

# FedNano: Toward Lightweight Federated Tuning for Pretrained Multimodal Large Language Models

Anonymous ACL submission

## Abstract

Multimodal Large Language Models (MLLMs) excel in tasks like multimodal reasoning and cross-modal retrieval but face deployment challenges in real-world scenarios due to distributed multimodal data and strict privacy requirements. Federated Learning (FL) offers a solution by enabling collaborative model training without centralizing data. However, realizing FL for MLLMs presents significant challenges, including high computational demands, limited client capacity, substantial communication costs, and heterogeneous client data. Existing FL methods assume client-side deployment of full models, an assumption that breaks down for large-scale MLLMs due to their massive size and communication demands. To address these limitations, we propose **FedNano**, the first FL framework that centralizes the LLM on the server while introducing *NanoEdge*, a lightweight module for client-specific adaptation. *NanoEdge* employs modality-specific encoders, connectors, and trainable *NanoAdapters* with low-rank adaptation. This design eliminates the need to access or modify the LLM on clients, reducing client-side storage by **95%** and communication overhead to just **0.01%** of model parameters. By transmitting only compact *NanoAdapter* updates, *FedNano* handles heterogeneous client data and resource constraints while preserving privacy. Experiments demonstrate that *FedNano* outperforms prior FL baselines, bridging the gap between MLLM scale and FL feasibility, and enabling scalable, decentralized multimodal AI systems.

## 1 Introduction

Multimodal Large Language Models (MLLMs) (Zhu et al., 2023; Liu et al., 2024b; Peng et al., 2023b; Alayrac et al., 2022; Li et al., 2023) excel in tasks like cross-modal retrieval (Yin et al., 2024), making them indispensable for applications such as visual question answering (VQA) (Antol et al., 2015). However, real-world deployment remains

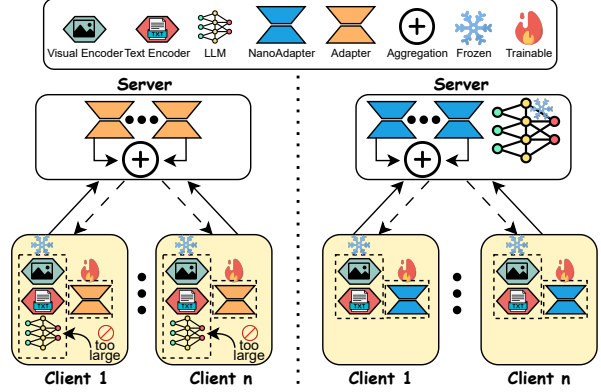


Figure 1: Comparison between traditional PEFT-based FL (left) and our proposed *FedNano* (right). *FedNano* keeps the LLM centralized on the server and performs lightweight tuning on clients, reducing both computation and communication overhead.

fundamentally constrained: multimodal data is inherently decentralized and privacy-sensitive, while the large parameter footprint of MLLMs renders on-device execution infeasible for edge clients.

Federated learning (FL) (McMahan et al., 2017) offers a promising solution for decentralized multimodal training. However, applying FL to MLLMs presents fundamental system-level challenges. First, although parameter-efficient fine-tuning (PEFT) (Houlsby et al., 2019; Lester et al., 2021; Zaken et al., 2021; Hu et al., 2021) reduces the number of trainable parameters, it still requires deploying the full MLLM—often exceeding 10 billion parameters—on each client, which is impractical for resource-constrained devices such as mobile phones or IoT systems. Second, PEFT methods typically insert adapters into internal layers of the language model, requiring structural access and full-model execution on clients, as seen in recent FL adaptations such as FedDPA-F (Yang et al., 2024), pFedLoRA (Yi et al., 2023), and FedIT (Zhang et al., 2024). Third, the resulting adapter updates remain sizable, imposing substantial communication overhead across training rounds. Finally, non-IID

client data introduces statistical heterogeneity that degrades global model convergence. These limitations collectively constrain the scalability and practicality of existing FL approaches for MLLMs.

To address these challenges, we propose *FedNano*, the first FL framework that enables MLLM adaptation without deploying LLM on clients. *FedNano* centralizes LLM on the server in a frozen state, and equips each client with *NanoEdge*—a lightweight adaptation module comprising modality-specific encoders, connectors, and trainable *NanoAdapters*. These adapters operate externally to LLM and are optimized using low-rank decomposition (Hu et al., 2021), minimizing both parameter size and transmission cost. This design removes the need for local LLM access, reduces client storage by over **95%** (Tab. 1), and enables deployment on edge devices. Only compact *NanoAdapter* updates are exchanged across training rounds, achieving **over 99% communication reduction** compared to PEFT-based FL methods (Chen et al., 2023; Yang et al., 2024), as illustrated in Fig. 1. By decoupling adaptation from the LLM, *FedNano* provides a scalable and communication-efficient solution for real-world MLLM deployment.

To address client heterogeneity, *FedNano* adapts Fisher Merging (Matena and Raffel, 2022) to align global updates with client-specific data distributions. This adaptation improves performance on non-IID datasets and outperforms traditional aggregation methods such as FedAvg (McMahan et al., 2017) and FedProx (Li et al., 2020). By integrating these innovations, *FedNano* effectively bridges the gap between the computational complexity of MLLMs and the constraints of FL, enabling efficient deployment in real-world scenarios.

Experiments across diverse MLLM and multimodal tasks demonstrate that *FedNano* not only outperforms existing methods but also significantly reduces resource and communication costs, enabling the scalable, efficient, and privacy-preserving deployment of MLLMs. This framework lays a strong foundation for advancing multimodal AI systems in decentralized real-world applications, including personalized healthcare, cross-device collaboration, and multimodal user interfaces.

The key contributions of this work are:

1. Novel FL Architecture for MLLMs: We propose *FedNano*, the first framework that centralizes the LLM on the server and enables lightweight client-side adaptation via *NanoEdge*, reducing client storage by over **95%** and enabling practical deployment

Approach	Client Params	Server Uploads
<i>FedNano</i>	304.55M (4.30%)	1.05M (0.01%)
FedDPA-F	7222.81M (100%)	180.89M (2.50%)
<b>Reduction</b>	<b>↓ 95.8%</b>	<b>↓ 99.4%</b>

Table 1: Comparison of parameter distribution and communication efficiency between *FedNano* and FedDPA-F (Yang et al., 2024) on LLaVA-1.5-7B (Liu et al., 2024b). *Client Params* denotes parameters retained on client devices, while *Server Uploads* denotes parameter updates sent to the server per round. Both use rank-64 adapters.

ment on resource-constrained devices.

2. Communication-Efficient Adaptation: *FedNano* employs low-rank decomposition in *NanoAdapters*, achieving an over 99% reduction in the number of transmitted parameters, allowing efficient deployment in bandwidth-constrained environments.

3. Improved Generalization on Non-IID Data: We adapt Fisher Merging for FL, aligning global updates with client-specific distributions to improve model performance on heterogeneous datasets.

4. Comprehensive Validation: Extensive experiments demonstrate the effectiveness and efficiency of *FedNano*, establishing it as a scalable solution for real-world MLLM deployment.

## 2 Related Work

### 2.1 Multimodal Large Language Models

MLLMs (Zhu et al., 2023; Liu et al., 2024b; Peng et al., 2023b; Alayrac et al., 2022; Li et al., 2023; Dai et al., 2023) extend LLMs (Touvron et al., 2023; Peng et al., 2023a; Bai et al., 2023) by integrating modality-specific encoders and connectors to process multimodal inputs. Recent works focus on efficient alignment, using lightweight connectors such as the linear projection in MiniGPT-4 (Zhu et al., 2023) or the MLP bridge in LLaVA (Liu et al., 2024b). However, these models assume full model access, which is incompatible with federated settings due to privacy and resource constraints. *FedNano* resolves this by freezing the LLM on the server and enabling lightweight client-side adaptation via *NanoAdapters*.

### 2.2 Parameter Efficient Fine-tuning

PEFT techniques (Houlsby et al., 2019; Lester et al., 2021; Zaken et al., 2021; Hu et al., 2021) adapt large pretrained models by updating only a small set of parameters, significantly reducing training costs. They include additive methods like

adapters (Houlsby et al., 2019) and soft prompts (Lester et al., 2021), selective tuning such as Bit-Fit (Zaken et al., 2021), and reparameterization methods like LoRA (Hu et al., 2021). While effective in centralized settings, PEFT-based FL methods (Chen et al., 2023; Wang et al., 2024; Zhang et al., 2024) assume the full model, including LLM, can be deployed on clients. This becomes impractical for MLLMs, where LLM accounts for the vast majority of parameters and cannot be hosted on resource-limited devices. To overcome this, *FedNano* introduces a new paradigm: the LLM is frozen and centralized on the server, while lightweight *NanoAdapters* are deployed on clients. This design eliminates the need for full-model access, reduces client overhead, and enables scalable FL for MLLM. Unlike conventional PEFT, which are inserted into LLM, *NanoAdapters* operate externally, interfacing solely through the modality connector. This allows adaptation without modifying or executing LLM on clients.

## 2.3 Multimodal Federated Learning

Multimodal FL has gained increasing attention for handling data heterogeneity and privacy constraints in real-world deployments. Prior work has focused on vision-language models, proposing strategies for modality imbalance (Yu et al., 2023; Che et al., 2024), non-IID distributions (Yang et al., 2024; Zhang et al., 2024; Chen and Zhang, 2024), and client personalization (Yi et al., 2023; Chen et al., 2023). Benchmarks like FedMultimodal (Feng et al., 2023) and FedMLLM (Xu et al., 2024) further standardize evaluation in heterogeneous multimodal settings. However, these methods still rely on client-side full model deployment. For MLLMs, this becomes infeasible due to their scale. Even with PEFT, deploying full MLLMs locally remains out of reach, and transmitting adapter updates still incurs significant communication overhead. *FedNano* departs from this design by keeping the LLM on the server and transmitting only compact *NanoAdapter* updates from clients. This makes it the first scalable FL framework tailored for large-scale MLLMs, enabling efficient multimodal collaboration without sacrificing practicality.

## 3 Methodology

### 3.1 Problem Definition

This work addresses federated fine-tuning for MLLMs in decentralized environments with sta-

tistical data heterogeneity. Each client  $k$  holds a private multimodal dataset  $D_k = \{(v_k^i, q_k^i, a_k^i)\}$ , comprising image-question-answer triplets. We assume complete modality availability and a shared model architecture across all clients; only data distributions differ (Chen et al., 2023). The marginal distributions of  $v_k^i$ ,  $q_k^i$ , and  $a_k^i$  vary across clients, resulting in shifts in both visual and textual representations, as well as answer semantics. Such heterogeneity poses challenges for achieving consistent generalization, as standard aggregation strategies struggle to align diverse local updates.

Our objective is to collaboratively fine-tune a shared global MLLM for VQA (Antol et al., 2015). Following (Liu et al., 2024a), we formulate this as an open-ended generation problem, where the model generates free-form answers given image-question pairs. Existing approaches assume that the full MLLM can be deployed on each client, which is infeasible in practice due to the massive size of LLM backbones. Client devices often lack sufficient compute, memory, and bandwidth to support such models, and privacy regulations further restrict centralized data access. These constraints call for a new FL framework that avoids client-side LLM deployment while enabling efficient adaptation and communication. To address these challenges, we propose *FedNano*, a parameter-efficient framework that centralizes the computationally intensive LLM on the server while enabling lightweight, client-specific tuning. In the following sections, we detail the design of *FedNano*, focusing on how it minimizes computational and communication overhead and addresses data heterogeneity.

### 3.2 Overview of *FedNano* Architecture

*FedNano* is designed to address the inherent challenges of deploying MLLMs in FL environments. As shown in Fig. 2, it introduces a new architecture that centralizes the computationally intensive LLM on the server, while clients maintain only lightweight *NanoEdge* modules for task-specific adaptation. *NanoEdge* freezes the modality encoders and connector, restricting training to task-specific *NanoAdapters*. This design eliminates the need to deploy the full model on resource-constrained devices, reducing client-side computation and enabling edge deployment on mobile or IoT systems. The complete training and aggregation algorithm is provided in Appendix A.

*FedNano* jointly addresses three key challenges in MLLM-based FL: high computation, commu-

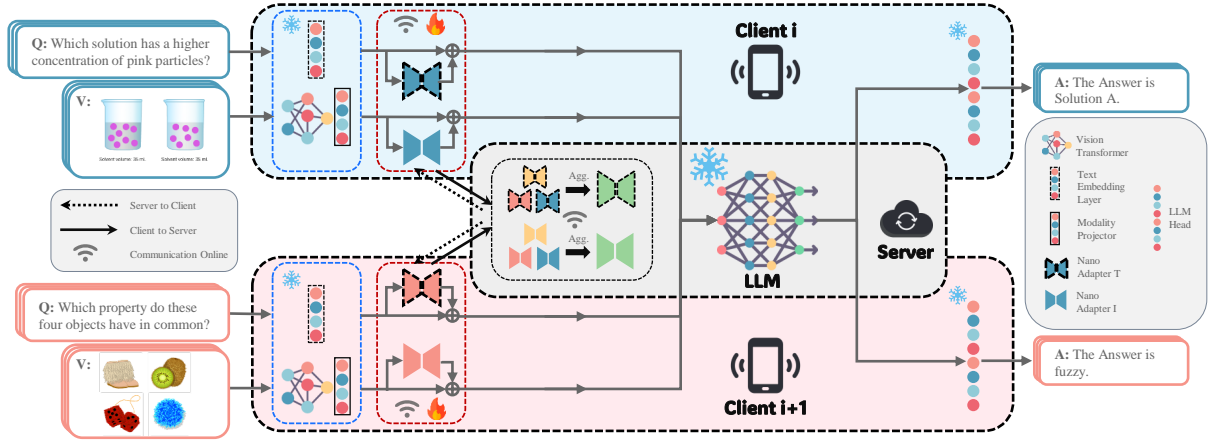


Figure 2: Overview of the *FedNano* framework. The server hosts the frozen LLM, while each client performs local tuning via *NanoEdge*, which includes *NanoAdapter-T* for text and *NanoAdapter-I* for vision. Clients upload low-rank adapter updates, which are aggregated on the server using Fisher merging. This design reduces client overhead and supports scalable, multimodal federated learning under data heterogeneity.

256 nication cost, data heterogeneity. By offloading  
 257 the LLM to the server, clients train only the *Na-*  
 258 *nanoEdge* module, which includes frozen encoders  
 259 and a connector, and optimizes a small set of  
 260 *NanoAdapters* for task-specific adaptation. The to-  
 261 tal client-side module accounts for less than 5% of  
 262 the model parameters, while the trainable *NanoAd-*  
 263 *apters* comprise only 0.01%. During aggregation,  
 264 only *NanoAdapters* updates are uploaded, signifi-  
 265 cantly reducing communication overhead. *NanoAd-*  
 266 *apters* are optimized via low-rank decomposition,  
 267 enabling expressive local tuning while preserv-  
 268 ing pretrained alignment with the frozen LLM.  
 269 This compact update mechanism supports low-  
 270 bandwidth environments and enhances training effi-  
 271 ciency. To address data heterogeneity, *FedNano* in-  
 272 tegrates Fisher Merging (Matena and Raffel, 2022)  
 273 into FL as an advanced aggregation strategy, lever-  
 274 aging client-specific posterior estimates to align  
 275 local updates with global objectives. By weight-  
 276 ing and combining *NanoAdapter* updates based  
 277 on their estimated importance, this method im-  
 278 proves robustness across diverse tasks and datasets,  
 279 even under non-IID conditions. Together with its  
 280 architectural and optimization designs, *FedNano*  
 281 bridges the gap between the computational barriers  
 282 of MLLM deployment and the practical constraints  
 283 of FL, offering a scalable, efficient, and privacy-  
 284 preserving solution for decentralized multimodal  
 285 learning.

### 3.3 *NanoEdge*: Client-Side Tuning Module

287 MLLMs are composed of three key components:  
 288 modality encoders, a connector, and a pretrained

289 LLM backbone. The modality encoders extract  
 290 embeddings from raw inputs, such as images and  
 291 text, while the connector aligns these embeddings  
 292 into a unified representation compatible with the  
 293 LLM. Together, these components enable MLLMs  
 294 to effectively handle diverse multimodal tasks by  
 295 leveraging their pretrained capabilities.

296 Building on this structure, *NanoEdge* introduces  
 297 *NanoAdapters* at the interface between the connec-  
 298 tor and the LLM to facilitate efficient task-specific  
 299 tuning while preserving the pretrained alignment  
 300 across modalities. By freezing the modality en-  
 301 coders and the connector, *NanoEdge* maintains  
 302 their alignment with the LLM, ensuring the foun-  
 303 dational structure of the pretrained model remains  
 304 intact. This design allows *NanoAdapters* to focus  
 305 solely on learning task-specific patterns from lo-  
 306 cal client data and integrating federated knowledge  
 307 updates, avoiding any disruption to the pretrained  
 308 alignment. By restricting training to the lightweight  
 309 *NanoAdapter* parameters, *NanoEdge* minimizes  
 310 client-side computational demands while enabling  
 311 efficient and privacy-preserving adaptation.

312 The *NanoAdapters* employ a low-rank decom-  
 313 position mechanism, inspired by LoRA (Hu et al.,  
 314 2021), consisting of a down-projection to reduce  
 315 embedding dimensionality and an up-projection  
 316 to restore it. This design balances parameter ef-  
 317 ficiency and adaptation capability, enabling *Na-*  
 318 *nanoEdge* to perform localized tuning and transmit  
 319 updates efficiently. Each modality is equipped with  
 320 a dedicated *NanoAdapter*— $\mathcal{A}_I$  for images and  $\mathcal{A}_T$   
 321 for text—capturing modality-specific patterns es-  
 322 sential for multimodal tasks. Unlike traditional

adapters that are inserted into LLM, *NanoAdapters* remain externally attached to the modality connector, requiring no structural access to or execution of LLM. This makes them uniquely compatible with server-hosted LLMs in federated environments.

### 3.4 Fisher-Guided Adaptive Aggregation

In FL, model aggregation can be interpreted as maximizing the joint posterior likelihood across clients. Traditional methods like FedAvg implicitly assume isotropic Gaussian posteriors (Matena and Raffel, 2022), which oversimplifies client uncertainty and leads to degraded performance under data heterogeneity. *FedNano* addresses this limitation by adopting Fisher Merging (Matena and Raffel, 2022), which leverages the Laplace approximation for more accurate posterior estimation. The global update is computed as:

$$\theta_{global} = \frac{\sum_{k=1}^K \frac{|D_k|}{\sum_{k=1}^K |D_k|} F_k \theta_k}{\sum_{k=1}^K \frac{|D_k|}{\sum_{k=1}^K |D_k|} F_k}, \quad (1)$$

where  $\theta_k$  denotes the *NanoAdapter* parameters of client  $k$ ,  $F_k$  is the Fisher Information Matrix (FIM), which serves as the precision matrix of the Laplace approximation, and  $D_k$  is the local dataset. This weighting improves the alignment of local updates with their estimated importance, enhancing generalization under non-IID data. To ensure scalability, *FedNano* approximates the full FIM with its diagonal (Kirkpatrick et al., 2017), and computes it efficiently from squared gradients during backpropagation (Wu et al., 2023). This reduces computation from  $O(|\theta|^2)$  to  $O(|\theta|)$  without sacrificing aggregation accuracy. Compared to uniform averaging, this method dynamically prioritizes impactful updates, achieving stronger global performance under statistical heterogeneity.

## 4 Experiment

### 4.1 Experimental Setup

We evaluate our approach on the Visual Question Answering (VQA) task using two established benchmarks: ScienceQA (Lu et al., 2022) and IconQA (Lu et al., 2021). These datasets were selected for their well-defined categorical structures and multimodal complexities, making them particularly suitable for assessing the performance of FL in non-IID settings. To simulate FL in non-IID setting, we partitioned the datasets using Dirichlet distributions following (Che et al., 2023; Lai et al.,

2022; Zhang et al., 2024) with a concentration parameter  $\alpha = 1$  to create strongly non-IID splits. Partitioning was guided by topic annotations in ScienceQA and skill annotations in IconQA, ensuring heterogeneous yet meaningful distributions across five simulated clients. Each partition, representing an individual client dataset, maintains consistent train-validation-test splits for evaluation. We evaluate our approach on MiniGPT-4 (Zhu et al., 2023) and LLaVA-1.5 (Liu et al., 2024b).

### 4.2 Implementation Details

**Baselines** To the best of our knowledge, *FedNano* is the first FL framework specifically designed to support MLLMs by centralizing the LLM on the server. This architectural shift renders existing PEFT-based FL methods inapplicable, as they assume full-model access and local integration with the LLM. Given the absence of prior work addressing this setting, we evaluate *FedNano* against three representative FL baselines: FedAvg (McMahan et al., 2017), a foundational aggregation method with limited handling of data heterogeneity; FedProx (Li et al., 2020), which mitigates client drift through a proximal term but lacks parameter-specific adaptation; and FedDPA-F (Yang et al., 2024), which integrates advanced alignment strategies but incurs high computational and communication overheads. We further include comparisons with a centralized model, representing the performance upper bound achieved with access to all data, and locally fine-tuned models, which operate in isolation without collaboration.

**Training Configurations** The training process includes 10 communication rounds ( $R = 10$ ), with each client performing one local epoch per round using a batch size of 8. All experiments were conducted on NVIDIA A100 80G GPUs.

### 4.3 Main Results

Results in Tab. 2 demonstrate that FL methods consistently outperform locally fine-tuned models (LocFT), emphasizing the benefit of global knowledge sharing in distributed, heterogeneous settings.

*FedNano* achieves the highest average performance among all FL methods, more effectively narrowing the gap to centralized training than existing baselines. While FedAvg performs competitively with simple weighted averaging, its inability to adapt to non-IID data results in suboptimal performance under heterogeneous distributions. FedProx

Backbone	Approach	ScienceQA (Clients)						IconQA (Clients)					
		C1	C2	C3	C4	C5	Avg	C1	C2	C3	C4	C5	Avg
MiniGPT-4	Centralized	73.70	88.34	89.83	84.52	87.41	84.76	80.76	86.62	81.16	82.74	85.36	83.33
	LocFT	67.74	74.69	77.42	72.46	74.07	73.28	67.70	73.48	70.63	70.86	77.53	72.04
	FedAvg	70.22	79.65	79.65	75.19	75.56	76.05	70.31	75.61	74.98	72.76	81.25	74.98
	FedProx	70.97	80.40	80.15	75.19	75.80	76.50	70.94	77.36	74.58	71.50	80.70	75.01
	FedDPA-F	<b>71.96</b>	78.41	<b>81.14</b>	76.42	75.80	76.75	70.94	<b>77.91</b>	74.51	73.08	80.30	75.35
	<i>FedNano</i>	68.98	<b>81.89</b>	80.89	<b>76.43</b>	<b>77.04</b>	<b>77.05</b>	<b>72.21</b>	77.28	<b>75.85</b>	<b>74.27</b>	<b>82.52</b>	<b>76.42</b>
LLaVA-1.5	Centralized	83.87	91.07	89.33	90.57	89.38	88.84	86.62	88.92	84.88	87.25	88.45	87.22
	LocFT	71.96	80.89	76.92	79.65	75.80	77.04	<b>75.93</b>	78.94	72.53	74.35	76.50	75.65
	FedAvg	73.20	<b>84.37</b>	83.62	82.13	<b>80.49</b>	80.76	71.18	79.89	76.80	<b>77.51</b>	<b>83.23</b>	77.72
	FedProx	73.95	<b>84.37</b>	83.87	81.39	80.00	80.71	70.23	80.13	76.72	<b>77.51</b>	82.36	77.39
	FedDPA-F	73.70	84.12	84.12	81.89	79.51	80.67	72.12	79.65	76.80	77.43	82.36	77.68
	<i>FedNano</i>	<b>74.94</b>	84.12	<b>84.86</b>	<b>82.88</b>	80.25	<b>81.41</b>	72.13	<b>80.44</b>	<b>77.36</b>	77.43	82.83	<b>78.04</b>

Table 2: Performance comparison. Results include centralized training, local fine-tuning (LocFT), and various federated approaches. *FedNano* achieves superior average performance on both datasets compared to other federated approaches, demonstrating its effectiveness in handling client heterogeneity.

Approach	$\alpha = 0.1$						$\alpha = 5$					
	C1	C2	C3	C4	C5	Avg	C1	C2	C3	C4	C5	Avg
LocFT	69.94	75.80	75.48	73.18	77.00	74.28	65.71	70.62	71.41	72.76	70.64	70.22
FedAvg	72.80	76.80	75.50	73.20	73.60	74.38	74.34	75.61	72.92	<b>76.08</b>	74.68	74.72
FedProx	71.54	74.79	74.15	69.72	70.06	73.05	68.48	70.30	70.15	70.15	71.04	70.02
FedDPA-F	70.25	76.40	74.10	72.50	<b>78.55</b>	74.27	71.52	<b>76.83</b>	<b>74.51</b>	73.24	<b>75.84</b>	74.38
<i>FedNano</i>	<b>73.85</b>	<b>78.22</b>	<b>80.14</b>	<b>76.28</b>	74.94	<b>76.68</b>	<b>74.90</b>	76.16	74.18	74.82	73.73	<b>74.75</b>

Table 3: Performance of MiniGPT-4 on IconQA under different data heterogeneity levels. *FedNano* consistently outperforms all baselines, with the largest gains under highly non-IID conditions.

mitigates client drift by constraining local updates toward the global model, but this rigid constraint limits flexibility, making it insufficient for complex multimodal tasks. FedDPA-F, though designed for personalization, requires careful tuning of global training epochs and risks overwriting the global adapter during local updates, potentially degrading performance due to catastrophic forgetting.

In contrast, the superior performance of *FedNano* is attributed to its novel design and optimization strategies. As shown in Tab. 2, *FedNano* achieves an average accuracy of 77.05% on ScienceQA and 76.42% on IconQA for MiniGPT-4, exceeding FedAvg and FedProx, indicating improved generalization in heterogeneous client environments. For LLaVA, *FedNano* attains 81.41% on ScienceQA and 78.04% on IconQA, surpassing FedDPA-F and FedProx, demonstrating enhanced robustness in multimodal FL. These results validate the effectiveness of *NanoAdapters* for modality-specific adaptation, while substantially reducing client-side computational and storage demands, enabling deployment on resource-limited devices.

Moreover, *FedNano* integrates Fisher Merging with a diagonal approximation of the FIM, allowing the system to prioritize critical parameter updates based on client-specific confidence. This results in more effective aggregation than uniform averaging, improving stability under non-IID distributions while reducing overfitting to local client noise. By balancing generalization and personalization, *FedNano* consistently delivers strong performance across diverse client settings, all while maintaining minimal communication overhead.

#### 4.4 Analysis

**Robustness under Data Heterogeneity** To assess the robustness of *FedNano* under varying levels of data heterogeneity, we evaluate its performance on IconQA using the MiniGPT-4 backbone across different Dirichlet concentration values ( $\alpha = 0.1$  and  $\alpha = 5$ ). As shown in Tab. 3, *FedNano* consistently achieves the highest average accuracy in the highly non-IID setting ( $\alpha = 0.1$ ), outperforming all FL baselines. This demonstrates the effectiveness of its Fisher-guided aggregation

Approach	C1	C2	C3	C4	C5	C6	C7	C8	C9	C10	Avg
LocFT	67.56	69.77	73.89	67.24	79.90	72.15	69.77	64.71	71.67	67.35	70.40
FedAvg	74.52	81.01	78.00	78.63	85.91	79.90	75.94	75.63	70.90	77.86	77.83
FedProx	73.89	76.74	77.37	75.63	84.01	76.58	73.41	71.36	78.79	72.29	76.00
FedDPA-F	74.52	81.01	78.00	78.63	85.91	79.90	75.94	75.63	70.90	77.86	77.83
<b>FedNano</b>	<b>77.03</b>	<b>82.77</b>	<b>78.22</b>	<b>79.67</b>	<b>88.57</b>	<b>80.35</b>	<b>81.34</b>	<b>72.84</b>	<b>73.77</b>	<b>79.47</b>	<b>78.86</b>

Table 4: Performance of MiniGPT-4 on IconQA with 10 simulated clients. *FedNano* achieves the best average accuracy, demonstrating strong scalability to larger federated setups.

Approach	C1	C2	C3	C4	Avg
FedAvg	34.35	28.83	29.00	29.53	30.86
FedProx	52.45	50.82	59.80	42.15	51.30
FedDPA-F	52.76	51.12	60.10	42.46	51.61
<b>FedNano</b>	<b>54.20</b>	<b>52.60</b>	<b>60.36</b>	<b>43.32</b>	<b>52.62</b>

Table 5: Performance under cross-task federated setup on MiniGPT-4. *FedNano* achieves the best average accuracy across clients with distinct VQA tasks.

Backbone	Variants	ScienceQA	IconQA
MiniGPT-4	$\mathcal{A}_T$	45.91	57.77
	$\mathcal{A}_I$	74.57	75.17
	$\mathcal{A}_T + \mathcal{A}_I$	<b>76.42</b>	<b>76.04</b>
LLaVA-1.5	$\mathcal{A}_T$	50.08	48.15
	$\mathcal{A}_I$	77.03	77.12
	$\mathcal{A}_T + \mathcal{A}_I$	<b>78.04</b>	<b>77.83</b>

Table 6: Performance of different adapters. Combining  $\mathcal{A}_T$  and  $\mathcal{A}_I$  consistently yields the best results across backbones, confirming their complementarity.

in aligning heterogeneous client updates. While the performance gap narrows under near-IID conditions ( $\alpha = 5$ ), *FedNano* remains competitive, indicating that its advantages are most pronounced in realistic heterogeneous federated scenarios.

**Scalability to Larger Client Populations** To evaluate the scalability of *FedNano*, we extend the number of clients from 5 to 10 on the IconQA dataset using the MiniGPT-4 backbone. As shown in Tab. 4, *FedNano* achieves the highest average accuracy, consistently outperforming all baselines. This demonstrates that the framework retains its effectiveness even as the federated environment becomes more fragmented. The results confirm that *FedNano* scales robustly with increasing client population, reinforcing its practicality for real-world large-scale federated deployments.

**Generalization under Cross-Task Client Distribution** We evaluate *FedNano* in a challenging cross-task setup where four clients are respectively assigned A-OKVQA, OK-VQA, IconQA, and GQA, introducing significant task-level heterogeneity. As shown in Tab. 5, *FedNano* achieves stable and strong performance across all clients. This robustness stems from its modular design and Fisher-guided aggregation, which enable effective alignment of heterogeneous updates and support generalization across semantically diverse tasks.

**The Necessity of Combining Both  $\mathcal{A}_T$  and  $\mathcal{A}_I$**  To evaluate the necessity of the textual adapter  $\mathcal{A}_T$

and the visual adapter  $\mathcal{A}_I$ , we conduct ablation experiments using three configurations:  $\mathcal{A}_T$  only,  $\mathcal{A}_I$  only, and both. For MiniGPT-4,  $\mathcal{A}_T$  achieves 45.91% on ScienceQA and 57.77% on IconQA, while  $\mathcal{A}_I$  improves to 74.57% and 75.17%. Their combination further boosts accuracy to 76.42% and 76.04%, outperforming  $\mathcal{A}_I$  alone by +1.85% and +0.87%. Similar trends are observed with LLaVA-1.5, confirming the robustness of combining both adapters. The poor performance of  $\mathcal{A}_T$  alone suggests that textual inputs provide insufficient task-relevant information in these vision-centric VQA tasks. These results validate the dual-adapter design of NanoEdge, where  $\mathcal{A}_I$  handles visual adaptation and  $\mathcal{A}_T$  enhances generalization.

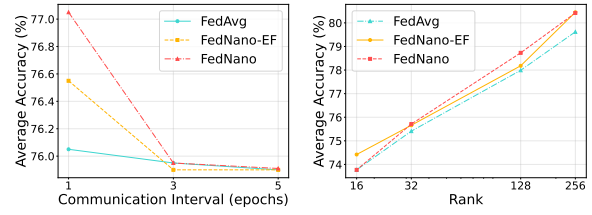
**Trade-offs in Fisher-Guided Adaptive Aggregation** FIM is specific to a particular set of model parameters and plays a key role in the ability of *FedNano* to achieve superior global alignment by capturing parameter importance. To compute the FIM precisely, *FedNano* employs additional forward and backward passes per communication round, ensuring accurate parameter estimation. While this enhances accuracy, it introduces modest computational overhead. To explore the trade-offs between precision and efficiency, we conduct an ablation study with *FedNano-EF*, a variant that approximates the FIM during standard training, elim-

Dataset	Variants	MiniGPT-4	LLaVA-1.5
ScienceQA	<i>FedNano</i>	<b>77.05</b>	<b>81.41</b>
	<i>FedNano-EF</i>	76.55	80.81
	FedAvg	76.05	80.76
	FedProx	76.50	80.71
	FedDPA-F	76.75	80.67
IconQA	<i>FedNano</i>	<b>76.42</b>	<b>78.04</b>
	<i>FedNano-EF</i>	76.04	77.83
	FedAvg	74.98	77.72
	FedProx	75.01	77.39
	FedDPA-F	75.35	77.68

Table 7: Performance comparison of *FedNano* and *FedNano-EF* on ScienceQA and IconQA. *FedNano* achieves the highest accuracy, while *FedNano-EF* offers a trade-off with reduced computational overhead, demonstrating strong performance across both datasets.

inating the need for additional computation steps. This modification reduces computational overhead to the level of FedAvg. Despite this simplification, *FedNano-EF* incurs only a slight accuracy trade-off and consistently outperforms baselines, as shown in Tab. 7. These results demonstrate the adaptability of *FedNano*: the standard version excels in accuracy-critical tasks by leveraging precise FIM computation to optimize alignment, while *FedNano-EF* provides a practical alternative for resource-constrained environments, achieving strong performance with reduced overhead.

**Frequent Communication Amplifies the Advantages of *FedNano*** As shown in Fig. 3a, reduced communication frequency leads to a general decline in global model performance across all methods due to increased parameter divergence, which hinders effective aggregation. Importantly, the results highlight that *FedNano* outperforms FedAvg by a larger margin when communication is more frequent. With shorter intervals, FIM mechanism of *FedNano* can better leverage aligned client parameters to prioritize impactful updates, amplifying its advantages in handling data heterogeneity. In contrast, FedAvg struggles with parameter divergence regardless of communication frequency, showing minimal improvement with more frequent updates. These findings underscore that while frequent communication benefits all methods, it significantly enhances the effectiveness of *FedNano*, reinforcing its superior ability to integrate client-specific updates and maintain robust performance in federated learning environments.



(a) Communication Freq.

(b) Adapter Rank

Figure 3: (a) Impact of communication frequency. *FedNano* outperforms FedAvg, with more frequent communication amplifying its advantages; (b) Effect of adapter rank. *FedNano* consistently achieves superior performance, demonstrating its ability to capture task-specific and client-specific information effectively.

**Higher Adapter Ranks Enhance *FedNano* Performance** Fig. 3b illustrates the impact of adapter rank, comparing *FedNano* with FedAvg on the ScienceQA dataset. As the adapter rank increases, accuracy improves due to the enhanced capacity to encode task-specific and client-specific information, which is particularly important in non-IID settings. However, higher ranks also incur greater communication costs, necessitating a trade-off between performance and resource efficiency in FL. *FedNano* consistently outperforms FedAvg across all ranks, with the performance gap widening at higher ranks. This improvement is driven by the FIM aggregation, which leverages richer client-specific updates at higher ranks to achieve better alignment between local contributions and the global model. In contrast, at lower ranks, the limited adapter capacity constrains the quality of updates, reducing the effectiveness of FIM aggregation.

## 5 Conclusion

This work introduced *FedNano*, an FL framework that tackles the unique challenges of deploying MLLMs in decentralized settings. By centralizing the LLM on the server and employing lightweight *NanoAdapters* on clients, *FedNano* achieves significant gains in both resource and communication efficiency, while effectively addressing data heterogeneity in non-IID environments. Comprehensive evaluations on ScienceQA and IconQA benchmarks demonstrate that *FedNano* consistently outperforms state-of-the-art FL baselines, further narrowing the gap between federated and centralized training. By combining scalable design with robust performance, *FedNano* offers a practical and privacy-preserving solution, advancing the real-world deployment of MLLMs.

## Limitation and Future Work

While FedNano demonstrates robust performance and efficiency, certain areas warrant further exploration to enhance its applicability and effectiveness. One limitation lies in the assumption that all clients possess similar hardware capabilities for managing NanoAdapters. This assumption may not hold in real-world scenarios characterized by highly heterogeneous devices. Future research could investigate adaptive mechanisms that dynamically tailor NanoAdapter configurations to match the computational resources and capabilities of individual clients, broadening FedNano usability across diverse environments.

Although FedNano effectively mitigates data heterogeneity, federated learning in real-world settings often involves extreme client disparities in data size, quality, and distribution. Addressing such scenarios may require dynamic strategies to adapt aggregation weights or incorporate more sophisticated representations of client-specific characteristics. These enhancements could further strengthen FedNano resilience and generalization capabilities in highly non-IID environments. Moreover, while the current framework supports vision and language modalities, extending it to incorporate audio, sensor data, or other modalities could unlock applications in areas such as autonomous systems, multimodal healthcare, and industrial IoT.

Deploying FedNano in noisy or incomplete federated datasets presents another promising avenue for research. Benchmarking its performance under these challenging conditions would not only provide valuable insights but also identify additional opportunities for optimization. Furthermore, integrating FedNano into federated multi-agent systems—where distinct agents collaborate to learn and share knowledge—could enable groundbreaking applications in fields like logistics and autonomous vehicles, highlighting the framework versatility.

Finally, while FedNano achieves strong privacy guarantees by transmitting only NanoAdapter updates, integrating advanced privacy-preserving methods such as differential privacy or secure multi-party computation could provide even stronger safeguards for sensitive client data. A critical future direction lies in achieving these enhanced privacy measures without compromising the computational and communication efficiency that underpins FedNano practicality.

In summary, while FedNano addresses many critical challenges in federated learning for MLLMs, these future directions highlight its potential for further innovation. By extending its capabilities to tackle more diverse environments, extreme heterogeneity, and advanced privacy requirements, FedNano can serve as a foundational framework that inspires continued advancements in federated learning research and applications.

## References

- Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. 2022. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. 2015. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pages 2425–2433.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, et al. 2023. Qwen technical report. *arXiv preprint arXiv:2309.16609*.
- Liwei Che, Jiaqi Wang, Xinyue Liu, and Fenglong Ma. 2024. [Leveraging foundation models for multi-modal federated learning with incomplete modality](#). *Preprint*, arXiv:2406.11048.
- Tianshi Che, Ji Liu, Yang Zhou, Jiaxiang Ren, Jiwen Zhou, Victor S Sheng, Huaiyu Dai, and Dejing Dou. 2023. Federated learning of large language models with parameter-efficient prompt tuning and adaptive optimization. *arXiv preprint arXiv:2310.15080*.
- Haokun Chen, Yao Zhang, Denis Krompass, Jindong Gu, and Volker Tresp. 2023. [Feddat: An approach for foundation model finetuning in multi-modal heterogeneous federated learning](#). *Preprint*, arXiv:2308.12305.
- Jiayi Chen and Aidong Zhang. 2024. On disentanglement of asymmetrical knowledge transfer for modality-task agnostic federated learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 11311–11319.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. [Instructblip: Towards general-purpose vision-language models with instruction tuning](#). *Preprint*, arXiv:2305.06500.
- Tiantian Feng, Digbalay Bose, Tuo Zhang, Rajat Hebbar, Anil Ramakrishna, Rahul Gupta, Mi Zhang, Salman Avestimehr, and Shrikanth Narayanan. 2023.

694	Fedmultimodal: A benchmark for multimodal federated learning. <i>Preprint</i> , arXiv:2306.09486.	747
695		748
696	Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morroni, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In <i>International conference on machine learning</i> , pages 2790–2799. PMLR.	749
697		750
698		751
699		752
700		753
701		
702	Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. <i>arXiv preprint arXiv:2106.09685</i> .	754
703		755
704		756
705		757
706		
707	James Kirkpatrick, Razvan Pascanu, Neil Rabinowitz, Joel Veness, Guillaume Desjardins, Andrei A Rusu, Kieran Milan, John Quan, Tiago Ramalho, Agnieszka Grabska-Barwinska, et al. 2017. Overcoming catastrophic forgetting in neural networks. <i>Proceedings of the national academy of sciences</i> , 114(13):3521–3526.	758
708		759
709		760
710		761
711		762
712		
713		
714	Fan Lai, Yinwei Dai, Sanjay Singapuram, Jiachen Liu, Xiangfeng Zhu, Harsha Madhyastha, and Mosharaf Chowdhury. 2022. Fedscale: Benchmarking model and system performance of federated learning at scale. In <i>International conference on machine learning</i> , pages 11814–11827. PMLR.	763
715		764
716		765
717		
718		
719		
720	Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. The power of scale for parameter-efficient prompt tuning. <i>arXiv preprint arXiv:2104.08691</i> .	766
721		767
722		768
723		769
724	Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. 2023. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In <i>International conference on machine learning</i> , pages 19730–19742. PMLR.	770
725		771
726		772
727		773
728		774
729	Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. 2020. Federated optimization in heterogeneous networks. <i>Proceedings of Machine learning and systems</i> , 2:429–450.	775
730		776
731		
732		
733	Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024a. Improved baselines with visual instruction tuning. In <i>Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition</i> , pages 26296–26306.	777
734		778
735		779
736		780
737		781
738	Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2024b. Visual instruction tuning. <i>Advances in neural information processing systems</i> , 36.	782
739		783
740		784
741	Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In <i>The 36th Conference on Neural Information Processing Systems (NeurIPS)</i> .	785
742		786
743		787
744		
745		
746		
	Pan Lu, Liang Qiu, Jiaqi Chen, Tony Xia, Yizhou Zhao, Wei Zhang, Zhou Yu, Xiaodan Liang, and Song-Chun Zhu. 2021. Iconqa: A new benchmark for abstract diagram understanding and visual language reasoning. In <i>The 35th Conference on Neural Information Processing Systems (NeurIPS) Track on Datasets and Benchmarks</i> .	788
		789
		790
		791
		792
	Michael S Matena and Colin A Raffel. 2022. Merging models with fisher-weighted averaging. <i>Advances in Neural Information Processing Systems</i> , 35:17703–17716.	793
		794
	Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Agüera y Arcas. 2017. Communication-efficient learning of deep networks from decentralized data. In <i>Artificial intelligence and statistics</i> , pages 1273–1282. PMLR.	795
		796
	Baolin Peng, Chunyuan Li, Pengcheng He, Michel Galley, and Jianfeng Gao. 2023a. Instruction tuning with gpt-4. <i>arXiv preprint arXiv:2304.03277</i> .	797
		798
	Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. 2023b. Kosmos-2: Grounding multimodal large language models to the world. <i>arXiv preprint arXiv:2306.14824</i> .	799
		800
	Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. 2023. Llama: Open and efficient foundation language models. <i>arXiv preprint arXiv:2302.13971</i> .	
	Ziyao Wang, Zheyu Shen, Yexiao He, Guoheng Sun, Hongyi Wang, Lingjuan Lyu, and Ang Li. 2024. Flora: Federated fine-tuning large language models with heterogeneous low-rank adaptations. <i>arXiv preprint arXiv:2409.05976</i> .	
	Chengyue Wu, Teng Wang, Yixiao Ge, Zeyu Lu, Ruisong Zhou, Ying Shan, and Ping Luo. 2023. $\pi$ -tuning: transferring multimodal foundation models with optimal multi-task interpolation. In <i>Proceedings of the 40th International Conference on Machine Learning</i> , pages 37713–37727.	
	Binqian Xu, Xiangbo Shu, Haiyang Mei, Guosen Xie, Basura Fernando, Mike Zheng Shou, and Jinhui Tang. 2024. Fedmllm: Federated fine-tuning mllm on multimodal heterogeneity data. <i>arXiv preprint arXiv:2411.14717</i> .	
	Yiyuan Yang, Guodong Long, Tao Shen, Jing Jiang, and Michael Blumenstein. 2024. Dual-personalizing adapter for federated foundation models. <i>arXiv preprint arXiv:2403.19211</i> .	
	Liping Yi, Han Yu, Gang Wang, and Xiaoguang Liu. 2023. Fedlora: Model-heterogeneous personalized federated learning with lora tuning. <i>arXiv preprint arXiv:2310.13283</i> .	

- Shukang Yin, Chaoyou Fu, Sirui Zhao, Ke Li, Xing Sun, Tong Xu, and Enhong Chen. 2024. A survey on multimodal large language models. *National Science Review*, page nwae403.
- Qiyang Yu, Yang Liu, Yimu Wang, Ke Xu, and Jingjing Liu. 2023. [Multimodal federated learning via contrastive representation ensemble](#). *Preprint*, arXiv:2302.08888.
- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. 2021. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. *arXiv preprint arXiv:2106.10199*.
- Jiayin Zhang, Saeed Vahidian, Martin Kuo, Chunyuan Li, Ruiyi Zhang, Tong Yu, Guoyin Wang, and Yiran Chen. 2024. Towards building the federatedgpt: Federated instruction tuning. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6915–6919. IEEE.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. 2023. Minigpt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*.

## A FedNano: Pseudocode Overview

---

**Algorithm 1** *FedNano*. The  $K$  clients are indexed by  $k$ ;  $R$  is the number of communication rounds, and  $T$  is the number of local steps.

---

**Server Update:**

- 1: Randomly initialize  $\mathcal{A}_I^0$  and  $\mathcal{A}_T^0$  in *NanoAdapter*, and distribute to clients
- 2: **for**  $r = 1$  to  $R$  **do**
- 3:   **for**  $k = 1$  to  $K$  **in parallel do**
- 4:      $\theta_k^r \leftarrow \text{ClientUpdate}(\theta_{global}^{r-1}, D_k)$
- 5:     Compute FIM  $F_k$
- 6:   **end for**
- 7:    $\theta_{global}^r \leftarrow \text{ServerAgg}(\{\theta_k^r, F_k^r\})$  ▷ Eq. 1
- 8: **end for**

**ClientUpdate** ( $\theta_{global}^{r-1}, D_k$ ):

- 1:  $\theta_k^{r-1} \leftarrow \theta_{global}^{r-1}$
  - 2: **for** local step  $t = 1$  to  $T$  **do**
  - 3:   Sample  $\{(v_k, q_k, a_k)\}$  from  $D_k$
  - 4:    $\theta_k^{r(t)} \leftarrow \text{Optimization}(\theta_k^{r(t-1)}, v_k, q_k, a_k)$
  - 5: **end for**
  - 6: **return**  $\theta_k^r$
-