

MLLM-LLaVA-FL: Multimodal Large Language Model Assisted Federated Learning

Jianyi Zhang¹, Hao Yang², Ang Li³, Xin Guo⁴, Pu Wang², Haiming Wang⁴,
Yiran Chen¹, Hai Li¹,

¹ Duke University, ² Johns Hopkins University,

³ University of Maryland College Park ⁴ Lenovo Research

Abstract

Previous studies on federated learning (FL) often encounter performance degradation due to data heterogeneity among different clients. In light of the recent advances in multimodal large language models (MLLMs), such as GPT-4v and LLaVA, which demonstrate their exceptional proficiency in multimodal tasks, such as image captioning and multimodal question answering. We introduce a novel federated learning framework, named Multimodal Large Language Model Assisted Federated Learning (MLLM-LLaVA-FL), which employs powerful MLLMs at the server end to address the heterogeneous and long-tailed challenges. Owing to the advanced cross-modality representation capabilities and the extensive open-vocabulary prior knowledge of MLLMs, our framework is adept at harnessing the extensive, yet previously underexploited, open-source data accessible from websites and powerful server-side computational resources. Hence, the MLLM-LLaVA-FL not only enhances the performance but also avoids increasing the risk of privacy leakage and the computational burden on local devices, distinguishing it from prior methodologies. Our framework has three key stages. Initially, we conduct global visual-text pretraining of the model. This pretraining is facilitated by utilizing the extensive open-source data available online, with the assistance of MLLMs. Subsequently, the pretrained model is distributed among various clients for local training. Finally, once the locally trained models are transmitted back to the server, a global alignment is carried out under the supervision of MLLMs to further enhance the performance. Experimental evaluations on established benchmarks, show that our framework delivers promising performance in the typical scenarios with data heterogeneity and long-tail distribution across different clients in FL.

1. Introduction

Data heterogeneity represents a significant challenge in federated learning [14, 22, 37, 42, 60, 66]. It largely stems from the fact that the data across participating clients are distributed independently, with each client having a different sample distribution. Due to the diversity in clients' datasets, these datasets often exhibit a long-tailed distribution, leading to client models that are biased toward the more common classes [42, 51]. This discrepancy often results in a drop in model accuracy. Although several approaches have been proposed [9, 10, 17–19, 26, 45, 57], the majority fail to strike an optimal balance between performance and mitigating two critical issues: 1. avoiding concerns of privacy leakage, and 2. preventing the imposition of extra computational loads on local edge devices. For instance, some contemporary methodologies necessitate the transmission of both gradients and parameters from local models to the server, which introduces substantial privacy risks. This is because attackers could potentially reverse-engineer the transmitted data to reconstruct client-specific images, as highlighted in various studies [12, 13, 72]. Alternatively, other approaches require the deployment of sizable models on local devices, which increases the memory and computational demands.

In light of the current popularity and exceptional proficiency of multimodal large language models (MLLMs) in tasks involving multimodalities, such as image captioning and multimodal question answering, we introduce a three-stage framework, named Multimodal Large Language Model Assisted Federated Learning (MLLM-LLaVA-FL), which utilizes multimodal large language models (MLLMs) to the FL performance on heterogeneous and long-tailed data. The adaptation of MLLMs in FL is supported by two main considerations. Firstly, beyond the heterogeneous and long-tailed distribution of client datasets, there exists an abundance of open-sourced and legally available data on the internet that can be utilized for training. This implies the potential of employing MLLMs to annotate

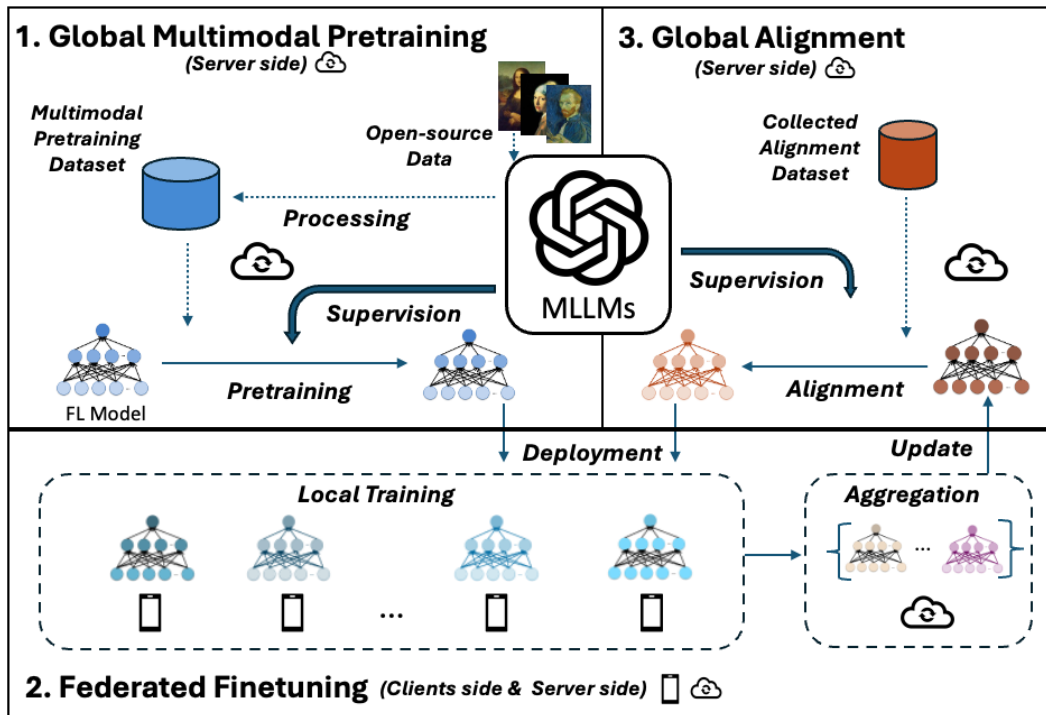


Figure 1. The workflow of MLLM-LLaVA-FL. The MLLM are utilized in the first stage Global Multimodal Pretraining and the third stage Global Alignment on the server side, to avoid extra computational load on devices.

unstructured, unlabeled online data, thereby augmenting FL performance. Secondly, in contrast to the limited computational resources available on client devices, the server-side capabilities are significantly more robust. This disparity opens up the possibility of deploying additional, more powerful MLLM on the server side to provide assistance to the FL system.

Our framework has three key stages. The initial stage, termed Global Multimodal Pretraining, first employs MLLMs to generate descriptions for unlabelled data collected from the internet. Then we develop a novel pretraining strategy, Dynamic Weighted Pretraining (DWP), which enables MLLMs to assist the compact FL models within the FL framework to conduct pertaining more efficiently on the open-sourced dataset. In the second stage, known as Federated Finetuning, we distribute the pretrained FL model to clients for local training on their datasets similar to with traditional FL approaches. This stage is highly flexible and compatible, allowing the integration of various previously designed FL methods. During the third stage, we perform Global Alignment on the server-side aggregated FL model under MLLM supervision. This process, similar to the idea of alignment in large language models, is aimed at further refining the model’s outputs to better align with task-specific requirements. Indeed, our framework is adaptable to a wide range of federated learning (FL) tasks. In this study, we specifically address the prevalent challenge of data heterogeneity in federated image classification tasks and the multimodal large language models we adopt here

are the large vision-language models (LVLM). During the pretraining stage, thanks to the extensive open-vocabulary prior knowledge embedded in large vision-language models, these models are capable of generating detailed descriptions for complex images found on the internet. This pretraining process of our FL model with the assistance of LVLMs on a large dataset of text-images, enables our FL model to develop enhanced image representation capabilities to better counteract the effects of data heterogeneity inherent in FL environments. Furthermore, the global alignment stage can also be designed to mitigate the issue of long-tailed distributions, which tend to bias client-side models towards more frequently occurring classes. Our contribution can be summarized as follows.

- Firstly, we pioneer the integration of the widely recognized multimodal large language model (MLLM) as an auxiliary tool in federated learning, aiming to enhance the utilization of previously underexplored internet data resources and server computational capabilities. Leveraging the formidable cross-modality representation capabilities and the vast open-vocabulary prior knowledge inherent in MLLMs, we introduce novel methodologies to address the challenges posed by long-tail distributions and data heterogeneity.
- In comparison to prevailing state-of-the-art approaches that address data heterogeneity in federated learning (FL), our methodology not only enhances privacy protection further but also significantly reduces the com-

Method	Multimodal Supervision	No Gradient Upload for Privacy	No Additional Computing Burden on Devices	Compatibility
CREFF [44], CLIP2FL [45]	✓		✓	
MLLM-LLaVA-FL (ours)	✓	✓	✓	✓

Table 1. Comparison between our methods and other status quo approaches for addressing long-tailed distribution challenges in FL

putational burden on client devices.

- Our extensive experimental results show that MLLM-LLaVA-FL can effectively handle heterogeneity and class-distribution imbalance, consistently surpassing the performance of existing state-of-the-art federated learning methodologies across a variety of datasets.

2. Related Work

2.1. Multimodal large language model

The introduction of GPT-4(Vision) [39] and Gemini [48] have demonstrated remarkable abilities in Multimodal understanding and generation, sparking a research fervor on Multimodal large language model. This enthusiasm extends to a variety of tasks, including image-text comprehension [27, 31, 71]; video-text understanding [28, 35]; and audio-text understanding [6]. Among them, recent studies in image-text comprehension with large vision-language models (LVLM) [27, 31, 71] have catalyzed notable advancements in harnessing the robust capabilities of large language models to tackle multimodal tasks effectively, such as crafting narratives from images and executing intricate reasoning tasks. Prominent instances include Visual ChatGPT [54], which amalgamates diverse visual foundational models for intricate visual tasks and instructions, employing iterative feedback to synchronize visual and textual modalities. In a similar way, MM-REACT [58] merges ChatGPT with visual models for multimodal undertakings, especially in the Visual Question Answering (VQA) framework. BLIP-2 [27], notable for its Q-former model, has shown encouraging outcomes in VQA tasks, both in zero-shot and fine-tuning settings. LLaMA-Adapter [11] enhances multimodal fine-tuning efficiency by integrating adaptation prompt vectors as adjustable parameters, showcasing versatility in multimodal contexts. MiniGPT-4 [71], derived from GPT-4 and incorporating elements from BLIP-2 and Vicuna [5], specializes in caption generation and model refinement through image-text pair fine-tuning. LLaVA [31], leveraging GPT-4, focuses on a broad spectrum of instruction fine-tuning data, ranging from multi-turn QA to image descriptions, adopting a dual-stage fine-tuning approach that prioritizes language model loss while keeping the visual model static.

2.2. Federated Learning with Heterogeneous Data

Current methodologies tackling the challenge of data heterogeneity fall into the following broad categories. Some approaches aim to simultaneously improve the models on both clients and server sides through optimization techniques. Key contributions in this area have been made by the work of [10, 17, 18, 26], who have investigated various optimization methods. Other strategies focus on enhancing the stability of local models via knowledge transfer, a technique that is model-agnostic and has been explored in the research by [4, 55, 61], aiming to mitigate data heterogeneity by spreading local training knowledge throughout the whole FL framework. Additional methods, such as those proposed by [2, 30], concentrate on improving model aggregation on the server side to address data heterogeneity. Certain strategies also regulate the scheduling of client participation to avoid biasing the FL model towards classes that are more prevalent, as explored in the studies by [56, 63]. While these approaches have advanced the handling of data heterogeneity, they often do not fully address the specific issues related to long-tailed distributions in FL.

A recent approach, CREFF [44], introduces a decoupling strategy to create balanced class-distribution federated features for the server model and to retrain the classifier with these features. Nonetheless, CREFF encounters two main limitations due to its reliance on generating federated features through client-side gradient information: 1) The one-to-many relationship between gradients and samples can result in the problem becoming ill-posed; 2) The absence of semantic guidance might lead to federated features that lack discriminative ability for their respective classes. The subsequent attempt, CLIP2FL [45], seeks to overcome these drawbacks by integrating a multimodal model to direct the federated learning process. However, it still has its own drawbacks. Firstly, deploying the sizable CLIP model on devices increases memory and computational demands. Secondly, transmitting both the gradient and parameters of local models to the server, as necessitated by both CLIP4FL and CREFF, raises significant privacy concerns, as attackers could potentially reconstruct client images through reverse engineering [12, 13, 72].

3. Methodology

In the conventional federated learning (FL) pipeline, the FL models are typically assigned to local clients for train-

ing on their heterogeneous datasets. These models are then sent back to the server for aggregation. This cycle continues repeatedly until the FL training concludes. To better align the above FL framework with practical requirements, we incorporate the following two additional stages. The first stage occurs before local training, involving pretraining on the server side, a strategy supported by previous work [3,49] which found that pretraining can accelerate the convergence of FL training and mitigate the effects of data heterogeneity on convergence. The second stage takes place after aggregation, where the FL model may undergo further training to meet the broader requirements of FL companies, such as the performance and safety considerations we discuss later.

Drawing inspiration from recent works in learning paradigms of NLP [8,21,25,40,49,52,59,62,65], we structure our work into three parts: global multimodal pretraining, federated local finetuning, and global alignment. We deploy multimodal approaches at the server side to assist both global multimodal pretraining and global alignment phases. In this section, we will introduce our comprehensive framework as follows: Section 3.1 will delve into the details of our global multimodal pretraining; Section 3.2 will cover federated local finetuning; Section 3.3 will discuss global alignment; and in Section 3.4, we will compare our method with previous approaches.

3.1. Global Multimodal Pretraining

Pretraining Dataset As discussed in Sec. 1, the wealth of open-source multimodal data, such as images and their captions, remains underutilized resources for pretraining FL models. Often, these datasets are noisy, unlabeled or contain elements that are too complex, making them unsuitable for straightforward pretraining of the compact FL models. However, given the current advanced capabilities in multimodal processing of MLLMs, we now have new, convenient methods to leverage such data for pretraining purposes. Utilizing GPT-4, akin to the approach used in LLaVA, we can transform complex image data collected from the internet into three main categories:

- *Conversation*: This category includes dialogues between an assistant and an individual seeking specific information about a photo. The assistant’s responses simulate observing the image directly, answering a variety of questions about the visual content, such as identifying objects, counting them, describing actions, pinpointing locations, and noting their spatial relations.
- *Detailed Description*: To gain a thorough understanding of an image, we formulated a series of questions designed to elicit detailed descriptions. Responses to these questions were generated using GPT-4, enriching our dataset with nuanced insights into the images.

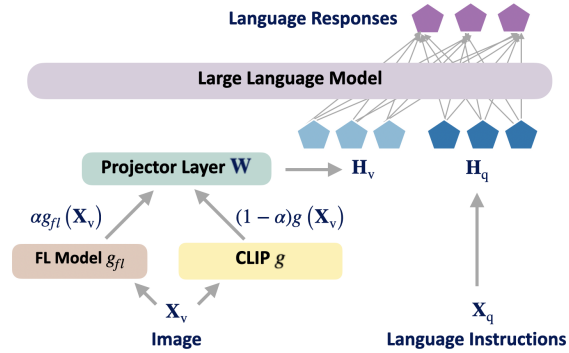


Figure 2. The visualization of our pretraining mechanism

- *Complex Reasoning*: This category focuses on more sophisticated reasoning questions based on the content of the images. Answering these questions involves a detailed logical breakdown, reflecting a deep comprehension of the images and the ability to reason through them.

Leveraging the aforementioned dataset formulations, we are equipped to facilitate the pretraining of FL models with the support of Multimodal Large Language Models.

Pretraining Mechanism Our pretraining mechanism draws inspiration from the structure of LLaVA, a highly effective and recent multimodal large language model. It consists of three key components: a visual encoder g , which is a frozen pretrained CLIP model; a projection layer designed to align the features of the visual model with the text domain embeddings, where the projection layer is a trainable matrix W ; and a part comprising a large language model (LLM), typically employing promising models such as Vicuna or LLaMA-2. The workflow of LLaVA proceeds as follows: For the data formats mentioned earlier, whether it be a conversation, detailed description, or complex reasoning, the input includes a text modality instruction X_q (e.g., "Could you provide a detailed description of this image?") and an image X_v . The instruction X_q passes through an embedding layer to obtain the text embedding H_q . As for the image X_v , it first goes through the visual encoder g to acquire the grid feature Z_v , which then passes through the projection layer to obtain the visual embedding H_v , aligned in dimension with H_q :

$$H_v = W \cdot Z_v, \text{ where } Z_v = g(X_v) \quad (1)$$

The text embedding H_q and the visual embedding H_v are concatenated and subsequently input into the LLM. The resulting output is the MLLM’s response to the given inputs. Throughout the LLaVA training process, both the projector and the LLM are trainable, whereas the visual encoder and the CLIP model remain fixed and are not subject to training.

Drawing from the LLaVA mechanism, our Global Multimodal Pretraining essentially integrates the compact FL model, designed for downstream image classification tasks, as a component within the LLaVA framework’s visual encoder. The FL model is made trainable and is denoted as g_{fl} . Inspired by the previous work in knowledge distillation [15, 20, 38, 64] and Bayesian Algorithms [1, 32, 53, 67–70], We have developed an approach termed Dynamic Weighted Distillation, which involves computing a weighted average of the visual features obtained from the FL model and those from the original visual encoder:

$$\mathbf{Z}_v = (1 - \alpha)g(\mathbf{X}_v) + \alpha g_{fl}(\mathbf{X}_v) \quad (2)$$

In this equation, \mathbf{Z}_v represents the weighted combined visual features, $g(\mathbf{X}_v)$ indicates the visual features from the original visual encoder, $g_{fl}(\mathbf{X}_v)$ refers to the visual features from the FL model, and α is the dynamic weighting factor that adjusts the influence of each feature. During the pretraining phase, the CLIP model g , the projector W , and the LLM are kept static, with only g_{fl} , the FL model component, being trainable. Initially, α is set to 0, and as pretraining progresses, it gradually increases to 1, where it remains for the duration of the pretraining. This approach is termed Dynamic Weighted Pretraining. The rationale behind this strategy stems from the typically smaller size of the FL model compared to the original CLIP visual encoder. This size discrepancy is due to the constraints imposed by subsequent local training on edge devices within the FL framework. Directly substituting the large CLIP model with a more compact FL model could significantly hamper the process of multimodal alignment, owing to the vast difference in capacity between the two visual models resulting from their size difference.

3.2. Federated Finetuning

In this subsection, we delve into the Federated Finetuning phase. Upon obtaining a pretrained FL model g_{fl} from the initial stage, we append classifier layers to it, tailoring the model for image classification tasks on local datasets on the client side. If we denote the set of all model parameters for the k -th client at the t -th local step as w_k^t , and the local data as \mathcal{D}^k , then the k -th client updates the received model in a manner similar to FedAvg:

$$w_k^{t+1} \leftarrow w_k^t - \eta \nabla_w L_{loc}(w_k^t; \mathcal{D}^k) \quad (3)$$

, where the L_{loc} represents the local loss function.

After local training, the model parameters w_k are sent back to the server for global aggregation, where we also utilize FedAvg:

$$w_{agg}^{t+1} = \sum_{k \in \Omega^t} \frac{|\mathcal{D}^k|}{\sum_{k \in \Omega^t} |\mathcal{D}^k|} w_k^{t+1} \quad (4)$$

, where the Ω^t is the set of clients selected at the t -th round.

This stage mirrors the traditional FL framework closely. Drawing inspiration from learning paradigms in NLP, we refer to this phase as Federated Finetuning, in light of the pretraining conducted in the preceding step. It’s important to note the flexibility and compatibility of our framework; we can substitute FedAvg with any other existing FL methods designed to enhance local training and global aggregation, aiming to boost final utility performance metrics like accuracy, or system performance aspects like speed or computational efficiency. Furthermore, this phase does not cause additional privacy concerns, and existing methods for privacy or safety protection can be seamlessly integrated.

3.3. Global Alignment

Recent studies on the alignment of large language models, such as Reinforcement Learning from Human Feedback (RLHF), are focused on refining the model’s outputs to more closely resonate with human-like understanding and reasoning. This enhancement significantly improves the model’s capability in tasks that require the interpretation and execution of complex instructions. In a similar vein, companies engaged in federated learning (FL) have analogous requirements for models after global aggregation. For instance, concerning safety requirements, an FL company must ensure that models trained via federated learning do not leak user information. Additionally, there are performance-related requirements, such as adjusting the model to prevent biases caused by long-tailed distributions or training the model on new datasets to acquire new skills. Typically, this involves constructing an alignment dataset \mathcal{D}_{align} and selecting a suitable global alignment function L_{align} :

$$w_{agg}^{new} \leftarrow w_{agg} - \eta \nabla_w L_{align}(w_{agg}; \mathcal{D}_{align}) \quad (5)$$

To address the issue of long-tailed distributions, one could design \mathcal{D}_{align} as a small, class-balanced dataset encompassing all categories, with L_{align} defined as follows:

$$L_{align} = L_{ce}(y, p) + \beta \cdot KL(q||p), \quad (6)$$

where $L_{ce}(\cdot, \cdot)$ is the cross-entropy loss. y is the label and p is the output logits vector of the FL models. Since the pretrained CLIP model in the pertaining mechanism has zero-shot image classification capability [41], q denotes the output logits vector of the CLIP model. KL is the Kullback-Leibler divergence and β is a hyperparameter balanced these two losses.

The idea of using a class-balanced dataset to alleviate the challenges of long-tailed distributions aligns with the concepts of data resampling in centralized training to handle class imbalance and client selection in federated learning. The feasibility of such an alignment dataset \mathcal{D}_{align} existing on the server side is justifiable in most cases because, in

Type	Method	CIFAR-10-LT			CIFAR-100-LT		
		IF=100	IF=50	IF=10	IF=100	IF=50	IF=10
Heterogeneity-oriented FL methods	FedAvg	56.17	59.36	77.45	30.34	36.35	45.87
	FedAvgM	52.03	57.11	70.81	30.80	35.33	44.66
	FedProx	56.92	60.89	76.53	31.67	36.30	46.10
	FedDF	55.15	58.74	76.51	31.43	36.22	46.19
	FedBE	55.79	59.55	77.78	31.97	36.39	46.25
	CCVR	69.53	71.89	78.48	33.43	36.98	46.88
	FedNova	57.79	63.91	77.79	32.64	36.62	46.75
Imbalance-oriented FL methods	Fed-Focal Loss	53.83	57.42	73.74	30.67	35.25	45.52
	Ratio Loss	59.75	64.77	78.14	32.95	36.88	46.79
	FedAvg+ τ -norm	49.95	51.41	72.08	26.22	33.71	43.65
Classifier-retraining	CReFF	70.55	73.08	80.71	34.67	37.64	47.08
SOTA	CLIP2FL	73.37	75.35	81.18	37.56	41.29	48.20
Our framework	MLLM-LLaVA-FL	75.49 ($\uparrow 2.12$)	76.11 ($\uparrow 1.24$)	81.45 ($\uparrow 0.27$)	39.50 ($\uparrow 1.94$)	42.34 ($\uparrow 1.05$)	48.87 ($\uparrow 0.67$)

Table 2. Top-1 classification accuracy(%) on CIFAR-10-LT and CIFAR-100-LT datasets with different FL methods, where th results are referred in [44, 45]. The best results are marked in bold.

practice, for global aggregation, the server typically pre-defines the categories for model classification and organizes them, which is essential for subsequent federated learning processes. Otherwise, the global aggregation of classifier layers would become chaotic. Knowing the categories, companies could feasibly collect data from the internet or generate data using powerful image-generation models like Stable Diffusion or Midjourney. However, we acknowledge that in some extreme cases, data collection can be challenging, necessitating the design of more specific L_{align} and \mathcal{D}_{align} , which we leave for future work.

4. Experiment

4.1. Experiment Setup

Dataset&Implementation We applied our MLLM-LLaVA-FL framework to three widely-used long-tailed datasets: CIFAR-10/100LT [24] and ImageNet-LT [33]. As for the first two datasets, we adopt the same sampling technique as previous studies [7] to create long-tailed distributions with various imbalance factors (IF = 100, 50, 10), and we follow CReFF [44] to use Dirichlet distribution with the key parameter α to generate the heterogeneous data partition among clients, where the value of α is set to 0.5 on CIFAR-10/100-LT. ImageNet-LT has 115.8 K images from 1000 classes and the number of images per class ranging from 1280 to 5, where the value of α is set to 0.1. We utilized ResNet-8 as the feature extractor for CIFAR-10/100-LT and ResNet-50 for ImageNet-LT, adding an MLP layer to each to align their feature dimensions with CLIP’s outputs. The number of clients is set to 20, and we select 40% at random for each training round. The client-side training batch size was uniform at 32 across Cifar-10/100 and imagenet. All the above settings are the same as the previous work in [45]. We employed the standard cross-entropy loss by default and executed 200

communication rounds. For the pertaining part, we adopt the pertaining dataset of LLaVA, CC-595K, and train the model for 4 epochs with a learning rate of $2e-3$ and a batch size of 128. During the first 2 epochs, the α in our pretraining mechanism increase from 0 to 1 following the cosine scheduler and then the value remains 1 for the following epochs. All the experiments were conducted using PyTorch on a single Nvidia A100 80G GPU.

Baselines We compare MLLM4FL with 13 FL methods: FedAvg [36], FedAvgM [16], FedProx [29], FedDF [46], FedBE [2], CCVR [34] and FedNova [50], Fed-Focal Loss [43], Ratio Loss [51] and FedAvg with τ -norm [23], CReFF [44] and CLIP2FL [45]. The first seven approaches are heterogeneity-oriented, and Fed-Focal Loss, Ratio Loss and FedAvg with τ -norm are imbalance-oriented.

4.2. Experimental Results

Results for CIFAR-10/100-LT are presented in Tab. 2, where we evaluate the performance of our CLIP2FL against a range of FL approaches on both CIFAR-10-LT and CIFAR-100-LT datasets. Notably, MLLM-LLaVA-FL outperforms other methods in terms of classification accuracy on both datasets. Specifically, at an Imbalance Factor (IF) of 100, which presents a severe imbalance, MLLM shows an improvement of 2.12% and 1.94% in classification accuracy over CLIP2FL for CIFAR-10-LT and CIFAR-100-LT, respectively. Under the condition of IF = 50 or 10, MLLM still manages to enhance performance by around 1%. This underscores MLLM-LLaVA-FL’s effectiveness and its outperforms over competing methods to deal with heterogeneity and long-tailed distributions.

In the context of ImageNet-LT, Tab. 3 presents a comparison of the accuracy achieved by our MLLM-LLaVA-FL framework against various FL approaches. The evaluation is segmented into four groups based on the number of sam-

Type	Method	ImageNet-LT			
		All	Many	Medium	Few
Heterogeneity-oriented FL methods	FedAvg	23.85	34.92	19.18	7.10
	FedAvgM	22.57	33.93	18.55	6.73
	FedProx	22.99	34.25	17.06	6.37
	FedDF	21.63	31.78	15.52	4.48
	CCVR	25.49	36.72	20.24	9.26
Imbalance-oriented FL methods	Fed-Focal Loss	21.60	31.74	15.77	5.52
	Ratio Loss	24.31	36.33	18.14	7.41
	FedAvg+ τ -norm	21.58	31.66	15.76	4.33
Classifier-retraining	CRFF	26.31	37.44	21.87	10.29
Our framework	MLLM-LLaVA-FL	27.53 (\uparrow 1.22)	30.85 (\downarrow 6.59)	25.89 (\uparrow 4.02)	25.58 (\uparrow 15.29)

Table 3. Top-1 accuracy(%) on ImageNet-LT dataset with different FL method

ples per class: “Many” (over 100 samples), “Medium” (20 to 100 samples), “Few” (less than 20 samples), and “All” (overall accuracy). While our method may not fully match the performance of CRFF in the “Many” categories, it excels in “Overall” accuracy and the “Medium” and “Few” category. These results underscore MLLM-LLaVA-FL’s capability not just in enhancing overall model performance but also in significantly improving classification outcomes for categories with fewer samples. The ImageNet-LT results highlight the effectiveness of MLLM-LLaVA-FL in tackling the inherent challenges of long-tailed data distributions.

4.3. Further Analysis

We conduct some further analysis to verify the effectiveness of our framework, especially the importance of global pertaining and global alignment.

4.3.1 Ablation studies on our pretraining mechanism

To evaluate the effectiveness of pretraining, we conducted a comparative analysis between a pretrained model and a non-pretrained model under a constrained training dataset scenario akin to few-shot learning, aiming to mirror real-world conditions where the amount of data available on each device is limited. We generated subsets of the CIFAR-10 and CIFAR-100 datasets at varying proportions and trained both models for 30 epochs. Our findings, detailed in Tab. 4, outline the number of epochs required by each model to reach predetermined accuracy thresholds with different training sample sizes, along with the highest accuracy achieved by each model within the 30-epoch span.

Specifically, within the CIFAR-10 context, to attain a target accuracy of 25%, the pretrained model consistently outperformed the non-pretrained model, requiring fewer epochs across all sample sizes. Similarly, for the CIFAR-

100 dataset, in pursuit of a target accuracy of 15%, the pretrained model proved more efficient, also necessitating fewer epochs. Fig. 3 further illustrates the superiority of our pretraining approach, demonstrating that models pretrained using our methodology surpass those trained from scratch in terms of accuracy.

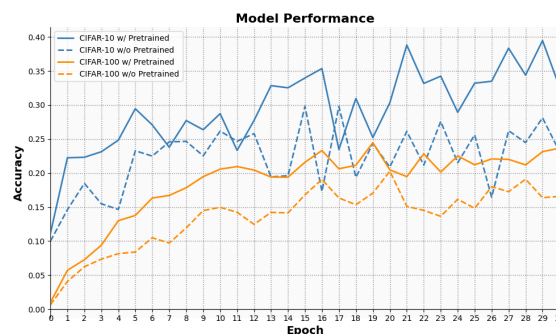


Figure 3. The comparative analysis of pretrained and non-pretrained models using 1% subsets of CIFAR-10/100 training data.

4.3.2 Ablation studies on our global alignment mechanism

In Fig. 4, we present the confusion matrix for our model equipped with global alignment within the CIFAR-10-LT dataset, characterized by an imbalance factor of 100. We normalize this confusion matrix by the volume of data in each class. Additionally, we illustrate a normalized confusion matrix for a baseline model devoid of global alignment within our Federated Learning (FL) framework. The visual representations indicate that, with alignment in place, data from each class can be accurately classified. In contrast, the absence of alignment yields inferior results, particularly for classes with few data. Compared to classes with abundant

Dataset	Initialization	0.4%	1%	2%
Cifar10	w/ Pretraining	(3, 37.62)	(5, 39.46)	(3, 37.83)
	w/o Pretraining	(6, 30.44)	(10, 29.8)	(14, 31.99)
Cifar100	w/ Pretraining	(5, 23.91)	(6, 24.46)	(6, 24.28)
	w/o Pretraining	(10, 18.61)	(15, 20.27)	(11, 19.42)

Table 4. Comparison between pretrained and non-pretrained models under constrained training dataset settings. The format is (number of epochs, highest accuracy)

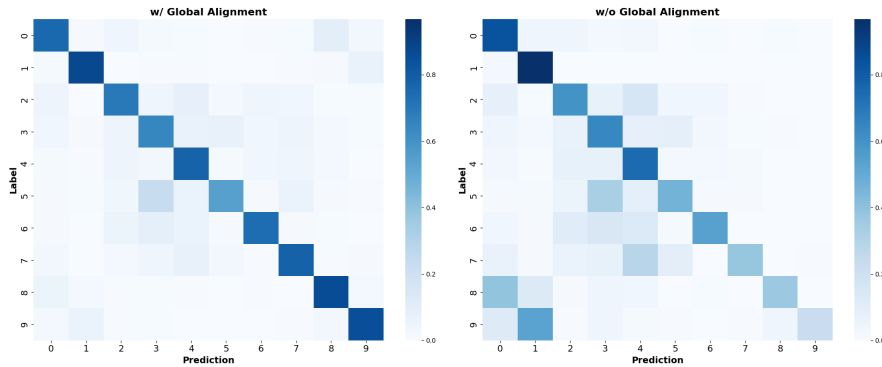


Figure 4. The comparative analysis of aligned and non-aligned models with normalized confusion matrices.

data, those with fewer instances often experience misclassification, with minority class data being inaccurately labeled as belonging to majority classes. This starkly underscores the significance of our global alignment strategy in enhancing both the performance and fairness of the FL system.

4.3.3 Discussion

Privacy: Our strategy sidesteps the conventional requirement for clients to send gradients back to the server, as seen in methods like CReFF [44] and CLIP2FL [45]. This aspect is vital because the transmission of gradients could enable the server to perform reverse engineering attacks [12, 13, 47, 72], potentially endangering client data confidentiality. By eliminating this step, our method diminishes the likelihood of leaking sensitive client information, promoting a more secure and privacy-centric learning environment.

Computational Efficiency: Our approach also stands out for its computational economy. Contrary to approaches like CLIP2FL, which necessitate deploying sizable multimodal models such as CLIP on client devices—demanding significant memory and potentially being unfeasible for edge devices with limited resources—our method positions the MLLM solely on the server side. At the client level, we deploy only the compact FL models. This resolution not only addresses memory constraints but also reduces the time and energy expenditure associated with federated local training. Consequently, our framework is rendered more practical and appealing for an extensive array of devices, particularly those with restricted storage capacities.

Compatibility: Our approach stands out for its adaptability, unlike specific methodologies like CReFF and CLIP2FL that impose unique requirements on federated local training and global aggregation. Our framework can be compatible with a wide array of existing FL algorithms. This includes but is not limited to, client selection strategies and various techniques aimed at further enhancing client privacy protection. For a more comprehensive comparison, please refer to the Tab. 1 highlighting the advantages of our method.

5. Conclusion

To overcome the challenges of federated learning in the context of heterogeneous and long-tailed data distributions, we introduced a novel framework, MLLM-LLaVA-FL. This framework is structured around three core stages: global pretraining, federated fine-tuning, and global alignment. This marks the inaugural integration of MLLMs into an FL system. Leveraging the strong multimodal capacities of MLLMs, our approach taps into the vast yet previously underutilized reservoir of open-source data available online, alongside substantial server-side computational resources. Crucially, our methodology does not compromise privacy nor impose additional computational demands on client devices. Experimental evidence verifies the efficacy of our framework, paving the way for future research on more multimodal tasks.

Acknowledgement The research performed at Duke University was supported in part by NSF 2112562 and ARO W911NF-23-2-0224.

References

- [1] Changyou Chen, Jianyi Zhang, Yi Xu, Liqun Chen, Jiali Duan, Yiran Chen, Son Tran, Belinda Zeng, and Trishul Chilimbi. Why do we need large batchsizes in contrastive learning? a gradient-bias perspective. *Advances in Neural Information Processing Systems*, 35:33860–33875, 2022. 5
- [2] Hong-You Chen and Wei-Lun Chao. Fedbe: Making bayesian model ensemble applicable to federated learning. *arXiv preprint arXiv:2009.01974*, 2020. 3, 6
- [3] Hong-You Chen, Cheng-Hao Tu, Ziwei Li, Han-Wei Shen, and Wei-Lun Chao. On pre-training for federated learning. *arXiv preprint arXiv:2206.11488*, 2022. 4
- [4] Yiqiang Chen, Xin Qin, Jindong Wang, Chaohui Yu, and Wen Gao. Fedhealth: A federated transfer learning framework for wearable healthcare. *IEEE Intelligent Systems*, 35(4):83–93, 2020. 3
- [5] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E. Gonzalez, Ion Stoica, and Eric P. Xing. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, March 2023. 3
- [6] Yunfei Chu, Jin Xu, Xiaohuan Zhou, Qian Yang, Shiliang Zhang, Zhijie Yan, Chang Zhou, and Jingren Zhou. Qwen-audio: Advancing universal audio understanding via unified large-scale audio-language models. *arXiv preprint arXiv:2311.07919*, 2023. 3
- [7] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 9268–9277, 2019. 6
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 4
- [9] Zhixu Du, Jingwei Sun, Ang Li, Pin-Yu Chen, Jianyi Zhang, Hai” Helen” Li, and Yiran Chen. Rethinking normalization methods in federated learning. In *Proceedings of the 3rd International Workshop on Distributed Machine Learning*, pages 16–22, 2022. 1
- [10] Chun-Mei Feng, Kai Yu, Nian Liu, Xinxing Xu, Salman Khan, and Wangmeng Zuo. Towards instance-adaptive inference for federated learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23287–23296, 2023. 1, 3
- [11] Peng Gao, Jiaming Han, Renrui Zhang, Ziyi Lin, Shijie Geng, Aojun Zhou, Wei Zhang, Pan Lu, Conghui He, Xiangyu Yue, et al. Llama-adapter v2: Parameter-efficient visual instruction model. *arXiv preprint arXiv:2304.15010*, 2023. 3
- [12] Jonas Geiping, Hartmut Bauermeister, Hannah Dröge, and Michael Moeller. Inverting gradients-how easy is it to break privacy in federated learning? *Advances in Neural Information Processing Systems*, 33:16937–16947, 2020. 1, 3, 8
- [13] Niv Haim, Gal Vardi, Gilad Yehudai, Ohad Shamir, and Michal Irani. Reconstructing training data from trained neural networks. *Advances in Neural Information Processing Systems*, 35:22911–22924, 2022. 1, 3, 8
- [14] Weituo Hao, Mostafa El-Khamy, Jungwon Lee, Jianyi Zhang, Kevin J Liang, Changyou Chen, and Lawrence Carin Duke. Towards fair federated learning with zero-shot data augmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3310–3319, 2021. 1
- [15] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015. 5
- [16] Tzu-Ming Harry Hsu, Hang Qi, and Matthew Brown. Measuring the effects of non-identical data distribution for federated visual classification. *arXiv preprint arXiv:1909.06335*, 2019. 6
- [17] Hua Huang, Fanhua Shang, Yuanyuan Liu, and Hongying Liu. Behavior mimics distribution: Combining individual and group behaviors for federated learning. *arXiv preprint arXiv:2106.12300*, 2021. 1, 3
- [18] Wenke Huang, Mang Ye, Zekun Shi, He Li, and Bo Du. Rethinking federated learning with domain shift: A prototype view. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16312–16322. IEEE, 2023. 1, 3
- [19] Yuqi Jia, Saeed Vahidian, Jingwei Sun, Jianyi Zhang, Vyacheslav Kungurtsev, Neil Zhenqiang Gong, and Yiran Chen. Unlocking the potential of federated learning: The symphony of dataset distillation via deep generative latents. In *European Conference on Computer Vision*, pages 18–33. Springer, 2025. 1
- [20] Xiao Jin, Baoyun Peng, Yichao Wu, Yu Liu, Jiaheng Liu, Ding Liang, Junjie Yan, and Xiaolin Hu. Knowledge distillation via route constrained optimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1345–1354, 2019. 5
- [21] Hailey Joren, Jianyi Zhang, Chun-Sung Ferng, Da-Cheng Juan, Ankur Taly, and Cyrus Rashtchian. Sufficient context: A new lens on retrieval augmented generation systems. *arXiv preprint arXiv:2411.06037*, 2024. 4
- [22] Peter Kairouz, H Brendan McMahan, Brendan Avent, Aurélien Bellet, Mehdi Bennis, Arjun Nitin Bhagoji, Kallista Bonawitz, Zachary Charles, Graham Cormode, Rachel Cummings, et al. Advances and open problems in federated learning. *Foundations and Trends® in Machine Learning*, 14(1–2):1–210, 2021. 1
- [23] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling representation and classifier for long-tailed recognition. *arXiv preprint arXiv:1910.09217*, 2019. 6
- [24] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). 6
- [25] Martin Kuo, Jianyi Zhang, and Yiran Chen. Dacbert: Leveraging dependency agreement for cost-efficient bert pretraining. *arXiv preprint arXiv:2311.04799*, 2023. 4
- [26] Bo Li, Mikkel N Schmidt, Tommy S Alstrøm, and Sebastian U Stich. On the effectiveness of partial variance reduction in federated learning with heterogeneous data. *arXiv preprint arXiv:2212.02191*, 2022. 1, 3

- [27] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023. 3
- [28] KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. Videochat: Chat-centric video understanding. *arXiv preprint arXiv:2305.06355*, 2023. 3
- [29] Tian Li, Anit Kumar Sahu, Manzil Zaheer, Maziar Sanjabi, Ameet Talwalkar, and Virginia Smith. Federated optimization in heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 2018. 6
- [30] Tao Lin, Lingjing Kong, Sebastian U Stich, and Martin Jaggi. Ensemble distillation for robust model fusion in federated learning. *Advances in Neural Information Processing Systems*, 33:2351–2363, 2020. 3
- [31] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning, 2023. 3
- [32] Qiang Liu and Dilin Wang. Stein variational gradient descent: A general purpose bayesian inference algorithm, 2019. 5
- [33] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2537–2546, 2019. 6
- [34] Mi Luo, Fei Chen, Dapeng Hu, Yifan Zhang, Jian Liang, and Jiashi Feng. No fear of heterogeneity: Classifier calibration for federated learning with non-iid data. *Advances in Neural Information Processing Systems*, 34:5972–5984, 2021. 6
- [35] Muhammad Maaz, Hanoona Rasheed, Salman Khan, and Fahad Shahbaz Khan. Video-chatgpt: Towards detailed video understanding via large vision and language models. *arXiv preprint arXiv:2306.05424*, 2023. 3
- [36] Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, and Blaise Aguera y Arcas. Communication-efficient learning of deep networks from decentralized data. In *Artificial intelligence and statistics*, pages 1273–1282. PMLR, 2017. 6
- [37] H Brendan McMahan, Eider Moore, Daniel Ramage, Seth Hampson, et al. Communication-efficient Learning of Deep Networks from Decentralized Data. *Artificial Intelligence and Statistics*, 2017. 1
- [38] Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. Improved knowledge distillation via teacher assistant. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 5191–5198, 2020. 5
- [39] OpenAI. Introducing chatgpt. <https://openai.com/blog/chatgpt/>, November 2022. 3
- [40] OpenAI. GPT-4 Technical Report. *arXiv e-prints*, page arXiv:2303.08774, Mar. 2023. 4
- [41] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. 5
- [42] Anit Kumar Sahu, Tian Li, Maziar Sanjabi, Manzil Zaheer, Ameet Talwalkar, and Virginia Smith. Federated optimization for heterogeneous networks. *arXiv preprint arXiv:1812.06127*, 1(2):3, 2018. 1
- [43] Dipankar Sarkar, Ankur Narang, and Sumit Rai. Fed-focal loss for imbalanced data classification in federated learning. *arXiv preprint arXiv:2011.06283*, 2020. 6
- [44] Xinyi Shang, Yang Lu, Gang Huang, and Hanzi Wang. Federated learning on heterogeneous and long-tailed data via classifier re-training with federated features. *arXiv preprint arXiv:2204.13399*, 2022. 3, 6, 8
- [45] Jiangming Shi, Shanshan Zheng, Xiangbo Yin, Yang Lu, Yuan Xie, and Yanyun Qu. Clip-guided federated learning on heterogeneous and long-tailed data. *arXiv preprint arXiv:2312.08648*, 2023. 1, 3, 6, 8
- [46] Dianbo Sui, Yubo Chen, Jun Zhao, Yantao Jia, Yuantao Xie, and Weijian Sun. Feded: Federated learning via ensemble distillation for medical relation extraction. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing, EMNLP 2020, Online, November 16-20, 2020*, pages 2118–2128, 2020. 6
- [47] Minxue Tang, Jianyi Zhang, Mingyuan Ma, Louis Di-Valentin, Aolin Ding, Amin Hassanzadeh, Hai Li, and Yiran Chen. Fade: Enabling large-scale federated adversarial training on resource-constrained edge devices. *arXiv preprint arXiv:2209.03839*, 2022. 8
- [48] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 3
- [49] Yuanyishu Tian, Yao Wan, Lingjuan Lyu, Dezhong Yao, Hai Jin, and Lichao Sun. Fedbert: When federated learning meets pre-training. *ACM Trans. Intell. Syst. Technol.*, 13(4), aug 2022. 4
- [50] Jianyu Wang, Qinghua Liu, Hao Liang, Gauri Joshi, and H Vincent Poor. Tackling the objective inconsistency problem in heterogeneous federated optimization. *arXiv preprint arXiv:2007.07481*, 2020. 6
- [51] Lixu Wang, Shichao Xu, Xiao Wang, and Qi Zhu. Addressing class imbalance in federated learning. *arXiv preprint arXiv:2008.06217*, 2020. 1, 6
- [52] Qinsi Wang, Saeed Vahidian, Hancheng Ye, Jianyang Gu, Jianyi Zhang, and Yiran Chen. Coreinfer: Accelerating large language model inference with semantics-inspired adaptive sparse activation. *arXiv preprint arXiv:2410.18311*, 2024. 4
- [53] Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 681–688, 2011. 5
- [54] Chenfei Wu, Shengming Yin, Weizhen Qi, Xiaodong Wang, Zecheng Tang, and Nan Duan. Visual chatgpt: Talking, drawing and editing with visual foundation models. *arXiv preprint arXiv:2303.04671*, 2023. 3
- [55] Chun-Mei Feng, Bangjun Li, Xinxing Xu, Yong Liu, and Huazhu Fu. Wangmeng Zuo. Learning federated visual prompt in null space for mri reconstruction. *arXiv preprint arXiv:2303.16181*, 2023. 3

- [56] Miao Yang, Akitanoshou Wong, Hongbin Zhu, Haifeng Wang, and Hua Qian. Federated learning with class imbalance reduction, 2020. [3](#)
- [57] Qian Yang, Jianyi Zhang, Weituo Hao, Gregory P. Spell, and Lawrence Carin. Flop: Federated learning on medical datasets using partial networks. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining*, KDD '21, page 3845–3853, New York, NY, USA, 2021. Association for Computing Machinery. [1](#)
- [58] Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Ehsan Azarnasab, Faisal Ahmed, Zicheng Liu, Ce Liu, Michael Zeng, and Lijuan Wang. Mm-react: Prompting chatgpt for multimodal reasoning and action. *arXiv preprint arXiv:2303.11381*, 2023. [3](#)
- [59] Jianyi Zhang, Yiran Chen, and Jianshu Chen. Join-chain network: A logical reasoning view of the multi-head attention in transformer. In *2022 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 1–11. IEEE, 2022. [4](#)
- [60] Jianyi Zhang, Zhixu Du, Jingwei Sun, Ang Li, Minxue Tang, Yuhao Wu, Zhihui Gao, Martin Kuo, Hai-Helen Li, and Yiran Chen. Next generation federated learning for edge devices: An overview. In *2022 IEEE 8th International Conference on Collaboration and Internet Computing (CIC)*, pages 10–15. IEEE, 2022. [1](#)
- [61] Jie Zhang, Song Guo, Xiaosong Ma, Haozhao Wang, Wenchao Xu, and Feijie Wu. Parameterized knowledge transfer for personalized federated learning. *Advances in Neural Information Processing Systems*, 34:10092–10104, 2021. [3](#)
- [62] Jianyi Zhang, Da-Cheng Juan, Cyrus Rashtchian, Chun-Sung Ferng, Heinrich Jiang, and Yiran Chen. Sled: Self logits evolution decoding for improving factuality in large language models. *arXiv preprint arXiv:2411.02433*, 2024. [4](#)
- [63] Jianyi Zhang, Ang Li, Minxue Tang, Jingwei Sun, Xiang Chen, Fan Zhang, Changyou Chen, Yiran Chen, and Hai Li. Fed-cbs: A heterogeneity-aware client sampling mechanism for federated learning via class-imbalance reduction. In *Proceedings of the 40th International Conference on Machine Learning*, 2023. [3](#)
- [64] Jianyi Zhang, Aashiq Muhamed, Aditya Anantharaman, Guoyin Wang, Changyou Chen, Kai Zhong, Qingjun Cui, Yi Xu, Belinda Zeng, Trishul Chilimbi, et al. Reaugkd: Retrieval-augmented knowledge distillation for pre-trained language models. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 1128–1136, 2023. [5](#)
- [65] Jingyang Zhang, Jingwei Sun, Eric Yeats, Yang Ouyang, Martin Kuo, Jianyi Zhang, Hao Frank Yang, and Hai Li. Min-k%++: Improved baseline for detecting pre-training data from large language models. *arXiv preprint arXiv:2404.02936*, 2024. [4](#)
- [66] Jianyi Zhang, Saeed Vahidian, Martin Kuo, Chunyuan Li, Ruiyi Zhang, Tong Yu, Guoyin Wang, and Yiran Chen. 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, 2024. [1](#)
- [67] Jianyi Zhang, Ruiyi Zhang, Lawrence Carin, and Changyou Chen. Stochastic particle-optimization sampling and the non-asymptotic convergence theory. In *International Conference on Artificial Intelligence and Statistics*, pages 1877–1887. PMLR, 2020. [5](#)
- [68] Jianyi Zhang, Yang Zhao, and Changyou Chen. Variance reduction in stochastic particle-optimization sampling. In *International Conference on Machine Learning*, pages 11307–11316. PMLR, 2020. [5](#)
- [69] Ruqi Zhang, Chunyuan Li, Jianyi Zhang, Changyou Chen, and Andrew Gordon Wilson. Cyclical stochastic gradient mcmc for bayesian deep learning. *arXiv preprint arXiv:1902.03932*, 2019. [5](#)
- [70] Yang Zhao, Jianyi Zhang, and Changyou Chen. Self-adversarially learned bayesian sampling, 2018. [5](#)
- [71] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. [3](#)
- [72] Ligeng Zhu, Zhijian Liu, and Song Han. Deep leakage from gradients. *Advances in neural information processing systems*, 32, 2019. [1](#), [3](#), [8](#)