

LEO-MINI: An Efficient Multimodal Large Language Model using Conditional Token Reduction and Mixture of Multi-Modal Experts

Yimu Wang*, Mozhgan Nasr Azadani*, Sean Sedwards, Krzysztof Czarnecki

University of Waterloo, Canada

{yimu.wang, mnasraza, sean.sedwards, k2czarne}@uwaterloo.ca


Abstract

Redundancy of visual tokens in multi-modal large language models (MLLMs) significantly reduces their computational efficiency. Recent approaches, such as resamplers and summarizers, have sought to reduce the number of visual tokens, but at the cost of visual reasoning ability. To address this, we propose LEO-MINI, a novel MLLM that significantly reduces the number of visual tokens and simultaneously boosts visual reasoning capabilities. For efficiency, LEO-MINI incorporates CoTR, a novel token reduction module to consolidate a large number of visual tokens into a smaller set of tokens, using the similarity between visual tokens, text tokens, and a compact learnable query. For effectiveness, to scale up the model’s ability with minimal computational overhead, LEO-MINI employs MMoE, a novel mixture of multi-modal experts module. MMoE employs a set of LoRA experts with a novel router to switch between them based on the input text and visual tokens instead of only using the input hidden state. MMoE also includes a general LoRA expert that is always activated to learn general knowledge for LLM reasoning. For extracting richer visual features, MMoE employs a set of vision experts trained on diverse domain-specific data. To demonstrate LEO-MINI’s improved efficiency and performance, we evaluate it against existing efficient MLLMs on various benchmark vision-language tasks.


1 Introduction

The development of multi-modal large language models (MLLMs) (Azadani et al., 2025; Shi et al., 2024; Liu et al., 2024c; Zhang et al., 2025; Dai et al., 2023; Lu et al., 2024; Wang et al., 2025a) has been significantly advanced by aligning vision models (Fang et al., 2024c; Liu et al., 2022; Lee et al., 2023) with large-scale pre-trained language models

*Equal contribution



(a) Improved effectiveness compared with token-reduction MLLMs.



(b) Improved efficiency of LEO-MINI with general and MoVE-based MLLMs.

Figure 1: Performance overview of LEO-MINI-Llama3-8B. **a:** Comparison between LEO-MINI (64 tokens) and existing token reduction MLLMs, where LEO-MINI outperforms all models. **b:** Comparison with other MLLMs on 11 vision-language tasks using Llama3-8B as the LLM. LEO-MINI achieves competitive performance while using only 64 visual tokens.

(LLMs) (Llama Team, 2024; Chiang et al., 2023). MLLMs, such as the LLaVA (Liu et al., 2024c,b), BLIP (Dai et al., 2023), and InternVL (Chen et al., 2024c), embed image patches into visual tokens through a vision expert (Radford et al., 2021; Kirillov et al., 2023; Fang et al., 2024c). Then, those visual tokens are input into the LLM for reason-

ing. This has led to strong performance in image and video understanding tasks (Fu et al., 2024; Liu et al., 2024d; Li et al., 2024a; Hudson and Manning, 2019; Lu et al., 2022; Yue et al., 2024; Li et al., 2023; Kembhavi et al., 2016), bridging the gap between vision and language models.

However, the substantial computational requirement of MLLMs presents a significant challenge to their efficiency. In MLLMs, the LLM predominantly drives computational costs, as the vision expert is smaller in comparison. For example, the commonly used vision expert, CLIP-L (Radford et al., 2021), has 0.3 billion parameters, whereas LLMs such as LLaMA (Llama Team, 2024) or Vicuna (Chiang et al., 2023), have 7–8 billion and 13 billion parameters, respectively. While the vision expert is relatively lightweight, its output, i.e., the visual tokens, which is fed into the LLM along with text instruction tokens, significantly increases the computational overhead. For instance, CLIP-L (Radford et al., 2021) encodes a single image into $24 \times 24 = 576$ visual tokens, whereas textual instructions typically consist of fewer than 100 tokens. And this becomes more challenging in high-resolution image understanding (Chen et al., 2024c; Azadani et al., 2025) or video understanding (Li et al., 2024b; Xu et al., 2016; Heilbron et al., 2015), which requires either more visual tokens per image or processing multiple images.

Reducing the number of visual tokens can thus be an effective strategy for enhancing the efficiency of MLLMs, either through training an efficient compressor (Zhang et al., 2025; Li et al., 2024c) or by using a training-free summarizer (Chen et al., 2024a; Zhang et al., 2024; Wen et al., 2025a). In this work, we focus on training-based methods. Recent training-based approaches (Zhang et al., 2025; Li et al., 2024c,d; Shang et al., 2024) reduce the number of visual tokens by selecting only the most informative ones (Chen et al., 2024a), rather than using all visual tokens. Notably, LLaVA-Mini (Zhang et al., 2025) achieves comparable performance to the full LLaVA-1.5 (Liu et al., 2024b), while using only a single visual token. However, aggressively reducing visual tokens may result in the loss of essential visual information, potentially degrading the model’s performance.

To extract informative visual features with improved efficiency, in this paper, we propose LEO-MINI, a new MLLM that incorporates a novel conditional token reduction module (CoTR) for increased efficiency and a novel mixture of multi-

modal experts (MMOE) module for greater effectiveness.

Efficiency. As the number of visual tokens has become a major bottleneck for MLLM efficiency, LEO-MINI introduces CoTR to reduce the number of tokens fed into the LLM. To focus on the most informative visual tokens based on the input instructions, CoTR aggregates visual tokens into a smaller set of consolidated tokens by their similarity to both visual tokens from other vision experts and text tokens. Moreover, a learnable query is employed to control the length of consolidated visual tokens, which can be adjusted according to the task or computational requirements. This significantly reduces the number of visual tokens, leading to improved training and inference efficiency.

Effectiveness. For a better understanding of visual features and improved reasoning ability, LEO-MINI incorporates MMOE, a novel mixture of multi-modal experts module consisting of MMOE-LLM and MMOE-Vision. Instead of conducting a full finetuning of the entire model after the pretraining, MMOE-LLM employs a mixture of LoRA experts (Hu et al., 2022; Dou et al., 2024; Wu et al., 2023; Tian et al., 2024) with a novel router and a general expert. In contrast to previous work (Hu et al., 2022; Dou et al., 2024; Wu et al., 2023; Tian et al., 2024) whose routers only take the hidden state as input to switch between experts, our router takes the text tokens and the visual tokens as additional input. This facilitates more effective switching between different LoRA experts. The general expert is continuously activated to learn general knowledge. To extract more informative visual features, MMOE-Vision incorporates multiple vision experts (Radford et al., 2021; Kirillov et al., 2023; Fang et al., 2024c), each trained on data from different domains. This boosts the model’s ability to understand visual information, leading to improved performance on vision-language tasks while maintaining minimal computational overhead, as both the vision experts and LoRA experts are substantially smaller than the LLM.

Our contributions can be summarized as follows:

- We propose a new MLLM, LEO-MINI, that incorporates a novel token reduction module (CoTR) to improve efficiency and a novel mixture of multi-modal experts module (MMOE) that increases effectiveness.
- To the best of our knowledge, CoTR is the first to exploit the similarity between visual

tokens from multiple vision experts, text tokens, and a small learnable query to focus on the most informative visual tokens.

- MMoE-LLM employs a novel router taking visual and text tokens as additional input for better switching between different experts, with a general expert for learning general knowledge. MMoE-Vision incorporates multiple vision experts for rich visual feature extraction.
- We demonstrate the effectiveness and improved efficiency of LEO-MINI on various vision-language tasks, as illustrated in Figure 1.

2 Related Work

Multi-modal large language models (MLLMs).

As a fundamental problem in multimodal learning (Fang et al., 2022, 2023, 2024b,a, 2025; Wang et al., 2021, 2020a, 2023, 2020b, 2025b; Liu et al., 2025; Wang and Shi, 2023; Li et al., 2023; Liu et al., 2024d), visual understanding has been revolutionized by recent developments in large language models (LLMs). Advancements in LLMs have fueled significant progress in MLLMs, enabling effective cross-modal reasoning through modality fusion and instruction following (Dai et al., 2023; Lin et al., 2024b; Liu et al., 2024b; Li et al., 2024f; Liu et al., 2024c). Early MLLMs struggled with complex visual understanding due to input resolution limits and the inefficiencies of single vision encoders. To address this, recent research has enhanced visual experts (Chen et al., 2024c; Zhai et al., 2023), incorporated higher-resolution inputs (Li et al., 2024e; Luo et al., 2024), and explored mixtures of vision experts (Azadani et al., 2025; Lu et al., 2024; Kar et al., 2024; Shi et al., 2024; Zong et al., 2024; Fan et al., 2024). Despite the success of these methods, a major challenge remains: the efficiency of MLLMs, as these approaches increase the number of visual tokens, leading to higher computational costs and scalability constraints.

Compressing visual tokens for MLLMs. The efficiency of MLLMs is constrained by the LLM backbone’s context length, as high-resolution images generate numerous vision tokens that quickly consume available space and increase computational cost. To address this challenge, recent methods have focused on reducing the number of visual tokens through both training-free (Chen et al., 2024a;

Huang et al., 2024; Wen et al., 2025b) and training-based (Zhang et al., 2025; Chen et al., 2024a; Li et al., 2024d; Shang et al., 2024) token reduction strategies. Focusing on training-based approaches, some models aggregate tokens based on visual feature similarities (Shang et al., 2024) or high-low resolution similarities (Li et al., 2024c), while others use attention distillation (Ye et al., 2024). MQT-LLaVA (Hu et al., 2024) employs a query transformer to process a random subset of latent query tokens per step. However, direct compression may lead to information loss. LLaMA-VID (Li et al., 2024d) integrates text tokens as contextual information and applies average pooling for efficient token reduction. More recently, LLaVA-Mini (Zhang et al., 2025) mitigates this by combining query-based reduction with prefusion of visual and text tokens.

Existing token reduction approaches (Zhang et al., 2025; Li et al., 2024d; Hu et al., 2024) tend to select tokens by a learnable query or saliency maps. To the best of our knowledge, we propose the first token reduction method that uses text tokens and visual tokens from other visual experts as context to perform attention and reduction over the current vision expert’s tokens.

Mixture of Experts (MoE) in LLMs. MoE is a model design that exploits multiple sparse experts to process different parts of the input space (Jacobs et al., 1991). Early works (Du et al., 2022; Fedus et al., 2022) demonstrated that sparse expert activation improves scalability and computational efficiency. Existing MoE-based LLMs (Jiang et al., 2024; Dai et al., 2024) typically incorporate MoE by replacing standard feed-forward networks with MoE layers, where each token is routed to a small subset of experts. More recent research (Liu et al., 2024a; Lin et al., 2024a; Zadouri et al., 2023; Wu et al., 2024) explores integrating MoE with LoRA to further reduce the parameter overhead of traditional MoE models. These methods make use of LoRA’s ability to fine-tune only a small subset of parameters (Wu et al., 2023; Dou et al., 2024; Chen et al., 2024b), enabling efficient expert selection and dynamic task adaptation.

In contrast to existing LoRA-based MoEs, which select experts based on the hidden state, our proposed MMoE-LLM employs a novel routing network that takes visual tokens and textual instructions as additional inputs, enabling expert selection based on multi-modal input. Moreover, MMoE-LLM employs a general expert to learn general




Figure 2: The overview of the proposed LEO-MINI. LEO-MINI is designed with MMoE (Section 3.3) to enhance visual comprehension and CoTR (Section 3.2) to reduce the number of visual tokens for efficiency.

knowledge.

3 Methodology

In Section 3.1, we introduce the overall framework of LEO-MINI. In Section 3.2, we present our proposed token reduction module, CoTR, which consolidates a large number of visual tokens into a smaller, more informative set. Finally, in Section 3.3, we introduce our mixture of multi-modal experts module, MMoE, which is designed to enhance the efficiency of fine-tuning MLLMs while preserving their strong performance.

3.1 Architecture

The overall architecture of LEO-MINI is presented in Figure 2. It follows the general design (vision expert-projector-LLM) of existing MLLMs (Shi et al., 2024; Azadani et al., 2025; Zhang et al., 2025), while incorporating a mixture of multi-modal experts and a visual feature compression module.

Specifically, MMoE introduces *multiple vision experts*, each of which is trained on a domain-specific vision task to extract diverse and informative visual features from the input image. These experts embed the input image into a group of visual tokens $\{I_i \in \mathcal{R}^{N_i^V \times d_i^V}\}_{i \in [m]}$, where m is the number of vision experts, N_i^V is the number of visual tokens generated by the i -th vision expert, and d_i^V is the feature dimension.

Visual feature compression is then applied to the group of visual tokens $\{I_i\}_{i \in [m]}$. First, as the visual tokens generated by different vision experts have different lengths, i.e., $\{N_i^V\}_{i \in [m]}$, the CoTR mod-




Figure 3: The overview of the proposed CoTR. The CoTR module takes a group of visual tokens, a learnable query, and text tokens as input, and outputs a consolidated group of visual tokens to reduce the number of visual tokens.

ule (Section 3.2) projects them into a group of consolidated visual tokens $\{\tilde{I}_i \in \mathcal{R}^{N^V \times d_i^V}\}_{i \in [m]}$ with the same length of N^V . N^V is much smaller than $\sum_{i \in [m]} N_i^V$, significantly reducing the number of visual tokens for efficiency. Then, the consolidated visual tokens are concatenated channel-wise to form the concatenated visual tokens $\tilde{I} \in \mathcal{R}^{N^V \times d^V}$, where $d^V = \sum_{i \in [m]} d_i^V$.

After that, a *visual projector* is applied to project the concatenated visual token $\tilde{I}$ to have the same dimension with the language model input, resulting in $\tilde{I} \in \mathcal{R}^{N^V \times d_{LLM}}$, where d_{LLM} is the feature dimension of the LLM’s input.

An LLM $f_{LLM}(\cdot)$ then takes the visual tokens $\tilde{I}$ and the textual instruction tokens T as input to generate the instruction-following response $Y = \{y_i\}_{i \in [L]}$ as,

$$p(Y|\tilde{I}, T) = \prod_{i=1}^L p(y_i|\tilde{I}, T, y_{<i}), \quad (1)$$

where L is the length of the response, and $y_{<i}$ is the previous tokens of y_i .

3.2 Conditional Token Reduction (CoTR)

The CoTR module, illustrated in Figure 3, is a query-based module that takes a group of visual tokens $\{I_i\}_{i \in [m]}$, text tokens T , and a group of query tokens $\{Q_i \in \mathcal{R}^{N^V \times d_i^V}\}_{i \in [m]}$ as input, and outputs a group of consolidated visual tokens $\{\tilde{I}_i\}_{i \in [m]}$. Specifically, for the visual tokens I_i generated by the i -th vision expert, the CoTR module computes an attention score $\alpha_i \in \mathcal{R}^{N^V \times N_i^V}$ using the query token $Q_i \in \mathcal{R}^{N^V \times d_i^V}$, the text tokens T , the visual tokens from other vision experts $\{I_j\}_{j \in [m] \setminus \{i\}}$, and the visual token I_i , as

$$s_i^{QUERY} = \hat{Q}_i \hat{I}_i^T \in \mathcal{R}^{N^V \times N_i^V}, \quad (2)$$




Figure 4: The overview of the proposed MMoE-LLM for LLM finetuning. MMoE (language) consists of a set of LoRA experts and a routing network that selects the appropriate expert based on the input visual tokens and textual instructions. Moreover, a general expert is employed to learn the general knowledge and improve the robustness of the model.

$$s_i^{\text{SELF}} = \mathbf{1} \hat{I}_i \hat{I}_i^\top \in \mathcal{R}^{1 \times N_i^V}, \quad (3)$$

$$s_i^{\text{CROSS}} = \sum_{j \in [m] \setminus \{i\}} \hat{I}_j \hat{I}_i^\top \in \mathcal{R}^{1 \times N_i^V}, \quad (4)$$

$$s_i^{\text{TEXT}} = \mathbf{1} \hat{T} \hat{I}_i^\top \in \mathcal{R}^{1 \times N_i^V}, \quad (5)$$

$$\alpha_i = \frac{\text{softmax}(s_i^{\text{QUERY}} + s_i^{\text{SELF}} + s_i^{\text{CROSS}} + s_i^{\text{TEXT}})}{\sqrt{d_i^V}}, \quad (6)$$

where $\hat{Q}_i$, $\hat{I}_i$, $\forall i \in [m]$, and $\hat{T}$ denote the query token, visual tokens, and text tokens projected by learnable linear projections, respectively. The term $\mathbf{1}$ is a vector of ones used to compute the self-attention score. The attention score α_i is then used to compute the consolidated visual tokens $\bar{I}_i$ as,

$$\bar{I}_i = \alpha_i \hat{I}_i \in \mathcal{R}^{N^V \times d_i^V}. \quad (7)$$

In this way, redundant visual tokens are aggregated into a compact set of consolidated visual tokens which significantly improves the efficiency. The length of the query tokens can be adjusted to control the number of visual tokens according to the specific task requirements. Moreover, as we use four different similarities, COTR can capture the complex relationships between multiple sources of features, leading to more informative consolidated visual tokens.

Finally, the consolidated visual tokens $\{\bar{I}_i\}_{i \in [m]}$ are concatenated channel-wise to form the concatenated visual tokens $\bar{I} \in \mathcal{R}^{N^V \times d^V}$, where $d^V = \sum_{i \in [m]} d_i^V$. These concatenated tokens are then projected using a vision projector to have the same dimension as the LLM input size, resulting in $\tilde{I} \in \mathcal{R}^{N^V \times d_{\text{LLM}}}$.

3.3 Mixture of Multi-modal Experts (MMoE)

Our mixture of multi-modal experts module, MMoE, incorporates multiple vision experts to

boost visual understanding and multiple LoRA language experts to enhance reasoning. MMoE comprises MMoE-Vision and MMoE-LLM.

Effective visual comprehension (MMoE-Vision). As described in Section 3.1, drawing inspiration from previous work (Shi et al., 2024; Azadani et al., 2025), visual tokens are generated by multiple vision experts, each extracting informative features from different perspectives to enrich visual understanding.

Effective reasoning ability (MMoE-LLM). To ensure efficient training with minimal computational overhead, we introduce MMoE for LLM tuning, following the mixture of LoRA experts (Hu et al., 2022), as shown in Figure 4. The vanilla mixture of LoRA experts consists of a set of experts $\{f_i^E(\cdot)\}_{i \in [E]}$, and a routing network $f^{\text{ROUTING}}(\cdot)$ that outputs the routing probability $R \in \mathcal{R}^E$ taking the hidden state as input, where E is the number of experts. Then, based on the routing probability, only the top-1 expert will be activated.

Different from the vanilla version, the routing network in MMoE-LLM takes the visual tokens $\tilde{I}$, the textual instruction tokens T , and the hidden state x as the input to compute the routing probability $R = \text{softmax}(f^{\text{ROUTING}}(\tilde{I}, T, x)) \in \mathcal{R}^E$, which facilitates better switching between the experts. Moreover, MMoE-LLM also employs a general expert $f_{\text{GEN}}^E(\cdot)$ to capture the general knowledge and improve the overall robustness of the model.

We select k experts with the highest routing probabilities, i.e., $E' = \text{Top}_k(R)$, and compute the output of the MMoE-LLM as,

$$\text{MMoE-LLM} = f_{\text{GEN}}^E(x) + \sum_{i \in E'} f_i^E(x)/k. \quad (8)$$

With the original linear layer $f^{\text{ORI}}(\cdot)$, the final

Model	# of Visual Tokens	General					OCR TextVQA	Knowledge		
		MME ^P	MMBench	SEED ^I	GQA	VizWiz		SQA	POPE	MMMU
VoCo-LLaMA (Ye et al., 2024)	1	1323.3	58.8	53.7	57.0	-	-	65.4	81.4	-
LLaMA-VID (Li et al., 2024d)	2	-	-	-	55.5	-	49.0	67.7	83.1	-
PruMerge (Shang et al., 2024)	32	1350.3	60.9	-	-	-	56.0	72.0	76.3	-
MQT-LLaVA (Hu et al., 2024)	64	1464.3	63.5	-	60.0	51.5	-	67.0	83.6	<u>34.4</u>
Token-Packer (Li et al., 2024c)	64	-	64.1	-	61.1	50.7	-	-	86.3	-
LLaVA-Mini (Zhang et al., 2025)	64	1476.8	67.5	60.2	61.8	<u>58.5</u>	59.1	69.7	85.3	-
LEO-MINI-Vicuna-7B	64	<u>1542.6</u>	<u>67.8</u>	<u>73.2</u>	<u>64.0</u>	51.4	<u>70.1</u>	<u>73.3</u>	<u>90.0</u>	34.1
LEO-MINI-Llama-8B	64	1583.0	77.0	75.8	64.5	69.3	75.1	84.5	90.3	38.8

Table 1: Comparison to token reduction methods on general, OCR, and knowledge-based tasks. Best in **Bold**. Second best in Underline.

Model	MME ^P	SEED ^I	GQA	SQA	MMMU	POPE	AI2D	TextVQA	ChartQA	OCRBench
LEO-MINI (64 tokens)	1583	75.8	64.5	84.5	38.8	90.3	75.7	75.1	80.5	62.4
w/ 1 visual token	1565	74.6	64.2	83.6	37.7	89.0	74.5	73.5	80.0	61.7
w/ 16 visual token	1577	75.4	64.4	84.3	38.2	90.1	75.4	74.2	80.2	<u>62.6</u>
w/ 256 visual token	<u>1584</u>	<u>76.1</u>	64.2	<u>85.5</u>	<u>39.0</u>	<u>90.7</u>	75.5	<u>75.5</u>	80.5	<u>63.0</u>
w/ 1.8m SFT data	1548	<u>76.6</u>	64.1	<u>85.9</u>	<u>39.4</u>	<u>90.3</u>	<u>77.9</u>	<u>75.7</u>	<u>80.9</u>	<u>63.1</u>

Table 2: Ablation studies of LEO-MINI. We explore the impact of the number of visual tokens, the amount of training data, and the MMoe. Numbers in green indicate the performance is improved compared to LEO-MINI (64 tokens).

output is computed as,

$$\underbrace{f^{\text{ORI}}(x)}_{\text{Original Linear}} + \underbrace{f^{\text{E}}_{\text{GEN}}(x)}_{\text{General Expert}} + \underbrace{\sum_{i \in E'} f^{\text{E}}_i(x)/k}_{\text{Selected Experts}}. \quad (9)$$

4 Experiments

4.1 Implementations and Benchmarks

Models. We use Vicuna-v1.5-7B (Chiang et al., 2023) and Llama3-8B (Llama Team, 2024) as the LLM. For vision experts, we follow the general design of EAGLE (Shi et al., 2024) and use CLIP (Radford et al., 2021), ConvNeXt (Liu et al., 2022), Pix2Struct (Lee et al., 2023) and EVA-02 (Fang et al., 2024c) for LEO-MINI-Llama-8B. Similarly, we add another vision expert SAM (Kirillov et al., 2023) for LEO-MINI-Vicuna-7B. The visual projector is a 2-layer MLP with the GELU activation function (Hendrycks and Gimpel, 2023). MMoe-LLM is only applied to the MLP in each block of the LLM. We use 3 experts for MMoe-LLM with $k = 1$ and 1 general expert being consistently activated. Each expert is a LoRA block (Hu et al., 2022) with rank of 16. The routing network is a 2-layer MLP with the GELU activation function.

More details on training and evaluation are deferred to the Appendix.

4.2 Main Results

We compare our LEO-MINI with several state-of-the-art token reduction MLLMs (table 1). LEO-MINI outperforms all the token reduction MLLMs on all tasks. Specifically, LEO-MINI with Llama3-8B achieves 1583.0 on MME (perception), 77.0 on MMBench, 75.8 on SeedBench, 64.5 on GQA, 69.3 on VizWiz, 75.1 on TextVQA, 84.5 on SQA, 90.3 on POPE, and 38.8 on MMMU, outperforming the best baseline by a large margin. This demonstrates the effectiveness of LEO-MINI in improving the performance of MLLMs on various tasks. Moreover, LEO-MINI with Vicuna-7B also outperforms the best baseline on all tasks except for GQA and MMMU, showing the generalization ability of LEO-MINI with different LLMs.

4.3 Ablation Studies

In this section, we conduct ablation studies to analyze the effectiveness of different components in LEO-MINI, the number of visual tokens, and the amount of SFT data using LEO-MINI-Llama-8B. Ablations using other token reduction modules and MoE are presented in the Appendix.

Is 1 visual token enough for MLLMs? To understand whether 1 visual token is enough for representing the visual information, we conduct the experiments and present the results in Table 2. LEO-MINI with 1 visual token ($N^V = 1$) shows a slight




Figure 5: Visualization of the expert choice using LEO-MINI-Llama-8B on TextVQA, ScienceQA, and GQA. Best viewed in color. “E” refers to Experts.

decrease across some metrics compared to LEO-MINI with 64 tokens, which is reasonable. For example, in the MME^P , the score dropped from 1583 to 1565, and in $SEED^I$, the performance slightly decreased from 75.8 to 75.6. Similarly, in the domains of GQA and SQA, the scores decrease from 64.5 to 64.2 and 84.5 to 83.6, respectively. This trend continues across other evaluated areas such as MMMU, POPE, AI2D, TextVQA, ChartQA, and OCRBench. While the use of only 1 visual token consistently leads to lower performance, it greatly improves model’s efficiency as shown in Table 4. To further understand how increasing the number of visual tokens impacts the performance, we conduct experiments with 16 and 256 visual tokens. The results show that the performance on most of the benchmarks is improved as the number of visual tokens increases. This is reasonable as more visual tokens can provide more detailed visual information to the model, which can help the model to better understand the visual information. However, we also observe that the performance on some benchmarks is slightly decreased, e.g., AI2D, which might be due to the overfitting issue.

Will more training data help on summarizing the visual information? To understand the impact of the amount of training data on the performance, we conduct the experiment with EAGLE-1.8M SFT data (Shi et al., 2024) in Table 2 with 64 tokens. The results show that the performance is improved on most of the benchmarks, i.e., $SeedBench^I$, SQA, MMU, POPE, AI2D, TextVQA, ChartQA, and OCRBench. For example, the performance on $SeedBench^{Image}$ is improved from 75.8 to 76.6, and the performance on SQA is improved from 84.5 to 85.9. This indicates that more training data leads to better performance.

How does MMOE-LLM switch between different LORA experts? We visualize the expert

choice in fig. 5 with GQA (General), TextVQA (OCR), and SQA (Knowledge). We observe that the model effectively switches between different LoRA experts based on the input data. For example, in the SQA task, the model mainly activates expert 2 at layer 0, while for TextVQA and GQA, the model evenly activates three experts. When it comes to the final layer, such as layer 29, the model evenly activates experts 0 and 1 for TextVQA and GQA, while for SQA, the model mainly activates expert 1, respectively. This indicates that MMOE-LLM effectively utilizes the multi-modal input instructions to switch between different experts. We also compare the expert choice between **LEO-MINI** and the **vanilla MoE**. Results are presented in Appendix A.3.

4.4 How Does LEO-MINI Compare to General MLLMs?

We also compare LEO-MINI to general and MoVE-based MLLMs (Table 3). With Vicuna-7B, LEO-MINI outperforms all the general and MoVE-based MLLMs on MME, POPE, AI2D, and OCRBench by 15, 1.2, 2.6, and 2.7 points, respectively, with only 64 visual tokens. LEO-MINI with Vicuna-7B also achieves the second best performance on $SeedBench^I$, SQA, and TextVQA. On the other side, with Llama3-8B, LEO-MINI achieves stronger performance as it outperforms all the general and MoVE-based MLLMs on MM-Bench, SQA, and CharQA by 1.1, 0.2, and 0.4 points, respectively, while reducing the number of visual tokens by 97.77% compared to LLaVA-NeXt and Mini-Gemini-HD. Moreover, compared to the general MLLMs, LEO-MINI with Llama3-8B achieves comparable performance on other benchmark datasets as the discrepancy is marginal.

LLM	Model	# V tokens	MME ^P	MMBench	SEED [†]	GQA	SQA	MMMU	POPE	AI2D	VizWiz	TextVQA	DocVQA	ChartQA	OCRBench
Vicuna-7B & Qwen-7B	InstructBLIP (Dai et al., 2023)	32	-	-	-	49.2	60.5	-	-	-	34.5	50.1	-	-	-
	LLaVA-1.5 (Liu et al., 2024b)	576	1511	64.3	66.1	62.0	66.8	-	85.9	-	50.0	58.2	-	-	-
	LLaVA-NeXT (Liu et al., 2024c)	2880	1519	-	-	64.2	70.1	35.1	-	66.6	57.6	64.9	74.4	54.8	-
	InternVL (Chen et al., 2024c)	1792	1525	64.3	-	62.9	-	-	86.4	-	52.5	57.0	-	-	-
	VILA (Lin et al., 2024b)	576	1533	68.9	61.1	62.3	68.2	-	85.5	-	57.8	64.4	-	-	-
	Monkey (Li et al., 2024f)	256	-	-	-	60.7	69.4	-	-	62.6	61.2	67.6	66.5	65.1	-
	MoVE-based MLLMs														
	LLaVA-HR (Luo et al., 2024)	1024	1554	-	64.2	64.2	65.1	-	87.6	-	48.7	67.1	-	-	-
	Brave-X5 (Kar et al., 2024)	160	-	-	52.7	-	-	-	87.6	-	54.2	-	-	-	-
	Mini-Gemini (Li et al., 2024e)	576	1523	65.8	-	64.5	71.1	36.1	-	-	-	65.2	-	-	-
	Mousi-X3 (Fan et al., 2024)	576	-	66.8	66.0	63.3	70.2	-	87.3	-	-	58.0	-	-	-
	LEO (Azadani et al., 2025)	512	-	72.9	72.2	64.8	78.5	36.4	88.0	69.6	<u>57.9</u>	68.8	80.1	71.0	-
	DeepSeek-VL (Lu et al., 2024)	576	-	73.2	70.4	-	-	36.6	88.1	-	-	-	-	-	45.6
	Eagle-X5 (Shi et al., 2024)	1024	1528	68.4	73.9	64.9	69.8	36.3	88.8	-	54.4	71.2	<u>78.6</u>	<u>67.8</u>	<u>52.9</u>
Llama-8B	LEO-MINI-Vicuna-7B	64 (<u>↓97.77%</u>)	1543	67.8	<u>73.2</u>	64.0	<u>73.3</u>	34.1	90.0	72.2	51.4	<u>70.1</u>	75.3	66.8	55.6
	Cambrian-1 (Tong et al., 2024)	576	1547	75.9	74.7	64.6	80.4	42.7	-	73.0	-	71.7	77.8	73.3	62.4
	LLaVA-NeXT (Liu et al., 2024c)	2880	<u>1604</u>	72.5	72.7	65.2	72.8	41.7	-	71.6	-	64.6	72.8	69.5	49.0
	MoVE-based MLLMs														
	Mini-Gemini-HD (Li et al., 2024e)	2880	1606	72.7	73.2	64.5	75.1	37.3	-	73.5	-	70.2	74.6	59.1	47.7
	Eagle-X4-Plus (Shi et al., 2024)	1024	1559	75.9	76.3	64.9	84.3	43.4	-	76.1	-	77.1	86.6	80.1	62.6
	LEO-MINI-Llama-8B	64 (<u>↓97.77%</u>)	1583	77.0	<u>75.8</u>	64.5	84.5	38.8	90.3	<u>75.7</u>	69.3	<u>75.1</u>	<u>86.3</u>	80.5	<u>62.4</u>

Table 3: Comparison to general and MoVE (mixture of vision experts)-based MLLMs. Best in **Bold**. Second best in Underline.




Figure 6: Qualitative results across four vision-language tasks demonstrating LEO-MINI’s detailed visual understanding. The images are taken from TextVQA (Singh et al., 2019), DocVQA (Mathew et al., 2021), and MMMU (Yue et al., 2024).

4.5 Qualitative Analysis

To understand the detailed visual understanding of LEO-MINI, we conduct a case study for qualitative analysis on four vision-language tasks (Singh et al., 2019; Mathew et al., 2021; Yue et al., 2024) as shown in Figure 6. We use LEO-MINI-Llama3-8B with 64 tokens.

Though the model only takes 64 visual tokens, LEO-MINI performs effectively in capturing visual details, such as accurately identifying the numbers and small book title in OCR. Moreover, LEO-MINI is able to understand the order and count numbers, as for document VQA, LEO-MINI precisely finds the correct results. For science and math questions, LEO-MINI also shows incredible reasoning ability. For science, LEO-MINI can effectively translate the computation diagram into mathematical equations. For math and accounting, LEO-MINI successfully finds the true expression and does calculations correctly.

4.6 Efficiency Analysis

To understand how efficient LEO-MINI is, we compare the number of visual tokens, FLOPs, and CUDA processing time of LEO-MINI with other. We also compare LEO-MINI to MLLMs (Liu et al., 2024c; Shi et al., 2024; Tong et al., 2024) using Llama3-8B. Specifically, we use **Pytorch Profiler** to measure the FLOPs, CUDA time, and GPU memory usage using the figure and “What is shown in this image?” as input. We use one A6000 GPU with 48 GB memory to inference the models.

The comparison is shown in Table 4. LEO-MINI with Llama3-8B is 97.77% more efficient in terms of visual tokens, 63.83% in terms of FLOPs, and 90.69% in terms of CUDA time compared to LLaVA-Next. Moreover, as shown in Figure 1, LEO-MINI outperforms LLaVA-Next by 6% on average performance with only 64 visual tokens. Even compared to the most powerful MLLM with similar performance, *i.e.*, Eagle-X4-Plus, LEO-

Models	# of VT	FLOPs (T)	CUDA Time (s)	GPU Memory (GB)
LLaVA-Next (Liu et al., 2024c)	2880	45.5	7.037	14.74
Eagle-X4-Plus (Shi et al., 2024)	1024	29.4	2.073	20.10
LEO-MINI-Llama-8B	1	13.7	0.593	19.28
LEO-MINI-Llama-8B	64	16.4	0.655	19.34
Δ	$\downarrow 97.77\%$	$\downarrow 63.83\%$	$\downarrow 90.69\%$	

Table 4: Efficiency analysis of LEO-MINI and other MLLMs using Llama3-8B. “VT” represents visual tokens. Δ indicates the difference between LEO-MINI and LLaVA-Next.

MINI is 93.75% more efficient in terms of visual tokens, 44.10% for FLOPs, and 68.40% for CUDA time. More excitingly and importantly, LEO-MINI with only 1 token achieves comparable performance to the state-of-the-art MLLMs with thousands of visual tokens with faster inference time and lower computational cost in terms of FLOPs.

5 Conclusion

In this paper, to address the redundancy of visual tokens in MLLMs, we propose a novel MLLM, LEO-MINI, that significantly reduces the number of visual tokens while boosting visual reasoning capabilities. LEO-MINI incorporates a novel token reduction module, CoTR, to consolidate a large number of visual tokens into a smaller set of tokens, using the similarity between visual tokens, text tokens, and a compact learnable query. However, simply reducing the number of visual tokens leads to an information loss. To avoid the loss and boost the visual comprehension ability with minimal computational overhead, LEO-MINI employs a novel mixture of multi-modal experts module, MMoE, that includes a set of language (LoRA) experts and a set of vision experts trained on diverse domain-specific data. For better switching between different LoRA experts, MMoE employs a new router that takes the text and visual tokens as additional inputs. MMoE also includes a general LoRA expert that is always activated to learn general knowledge. We evaluate LEO-MINI on various vision-and-language tasks, showcasing its potential for practical applications with improved efficiency and performance compared.

6 Limitations

First, the proposed CoTR needs training. Introducing a training-free token reduction module might be a promising direction. Second, the proposed MMoE is designed to be efficient and scalable, but it may not be optimal for all tasks. For vision experts, due to the computational limitations, we

use a fixed set of experts. It would be interesting to explore more vision experts from a wide range of domains and introduce more advanced expert selection mechanisms, such as dynamic vision expert routing. Last, due to the computational limitations, we did not test the efficiency of LEO-MINI using bigger LLMs with 13B or 67B parameters.

7 Potential Risks

LEO-MINI is an MLLM that can be used for various tasks, including visual question answering, image captioning, and document understanding. However, as an inherent risk of (multi-modal) LLMs, LEO-MINI may generate biased or harmful content, especially when the query data contains sensitive information. We recommend that users carefully review the generated content and apply appropriate filters to mitigate potential risks.

Acknowledgement

This work was supported by the Natural Sciences and Engineering Research Council of Canada (NSERC)-CSE Research Community project entitled “An End-to-End Approach to Safe and Secure AI Systems” and NSERC’s Postdoctoral Fellowship. Researchers funded through the NSERC-CSE Research Communities Grants do not represent the Communications Security Establishment Canada or the Government of Canada. Any research, opinions, or positions they produce as part of this initiative do not represent the official views of the Government of Canada.

References

- Mozhgan Nasr Azadani, James Riddell, Sean Sedwards, and Krzysztof Czarnecki. 2025. Leo: Boosting mixture of vision encoders for multimodal large language models. *arXiv preprint arXiv:2501.06986*.
- Liang Chen, Haozhe Zhao, Tianyu Liu, Shuai Bai, Junyang Lin, Chang Zhou, and Baobao Chang. 2024a. An Image is Worth 1/2 Tokens After Layer 2: Plug-and-Play Inference Acceleration for Large Vision-Language Models. In *Computer Vision – ECCV 2024*, pages 19–35.
- Shaoxiang Chen, Zequn Jie, and Lin Ma. 2024b. Llavamole: Sparse mixture of lora experts for mitigating data conflicts in instruction finetuning mllms. *arXiv preprint arXiv:2401.16160*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, and 1 others. 2024c. Internvl:

- Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24185–24198.
- 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. 2023. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality.
- Damai Dai, Chengqi Deng, Chenggang Zhao, RX Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y Wu, and 1 others. 2024. Deepseek-moe: Towards ultimate expert specialization in mixture-of-experts language models. *arXiv preprint arXiv:2401.06066*.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. 2023. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*.
- Shihan Dou, Enyu Zhou, Yan Liu, Songyang Gao, Wei Shen, Limao Xiong, Yuhao Zhou, Xiao Wang, Zhiheng Xi, Xiaoran Fan, Shiliang Pu, Jiang Zhu, Rui Zheng, Tao Gui, Qi Zhang, and Xuanjing Huang. 2024. LoRAMoE: Alleviating World Knowledge Forgetting in Large Language Models via MoE-Style Plugin. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1932–1945.
- Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, and 1 others. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *International Conference on Machine Learning*, pages 5547–5569. PMLR.
- Xiaoran Fan, Tao Ji, Changhao Jiang, Shuo Li, Senjie Jin, Sirui Song, Junke Wang, Boyang Hong, Lu Chen, Guodong Zheng, and 1 others. 2024. Mousi: Poly-visual-expert vision-language models. *arXiv preprint arXiv:2401.17221*.
- Xiang Fang, Wanlong Fang, Wei Ji, and Tat-Seng Chua. 2025. Turing patterns for multimedia: Reaction-diffusion multi-modal fusion for language-guided video moment retrieval. In *ACM International Conference on Multimedia*.
- Xiang Fang, Wanlong Fang, Daizong Liu, Xiaoye Qu, Jianfeng Dong, Pan Zhou, Renfu Li, Zichuan Xu, Lixing Chen, Panpan Zheng, and 1 others. 2024a. Not all inputs are valid: Towards open-set video moment retrieval using language. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 28–37.
- Xiang Fang, Daizong Liu, Wanlong Fang, Pan Zhou, Zichuan Xu, Wenzheng Xu, Junyang Chen, and Renfu Li. 2024b. Fewer steps, better performance: Efficient cross-modal clip trimming for video moment retrieval using language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 1735–1743.
- Xiang Fang, Daizong Liu, Pan Zhou, and Yuchong Hu. 2022. Multi-modal cross-domain alignment network for video moment retrieval. *IEEE Transactions on Multimedia*, 25:7517–7532.
- Xiang Fang, Daizong Liu, Pan Zhou, and Guoshun Nan. 2023. You can ground earlier than see: An effective and efficient pipeline for temporal sentence grounding in compressed videos. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2448–2460.
- Yuxin Fang, Quan Sun, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. 2024c. Eva-02: A visual representation for neon genesis. *Image and Vision Computing*, page 105171.
- William Fedus, Barret Zoph, and Noam Shazeer. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39.
- Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, Yunsheng Wu, and Rongrong Ji. 2024. MME: A Comprehensive Evaluation Benchmark for Multimodal Large Language Models. *arXiv preprint. ArXiv:2306.13394 [cs]*.
- Danna Gurari, Qing Li, Abigale J. Stangl, Anhong Guo, Chi Lin, Kristen Grauman, Jiebo Luo, and Jeffrey P. Bigham. 2018. VizWiz Grand Challenge: Answering Visual Questions from Blind People. In *2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3608–3617.
- Fabian Caba Heilbron, Victor Escorcia, Bernard Ghanem, and Juan Carlos Nieves. 2015. ActivityNet: A large-scale video benchmark for human activity understanding. In *IEEE conference on computer vision and pattern recognition, CVPR 2015, boston, MA, USA, june 7-12, 2015*, pages 961–970.
- Dan Hendrycks and Kevin Gimpel. 2023. [Gaussian error linear units \(gelu\)](#). *Preprint*, arXiv:1606.08415.
- Edward J Hu, yelong shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Wenbo Hu, Zi-Yi Dou, Liunian Harold Li, Amita Kamath, Nanyun Peng, and Kai-Wei Chang. 2024. Matryoshka query transformer for large vision-language models. *arXiv preprint arXiv:2405.19315*.
- Wenxuan Huang, Zijie Zhai, Yunhang Shen, Shaosheng Cao, Fei Zhao, Xiangfeng Xu, Zheyu Ye, and Shao-hui Lin. 2024. [Dynamic-LLaVA: Efficient Multi-modal Large Language Models via Dynamic Vision-language Context Sparsification](#).

- Drew A. Hudson and Christopher D. Manning. 2019. GQA: A New Dataset for Real-World Visual Reasoning and Compositional Question Answering. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6693–6702.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, and 1 others. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Oğuzhan Fatih Kar, Alessio Tonioni, Petra Poklucar, Achin Kulshrestha, Amir Zamir, and Federico Tombari. 2024. Brave: Broadening the visual encoding of vision-language models. *arXiv preprint arXiv:2404.07204*.
- Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. 2016. A Diagram is Worth a Dozen Images. In *Computer Vision – ECCV 2016*, pages 235–251.
- Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C. Berg, Wan-Yen Lo, Piotr Dollar, and Ross Girshick. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4015–4026.
- Kenton Lee, Mandar Joshi, Iulia Raluca Turc, Hexiang Hu, Fangyu Liu, Julian Martin Eisenschlos, Urvasi Khandelwal, Peter Shaw, Ming-Wei Chang, and Kristina Toutanova. 2023. Pix2Struct: Screenshot parsing as pretraining for visual language understanding. In *Proceedings of the 40th International Conference on Machine Learning*, pages 18893–18912.
- Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. 2024a. SEED-Bench: Benchmarking Multimodal Large Language Models. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13299–13308.
- KunChang Li, Yinan He, Yi Wang, Yizhuo Li, Wenhai Wang, Ping Luo, Yali Wang, Limin Wang, and Yu Qiao. 2024b. VideoChat: Chat-Centric Video Understanding.
- Wentong Li, Yuqian Yuan, Jian Liu, Dongqi Tang, Song Wang, Jie Qin, Jianke Zhu, and Lei Zhang. 2024c. Tokenpacker: Efficient visual projector for multimodal llm. *arXiv preprint arXiv:2407.02392*.
- Yanwei Li, Chengyao Wang, and Jiaya Jia. 2024d. Llama-vid: An image is worth 2 tokens in large language models. In *European Conference on Computer Vision*, pages 323–340. Springer.
- Yanwei Li, Yuechen Zhang, Chengyao Wang, Zhisheng Zhong, Yixin Chen, Ruihang Chu, Shaoteng Liu, and Jiaya Jia. 2024e. Mini-gemini: Mining the potential of multi-modality vision language models. *arXiv preprint arXiv:2403.18814*.
- Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Xin Zhao, and Ji-Rong Wen. 2023. Evaluating Object Hallucination in Large Vision-Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 292–305.
- Zhang Li, Biao Yang, Qiang Liu, Zhiyin Ma, Shuo Zhang, Jingxu Yang, Yabo Sun, Yuliang Liu, and Xiang Bai. 2024f. Monkey: Image resolution and text label are important things for large multi-modal models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26763–26773.
- Bin Lin, Zhenyu Tang, Yang Ye, Jiayi Cui, Bin Zhu, Peng Jin, Junwu Zhang, Munan Ning, and Li Yuan. 2024a. Moe-llava: Mixture of experts for large vision-language models. *arXiv preprint arXiv:2401.15947*.
- Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. 2024b. Vila: On pre-training for visual language models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26689–26699.
- Dongxu Liu, Bing Xu, Yinzhuo Chen, Bufan Xu, Wenpeng Lu, Muyun Yang, and Tiejun Zhao. 2024a. Pmol: Parameter efficient moe for preference mixing of llm alignment. *arXiv preprint arXiv:2411.01245*.
- Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. 2024b. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26296–26306.
- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024c. Llava-next: Improved reasoning, ocr, and world knowledge.
- Xuye Liu, Yimu Wang, and Jian Zhao. 2025. ELIOT: Zero-shot video-text retrieval through relevance-boosted captioning and structural information extraction. In *SRW of NAACL*, pages 381–391.
- Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. 2024d. MMBench: Is Your Multi-modal Model an All-Around Player? In *Computer Vision – ECCV 2024*, pages 216–233.
- Yuliang Liu, Zhang Li, Mingxin Huang, Biao Yang, Wenwen Yu, Chunyuan Li, Xu-Cheng Yin, Cheng-Lin Liu, Lianwen Jin, and Xiang Bai. 2024e. **OCR-Bench: on the hidden mystery of OCR in large multi-modal models**. *Science China Information Sciences*, 67(12). Publisher: Springer Science and Business Media LLC.

- Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. 2022. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11976–11986.
- AI @ Meta1 Llama Team. 2024. *The llama 3 herd of models*. Preprint, arXiv:2407.21783.
- Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Hao Yang, and 1 others. 2024. Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*.
- Pan Lu, Swaroop Mishra, Tony Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *The 36th conference on neural information processing systems (NeurIPS)*.
- Gen Luo, Yiyi Zhou, Yuxin Zhang, Xiawu Zheng, Xiaoshuai Sun, and Rongrong Ji. 2024. Feast your eyes: Mixture-of-resolution adaptation for multimodal large language models. *arXiv preprint arXiv:2403.03003*.
- Ahmed Masry, Xuan Long Do, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. 2022. ChartQA: A Benchmark for Question Answering about Charts with Visual and Logical Reasoning. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2263–2279.
- Minesh Mathew, Dimosthenis Karatzas, and C. V. Jawahar. 2021. DocVQA: A Dataset for VQA on Document Images. In *2021 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 2199–2208.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. 2021. Learning transferable visual models from natural language supervision. In *Proceedings of the 38th International Conference on Machine Learning*, pages 8748–8763.
- Samyam Rajbhandari, Jeff Rasley, Olatunji Ruwase, and Yuxiong He. 2020. Zero: memory optimizations toward training trillion parameter models. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press.
- Yuzhang Shang, Mu Cai, Bingxin Xu, Yong Jae Lee, and Yan Yan. 2024. Llava-prumerge: Adaptive token reduction for efficient large multimodal models. *arXiv preprint arXiv:2403.15388*.
- Noam Shazeer, *Azalia Mirhoseini, *Krzysztof Maziarczyk, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. 2017. *Outrageously large neural networks: The sparsely-gated mixture-of-experts layer*. In *International Conference on Learning Representations*.
- Min Shi, Fuxiao Liu, Shihao Wang, Shijia Liao, Subhashree Radhakrishnan, De-An Huang, Hongxu Yin, Karan Sapra, Yaser Yacoob, Humphrey Shi, and 1 others. 2024. Eagle: Exploring the design space for multimodal llms with mixture of encoders. *arXiv preprint arXiv:2408.15998*.
- Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. 2019. Towards VQA Models That Can Read. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8309–8318.
- Chunlin Tian, Zhan Shi, Zhijiang Guo, Li Li, and Chengzhong Xu. 2024. HydraLoRA: An Asymmetric LoRA Architecture for Efficient Fine-Tuning. *arXiv preprint*.
- Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, and 1 others. 2024. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*.
- Yimu Wang, Mozhgan Nasr Azadani, Sean Sedwards, and Krzysztof Czarnecki. 2025a. HAWAII: Hierarchical visual knowledge transfer for efficient vision-language models. In *NeurIPS*.
- Yimu Wang, Xiangru Jian, and Bo Xue. 2023. Balance act: Mitigating hubness in cross-modal retrieval with query and gallery banks. In *EMNLP*, pages 10542–10567.
- Yimu Wang, Shiyin Lu, and Lijun Zhang. 2020a. Searching Privately by Imperceptible Lying: A Novel Private Hashing Method with Differential Privacy. In *ACM International Conference on Multimedia*, pages 2700–2709.
- Yimu Wang and Peng Shi. 2023. Video-text retrieval by supervised sparse multi-grained learning. In *Findings of EMNLP*, pages 633–649.
- Yimu Wang, Xiu-Shen Wei, Bo Xue, and Lijun Zhang. 2020b. Piecewise Hashing: A Deep Hashing Method for Large-Scale Fine-Grained Search. In *Pattern Recognition and Computer Vision*, pages 432–444.
- Yimu Wang, Bo Xue, Quan Cheng, Yuhui Chen, and Lijun Zhang. 2021. Deep Unified Cross-Modality Hashing by Pairwise Data Alignment. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 1129–1135.
- Yimu Wang, Shuai Yuan, Bo Xue, Xiangru Jian, Wei Pang, Mushi Wang, and Ning Yu. 2025b. DREAM: Improving video-text retrieval through relevance-based augmentation using large foundation models. In *NAACL*, pages 3037–3056.

- Zichen Wen, Yifeng Gao, Shaobo Wang, Junyuan Zhang, Qintong Zhang, Weijia Li, Conghui He, and Linfeng Zhang. 2025a. [Stop Looking for Important Tokens in Multimodal Language Models: Duplication Matters More](#). *arXiv preprint*. ArXiv:2502.11494 [cs].
- Zichen Wen, Yifeng Gao, Shaobo Wang, Junyuan Zhang, Qintong Zhang, Weijia Li, Conghui He, and Linfeng Zhang. 2025b. [Stop looking for important tokens in multimodal language models: Duplication matters more](#). *Preprint*, arXiv:2502.11494.
- Jialin Wu, Xia Hu, Yaqing Wang, Bo Pang, and Radu Soricut. 2024. Omni-smola: Boosting generalist multimodal models with soft mixture of low-rank experts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14205–14215.
- Xun Wu, Shaohan Huang, and Furu Wei. 2023. Mixture of LoRA Experts.
- Jun Xu, Tao Mei, Ting Yao, and Yong Rui. 2016. MSR-VTT: A Large Video Description Dataset for Bridging Video and Language. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5288–5296.
- Xubing Ye, Yukang Gan, Xiaoke Huang, Yixiao Ge, Ying Shan, and Yansong Tang. 2024. Voco-llama: Towards vision compression with large language models. *arXiv preprint arXiv:2406.12275*.
- Xiang Yue, Yuansheng Ni, Tianyu Zheng, Kai Zhang, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, and 3 others. 2024. MMMU: A Massive Multi-Discipline Multimodal Understanding and Reasoning Benchmark for Expert AGI. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9556–9567.
- Ted Zadori, Ahmet Üstün, Arash Ahmadian, Beyza Ermiş, Acyr Locatelli, and Sara Hooker. 2023. Pushing mixture of experts to the limit: Extremely parameter efficient moe for instruction tuning. *arXiv preprint arXiv:2309.05444*.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. 2023. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986.
- Qizhe Zhang, Aosong Cheng, Ming Lu, Zhiyong Zhuo, Minqi Wang, Jiajun Cao, Shaobo Guo, Qi She, and Shanghang Zhang. 2024. [\[CLS\] Attention is All You Need for Training-Free Visual Token Pruning: Make VLM Inference Faster](#). *arXiv preprint*. ArXiv:2412.01818 [cs].
- Shaolei Zhang, Qingkai Fang, Zhe Yang, and Yang Feng. 2025. Llava-mini: Efficient image and video large multimodal models with one vision token. *arXiv preprint arXiv:2501.03895*.
- Zhuofan Zong, Bingqi Ma, Dazhong Shen, Guanglu Song, Hao Shao, Dongzhi Jiang, Hongsheng Li, and Yu Liu. 2024. MoVA: Adapting Mixture of Vision Experts to Multimodal Context. *arXiv preprint*. ArXiv:2404.13046 [cs].

A Experiments

A.1 Benchmark Datasets

We evaluate our method on the following benchmark datasets: MME (Fu et al., 2024), MMBench (Liu et al., 2024d), Seed-Bench (Li et al., 2024a), GQA (Hudson and Manning, 2019), SQA (Lu et al., 2022), MMMU (Yue et al., 2024), POPE (Li et al., 2023), AI2D (Kembhavi et al., 2016), VizWiz (Gurari et al., 2018), TextVQA (Singh et al., 2019), DocVQA (Mathew et al., 2021), ChartQA (Masry et al., 2022), and OCRBench (Liu et al., 2024e).

MME (Fu et al., 2024). The MME benchmark is designed to rigorously evaluate a model’s perceptual and cognitive abilities through 14 subtasks. It employs carefully constructed instruction-answer pairs and concise instructions to minimize data leakage and ensure fair evaluation. This setup provides a robust measure of a model’s performance across various tasks.

MMBench (Liu et al., 2024d). MMBench offers a hierarchical evaluation framework, categorizing model capabilities into three levels. The first level (L-1) focuses on perception and reasoning. The second level (L-2) expands this to six subabilities, while the third level (L-3) further refines these into 20 specific dimensions. This structured approach allows for a nuanced and comprehensive assessment of a model’s multifaceted abilities.

Seed-Bench (Li et al., 2024a). SEED-Bench consists of 19K multiple-choice questions with accurate human annotations, covering 12 evaluation dimensions including both the spatial and temporal understanding.

GQA (Hudson and Manning, 2019). GQA is structured around three core components: scene graphs, questions, and images. It includes not only the images themselves but also detailed spatial features and object-level attributes. The questions are crafted to assess a model’s ability to comprehend visual scenes and perform reasoning tasks based on the image content.

ScienceQA (Lu et al., 2022). ScienceQA spans a wide array of domains, including natural, language, and social sciences. Questions are hierarchically categorized into 26 topics, 127 categories, and 379 skills, providing a diverse and comprehensive testbed for evaluating multimodal understanding, multi-step reasoning, and interpretability.

MMMU (Yue et al., 2024). MMMU includes 11.5K meticulously collected multimodal questions

from college exams, quizzes, and textbooks, covering six core disciplines: Art & Design, Business, Science, Health & Medicine, Humanities & Social Science, and Tech & Engineering. These questions span 30 subjects and 183 subfields, comprising 30 highly heterogeneous image types, such as charts, diagrams, maps, tables, music sheets, and chemical structures.

POPE (Li et al., 2023). POPE is tailored to assess object hallucination in models. It presents a series of binary questions about the presence of objects in images, using accuracy, recall, precision, and F1 score as metrics. This approach offers a precise evaluation of hallucination levels under different sampling strategies.

AI2D (Kembhavi et al., 2016). AI2D is a dataset of over 5000 grade school science diagrams with over 150000 rich annotations, their ground truth syntactic parses, and more than 15000 corresponding multiple choice questions.

VizWiz (Gurari et al., 2018). VizWiz consists of over 31,000 visual questions originating from blind people who each took a picture using a mobile phone and recorded a spoken question about it, together with 10 crowdsourced answers per visual question.

TextVQA (Singh et al., 2019). TextVQA emphasizes the integration of textual information within images. It evaluates a model’s proficiency in reading and reasoning about text embedded in visual content, requiring both visual and textual comprehension to answer questions accurately.

DocVQA (Mathew et al., 2021). DocVQA consists of 50,000 questions defined on 12,000+ document images.

ChartQA (Masry et al., 2022). ChartQA is a large benchmark covering 9.6K human-written questions as well as 23.1K questions generated from human-written chart summaries.

OCRBench (Liu et al., 2024e). OCRBench is a comprehensive benchmark for evaluating the OCR capabilities of multi-modal language models across five key tasks: text recognition, scene text-centric and document-oriented VQA, key information extraction, and handwritten mathematical expression recognition.

A.2 Training Details

The training of LEO-MINI consists of three stages, as shown in Table 6. **Stage 1: Warming up for the visual projector.** This stage pre-trains the visual projector while keeping the LLM and vision

Token Reduction	MoE	MME ^P	SEED ^I	GQA	SQA	MMMU	POPE	AI2D	TextVQA	ChartQA	OCRBench
CoTR	MMoE-Vision + MMoE-LLM	1583	75.8	64.5	84.5	38.8	90.3	75.7	75.1	80.5	62.4
MQT	MMoE-Vision + MMoE-LLM	1537	75.3	64.4	84.2	37.0	90.4	75.3	74.2	80.4	61.6
MQT	-	1435	-	61.6	67.6	34.8	84.4	-	-	-	-
CoTR	MMoE-Vision + LoRA-MoE	1521	75.3	64.5	83.4	38.2	90.3	75.4	74.0	80.0	61.7

Table 5: Ablation study on the effectiveness of CoTR and MMoE-LLM. We use Llama3-8B as the backbone.




Figure 7: Visualization of the expert choice using LEO-MINI-Llama-8B on TextVQA, ScienceQA, and GQA. Best viewed in color.

	Stage 1	Stage 2	Stage 3
LLM	❄️	🔥	❄️
Visual Experts	❄️	🔥	❄️
Visual Projector	🔥	🔥	🔥
CoTR	-	-	🔥
MMoE-LLM	-	-	🔥

Table 6: The training stages of LEO-MINI. CoTR and MMoE-LLM are added and fine-tuned in the third stage, while the LLM and vision experts (MMoE-Vision) remain frozen.

experts frozen. The LLM and vision experts are initialized from the base models, while the vision projector is randomly initialized. **Stage 2: Supervised fine-tuning.** In this stage, we fine-tune the entire model, including the LLM, vision experts, and visual projector. **Stage 3: Supervised fine-tuning for reducing the number of visual tokens.** In this stage, we introduce CoTR to the model and perform LoRA fine-tuning (MMoE-LLM) for efficiency. Specifically, CoTR and MMoE-LLM are randomly initialized and fine-tuned, while the LLM and vision experts remain frozen. To avoid the risk of routing collapse (Shazeer et al., 2017), we employ a balanced loss for MMoE with a hyperparameter $\lambda = 0.05$, following previous works (Shazeer et al., 2017; Tian et al., 2024).

LEO-MINI uses the same training data as LLaVA-v1.5 (Liu et al., 2024b) in stage 3 with 665K instruction data. For stage 1 and stage 2, we follow the same data as EAGLE (Shi et al., 2024). The training was conducted on 8 A6000 GPUs (48 GB) using DeepSpeed’s Zero2 strategy (Rajbhan-

dari et al., 2020).

We expect the model to learn a comprehensive understanding of the input images in the first two stages, and then in the third stage, the model is guided to focus on the most important information for improved efficiency.

A.3 Ablation Studies

How do other token reduction methods work?

To understand how our proposed token reduction module, *i.e.*, CoTR, works, we compare it with a representative token reduction method, *i.e.*, MQT-LLaVA (Hu et al., 2024), as shown in Table 5, while keeping the MMoE unchanged. The results show that CoTR outperforms MQT on most benchmarks, demonstrating the effectiveness of our proposed method. Moreover, compared to MQT-LLaVA (the third row), MMoE improves the performance of MQT-LLaVA on most benchmarks, showing the importance of introducing multiple experts for extracting informative tokens.

How does MMoE-LLM improve the performance?

We conduct an ablation study to investigate the effectiveness of MMoE-LLM in Table 5. Results show that MMoE-LLM significantly improves the performance on most benchmarks, compared to the baseline with MMoE-Vision and LoRA-MoE (Wu et al., 2023). MMoE-LLM improves performance by 62, 0.5, 1.1, 0.6, 0.3, 1.1, 0.5, and 0.7 on MME^P, SEEDBench^I, SQA, MMMU, AI2D, TextVQA, ChartQA, and OCRBench, respectively. This demonstrates the effectiveness of MMoE-LLM in better switching between different experts to better understand visual information.




Figure 8: Visualization of the expert choice using LEO-MINI-Llama-8B on TextVQA, ScienceQA, and GQA compared to the vanilla MoE. Best viewed in color.




Figure 9: Visualization of the expert choice using the vanilla MoE on TextVQA, ScienceQA, and GQA. Best viewed in color.

How does MMoE-LLM switch between different LORA experts (full results)? The full results are presented in Figure 7. The visualization shows that MMoE-LLM can effectively switch between different experts to better understand visual information. We notice that for some layers, such as layers 0, 3, 5, 6, 23, 25, 29, and 31, MMoE-LLM selects different experts for different benchmarks, while for some layers, such as layers 1, 2, 17, and 27, the model selects the similar experts for different benchmarks. This demonstrates the effectiveness of MMoE-LLM in better understanding visual information by switching between different experts.

How does MMoE-LLM switch between different LORA experts compared to the vanilla MoE? To understand how MMoE-LLM switches using the additional visual and input signals, we compare it with the vanilla MoE in Figure 8. The full routing results of the vanilla MoE are shown in Figure 9, while the results of MMoE-LLM are shown in Figure 7. The results show that at different layers, MMoE-LLM shows different expert preferences for different benchmarks, while the vanilla MoE shows similar expert preferences for different benchmarks.