devices. At each communication round, the sampled clients only download a copy of the global model's LoRA parameters. If the sampled clients do not have initialized LoRA parameters for their local model, they initialize fresh LoRA parameters. Then all sampled clients replace the weights of their LoRA parameters with the weights of the downloaded LoRA parameters. The sampled clients train only the weights in the LoRA parameters, freezing the remaining LLM weights as described in section 2.2. Once local training is complete, the sampled clients send only the LoRA parameters back to the central server where they are aggregated into new global LoRA parameters.

Using LoRA for FL is particularly useful as it enables efficient on-device training. Training the full LLM at the client level would be difficult given memory and computation constraints. In cross-device FL settings, client devices tend to be especially resource-constrained. As such, reducing the computational load of FL training is important in this setting.

4.2 Sparsity for LoRA

While LoRA enables on-device training by significantly reducing VRAM usage, LoRA parameters can still be expensive to communicate. LoRA adapters are inserted for each weight matrix in the MSA layer and for each transformer block in the model. Communicating all these parameters is difficult on a wireless network where communication, and specifically the upload phase, is slow and time-consuming. There can also be a large number of clients and communication rounds, necessitating faster communication. Thus, it is important to retain the on-device computational benefits of LoRA training while reducing communication latency for this training procedure. We reduce communication cost by applying top- k sparsity to LoRA during the download and upload phase of FL training.

To perform top- k sparsification consider a weight matrix $W \in \mathbb{R}^{d \times k}$ with n nonzero entries. We define a sparsity ratio $\alpha \in [0, 1]$. We apply the top- k function to $|W|$ where $k = \alpha \cdot n$ and retrieve the indices of these top- k values. We call this set of indices $I = \{(y_i, x_i) \mid y_i \in [0, d - 1], x_i \in [0, k - 1]\}$. We then define a binary mask $M \in \mathbb{R}^{d \times k}$ where $M[y_i, x_i] = 1$ if $(y_i, x_i) \in I$ and $M[y_i, x_i] = 0$ otherwise. W is then updated as $W \odot M$. This functions to mask the smallest weights in W .

In our training scheme, top- k sparsity is applied to the LoRA adapters prior to download and prior to upload. In this way, fewer weights are communicated in each phase, reducing communication latency as a result. There are a few important design considerations in determining how sparsity should be specifically applied to LoRA. First, the sparsity ratio can be different for download and upload. This is because download bandwidth is typically much greater than upload bandwidth (the download bandwidth for a wireless network can be $10\times$ greater than the upload bandwidth for a wireless network). Therefore, being able to configure separate sparsity ratios is important as sparser downloads may be less helpful than sparser uploads. Second, unstructured sparsity is more useful in our setup than structured sparsity. An important distinction is that FLoSS does not retain sparsity during client local training. This is because sparsity would be unhelpful in accelerating training on devices such as smartphones or IoT devices where the hardware is not specifically designed to handle sparse computation. Therefore, dense training would yield results just as quickly and much more accurately than sparse training at the client level. Since unstructured sparsity has been shown to boost performance relative to structured sparsity, and training acceleration is not a consideration, unstructured sparsity is the preferred sparsity method for communication in FLoSS.

4.3 FLoSS Algorithm

FLoSS combines federated LoRA with sparsity to perform communication-efficient LLM training. At each communication round, the central server will sparsify the LoRA parameters of the global model according to a download sparsity ratio d_{down} . Sampled clients then download these sparse adapters and train them locally without any sparsity applied. Before uploading the locally trained adapters, clients sparsify the adapters according to an upload sparsity ratio d_{up} . The com-




Figure 4.1: FLoSS procedure involves (1) sparsifying adapters prior to download (2) training dense adapters (3) sparsifying adapters prior to upload (4) aggregating using FedAdam.

munication round ends with the central server aggregating these sparse adapters.

FLoSS offers a few key advantages for federated LLM training. First, FLoSS enables more efficient on-device training for LLMs. It does so by training only LoRA adapter weights and limiting local training to a single epoch in each communication round for all clients. Since cross-device settings usually involve many clients where each client has limited data, we argue that this training scheme is feasible at the client level. Second, FLoSS reduces communication latency despite the large number of adapter weights being communicated in each round. By sparsifying adapters prior to communication, fewer weights are communicated, and communication can occur more quickly. Finally, FLoSS retains model performance by allowing dense fine-tuning at the client level. We demonstrate the importance of this last feature in an ablation study that identifies the consequence of various pruning methods on model performance [30]. Ultimately, FLoSS achieves high communication efficiency without sacrificing model performance.

4.4 Benchmarks

We present experiments on three datasets: CIFAR10, 20NewsGroups, and Reddit. We resize the CIFAR10 images to 224×224 to match ImageNet, the pretraining dataset for the ViT model architecture we chose. We use the GPT2 tokenizer to preprocess the examples of 20NewsGroups and Reddit into sequences with length 128 and 25 respectively. We partition CIFAR10 and 20NewsGroups across the client devices. As described in the Results section, we test both I.I.D. clients as well as non-I.I.D. clients for partitioning both datasets. The Reddit comments are naturally partitioned by user.

In all experiments, we sample 5 clients at each round and perform one epoch of local training with a batch size of 16. We fine-tune all models for 200 rounds. For the pretrained models, we used ViT-B-16 (85M params) and GPT2-Small (124M params). For all datasets, we report the accuracy on the validation partition. More details on the task setups can be found in Table 4.1.

Dataset	Backbone	Task	#Clients	#Examples	#Classes
CIFAR10	ViT-B-16	Image Classification	500	50K	10
20NewsGroups	GPT2-Small	Sequence Classification	350	20K	20
Reddit	GPT2-Small	Next Token Prediction	40K	1.1M	50257

Table 4.1: Statistics of the datasets used in the experiments.

We compare FLoSS against two other sparse LoRA methods, SparseAdapter and AdapterLTH. Both these methods are described in section 3.3. To use AdapterLTH in FL, we consider training LoRA weights A and B using FedAdam. After each aggregation round, we apply increasingly sparse masks to the LoRA weights. We use the efficient “fine-tuning” version of LTH which continues training from the pruned state rather than rewinding the weights after pruning. This allows the model to recover from pruning within fewer rounds and is necessary to keep communication costs competitive with the dense LoRA baseline. For both of these methods our choice of scoring function is top- k applied to the magnitude of the weight. This allows for a direct comparison to FLoSS which uses the same unstructured scoring function but only sparsifies communication as opposed to sparsifying both communication and model training. For these pruning methods, we perform an initial round of dense LoRA training, so the B adapter is non-zero and can be effectively pruned using our score function.

Algorithm 1: PyTorch-like LoRA training with FedAdam and sparse communication

```

1 Require:  $d_{\text{down}}, d_{\text{up}}$  (download and upload sparsity ratio)
2  $P \leftarrow$  Initialize LoRA parameters
3  $\text{optim} \leftarrow \text{torch.nn.optim.Adam}(\text{params}=P)$ 
4 for  $r = 1, \dots, R$  do
5      $M_{\text{down}} \leftarrow$  mask of top  $d_{\text{down}}$  fraction entries of  $P$  by magnitude
6     Sample clients  $c_1, \dots, c_n$  uniformly at random without replacement
7     for  $i = 1, \dots, n$  in parallel do
8          $P_i = P \odot M_{\text{down}}$  # sparse download
9          $P'_i \leftarrow$  update  $P_i$  with 1 SGD epoch on data of  $c_i$  # fine-tuning all entries of  $P_i$ 
10         $\Delta P_i \leftarrow P_i - P'_i$ 
11         $M_{\text{up},i} \leftarrow$  mask of top  $d_{\text{up}}$  fraction entries of  $\Delta P_i$  by magnitude
12         $\Delta P_i \leftarrow \Delta P_i \odot M_{\text{up},i}$  # sparse upload
13     $\text{optim.grad} \leftarrow \frac{1}{n} \sum_{i=1}^n \Delta P_i$  # set Adam pseudo-gradient
14     $\text{optim.step}()$  (updates  $P$ ) # take one step of Adam

```

Chapter 5

Results

We describe our key findings when running FL training experiments with FLoSS. There are key metrics by which we can measure the effectiveness of this training procedure. Concretely, we determine improvements to model performance, communication efficiency, and privacy over other efficient adapter methods for FL. The first section compares the performance of models trained with FLoSS with both dense training as well as other pruning benchmarks. We highlight that FLoSS achieves performance comparable to dense training, even at high levels of sparsity, while outperforming other pruning benchmarks. The second section shows that we can retain this high level of performance across multiple heterogeneous settings. The third section serves to demonstrate the communication benefits of FLoSS by reducing communication time in realistic wireless networks. Finally, we use an implementation of DP-FedAdam [36] to show that FLoSS achieves a desirable privacy-utility tradeoff.

5.1 Model Performance with Fixed Communication

We compare the performance of models trained with FLoSS to dense LoRA training and other pruning methods in FL across our various configurations of datasets and model architectures. While dense adapter training will necessarily communicate the entire adapter, in order to compare model performance across communication-efficient pruning methods we fix a single rank and download/upload sparsity configuration across methods. For FLoSS we fix both the download and upload sparsity to be 0.25. Similarly, we have SparseAdapter prune to 0.25 sparsity at initialization. AdapterLTH is harder to fix in this same way as it iteratively prunes throughout training. As we train for 200 communication rounds, we set AdapterLTH to use an incremental sparsity ratio of 0.98. This way, on average across 200 communication rounds, AdapterLTH is communicating approximately 0.25 of the weights. Comparing in this way allows us to test model performance across the various training methods with fixed communication. For this reason, we also include results for dense LoRA with rank that is $\frac{1}{4}$ of the original rank as this is equivalent to communicating 0.25 of the weights with our original LoRA rank. While these numbers appear arbitrary, we note that the method is robust to multiple settings of these hyperparameters. Section 5.3 illustrates that FLoSS is effective even at more extreme rank and sparsity configurations.

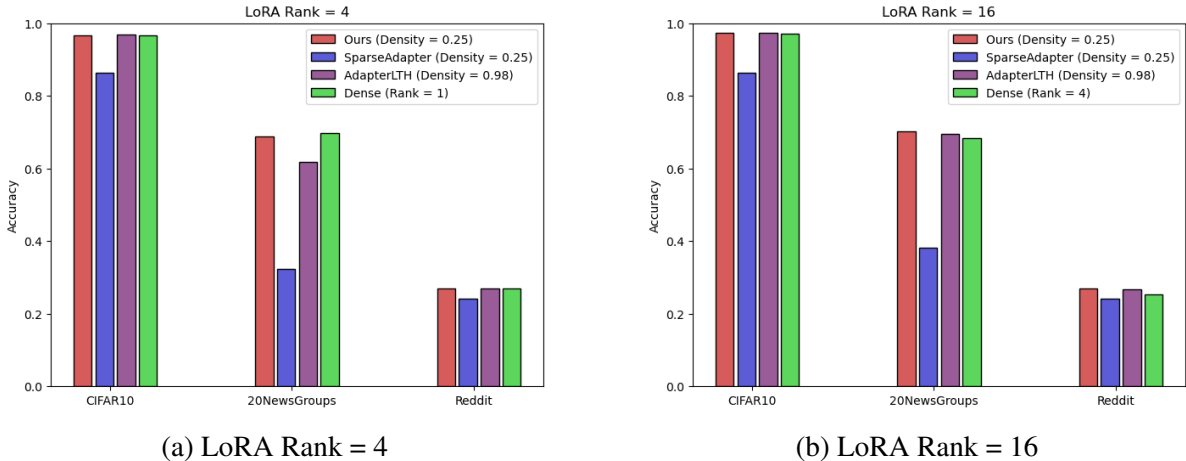


Figure 5.1: Comparison of communication-efficient LoRA methods in FL.

The results in figure 5.1 demonstrate that FLoSS outperforms other efficient fine-tuning methods in a federated setting. Specifically, FLoSS performs significantly better than the one-shot pruning method specified in SparseAdapter and marginally better than AdapterLTH and dense training of a LoRA adapter with $\frac{1}{4}$ rank. These findings are consistent in experiments in CIFAR10, 20NewsGroups, and Reddit, demonstrating that the method of sparsifying an adapter prior to download and upload is an effective way to reduce communication costs for multiple tasks. Additionally, this sparsification works even with smaller ranks (e.g. rank=4) meaning that we can get especially efficient communication by using a combination of LoRA where the rank is significantly smaller than the embedding dimension in conjunction with sparse updates. While we note that some of these performance benefits appear to be marginal at first, we show in later sections that FLoSS offers significant communication benefits in comparison to the other methods described as performance does not degrade with extreme levels of sparsity and different download/upload sparsity ratios.

5.2 Robustness to Statistical Heterogeneity

While Reddit naturally has non-I.I.D. clients as a consequence of being partitioned by user, CIFAR10 and 20NewsGroups do not have metadata regarding the user that produced certain images. This means that if the datasets are partitioned across clients in a random fashion, the clients will generally have I.I.D. data. This may not be representative of real-world FL environments where clients may have non-I.I.D. data, and training can be affected by statistical heterogeneity. To incorporate statistical heterogeneity for CIFAR10 and 20NewsGroups we partition the datasets across clients based on the ground-truth label for each image. We follow the procedure described in [22] using Dirichlet distributions with various α parameters to generate synthetic, non-I.I.D. client data from CIFAR10 and 20NewsGroups. In this context α refers to a concentration parameter that controls how identical sample distributions are across clients. A smaller α results in higher levels of statistical heterogeneity whereas a larger α means clients largely have




Figure 5.2: Comparison of communication-efficient LoRA methods in FL with non-I.I.D. clients.

identical data distributions.

Based on the results in figure 5.2, FLoSS performs well even in the presence of extreme statistical heterogeneity. For example, even at $\alpha = 0.1$ and $\alpha = 0.01$ where clients are primarily sampling examples from a single label, there is little degradation in the model’s utility in comparison to the I.I.D. setting. Another important observation is that [22] reports that statistical heterogeneity has an adverse impact on a CNN architecture trained on CIFAR10 for FL. In comparison, we find that fine-tuning a pretrained LLM on a similarly heterogeneous CIFAR10 partitioning for FL is not hindered by the same performance losses. This suggests that LLMs may offer a potential solution to the performance problems that arise as a result of statistical heterogeneity in FL.

5.3 Improvements to Communication Efficiency

While FLoSS marginally outperforms other efficient training methods on evaluation data, FLoSS improves the communication-efficiency of training significantly in comparison to other methods.

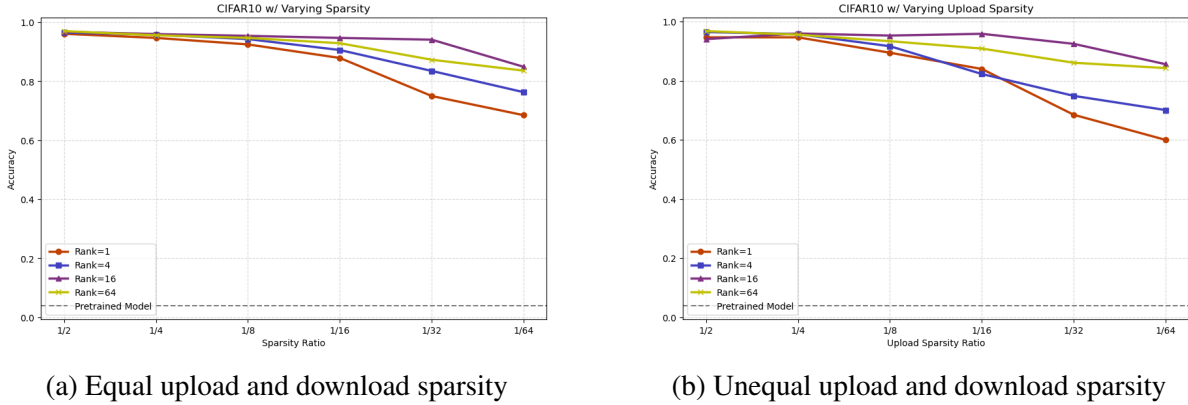


Figure 5.3: Performance of FLoSS on CIFAR10 with varying rank and sparsity configurations.

There are two critical reasons why FLoSS is able to reduce communication costs considerably. First, the method is able to utilize extremely sparse communication (e.g. approximately 0.01 sparsity ratio) and retain performance. In comparison, SparseAdapter noticeably degrades in performance even at sparsity ratios of 0.25 as demonstrated in section 5.1. AdapterLTH relies on pruning very few additional parameters at every communication round. However, this iterative pruning technique is not well-suited for extreme sparsity. In order for AdapterLTH to produce an average sparsity ratio of 0.01 across 200 communication rounds, the method would have to iteratively prune 50% of the weights at every communication round. In the last communication round, only $1.24e-60$ of the LoRA weights would remain, leading to a model that is no more accurate than the pretrained model. Figure 5.3a demonstrates that across multiple ranks, FLoSS can support extreme sparsity and result in a model significantly better than the pretrained model.

The second reason why FLoSS is able to significantly cut down on communication costs is because the method is able to define separate download and upload sparsity ratios. In wireless networks, upload is significantly more expensive than download (with upload sometimes being up to 10x slower) [42]. In cases where upload speed is significantly reduced in comparison to download speed, methods like AdapterLTH and SparseAdapter are too rigid in that AdapterLTH can only prune prior to download and SparseAdapter can only prune prior to model training. Figure 5.3b demonstrates the model performance of FLoSS trained with a fixed download sparsity of 0.25 and varying upload sparsity. In real-world settings where upload is much slower than download, we can use stricter sparsification of uploaded parameters. This prevents lags and delays during the upload phase of communication. Our results demonstrate that we can retain model performance with uneven download and upload sparsity configurations making FLoSS better suited for real-world FL configurations that have disparate download and upload speeds.

5.4 Privacy-Preserving Variations

A core goal of FL applications is to preserve the private information of clients. FL aims to mitigate the privacy and security concerns that arise with data aggregation, such as widespread

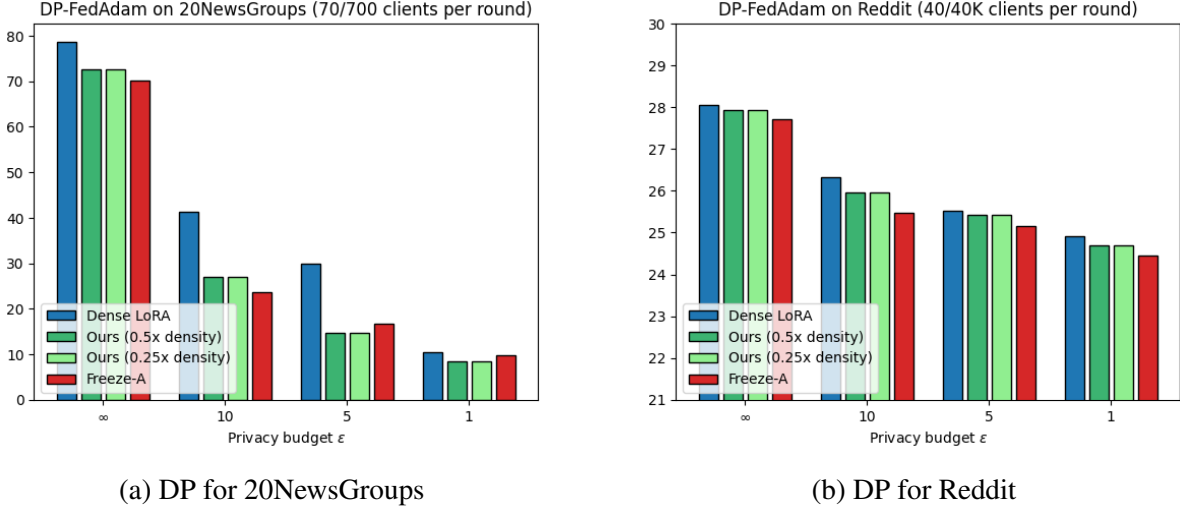


Figure 5.4: Performance of FLoSS with varying degrees of DP.

leaks of user data. For this reason, it is important to consider notions of privacy in our efficient LLM training method for FL. We apply the definition of differential privacy (DP) used in [36]. DP ensures user-level privacy, as opposed to example-level privacy, so that the parameters for the model can be publicly released with the guarantee that an adversary has a limited ability to learn about the data used to train the model. This is particularly important in FL as the model parameters are communicated between clients and a central server. An adversarial client or an adversary that has access to the server’s global model must be restricted in what they can learn about other clients’ private data.

We implement the DP aggregation algorithm described in [38]. This process includes two critical steps. First, when clients upload their local Δ_i , Δ_i is clipped by scaling each value by $\frac{1}{\max(1, \frac{\|\Delta_i\|_2}{C})}$ where C is the clipping norm. Noise is then added to this clipped delta to obscure the original weights. This method ensures user-level DP with certain guarantees based on the amount of noise added and the norm of the uploaded parameters. However, while smaller norms and more noise enforce stricter privacy guarantees, they can diminish the model’s utility. We analyze the performance of FLoSS using DP-FedAdam with varying levels of privacy guarantees. We further compare the method to FFA [44] which aims to preserve privacy for federated LoRA by freezing the LoRA A parameters at initialization. This method approximately halves the communication cost of LoRA and reduces the norm of uploaded model updates.

Figure 5.4 demonstrates that FLoSS can effectively preserve privacy while retaining performance with a variety of privacy budgets. Even though the method is not specifically designed for privacy-preserving FL, the method can be easily integrated with DP-FedAdam to ensure user-level privacy.

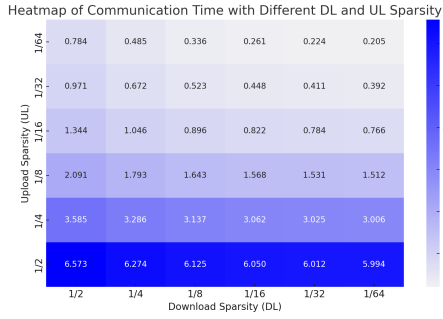
Chapter 6

Discussion and Future Work

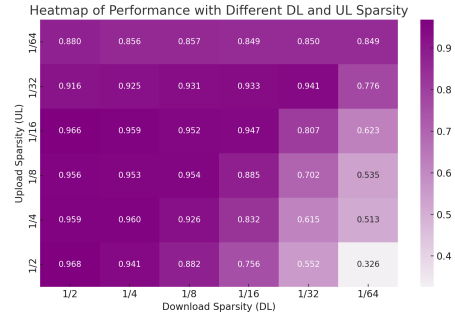
This thesis examines the problem of training LLMs in a federated environment. The scale of LLMs presents a problem in FL settings due to the limited compute on client devices and limited communication across the network. We present a few critical steps that can be taken to tackle these bottlenecks and enable efficient training. First, we employ adapter methods, specifically LoRA, to reduce the computational load of local model training. LoRA is especially useful in FL as client devices do not have to train full model parameters for the LLM, and the adapters can be merged back into the backbone of the model to eliminate additional storage and inference costs after training. We augment the communication efficiency of LoRA using sparsity applied only to communication. Crucially, we find that applying sparsity in this way is important to retain performance. Alternative methods that prune adapters at initialization or iteratively prune adapters throughout training do not perform as well with limited communication budgets or cases where upload speed is slower than download speed. These limitations suggest that sparsity must be carefully applied during the training process to ensure that communication efficiency does not degrade model utility. A few crucial directions remain to be explored in the context of efficient LLM training for FL. We describe some of these future directions below.

Local Compute While LoRA certainly makes local computation easier for client devices, methods to further enable LLM training in resource-constrained environments may improve the functionality of this method. Some recent research has focused on the importance of efficient on-device training [33]. A few methods are natural choices for improving the on-device efficiency of FLoSS. For example, approaches like QLoRA [12] and LoftQ [32] significantly reduce memory usage in LoRA training by quantizing the model backbone. These methods could be integrated into FLoSS to further augment the on-device efficiency of the method.

Rank and Sparsity Configurations Currently, FLoSS requires manual selection of the rank and sparsity configurations used in testing. Methods to automatically configure the rank and sparsity ratios for FLoSS are important in preventing expensive FL hyperparameter tuning. Our experiments suggest that certain rank and sparsity configurations do not perform as well as others within the same communication budget. For example, Figure 6.1 demonstrates that configurations with similar performance can have vastly overall communication times. This is especially true in our simulated procedure where upload speed is $10\times$ slower than download speed (20



(a) Communication time heatmap



(b) Performance heatmap

Figure 6.1: Comparison of communication time vs. performance for FLoSS on CIFAR10

MBps vs. 200 MBps). Therefore, being able to automatically determine the optimal value for these hyperparameters prior to training may be crucial in extracting the full value of LLMs in this setting.

FL is an increasingly important paradigm in the context of privacy, personalization, and scalability for ML models. As LLMs continue to outperform other methods in a variety of domains, exploring ways LLMs can be effectively utilized in FL is important in handling issues of model utility and statistical heterogeneity. This thesis suggests concrete first steps towards achieving this goal. In the future, we aim to investigate such questions and design methods to make high-quality models more accessible to low-resource users.

Bibliography

- [1] Ahmed A. Al-Saedi, Veselka Boeva, and Emiliano Casalicchio. Reducing communication overhead of federated learning through clustering analysis. In *2021 IEEE Symposium on Computers and Communications (ISCC)*, pages 1–7, 2021. doi: 10.1109/ISCC53001.2021.9631391. 3.1
- [2] Mohammad Mohammadi Amiri, Deniz Gunduz, Sanjeev R Kulkarni, and H Vincent Poor. Federated learning with quantized global model updates. *arXiv preprint arXiv:2006.10672*, 2020. 3.1
- [3] Elad Ben Zaken, Yoav Goldberg, and Shauli Ravfogel. BitFit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In Smaranda Muresan, Preslav Nakov, and Aline Villavicencio, editors, *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1–9, Dublin, Ireland, May 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.acl-short.1. URL <https://aclanthology.org/2022.acl-short.1>. 3.2
- [4] Neelkamal Bhuyan, Sharayu Moharir, and Gauri Joshi. Multi-model federated learning with provable guarantees. In Esa Hyttiä and Veeraruna Kavitha, editors, *Performance Evaluation Methodologies and Tools*, pages 207–222, Cham, 2023. Springer Nature Switzerland. ISBN 978-3-031-31234-2. 2.1.1
- [5] Keith Bonawitz, Vladimir Ivanov, Ben Kreuter, Antonio Marcedone, H. Brendan McMahan, Sarvar Patel, Daniel Ramage, Aaron Segal, and Karn Seth. Practical secure aggregation for privacy-preserving machine learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS '17*, page 1175–1191, New York, NY, USA, 2017. Association for Computing Machinery. ISBN 9781450349468. doi: 10.1145/3133956.3133982. URL <https://doi.org/10.1145/3133956.3133982>. 1
- [6] Keith Bonawitz, Hubert Eichner, Wolfgang Grieskamp, Dzmitry Huba, Alex Ingerman, Vladimir Ivanov, Chloé Kiddon, Jakub Konečný, Stefano Mazzocchi, Brendan McMahan, Timon Van Overveldt, David Petrou, Daniel Ramage, and Jason Roselander. Towards federated learning at scale: System design. In A. Talwalkar, V. Smith, and M. Zaharia, editors, *Proceedings of Machine Learning and Systems*, volume 1, pages 374–388, 2019. URL https://proceedings.mlsys.org/paper_files/paper/2019/file/7b770da633baf74895be22a8807f1a8f-Paper.pdf. 1
- [7] Christopher Briggs, Zhong Fan, and Peter Andras. Federated learning with hierarchical clustering of local updates to improve training on non-iid data. In *2020 International Joint*

Conference on Neural Networks (IJCNN), pages 1–9. IEEE, 2020. 3.1

- [8] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf. 1, 2.2
- [9] Mingzhe Chen, Nir Shlezinger, H. Vincent Poor, Yonina C. Eldar, and Shuguang Cui. Communication-efficient federated learning. *Proceedings of the National Academy of Sciences*, 118(17):e2024789118, 2021. doi: 10.1073/pnas.2024789118. URL <https://www.pnas.org/doi/abs/10.1073/pnas.2024789118>. 3.1
- [10] Hongrong Cheng, Miao Zhang, and Javen Qinfeng Shi. A survey on deep neural network pruning-taxonomy, comparison, analysis, and recommendations. *arXiv preprint arXiv:2308.06767*, 2023. 2.3
- [11] Don Kurian Dennis, Tian Li, and Virginia Smith. Heterogeneity for the win: One-shot federated clustering. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 2611–2620. PMLR, 18–24 Jul 2021. URL <https://proceedings.mlr.press/v139/dennis21a.html>. 3.1
- [12] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms. *Advances in Neural Information Processing Systems*, 36, 2024. 6
- [13] 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. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>. 2.2
- [14] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel Roy, and Michael Carbin. Pruning neural networks at initialization: Why are we missing the mark? In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=Ig-VyQc-MLK>. 2.3
- [15] Avishek Ghosh, Jichan Chung, Dong Yin, and Kannan Ramchandran. An efficient framework for clustered federated learning. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 19586–19597. Curran Associates, Inc., 2020. URL https://proceedings.neurips.cc/paper_files/paper/

2020/file/e32cc80bf07915058ce90722ee17bb71-Paper.pdf. 2.1.1

- [16] Neel Guha, Ameet Talwalkar, and Virginia Smith. One-shot federated learning. *arXiv preprint arXiv:1902.11175*, 2019. 3.1
- [17] Jenny Hamer, Mehryar Mohri, and Ananda Theertha Suresh. FedBoost: A communication-efficient algorithm for federated learning. In Hal Daumé III and Aarti Singh, editors, *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pages 3973–3983. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/hamer20a.html>. 3.1
- [18] Zeyu Han, Chao Gao, Jinyang Liu, Sai Qian Zhang, et al. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*, 2024. 2.2
- [19] Shwai He, Liang Ding, Daize Dong, Jeremy Zhang, and Dacheng Tao. SparseAdapter: An easy approach for improving the parameter-efficiency of adapters. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang, editors, *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2184–2190, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.findings-emnlp.160. URL <https://aclanthology.org/2022.findings-emnlp.160>. 3.3
- [20] Agrin Hilmkil, Sebastian Callh, Matteo Barbieri, Leon René Sütfield, Edvin Listo Zec, and Olof Mogren. Scaling federated learning for fine-tuning of large language models. In Elisabeth Métais, Farid Meziane, Helmut Horacek, and Epaminondas Kapetanios, editors, *Natural Language Processing and Information Systems*, pages 15–23, Cham, 2021. Springer International Publishing. ISBN 978-3-030-80599-9. 3.2
- [21] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Larousilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for NLP. In Kamalika Chaudhuri and Ruslan Salakhutdinov, editors, *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/houlsby19a.html>. 1, 3.2
- [22] Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. 2019. URL <https://arxiv.org/abs/1909.06335>. 5.2, 5.2
- [23] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022. URL <https://openreview.net/forum?id=nZeVKeeFYf9>. 1, 2.2, 3.2
- [24] Dzmitry Huba, John Nguyen, Kshitiz Malik, Ruiyu Zhu, Mike Rabbat, Ashkan Yousefpour, Carole-Jean Wu, Hongyuan Zhan, Pavel Ustinov, Harish Srinivas, et al. Papaya: Practical, private, and scalable federated learning. *Proceedings of Machine Learning and Systems*, 4: 814–832, 2022. 1
- [25] Sohei Itahara, Takayuki Nishio, Yusuke Koda, Masahiro Morikura, and Koji Yamamoto. Distillation-based semi-supervised federated learning for communication-efficient collaborative training with non-iid private data. *IEEE Transactions on Mobile Computing*, 22

- (1):191–205, January 2023. ISSN 2161-9875. doi: 10.1109/tmc.2021.3070013. URL <http://dx.doi.org/10.1109/TMC.2021.3070013>. 3.1
- [26] Yuang Jiang, Shiqiang Wang, Víctor Valls, Bong Jun Ko, Wei-Han Lee, Kin K. Leung, and Leandros Tassiulas. Model pruning enables efficient federated learning on edge devices. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12):10374–10386, 2023. doi: 10.1109/TNNLS.2022.3166101. 3.1
- [27] Zhida Jiang, Yang Xu, Hongli Xu, Zhiyuan Wang, Chunming Qiao, and Yangming Zhao. Fedmp: Federated learning through adaptive model pruning in heterogeneous edge computing. In *2022 IEEE 38th International Conference on Data Engineering (ICDE)*, pages 767–779, 2022. doi: 10.1109/ICDE53745.2022.00062. 3.1
- [28] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 2.2
- [29] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacL-HLT*, volume 1, page 2, 2019. 1, 2.2
- [30] Kevin Kuo, Arian Raje, Kousik Rajesh, and Virginia Smith. Sparsity for communication-efficient loRA. In *5th Workshop on practical ML for limited/low resource settings*, 2024. URL <https://openreview.net/forum?id=wibit67d29>. 4.3
- [31] Brian Lester, Rami Al-Rfou, and Noah Constant. The power of scale for parameter-efficient prompt tuning. In Marie-Francine Moens, Xuanjing Huang, Lucia Specia, and Scott Wen-tau Yih, editors, *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.emnlp-main.243. URL <https://aclanthology.org/2021.emnlp-main.243>. 3.2
- [32] Yixiao Li, Yifan Yu, Chen Liang, Nikos Karampatziakis, Pengcheng He, Weizhu Chen, and Tuo Zhao. Loftq: LoRA-fine-tuning-aware quantization for large language models. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=LzPWWPAdY4>. 6
- [33] Ji Lin, Ligeng Zhu, Wei-Ming Chen, Wei-Chen Wang, Chuang Gan, and Song Han. On-device training under 256kb memory. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NeurIPS '22*, Red Hook, NY, USA, 2024. Curran Associates Inc. ISBN 9781713871088. 6
- [34] Xiaonan Liu, Shiqiang Wang, Yansha Deng, and Arumugam Nallanathan. Adaptive federated pruning in hierarchical wireless networks. *IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS*, page 1, 2023. ISSN 1536-1276. doi: 10.1109/TWC.2023.3329450. Publisher Copyright: IEEE. 3.1
- [35] H. Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. 1, 2.1.1, 2.1.2, 3.1

- [36] H. Brendan McMahan, Daniel Ramage, Kunal Talwar, and Li Zhang. Learning differentially private recurrent language models. In *International Conference on Learning Representations*, 2018. URL <https://openreview.net/forum?id=BJ0hF1Z0b>. 5, 5.4
- [37] Michela Paganini and Jessica Forde. On iterative neural network pruning, reinitialization, and the similarity of masks. *arXiv preprint arXiv:2001.05050*, 2020. 2.3
- [38] Natalia Ponomareva, Hussein Hazimeh, Alex Kurakin, Zheng Xu, Carson Denison, H. Brendan McMahan, Sergei Vassilvitskii, Steve Chien, and Abhradeep Guha Thakurta. How to dp-fy ml: A practical guide to machine learning with differential privacy. *Journal of Artificial Intelligence Research*, 77:1113–1201, July 2023. ISSN 1076-9757. doi: 10.1613/jair.1.14649. URL <http://dx.doi.org/10.1613/jair.1.14649>. 5.4
- [39] Libo Qin, Qiguang Chen, Yuhang Zhou, Zhi Chen, Yinghui Li, Lizi Liao, Min Li, Wanxiang Che, and Philip S Yu. Multilingual large language model: A survey of resources, taxonomy and frontiers. *arXiv preprint arXiv:2404.04925*, 2024. 2.2
- [40] Sashank Reddi, Zachary Burr Charles, Manzil Zaheer, Zachary Garrett, Keith Rush, Jakub Konečný, Sanjiv Kumar, and Brendan McMahan, editors. *Adaptive Federated Optimization*, 2021. URL <https://openreview.net/forum?id=LkFG31B13U5>. 2.1.2
- [41] Amirhossein Reiszadeh, Aryan Mokhtari, Hamed Hassani, Ali Jadbabaie, and Ramtin Pedarsani. Fedpaq: A communication-efficient federated learning method with periodic averaging and quantization. In Silvia Chiappa and Roberto Calandra, editors, *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2021–2031. PMLR, 26–28 Aug 2020. URL <https://proceedings.mlr.press/v108/reiszadeh20a.html>. 3.1
- [42] Osama Shahid, Seyedamin Pouriyeh, Reza M. Parizi, Quan Z. Sheng, Gautam Srivastava, and Liang Zhao. Communication efficiency in federated learning: Achievements and challenges, 2021. 5.3
- [43] Guangyu Sun, Umar Khalid, Matias Mendieta, Taojiannan Yang, and Chen Chen. Conquering the communication constraints to enable large pre-trained models in federated learning. *arXiv preprint arXiv:2210.01708*, 2024. 3.2
- [44] Youbang Sun, Zitao Li, Yaliang Li, and Bolin Ding. Improving loRA in privacy-preserving federated learning. In *The Twelfth International Conference on Learning Representations*, 2024. URL <https://openreview.net/forum?id=NLPzL6HWN1>. 5.4
- [45] Alysa Ziyang Tan, Han Yu, Lizhen Cui, and Qiang Yang. Towards personalized federated learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(12): 9587–9603, December 2023. ISSN 2162-2388. doi: 10.1109/tnnls.2022.3160699. URL <http://dx.doi.org/10.1109/TNNLS.2022.3160699>. 1
- [46] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023. 2.2

- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf. 1, 2.2
- [48] Luping WANG, Wei WANG, and Bo LI. Cmfl: Mitigating communication overhead for federated learning. In *2019 IEEE 39th International Conference on Distributed Computing Systems (ICDCS)*, pages 954–964, 2019. doi: 10.1109/ICDCS.2019.00099. 3.1
- [49] Wei Wang and Liqiang Zhu. Structured feature sparsity training for convolutional neural network compression. *Journal of Visual Communication and Image Representation*, 71:102867, 2020. ISSN 1047-3203. doi: <https://doi.org/10.1016/j.jvcir.2020.102867>. URL <https://www.sciencedirect.com/science/article/pii/S1047320320301176>. 2.3
- [50] Wenxiao Wang, Wei Chen, Yicong Luo, Yongliu Long, Zhengkai Lin, Liye Zhang, Binbin Lin, Deng Cai, and Xiaofei He. Model compression and efficient inference for large language models: A survey. *arXiv preprint arXiv:2402.09748*, 2024. 2.3
- [51] Yaqing Wang, Sahaj Agarwal, Subhabrata Mukherjee, Xiaodong Liu, Jing Gao, Ahmed Hassan Awadallah, and Jianfeng Gao. AdaMix: Mixture-of-adaptations for parameter-efficient model tuning. In Yoav Goldberg, Zornitsa Kozareva, and Yue Zhang, editors, *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5744–5760, Abu Dhabi, United Arab Emirates, December 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.emnlp-main.388. URL <https://aclanthology.org/2022.emnlp-main.388>. 1
- [52] Ziheng Wang, Jeremy Wohlwend, and Tao Lei. Structured pruning of large language models. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Association for Computational Linguistics, 2020. doi: 10.18653/v1/2020.emnlp-main.496. URL <http://dx.doi.org/10.18653/v1/2020.emnlp-main.496>. 2.3
- [53] Chuhan Wu, Fangzhao Wu, Lingjuan Lyu, Yongfeng Huang, and Xing Xie. Communication-efficient federated learning via knowledge distillation. *Nature Communications*, 13(1), April 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-29763-x. URL <http://dx.doi.org/10.1038/s41467-022-29763-x>. 3.1
- [54] Jiarun Wu and Qingliang Chen. Pruning adapters with lottery ticket. *Algorithms*, 15(2), 2022. ISSN 1999-4893. doi: 10.3390/a15020063. URL <https://www.mdpi.com/1999-4893/15/2/63>. 3.3
- [55] Guang Yang, Ke Mu, Chunhe Song, Zhijia Yang, and Tierui Gong. Ringfed: Reducing communication costs in federated learning on non-iid data. *arXiv preprint arXiv:2107.08873*, 2021. 3.1
- [56] Jingfeng Yang, Hongye Jin, Ruixiang Tang, Xiaotian Han, Qizhang Feng, Haoming Jiang, Shaochen Zhong, Bing Yin, and Xia Hu. Harnessing the power of llms in practice: A

survey on chatgpt and beyond. *ACM Trans. Knowl. Discov. Data*, 18(6), apr 2024. ISSN 1556-4681. doi: 10.1145/3649506. URL <https://doi.org/10.1145/3649506>. 2.2

- [57] Zhengwu Yang and Han Zhang. Comparative analysis of structured pruning and unstructured pruning. In Jason C. Hung, Neil Y. Yen, and Jia-Wei Chang, editors, *Frontier Computing*, pages 882–889, Singapore, 2022. Springer Nature Singapore. ISBN 978-981-16-8052-6. 2.3
- [58] Zhuo Zhang, Yuanhang Yang, Yong Dai, Qifan Wang, Yue Yu, Lizhen Qu, and Zenglin Xu. FedPETuning: When federated learning meets the parameter-efficient tuning methods of pre-trained language models. In Anna Rogers, Jordan Boyd-Graber, and Naoaki Okazaki, editors, *Findings of the Association for Computational Linguistics: ACL 2023*, pages 9963–9977, Toronto, Canada, July 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-acl.632. URL <https://aclanthology.org/2023.findings-acl.632>. 3.2
- [59] Yong Zhou, Yuanming Shi, Haibo Zhou, Jingjing Wang, Liqun Fu, and Yang Yang. Toward scalable wireless federated learning: Challenges and solutions. *IEEE Internet of Things Magazine*, 6(4):10–16, 2023. doi: 10.1109/IOTM.001.2300099. 1
- [60] Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. A survey on model compression for large language models. *arXiv preprint arXiv:2308.07633*, 2023. 2.3
- [61] Yuchen Zhuang, Yue Yu, Kuan Wang, Haotian Sun, and Chao Zhang. ToolQA: A dataset for LLM question answering with external tools. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. URL <https://openreview.net/forum?id=pV1xV2RK6I>. 2.2