

Federated Sketching LoRA: A Flexible Framework for Heterogeneous Collaborative Fine-Tuning of LLMs

Wenzhi Fang¹ Dong-Jun Han² Liangqi Yuan¹ Seyyedali Hosseinalipour³ Christopher G. Brinton¹

Abstract

Fine-tuning large language models (LLMs) on resource-constrained clients remains a challenging problem. Recent works have fused low-rank adaptation (LoRA) techniques with federated fine-tuning to mitigate challenges associated with client model sizes and data scarcity. Still, the heterogeneity of resources remains a critical bottleneck: while higher-rank modules generally enhance performance, varying client capabilities constrain LoRA’s feasible rank range. Existing approaches attempting to resolve this issue either lack analytical justification or impose additional computational overhead, leaving a wide gap for efficient and theoretically-grounded solutions. To address these challenges, we propose federated sketching LoRA (FSLoRA), which leverages a sketching mechanism to enable clients to selectively update submatrices of global LoRA modules maintained by the server. By adjusting the sketching ratios, which determine the ranks of the submatrices on the clients, FSLoRA flexibly adapts to client-specific communication and computational constraints. We provide a rigorous convergence analysis of FSLoRA that characterizes how the sketching ratios affect the convergence rate. Through extensive experiments, we demonstrate that FSLoRA outperforms baselines and significantly improves training efficiency while preserving stable convergence.

1. Introduction

Lightweight client-side large language models (LLMs) have recently gained significant attention as a promising complement to cloud-based LLMs (Fan et al., 2024). They align

¹Purdue University ²Yonsei University ³University at Buffalo-SUNY. Correspondence to: Wenzhi Fang <fang375@purdue.edu>.

with the typical paradigm of LLMs: starting from a base model pre-trained on large-scale datasets to learn general linguistic patterns, semantics, and context, and then undergoing fine-tuning on task-specific data to enhance performance on specialized or domain-specific applications. However, an LLM fine-tuned on a single client often achieves unsatisfactory performance due to the limited data. Federated learning (FL) (McMahan et al., 2017; Yao et al., 2024) has been investigated as a potential solution here, enabling the model to be fine-tuned across a group of distributed clients within the same task domain, without any raw data sharing.

However, federated LLM fine-tuning is costly in both computation and communication due to the massive parameter volume. Importantly, many parameter-efficient fine-tuning methods have been proposed (Lester et al., 2021; Li & Liang, 2021; Hu et al., 2021) to reduce the model adaptation cost. Among them, low-rank adaptation (LoRA) (Hu et al., 2021) stands out as a particularly effective approach due to its flexibility. In particular, LoRA enables efficient fine-tuning by approximating weight updates $\Delta \mathbf{W}$ through a low-rank decomposition $\Delta \mathbf{W} = \mathbf{B}\mathbf{A}$, where matrices $\mathbf{B}$ and $\mathbf{A}$ contain significantly fewer trainable parameters than the original weight matrix. Building on this foundation, recent works incorporated LoRA with federated averaging (FedAvg) to support collaborative LLM fine-tuning (Zhang et al., 2024; Ye et al., 2024; Zou et al., 2025), significantly reducing the fine-tuning cost by cutting down the number of parameters that need to be synchronized across distributed clients.

Challenges. While incorporating LoRA into federated LLM fine-tuning reduces the number of trainable parameters, *computation and communication costs are still forced to increase with the LoRA rank*. This poses challenges when complex tasks demand higher-rank LoRA modules, particularly on resource-constrained clients. Furthermore, the *heterogeneity in resource availability across distributed clients makes a uniform rank adopted in federated LoRA inefficient*: a fixed rank r may be too large for some constrained clients, while being too small for more powerful ones, resulting in underutilized resources. Consequently, an approach that further reduces computation and communication overhead while adapting LoRA ranks to heterogeneous client capabilities is highly desirable. Although some existing approaches

(Cho et al., 2024; Bai et al., 2024; Wang et al., 2024) have attempted to provide a solution, they either lack theoretical justification or impose additional computational overhead, leaving a gap for an efficient and theoretically-grounded solution. A detailed discussion of these limitations is provided in Section 2.2. Overall, a comprehensive approach that preserves the analytical and practical benefits of LoRA while allowing heterogeneous collaborative fine-tuning under tight client resource constraints remains elusive.

1.1. Contributions

Motivated by these limitations, this work develops a methodology for efficient federated LLM fine-tuning that (i) retains the flexibility of LoRA, (ii) provides theoretical convergence guarantees, and (iii) addresses the challenges posed by heterogeneous and constrained resources across distributed clients. As depicted in Figure 1, our key idea is to introduce a sketching-based LoRA update to the local fine-tuning, which allows clients to selectively update a subset of columns and rows of the LoRA modules during each round, reducing the computation and communication consumption. Additionally, our method customizes the fine-tuning process by adjusting the sparsity level of the sketching matrix, i.e., the size of the updated submatrices for each client in each iteration. As we will see, the impact of the introduced sketching mechanism on the overall optimization landscape requires careful modeling consideration, posing additional challenges for the theoretical analysis that we address in this work.

In summary, we make the following contributions:

- We propose federated sketching LoRA (FSLoRA), which leverages a sketching mechanism to enable clients to selectively update submatrices of global LoRA modules maintained by the server. By adjusting the sketching ratios, which determine the ranks of the submatrices on clients, FSLoRA effectively adapts to client-specific communication and computational constraints.
- We present a rigorous convergence analysis of FSLoRA under non-uniform submatrix updates across clients (i.e., heterogeneous LoRA configurations), revealing how the sketching ratios (i.e., resource heterogeneity) affect the convergence rate through scaled smoothness constants. Our analysis establishes convergence under arbitrary sketching ratios and recovers the convergence rate of FedAvg in the limiting case.
- We conduct extensive experiments across multiple datasets and LLM models with diverse parameter settings, demonstrating FSLoRA’s superior performance compared to various baselines in accuracy, training efficiency, and resource utilization. Our ablation studies further validate




Figure 1. An illustration of our proposed methodology where the server maintains a pair of global LoRA modules while the clients adaptively update submatrices of the global LoRA modules through sketching in each round.

the effectiveness of the sketching mechanism and the ability of clients to exploit larger global ranks under FSLoRA.

1.2. Related Works

Collaborative fine-tuning via federated LoRA: Federated LoRA is an efficient approach for collaborative LLM fine-tuning among distributed clients (Yao et al., 2024; Sun et al., 2024; Guo et al., 2025; Chen et al., 2025; Guo et al., 2024). Building on this foundation, Kuo et al. (2024) proposed integrating communication compression with federated LoRA to further reduce communication overhead. Meanwhile, Bai et al. (2024); Cho et al. (2024); Byun & Lee (2024); Wang et al. (2024); Koo et al. (2024) explored the challenges of resource heterogeneity across distributed clients and introduced heterogeneous LoRA as a solution. However, the approaches proposed in (Cho et al., 2024; Koo et al., 2024; Byun & Lee, 2024; Bai et al., 2024) lack a theoretical foundation. Moreover, the FlexLoRA method introduced in (Bai et al., 2024) incurs additional computational overhead due to its reliance on singular value decomposition (SVD). Furthermore, the FLoRA algorithm proposed in (Wang et al., 2024) requires the clients to merge the LoRA modules into the base model, thereby compromising the inherent flexibility of LoRA. Overall, there is still a lack of a systematic and theoretically grounded solution that can effectively tackle heterogeneous collaborative LLM fine-tuning.

Enhancing adaptability via higher-rank LoRA modules:

The foundational study by Hu et al. (2021) demonstrated that small ranks can be sufficient for certain tasks; however, they also acknowledge that small rank LoRA modules may not work universally, especially when the downstream task differs significantly from pretraining. Following this,

several works explored the effect of increasing the rank in LoRA modules. In a centralized setup, Kalajdzievski (2023) and Shuttleworth et al. (2024) showed that higher-rank LoRA models can closely approximate full fine-tuning under rsLoRA. In a federated LLM fine-tuning regime, Bai et al. (2024) demonstrated improved performance with larger ranks under FlexLoRA. Similarly, Cho et al. (2024) reported that, with proper overfitting control, HeteroLoRA can also benefit from larger ranks. Orthogonal to these rank-scaling approaches, Uni-LoRA (Li et al., 2026) reduces the trainable parameter space by reconstructing LoRA parameters from a projected low-dimensional subspace, highlighting another direction for improving LoRA efficiency. Overall, while small ranks may suffice for simpler tasks or strong base models, higher-rank modules are necessary to compensate for limited backbone capability, such as in lightweight LLMs, and to enable effective adaptation to more complex downstream tasks.

Sketching-based optimization: Sketching is an efficient technique for mitigating the complexity of high-dimensional optimization, with its earliest applications in least-squares regression (Sarlos, 2006; Wang et al., 2022). Beyond this, gradient sketching has been employed to construct preconditioners for gradient descent methods (Gower & Richtárik, 2015). Building on these foundations, recent work has applied sketching to distributed optimization. In particular, Charalambides & Mazumdar (2024) proposed hybrid local-global sketching for distributed least-squares, while Demidovich et al. (2023) developed a distributed sparsified training framework based on sketching. Shrivastava et al. (2024) demonstrated that sketching substantially reduces communication in distributed training of overparameterized deep models without sacrificing accuracy. More recently, Nicolas et al. (2026) investigated sketching-based differential privacy and demonstrated its compatibility with secure aggregation. Despite these advances, sketching strategies tailored to structured low-rank adaptation modules such as LoRA remain largely unexplored.

2. Problem Background

2.1. LoRA-based Federated LLM Fine-tuning

The federated LoRA fine-tuning problem can be formulated as

$$\min_{\mathbf{B}, \mathbf{A}} f(\mathbf{B}, \mathbf{A}) := \frac{1}{N} \sum_{i=1}^N f_i(\mathbf{B}, \mathbf{A}) \quad (1)$$

$$f_i(\mathbf{B}, \mathbf{A}) := \mathbb{E}_{\xi \sim \mathcal{D}_i} [\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{A}, \xi)],$$

where $\mathbf{W}_0$ denotes the frozen base model, $\mathbf{B} \in \mathbb{R}^{m \times r}$, $\mathbf{A} \in \mathbb{R}^{r \times n}$ are LoRA modules, N denotes the number of clients, ξ denotes a data sample, and $\mathcal{D}_i$ is the local dataset on client i . ℓ , f_i , and f are the sample loss function, the local loss for client i , and the global loss, respectively.

Problem (1) aligns with the conventional federated optimization formulation, which thus can be solved using the FedAvg algorithm. Based on the FedAvg framework, Zhang et al. (2024) developed federated LoRA, which applies a uniform rank r across all clients, overlooking resource heterogeneity. This one-size-fits-all approach leads to resource mismatches, where computationally constrained clients may struggle, while more powerful clients remain underutilized with a fixed rank.

2.2. Aren't the Existing Solutions Good Enough?

To address this issue, researchers have proposed heterogeneous federated LoRA approaches, where clients maintain non-uniform LoRA modules with varying ranks. They also introduce mechanisms to overcome the challenges of directly aggregating matrices with different dimensions. However, these methods often lack theoretical foundation or incur additional computational and memory overhead.

HeteroLoRA (Cho et al., 2024) lets the server pad the updates from the clients with smaller ranks to match the size of the largest rank during aggregation. During model dissemination, clients receive a truncated version of the global LoRA modules from the server. Although easy to implement, HeteroLoRA is primarily heuristic in nature and lacks a rigorous theoretical foundation, potentially limiting its performance, as we will see in Section 5.

FlexLoRA (Bai et al., 2024) requires the server to collect the individual LoRA matrices $\mathbf{B}_i$ and $\mathbf{A}_i$ from the clients and then computes their product $\mathbf{B}_i \mathbf{A}_i$. To support the initialization of non-uniform LoRA modules, the server applies truncated SVD to the averaged product $\frac{1}{N} \sum_{i=1}^N \mathbf{B}_i \mathbf{A}_i$. However, this approach introduces extra computational and memory overhead on the server due to truncated SVD, and the associated error can limit the performance as demonstrated in Section 5.

FLoRA (Wang et al., 2024) introduces a stacking mechanism where the server concatenates LoRA modules from the clients. The concatenated matrices are then sent back to the clients, which compute their product and merge it into the base model before initializing new LoRA modules for the next fine-tuning round. However, this approach increases communication complexity linearly with the number of clients, imposes higher computation and memory demands on the clients, and compromises LoRA's flexibility to support multiple adapters for different tasks.

More detailed comparisons on computation, memory, and communication are presented in Appendix A. In summary, a theoretically-grounded solution that preserves LoRA's benefits while effectively addressing resource heterogeneity across distributed clients remains lacking.

3. Federated Sketching LoRA

Motivated by the limitations of existing methods, we propose a reformulation for federated LoRA. Building on this foundation, we develop FSLoRA, a heterogeneous LoRA algorithm that preserves LoRA’s flexibility while accommodating client resource heterogeneity.

3.1. Our Formulation

We propose a sketching-based LoRA formulation for collaborative LLM fine-tuning as follows:

$$\min_{\mathbf{B}, \mathbf{A}} f^{\mathcal{S}}(\mathbf{B}, \mathbf{A}) := \frac{1}{N} \sum_{i=1}^N f_i^{\mathcal{S}}(\mathbf{B}, \mathbf{A}) \quad (2)$$

$$f_i^{\mathcal{S}}(\mathbf{B}, \mathbf{A}) := \mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i; \xi \sim \mathcal{D}_i} [\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)],$$

where $\mathbf{B} \in \mathbb{R}^{m \times r}$, $\mathbf{A} \in \mathbb{R}^{r \times n}$ are LoRA modules, $f_i^{\mathcal{S}}$ is the local loss function at client i with sketching, and $\mathbf{S}$ denotes a sketching matrix randomly sampled from the diagonal matrix set $\mathcal{S}_i = \mathcal{S}(r, k_i)$. The set $\mathcal{S}(r, k_i)$ comprises diagonal matrices of size $r \times r$ with exactly k_i non-zero entries. The formal definition of $\mathcal{S}(r, k)$ is provided below:

Definition 3.1 (Random- k sketching). A random- k diagonal matrix set is defined as:

$$\mathcal{S}(r, k) = \left\{ \mathbf{S} \mid \mathbf{S} = \frac{r}{k} \sum_{j \in \mathcal{I}} \mathbf{e}_j \mathbf{e}_j^\top, \mathcal{I} \subseteq \{1, \dots, r\}, |\mathcal{I}| = k \right\},$$

where $\mathbf{e}_1, \dots, \mathbf{e}_r \in \mathbb{R}^r$ are standard unit basis vectors and index set $\mathcal{I}$ is a random subset of $[r] := \{1, 2, \dots, r\}$ sampled uniformly from all subsets of $[r]$ with cardinality k .

With $\mathbf{S}$ being a matrix sampled from $\mathcal{S}_i$, we have $\mathbf{B}\mathbf{S}\mathbf{A} = \frac{r}{k_i} \sum_{j \in \mathcal{I}_i} \mathbf{B}\mathbf{e}_j \mathbf{e}_j^\top \mathbf{A}$, where $\mathcal{I}_i$ corresponds to the index set of non-zero diagonal entries of $\mathbf{S}$. $\mathbf{B}\mathbf{e}_j$ extracts the j -th column of $\mathbf{B}$ while $\mathbf{e}_j^\top \mathbf{A}$ extracts the j -th row of $\mathbf{A}$. In other words, only k_i columns and rows in the LoRA modules $\mathbf{B}$ and $\mathbf{A}$ are activated by the sketching matrix in the loss $\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ at client i . On the other hand, the sketching matrix $\mathbf{S}$ satisfies $\mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i} [\mathbf{S}] = \mathbf{I}_r$ where $\mathbf{I}_r$ is a r -dimensional identity matrix. Based upon this property, $\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}$ can be treated as an unbiased estimate of $\mathbf{W}_0 + \mathbf{B}\mathbf{A}$.

Intuition: A larger rank allows LoRA modules to be more expressive, leading to better performance (Bai et al., 2024; Kalajdziewski, 2023; Shuttleworth et al., 2024). However, resource-constrained clients cannot afford the computational or communication demands of large-rank modules. Our formulation (2) leverages the sketching matrix to balance the expressiveness of high-rank LoRA modules with the resource constraints of different clients. With the sketching mechanism introduced, the local gradients with respect to the LoRA modules on the clients will exhibit structured

sparsity. By adjusting the sketching ratio k_i/r , we can tailor the sparsity of the gradient to match the capabilities of each client, ensuring affordable training while maintaining performance across heterogeneous systems, as elaborated in the following subsection. Overall, compared to (1), our formulation offers a more flexible framework, tailored to address the diverse capabilities of heterogeneous clients.

3.2. Sparsity in the Gradients

In this subsection, we analyze the gradient structure of LoRA modules and highlight the gradients’ sparsity properties under a given sketching matrix. To begin, we present the gradient expressions for the LoRA modules in the following lemma. The proof is provided in Appendix J.2.

Lemma 3.2 (Gradient Formulation). *For a given sketching matrix $\mathbf{S}$, the gradients of $\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ with respect to $\mathbf{B}$ and $\mathbf{A}$ take the following form*

$$\begin{aligned} \nabla_{\mathbf{B}} \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi) &= \nabla \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi) \mathbf{A}^\top \mathbf{S}^\top \\ \nabla_{\mathbf{A}} \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi) &= \mathbf{S}^\top \mathbf{B}^\top \nabla \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi), \end{aligned} \quad (3)$$

where $\nabla_{\mathbf{B}} \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$, $\nabla_{\mathbf{A}} \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$, and $\nabla \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ denote the gradients of $\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ to $\mathbf{B}$, $\mathbf{A}$, and $\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}$, respectively.

In particular, a random- k diagonal sketching matrix selectively samples k rows or columns of a matrix through left product or right product, respectively. With $\mathbf{S}$ being a random- k diagonal matrix, the gradients of $\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ with respect to LoRA modules $\mathbf{B}$ and $\mathbf{A}$, as shown in (3), naturally become structurally sparse matrices. This sparsity reduces computational and memory overhead during training, enabling faster gradient computation and parameter updates, while alleviating communication overhead across distributed clients by transmitting only non-zero elements.

Remark 3.3 (Sparsity Level Control). A key advantage of our formulation is its flexible control over the sparsity level of local gradients, achieved by configuring the parameter k_i of the sketching matrix set $\mathcal{S}_i = \mathcal{S}(r, k_i)$. This mechanism allows each client to tailor its local updates according to its communication and computation resource constraints, ensuring efficient and scalable fine-tuning in heterogeneous federated systems. Lowering k_i helps resource-constrained clients reduce computation and communication overhead, while more capable clients can increase k_i to conduct more informative local updates.

Remark 3.4 (Justification for the Choice of Random- k Sketching). We adopt Random- k sketching because it is unbiased, induces structured sparsity, and exhibits strong empirical performance (see Section 5). In contrast, identifying an appropriate metric for quantifying the importance of individual LoRA rows or columns remains an open problem,

even in the centralized LoRA setting. This lack of a principled importance measure prevents the reliable development of Top- k -style or importance-based sketching strategies. We empirically evaluate several heuristic importance-based sketching variants and find that they consistently underperform Random- k sketching. Detailed theoretical intuition and empirical results can be found in Appendix B.

3.3. FSLoRA Algorithm

Based on the formulation in (2), we propose a resource-adaptive algorithm, termed FSLoRA. In each round, FSLoRA enables each client to update only a client-specific pair of submatrices from the global LoRA modules $\mathbf{B}$ and $\mathbf{A}$. The server maintains the global LoRA modules $\mathbf{B}$ and $\mathbf{A}$ and periodically updates them by aggregating sparse local updates from distributed clients. The per-round FSLoRA procedure is detailed below.

- **Sketch Sampling & Broadcast.** The server begins by sampling sketching matrices $\{\mathbf{S}_i^t \sim \mathcal{S}_i\}_{i=1}^N$ for all clients, where $\mathcal{S}_i$ represents the set of possible sketching matrices for client i . These sketches are then sent to the corresponding clients. Additionally, the server broadcasts the current global LoRA modules $[\mathbf{B}^t; \mathbf{A}^t]$ to all clients. Note that the communication load introduced by transmitting the sketching matrix is negligible compared to global LoRA modules, as it involves only *binary sketching indices*.
- **Local Sparse Training.** Clients initialize local LoRA modules from the global model and perform sketch-guided updates using $\mathbf{S}_i^t$, given by:

$$\begin{bmatrix} \mathbf{B}_i^{t,h+1} \\ \mathbf{A}_i^{t,h+1} \end{bmatrix} = \begin{bmatrix} \mathbf{B}_i^{t,h} \\ \mathbf{A}_i^{t,h} \end{bmatrix} - \gamma \begin{bmatrix} \Delta \mathbf{B}_i^{t,h} (\mathbf{S}_i^t)^\top \\ (\mathbf{S}_i^t)^\top \Delta \mathbf{A}_i^{t,h} \end{bmatrix}, \quad (4)$$

where γ denotes the learning rate and $[\Delta \mathbf{B}_i^{t,h}; \Delta \mathbf{A}_i^{t,h}]$ is a shorthand representation for:

$$\begin{bmatrix} \Delta \mathbf{B}_i^{t,h} \\ \Delta \mathbf{A}_i^{t,h} \end{bmatrix} = \begin{bmatrix} \nabla \ell(\mathbf{W}_0 + \mathbf{B}_i^{t,h} \mathbf{S}_i^t \mathbf{A}_i^{t,h}, \xi_i^{t,h}) (\mathbf{A}_i^{t,h})^\top \\ (\mathbf{B}_i^{t,h})^\top \nabla \ell(\mathbf{W}_0 + \mathbf{B}_i^{t,h} \mathbf{S}_i^t \mathbf{A}_i^{t,h}, \xi_i^{t,h}) \end{bmatrix}.$$

The update direction in (4) corresponds to the negative stochastic gradient of $\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ with respect to $[\mathbf{B}; \mathbf{A}]$ for a given sketch $\mathbf{S}_i^t$, as established in Lemma 3.2. The total update for client i in one round of training, consisting of H local steps, can be expressed as follows:

$$\begin{bmatrix} \mathbf{B}_i^{t,H} - \mathbf{B}_i^{t,0} \\ \mathbf{A}_i^{t,H} - \mathbf{A}_i^{t,0} \end{bmatrix} = \begin{bmatrix} \gamma \left(\sum_{h=0}^{H-1} \Delta \mathbf{B}_i^{t,h} \right) (\mathbf{S}_i^t)^\top \\ \gamma (\mathbf{S}_i^t)^\top \left(\sum_{h=0}^{H-1} \Delta \mathbf{A}_i^{t,h} \right) \end{bmatrix}.$$

From the above equation, we see that only the columns of $\mathbf{B}$ and the rows of $\mathbf{A}$ corresponding to the nonzero entries of $\mathbf{S}_i^t$ are updated during the t -th round at client i . In essence, $\mathbf{S}_i^t$ *selectively activates* specific columns of $\mathbf{B}$ and rows of $\mathbf{A}$ for each round. Afterward, clients

transmit these nonzero columns and rows of the sparse local updates to the server.

- **Sparse Update Reconstruction & Aggregation.** Using the sketch information, the server reconstructs the full-size sparse updates from the received columns and rows via zero padding, and aggregates them to update the global LoRA modules:

$$\begin{bmatrix} \mathbf{B}^{t+1} \\ \mathbf{A}^{t+1} \end{bmatrix} = \begin{bmatrix} \mathbf{B}^t \\ \mathbf{A}^t \end{bmatrix} + \frac{1}{N} \sum_{i=1}^N \begin{bmatrix} \mathbf{B}_i^{t,H} - \mathbf{B}_i^{t,0} \\ \mathbf{A}_i^{t,H} - \mathbf{A}_i^{t,0} \end{bmatrix}. \quad (5)$$

Algorithm 1 summarizes the procedure of FSLoRA under full participation, while the extension to partial participation is provided in Appendix E.

Intuitively, Random- k sketching ensures uniform expected exposure of all columns and rows of the global LoRA modules to each client, closely mirroring the behavior of vanilla Federated LoRA. As training progresses, every component in the global LoRA is sampled with sufficient frequency in expectation, and thus adequately optimized. The convergence of FSLoRA and the impact of Random- k sketching on convergence are analyzed in Section 4.

Remark 3.5 (Computation, memory, and communication). Compared to vanilla Federated LoRA (Zhang et al., 2024), FSLoRA introduces only lightweight additional operations, including Random- k sketch sampling, sketch index broadcasting, and zero padding. Consequently, FSLoRA incurs negligible additional overhead at the server, while maintaining comparable or lower overall complexity than existing heterogeneous LoRA baselines (Bai et al., 2024; Wang et al., 2024). More detailed server- and client-side comparisons with baselines are provided in Appendix A.

Remark 3.6 (Aggregation). Existing works on federated LoRA primarily adopt two aggregation strategies: (1) aggregating the LoRA modules as $[\mathbf{B}; \mathbf{A}]$ (e.g., vanilla Federated LoRA (Zhang et al., 2024)), and (2) aggregating the product $\mathbf{B}\mathbf{A}$ (e.g., FlexLoRA (Bai et al., 2024)). Both methods have demonstrated effectiveness, as evidenced by their promising performance in prior studies. In this work, we adopt the former, as it introduces minimal overhead and retains the simplicity of LoRA. Additionally, we prove that FSLoRA is compatible with secure aggregation in Appendix D.

3.4. Comparison with Communication Compression

Although both the sketching approach in FSLoRA and communication compression (Kuo et al., 2024) reduce communication overhead, the sketching approach fundamentally differs from traditional compression techniques. Compression methods focus solely on reducing the transmission load, leaving the gradient computation and model updates unchanged from the vanilla Federated LoRA. FSLoRA goes beyond communication savings by also reducing gradient computation and model update overhead through sparse

Algorithm 1 Federated Sketching LoRA (FSLoRA)

Require: Base model $\mathbf{W}_0$, LoRA modules $\mathbf{B}^0$ and $\mathbf{A}^0$, learning rate γ , and sketching set $\{\mathcal{S}_i\}_{i=1}^N$

- 1: **for** $t = 0, 1, \dots, T-1$ **do**
- 2: Server samples sketching matrices $\{\mathbf{S}_i^t \sim \mathcal{S}_i\}_{i=1}^N$ and sends them back to the clients
- 3: Server broadcasts the current global LoRA modules to the clients
- 4: **for** $h = 0, 1, \dots, H-1$ **do**
- 5: Clients update the local LoRA modules via (4)
- 6: **end for**
- 7: Clients upload the non-zero columns of $(\mathbf{B}_i^{t,H} - \mathbf{B}_i^{t,0})$ and the non-zero rows $(\mathbf{A}_i^{t,H} - \mathbf{A}_i^{t,0})$
- 8: Server updates the global LoRA modules via (5)
- 9: **end for**

training. Notably, these two methods are orthogonal and can be combined to achieve greater efficiency. We validate the effectiveness of this combination in Appendix G.5.

4. Analysis

Our theoretical analysis focuses on the full-participation setting in order to isolate and clearly characterize the effect of the *sketching mechanism*. As sketching is independent of client sampling, the impact of partial participation on FSLoRA is analogous to that on FedAvg. Accordingly, extending our analysis from full to partial participation follows standard techniques in FL (Yang et al., 2021).

We conduct our analysis under the following notations and assumptions.

Notations: We define $\tilde{\ell}(\mathbf{B}, \mathbf{A}, \xi; \mathbf{S}) = \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ and $\tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) = \mathbb{E}_{\xi \sim \mathcal{D}_i} [\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)]$ for a given $\mathbf{S}$ and $f_i^S(\mathbf{B}, \mathbf{A}) = \mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i} [\tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S})]$. We denote $\mathbf{X} = [\mathbf{B}; \mathbf{A}]$ and rewrite $f(\mathbf{B}, \mathbf{A})$, $f_i(\mathbf{B}, \mathbf{A})$, $f^S(\mathbf{B}, \mathbf{A})$, $f_i^S(\mathbf{B}, \mathbf{A})$, $\tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S})$, and $\ell(\mathbf{B}, \mathbf{A}, \xi; \mathbf{S})$ as $f(\mathbf{X})$, $f_i(\mathbf{X})$, $f^S(\mathbf{X})$, $f_i^S(\mathbf{X})$, $\tilde{f}_i(\mathbf{X}; \mathbf{S})$, and $\ell(\mathbf{X}, \xi; \mathbf{S})$ respectively. Additionally, we use $\|\cdot\|$ to represent the Frobenius norm.

Assumption 4.1. During fine-tuning, the iterates remain in a bounded region $\mathcal{C}$. For each client i , $f_i(\mathbf{X})$ is differentiable and L -smooth on $\mathcal{C}$, i.e., there exists a positive constant L such that for any $\mathbf{X}, \mathbf{Y} \in \mathcal{C}$,

$$\|\nabla f_i(\mathbf{X}) - \nabla f_i(\mathbf{Y})\| \leq L\|\mathbf{X} - \mathbf{Y}\|, \quad \forall i.$$

Assumption 4.2. $\nabla_{\mathbf{X}} \tilde{\ell}(\mathbf{X}, \xi; \mathbf{S})$ is an unbiased estimate of $\nabla_{\mathbf{X}} f_i^S(\mathbf{X})$ and its variance is bounded as

$$\mathbb{E} \|\nabla_{\mathbf{X}} \tilde{\ell}(\mathbf{X}, \xi; \mathbf{S}) - \nabla_{\mathbf{X}} f_i^S(\mathbf{X})\|^2 \leq \rho \|\nabla_{\mathbf{X}} f_i^S(\mathbf{X})\|^2 + \sigma^2$$

where the expectation is computed over $\xi \sim \mathcal{D}_i$ and $\mathbf{S} \sim \mathcal{S}_i$.

Assumption 4.3. The gradient dissimilarity between the

global loss $f^S(\mathbf{X})$ and each local loss $f_i^S(\mathbf{X})$ satisfies

$$\|\nabla_{\mathbf{X}} f_i^S(\mathbf{X}) - \nabla_{\mathbf{X}} f^S(\mathbf{X})\|^2 \leq c_h \|\nabla_{\mathbf{X}} f^S(\mathbf{X})\|^2 + \delta_h^2,$$

where $c_h \geq 0$ and $f^S(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N f_i^S(\mathbf{X})$.

Assumption 4.1 is standard in optimization literature (Bottou et al., 2018; Fang et al., 2024). Assumptions 4.2 and 4.3 are commonly adopted in FL to bound sampling randomness and data heterogeneity (Fang et al., 2022; Yi et al., 2022). We further discuss Assumption 4.1 and provide an empirical validation for Assumptions 4.2 and 4.3 in Appendix C, showing that these assumptions are reasonable within the LLM fine-tuning scenario. Building on these assumptions, we analyze the convergence behavior of FSLoRA and the results are summarized in the following theorem.

Theorem 4.4 (Convergence of FSLoRA). *Suppose that Assumptions 4.1-4.3 hold, then there exists a learning rate*

$$\gamma \leq \min \left\{ \frac{N}{24\rho(c_h+1)HL}, \frac{1}{8\sqrt{\tilde{L}L(\rho+1)(c_h+1)H}}, \frac{1}{H} \right\}$$

such that the iterates $\{\mathbf{X}^t\}$ generated by FSLoRA satisfy

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla_{\mathbf{X}} f^S(\mathbf{X}^t)\|^2 \leq 8 \frac{\sqrt{\tilde{L}\mathcal{F}_0\sigma_\rho^2}}{\sqrt{NTH}} + 10(\tilde{L}L)^{\frac{1}{3}} \left(\frac{\mathcal{F}_0\sigma_\rho}{T} \right)^{\frac{2}{3}} + \frac{4\mathcal{F}_0}{T}, \quad (6)$$

where $\sigma_\rho^2 = \sigma^2 + 3(\rho+1)\sigma_h^2$, $\bar{L} = \left(\frac{1}{N} \sum_{i=1}^N \frac{r}{k_i} \right) L$, $\tilde{L} = \left(\frac{1}{N} \sum_{i=1}^N \frac{r^2}{k_i^2} \right) L$ and $\mathcal{F}_0 = f^S(\mathbf{X}^0) - f^*$ with f^* denoting the lower bound of $f^S(\mathbf{X})$.

Technical highlights of Theorem 4.4: A key step in the proof of Theorem 4.4 is characterizing the impact of the sketching mechanism on the optimization landscape. Our analysis reveals how the sketching operation modifies the smoothness properties of the objective, introducing scaled smoothness constants, $\frac{r}{k_i}L$ and $\frac{r^2}{k_i^2}L$, which directly influence the convergence behavior. Further details are presented in Appendix J.3.

Discussion: Theorem 4.4 establishes an upper bound on the convergence rate of the proposed FSLoRA algorithm. The parameters $\bar{L}$ and $\tilde{L}$ provide insight into how client resource heterogeneity influences the convergence. When clients possess sufficient computational resources, they can adopt larger ranks for their local LoRA modules (i.e., larger k_i), which leads to faster convergence of FSLoRA (i.e., smaller $\bar{L}$ and $\tilde{L}$). This effect is also empirically validated in Section 5. Moreover, the derived upper bound vanishes as $T \rightarrow \infty$ for any $k_i > 0$, implying that FSLoRA converges to a stationary point under general sketching ratios.

Table 1. Testing accuracy over 3 independent runs for fine-tuning the RoBERTa model on the GLUE benchmark. FSLoRA achieves a notable improvement in average performance compared to the baselines.

Method	GPU hours	QNLI	MRPC	CoLA	MNLI	RTE	SST-2	QQP	Avg.
Base model	-	49.5	52.0	30.9	32.6	47.3	50.9	63.2	46.6
HeteroLoRA	10.7h	87.5 $\pm$ 0.5	84.4 $\pm$ 0.9	75.3 $\pm$ 1.2	66.3 $\pm$ 0.8	69.0 $\pm$ 1.7	89.5 $\pm$ 0.0	85.3 $\pm$ 0.1	79.6
FlexLoRA	12.6h	88.5 $\pm$ 0.2	81.2 $\pm$ 0.4	77.5 $\pm$ 1.2	63.0 $\pm$ 0.5	62.2 $\pm$ 1.9	92.8 $\pm$ 0.4	87.4 $\pm$ 0.1	78.9
FLoRA	12.3h	87.2 $\pm$ 0.3	78.1 $\pm$ 0.7	77.4 $\pm$ 1.7	74.6 $\pm$ 0.5	54.4 $\pm$ 2.1	93.4 $\pm$ 0.1	87.1 $\pm$ 0.3	78.9
RAVAN	12.3h	83.4 $\pm$ 1.1	76.2 $\pm$ 1.6	71.8 $\pm$ 1.4	66.3 $\pm$ 3.6	63.1 $\pm$ 2.8	92.1 $\pm$ 0.9	83.4 $\pm$ 1.7	76.6
FSLoRA	10.9h	88.0 $\pm$ 0.3	87.3 $\pm$ 0.2	82.2 $\pm$ 0.5	76.4 $\pm$ 0.2	69.8 $\pm$ 1.2	93.5 $\pm$ 0.1	85.8 $\pm$ 0.1	83.3

Table 2. Testing accuracy over 3 independent runs for fine-tuning the LLaMA-3.2-3B-Instruct on the commonsense reasoning benchmark. FSLoRA demonstrates consistent performance improvement across these tasks compared to baselines.

Method	GPU hours	ARC-c	ARC-e	BoolQ	HellaSwag	OBQA	PIQA	SIQA	WinoGrande	Avg.
Base model	-	51.5	61.4	47.2	55.8	52.4	21.3	54.4	4.6	43.6
HeteroLoRA	43.7h	73.4 $\pm$ 0.3	86.6 $\pm$ 0.2	65.8 $\pm$ 0.5	73.0 $\pm$ 0.5	71.4 $\pm$ 0.3	80.9 $\pm$ 0.7	73.8 $\pm$ 0.3	72.0 $\pm$ 0.3	74.6
FlexLoRA	68.3h	74.2 $\pm$ 0.3	86.7 $\pm$ 0.6	68.6 $\pm$ 0.8	79.4 $\pm$ 0.7	75.8 $\pm$ 0.4	81.0 $\pm$ 0.3	75.9 $\pm$ 0.4	77.9 $\pm$ 0.3	77.4
FLoRA	49.8h	68.3 $\pm$ 0.6	83.1 $\pm$ 0.5	65.8 $\pm$ 0.9	77.2 $\pm$ 0.5	74.2 $\pm$ 0.3	80.5 $\pm$ 0.6	76.1 $\pm$ 0.5	71.5 $\pm$ 0.5	74.6
RAVAN	51.2h	71.9 $\pm$ 0.9	81.6 $\pm$ 1.3	65.7 $\pm$ 3.2	76.7 $\pm$ 1.4	75.6 $\pm$ 1.5	77.1 $\pm$ 1.2	74.8 $\pm$ 0.6	73.8 $\pm$ 0.6	74.7
FSLoRA	44.3h	76.1 $\pm$ 0.4	87.2 $\pm$ 0.5	69.3 $\pm$ 0.7	82.2 $\pm$ 1.1	80.7 $\pm$ 0.6	84.0 $\pm$ 0.2	76.8 $\pm$ 0.0	79.1 $\pm$ 0.2	79.4

This highlights the robustness of FSLoRA to heterogeneous client resource constraints. In particular, as the sketching ratio $k_i/r \rightarrow 1$ (i.e., $\bar{L} = L$), the decay rate of the dominating term (i.e., the first term on the right-hand-side of (6)) recovers the convergence rate of FedAvg (Yu et al., 2019; Karimireddy et al., 2020). This demonstrates the tightness of our analysis and shows that FSLoRA retains the convergence guarantees of vanilla Federated LoRA in the limiting case.

Remark 4.5 (Analysis Extension to Importance-Based Sketching). To this end, we provide a theoretically grounded framework for heterogeneous LoRA fine-tuning. Although our algorithm and analysis are based on Random- k sketching, the framework naturally extends to importance-based sketching when a reliable importance metric is available. We investigate this extension in Appendix J.5.

5. Experiments

Our experiments focus on RoBERTa (125M) (Liu, 2019) and LLaMA-3.2-3B-Instruct models (Dubey et al., 2024), which represent typical model sizes suitable for client-side deployment, as well as the LLaMA-7B to reflect large-scale scenarios. For RoBERTa and LLaMA-3.2-3B-Instruct models, we fine-tune and evaluate them on the GLUE (Wang, 2018) and commonsense reasoning benchmarks (Hu et al., 2023), respectively. Similar to (Zhang et al., 2024; Wang et al., 2024), we adopt Dirichlet-based partitioning for dataset splits. The number of clients is set to 20 in the main manuscript and increased to 50 and 100 in Appendix E. Further details are provided in the Appendix

I. The implementation code for this project is available at <https://github.com/wenzhifang/Federated-Sketching-LoRA-Implementation>.

5.1. Main Results Under Heterogeneous LoRA Setup

Performance comparison with baselines: We consider four state-of-the-art baselines: three representative methods discussed in Section 2.2, and RAVAN (Raje et al., 2026), a recent federated LoRA fine-tuning method that addresses resource heterogeneity through partial LoRA head selection/freezing. We also report the model’s accuracy before fine-tuning (i.e., the base model). For FSLoRA, the rank of the global LoRA modules is fixed to $r = 64$, while the sketching ratio for client i is sampled from $\{0.125, 0.25, 0.5, 0.75\}$. For fair comparison, the same rank configuration is applied to all baseline methods. Tables 1 and 2 present the performance of FSLoRA and the baselines. Across both benchmarks, FSLoRA consistently outperforms the base model in a wide margin, confirming the necessity of collaborative LLM fine-tuning. On the GLUE benchmark with RoBERTa, FSLoRA achieves the highest average accuracy among all methods. On the more challenging commonsense reasoning benchmark with LLaMA, FSLoRA still delivers the best overall performance. Notably, FSLoRA maintains computational efficiency comparable to HeteroLoRA in terms of GPU hours, highlighting its advantage in heterogeneous LoRA fine-tuning.

Evaluation under broader heterogeneity, increased number of clients, and other models: In Appendix E, we extend our evaluation to 50 and 100 clients under partial client participation, incorporating greater diversity in clients’ com-





Figure 2. Comparison between FSLoRA with and without sketching (the latter equivalent to Federated LoRA) where the upload budget for clients is set to $100\times$ the size of the global LoRA modules at each rank. FSLoRA obtains a better performance, validating its communication efficiency.



(a) Comparison of FSLoRA with and without sketching, with an upload budget $80\times$ the size of the global LoRA modules.

(b) Impact of the rank of global LoRA modules on FSLoRA, given a fixed rank k_i for the updated submatrices at the clients.

(c) Impact of the sketching ratio on FSLoRA’s performance under a fixed rank $r = 64$ for the global LoRA modules.

Figure 3. Fine-tuning the LLaMA-3.2-3B-Instruct model on the commonsense reasoning benchmark. The results are averaged over eight tasks, illustrating FSLoRA’s ability to maintain strong performance while adapting to different rank and sketching configurations.

munication and computation capabilities, as well as varying levels of data heterogeneity. In Appendix F, we further assess the effectiveness of our method on Qwen2.5-1.5B-Instruct and LLaMA-7B models.

5.2. Ablation Study

Impact of sketching: In Figures 2 and 3(a), we compare the performance of FSLoRA with and without sketching on fine-tuning RoBERTa and LLaMA-3.2-3B-Instruct, respectively. Note that FSLoRA without sketching is equivalent to the vanilla Federated LoRA. For FSLoRA with sketching, we apply a uniform sketching ratio of $k_i/r = 0.5$ across all clients. The upload budget for each client is set to 100 and 80 times the size of the full global LoRA modules at the corresponding rank for the RoBERTa and the LLaMA-3.2-3B-Instruct models, respectively. As shown in Figures 2 and 3(a), both FSLoRA with and without sketching achieve higher accuracy when the rank r increases due to the availability of more tunable parameters. In addition, FSLoRA consistently outperforms its non-sketched counterpart across all the ranks and datasets. The use of sketching increases the communication frequency for clients under the same communication budget, thereby facilitating the optimization process and enhancing fine-tuning efficiency.

Impact of the global rank: In Figure 3(b), we investigate the impact of the rank of the global LoRA modules on FSLoRA’s performance. We vary the rank of the global LoRA modules while keeping the rank of submatrices updated by the clients to be consistent (i.e., $k_i = 8$). This ensures that the communication and computational resources on the client side remain unchanged. As illustrated in Figure 3(b), FSLoRA maintains stable convergence across all the configurations. Moreover, FSLoRA demonstrates improved performance as the global rank increases. This observation confirms that the proposed sketching mechanism enables resource-constrained systems to reap the benefits of a higher global rank, striking an effective balance between efficiency and performance.

Impact of sketching ratio: We investigate the impact of the sketching ratio on FSLoRA’s performance by maintaining a constant global LoRA rank $r = 64$ while varying the sketching ratio k_i/r in the range $\{0.125, 0.5, 1\}$. As shown in Figure 3(c), there is a slight performance degradation as the sketching ratio decreases, which is consistent with our theoretical analysis. This also reflects a tradeoff: while a larger sketching ratio improves convergence and accuracy, a smaller ratio reduces both computational and communication overhead. Notably, the observed degradation remains minor, highlighting FSLoRA’s ability to maintain strong

performance even under constrained resources. This demonstrates its effectiveness in balancing efficiency and accuracy, making it well-suited for resource-limited scenarios.

Further experiments: Results with more clients under broader heterogeneity, as well as with other models, are reported in Appendix E and Appendix F, respectively. Appendix G.1 provides detailed per-task comparisons on the commonsense reasoning benchmark corresponding to Figures 3(a) and 3(b). The impact of varying the number of local updates H is studied in Appendix G.2, while the empirical extension to dynamic sketching ratios is presented in Appendix G.3. Finally, Appendix G.5 demonstrates the synergistic effect of integrating compression with sketching.

6. Conclusion

We proposed FSLoRA, a novel collaborative LLM fine-tuning framework that introduces a sketching mechanism to enhance training efficiency in resource-constrained systems. By maintaining large-rank LoRA modules on the server and allowing clients to selectively update submatrices based on the sketching ratios, FSLoRA effectively adapts to heterogeneous communication and computational constraints. We provided a rigorous convergence analysis of FSLoRA that characterizes how the sketching ratios (i.e., resource heterogeneity) affect the convergence rate. Finally, we confirmed the effectiveness of FSLoRA through extensive experiments across multiple datasets and models.

Limitations

One limitation of our work is that, although we provide theory and extensive experiments, the analysis focuses on a standard federated setting and does not yet cover more realistic system effects such as asynchronous participation, unstable connections, client stragglers, and latency fluctuations. In addition, we focus on random- k sketching because it is unbiased and theoretically tractable; more adaptive sketching strategies may be interesting future work if reliable importance measures for LoRA components can be developed.

Impact Statement

This paper makes contributions to collaborative LLM fine-tuning by developing a resource-adaptive algorithm for resource-constrained distributed LLM systems. The focus of this work is on the technical advancement of LLM fine-tuning algorithms. While this research has potential societal impacts, it primarily addresses technical challenges and does not necessitate a specific discussion on societal consequences.

Acknowledgments

This work was supported in part by the National Science Foundation (NSF) under Grants CNS-2146171, CNS-2212565, and ECCS-2543754; by the Office of Naval Research (ONR) under Grant N00014-22-1-2305; by the Air Force Office of Scientific Research (AFOSR) under Grant FA9550-24-1-0083; and by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. RS-2026-25470983).

References

- Bai, J., Chen, D., Qian, B., Yao, L., and Li, Y. Federated fine-tuning of large language models under heterogeneous tasks and client resources. *arXiv preprint arXiv:2402.11505*, 2024.
- Bisk, Y., Zellers, R., Gao, J., Choi, Y., et al. Piqa: Reasoning about physical commonsense in natural language. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 7432–7439, 2020.
- Bottou, L., Curtis, F. E., and Nocedal, J. Optimization methods for large-scale machine learning. *SIAM review*, 60(2):223–311, 2018.
- Byun, Y. and Lee, J. Towards federated low-rank adaptation of language models with rank heterogeneity. *arXiv preprint arXiv:2406.17477*, 2024.
- Charalambides, N. and Mazumdar, A. Distributed hybrid sketching for l2-embeddings. *arXiv preprint arXiv:2412.20301*, 2024.
- Chen, S., Guo, Y., Ju, Y., Dalal, H., Zhu, Z., and Khisti, A. Robust federated finetuning of llms via alternating optimization of lora. *arXiv preprint arXiv:2502.01755*, 2025.
- Cho, Y. J., Liu, L., Xu, Z., Fahrezi, A., and Joshi, G. Heterogeneous LoRA for federated fine-tuning of on-device foundation models. *arXiv preprint arXiv:2401.06432*, 2024.
- Clark, C., Lee, K., Chang, M.-W., Kwiatkowski, T., Collins, M., and Toutanova, K. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*, 2019.
- Clark, P., Cowhey, I., Etzioni, O., Khot, T., Sabharwal, A., Schoenick, C., and Tafjord, O. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*, 2018.
- Demidovich, Y., Malinovsky, G., Shulgin, E., and Richtárik, P. Mast: Model-agnostic sparsified training. *arXiv preprint arXiv:2311.16086*, 2023.

- Dubey, A., Jauhri, A., Pandey, A., Kadian, A., Al-Dahle, A., Letman, A., Mathur, A., Schelten, A., Yang, A., Fan, A., et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.
- Fan, D., Messmer, B., and Jaggi, M. On-device collaborative language modeling via a mixture of generalists and specialists. *arXiv preprint arXiv:2409.13931*, 2024.
- Fang, W., Yu, Z., Jiang, Y., Shi, Y., Jones, C. N., and Zhou, Y. Communication-efficient stochastic zeroth-order optimization for federated learning. *IEEE Transactions on Signal Processing*, 70:5058–5073, 2022.
- Fang, W., Han, D.-J., Chen, E., Wang, S., and Brinton, C. G. Hierarchical federated learning with multi-timescale gradient correction. *arXiv preprint arXiv:2409.18448*, 2024.
- Gower, R. M. and Richtárik, P. Randomized iterative methods for linear systems. *SIAM Journal on Matrix Analysis and Applications*, 36(4):1660–1690, 2015.
- Guo, P., Zeng, S., Wang, Y., Fan, H., Wang, F., and Qu, L. Selective aggregation for low-rank adaptation in federated learning. *arXiv preprint arXiv:2410.01463*, 2024.
- Guo, P., Zeng, S., Wang, Y., Fan, H., Wang, F., and Qu, L. Selective aggregation for low-rank adaptation in federated learning. *arXiv preprint arXiv:2410.01463*, 2025.
- Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. LoRA: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- Hu, Z., Wang, L., Lan, Y., Xu, W., Lim, E.-P., Bing, L., Xu, X., Poria, S., and Lee, R. K.-W. LLM-adapters: An adapter family for parameter-efficient fine-tuning of large language models. *arXiv preprint arXiv:2304.01933*, 2023.
- Kalajdzievski, D. A rank stabilization scaling factor for fine-tuning with LoRA. *arXiv preprint arXiv:2312.03732*, 2023.
- Karimireddy, S. P., Kale, S., Mohri, M., Reddi, S., Stich, S., and Suresh, A. T. Scaffold: Stochastic controlled averaging for federated learning. In *International conference on machine learning*, pp. 5132–5143. PMLR, 2020.
- Koloskova, A., Loizou, N., Boreiri, S., Jaggi, M., and Stich, S. A unified theory of decentralized sgd with changing topology and local updates. In *International conference on machine learning*, pp. 5381–5393. PMLR, 2020.
- Koo, J., Jang, M., and Ok, J. Towards robust and efficient federated low-rank adaptation with heterogeneous clients. *arXiv preprint arXiv:2410.22815*, 2024.
- Kuo, K., Raje, A., Rajesh, K., and Smith, V. Federated LoRA with sparse communication. *arXiv preprint arXiv:2406.05233*, 2024.
- Lester, B., Al-Rfou, R., and Constant, N. The power of scale for parameter-efficient prompt tuning. *arXiv preprint arXiv:2104.08691*, 2021.
- Li, K., Han, S., Su, Q., Li, W., Cai, Z., and Ji, S. Unilora: One vector is all you need. *Advances in Neural Information Processing Systems*, 38:43816–43841, 2026.
- Li, X. L. and Liang, P. Prefix-tuning: Optimizing continuous prompts for generation. *arXiv preprint arXiv:2101.00190*, 2021.
- Liu, Y. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 364, 2019.
- McMahan, B., Moore, E., Ramage, D., Hampson, S., and y Arcas, B. A. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pp. 1273–1282. PMLR, 2017.
- Mihaylov, T., Clark, P., Khot, T., and Sabharwal, A. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*, 2018.
- Nicolas, J., Maouche, M., Mokhtar, S. B., and Coates, M. Dome: Improving signal-to-noise in stochastic gradient descent via sharp-direction subspace filtering. *arXiv preprint arXiv:2507.03545*, 2026.
- Raje, A., Askin, B., Jhunjhunwala, D., and Joshi, G. Ravan: Multi-head low-rank adaptation for federated fine-tuning. *Advances in Neural Information Processing Systems*, 38: 48564–48592, 2026.
- Sakaguchi, K., Bras, R. L., Bhagavatula, C., and Choi, Y. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106, 2021.
- Sap, M., Rashkin, H., Chen, D., LeBras, R., and Choi, Y. Socialliqa: Commonsense reasoning about social interactions. *arXiv preprint arXiv:1904.09728*, 2019.
- Sarlos, T. Improved approximation algorithms for large matrices via random projections. In *2006 47th annual IEEE symposium on foundations of computer science (FOCS'06)*, pp. 143–152. IEEE, 2006.
- Shrivastava, M., Isik, B., Li, Q., Koyejo, S., and Banerjee, A. Sketching for distributed deep learning: A sharper analysis. *Advances in Neural Information Processing Systems*, 37:6417–6447, 2024.

- Shuttleworth, R., Andreas, J., Torralba, A., and Sharma, P. LoRA vs full fine-tuning: An illusion of equivalence. *arXiv preprint arXiv:2410.21228*, 2024.
- Sun, Y., Li, Z., Li, Y., and Ding, B. Improving LoRA in privacy-preserving federated learning. *arXiv preprint arXiv:2403.12313*, 2024.
- Wang, A. Glue: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461*, 2018.
- Wang, J., Liu, Q., Liang, H., Joshi, G., and Poor, H. V. Tackling the objective inconsistency problem in heterogeneous federated optimization. *Advances in neural information processing systems*, 33:7611–7623, 2020.
- Wang, J., Liu, Q., Liang, H., Joshi, G., and Poor, H. V. A novel framework for the analysis and design of heterogeneous federated learning. *IEEE Transactions on Signal Processing*, 69:5234–5249, 2021.
- Wang, R., Ouyang, Y., and Xu, W. Iterative double sketching for faster least-squares optimization. In *International Conference on Machine Learning*, pp. 22935–22963. PMLR, 2022.
- Wang, Z., Shen, Z., He, Y., Sun, G., Wang, H., Lyu, L., and Li, A. FLoRA: Federated fine-tuning large language models with heterogeneous low-rank adaptations. *arXiv preprint arXiv:2409.05976*, 2024.
- Yang, H., Fang, M., and Liu, J. Achieving linear speedup with partial worker participation in non-iid federated learning. *arXiv preprint arXiv:2101.11203*, 2021.
- Yao, Y., Zhang, J., Wu, J., Huang, C., Xia, Y., Yu, T., Zhang, R., Kim, S., Rossi, R., Li, A., et al. Federated large language models: Current progress and future directions. *arXiv preprint arXiv:2409.15723*, 2024.
- Ye, R., Wang, W., Chai, J., Li, D., Li, Z., Xu, Y., Du, Y., Wang, Y., and Chen, S. Openfedllm: Training large language models on decentralized private data via federated learning. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 6137–6147, 2024.
- Yi, X., Zhang, S., Yang, T., and Johansson, K. H. Zeroth-order algorithms for stochastic distributed nonconvex optimization. *Automatica*, 142:110353, 2022.
- Yu, H., Yang, S., and Zhu, S. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pp. 5693–5700, 2019.
- Zellers, R., Holtzman, A., Bisk, Y., Farhadi, A., and Choi, Y. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*, 2019.
- Zhang, J., Vahidian, S., Kuo, M., Li, C., Zhang, R., Yu, T., Wang, G., and Chen, Y. Towards building the federatedgpt: Federated instruction tuning. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 6915–6919. IEEE, 2024.
- Zou, Y., Wang, Z., Shi, Y., and Zhou, Y. Joint rank optimization and bandwidth allocation for heterogeneous federated lora fine-tuning. *IEEE Transactions on Vehicular Technology*, 2025.

Appendix

A	Comparison of Computation, Memory, and Communication	13
B	Justification for Random-k Sketching	14
C	Justification for the Assumptions in the Convergence Analysis	15
D	Compatibility of FSLoRA with Secure Aggregation	16
E	Evaluation under Broader Heterogeneity and Increased Number of Clients with Partial Participation	16
E.1	Increasing Resource Heterogeneity and the Number of Clients	17
E.2	Further Increasing the Number of Clients	18
E.3	Varying the Level of Data Heterogeneity	18
F	Experiments on Qwen2.5-1.5B-Instruct and LLaMA-7B	18
G	Further Experiments	19
G.1	Further Details on the Ablation Study	19
G.2	Impact of Local Updates	20
G.3	Dynamic Sketching Ratios	20
G.4	Comparison with FFA-LoRA	21
G.5	Integration of Sketching and Top-k Compression	22
H	LoRA structure over training	23
I	Implementation Details	23
I.1	Details on Hyperparameters	23
I.2	Details on Datasets	24
I.2.1	GLUE Benchmark	24
I.2.2	Commonsense Reasoning Benchmark	26
J	Proof of the Theoretical Results	26
J.1	Preliminaries	26
J.2	Proof of Lemma 3.2	27
J.3	Proof of Theorem 4.4	28
J.4	Proof of Proposition J.5	31
J.5	Extension to Importance-Based Sketching	32

A. Comparison of Computation, Memory, and Communication

Computation and memory. Let P and q denote the memory cost of the full model and the global LoRA module (rank r), respectively. The computational cost is expressed with the big O notation. Forward and backward computations, as well as activation memory, are omitted as they are identical across all the considered methods. The results are summarized in Tables 3 and 4, where m and n denote the shape of the base model, k_i denotes the LoRA rank for client i , H denotes the number of iterations per round, and N is the number of clients. Additionally, the results for the vanilla Federated LoRA, denoted as FedLoRA, are reported under the case of homogeneous LoRA ranks, i.e., $k_i = r$.

Table 3. Client-side computation load and memory usage comparison.

Method	Memory	Computation (per round)
FedLoRA	$P + q$	$\mathcal{O}(Hr(m + n))$
HeteroLoRA	$P + \frac{k_i}{r}q$	$\mathcal{O}(Hk_i(m + n))$
FlexLoRA	$P + \frac{k_i}{r}q$	$\mathcal{O}(Hk_i(m + n))$
FLoRA	$P + \max\left\{\sum_{i=1}^N \frac{k_i}{r}q, P\right\}$	$\mathcal{O}\left(Hk_i(m + n) + (\sum_{i=1}^N k_i)mn + mn\right)$
FSLoRA	$P + \frac{k_i}{r}q$	$\mathcal{O}(Hk_i(m + n))$

Table 4. Server-side computation load and memory usage comparison.

Method	Memory	Computation (per round)
FedLoRA	Nq	$\mathcal{O}(N(m + n)r)$
HeteroLoRA	$\sum_{i=1}^N \frac{k_i}{r}q$	$\mathcal{O}(N(m + n)r)$
FlexLoRA	$\max\left\{\sum_{i=1}^N \frac{k_i}{r}q, 2P\right\}$	$\mathcal{O}\left((\sum_{i=1}^N k_i)mn + Nmn + \min\{m, n\}mn\right)$
FLoRA	$\sum_{i=1}^N \frac{k_i}{r}q$	$\mathcal{O}\left((\sum_{i=1}^N k_i)(m + n)\right)$
FSLoRA	$\sum_{i=1}^N \frac{k_i}{r}q$	$\mathcal{O}(N(m + n)r)$

As shown in Tables 3 and 4, FSLoRA matches HetLoRA in both computation and memory cost. FLoRA introduces additional client-side overhead due to merging LoRA modules. FlexLoRA incurs extra server-side costs from conducting SVD on the full model. In summary, FSLoRA guarantees convergence with minimum overhead.

Communication. We detailed the communication load for baselines and our methods in Table 5, where q denotes the communication cost of a global LoRA module with rank r , k_i denotes the local LoRA rank for client i , m and n denote the shape of the base model, and N denotes the number of clients.

Table 5. Communication complexity, assuming float 32 parameters and binary sketching indices.

	FedLoRA	HeteroLoRA	FlexLoRA	FLoRA	FSLoRA
Uplink	q	$\frac{k_i}{r}q$	$\frac{k_i}{r}q$	$\frac{k_i}{r}q$	$\frac{k_i}{r}q$
Downlink	q	q	q	$\sum_{i=1}^N \frac{k_i}{r}q$	$q(1 + \frac{Nr}{32mn})$

For the uplink, all four heterogeneous LoRA algorithms incur the same communication overhead for transmitting updated local LoRA modules, which is lower than that of FedLoRA. For the downlink, FLoRA requires broadcasting the stacked LoRA modules, while HeteroLoRA and FlexLoRA broadcast the updated global LoRA modules. FSLoRA, on the other hand, broadcasts both the global LoRA modules and additional sketching matrices. The extra communication introduced by the sketching matrices is negligible compared to that of the global LoRA modules, as it consists only of *binary sketching indices* (i.e., the diagonal elements of the sketching matrix). For instance, in the case of the LLaMA-3.2-3B-Instruct model under our experimental LoRA configuration, the global LoRA modules contain 66,060,288 parameters, equivalent to approximately 252 MB when using float32. With a global rank of $r = 64$, the sketching indices require only 64 bits per client, covering all LoRA layers. Even with 100 clients, the total sketching overhead is merely 0.78 KB, which accounts for only 0.0003% of the global LoRA modules.

Resource-Efficient Implementation Details. We use a sparse full-rank formulation for analysis in Section 3, but the implementation can be written equivalently in dense local form with submatrix. Consider

$$l(\mathbf{W}_0 + \mathbf{BSA}), \quad \mathbf{BSA} = \frac{r}{k} \sum_{j \in I} \mathbf{B} \mathbf{e}_j \mathbf{e}_j^\top \mathbf{A}, \quad (7)$$

where $I = \{j_1, \dots, j_k\}$ is the selected index set. Let $\mathbf{E}_I = [\mathbf{e}_{j_1}, \dots, \mathbf{e}_{j_k}]$, $\hat{\mathbf{B}} = \mathbf{B} \mathbf{E}_I$, and $\hat{\mathbf{A}} = \mathbf{E}_I^\top \mathbf{A}$, then $\hat{\mathbf{B}}$ and $\hat{\mathbf{A}}$ are exactly the selected local submatrices in dense form, and

$$\mathbf{BSA} = \frac{r}{k} \hat{\mathbf{B}} \hat{\mathbf{A}}.$$

Hence, updating the sparse full-rank parameters is equivalent to updating the dense local submatrices with loss

$$l(\mathbf{W}_0 + \frac{r}{k} \hat{\mathbf{B}} \hat{\mathbf{A}}). \quad (8)$$

The forward passes in (8) and (7) are identical, and the gradients on the selected entries are also identical because both parameterizations refer to the same variables through indexing. In other words, FSLoRA can be implemented in a dense form with submatrix, which does not require clients to maintain identical fixed-rank local LoRA adapters in the implementation.

B. Justification for Random- k Sketching

FSLoRA is built upon Random- k diagonal sketching due to two key properties:

- **Submatrix selection.** Given a sparse diagonal matrix $\mathbf{S}_i$, we have

$$\mathbf{BS}_i \mathbf{A} = \sum_{j \in \mathcal{I}_i} s_j \mathbf{b}_j \mathbf{a}_j^\top = [\mathbf{b}_j]_{j \in \mathcal{I}_i} \text{diag}\{s_j\}_{j \in \mathcal{I}_i} [\mathbf{a}_j^\top]_{j \in \mathcal{I}_i}, \quad (9)$$

where $\mathcal{I}_i$ corresponds to the index set of non-zero diagonal entries of $\mathbf{S}_i$ and $\mathbf{b}_j$ and $\mathbf{a}_j^\top$ denote the j -th column of module $\mathbf{B}$ and the j -th row of module $\mathbf{A}$, respectively. In other words, with Random- k sketching, only a subset of $\mathbf{B}$'s columns and $\mathbf{A}$'s rows are activated for client i . The *sparse diagonal structure* effectively reduces local training cost for each client.

- **Unbiasedness for convergence.** When $\mathbf{S}_i$ is a Random- k diagonal matrix with k_i nonzero diagonal entries $s_j = \frac{r}{k_i}$, we have

$$\mathbb{E}[\mathbf{BS}_i \mathbf{A}] = \mathbf{BA}. \quad (10)$$

This unbiasedness is critical for our convergence analysis.

Random- k sketching is easy to implement, which ensures uniform expected exposure of all columns and rows of the global LoRA modules to each client, closely mirroring the behavior of the standard LoRA. As training progresses, every component in the global LoRA is repeatedly sampled and thus sufficiently optimized in expectation. In particular, we theoretically and empirically validated its effectiveness in Sections 4 and 5.

In contrast, identifying a principled metric to quantify the importance of individual LoRA rows or columns remains an open problem, even in the centralized LoRA setting. This limitation hinders the reliable development of Top- k or importance-based sketching strategies. Existing approaches typically rely on performing an SVD of the composite matrix $\mathbf{BA}$ (e.g., $\mathbf{BA} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top$) and defining importance scores for the columns of $\mathbf{U}$ and the rows of $\mathbf{V}^\top$. However, this departs from our intended setting, as we operate directly in the original $\mathbf{B}$ and $\mathbf{A}$ parameter space.

In addition, we empirically compare Random- k sketching with several heuristic importance-based sketching strategies, as a rigorous importance metric is currently unavailable. Specifically, we consider two heuristic importance scores for each LoRA component $\mathbf{b}_j \mathbf{a}_j^\top$: (i) the spectral norm $\|\mathbf{b}_j\|_2 \cdot \|\mathbf{a}_j\|_2$ inspired by matrix analysis, and (ii) a simpler magnitude-based metric $\|\mathbf{b}_j\|_2 + \|\mathbf{a}_j\|_2$. Using these scores, we sample (sketch) k_i components from $\{\mathbf{b}_j \mathbf{a}_j^\top\}_{j=1}^r$ with probabilities proportional to the corresponding importance scores.

The results are reported in Table 6. We observe that Random- k sketching consistently outperforms these heuristic importance-based alternatives. A potential reason is that these importance measures do not reliably capture the true contribution of

individual components. As a result, importance-based sketching may systematically under-sample certain components that are critical for optimization but poorly reflected by such heuristics, leading to biased training. On the other hand, Random- k sketching provides uniform expected exposure of all rank-one components throughout training, which prevents under-training of any individual column or row. While more principled importance-based sketching strategies may further improve performance and remain an interesting direction for future work, our current empirical evidence favors Random- k sketching.

Table 6. Accuracy comparison of Random- k sketching and importance-based sketching under the commonsense reasoning benchmark with the LLaMA-3.2-3B-Instruct model. The experimental settings are the same as those in Table 2. Random- k sketching achieves better performance.

Importance metric	ARC-c	ARC-e	BoolQ	HSwag	OBQA	PIQA	SIQA	Wino	Avg.
Spectral Norm: $\ a\ _2\ b\ _2$	73.4	86.2	60.5	77.9	75.8	81.6	74.6	73.1	75.4
Magnitude Sum: $\ a\ _2 + \ b\ _2$	72.1	86.5	57.4	75.3	73.5	81.3	72.6	69.9	73.6
Random- k	75.8	86.7	69.7	81.4	80.4	83.9	76.2	78.8	79.1

Notably, although our algorithm and analysis are based on Random- k sketching, the framework naturally extends to importance-based sketching when a reliable importance metric is available. We theoretically discussed this extension in Appendix J.5.

C. Justification for the Assumptions in the Convergence Analysis

Smoothness. Assumption 4.1, although global L -smoothness may be difficult to guarantee for the LoRA fine-tuning objective, it is reasonable to assume local L -smoothness within the neighborhood explored during fine-tuning. Since LoRA updates are typically low-rank and initialized near zero, the optimization trajectory often remains in a local region around the pretrained model, where the loss function can be approximated as smooth.

Variance and Data Heterogeneity. In the context of LLM fine-tuning, both the magnitude of stochastic gradients and gradients in LLM fine-tuning are in a mild range since:

- Transformer-based LLMs use *LayerNorm* and *scaled softmax attention*, which stabilize activations and suppress gradient spikes.
- Fine-tuning starts from a well-pretrained model already near a local minimum, leading to smaller gradients.
- The fine-tuning dataset does not typically contain strong contradictory signals to what the model already knows, resulting in a relatively flat loss surface.

To further support this empirically, we report the statistics of the expected norm of the stochastic gradients $\mathbb{E}\|\nabla_{\mathbf{X}}\tilde{\ell}(\mathbf{X}, \xi; \mathbf{S})\|$ over approximately 4500 samples for 4 representative clients with different sketching ratios. The table below reports the minimum and maximum expected norm among 30 randomly sampled model states $\mathbf{X} = [\mathbf{B}; \mathbf{A}]$.

Table 7. Statistics of the expected norm of stochastic gradients across clients.

Client ID	Number of samples	Rank	Min	Max
1	4580	4	0.1286	0.9499
2	4216	19	0.1284	0.7069
3	4873	9	0.1237	0.5774
4	5124	32	0.1499	0.8066

As we can see from Table 7, the expected norm, i.e., $\mathbb{E}_{\xi \sim \mathcal{D}_i, \mathbf{S} \sim \mathcal{S}_i} \|\nabla_{\mathbf{X}}\tilde{\ell}(\mathbf{X}, \xi; \mathbf{S})\|$, is in a moderate range. Notably, the variance is upper-bounded by the expected squared gradient norm:

$$\mathbb{E}_{\xi \sim \mathcal{D}_i, \mathbf{S} \sim \mathcal{S}_i} \|\nabla_{\mathbf{X}}\tilde{\ell}(\mathbf{X}, \xi; \mathbf{S}) - \nabla_{\mathbf{X}}f_i^{\mathbf{S}}(\mathbf{X})\|^2 \leq \mathbb{E}_{\xi \sim \mathcal{D}_i, \mathbf{S} \sim \mathcal{S}_i} \|\nabla_{\mathbf{X}}\tilde{\ell}(\mathbf{X}, \xi; \mathbf{S})\|^2.$$

Therefore, it is generally not hard to find a σ and ρ to make Assumption 4.2 hold.

On the other hand, we have

$$\begin{aligned}
 & \|\nabla_{\mathbf{X}} f_i^S(\mathbf{X}) - \nabla_{\mathbf{X}} f^S(\mathbf{X})\|^2 \\
 & \leq 2 \|\nabla_{\mathbf{X}} f_i^S(\mathbf{X})\|^2 + 2 \|\nabla_{\mathbf{X}} f^S(\mathbf{X})\|^2 \\
 & \leq 2 \|\nabla_{\mathbf{X}} f_i^S(\mathbf{X})\|^2 + 2 \frac{1}{N} \sum_{i=1}^N \|\nabla_{\mathbf{X}} f_i^S(\mathbf{X})\|^2 \\
 & \leq 2 \mathbb{E}_{\xi \sim \mathcal{D}_i, \mathbf{S} \sim \mathcal{S}_i} \|\nabla_{\mathbf{X}} \tilde{\ell}(\mathbf{X}, \xi; \mathbf{S})\|^2 + 2 \frac{1}{N} \sum_{i=1}^N \mathbb{E}_{\xi \sim \mathcal{D}_i, \mathbf{S} \sim \mathcal{S}_i} \|\nabla_{\mathbf{X}} \tilde{\ell}(\mathbf{X}, \xi; \mathbf{S})\|^2,
 \end{aligned} \tag{11}$$

where the first inequality follows Cauchy-Schwarz inequality, while the last inequality follows Jensen's inequality. Thus, the deviation can be controlled by the expected gradient norm, which we empirically found to be moderate (see Table 7). Hence, it is reasonable to impose an upper bound on $\|\nabla_{\mathbf{X}} f_i^S(\mathbf{X}) - \nabla_{\mathbf{X}} f^S(\mathbf{X})\|^2$ as in Assumption 4.3.

D. Compatibility of FSLoRA with Secure Aggregation

The aggregation of FSLoRA is compatible with secure aggregation. Taking the aggregation of module **B** as an example, we illustrate this below.

In FSLoRA, client updates are sparse matrices with non-zero values only in columns indexed by $\mathcal{I}_i \subset [r]$, size $|\mathcal{I}_i| = k_i$. With secure aggregation, each client apply *additive masking*:

$$\tilde{\mathbf{B}}_i = \Delta \mathbf{B}_i + \mathbf{R}_i, \tag{12}$$

where mask $\mathbf{R}_i$ satisfies $\text{supp}(\mathbf{R}_i) \subseteq (u, v) : u \in [m], v \in \mathcal{I}_i$ and $\sum_{i=1}^N \mathbf{R}_i = 0$. That is, the mask has non-zero entries only in the client's active columns, and all masks together sum to zero to preserve correctness. Such masks can be constructed following the classical protocol: for each pair of clients (i, j) , define a random matrix

$$\mathbf{M}_{ij} = -\mathbf{M}_{ji} \in \mathbb{R}^{m \times r}, \quad \text{supp}(\mathbf{M}_{ij}) \subseteq (u, v) \mid u \in [m], v \in \mathcal{I}_i \cap \mathcal{I}_j, \tag{13}$$

and then construct its total mask

$$\mathbf{R}_i = \sum_{j>i} \mathbf{M}_{ij} - \sum_{j<i} \mathbf{M}_{ji}. \tag{14}$$

During uploading, client i sends the masked k_i non-zero columns of $\tilde{\mathbf{B}}_i$, and then the server adds the corresponding padding and averages them as:

$$\sum_{i=1}^N \tilde{\mathbf{B}}_i = \sum_{i=1}^N (\Delta \mathbf{B}_i + \mathbf{R}_i) = \sum_{i=1}^N \Delta \mathbf{B}_i, \tag{15}$$

which matches the aggregation of module **B** in (5). The aggregation of module **B** in FSLoRA is thus compatible with secure aggregation.

We can draw the same conclusion for module **A** under the same derivation.

E. Evaluation under Broader Heterogeneity and Increased Number of Clients with Partial Participation

To accommodate a larger number of clients, we extend FSLoRA (Algorithm 1) to support *partial client participation*. Specifically, at each round, the server samples a subset of clients, distributes the current global LoRA modules to them, and aggregates only the updates from these clients. The detailed algorithm procedure is summarized in Algorithm 2.

Our theoretical analysis focuses on the full-participation setting to isolate the effect of the sketching mechanism. Since sketching is independent of client sampling, partial participation affects FSLoRA in the same manner as FedAvg. Extending the analysis therefore follows standard techniques widely studied in the FL literature (Yang et al., 2021; Wang et al., 2020), and we empirically validate FSLoRA in this setting.

Throughout this section, we fix the partial participation size to $M = 10$, i.e., 10 clients are sampled in each round.

Algorithm 2 Federated Sketching LoRA (FSLoRA) with Partial Participation

Require: Base model $\mathbf{W}_0$, LoRA modules $\mathbf{B}^0$ and $\mathbf{A}^0$, learning rate γ , sketching sets $\{\mathcal{S}_i\}_{i=1}^N$, client sampling size M

- 1: **for** $t = 0, 1, \dots, T - 1$ **do**
- 2: Server samples a participating client set $\mathcal{C}_t \subseteq \{1, \dots, N\}$ with $|\mathcal{C}_t| = M$
- 3: Server samples sketching matrices $\{\mathbf{S}_i^t \sim \mathcal{S}_i\}_{i \in \mathcal{C}_t}$ and sends them to participating clients
- 4: Server broadcasts the current global LoRA modules $(\mathbf{B}^t, \mathbf{A}^t)$ to all $i \in \mathcal{C}_t$
- 5: **for all** clients $i \in \mathcal{C}_t$ **in parallel do**
- 6: Initialize local LoRA modules: $\mathbf{B}_i^{t,0} \leftarrow \mathbf{B}^t, \mathbf{A}_i^{t,0} \leftarrow \mathbf{A}^t$
- 7: **for** $h = 0, 1, \dots, H - 1$ **do**
- 8: Client i updates local LoRA modules: $[\mathbf{B}_i^{t,h+1}; \mathbf{A}_i^{t,h+1}] \leftarrow [\mathbf{B}_i^{t,h}; \mathbf{A}_i^{t,h}] - \gamma [\Delta \mathbf{B}_i^{t,h} (\mathbf{S}_i^t)^\top; (\mathbf{S}_i^t)^\top \Delta \mathbf{A}_i^{t,h}]$
- 9: **end for**
- 10: Client i uploads the non-zero columns of $(\mathbf{B}_i^{t,H} - \mathbf{B}_i^{t,0})$ and the non-zero rows of $(\mathbf{A}_i^{t,H} - \mathbf{A}_i^{t,0})$
- 11: **end for**
- 12: Server updates the global LoRA modules: $[\mathbf{B}^{t+1}; \mathbf{A}^{t+1}] \leftarrow [\mathbf{B}^t; \mathbf{A}^t] + \frac{1}{M} \sum_{i \in \mathcal{C}_t} [\mathbf{B}_i^{t,H} - \mathbf{B}_i^{t,0}; \mathbf{A}_i^{t,H} - \mathbf{A}_i^{t,0}]$.
- 13: **end for**

E.1. Increasing Resource Heterogeneity and the Number of Clients

We extend our experiments on LLaMA-3.2-3B-Instruct with the commonsense reasoning benchmark to 50 clients. We adopt Dirichlet-based partitioning for dataset splits. Specifically, the commonsense reasoning benchmark includes 8 tasks, and we partitioned them based on the Dirichlet distribution to construct task heterogeneity among 50 clients. The Dirichlet concentration parameter is set to $\alpha = 0.1$. We simulate client resource heterogeneity via different LoRA rank distributions (beyond the limited sketching ratio considered in Section 5). More capable clients are assigned higher ranks, reflecting varying compute capacities. We consider two different rank distributions: normal and heavy-tail distributions in the range $[4, 64]$.

Normal distribution: Ranks are sampled from a normal distribution with mean $\mu = \frac{a+b}{2}$ and standard deviation $\sigma = \frac{b-a}{6}$, where $a = 4$ and $b = 64$. This models a balanced distribution of client capabilities centered around the middle of the range.

Heavy-tail distribution: We sample ranks using an inverse log-normal distribution. Specifically, we draw $x_i \sim \text{LogNormal}(\mu, \sigma)$ with $\mu = \log(\frac{a+b}{4})$ and $\sigma = 1.0$, then set $k_i = 1/x_i$ and apply min-max normalization to scale values into the range $[a, b]$. This results in a heavy-tailed distribution where most clients receive low ranks, reflecting a scenario with many low-capability clients and a few high-capability ones.

Table 8. Accuracy comparison under different client heterogeneity settings. FSLoRA outperforms baseline methods across both normal and heavy-tail LoRA rank distributions.

Rank setup	Method	ARC-c	ARC-e	BoolQ	HellaSwag	OBQA	PIQA	SIQA	WinoGrande	Avg.
Normal	HeteroLoRA	73.38	85.82	62.17	71.23	77.40	80.14	74.72	72.53	74.67
	FlexLoRA	74.23	87.84	68.37	79.77	76.00	82.97	75.90	78.13	77.90
	FLoRA	68.17	83.75	64.93	75.67	71.40	77.20	71.24	70.09	72.81
	FSLoRA	75.77	86.95	69.67	81.53	80.60	84.06	76.20	78.85	79.20
Heavy-tail	HeteroLoRA	72.44	86.78	63.60	73.10	72.00	81.34	71.65	68.75	73.71
	FlexLoRA	73.04	86.70	62.23	75.57	78.00	81.12	74.77	73.32	75.59
	FLoRA	67.92	81.90	64.90	72.77	74.00	80.41	75.28	70.24	73.43
	FSLoRA	75.77	86.70	69.67	81.40	80.40	83.90	76.15	78.77	79.10

As shown in Table 8, FSLoRA outperforms other methods under both heterogeneity settings. As we move from normal to heavy-tail, where more clients are low-resource, overall performance decreases for all methods. However, FSLoRA exhibits the smallest drop, demonstrating stronger robustness to extreme client heterogeneity.

In Figure 4, we compare the convergence behavior of FSLoRA and three baseline methods under the aforementioned two types of client heterogeneity. Under the normal distribution, FlexLoRA exhibits fast initial progress but falls behind FSLoRA in final accuracy, likely due to approximation errors introduced by truncated SVD. This issue is exacerbated in the heavy-tail distribution, where low-rank clients dominate and SVD truncation causes greater distortion, further degrading FlexLoRA’s performance. Similarly, HeteroLoRA’s reliance on zero-padding reduces optimization efficiency, preventing it




Figure 4. Convergence behavior of FSLoRA and baselines on the commonsense reasoning benchmark with the LLaMA-3.2-3B-Instruct model. Notably, FSLoRA’s per-round communication cost is at most equal to the baselines (as detailed in Appendix A). Testing accuracy is averaged over eight tasks.

from achieving higher accuracy. FLoRA fails to show steady improvement as communication progresses. One potential reason is that frequent model merging and random reinitialization of LoRA modules in each round disrupt the convergence continuity. In contrast, FSLoRA demonstrates robust and stable convergence across both scenarios, achieving the highest overall accuracy.

E.2. Further Increasing the Number of Clients

We further evaluated the performance of FSLoRA by increasing the number of clients to 100. The results are presented in Table 9. In this setting, local ranks follow a heavy-tailed distribution as described in the previous subsection, and all other experimental configurations remain unchanged. As shown in the table, FSLoRA maintains its advantage in terms of the average performance when scaling to more clients.

Table 9. Accuracy comparison when the number of clients is $N = 100$. FSLoRA maintains its advantage in terms of the average accuracy.

Method	ARC-c	ARC-e	BoolQ	HellaSwag	OBQA	PIQA	SIQA	WinoGrande	Avg.
HeteroLoRA	71.76	86.24	62.57	68.07	76.60	79.38	74.10	69.69	73.55
FlexLoRA	73.38	87.54	69.03	75.27	78.60	80.47	74.16	73.80	76.53
FLoRA	69.97	83.25	67.10	71.67	73.60	78.94	72.21	70.80	73.44
FSLoRA	74.40	87.54	70.13	79.90	79.40	83.57	76.51	78.93	78.80

E.3. Varying the Level of Data Heterogeneity

In Table 10, we investigate the impact of the degree of data heterogeneity on performance. We increase the heterogeneity by decreasing the Dirichlet concentration parameter from $\alpha = 1$ to $\alpha = 0.1$. The local ranks follow the heavy-tail distribution described in the previous subsection, and all other experimental configurations remain consistent with Appendix E.1. As observed from Table 10, the performance of all methods degrades as heterogeneity increases. FSLoRA consistently achieves higher accuracy.

F. Experiments on Qwen2.5-1.5B-Instruct and LLaMA-7B

We further extended our evaluation to the Qwen model. Specifically, we fine-tuned Qwen2.5-1.5B-Instruct on the commonsense reasoning benchmark using the same setup as in Table 2. As shown in Table 11, FSLoRA achieves competitive performance compared with the heterogeneous LoRA baselines, suggesting its effective across different model architectures.

Table 10. Accuracy comparison under different data heterogeneity settings. FSLoRA maintains its advantage over the baselines as the data heterogeneity increases. The number of clients is set to 50.

Data setup	Method	ARC-c	ARC-e	BoolQ	HellaSwag	OBQA	PIQA	SIQA	WinoGrande	Avg.
Dir(1)	HeteroLoRA	72.18	86.11	62.57	73.10	77.60	79.82	74.26	69.46	74.39
	FlexLoRA	74.06	87.25	65.67	74.90	78.80	81.01	74.16	74.27	76.27
	FLoRA	70.14	83.29	67.27	71.60	73.60	78.73	72.16	70.96	73.47
	FSLoRA	75.85	87.50	70.93	81.47	81.00	82.86	76.66	78.53	79.35
Dir(0.1)	HeteroLoRA	72.44	86.78	63.60	73.10	72.00	81.34	71.65	68.75	73.71
	FlexLoRA	73.04	86.70	62.23	75.57	78.00	81.12	74.77	73.32	75.59
	FLoRA	67.92	81.90	64.90	72.77	74.00	80.41	75.28	70.24	73.43
	FSLoRA	75.77	86.70	69.67	81.40	80.40	83.90	76.15	78.77	79.10

Table 11. Performance comparison on the Qwen2.5-1.5B-Instruct commonsense reasoning benchmark. FSLoRA consistently outperforms the baselines, demonstrating robust effectiveness across model architectures.

Method	ARC-c	ARC-e	BoolQ	HellaSwag	OBQA	PIQA	SIQA	WinoGrande	Avg.
HeteroLoRA	70.36	86.53	65.50	68.55	79.01	74.93	68.25	68.04	72.65
FlexLoRA	73.41	88.07	63.32	71.30	76.44	78.02	70.94	68.56	73.76
FLoRA	70.28	84.51	60.90	68.82	69.58	73.45	67.53	66.27	70.17
FSLoRA	75.82	88.16	65.08	74.93	79.82	80.72	72.62	71.29	76.06

Although our primary focus is on models suitable for client-side deployment, such as RoBERTa and LLaMA-3.2-3B-Instruct models, we also include experiments on the larger LLaMA-7B model to demonstrate the scalability of FSLoRA in more complex models. Specifically, we fine-tune the LLaMA-7B model on the Wizard, Dolly-15k, and Alpaca datasets and evaluate it on 1444 MMLU samples (available at: <https://github.com/ATP-1010/FederatedLLM>). For Wizard and Dolly-15k, we adopt the same heterogeneous data partitioning as (Wang et al., 2024). Since the Alpaca dataset lacks a clear task or domain structure, we apply a uniform random partitioning strategy to distribute the data across clients. We tune the q_proj and v_proj modules and set the local LoRA ranks $k_i = [64, 32, 16, 16, 8, 8, 4, 4, 4, 4]$ for 10 clients. The parameter settings are aligned with those in (Wang et al., 2024). As shown in Table 12, FSLoRA achieves the highest average performance across all three datasets compared to baselines. These results demonstrate FSLoRA’s potential for effective fine-tuning with the large-scale LLaMA-7B model under heterogeneous client settings.

G. Further Experiments

In this section, we provide additional results, including detailed per-task comparisons from the commonsense reasoning benchmark corresponding to Figures 3(a) and 3(b). In addition, we further investigate the impact of the number of local updates H on the convergence, the robustness of FSLoRA under dynamic sketching ratio, and the integration of communication compression and sketching.

G.1. Further Details on the Ablation Study

Impact of sketching: In Figure 5, we compare the performance of FSLoRA with and without sketching on eight tasks from the commonsense reasoning benchmark using the LLaMA-3.2-3B-Instruct model. Notably, FSLoRA without sketching is equivalent to the vanilla Federated LoRA. For FSLoRA with sketching, we apply a uniform sketching ratio of $k_i/r = 0.5$ across all distributed clients. The uploading budget for each client is set to 200 times the size of the full global LoRA modules at the corresponding rank. It is clear that FSLoRA with sketching consistently outperforms its non-sketched counterpart across these eight tasks, demonstrating the effectiveness of sketching in improving performance.

Impact of the global rank: In Figure 6, we present the impact of the rank of the global LoRA modules on FSLoRA’s performance across eight tasks from the commonsense reasoning benchmark. We consider four configurations: 1) $r = 8$, $k_i/r = 1$, 2) $r = 16$, $k_i/r = 0.5$, 3) $r = 32$, $k_i/r = 0.25$, and 4) $r = 64$, $k_i/r = 0.125$. The rank of submatrices updated by the clients at each iteration remains consistent across all configurations (i.e., $k_i = 8$), ensuring that the communication and computational resources on the client side are kept fixed for all cases. In the ARC-Easy task, performance decreases as the rank increases to 64, potentially due to overfitting. In general, FSLoRA shows improved performance as the rank increases.

Table 12. Performance comparison on the LLaMA-7B model across Wizard, Dolly-15k, and Alpaca datasets. FSLoRA achieves the highest average performance, validating its scalability to larger and more complex models.

Method	Wizard	Dolly-15k	Alpaca	Avg
HeteroLoRA	27.15	26.70	28.74	27.53
FlexLoRA	28.25	35.60	30.40	31.42
FLoRA	27.91	28.50	29.54	28.65
FSLoRA	30.33	40.79	30.68	33.93




Figure 5. Comparison of FSLoRA with and without sketching, with an upload budget $200\times$ the global LoRA module size at each rank. This is based on the commonsense reasoning benchmark and the LLaMA-3.2-3B-Instruct model. We observe that the sketching mechanism improves performance across all considered tasks. The average accuracy of the eight tasks is shown in Figure 3(a).

G.2. Impact of Local Updates

Based on the commonsense reasoning benchmark and the LLaMA-3.2-3B-Instruct model, we evaluated the convergence behavior of FSLoRA under varying numbers of local updates (i.e., H). The experimental results are presented in Figure 7. In the low-to-moderate regime of local updates (i.e., $H \in \{10, 20, 100\}$), FSLoRA demonstrates a clear acceleration in convergence as H increases. For example, moving from $H = 10$ to $H = 20$ substantially reduces the number of communication rounds required to reach the same level of testing accuracy, while further increasing H to 100 yields even faster progress toward convergence. This observation indicates that a moderate increase in local updates allows clients to improve communication efficiency. However, when the number of local updates is pushed beyond this range (e.g., $H = 200$), no additional convergence gain is observed. These findings align well with our theoretical analysis in Section 4, which shows that FSLoRA can achieve a convergence speedup under certain conditions on H .

G.3. Dynamic Sketching Ratios

While our primary focus is on developing a heterogeneous federated LoRA method under a standard static setup, following prior works (Wang et al., 2024; Cho et al., 2024; Bai et al., 2024), the proposed FSLoRA algorithm can be naturally extended to dynamic, time-varying resource environments. The modification is straightforward: we allow the sparsity levels, corresponding to the sketching ratios of FSLoRA, of the matrices in the sketching set $\mathcal{S}_i$ in Algorithm 1 to become time-varying, while keeping the remaining steps unchanged.

We empirically validate the effectiveness of FLoRA under this dynamic setting. In the simulation, we group clients into three




Figure 6. Impact of the rank of global LoRA modules on FSLoRA, given a fixed rank for the updated submatrices at the clients. This is based on the commonsense reasoning benchmark and the LLaMA-3.2-3B-Instruct model. Overall, FSLoRA demonstrates improved performance as the global rank increases. The average accuracy of the eight tasks is shown in Figure 3(b).




Figure 7. Impact of the number of local updates on FSLoRA’s convergence. This is based on the commonsense reasoning benchmark and the LLaMA-3.2-3B-Instruct model. In a certain range, i.e., from 10 to 100, FSLoRA achieves a fast convergence as H increases.

capability levels *low*, *medium*, and *high*, assigned sketching ratio ranges $[0.125, 0.25]$, $[0.25, 0.5]$, and $[0.5, 1.0]$, respectively, to balance local training latencies across groups. Within each range, the sketching ratios are allowed to vary dynamically. The results, reported in Table 13, show that FSLoRA maintains comparable performance when moving from the static to the dynamic case, demonstrating its robustness under time-varying sketching ratios.

G.4. Comparison with FFA-LoRA

FFA-LoRA (Sun et al., 2024) freezes the LoRA matrix $\mathbf{A}$ and tunes only matrix $\mathbf{B}$, which reduces training cost. However, FFA-LoRA requires all clients to use the same LoRA rank as the server, and therefore does not apply to our targeted heterogeneous-rank setting, where clients may use different local ranks. Hence, we compare the proposed FSLoRA with FFA-LoRA under the homogeneous setting based on Llama-3.2-3B. Since FFA-LoRA requires the client rank to be the same as the server rank, while FSLoRA allows different local and global ranks, we carefully configured the local/global rank combinations so that FSLoRA, FFA-LoRA, and FedLoRA (i.e., vanilla federated LoRA) have the same number of locally trainable parameters for a fair comparison.

Table 13. The performance of FSLoRA under static and dynamic sketching ratios. This is based on the commonsense reasoning benchmark and the LLaMA-3.2-3B-Instruct model. FSLoRA maintains comparable performance when moving from the static to the dynamic case.

Ratios	ARC-c	ARC-e	BoolQ	HSwag	OBQA	PIQA	SIQA	Wino	Avg.
Static	76.1	87.1	70.0	81.7	81.4	82.6	76.4	78.9	79.3
Dynamic	75.5	87.7	69.2	81.3	81.2	82.2	76.0	78.8	79.0

Table 14. Comparison with FFA-LoRA under the homogeneous-rank setting on Llama-3.2-3B

Method	ARC-c	ARC-e	BoolQ	HellaSwag	OBQA	PIQA	SIQA	WinoGrande	Avg.
FFA-LoRA, $k = 8; r = 8$	68.17	84.47	64.07	64.20	69.20	78.51	67.40	63.30	69.92
FedLoRA, $k = 4; r = 4$	74.66	87.33	67.57	77.53	77.26	81.50	75.49	75.93	77.16
FSLoRA, $k = 4; r = 8$	75.00	87.75	68.90	77.70	77.60	82.32	75.69	76.16	77.64

As shown in Table 14, FSLoRA outperforms FFA-LoRA under the same local budget in the homogeneous setting. The potential reason is that freezing one LoRA factor limits the expressive power of the adaptation module. In contrast, FSLoRA randomly selects partial rows and columns for training in each round, so that all rows and columns are trained with approximately uniform exposure across clients over the course of training.

G.5. Integration of Sketching and Top-k Compression

Building on the commonsense reasoning benchmark and the LLaMA-3.2-3B-Instruct model, we further explore the integration of two orthogonal techniques, sketching and top-k compression, to further reduce the uplink communication overhead of clients in FSLoRA.




Figure 8. Testing accuracy versus communication overhead using float 32 precision. Lower sketching ratios achieve higher accuracy at the same communication cost, demonstrating that combining sketching with top-k compression leads to more communication-efficient training.

Specifically, with sketching, each client activates and updates submatrices of the global LoRA weights, $[\mathbf{b}_j]_{j \in \mathcal{I}_i}, [\mathbf{a}_j^\top]_{j \in \mathcal{I}_i}$, which are selected at the beginning of each round:

$$\mathbf{BS}_i \mathbf{A} = \sum_{j \in \mathcal{I}_i} \frac{r}{k_i} \mathbf{b}_j \mathbf{a}_j^\top = \frac{r}{k_i} [\mathbf{b}_j]_{j \in \mathcal{I}_i} [\mathbf{a}_j^\top]_{j \in \mathcal{I}_i}, \quad (16)$$

where $\mathbf{b}_j$ and $\mathbf{a}_j^\top$ denote the j -th column of module $\mathbf{B}$ and the j -th row of module $\mathbf{A}$, respectively. By limiting updates to submatrices $[\mathbf{b}_j]_{j \in \mathcal{I}_i}$ and $[\mathbf{a}_j^\top]_{j \in \mathcal{I}_i}$, FSLoRA reduces communication and computation. To further reduce communication cost, we can apply Top- k compression to the uploading stage. For instance, instead of sending the full $\Delta[\mathbf{b}_j]_{j \in \mathcal{I}_i}$, each client transmits the compressed update $\text{Top}_k(\Delta[\mathbf{b}_j]_{j \in \mathcal{I}_i})$. Sketching selects the update submatrix at the beginning of each round, while compression further reduces its transmission cost at the uploading stage. These two techniques operate at different stages and can jointly improve communication efficiency.

Table 15. Stable-rank dynamics of the server-side global LoRA update at layer 15 of Llama-3.2-3B with global rank $r = 32$.

L15, q_proj										
Rounds	20	40	60	80	100	120	140	160	180	200
HeterLoRA	1.158	1.115	1.097	1.098	1.102	1.105	1.107	1.112	1.121	1.124
FlexLoRA	1.348	1.308	1.109	1.030	1.016	1.008	1.009	1.006	1.005	1.002
FSLoRA	1.493	1.694	1.781	1.870	1.958	2.047	2.102	2.124	2.131	2.137

L15, down_proj										
Rounds	20	40	60	80	100	120	140	160	180	200
HeterLoRA	1.402	1.327	1.423	1.532	1.665	1.790	1.911	2.009	2.105	2.183
FlexLoRA	1.243	1.167	1.093	1.028	1.012	1.006	1.007	1.006	1.005	1.002
FSLoRA	1.608	1.859	2.032	2.153	2.272	2.421	2.493	2.517	2.531	2.536

In our setup, the compression ratio is fixed at 0.5 for all methods, while the sketching ratio k_i/r varies over $\{0.125, 0.25, 0.5, 1\}$. Notably, FSLoRA with sketching ratio $k_i/r = 1$ corresponds to the vanilla Federated LoRA (i.e., without sketching). Figure 8 plots testing accuracy versus communication overhead, where the x-axis represents the amount of data uploaded per client (in MB), assuming parameters are stored in float 32 precision. The results clearly show that integrating sketching with top-k compression further improves communication efficiency: methods with lower sketching ratios consistently achieve higher accuracy under the same communication budget, highlighting the potential of FSLoRA for scalable and communication-efficient collaborative LLM fine-tuning.

H. LoRA structure over training

To characterize the learned LoRA structure, we compare the stable-rank dynamics of the LoRA adapter under the proposed FSLoRA and representative baselines. Specifically, we analyze the stable rank of the server-side global LoRA update on Llama-3.2-3B under the same heavy-tail setting as in Table 8. The stable rank is defined as

$$\text{srank}(\Delta \mathbf{W}_t) = \frac{\|\Delta \mathbf{W}_t\|_F^2}{\|\Delta \mathbf{W}_t\|_2^2}, \quad \Delta \mathbf{W}_t = \mathbf{B}_t \mathbf{A}_t. \quad (17)$$

For representative illustration, we report a middle layer (L15, q_proj and down_proj) with global rank $r = 32$.

As we can see from Table 15, the stable-rank trajectories differ markedly across methods. FlexLoRA quickly collapses toward rank 1, while FSLoRA maintains a richer higher-rank adaptation structure throughout training.

I. Implementation Details

I.1. Details on Hyperparameters

Unless stated otherwise, the hyperparameters used in this work are as follows.

Table 16. The hyperparameters for RoBERTa & GLUE and LLaMA-3.2-3B-Instruct & commonsense reasoning benchmarks.

Hyperparameter	RoBERTa & GLUE	LLaMA-3.2-3B-Instruct & commonsense reasoning
Dirichlet parameter	0.1	0.1
Batch size	16	16
LoRA dropout rate	0.1	0.1
Learning rate, γ	5e-4	3e-4
Communication round, T	200	200
Local iteration number, H	50	20
Target module	["query", "value", "classification head"]	["q_proj", "k_proj", "v_proj", "up_proj", "down_proj"]

I.2. Details on Datasets

I.2.1. GLUE BENCHMARK

GLUE is a widely recognized benchmark designed to assess the natural language understanding capabilities of language models (Wang, 2018).

- **CoLA** focuses on whether a given sentence is acceptable according to linguistic rules. It evaluates a model’s ability to recognize well-formed sentences.
 - ▷ Input: A single sentence.
 - ☆ Output: A label indicating whether the sentence is acceptable or unacceptable.
- **SST-2** is designed for sentiment classification on movie reviews or short texts. It tests whether a model can correctly identify positive or negative sentiment in a given sentence.
 - ▷ Input: A single sentence.
 - ☆ Output: A label indicating positive or negative sentiment.
- **MRPC** checks if two sentences are paraphrases of each other, i.e., if they mean the same thing.
 - ▷ Input: Two sentences (‘sentence1’ and ‘sentence2’).
 - ☆ Output: A label indicating either equivalent or not equivalent.
- **QQP** tests a model’s ability to determine if two questions ask the same thing.
 - ▷ Input: Two questions.
 - ☆ Output: A label indicating duplicate or not duplicate.
- **MNLI** tests whether a given hypothesis is entailed, contradicted, or neutral with respect to a premise.
 - ▷ Input: A premise (first sentence) and a hypothesis (second sentence).
 - ☆ Output: A label indicating entailment, contradiction, or neutral.
- **QNLI** aims to determine if a context sentence correctly answers a given question.
 - ▷ Input: A question and a sentence.
 - ☆ Output: A label indicating the sentence answers the question or it does not.
- **RTE** provides pairs of sentences to see if one implies the other.
 - ▷ Input: Two sentences (‘sentence1’ and ‘sentence2’)
 - ☆ Output: A label indicating whether the meaning of one sentence is entailed from the other one.

Table 17. The prompt template of the commonsense reasoning datasets (Hu et al., 2023).

Dataset	Input Template
ARC-c/e	Please choose the correct answer to the question: [QUESTION] Answer1: [ANSWER_1] Answer2: [ANSWER_2] Answer3: [ANSWER_3] Answer4: [ANSWER_4] Answer format: answer1/answer2/answer3/answer4 the correct answer is [ANSWER]
BoolQ	Please answer the following question with true or false, question: [QUESTION] Answer format: true/false the correct answer is [ANSWER]
HellaSwag	Please choose the correct ending to complete the given sentence: [ACTIVITY_LABEL]: [CONTEXT] Ending1: [ENDING_1] Ending2: [ENDING_2] Ending3: [ENDING_3] Ending4: [ENDING_4] Answer format: ending1/ending2/ending3/ending4 the correct answer is [ANSWER]
OBQA	Please choose the correct answer to the question: [QUESTION] Answer1: [ANSWER_1] Answer2: [ANSWER_2] Answer3: [ANSWER_3] Answer4: [ANSWER_4] Answer format: answer1/answer2/answer3/answer4 the correct answer is [ANSWER]
PIQA	Please choose the correct solution to the question: [QUESTION] Solution1: [SOLUTION_1] Solution2: [SOLUTION_2] Answer format: solution1/solution2 the correct answer is [ANSWER]
SIQA	Please choose the correct answer to the question: [QUESTION] Answer1: [ANSWER_1] Answer2: [ANSWER_2] Answer3: [ANSWER_3] Answer format: answer1/answer2/answer3 the correct answer is [ANSWER]
WinoGrande	Please choose the correct answer to fill in the blank to complete the given sentence: [SENTENCE] Option1: [OPTION_1] Option2: [OPTION_2] the correct answer is [ANSWER]

I.2.2. COMMONSENSE REASONING BENCHMARK

The training set of the commonsense reasoning benchmark is a mixture of multiple datasets including about 170K training samples from ARC-c/e (Clark et al., 2018), BoolQ (Clark et al., 2019), HellaSwag (Zellers et al., 2019), OBQA (Mihaylov et al., 2018), PIQA (Bisk et al., 2020), SIQA (Sap et al., 2019), and WinoGrande (Sakaguchi et al., 2021) datasets.

- **ARC-c/e** contains the challenge and easy question set from the ARC dataset of genuine grade-school level, multiple-choice science questions.
- **BoolQ** is a question-answering dataset with yes/no questions derived from natural, real-world scenarios.
- **HellaSwag** includes questions for commonsense natural language inference, where a context and multiple endings are given, requiring the most coherent ending to be selected.
- **OBQA** involves multi-step problem-solving that combines commonsense knowledge, reasoning, and comprehension of accompanying textual information.
- **PIQA** focuses on questions requiring physical commonsense to solve. Each question offers two answer choices.
- **SIQA** targets reasoning about human actions and their social implication.
- **WinoGrande** is designed as a binary-choice fill-in-the-blank task, this dataset evaluates the ability to resolve ambiguous sentences through commonsense reasoning.

The input template, i.e., prompt format for these datasets is detailed in Table 17.

J. Proof of the Theoretical Results

J.1. Preliminaries

Before presenting the proof of the main results, we first introduce some preliminary facts that will be used later. Throughout this work, $\|\cdot\|$ denotes the Frobenius norm when applied to a matrix and the ℓ_2 norm when applied to a vector.

Lemma J.1. *Suppose a sequence of independent random matrices $\{\mathbf{P}_i\}_{i=1}^N$ satisfy $\mathbb{E}[\mathbf{P}_i] = \mathbf{0}, \forall i$. Then,*

$$\mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \mathbf{P}_i \right\|^2 = \frac{1}{N^2} \sum_{i=1}^N \mathbb{E} \|\mathbf{P}_i\|^2. \quad (18)$$

Lemma J.2. (Wang et al., 2021, Lemma 2) *Suppose a sequence of random matrices $\{\mathbf{P}_i\}_{i=1}^N$ satisfy $\mathbb{E}[\mathbf{P}_i | \mathbf{P}_{i-1}, \mathbf{P}_{i-2}, \dots, \mathbf{P}_1] = \mathbf{0}, \forall i$. Then,*

$$\mathbb{E} \left[\left\| \sum_{i=1}^N \mathbf{P}_i \right\|^2 \right] = \sum_{i=1}^N \mathbb{E} [\|\mathbf{P}_i\|^2]. \quad (19)$$

Lemma J.3. (Koloskova et al., 2020, Lemma 17) *For any $a_0 \geq 0, b \geq 0, c \geq 0, d > 0$, there exist a constant $\eta \leq \frac{1}{d}$ such that*

$$\frac{a_0}{T\eta} + b\eta + c\eta^2 \leq 2 \left(\frac{a_0 b}{T} \right)^{\frac{1}{2}} + 2c^{\frac{1}{3}} \left(\frac{a_0}{T} \right)^{\frac{2}{3}} + \frac{da_0}{T}. \quad (20)$$

Lemma J.4 (Random sketching bounds). *Let $\mathbf{S}$ be a random diagonal sketching matrix of the form*

$$\mathbf{S} = \frac{r}{k} \sum_{j \in \mathcal{I}} \mathbf{e}_j \mathbf{e}_j^\top, \quad (21)$$

where $\mathbf{e}_1, \dots, \mathbf{e}_r \in \mathbb{R}^r$ are standard unit basis vectors and $\mathcal{I} \subseteq \{1, \dots, r\}$ is chosen uniformly at random with $|\mathcal{I}| = k$. Then for any matrix $\mathbf{X}$, we have

$$\|\mathbf{X}\mathbf{S}\|^2 \leq \frac{r^2}{k^2} \|\mathbf{X}\|^2, \quad (22)$$

and in expectation we have

$$\mathbb{E}_{\mathbf{S}}[\|\mathbf{X}\mathbf{S}\|^2] \leq \frac{r}{k} \|\mathbf{X}\|^2. \quad (23)$$

Proof. Since $\mathbf{S}$ is diagonal with exactly k diagonal entries equal to $\frac{r}{k}$ and the rest zero, its largest eigenvalue is $\frac{r}{k}$. Squaring gives

$$\mathbf{S}\mathbf{S}^\top = \mathbf{S}^2 \preceq \frac{r^2}{k^2} \mathbf{I}, \quad (24)$$

Equivalently,

$$\mathbf{x}^\top (\mathbf{S}\mathbf{S}^\top) \mathbf{x} \leq \frac{r^2}{k^2} \|\mathbf{x}\|^2, \forall \mathbf{x}. \quad (25)$$

Setting $\mathbf{x} = \mathbf{x}_j$ to be the j -th column of $\mathbf{X}$ and summing over j implies

$$\|\mathbf{X}\mathbf{S}\|^2 = \sum_j \|\mathbf{S}^\top \mathbf{x}_j\|^2 = \sum_j \mathbf{x}_j^\top (\mathbf{S}\mathbf{S}^\top) \mathbf{x}_j \leq \frac{r^2}{k^2} \sum_j \|\mathbf{x}_j\|^2 = \frac{r^2}{k^2} \|\mathbf{X}\|^2, \quad (26)$$

which proves (22).

For the expected bound (23), note that each diagonal index $j \in \{1, \dots, r\}$ is included in $\mathcal{I}$ with probability $\frac{k}{r}$. Hence the expectation of $\mathbf{S}^2$ satisfies

$$\mathbb{E}_{\mathbf{S}}[\mathbf{S}^2] = \frac{r^2}{k^2} \mathbb{E}\left[\sum_{j \in \mathcal{I}} \mathbf{e}_j \mathbf{e}_j^\top\right] = \frac{r^2}{k^2} \frac{k}{r} \mathbf{I} = \frac{r}{k} \mathbf{I}. \quad (27)$$

Thus for any vector $\mathbf{x}$,

$$\mathbb{E}_{\mathbf{S}}[\|\mathbf{S}^\top \mathbf{x}\|^2] = \mathbb{E}_{\mathbf{S}}[\mathbf{x}^\top \mathbf{S}\mathbf{S}^\top \mathbf{x}] = \mathbf{x}^\top (\mathbb{E}[\mathbf{S}^2]) \mathbf{x} = \frac{r}{k} \|\mathbf{x}\|^2. \quad (28)$$

Summing over columns of $\mathbf{X}$ again establishes

$$\mathbb{E}_{\mathbf{S}}[\|\mathbf{X}\mathbf{S}\|^2] = \sum_j \mathbb{E}_{\mathbf{S}}[\|\mathbf{S}^\top \mathbf{x}_j\|^2] = \sum_j \mathbf{x}_j^\top (\mathbb{E}[\mathbf{S}^2]) \mathbf{x}_j = \frac{r}{k} \|\mathbf{X}\|^2. \quad (29)$$

This completes the proof of Lemma J.4. $\square$

J.2. Proof of Lemma 3.2

From the chain rule for matrix calculus, we know that:

$$\nabla_{\mathbf{Y}} g(\mathbf{X}\mathbf{Y}) = \mathbf{X}^\top \nabla g(\mathbf{X}\mathbf{Y}), \quad \nabla_{\mathbf{X}} g(\mathbf{X}\mathbf{Y}) = \nabla g(\mathbf{X}\mathbf{Y}) \mathbf{Y}^\top, \quad (30)$$

where $\nabla g(\mathbf{X}\mathbf{Y})$ denotes the gradient of g to $\mathbf{X}\mathbf{Y}$. Applying this to $\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$, we proceed as follows: To compute the gradient with respect to $\mathbf{B}$, set $\mathbf{X} = \mathbf{B}$ and $\mathbf{Y} = \mathbf{S}\mathbf{A}$:

$$\nabla_{\mathbf{B}} \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi) = \nabla \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi) (\mathbf{S}\mathbf{A})^\top. \quad (31)$$

Similarly, to compute the gradient with respect to $\mathbf{A}$, set $\mathbf{X} = \mathbf{B}\mathbf{S}$ and $\mathbf{Y} = \mathbf{A}$:

$$\nabla_{\mathbf{A}} \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi) = \mathbf{S}^\top \mathbf{B}^\top \nabla \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi). \quad (32)$$

J.3. Proof of Theorem 4.4

The proof of Theorem 4.4 relies on the following proposition.

Proposition J.5. *Under Assumption 4.1, $\tilde{f}_i(\mathbf{X}; \mathbf{S}) = f_i(\mathbf{B}\mathbf{S}, \mathbf{A})$, $\mathbf{S} \in \mathcal{S}_i$, $f_i^{\mathcal{S}}(\mathbf{X}) = \mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i}[\tilde{f}_i(\mathbf{X}; \mathbf{S})]$, and $f^{\mathcal{S}}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N f_i^{\mathcal{S}}(\mathbf{X})$ are smooth with parameters $L \frac{r^2}{k_i^2}$, $L \frac{r}{k_i}$, and $\left(\frac{1}{N} \sum_{i=1}^N \frac{r}{k_i}\right) L$, respectively.*

The proof of Proposition J.5 is deferred to Appendix J.4. With this proposition, we are ready to prove Theorem 4.4.

In FSLoRA, the update direction in (4) corresponds to the negative stochastic gradient of $\ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$ with respect to $[\mathbf{B}; \mathbf{A}]$ for a given sketch $\mathbf{S}_i^t$. We have defined $\tilde{\ell}(\mathbf{X}, \xi; \mathbf{S}) = \ell(\mathbf{W}_0 + \mathbf{B}\mathbf{S}\mathbf{A}, \xi)$. The iterative equation for the proposed FSLoRA algorithm thus can be written as

$$\mathbf{X}^{t+1} = \mathbf{X}^t - \gamma \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} \tilde{\ell}(\mathbf{X}_i^{t,h}, \xi_i^{t,h}; \mathbf{S}_i^t), \quad (33)$$

where $\mathbf{g}_i^{t,h}$ denotes the stochastic gradient $\nabla_{\mathbf{X}} \tilde{\ell}(\mathbf{X}_i^{t,h}, \xi_i^{t,h}; \mathbf{S}_i^t)$. Based on the smoothness of $f^{\mathcal{S}}(\mathbf{X})$, i.e., Proposition J.5, we have

$$\mathbb{E}[f^{\mathcal{S}}(\mathbf{X}^{t+1})] \leq \mathbb{E}[f^{\mathcal{S}}(\mathbf{X}^t)] - \underbrace{\mathbb{E} \left\langle \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t), \gamma \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \mathbf{g}_i^{t,h} \right\rangle}_{T_1} + \underbrace{\frac{\gamma^2 \bar{L}}{2} \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \mathbf{g}_i^{t,h} \right\|^2}_{T_2}, \quad (34)$$

where $\bar{L} = \left(\frac{1}{N} \sum_{i=1}^N \frac{r}{k_i}\right) L$.

For T_1 , we have

$$\begin{aligned} T_1 &= -H \mathbb{E} \left\langle \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t), \gamma \frac{1}{NH} \sum_{i=1}^N \sum_{h=0}^{H-1} \mathbf{g}_i^{t,h} \right\rangle \\ &= -H \mathbb{E} \left\langle \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t), \gamma \frac{1}{NH} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\rangle \\ &= -\frac{\gamma H}{2} \mathbb{E} \|\nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t)\|^2 - \frac{\gamma H}{2} \mathbb{E} \left\| \frac{1}{NH} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 \\ &\quad + \frac{\gamma H}{2} \mathbb{E} \left\| \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) - \frac{1}{NH} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 \\ &\leq -\frac{\gamma H}{2} \mathbb{E} \|\nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t)\|^2 - \frac{\gamma H}{2} \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 \\ &\quad + \frac{\gamma}{2} \sum_{h=0}^{H-1} \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}^t) - \frac{1}{N} \sum_{i=1}^N \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 \\ &\leq -\frac{\gamma H}{2} \mathbb{E} \|\nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t)\|^2 - \frac{\gamma}{2H} \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 \\ &\quad + \frac{\gamma H L^2}{2} \frac{1}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} \mathbb{E} \|\mathbf{X}_i^{t,h} - \mathbf{X}^t\|^2, \end{aligned} \quad (35)$$

where the last inequalities follow Jensen's inequality and Proposition J.5.

For T_2 , we have

$$\begin{aligned} T_2 &= \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} (\mathbf{g}_i^{t,h} \mp \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h})) \right\|^2 \\ &\leq \frac{2}{N^2} \sum_{i=1}^N \mathbb{E} \left\| \sum_{h=0}^{H-1} (\mathbf{g}_i^{t,h} - \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h})) \right\|^2 + 2 \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2, \end{aligned} \quad (36)$$

where the inequality follows the fact that $\mathbb{E}[\sum_{h=0}^{H-1} (\mathbf{g}_i^{t,h} - \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}))] = 0$ and the independence between clients.

Furthermore, we bound the first term on the right-hand side of the above inequality as

$$\mathbb{E} \left\| \sum_{h=0}^{H-1} (\mathbf{g}_i^{t,h} - \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h})) \right\|^2 = \sum_{h=0}^{H-1} \mathbb{E} \left\| \mathbf{g}_i^{t,h} - \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 \leq H\sigma^2 + \rho \sum_{h=0}^{H-1} \mathbb{E} \left\| \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2,$$

where the equality follows Lemma J.2 and the inequality follows Assumption 4.2. For $\left\| \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2$, utilizing Assumption 4.3 and Proposition J.5, we have

$$\begin{aligned} \left\| \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 &= \left\| \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \mp \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}^t) \mp \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 \\ &\leq 3 \left\| \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) - \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 + 3 \left\| \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}^t) - \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 + 3 \left\| \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 \\ &\leq 3 \frac{r^2}{k_i^2} L^2 \left\| \mathbf{X}_i^{t,h} - \mathbf{X}^t \right\|^2 + 3(c_h + 1) \left\| \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 + 3\rho\delta_h^2. \end{aligned} \quad (37)$$

Combining the above three inequalities gives rise to

$$\begin{aligned} T_2 &\leq \frac{2H}{N} (\sigma^2 + 3\rho\delta_h^2) + 2 \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 + \frac{6\rho(c_h + 1)H}{N} \mathbb{E} \left\| \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 \\ &\quad + \frac{6\rho HL^2}{N} T_3, \end{aligned} \quad (38)$$

where $T_3 = \frac{1}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} \mathbb{E} \left\| \mathbf{X}_i^{t,h} - \mathbf{X}^t \right\|^2$. Combining (34), (35), and (38) yields

$$\begin{aligned} \mathbb{E}[f^{\mathcal{S}}(\mathbf{X}^{t+1})] &\leq \mathbb{E}[f^{\mathcal{S}}(\mathbf{X}^t)] - \left(\frac{\gamma H}{2} - 3\gamma^2 \rho(c_h + 1) \frac{H}{N} \bar{L} \right) \mathbb{E} \left\| \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 + \gamma^2 \bar{L} \frac{H}{N} (\sigma^2 + 3\rho\sigma_h^2) \\ &\quad - \left(\frac{\gamma}{2H} - \gamma^2 \bar{L} \right) \mathbb{E} \left\| \frac{1}{N} \sum_{i=1}^N \sum_{h=0}^{H-1} \nabla_{\mathbf{X}} f_i^{\mathcal{S}}(\mathbf{X}_i^{t,h}) \right\|^2 + \left(\frac{\gamma HL^2}{2} + 3\gamma^2 \rho \bar{L} L^2 \frac{H}{N} \right) T_3, \end{aligned} \quad (39)$$

where $\bar{L} = \left(\frac{1}{N} \sum_{i=1}^N \frac{r}{k_i} \right) L$. Let $\gamma \leq \min\{\frac{N}{24\rho(c_h+1)HL}, \frac{1}{2H\bar{L}}, \frac{N}{6\rho\bar{L}}\}$, we have

$$\mathbb{E}[f^{\mathcal{S}}(\mathbf{X}^{t+1})] \leq \mathbb{E}[f^{\mathcal{S}}(\mathbf{X}^t)] - \frac{3\gamma H}{8} \mathbb{E} \left\| \nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t) \right\|^2 + \gamma^2 \bar{L} \frac{H}{N} (\sigma^2 + 3\rho\sigma_h^2) + \frac{5\gamma}{8} HL^2 T_3. \quad (40)$$

For T_3 , we have

$$\begin{aligned}
 T_3 &= \frac{1}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} \mathbb{E} \left\| \gamma \sum_{\tau=0}^{h-1} \mathbf{g}_i^{t,\tau} \right\|^2 \\
 &= \gamma^2 \frac{1}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} \mathbb{E} \left\| \sum_{\tau=0}^{h-1} (\mathbf{g}_i^{t,\tau} \mp \nabla_{\mathbf{x}} f_i^S(\mathbf{X}_i^{t,\tau})) \right\|^2 \\
 &\leq 2\gamma^2 \frac{1}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} \sum_{\tau=0}^{h-1} \mathbb{E} \|\mathbf{g}_i^{t,\tau} - \nabla_{\mathbf{x}} f_i^S(\mathbf{X}_i^{t,\tau})\|^2 + 2\gamma^2 \frac{1}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} h \sum_{\tau=0}^{h-1} \mathbb{E} \|\nabla_{\mathbf{x}} f_i^S(\mathbf{X}_i^{t,\tau})\|^2 \\
 &\leq 2\gamma^2 H \sigma^2 \left(\frac{1}{N} \sum_{i=1}^N \frac{r^2}{k_i^2} \right) + \frac{2\rho\gamma^2}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} \sum_{\tau=0}^{h-1} \mathbb{E} \|\nabla_{\mathbf{x}} f_i^S(\mathbf{X}_i^{t,\tau})\|^2 \\
 &\quad + \frac{2\gamma^2}{NH} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} h \sum_{\tau=0}^{h-1} \mathbb{E} \|\nabla_{\mathbf{x}} f_i^S(\mathbf{X}_i^{t,\tau})\|^2 \\
 &\leq 2\gamma^2 H \sigma^2 \left(\frac{1}{N} \sum_{i=1}^N \frac{r^2}{k_i^2} \right) + \frac{2(\rho+1)\gamma^2 H}{N} \sum_{i=1}^N \frac{r^2}{k_i^2} \sum_{h=0}^{H-1} \mathbb{E} \|\nabla_{\mathbf{x}} f_i^S(\mathbf{X}_i^{t,\tau})\|^2. \tag{41}
 \end{aligned}$$

Plugging inequality (37) into inequality (41) yeilds

$$\begin{aligned}
 T_3 &\leq 2\gamma^2 H \left(\frac{1}{N} \sum_{i=1}^N \frac{r^2}{k_i^2} \right) \sigma^2 + 6(\rho+1)\gamma^2 H^2 \left(\frac{1}{N} \sum_{i=1}^N \frac{r^2}{k_i^2} \right) \sigma_h^2 \\
 &\quad + 6(\rho+1)\gamma^2 L^2 H^2 T_3 + 6(\rho+1) \left(\frac{1}{N} \sum_{i=1}^N \frac{r^2}{k_i^2} \right) (c_h+1)\gamma^2 H^2 \mathbb{E} \|\nabla_{\mathbf{x}} f^S(\mathbf{X}^t)\|^2. \tag{42}
 \end{aligned}$$

Denote $\kappa = \frac{1}{N} \sum_{i=1}^N \frac{r^2}{k_i^2}$, we simplify the above inequality as

$$(1 - 6(\rho+1)\gamma^2 L^2 H^2) T_3 \leq 2\kappa\gamma^2 H^2 (\sigma_g^2 + 3(\rho+1)\sigma_h^2) + 6\kappa(\rho+1)(c_h+1)\gamma^2 H^2 \mathbb{E} \|\nabla_{\mathbf{x}} f^S(\mathbf{X}^t)\|^2.$$

Let $\gamma \leq \frac{1}{\sqrt{12(\rho+1)HL}}$, we get the bound for T_3

$$T_3 \leq 4\kappa\gamma^2 H^2 (\sigma^2 + 3(\rho+1)\sigma_h^2) + 12\kappa(\rho+1)(c_h+1)\gamma^2 H^2 \mathbb{E} \|\nabla_{\mathbf{x}} f^S(\mathbf{X}^t)\|^2. \tag{43}$$

Plugging the bound for T_3 into inequality (40) gives rise to

$$\begin{aligned}
 \mathbb{E}[f^S(\mathbf{X}^{t+1})] &\leq \mathbb{E}[f^S(\mathbf{X}^t)] - \left(\frac{3\gamma H}{8} - \frac{5\gamma H}{8} L^2 (12\kappa(\rho+1)(c_h+1)\gamma^2 H^2) \right) \mathbb{E} \|\nabla_{\mathbf{x}} f^S(\mathbf{X}^t)\|^2 \\
 &\quad + \gamma^2 \bar{L} \frac{H}{N} (\sigma^2 + 3\rho\sigma_h^2) + \frac{5\gamma}{8} H L^2 \cdot 4\kappa\gamma^2 H^2 (\sigma^2 + 3(\rho+1)\sigma_h^2). \tag{44}
 \end{aligned}$$

Let $\gamma \leq \frac{1}{8\sqrt{\kappa(\rho+1)(c_h+1)HL}}$, we obtain

$$\mathbb{E}[f^S(\mathbf{X}^{t+1})] \leq \mathbb{E}[f^S(\mathbf{X}^t)] - \frac{\gamma H}{4} \mathbb{E} \|\nabla_{\mathbf{x}} f^S(\mathbf{X}^t)\|^2 + \gamma^2 \bar{L} \frac{H}{N} \sigma_\rho^2 + \frac{5}{2} \kappa \gamma^3 H^3 L^2 \sigma_\rho^2, \tag{45}$$

where $\sigma_\rho^2 = \sigma^2 + 3(\rho+1)\sigma_h^2$.

Telescoping the above inequality from $t = 0$ to $T - 1$, we have

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla_{\mathbf{x}} f^S(\mathbf{X}^t)\|^2 \leq 4 \frac{f^S(\mathbf{X}^0) - f^*}{\gamma TH} + \gamma \frac{4\bar{L}}{N} \sigma_\rho^2 + 10\gamma^2 H^2 \bar{L} L \sigma_\rho^2, \tag{46}$$

where f^* denotes the lower bound of $f^S(\mathbf{X})$ and $\tilde{L} = \kappa L$.

Applying Lemma J.3 to the above inequality and letting $d = H$, it follows that there exists a learning rate

$$\gamma \leq \min\left\{\frac{N}{24\rho(c_h + 1)H\tilde{L}}, \frac{1}{8\sqrt{\tilde{L}L(\rho + 1)(c_h + 1)H}}, \frac{1}{H}\right\}$$

such that

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla_{\mathbf{X}} f^S(\mathbf{X}^t)\|^2 \leq 8 \frac{\sqrt{\tilde{L}\mathcal{F}_0\sigma_\rho^2}}{\sqrt{NTH}} + 10(\tilde{L}L)^{\frac{1}{3}} \left(\frac{\mathcal{F}_0\sigma_\rho}{T}\right)^{\frac{2}{3}} + \frac{4\mathcal{F}_0}{T}. \quad (47)$$

This completes the proof of Theorem 4.4.

J.4. Proof of Proposition J.5

i) For illustration, we need to recover $\mathbf{X}$ to $[\mathbf{B}; \mathbf{A}]$ in this proof. According to the definition of $\tilde{f}_i(\mathbf{X}; \mathbf{S})$ and $f_i(\mathbf{B}, \mathbf{A})$, we have

$$\tilde{f}_i(\mathbf{X}; \mathbf{S}) = \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) \quad (48)$$

$$\begin{aligned} &= \mathbb{E}_{\xi \sim \mathcal{D}_i} [\ell(\mathbf{W}_0 + \mathbf{BSA}, \xi)] \\ &= f_i(\mathbf{BS}, \mathbf{A}). \end{aligned} \quad (49)$$

As $f_i(\mathbf{B}, \mathbf{A})$ is L -smooth, we have

$$f_i(\mathbf{BS} + \Delta\mathbf{BS}, \mathbf{A} + \Delta\mathbf{A}) \leq f_i(\mathbf{BS}, \mathbf{A}) + \left\langle \begin{bmatrix} \nabla_{\mathbf{BS}} f_i(\mathbf{BS}, \mathbf{A}) \\ \nabla_{\mathbf{A}} f_i(\mathbf{BS}, \mathbf{A}) \end{bmatrix}, \begin{bmatrix} \Delta\mathbf{BS} \\ \Delta\mathbf{A} \end{bmatrix} \right\rangle + \frac{L}{2} \left\| \begin{bmatrix} \Delta\mathbf{BS} \\ \Delta\mathbf{A} \end{bmatrix} \right\|^2. \quad (50)$$

According to (48) and (49), we have $\tilde{f}_i(\mathbf{B} + \Delta\mathbf{B}, \mathbf{A} + \Delta\mathbf{A}; \mathbf{S}) = f_i(\mathbf{BS} + \Delta\mathbf{BS}, \mathbf{A} + \Delta\mathbf{A})$ and $\tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) = f_i(\mathbf{BS}, \mathbf{A})$. Combining these with (50) gives rise to

$$\tilde{f}_i(\mathbf{B} + \Delta\mathbf{B}, \mathbf{A} + \Delta\mathbf{A}; \mathbf{S}) \leq \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) + \left\langle \begin{bmatrix} \nabla_{\mathbf{BS}} f_i(\mathbf{BS}, \mathbf{A}) \\ \nabla_{\mathbf{A}} f_i(\mathbf{BS}, \mathbf{A}) \end{bmatrix}, \begin{bmatrix} \Delta\mathbf{BS} \\ \Delta\mathbf{A} \end{bmatrix} \right\rangle + \frac{L}{2} \left\| \begin{bmatrix} \Delta\mathbf{BS} \\ \Delta\mathbf{A} \end{bmatrix} \right\|^2. \quad (51)$$

We denote

$$L(\mathbf{W}_0 + \mathbf{BSA}) = \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) = \mathbb{E}_{\xi \sim \mathcal{D}_i} [\ell(\mathbf{W}_0 + \mathbf{BSA}, \xi)]. \quad (52)$$

Note that $\nabla_{\mathbf{BS}} f_i(\mathbf{BS}, \mathbf{A}) = \nabla L(\mathbf{W}_0 + \mathbf{BSA}) \mathbf{A}^\top$ and $\nabla_{\mathbf{A}} f_i(\mathbf{BS}, \mathbf{A}) = \mathbf{S}^\top \mathbf{B}^\top \nabla L(\mathbf{W}_0 + \mathbf{BSA})$. We thus have

$$\begin{aligned} \left\langle \begin{bmatrix} \nabla_{\mathbf{BS}} f_i(\mathbf{BS}, \mathbf{A}) \\ \nabla_{\mathbf{A}} f_i(\mathbf{BS}, \mathbf{A}) \end{bmatrix}, \begin{bmatrix} \Delta\mathbf{BS} \\ \Delta\mathbf{A} \end{bmatrix} \right\rangle &= \left\langle \begin{bmatrix} \nabla L(\mathbf{W}_0 + \mathbf{BSA}) \mathbf{A}^\top \\ \mathbf{S}^\top \mathbf{B}^\top \nabla L(\mathbf{W}_0 + \mathbf{BSA}) \end{bmatrix}, \begin{bmatrix} \Delta\mathbf{BS} \\ \Delta\mathbf{A} \end{bmatrix} \right\rangle \\ &= \left\langle \begin{bmatrix} \nabla L(\mathbf{W}_0 + \mathbf{BSA}) \mathbf{A}^\top \mathbf{S}^\top \\ \mathbf{S}^\top \mathbf{B}^\top \nabla L(\mathbf{W}_0 + \mathbf{BSA}) \end{bmatrix}, \begin{bmatrix} \Delta\mathbf{B} \\ \Delta\mathbf{A} \end{bmatrix} \right\rangle \\ &= \left\langle \begin{bmatrix} \nabla_{\mathbf{B}} \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) \\ \nabla_{\mathbf{A}} \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) \end{bmatrix}, \begin{bmatrix} \Delta\mathbf{B} \\ \Delta\mathbf{A} \end{bmatrix} \right\rangle, \end{aligned} \quad (53)$$

where the last equality follows the fact that $\tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) = L(\mathbf{W}_0 + \mathbf{BSA})$ defined in (52) and

$$\begin{bmatrix} \nabla_{\mathbf{B}} \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) \\ \nabla_{\mathbf{A}} \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) \end{bmatrix} = \begin{bmatrix} \nabla L(\mathbf{W}_0 + \mathbf{BSA}) \mathbf{A}^\top \mathbf{S}^\top \\ \mathbf{S}^\top \mathbf{B}^\top \nabla L(\mathbf{W}_0 + \mathbf{BSA}) \end{bmatrix}. \quad (54)$$

Plugging (53) into (51) gives rise to

$$\tilde{f}_i(\mathbf{B} + \Delta\mathbf{B}, \mathbf{A} + \Delta\mathbf{A}; \mathbf{S}) \leq \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) + \left\langle \begin{bmatrix} \nabla_{\mathbf{B}} \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) \\ \nabla_{\mathbf{A}} \tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S}) \end{bmatrix}, \begin{bmatrix} \Delta\mathbf{B} \\ \Delta\mathbf{A} \end{bmatrix} \right\rangle + \frac{L}{2} \left\| \begin{bmatrix} \Delta\mathbf{BS} \\ \Delta\mathbf{A} \end{bmatrix} \right\|^2. \quad (55)$$

In particular, $\left\| \begin{bmatrix} \Delta \mathbf{B} \mathbf{S} \\ \Delta \mathbf{A} \end{bmatrix} \right\|^2 = \|\Delta \mathbf{B} \mathbf{S}\|^2 + \|\Delta \mathbf{A}\|^2$. From (22), we know $\|\Delta \mathbf{B} \mathbf{S}\|^2 \leq \frac{r^2}{k_i^2} \|\Delta \mathbf{B}\|^2$. Therefore, we have $\left\| \begin{bmatrix} \Delta \mathbf{B} \mathbf{S} \\ \Delta \mathbf{A} \end{bmatrix} \right\|^2 = \frac{r^2}{k_i^2} \left\| \begin{bmatrix} \Delta \mathbf{B} \\ \Delta \mathbf{A} \end{bmatrix} \right\|^2$. As a result, $\tilde{f}_i(\mathbf{B}, \mathbf{A}; \mathbf{S})$ (i.e., $\tilde{f}_i(\mathbf{X}, \mathbf{S})$) is $L \frac{r^2}{k_i^2}$ -smooth.

ii) Note that $f_i^{\mathcal{S}}(\mathbf{X}) = \mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i}[\tilde{f}_i(\mathbf{X}, \mathbf{S})]$. Therefore, we further take expectation for (55) over $\mathbf{S} \sim \mathcal{S}_i$, leading to

$$f_i^{\mathcal{S}}(\mathbf{B} + \Delta \mathbf{B}, \mathbf{A} + \Delta \mathbf{A}) \leq f_i^{\mathcal{S}}(\mathbf{B}, \mathbf{A}) + \left\langle \begin{bmatrix} \nabla_{\mathbf{B}} f_i^{\mathcal{S}}(\mathbf{B}, \mathbf{A}) \\ \nabla_{\mathbf{A}} f_i^{\mathcal{S}}(\mathbf{B}, \mathbf{A}) \end{bmatrix}, \begin{bmatrix} \Delta \mathbf{B} \\ \Delta \mathbf{A} \end{bmatrix} \right\rangle + \frac{L}{2} \mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i} \left\| \begin{bmatrix} \Delta \mathbf{B} \mathbf{S} \\ \Delta \mathbf{A} \end{bmatrix} \right\|^2. \quad (56)$$

In particular, $\mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i} \left\| \begin{bmatrix} \Delta \mathbf{B} \mathbf{S} \\ \Delta \mathbf{A} \end{bmatrix} \right\|^2 = \mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i} \|\Delta \mathbf{B} \mathbf{S}\|^2 + \|\Delta \mathbf{A}\|^2$. From (23), we know $\mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i} \|\Delta \mathbf{B} \mathbf{S}\|^2 \leq \frac{r}{k_i} \|\Delta \mathbf{B}\|^2$. In other words, $\mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i} \left\| \begin{bmatrix} \Delta \mathbf{B} \mathbf{S} \\ \Delta \mathbf{A} \end{bmatrix} \right\|^2 \leq \frac{r}{k_i} \left\| \begin{bmatrix} \Delta \mathbf{B} \\ \Delta \mathbf{A} \end{bmatrix} \right\|^2$. We thus claim that $f_i^{\mathcal{S}}(\mathbf{B}, \mathbf{A})$ (i.e., $f_i^{\mathcal{S}}(\mathbf{X})$) is $L \frac{r}{k_i}$ -smooth.

iii) Finally, for $f^{\mathcal{S}}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N f_i^{\mathcal{S}}(\mathbf{X})$, we have

$$\nabla f^{\mathcal{S}}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N \nabla f_i^{\mathcal{S}}(\mathbf{X}). \quad (57)$$

Since $f_i^{\mathcal{S}}(\mathbf{X})$ is $L \frac{r}{k_i}$ -smooth, we thus have

$$\|\nabla f_i^{\mathcal{S}}(\mathbf{X}) - \nabla f_i^{\mathcal{S}}(\mathbf{Y})\| \leq L \frac{r}{k_i} \|\mathbf{X} - \mathbf{Y}\|, \quad \forall \mathbf{X}, \mathbf{Y}. \quad (58)$$

To find the Lipschitz constant of $f^{\mathcal{S}}(\mathbf{X})$, we analyze the difference between the gradients at two points $\mathbf{X}$ and $\mathbf{Y}$:

$$\begin{aligned} \|\nabla f^{\mathcal{S}}(\mathbf{X}) - \nabla f^{\mathcal{S}}(\mathbf{Y})\| &= \left\| \frac{1}{N} \sum_{i=1}^N (\nabla f_i^{\mathcal{S}}(\mathbf{X}) - \nabla f_i^{\mathcal{S}}(\mathbf{Y})) \right\| \\ &\leq \frac{1}{N} \sum_{i=1}^N \|\nabla f_i^{\mathcal{S}}(\mathbf{X}) - \nabla f_i^{\mathcal{S}}(\mathbf{Y})\| \\ &\leq \left(\frac{1}{N} \sum_{i=1}^N \frac{r}{k_i} L \right) \|\mathbf{X} - \mathbf{Y}\|. \end{aligned} \quad (59)$$

Therefore, $f^{\mathcal{S}}(\mathbf{X})$ is $\left(\frac{1}{N} \sum_{i=1}^N \frac{r}{k_i} L \right)$ -smooth.

J.5. Extension to Importance-Based Sketching

In this work, we provide a theoretically grounded framework for heterogeneous LoRA fine-tuning. Although our algorithm and analysis are based on Random- k sketching, the framework can be extended to importance-based sketching when reliable importance scores for the rows and columns of the global LoRA modules are available.

Without loss of generality, we assume that each column-row pair $(\mathbf{b}_i, \mathbf{a}_i)$ of $\mathbf{B}$ and $\mathbf{A}$ is associated with a static importance weight w_i , for $i = 1, \dots, r$, where $w_1 \geq w_2 \geq \dots \geq w_r$. This assumption is reasonable because permuting the column-row pairs does not change the product $\mathbf{B}\mathbf{A}$, as long as the same permutation is applied to both $\mathbf{B}$ and $\mathbf{A}$. Consequently, the global LoRA modules can always be reordered so that the components are indexed in descending order of importance prior to sketching in each training round.

With the aforementioned importance scores, we sample k indices from $\{1, \dots, r\}$ without replacement according to the weights $\{\frac{w_i}{\sum_{j=1}^r w_j}\}_{i=1}^r$, and denote the selected index set by $\mathcal{I}$. The diagonal entries of the sketching matrix $\mathbf{S}$ are defined as

$$s_i = \begin{cases} \frac{1}{\pi_i}, & i \in \mathcal{I}, \\ 0, & i \notin \mathcal{I}, \end{cases} \quad \text{where} \quad \pi_i = \frac{k w_i}{\sum_{j=1}^r w_j}. \quad (60)$$

Here, π_i denotes the marginal probability that index i is selected. With this construction, the sketch remains unbiased:

$$\mathbb{E}[\mathbf{BSA}] = \mathbf{B} \mathbb{E}[\mathbf{S}] \mathbf{A} = \mathbf{BA}.$$

The *sparsity* of the stochastic sketching matrix can be controlled by the choice of k .

Consequently, the convergence analysis for Random- k sketching extends directly to the importance-based setting, with the sketch-dependent quantities modified to reflect the non-uniform sampling probabilities. In particular, Lemma J.4 is updated as follows.

Lemma J.6. *Let $\mathbf{S}$ be a stochastic diagonal sketching matrix defined in (60). Then for any matrix $\mathbf{X}$, we have*

$$\|\mathbf{XS}\|^2 \leq \left(\frac{\sum_{j=1}^r w_j}{kw_r} \right)^2 \|\mathbf{X}\|^2 \quad \text{and} \quad \mathbb{E}_{\mathbf{S}}[\|\mathbf{XS}\|^2] \leq \frac{\sum_{j=1}^r w_j}{kw_r} \|\mathbf{X}\|^2. \quad (61)$$

Building on Lemma J.6, the smoothness characterization in Proposition J.5 is correspondingly updated as follows.

Proposition J.7. *Let $\mathcal{S}_i$ denote the set of sketching matrices, where each matrix is defined as in (60) and has k_i non-zero diagonal entries. Under Assumption 4.1, $\tilde{f}_i(\mathbf{X}; \mathbf{S}) = f_i(\mathbf{BS}, \mathbf{A})$, $f_i^{\mathcal{S}}(\mathbf{X}) = \mathbb{E}_{\mathbf{S} \sim \mathcal{S}_i}[\tilde{f}_i(\mathbf{X}; \mathbf{S})]$, and $f^{\mathcal{S}}(\mathbf{X}) = \frac{1}{N} \sum_{i=1}^N f_i^{\mathcal{S}}(\mathbf{X})$ are smooth with parameters $L \left(\frac{\sum_{j=1}^r w_j}{kw_r} \right)^2$, $L \frac{\sum_{j=1}^r w_j}{kw_r}$, and $\left(\frac{1}{N} \sum_{i=1}^N \frac{\sum_{j=1}^r w_j}{kw_r} \right) L$, respectively.*

With Proposition J.7, we can obtain a convergence bound analogous to Theorem 4.4, in which the sketching-dependent coefficients $\frac{r}{k_i}$ and $\frac{r^2}{k_i^2}$ are replaced by $\frac{\sum_{j=1}^r w_j}{k_i w_r}$ and $\left(\frac{\sum_{j=1}^r w_j}{k_i w_r} \right)^2$, respectively.

Theorem J.8. *Suppose that Assumptions 4.1-4.3 hold, then there exists a learning rate such that the iterates $\{\mathbf{X}^t\}$ generated by FSLoRA satisfy*

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \|\nabla_{\mathbf{X}} f^{\mathcal{S}}(\mathbf{X}^t)\|^2 \leq 8 \frac{\sqrt{\bar{L}' \mathcal{F}_0 \sigma_{\rho}^2}}{\sqrt{NTH}} + 10(\tilde{L}' L)^{\frac{1}{3}} \left(\frac{\mathcal{F}_0 \sigma_{\rho}}{T} \right)^{\frac{2}{3}} + \frac{4\mathcal{F}_0}{T}, \quad (62)$$

where $\sigma_{\rho}^2 = \sigma^2 + 3(\rho + 1)\sigma_h^2$, $\bar{L}' = \left(\frac{1}{N} \sum_{i=1}^N \frac{\sum_{j=1}^r w_j}{k_i w_r} \right) L$, $\tilde{L}' = \left(\frac{1}{N} \sum_{i=1}^N \left(\frac{\sum_{j=1}^r w_j}{k_i w_r} \right)^2 \right) L$ and $\mathcal{F}_0 = f^{\mathcal{S}}(\mathbf{X}^0) - f^*$ with f^* denoting the lower bound of $f^{\mathcal{S}}(\mathbf{X})$.

The proofs of Lemma J.6 and Proposition J.7, and Theorem J.8 follow those of Lemma J.4, Proposition J.5, and Theorem 4.4, respectively, with the modified sketching constants incorporated accordingly. This completes the extension to the importance-aware sketching.