

# LLaVA-KD: A Framework of Distilling Multimodal Large Language Models

Yuxuan Cai<sup>1\*</sup> Jiangning Zhang<sup>2,3\*</sup> Haoyang He<sup>2</sup> Xinwei He<sup>4</sup> Ao Tong<sup>1</sup> Zhenye Gan<sup>3</sup>  
Chengjie Wang<sup>3</sup> Zhucun Xue<sup>2</sup> Yong Liu<sup>2</sup> Xiang Bai<sup>1†</sup>

<sup>1</sup>Huazhong University of Science and Technology <sup>2</sup>Zhejiang University

<sup>3</sup>Tencent YouTu Lab <sup>4</sup>Huazhong Agricultural University

{cyx.hust, xbai}@hust.edu.cn, 186368@zju.edu.cn

## Abstract

*The success of Large Language Models (LLMs) has inspired the development of Multimodal Large Language Models (MLLMs) for unified understanding of vision and language. However, the increasing model size and computational complexity of large-scale MLLMs (l-MLLMs) limit their use in resource-constrained scenarios. Although small-scale MLLMs (s-MLLMs) are designed to reduce computational costs, they typically suffer from performance degradation. To mitigate this limitation, we propose a novel LLaVA-KD framework to transfer knowledge from l-MLLMs to s-MLLMs. Specifically, we introduce Multimodal Distillation (MDist) to transfer teacher model’s robust representations across both visual and linguistic modalities, and Relation Distillation (RDist) to transfer teacher model’s ability to capture visual token relationships. Additionally, we propose a three-stage training scheme to fully exploit the potential of the proposed distillation strategy: 1) Distilled Pre-Training to strengthen the alignment between visual-linguistic representations in s-MLLMs, 2) Supervised Fine-Tuning to equip the s-MLLMs with multimodal understanding capacity, and 3) Distilled Fine-Tuning to refine s-MLLM’s knowledge. Our approach significantly improves s-MLLMs performance without altering the model architecture. Extensive experiments and ablation studies validate the effectiveness of each proposed component. Code will be available at <https://github.com/Fantasyele/LLaVA-KD>.*

## 1. Introduction

Inspired by the significant achievements of Large Language Models (LLM) in the field of Natural Language Processing, an emerging and rapidly developing research area is focusing on the development of Multimodal Large Language Models (MLLM). These models [3, 24, 45] integrate visual encoder, feature projector, and LLM to achieve unified

understanding of visual and linguistic information. However, the success of MLLMs benefits from the scaling law, which significantly increases the model size. The large-scale model and high-cost inference limit the application of MLLMs in resource-constrained scenarios.

Some works [6, 53] address this challenge by directly adopting lightweight LLM backbones to build small-scale MLLMs (s-MLLMs), while preserving the conventional two-stage paradigm of large MLLMs (l-MLLMs) for model training, as illustrated in Fig. 1(a). This paradigm comprises two stages: Pre-Training (PT) and Supervised Fine-Tuning (SFT). The PT stage establishes cross-modal alignment by projecting visual features to the text embedding space, while the SFT stage enhances multimodal understanding capabilities. Nevertheless, directly applying this paradigm—originally designed for l-MLLMs to s-MLLMs often yields suboptimal performance due to inherent capacity limitations [18], preventing s-MLLMs from effectively learning the complex knowledge that l-MLLMs can capture. Recent efforts compensate for this issue by optimizing model structure and improving the quality of training data: MoE-LLaVA [23] introduces the Mixture-of-Experts [17] (MoE) to enhance the model’s ability for understanding complex multimodal information while maintaining the computational cost, and Bunny [13] improves the training data quality by removing redundant data. Unlike these methods, we propose investigating training paradigm optimization as a novel pathway to enhance s-MLLMs performance, without architectural modifications or data engineering.

Knowledge distillation (KD) transfers knowledge from powerful teacher models to small student networks, significantly improves the performance of students with less computation. Although KD has achieved remarkable success in traditional vision tasks, its application to MLLMs remains underexplored. Given that MLLMs primarily rely on language models to achieve robust cross-modal understanding, they present a challenge for distillation: *how to effectively transfer multimodal understanding capabilities via language model distillation*. However, existing LLM

\*Equal Contribution. Work done during internship at Tencent YouTu Lab. †Correspondence to Xiang Bai<[xbai@hust.edu.cn](mailto:xbai@hust.edu.cn)>.

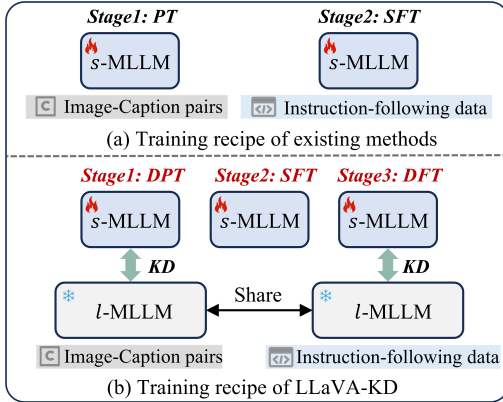


Figure 1. To train a small-scale MLLM, (a) the existing methods follow a two-stage training scheme, including Pre-Training (PT) and Supervised Fine-Tuning (SFT). (b) Our LLaVA-KD proposes a three-stage scheme to exploit the potential of *s*-MLLM, including Distilled Pre-Training (DPT) to align visual-textual representation, SFT to equip the model with multimodal understanding, and Distilled Fine-Tuning (DFT) to refine *s*-MLLM’s capacities.

distillation research focuses solely on knowledge transfer in the text modality, promoting the student model to mimic teacher-generated responses. Visual modality, however, play an equally crucial role in achieving comprehensive cross-modal understanding. To address this limitation, we propose two key innovations: (1) **Multimodal Distillation (MDist)**, which extends distillation optimization to both visual and linguistic modalities, ensuring a holistic transfer of multimodal representations, and (2) **Relation Distillation (RDist)**, which is predicated on the observation that structural relationships among visual tokens encode essential spatial and semantic dependencies for accurately understanding complex visual scenes, such as object positioning and inter-object interactions. We propose RDist to explicitly transfer the teacher model’s capability to capture visual token relationships to the *s*-MLLM.

A straightforward approach is to introduce the above distillation strategy during the SFT stage to enhance *s*-MLLM multimodal understanding capacity. However, we observe that this approach exhibits limited benefits. In this paper, motivated by two key observations, we propose an optimized three-stage training framework, as displayed in Fig. 1(b). **First**, effective cross-modal alignment between visual and textual representations forms the foundation for multimodal understanding. To strengthen this alignment, we incorporate the Distillation during the PT stage (*DPT*), enabling the *s*-MLLM to optimize the visual-linguistic joint embedding space through emulating the multimodal supervision signals from the *l*-MLLM. **Second**, employing knowledge distillation exclusively during SFT is insufficient for comprehensive capability transfer. We therefore propose a *SFT-DFT* (Supervised Fine-Tuning followed by Distilled Fine-Tuning) paradigm where the *s*-MLLM first

establishes baseline understanding through standard SFT, then undergoes dedicated DFT to refine its knowledge.

Our method surpasses existing *s*-MLLM approaches across various multimodal benchmarks. We summarize our contributions as follows:

- We introduce LLaVA-KD, a novel MLLM-oriented distillation framework to transfers the knowledge from large-scale MLLM to the small-scale MLLM. Specifically, it contains a three-stage training scheme, including Distilled Pre-Training (DPT) for enhanced multimodal alignment, Supervised Fine-Tuning (SFT) for knowledge acquisition, and Distilled Fine-Tuning (DFT) for comprehensive knowledge refinement.
- We propose an innovative distillation strategy that integrates Multimodal Distillation (MDist) to optimize both visual and linguistic modalities with Relational Distillation (RDist), enhancing the *s*-MLLM’s ability to capture relationships among visual tokens.
- Comprehensive evaluations across popular multimodal benchmarks demonstrate LLaVA-KD’s superiority over recent *s*-MLLM advances.

## 2. Related Works

### 2.1. Multimodal Large Language Model

The emergence of large language models (LLMs) has promoted significant advances in MLLM for vision-language understanding. Early approaches like BLIP-2 [20] employs a querying transformer architecture to establish cross-modal alignment for multimodal information understanding. In contrast, Flamingo [2] introduces gated cross-attention layers to fuse visual features into LLMs. Recent methods [3, 24, 25] typically employ the lightweight projector such as Multi-Layer Perceptron (MLP) to align visual features with the text embedding space, followed by supervised instruction tuning on large-scale instruction dataset to enhance MLLMs’s understanding capabilities. Most recently, a prominent research direction aims to enhance MLLMs’ fine-grained visual perception through high-resolution input processing [22, 28], enabling diverse downstream applications including image segmentation [39, 50, 51] and grounding [29, 49]. Despite these advancements, the large model size and computational cost of current MLLMs greatly limit their application in resource-constrained scenarios, such as mobile devices.

**Lightweight Multimodal Large Language Model.** Existing lightweight MLLMs primarily reduce model parameters by employing lightweight LLM backbones. For example, LLaVA-Phi [53] adapts the LLaVA1.5 [24] framework by substituting its foundation model with the parameter-efficient Phi-2 backbone. Recent studies demonstrate that optimizing model structure and training data can compensate for performance degradation caused by reduced

model capacity: MoE-LLaVA [23] incorporates mixture-of-experts (MoE) into LLMs, achieving competitive multimodal understanding with merely 3B activated parameters. Bunny [13] curates training data by clustering image embeddings via K-Means, followed by graph-based similarity pruning to reduce data size while maintaining diversity.

Unlike these methods, our approach focuses on improving the training paradigm for MLLMs. In this paper, we propose a three-stage training framework based on knowledge distillation. By transferring the knowledge of large MLLMs to lightweight MLLMs, the light MLLMs’ capabilities will be significantly enhanced.

## 2.2. Knowledge Distillation

Knowledge Distillation (KD) [14] aims to transfer the knowledge from a large, complex teacher model to a lightweight, simple student model. This technique can significantly improve the performance of small models with fewer parameters, less computation, and faster speed. Knowledge distillation has been successfully applied in visual tasks and has achieved success in many fields, typically in the domain of image classification. For example, Early work [14] establishes soft target distillation using teacher logits as supervisory signals. Subsequent innovations like DKMF [37] and FNKD [41] reveal that intermediate feature alignment between teacher and student architectures yields superior classification accuracy. DGKD [34] further improves the student model’s predictions by integrating multiple teacher models for guidance.

**KD for LLM.** With the successful release of ChatGPT and its significant application value, LLM has gradually attracted attention and achieved numerous research progress in recent years [1, 4]. However, to achieve better results, the model size has also become increasingly larger which follows scaling law [18], limiting its application in resource-constrained scenarios. This has promoted research into KD techniques for compressing LLMs while preserving their emergent capabilities.

MiniLLM [11] and DistiLLM [19] optimize the distillation process, proposing reverse Kullback-Leibler divergence (KLD) and skewed KLD objectives respectively, mitigating student overfitting to teacher distribution tails. Wu *et al.* [40] introduce an adaptive strategy that dynamically balances the weighting between standard KLD and reverse KLD objectives. Parallel efforts [15, 30, 35] leverage the Chain-of-Thought (CoT) capability of large LLMs to establish causal dependencies, thereby enriching training data.

**KD for MLLM.** Most recently, LLAVA-MoD [32] advances knowledge distillation for training  $s$ -MLLMs. The method first optimizes the structure of the  $s$ -MLLM by integrating MoE [17, 23] into the LLM, thereby enhancing the model’s expressive ability. During training, it firstly employs standard KLD to align the output distribution between the  $s$ -MLLM and  $l$ -MLLM. Additionally, it introduces a

preference distillation process to improve the  $s$ -MLLM’s discriminative capability, thereby reducing hallucinations. LLaVADI [42] is another  $s$ -MLLM work based on distillation, which reveals that most training strategies designed for LLMs do not bring additional benefits to the MLLMs, suggesting the need for specialized optimization approaches in  $s$ -MLLM development.

Unlike existing LLM/MLLM distillation methods that rely on complex constraint design, or elaborate architectural modifications, we focus on optimizing training schemes and developing multimodal distillation strategies. This enables efficient performance improvement for  $s$ -MLLMs while maintaining architectural simplicity.

## 3. LLaVA-KD

The deployment of lightweight MLLMs is crucial for resource-constrained environments. However, small-scale MLLMs trained using conventional paradigm often yield suboptimal results. For instance, while a 4B TinyLLaVA model achieves 65.0% performance, reducing the model size to 0.5B results in a significant performance drop to 54.7%. To mitigate this issue, we propose an innovative three-stage training scheme, incorporating a novel distillation strategy, termed LLaVA-KD, as illustrated in Fig. 2.

### 3.1. Composition of Distilled MLLM Architecture

Fig. 2(Left) illustrates the distillation process for MLLM, which includes a large-scale  $l$ -MLLM as the teacher model and a small-scale  $s$ -MLLM as the student model. Both them follow the architecture design of LLaVA-1.5 [24], and each includes three main components:

**Frozen Visual Encoder** is used to obtain robust visual features. Specifically, given an input image  $X_v \in \mathbb{R}^{H \times W \times 3}$ , it is first sequenced to 2D patches  $P_v \in \mathbb{R}^{N_p \times S_p^2 \times 3}$  with  $S_p$  and  $N_p$  representing patch size and its number, respectively. The subsequent transformer layers will project these patches to visual features  $Z_v \in \mathbb{R}^{N_p \times C}$ , where  $C$  denotes the dimension. Both teacher and student models use the same visual encoder, and we employ the pre-trained SigLIP [48] following previous success [13, 36].

**Visual Projector** contains two MLP layers with a GELU activation function to project visual features  $Z_v$  into the text embedding space  $H_v \in \mathbb{R}^{N_p \times D}$ , where  $D$  denotes the embedding dimensions.

**Large Language Model (LLM)** The LLM integrates multimodal information through joint processing of visual embeddings  $H_v$  and text embeddings  $H_t$ . These representations are concatenated as  $H = [H_v, H_t]$  and processed to generate sequential outputs:  $\mathbf{y} = [\mathbf{y}_p, \mathbf{y}_v, \mathbf{y}_r] = \{y_t\}_{t=1}^T$ , where  $\mathbf{y}_p$ ,  $\mathbf{y}_v$ , and  $\mathbf{y}_r$  respectively denote prompt, visual, and response token sequences, and  $T$  denotes the length of all prediction tokens.

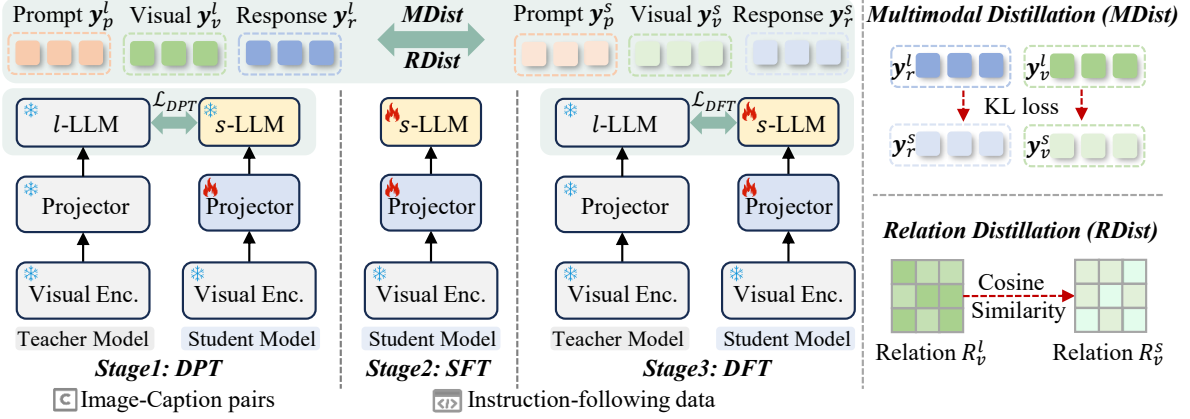


Figure 2. **Overview of our LLaVA-KD** that contains three stages for effect training: **1)** Distilled Pre-Training (DPT) to align visual and text information as *l*-MLLM. **2)** Supervised Fine-Tuning (SFT) to enable *s*-MLLM with multimodal understanding capacity. **3)** Distilled Fine-Tuning (DFT) to refine *s*-MLLM’s capacity by transferring *l*-MLLM’s knowledge. During the training phase, we employ Multimodal Distillation (MDist) and Relation Distillation (RDist) in both DPT and DFT stages.

### 3.2. Training Scheme of Teacher Model *l*-MLLM

This section presents the common training scheme for powerful *l*-MLLMs, establishing performance upper bounds for *s*-MLLMs. This scheme consists of two stages:

**Pre-Training.** The *Visual Encoder* and *l*-LLM are kept frozen, and only the *Projector* is optimized to align visual features with textual features. The training process employs image-caption pairs with the objective:

$$\mathcal{L}_{reg} = - \sum_{m=1}^M \log \phi_l(y_m | \mathbf{y}_{<m}), \quad (1)$$

where  $M$  denotes the length of predicted response tokens, while  $\phi_l(y_m | \mathbf{y}_{<m})$  denotes the distribution for the response token  $y_m$  based on previous predictions  $\mathbf{y}_{<m}$ .

**Supervised Fine-Tuning.** This stage keeps the *Visual Encoder* frozen, aiming at jointly optimizing *Projector* and *l*-LLM to enhance understanding and instruction-following capacities of the teacher model. Training employs high-quality instruction datasets with the objective function  $\mathcal{L}_{SFT}$  following the autoregressive training paradigm defined in Equation 1.

### 3.3. Framework of LLaVA-KD

We first adopt the previous training strategy (Sec. 3.2) to develop the *l*-MLLM. For training *s*-MLLM, we propose a novel distillation strategy tailored for multimodal information learning (Sec. 3.3.1), and we further design a three-stage training scheme (Sec. 3.3.2).

#### 3.3.1. MLLM-Oriented KD Strategy

**Multimodal Distillation (MDist).** Considering that MLLM essentially leverages LLM for multimodal information understanding and reasoning, we follow the naive distillation method of LLM that uses standard

Kullback-Leibler Divergence (KLD) to distill the response predictions. The training objective can be defined as:

$$\begin{aligned} \mathcal{L}_{res} &= \sum_{m=1}^M \text{KLD}(\phi_l(y_m | \mathbf{y}_{<m}), \phi_s(y_m | \mathbf{y}_{<m})), \\ &= \sum_{m=1}^M \sum_{j=1}^V \phi_l(Y_j | \mathbf{y}_{<m}) \log \left( \frac{\phi_l(Y_j | \mathbf{y}_{<m})}{\phi_s(Y_j | \mathbf{y}_{<m})} \right), \end{aligned} \quad (2)$$

where  $M$  represents the length of response tokens and  $V$  denotes the vocabulary space.  $\phi_l$  and  $\phi_s$  denote the model parameters of *l*-MLLM and *s*-MLLM, respectively,  $\phi_l(Y_j | \mathbf{y}_{<m})$  and  $\phi_s(Y_j | \mathbf{y}_{<m})$  denote the probability of vocabulary  $Y_j$  in the response token  $y_m$ , as predicted by *l*-MLLM and *s*-MLLM.

Meanwhile, the visual representation also plays a critical role for multimodal understanding of LLM. To leverage this property, we further optimize the output distributions of visual tokens from the teacher and student:

$$\mathcal{L}_{vis} = \sum_{k=1}^K \sum_{j=1}^V \phi_l(Y_j | \mathbf{y}_{<k}) \log \left( \frac{\phi_l(Y_j | \mathbf{y}_{<k})}{\phi_s(Y_j | \mathbf{y}_{<k})} \right), \quad (3)$$

where  $K$  denotes the length of visual tokens,  $V$  denotes the vocabulary space,  $\phi_l(Y_j | \mathbf{y}_{<k})$  and  $\phi_s(Y_j | \mathbf{y}_{<k})$  represent the probability of vocabulary  $Y_j$  in the visual token  $y_k$ , as predicted by *l*-MLLM and *s*-MLLM respectively.

**Relation Distillation (RDist).** To empower the student model with fine-grained relational reasoning capabilities for complex visual scenes understanding, we construct a self-correlation matrix from the LLM-generated visual tokens. Through knowledge distillation by minimizing the matrix divergence, the student model inherits the teacher model’s ability to comprehend the intricate relationships among vi-

visual tokens. To achieve this, we first compute the self-correlation matrices as follows:

$$\begin{cases} R_v^s = \mathbf{y}_v^s \otimes \mathbf{y}_v^s \in \mathbb{R}^{N_p \times N_p}, \\ R_v^t = \mathbf{y}_v^t \otimes \mathbf{y}_v^t \in \mathbb{R}^{N_p \times N_p}, \end{cases} \quad (4)$$

where  $\otimes$  represents matrix multiplication,  $\mathbf{y}_v^s$  and  $\mathbf{y}_v^t$  denote the visual tokens of the student and teacher, and  $N_p$  denotes the number of visual tokens. Subsequently, we maximum the cosine similarity between  $R_v^s$  and  $R_v^t$  that is defined as:

$$\mathcal{L}_{rel} = 1 - \text{Cos}(R_v^s, R_v^t) = 1 - \frac{R_v^s \cdot R_v^t}{\|R_v^s\| \|R_v^t\|}, \quad (5)$$

where  $\text{Cos}(\cdot)$  denotes the cosine similarity.

### 3.3.2. Three-Stage Distillation Scheme

Motivated by two critical observations, we propose a novel three-stage training paradigm, as illustrated in Figure 2. First, effective cross-modal alignment between visual and textual representations serves as the foundation for comprehensive multimodal understanding. To achieve this, we introduce Distilled Pre-Training (DPT) that enhances visual feature quality through targeted knowledge transfer, thereby promoting robust visual-textual alignment. Second, we observe that directly applying distillation during SFT yields limited benefits. To address this limitation, we propose a paradigm where the student model first develops baseline comprehension through standard SFT, then undergoes refinement through Distilled Fine-Tuning (DFT). This phased approach enables progressive knowledge acquisition and refinement and ultimately enhances the model’s overall multimodal understanding capacity.

**Distilled Pre-Training (DPT).** The main purpose of this stage is to project visual features to the text embedding space. Previous methods [24, 53] primarily achieve this by optimizing the autoregressive prediction process of LLM (Eq. 1). In contrast, LLaVA-KD further incorporates the distillation strategy to enhance cross-modal alignment.

Specifically, the visual encoder and language model components from the  $s$ -MLLM remain frozen during training, with optimization exclusive to the projection module. MDist is utilized to minimize the discrepancy between the student and the teacher model in terms of the output distribution of visual and response, while RDist is incorporated to enable the student to learn from the teacher model’s ability to handle complex visual interactions. The combined MDist and RDist further optimize the quality of visual features, promoting the alignment between projected visual features and textual embeddings.

Overall, in addition to optimizing the autoregressive prediction results, we also utilize a multimodal distillation and relation distillation procedure. The objective is defined as:

$$\mathcal{L}_{DPT} = \mathcal{L}_{reg} + \alpha \mathcal{L}_{res} + \beta \mathcal{L}_{vis} + \gamma \mathcal{L}_{rel}, \quad (6)$$

where  $\mathcal{L}_{reg}$  denotes the auto-regressive prediction loss,  $\alpha, \beta$ , are weights for visual and response distillation in MDist, and  $\gamma$  is the weight for RDist.

**Supervised Fine-Tuning (SFT).** At this stage, we follow the common SFT procedure of the  $l$ -MLLM training phase (Sec. 3.2). By jointly training the Projector and  $s$ -LLM, the model is initialized with multimodal information understanding ability and instruction-following capability. The training objective can be defined as Eq. 1, denoted as  $\mathcal{L}'_{SFT}$ .

**Distilled Fine-Tuning (DFT).** The main objective of this stage is to further refine the capacity of the  $s$ -MLLM. Specifically, we integrate the distillation strategies of MDist and RDist during this training phase to effectively acquire knowledge from the  $l$ -MLLM. MDist enables the  $s$ -MLLM to directly emulate the  $l$ -MLLM’s sophisticated multimodal understanding capabilities, whereas RDist enhances the  $s$ -MLLM’s visual representations by learning to capture visual token relationships, promoting a more context-aware understanding in the  $s$ -MLLM. Overall, the training objective can be defined as:

$$\mathcal{L}_{DFT} = \mathcal{L}_{reg} + \alpha' \mathcal{L}_{res} + \beta' \mathcal{L}_{vis} + \gamma' \mathcal{L}_{rel}, \quad (7)$$

where  $\mathcal{L}_{reg}$  denotes the auto-regressive prediction loss,  $\alpha', \beta'$  are weights for visual and response distillation in MDist, and  $\gamma'$  is the weight for RDist.

## 4. Experimental Results

### 4.1. Setup

**Implementation Details.** For both the large/small-scale MLLMs, we employ the pre-trained SigLIP-B/14@384px [48] as the Visual Encoder and a two-layer MLP with a GELU activation layer as the Projector, while adopting Qwen1.5 [43] and Qwen2.5 families as LLM models. During training, we utilize the LLaVA1.5-558k [24] for the DPT stage, and LLaVA-mix-665k [24] for both SFT and DFT stages. The loss weights ( $\{\alpha, \beta, \gamma\}$  and  $\{\alpha', \beta', \gamma'\}$ ) during DPT and DFT stages are set to 1.0, 1.0, and 0.5, respectively. The detailed training hyperparameter configurations are illustrated in Appendix.

**Details of Comparison Methods.** We primarily compare with recent efforts focused on small-scale MLLMs, including Imp [31], Bunny [13], Mini-Gemini [22], MoE-LLaVA [23], SPHINX [9], TinyLLaVA [52] and LLaVA-MoD [32]. Additionally, we also compare our LLaVA-KD with current state-of-the-art MLLMs, such as InstructBLIP [7], mPLUG-Owl2 [46], LLaVA1.5 [24], Qwen-VL [3], MobileVLM [6], MiniCPM-V [45].

**Benchmark Settings.** General VQA requires models to generate answers based on the image and related question, necessitating the ability to understand how visual and textual information interrelate. For general VQA,

Method	LLM	#Samples	Image Question Answering					Benchmarks					Avg <sub>7</sub>	Avg <sub>10</sub>
			VQAv2	GQA	VisWiz	SciQA	TextVQA	MME	MMB	MMB <sup>CN</sup>	POPE	MMMU		
LLaVA-1.5	Vicuna-7B	1.2 M	78.5	62.0	50.0	66.8	58.2	75.5	64.3	58.3	85.9	34.4	62.2	63.4
InstructBLIP	Vicuna-7B	130 M	-	49.2	-	60.5	50.1	-	36.0	-	-	-	-	-
Qwen-VL	Qwen-7B	1500 M	78.8	59.3	35.2	67.1	63.8	-	38.2	7.4	-	-	-	-
Qwen-VL-Chat	Qwen-7B	1500 M	78.2	57.5	38.9	68.2	61.5	74.4	60.6	56.7	-	35.9	59.7	-
mPLUG-Owl2	LLaMA2-7B	400 M	79.4	56.1	54.5	68.7	54.3	72.5	66.5	-	85.8	32.7	62.1	-
TinyLLaVA <sup>†</sup>	Qwen1.5-4B	1.2 M	79.9	63.4	46.3	72.9	59.0	69.3	67.9	67.1	85.2	38.9	63.7	65.0
TinyLLaVA <sup>†</sup>	Qwen2.5-3B	1.2 M	80.4	63.2	38.7	76.0	61.5	73.9	71.8	69.5	86.4	40.3	64.9	66.2
TinyLLaVA	Phi2-2.7B	1.2 M	79.9	62.0	-	69.1	59.1	73.2	66.9	-	86.4	38.4	-	-
Bunny	Phi2-2.7B	2.6 M	79.8	62.5	43.8	70.9	56.7	74.4	68.6	37.2	-	38.2	59.2	-
Imp-3B	Phi2-2.7B	1.5 M	-	63.5	54.1	72.8	59.8	-	72.9	46.7	-	-	-	-
MobileVLM	MLLaMA-2.7B	1.2 M	-	59.0	-	61.0	47.5	64.4	59.6	-	84.9	-	-	-
MoE-LLaVA	Phi2-2.7B	2.2 M	79.9	62.6	-	70.3	57.0	-	68.0	-	85.7	-	-	-
MiniCPM-V	MiniCPM-2.4B	570 M	-	51.5	50.5	74.4	56.6	68.9	64.0	62.7	79.5	-	61.2	-
LLaVADI	MLLaMA-2.7B	1.2 M	-	61.4	-	64.1	50.7	68.8	62.5	-	86.7	-	-	-
Imp-2B	Qwen1.5-1.8B	1.5 M	79.2	61.9	39.6	66.1	54.5	65.2	63.8	61.3	86.7	-	58.9	-
Bunny-2B	Qwen1.5-1.8B	2.6 M	76.6	59.6	34.2	64.6	53.2	65.0	59.1	58.5	85.8	-	56.3	-
Mini-Gemini-2B	Gemma-2B	2.7 M	-	60.7	41.5	63.1	56.2	67.0	59.8	51.3	85.6	31.7	57.1	-
MoE-LLaVA-2B	Qwen-1.5-1.8B	2.2 M	76.2	61.5	32.6	63.1	48.0	64.6	59.7	57.3	<b>87.0</b>	-	55.3	-
TinyLLaVA <sup>†</sup>	Qwen2.5-1.5B	1.2 M	78.8	62.0	43.2	<b>72.0</b>	57.4	<b>72.5</b>	68.6	63.0	85.5	37.0	62.7	64.0
TinyLLaVA <sup>†</sup>	Qwen1.5-1.8B	1.2 M	73.1	55.5	34.9	65.3	47.7	61.2	57.1	55.5	83.4	34.1	53.9	56.8
LLaVA-MoD	Qwen1.5-1.8B	5 M	-	58.7	39.2	68.0	58.5	66.7	66.3	61.9	<b>87.0</b>	-	59.9	-
LLaVA-KD	Qwen1.5-1.8B	1.2 M	79.0	<u>62.3</u>	<u>44.7</u>	64.7	53.4	69.1	64.0	63.7	86.3	33.6	60.3	62.1
LLaVA-KD	Qwen2.5-1.5B	1.2 M	<b>80.3</b>	<b>62.5</b>	<b>46.0</b>	71.6	<b>59.7</b>	<b>70.0</b>	<b>71.0</b>	<b>66.6</b>	86.7	<b>35.8</b>	<b>63.9</b>	<b>65.0</b>
SPHINX-Tiny	TinyLlama-1.1B	15 M	74.7	58.0	49.2	21.5	57.8	63.1	52.3	56.6	82.2	-	51.2	-
TinyLLaVA <sup>†</sup>	Qwen1.5-0.5B	1.2 M	73.9	57.4	24.9	60.9	47.4	59.8	55.0	52.4	83.7	31.6	51.1	54.7
TinyLLaVA <sup>†</sup>	Qwen2.5-0.5B	1.2 M	74.8	58.3	28.9	59.1	49.2	61.5	58.9	54.2	86.1	33.6	52.9	56.5
LLaVADI	MLLaMA-1.4B	1.2 M	-	55.4	-	56.0	45.3	58.9	55.0	-	84.7	-	-	-
LLaVA-MoD	Qwen1.5-0.5B	5 M	-	56.2	31.6	62.8	53.9	65.3	58.8	50.4	-	-	54.1	-
LLaVA-KD	Qwen1.5-0.5B	1.2 M	77.0	59.6	35.9	60.6	49.9	64.5	60.1	55.5	85.9	30.2	55.2	57.9
LLaVA-KD	Qwen2.5-0.5B	1.2 M	<b>77.7</b>	<b>59.8</b>	41.5	60.6	52.0	64.7	<b>61.3</b>	<b>57.0</b>	<b>86.4</b>	28.3	<b>56.7</b>	<b>58.9</b>

Table 1. **Benchmarked results with SoTA MLLMs.** Compared with counterparts, our LLaVA-KD achieves highly competitive results than current small-scale MLLM models. Optimal and sub-optimal results are in **bold** and underline. **grey**, **orange** and **blue** backgrounds represent ours baseline, the most direct MLLM distillation methods and our approach, respectively. Avg<sub>7</sub>: The average of the seven benchmarks for direct comparison with existing MLLM distillation methods, excluding VQAv2, POPE, MMMU. Avg<sub>10</sub>: The average across all benchmarks for comprehensive comparison. †: reproduced results using the official code.

we evaluate LLaVA-KD on four benchmarks including VQAv2 [10], GQA [16], VizWiz [12], and ScienceQA (Image set) [27]. Scene Text-centric VQA (TextVQA [33]) requires the model recognize and understand textual information in an image. Additionally, we use five popular benchmarks for comprehensive evaluation including MME [8], MMB [26], MMB<sup>CN</sup> [26], POPE [21], and MMMU [47].

## 4.2. Benchmarked Results with the SoTAs

As shown in Table 1, our LLaVA-KD demonstrates superior performance across both 1B and 2B model scales. Specifically, for the 1B configuration, when initialized with Qwen1.5-0.5B, LLaVA-KD achieves a 4.0% average improvement in Avg<sub>7</sub> over SPHINX-Tiny [9], using only 1M training samples compared to SPHINX-Tiny’s 15M. Meanwhile, we outperform the recent MLLM distillation art LLaVA-MoD [32] by 1.1% in Avg<sub>7</sub>, while requiring less training data (1.2M vs. 5M). Notably, LLaVA-MoD introduces a MoE structure in the *s*-MLLM, resulting in large total parameters. Further employing Qwen2.5-0.5B improves LLaVA-KD’s performance. Moreover, compared to the TinyLLaVA baseline (Qwen1.5-0.5B/Qwen2.5-0.5B) using conventional PT-SFT training scheme, LLaVA-KD achieves improvements of 3.2% and 2.4% in Avg<sub>10</sub> respectively across all ten benchmarks.

Our advantage extends to the 2B model scale, achieving the leading performance compared to existing small-scale MLLMs. When built upon Qwen1.5-1.8B, our model out-

performs both LLaVA-MoD [32] and Imp-2B [31] by 0.4% and 1.4% in Avg<sub>7</sub> respectively, while maintaining greater data efficiency. Similar to the 1B observations, LLaVA-KD-2B demonstrates comprehensive improvements over TinyLLaVA baseline (Qwen1.5-1.8B/Qwen2.5-1.5B) with gains of 5.3% and 1.0% in Avg<sub>10</sub>, respectively.

## 4.3. Ablation Study and Analysis

**Three-Stage Training Recipe.** In Table 2, we investigate the impact of different training stages. Initially, we follow the conventional Pre-Training (PT) followed by Supervised Fine-Tuning (SFT) recipe to train the small MLLM, which serves as our baseline (Row 1) with an average performance of 54.7%. Incorporating DPT alongside SFT (Row 2) results in a 0.9% improvement (55.6% vs. 54.7%), suggesting that DPT enhances the alignment between visual feature and text embedding and thereby facilitates the LLM’s understanding of multimodal information. When we further employ DFT (DPT-SFT-DFT, Row 6), the model’s performance improves significantly by 2.3%, achieving 57.9% in Avg<sub>10</sub>. The result indicates that DFT effectively enables the *s*-MLLM to learn knowledge from the *l*-MLLM, substantially enhancing its understanding capabilities. In contrast, omitting the SFT stage (DPT-DFT, Row 4) results in a notable performance drop to 55.9%, underscoring the necessity of the SFT stage for knowledge acquisition. Notably, the DPT-DFT achieves optimal results in the two-stage training paradigm. Introducing an additional DFT

Training Scheme	Image Question Answering					Benchmarks					Avg <sub>10</sub>
	VQAv2	GQA	VizWiz	SciQA	TextVQA	MME	MMB	MMB <sup>GN</sup>	POPE	MMMU	
PT-SFT	73.9	57.4	24.9	60.9	47.4	59.8	55.0	52.4	83.7	<b>31.6</b>	54.7
DPT-SFT	74.6	57.8	28.8	<b>61.2</b>	49.1	59.9	56.9	51.6	84.3	31.4	55.6
PT-DFT	75.1	57.0	29.5	60.9	49.2	59.6	57.3	55.0	85.5	29.6	55.8
DPT-DFT	75.5	58.0	27.5	59.7	49.3	60.6	57.7	54.7	85.4	30.3	55.9
PT-SFT-DFT	76.6	59.4	32.6	60.4	48.4	60.9	57.8	54.0	84.9	31.3	56.6
DPT-SFT-DFT	<b>77.0</b>	<b>59.6</b>	<b>35.9</b>	60.6	<b>49.9</b>	<b>64.5</b>	<b>60.1</b>	<b>55.5</b>	<b>85.9</b>	30.2	<b>57.9</b>
DPT-DFT-DFT	77.5	60.3	37.6	61.1	49.3	63.2	59.4	54.9	86.0	31.0	58.0

Table 2. Ablation studies of different training stages.

DPT		SFT	DFT		Avg <sub>10</sub>
MDist	RDist		MDist	RDist	
✓	✓	✓	✓	✓	55.5
✓	✓	✓	✓	✓	55.1
✓	✓	✓	✓	✓	<b>55.6</b>
✓	✓	✓	✓	✓	57.0
✓	✓	✓	✓	✓	57.7
✓	✓	✓	✓	✓	<b>57.9</b>

Table 3. Ablation Study on MDist and RDist.

Training Stage	Response tokens	Prompt tokens	Visual tokens	Avg <sub>10</sub>
DPT	✓	✓	✓	54.9
	✓	✓	✓	55.0
	✓	✓	✓	55.1
	✓	✓	✓	54.6
DFT	✓	✓	✓	57.2
	✓	✓	✓	56.9
	✓	✓	✓	57.7
	✓	✓	✓	57.1

Table 4. Ablation Study on Distillation Targets.

phase (Row 7) yields comparable performance to DPT-SFT-DFT with 58.0%, but requires increased computational overhead (120 GPU hours). This indicates that DPT-SFT-DFT achieves the best performance-computation trade-off.

Additionally, we observe that the performance of PT-DFT (Row 3) and DPT-DFT is comparable. Meanwhile, when the SFT stage is incorporated, the DPT-based pipeline (DPT-SFT-DFT, Row 6) outperforms PT-based pipeline (PT-SFT-DFT, Row 5) by 1.3%, emphasizing the advantage provided by DPT. The comparable results obtained in the absence of SFT highlight the critical role of DFT, which enhances the model’s comprehension ability regardless of the pretraining strategy. Although DPT establishes superior multimodal alignment during initial pretraining, DFT’s robust knowledge transfer compensates for the limitations of PT, ultimately leading to similar performance between PT-DFT and DPT-DFT.

**Distillation Strategy.** Table 3 explores the influence of different distillation strategies, including MDist and RDist, during both the DPT and DFT stages. First, we report the results of DPT using different distillation strategies, followed by SFT (Rows 1-3). The results show that using RDist alone is more effective than using MDist alone, with this improvement potentially attributed to RDist’s effectiveness in preserving relationships in visual representations, thereby enhancing feature quality and promoting effective visual-textual alignment for improved cross-modal understanding in the *s*-MLLM. Please note that while more sophisticated methods could capture finer visual relations, they typically introduce greater complexity. As such, our choice of RDist represents a practical balance between complexity and ef-

fectiveness. During the DFT stage (Rows 4-6), MDist proves to be more effective than RDist. The advantage likely arises from the direct alignment of the output multimodal distribution of large MLLMs, which enhances the understanding capabilities of small MLLMs. Across both distillation stages, combining MDist and RDist achieves optimal performance. The results demonstrate that integrating both MDist and RDist helps to comprehensively transfer the knowledge from large MLLMs to small MLLMs.

**Distillation Targets.** As shown in Table 4, we validate the effectiveness of different distillation targets in MDist during both the DPT stage and DFT stage. The results indicate that, unlike most existing methods that focus solely on distilling the response, incorporating visual distillation achieves the best results, whether in the DPT or DFT stage. This improvement can be attributed to that, in the DPT stage, adding visual constraints helps improve the quality of the projected visual space in the *s*-MLLM, thereby promoting the alignment of visual and language information, facilitating unified understanding by the LLM. Furthermore, during the DFT stage, direct imitation of visual-textual cross-modal comprehension in *l*-MLLM, further enhances the *s*-MLLM’s understanding and reasoning capabilities.

**Comparison of Distillation strategy of LLM.** We further evaluate our proposed MDist and RDist distillation strategy against existing LLM distillation approaches under the PT-DFT training recipe, as illustrated in Table 5. Our method achieves the best results across seven benchmarks, demonstrating an average improvement of 0.3% compared to Forward Kullback-Leibler Divergence (FKL) and Jensen-Shannon Divergence [38] (JSD), and an average improvement of 0.8% compared to Reverse Kullback-Leibler Divergence [11] (RKL).

Unlike existing LLM distillation methods that primarily focus on the optimization of textual information, our distillation strategy incorporates the role of visual information in multimodal understanding tasks. This approach enables small-scale MLLMs to better align with the multimodal representations of large-scale MLLMs, as well as improving the ability to capture the relationships across visual tokens. By refining the distillation process to emphasize visual-token alignment, the *s*-MLLM generates more context-aware predictions, leading to superior overall performance.

**Comparison with the distillation pipeline in LLaVA-MOD.** LLaVA-MOD [32] introduces a combined distillation pipeline of Mimic Distillation (MD) and Preference

Distillation strategy	Image Question Answering					Benchmarks					$Avg_{10}$
	VQAv2	GQA	VizWiz	SciQA	TextVQA	MME	MMB	MMB <sup>CN</sup>	POPE	MMMU	
FKL	74.3	56.1	31.7	59.4	49.0	58.9	57.4	54.0	84.4	29.8	55.5
RKL [111]	74.3	56.6	26.7	60.8	49.1	57.8	56.8	53.7	84.7	<b>30.0</b>	55.0
JSD [38]	73.8	54.9	<b>32.3</b>	60.3	48.7	57.6	<b>57.8</b>	54.3	85.1	29.8	55.5
Ours	<b>75.1</b>	<b>57.0</b>	29.5	<b>60.9</b>	<b>49.2</b>	<b>59.6</b>	57.3	<b>55.0</b>	<b>85.5</b>	29.6	<b>55.8</b>

Table 5. Comparison with Distillation Strategies in LLMs.

Distillation Strategy	Image Question Answering					Benchmarks					$Avg_{10}$
	VQAv2	GQA	VisWiz	SciQA	TextVQA	MME	MMB	MMB <sup>CN</sup>	POPE	MMMU	
MD	76.3	58.5	31.6	58.4	<b>51.7</b>	60.6	59.6	<b>55.8</b>	86.2	30.2	56.9
MD+PD	74.4	57.1	22.7	58.4	47.7	58.4	58.8	54.5	<b>86.6</b>	<b>32.1</b>	55.1
Ours	<b>77.0</b>	<b>59.6</b>	<b>35.9</b>	<b>60.6</b>	49.9	<b>64.5</b>	<b>60.1</b>	55.5	85.9	30.2	<b>57.9</b>

Table 6. Comparison with distillation pipeline in LLaVA-MOD. MD and PD denote the Mimic Distillation and Preference Distillation.

LLM of the Teacher	LLM of the Student	Training Recipe	$Avg_{10}$
MLLaMA 2.7B	/	PT-SFT	55.2
/	MLLaMA 1.7B	PT-SFT	48.8
MLLaMA 2.7B	MLLaMA 1.7B	DPT-SFT	50.5
MLLaMA 2.7B	MLLaMA 1.7B	DPT-SFT-DFT	53.4

Table 7. Verification on MobileVLM.

LLM of the Teacher	LLM of the Student	Training Recipe	$Avg_{10}$
Qwen1.5-4B	/	PT-SFT	65.0
Qwen1.5-7B	/	PT-SFT	65.7
/	Qwen2.5-0.5B	PT-SFT	54.7
Qwen1.5-4B	Qwen1.5-0.5B	DPT-SFT-DFT	57.9
Qwen1.5-7B	Qwen1.5-0.5B	DPT-SFT-DFT	57.4
Qwen2.5-3B	/	PT-SFT	66.2
Qwen2.5-7B	/	PT-SFT	69.3
/	Qwen2.5-0.5B	PT-SFT	56.5
Qwen2.5-3B	Qwen2.5-0.5B	DPT-SFT-DFT	58.9
Qwen2.5-7B	Qwen2.5-0.5B	DPT-SFT-DFT	58.3

Table 8. Ablation study on teacher models with different sizes.

Distillation (PD). MD enables the  $s$ -MLLM to emulate the understanding of the  $l$ -MLLM. In contrast, PD enhances the  $s$ -MLLM’s discriminative capacity by leveraging the preference signals from the  $l$ -MLLM. To compare with LLaVA-MOD, we replace our DFT stage in the DPT-SFT-DFT paradigm with MD and PD. As illustrated in Table 6, when employing the MD, our method demonstrates an average improvement of 1.0%. Moreover, integrating PD alongside MD reduces overall performance to 55.1%, with our method showing an advantage of 2.8%. This aligns with LLaVA-MoD’s observation that PD primarily reduces the hallucination of  $s$ -MLLM, but this also results in a decline in the model’s comprehension capability.

**Verification on another MLLM framework.** We additionally conduct experiments on the MobileVLM [6] framework to evaluate the generalization of our training algorithm, as illustrated in Table 7. The  $l$ -MLLM and  $s$ -MLLM employ MobileLLaMA-2.7B and -1.7B respectively, achieving 55.2% and 48.8% average performance across ten benchmarks. Application of DPT-SFT to  $s$ -MLLM yields 50.5% performance. The complete DPT-SFT-DFT paradigm achieves 53.4%, achieving 4.6% improvement over conventional PT-SFT.

**Impact of teacher models with different sizes.** In Table 8, we investigate the impact of teacher model size on LLaVA-KD by employing different variants of LLMs (Qwen1.5-4B/7B and Qwen2.5-3B/7B). The results demonstrate that our LLaVA-KD with either teacher variant significantly enhances the capabilities of  $s$ -MLLM compared to the conventional PT-SFT paradigm. However, counter-intuitively, larger teachers (Qwen1.5-7B and Qwen2.5-7B) exhibit lower distillation effectiveness than the smaller ones

(Qwen1.5-4B and Qwen2.5-3B).

This finding aligns with recent knowledge distillation studies [5, 44], which suggest that beyond a critical capability threshold, increasing teacher capacity yields diminishing returns for student performance. Specifically, excessive capability gaps between teachers and students create an alignment bottleneck during distillation, limiting the student model’s expressive capacity. This insight highlights the importance of optimizing teacher-student capacity matching in MLLMs knowledge distillation, suggesting an open issue for future research.

**Futher Exploration.** It should be noted that in our framework, to ensure that the  $s$ -MLLM can effectively learn from the  $l$ -MLLM, both  $l$ -MLLM and  $s$ -MLLM need to employ the same series of LLMs to maintain consistency in the vocabulary space. Future research could address both architectural limitations and capacity mismatch simultaneously, potentially enabling integration of heterogeneous MLLMs to leverage complementary knowledge sources. This dual-axis improvement pathway may yield more powerful teacher models while further enhancing  $s$ -MLLM performance through optimized knowledge transfer.

## 5. Conclusion

This paper introduced LLaVA-KD, a framework that transfers knowledge from  $l$ -MLLM to  $s$ -MLLM. This approach effectively reduces model size and computational complexity while enabling the  $s$ -MLLM to maintain the capabilities. LLaVA-KD introduces a distillation strategy, including MDist and RDist. MDist simultaneously transfer the visual-textual cross-modal information, whereas RDist transfers  $l$ -MLLM’s capacity to capture correlations among visual representations. Additionally, we propose a simple yet effective three-stage training framework to fully exploit the potential of distillation: DPT to promote the alignment between visual-textual representations in the  $s$ -MLLM, SFT to equip model with multimodal understanding, and DFT to further refine  $s$ -MLLM’s capacity. Comprehensive experiments reveal the effectiveness of our framework.

## Acknowledgments

This work was supported by the NSFC (62225603).

## References

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023. 3
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022. 2
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 1, 2, 5
- [4] Tom B Brown. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020. 3
- [5] Dan Busbridge, Amitis Shidani, Floris Weers, Jason Ramapuram, Etai Littwin, and Russ Webb. Distillation scaling laws. *arXiv preprint arXiv:2502.08606*, 2025. 8
- [6] Xiangxiang Chu, Limeng Qiao, Xinyang Lin, Shuang Xu, Yang Yang, Yiming Hu, Fei Wei, Xinyu Zhang, Bo Zhang, Xiaolin Wei, et al. Mobilevlm: A fast, reproducible and strong vision language assistant for mobile devices. *arXiv preprint arXiv:2312.16886*, 2023. 1, 5, 8
- [7] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning, 2023. 5
- [8] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 6
- [9] Peng Gao, Renrui Zhang, Chris Liu, Longtian Qiu, Siyuan Huang, Weifeng Lin, Shitian Zhao, Shijie Geng, Ziyi Lin, Peng Jin, et al. Sphinx-x: Scaling data and parameters for a family of multi-modal large language models. *arXiv preprint arXiv:2402.05935*, 2024. 5, 6
- [10] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017. 6
- [11] Yuxian Gu, Li Dong, Furu Wei, and Minlie Huang. Minillm: Knowledge distillation of large language models. In *The Twelfth International Conference on Learning Representations*, 2024. 3, 7, 8
- [12] Danna Gurari, Qing Li, Abigale J Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P Bigham. Vizwiz grand challenge: Answering visual questions from blind people. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3608–3617, 2018. 6
- [13] Muyang He, Yexin Liu, Boya Wu, Jianhao Yuan, Yueze Wang, Tiejun Huang, and Bo Zhao. Efficient multi-modal learning from data-centric perspective. *arXiv preprint arXiv:2402.11530*, 2024. 1, 3, 5
- [14] Geoffrey Hinton. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 3
- [15] Cheng-Yu Hsieh, Chun-Liang Li, Chih-Kuan Yeh, Hootan Nakhost, Yasuhisa Fujii, Alexander Ratner, Ranjay Krishna, Chen-Yu Lee, and Tomas Pfister. Distilling step-by-step! outperforming larger language models with less training data and smaller model sizes. *arXiv preprint arXiv:2305.02301*, 2023. 3
- [16] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019. 6
- [17] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991. 1, 3
- [18] Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. 1, 3
- [19] Jongwoo Ko, Sungnyun Kim, Tianyi Chen, and Se-Young Yun. Distillm: Towards streamlined distillation for large language models. *arXiv preprint arXiv:2402.03898*, 2024. 3
- [20] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. 2
- [21] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023. 6
- [22] Yanwei Li, Yuechen Zhang, Chengyao Wang, Zhisheng Zhong, Yixin Chen, Ruihang Chu, Shaoteng Liu, and Jiaya Jia. Mini-gemini: Mining the potential of multi-modality vision language models. *arXiv preprint arXiv:2403.18814*, 2024. 2, 5
- [23] Bin Lin, Zhenyu Tang, Yang Ye, Jiayi Cui, Bin Zhu, Peng Jin, Junwu Zhang, Munan Ning, and Li Yuan. Moe-llava: Mixture of experts for large vision-language models. *arXiv preprint arXiv:2401.15947*, 2024. 1, 3, 5
- [24] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306, 2024. 1, 2, 3, 5
- [25] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36, 2024. 2
- [26] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an

- all-around player? *arXiv preprint arXiv:2307.06281*, 2023. 6
- [27] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. *Advances in Neural Information Processing Systems*, 35:2507–2521, 2022. 6
- [28] Gen Luo, Yiyi Zhou, Yuxin Zhang, Xiawu Zheng, Xiaoshuai Sun, and Rongrong Ji. Feast your eyes: Mixture-of-resolution adaptation for multimodal large language models. *arXiv preprint arXiv:2403.03003*, 2024. 2
- [29] Chuofan Ma, Yi Jiang, Jiannan Wu, Zehuan Yuan, and Xiaojuan Qi. Groma: Localized visual tokenization for grounding multimodal large language models. In *European Conference on Computer Vision*, pages 417–435. Springer, 2024. 2
- [30] Leonardo Ranaldi and Andre Freitas. Aligning large and small language models via chain-of-thought reasoning. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1812–1827, 2024. 3
- [31] Zhenwei Shao, Zhou Yu, Jun Yu, Xuecheng Ouyang, Lihao Zheng, Zhenbiao Gai, Mingyang Wang, and Jiajun Ding. Imp: Highly capable large multimodal models for mobile devices. *arXiv preprint arXiv:2405.12107*, 2024. 5, 6
- [32] Fangxun Shu, Yue Liao, Le Zhuo, Chenning Xu, Guanghao Zhang, Haonan Shi, Long Chen, Tao Zhong, Wanggui He, Siming Fu, et al. Llava-mod: Making llava tiny via moe knowledge distillation. *arXiv preprint arXiv:2408.15881*, 2024. 3, 5, 6, 7
- [33] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019. 6
- [34] Wonchul Son, Jaemin Na, Junyong Choi, and Wonjun Hwang. Densely guided knowledge distillation using multiple teacher assistants. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9395–9404, 2021. 3
- [35] Yijun Tian, Yikun Han, Xiushi Chen, Wei Wang, and Nitesh V Chawla. Tinyllm: Learning a small student from multiple large language models. *arXiv preprint arXiv:2402.04616*, 2024. 3
- [36] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*, 2024. 3
- [37] Guo-Hua Wang, Yifan Ge, and Jianxin Wu. Distilling knowledge by mimicking features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(11):8183–8195, 2021. 3
- [38] Yuqiao Wen, Zichao Li, Wenyu Du, and Lili Mou. f-divergence minimization for sequence-level knowledge distillation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 10817–10834, 2023. 7, 8
- [39] Jiannan Wu, Muyan Zhong, Sen Xing, Zeqiang Lai, Zhaoyang Liu, Zhe Chen, Wenhai Wang, Xizhou Zhu, Lewei Lu, Tong Lu, et al. Visionllm v2: An end-to-end generalist multimodal large language model for hundreds of vision-language tasks. *Advances in Neural Information Processing Systems*, 37:69925–69975, 2025. 2
- [40] Taiqiang Wu, Chaofan Tao, Jiahao Wang, Zhe Zhao, and Ngai Wong. Rethinking kullback-leibler divergence in knowledge distillation for large language models. *arXiv preprint arXiv:2404.02657*, 2024. 3
- [41] Kunran Xu, Lai Rui, Yishi Li, and Lin Gu. Feature normalized knowledge distillation for image classification. In *European conference on computer vision*, pages 664–680. Springer, 2020. 3
- [42] Shilin Xu, Xiangtai Li, Haobo Yuan, Lu Qi, Yunhai Tong, and Ming-Hsuan Yang. Llavadi: What matters for multimodal large language models distillation. *arXiv preprint arXiv:2407.19409*, 2024. 3
- [43] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*, 2024. 5
- [44] Chuanguang Yang, Zhulin An, Libo Huang, Junyu Bi, Xinqiang Yu, Han Yang, Boyu Diao, and Yongjun Xu. Clipkd: An empirical study of clip model distillation. In *CVPR*, pages 15952–15962, 2024. 8
- [45] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024. 1, 5
- [46] Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, and Fei Huang. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13040–13051, 2024. 5
- [47] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567, 2024. 6
- [48] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986, 2023. 3, 5
- [49] Hao Zhang, Hongyang Li, Feng Li, Tianhe Ren, Xueyan Zou, Shilong Liu, Shijia Huang, Jianfeng Gao, Leizhang, Chunyuan Li, et al. Llava-grounding: Grounded visual chat with large multimodal models. In *European Conference on Computer Vision*, pages 19–35. Springer, 2024. 2
- [50] Tao Zhang, Xiangtai Li, Hao Fei, Haobo Yuan, Shengqiong Wu, Shunping Ji, Chen Change Loy, and Shuicheng Yan. Omg-llava: Bridging image-level, object-level, pixel-level reasoning and understanding. *Advances in Neural Information Processing Systems*, 37:71737–71767, 2025. 2

- [51] Zheng Zhang, Yeyao Ma, Enming Zhang, and Xiang Bai. Psalm: Pixelwise segmentation with large multi-modal model. In *European Conference on Computer Vision*, pages 74–91. Springer, 2024. [2](#)
- [52] Baichuan Zhou, Ying Hu, Xi Weng, Junlong Jia, Jie Luo, Xien Liu, Ji Wu, and Lei Huang. Tynyllava: A framework of small-scale large multimodal models. *arXiv preprint arXiv:2402.14289*, 2024. [5](#)
- [53] Yichen Zhu, Minjie Zhu, Ning Liu, Zhicai Ou, Xiaofeng Mou, and Jian Tang. Llava-phi: Efficient multi-modal assistant with small language model. *arXiv preprint arXiv:2401.02330*, 2024. [1](#), [2](#), [5](#)