

FROM CLIP TO DINO: VISUAL ENCODERS SHOUT IN MULTI-MODAL LARGE LANGUAGE MODELS

Anonymous authors

Paper under double-blind review

ABSTRACT

Multi-modal Large Language Models (MLLMs) have made significant strides in expanding the capabilities of Large Language Models (LLMs) through the incorporation of visual perception interfaces. Despite the emergence of exciting applications and the availability of diverse instruction tuning data, existing approaches often rely on CLIP or its variants as the visual branch, and merely extract features from the deep layers. However, these methods lack a comprehensive analysis of the visual encoders in MLLMs. In this paper, we conduct an extensive investigation into the effectiveness of different vision encoders within MLLMs. Our findings reveal that the shallow layer features of CLIP offer particular advantages for fine-grained tasks such as grounding and region understanding. Surprisingly, the vision-only model DINO, which is not pretrained with text-image alignment, demonstrates promising performance as a visual branch within MLLMs. By simply equipping it with an MLP layer for alignment, DINO surpasses CLIP in fine-grained related perception tasks. Building upon these observations, we propose a simple yet effective feature merging strategy, named **COMM**, that integrates **CLIP** and **DINO** with **Multi-level features Merging**, to enhance the visual capabilities of MLLMs. We evaluate **COMM** through comprehensive experiments on a wide range of benchmarks, including image captioning, visual question answering, visual grounding, and object hallucination. Experimental results demonstrate the superior performance of **COMM** compared to existing methods, showcasing its enhanced visual capabilities within MLLMs.

1 INTRODUCTION

Large Language Models (LLMs) (OpenAI, 2023a;b; Touvron et al., 2023a;b; Taori et al., 2023; Chiang et al.) have made significant strides in the domains of language understanding and generation, achieving remarkable progress recently. Through instruction tuning (Wei et al., 2021; Wang et al., 2022), existing LLMs demonstrate their versatility as general-purpose models capable of handling a wide range of tasks. This capability unlocks their potential zero-shot learning ability, enabling seamless task switching guided by instructions. Building upon the promising performance of LLMs, researchers are now motivated to enhance their capabilities by incorporating visual signals as inputs. This extension allows for the generation of textual outputs that are closely related to visual content, opening up exciting possibilities in the realm of vision-language understanding and generation.

To this end, Flamingo (Alayrac et al., 2022) and BLIP2 (Li et al., 2023c) align the powerful LLMs with a frozen visual encoder to understand visual inputs and perform various vision-language tasks. A series of following works, LLaVA (Liu et al., 2023a), InstructBLIP (Dai et al., 2023), MiniGPT-4 (Zhu et al., 2023) and mPLUG-OWL (Ye et al., 2023) further improve the ability to follow human instructions by constructing multi-modal instruction-following datasets for training. However, these methods are built on image-level alignments, which suffer from the limited fine-grained understanding (such as region description (Liu et al., 2017) and reasoning (Zellers et al., 2019)) and severe object hallucination problem (Li et al., 2023d). To this end, GPT4ROI (Zhang et al., 2023b) proposes instruction tuning on region-of-interest and unlocks the region-level multimodal capacities. Kosmos-2 (Peng et al., 2023) and Shikra (Chen et al., 2023) further integrate the grounding abilities into LLMs and unlock the referential ability in dialogue, *i.e.*, enable the user to point to the object or region as input and the model responds with spatial coordinates of bounding boxes. Such grounding capacity can fulfill numerous vision-language tasks, which is a great progress in MLLMs.

Despite a wide variety of exciting methods and applications, most of existing multi-modal LLMs employ CLIP (Radford et al., 2021) or its variants (Sun et al., 2023) as the visual branch, where the features output from the deep layers (*e.g.*, the penultimate layer) are usually employed as inputs to the language decoders. However, it still lacks analysis that: *Whether using the Vanilla CLIP features as visual encoder is the best way for MLLMs?* Though the visual encoder of CLIP is apparently well aligned with the word embedding space by image-text contrastive learning, it fails to learn more detailed pixel-level information such as color and positioning due to the global supervision of image captions, which might hinder the fine-grained perception ability in MLLMs. Besides, existing MLLMs have quite unbalanced visual and language encoders (*e.g.*, ViT-Large-300M vs. Vicuna-7B/13B). Since the language models have succeeded in scaling up the model size with progressively powerful language abilities, the short plate of the Buckets Effect for MLLMs lies in the visual models, which fails to perform in-context learning, and suffer from domain gap and limited zero-shot ability. Consequently, it is critical to enhance the visual capabilities for boosting MLLMs.

This paper presents an extensive investigation into different visual encoders for MLLMs. Four typical visual foundation models are considered, *i.e.*, image-text contrastive learning CLIP, image-only contrastive learning DINOv2 (Oquab et al., 2023), masked image modeling MAE (He et al., 2022) and supervised learning DeiT (Touvron et al., 2021). We evaluate the performance on commonly-used vision-language tasks including visual grounding, object hallucination, visual question answering, image captioning and MME benchmark. Our analysis reveals that different layers of features exhibit varying biases towards local and global patterns. Shallow layer features containing low-level detailed information prove beneficial for fine-grained perception tasks such as grounding and positioning ability, while deep layer features are superior at global understanding. To enhance representation, we propose a multi-level feature merging strategy that incorporates both low-level and high-level features. Surprisingly, when equipped with an MLP layer for alignment, the vision-only model DINOv2 shows promise as a visual branch for MLLMs. We attribute this to the fine-grained localization information captured by DINOv2. Conversely, MAE and DeiT perform inferiorly as visual branches for MLLMs. MAE learns limited semantic information, while DeiT’s strong supervised training makes the alignment with the textual space challenging. Based on the above observations, we propose a fusion strategy that integrates CLIP and DINO with Multi-level features Merging), dubbed as **COMM**, for boosting the visual branches of MLLMs. Experimental results demonstrate clear advantages of the proposed model over existing approaches and highlight the enhanced visual capabilities brought by **COMM**.

In a nutshell, the contributions of this paper are summarized as follows:

- We are the first to extensively investigate the effectiveness of different visual encoders for MLLMs. Based on the analysis that shallow layer features contain low-level detailed information which is helpful for fine-grained tasks, we propose a multi-level feature fusion strategy to incorporate low-level and high-level features for improving representation.
- Our analysis indicates that vision-only DINOv2 achieves promising results in MLLMs with only an MLP layer for alignment. Considering fine-grained pixel information in DINOv2 and global semantic information in CLIP, we propose **COMM** to fuse the visual embeddings of these two models to enhance visual capabilities for boosting MLLMs.
- Extensive experiments on a wide range of tasks including visual grounding, referring expression generation, object hallucination, visual question answering and image captioning demonstrate the superiority of **COMM** over existing works.

2 RELATED WORK

Multi-modal Large Language Model. LLMs (Dai et al., 2019; Brown et al., 2020) have garnered significant attention in both academia and industry due to their remarkable understanding and generative abilities. The success of LLMs has motivated researchers to explore the integration of vision into these models, leading to the development of powerful multi-modal LLMs (MLLMs). Flamingo (Alayrac et al., 2022) employs a cross-attention module to extract visual contexts, which are concatenated with text token as input for LLMs. LLaVA (Liu et al., 2023b) and FROMAGE (Koh et al., 2023) leverage the vision encoder of CLIP to extract visual features, which is aligned to text features using a single linear layer and then input to LLMs. Models such as BLIP-2 (Li et al., 2023b), mPLUG-OWL (Ye et al., 2023), MiniGPT-4 (Zhu et al., 2023) and InstructBLIP (Dai et al.,




Figure 1: Feature correspondence visualization by computing the cosine similarity of different visual tokens extracted from the shallow and deep layers of CLIP and DINOv2.

2023) employ Q-former to extract text-aligned visual features for LLMs. Recently, some interesting works extend LLMs to image retrieval (Koh et al., 2023), video understanding (Zhang et al., 2023a), audio (Su et al., 2023), biomedical analysis (Li et al., 2023a), control systems Driess et al. (2023).

In recent studies, there has been a growing interest in extending MLLMs to improve their fine-grained understanding abilities through region-level image-text alignment. Kosmos-2 (Peng et al., 2023) addresses this by constructing a large-scale dataset of grounded region-text pairs, enabling the integration of grounding abilities into LLMs. GPT4RoI (Zhang et al., 2023b) reformulates the bounding box as a spatial instruction format and extracts visual features based on region-of-interest, facilitating region-level multi-modal understanding. Shikra (Chen et al., 2023) proposes a unified model that handles spatial coordinates to possess referential abilities in dialogue contexts. Additionally, Qwen (Bai et al., 2023) presents a set of MLLMs that demonstrate remarkable performance across various tasks. However, previous works have predominantly focused on extracting visual features solely from the last few layers of the CLIP model, resulting in an emphasis on global image properties. In this study, we draw attention to the fact that features extracted from shallower layers exhibit a stronger focus on localized properties, which we argue can be more potent in comprehending object locations and image details. Additionally, while CLIP primarily learns globally aligned features, advanced vision-alone models such as DINOv2 excel in capturing more fine-grained vision features. We posit that leveraging these fine-grained vision features can effectively enhance the capabilities of MLLMs, as demonstrated in our analysis. To further advance this line of research, we introduce a novel fusion module that expands and enhances the visual branches, thereby aiming to significantly improve the performance of MLLMs.

Large Vision Foundation Model. Recent progresses in training vision foundation models with large-scale image data focus on contrastive learning, masked image modeling and supervised training. For one thing, contrastive learning can be conducted in an image-only or image-text manner. DINOv2 (Oquab et al., 2023) pretrains the image encoder on large curated image data, which shows a superior understanding of object parts and scene geometry across image domains. Image-text contrastive learning as CLIP (Radford et al., 2021) and EVA-CLIP (Sun et al., 2023) employs the natural language as weak supervision to guide the learning of visual features. For another, BEiT (Bao et al., 2021) predicts discrete tokens based on a pre-trained image tokenizer while iBOT (Zhou et al., 2021) proposes an online image tokenizer. MAE (He et al., 2022) proposes a masked autoencoder for reconstructing image pixels. Besides, DeiT III (Touvron et al., 2022) proposes a training recipe to achieve promising performance. Recent MLLMs employ the vision encoder of CLIP/EVA-CLIP without considering the properties of specific visual models. In this paper, we are the first to re-examine the effectiveness of existing visual models in MLLMs and propose a simple yet effective fusion strategy for boosting visual capabilities.

3 ANALYSIS OF THE VISUAL BRANCH IN MLLMS

Previous MLLMs usually utilize the vision encoder of CLIP as their visual branch. Typically, these models extract features from the last few layers, such as the penultimate layer, which are then fed into an alignment network. Subsequently, the aligned features are concatenated with text tokens to serve as input for the LLMs. While the image-text pretraining of CLIP aligns well with the language model, it primarily learns image-level features but overlooks the richer pixel-level features due to the constraint of limited fine-grained information in the guided captions. Moreover, the deep-layer features primarily focus on global image properties and inadequately explore the intricate details of local object parts. As depicted in Fig. 1, the visual features extracted from the shallow layers of CLIP and the deep visual features obtained from the visual-only model DINOv2 contain more




Figure 2: Average REC, POPE accuracy and REG CIDEr for using different layers of features extracted from various vision models (CLIP, DINOv2 and MAE), as input to MLLMs. Shikra uses the 23rd layer features of CLIP and we reproduce its results with fewer iterations (denoted as Shikra*).

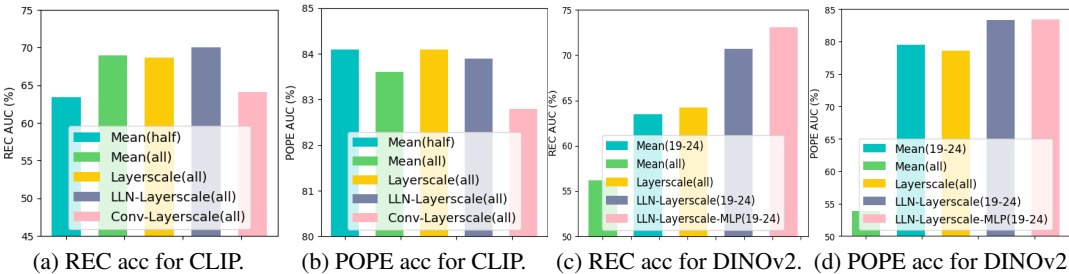


Figure 3: Average REC and POPE accuracy for merging different layers of features with multi-feature merging (MFM) strategies as input to MLLMs for visual backbones of CLIP and DINOv2.

detailed information regarding local objects, such as shape or texture. Leveraging these detailed features may enhance the MLLMs’ fine-grained perception abilities.

Evaluation Settings. For further analysis, we conduct a series of quantitative experiments using different kinds of visual models, *i.e.*, image-text contrastive learning CLIP, image-only contrastive learning DINOv2, masked image modeling MAE and supervised learning DeiT. In specific, the visual features extracted from different layers of visual models (based on ViT-Large) are aligned using a linear projection layer and then concatenated with text tokens as the input for LLMs (here we use Vicuna-7B (Chiang et al.)). The overall architecture and training process follow Shikra (Chen et al., 2023) but with fewer iterations (9400 iterations, batch size 16 on 4 A800) to save the computation cost. Then, we measure the capability of the trained MLLMs on referring expression comprehension (REC) (Chen et al., 2023), referring expression generation (REG) (Peng et al., 2023) and object hallucination benchmark (POPE) (Li et al., 2023d). Detailed descriptions of these tasks can be referred to Sec. 5.

CLIP as the Visual Branch of MLLMs. As depicted in Fig. 2, we observe that different layers of features exhibit varying biases towards grounding and understanding abilities. For instance, the shallow features demonstrate relatively higher accuracy in terms of REC and reach their optimal value at layer 12. Conversely, the deep features achieve higher accuracy in terms of POPE, indicating superior understanding ability. Notably, the relatively deep features (layer 16) display the best REG CIDEr score, showcasing promising region understanding capabilities. Consequently, instead of solely relying on deep features as done in previous works, we argue that integrating both shallow and deep features is crucial for MLLMs with improved overall performance.

We further explore various merging modes of low-level and high-level features. Denoting the output features from each transformer layer of ViT with a depth of N as $\mathbf{z} = [z_1, \dots, z_i, \dots, z_N]$, we discuss several multi-level feature merging (MFM) strategies for combining shallow and deep features, namely:

- *Mean(half)*: averaging output patch token features in the second half of the backbone as $z = (z_{N/2} + \dots + z_N)/(N/2)$.
- *Mean(all)*: averaging features output by all layers as $z = (z_1 + \dots + z_N)/N$.

Table 1: Comparison of the visual model using CLIP, DINOv2 with Multi-level Feature Merging (MFM) and **COMM** to incorporate visual embeddings of both models on VL tasks. CLIP baseline use the 23rd layer features, which follows Shikra but with fewer training iterations. MME CS and PS indicate the cognition score and perception score, respectively.

Visual Model	Avg REC	Avg POPE	COCO	Flickr30k	MME CS	MME PS	VQAv2	OK-VQA
CLIP	47.3	82.3	125.0	80.7	209.6	1107.8	68.8	44.2
DINOv2	54.8	78.3	118.0	68.9	261.8	930.5	63.1	41.9
CLIP w/ MFM	70.0	83.6	125.8	79.3	296.6	1164.4	69.5	44.7
DINOv2 w/ MFM	73.1	83.3	123.4	76.3	252.9	1086.8	68.0	42.1
COMM	72.8	83.6	127.3	81.9	360.4	1234.9	70.1	45.0

- *Layerscale(all)*: learning a scale parameter as the weight to sum features output by all layers as $z = w_1 z_1 + \dots + w_N z_N$, where w_i refers to the weight assigned to the i -th layer feature and all these weights are dynamically updated and summed up to 1.
- *LLN-Layerscale(all)*: using a linear-layer norm module to align the feature space between different layers' features and then summed by *Layerscale* as $z = w_1 \text{LLN}(z_1) + \dots + w_N \text{LLN}(z_N)$.
- *Conv-Layerscale(all)*: using a convolution and bn module to align the feature space between different layers' features and then summed by *Layerscale* as $z = w_1 \text{Conv}(z_1) + \dots + w_N \text{Conv}(z_N)$.

Fig. 3 (a) and (b) shows that simply averaging all shallow and deep features of CLIP can *de facto* achieve a satisfactory accuracy and *LLN-Layerscale* strategy further improves performance. With *LLN-Layerscale* as MFM module, the performance of CLIP can be evidently improved on commonly-used vision-language tasks as shown in Table 1.

DINOv2 as the visual Branch of MLLMs. To leverage the rich fine-grained visual information present in DINOv2, but not inherently aligned with text, we employ a non-linear Multi-Layer Perceptron (MLP) module to align the image features with the word embedding space. Fig. 2 demonstrates that the deep-layer features of DINOv2 exhibit superior grounding abilities, as evidenced by higher REC accuracy, and display satisfactory understanding abilities, as indicated by favorable POPE and REG results. Additionally, we explore the efficacy of multi-level feature merging to enhance performance. In contrast to CLIP, the merging of shallow features from DINOv2 leads to a significant performance degradation. Specifically, in Fig. 3(c) and (d), it is evident that *Mean(all)* performs notably worse than *Mean(19-24)* in terms of both REC and POPE accuracy, indicating that the shallow representations lack sufficient semantic information. Building upon the *LLN-Layerscale* approach, the incorporation of the MLP module for a more potent connection between the visual and text spaces demonstrates a clear improvement in performance. Table 1 showcases the substantial performance gains achieved by employing *LLN-Layerscale-MLP* as Multi-Level Feature Merging (MFM) module across various vision language tasks. Further detailed ablation studies on the MLP module can be found in the appendix.

MAE and DeiT as the Visual Branch of MLLMs. Fig. 2 shows that MAE features achieve acceptable REC accuracy, but suffers large performance drop on POPE and REG evaluation. This is because MAE features lack sufficient semantic information for global or regional understanding. Therefore, MAE is not suitable as the visual branch for MLLMs. DeiT performs even worse than MAE (details in the appendix). We speculate that this is because supervised training is too strong, which learns a specialized visual space that is difficult to align with the word embedding space.

4 COMM: COMBINING CLIP AND DINO WITH MULTI-LEVEL FEATURE MERGING

Architecture Overview. In this section, we introduce the proposed **COMM**, that integrates CLIP and DINO with Multi-level features Merging to enhance the visual capabilities of MLLMs. The overall framework is illustrated in Fig. 4, **COMM** is incorporated into a vision-language instruction following model built upon the recent advanced language and vision-language foundation models. Following the input instructions, our model takes vision and language as inputs to generate text responses following the input instructions. Specifically, we adopt the visual encoder of CLIP and DINOv2 (based on ViT-Large) with our proposed fusion strategy as the visual branch, and Vi-

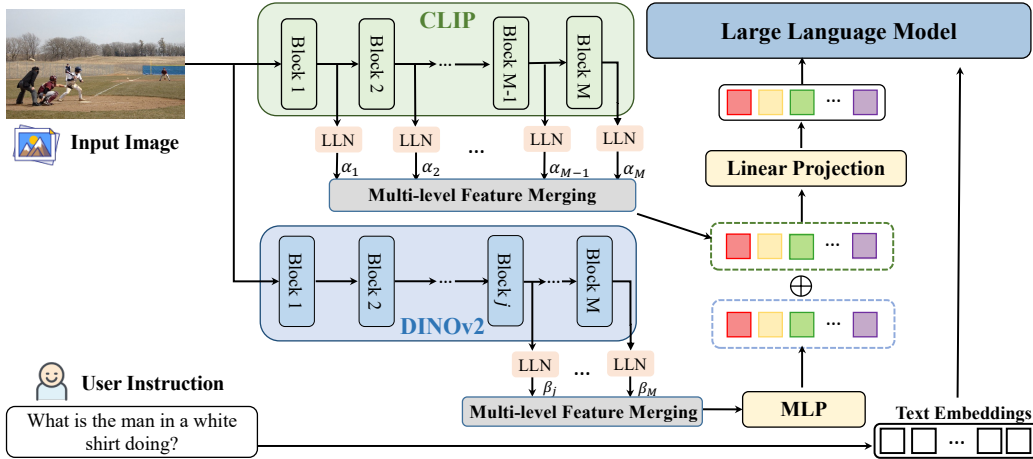


Figure 4: Overview of our **COMM**. The image is input to the vision encoder of CLIP and DINOv2, and the features from shallow and deep layers are incorporated by multi-level feature merging. The features of DINOv2 are aligned with an MLP and concatenated with features of CLIP, which are input to a linear layer. Then the fused features are concatenated with text tokens as input to LLMs.

cuna (Chiang et al.) (7B/13B) as language decoder. The visual encoder is downsampled with rate 14, meaning that an image with resolution $H \times W$ will be represented by a sequence of $\frac{H}{14} \times \frac{W}{14}$ tokens. The fused token features are projected using a linear layer and then concatenated with the instruction tokens as inputs to the language decoder, which is a generic interface to unify various vision-language tasks as text generation task.

Specifically, denote the visual encoder of CLIP and DINOv2 (ViT Large used) as f_1 and f_2 , respectively. Given an input image x , we extract the patch token features output by all layers of CLIP as $f_1(x) = [v_1^1, \dots, v_1^i, \dots, v_1^{24}]$, where $v_1^i \in \mathbf{R}^{N \times D}$, N is the number of patch tokens and D is the embedding dimension. The features output by the deep layers of DINOv2 are $f_2(x) = [v_2^{19}, \dots, v_2^i, \dots, v_2^{24}]$. Then we concatenate the features output by these two models as $\mathbf{v} = [v_1^1, \dots, v_1^{24}, v_2^{19}, \dots, v_2^{24}]$. A linear-layer norm module is leveraged to align the feature space between different layers' features and layerscale is employed for merging the multiple layer features as,

$$\bar{v}_1 = \sum_{i=1}^{24} \alpha_i \cdot \text{Linear}(\text{LN}(v_1^i)), \quad \bar{v}_2 = \sum_{j=19}^{24} \beta_j \cdot \text{Linear}(\text{LN}(v_2^j)) \quad (1)$$

where α and β are the learnable scaling parameter. Then, we employ an MLP layer to project the features of DINOv2 and concatenate the output features with that of CLIP as $\bar{\mathbf{v}} = [\bar{v}_1, \text{MLP}(\bar{v}_2)]$. Then, a linear layer is employed to match the dimension of visual features to that of text features as $\hat{\mathbf{v}} = \text{Linear}(\bar{\mathbf{v}})$. Finally, fused visual features $\hat{\mathbf{v}}$ are concatenated with text tokens as input to LLMs.

5 EXPERIMENTS

In this section, we conduct extensive evaluation on four kinds of vision-language tasks to comprehensively evaluate the visual understanding ability of our model, namely, Referring Expression Comprehension, Referring Expression Generation, Object Hallucination Benchmark, and Visual Question Answering and Image Captioning.

Training Details. Similar to previous MLLM methods, **COMM** is trained in two stages. In the first pretraining stage, we train the model on the reorganized vision-language dataset as (Chen et al., 2023), including public VQA, Image Captioning dataset and several datasets containing positional annotation RefCOCO, visual gemome (Krishna et al., 2017) and Visual-7W (Mani et al., 2020). The first pretraining stage is conducted for 100K steps. In the second instruction tuning stage, we set the sampling ratio to 50% on LLaVA-Instruct-150K (Liu et al., 2023a) and Shikra-RD (Chen et al., 2023). Instead of 224×224 resolution currently used by existing MLLMs, we use 336×336 resolution to reduce the information loss caused by image down-sampling and promote the

Table 2: Results on standard referring expression comprehension (REC) task. Generalist VL models can perform various vision-language tasks. Specialist models are designed specifically for localization tasks or generalist pretraining models that have undergone finetuning. The results of Shikra and Qwen are taken from their papers.

Model type	Model	RefCOCO			RefCOCO+			RefCOCog	
		val	test-A	test-B	val	test-A	test-B	val-u	test-u
Generalist VL SOTAs (w/o finetuning)	OFA-L*	79.96	83.67	76.39	68.29	76.00	61.75	67.57	67.58
	OFASys	-	80.10	-	-	-	-	-	-
	VisionLLM-H	-	86.70	-	-	-	-	-	-
	Shikra-7B	87.01	90.61	80.24	81.60	87.36	72.12	82.27	82.19
	Shikra-13B	87.83	91.11	81.81	82.89	87.79	74.41	82.64	83.16
	Qwen-VL-7B	89.36	92.26	85.34	83.12	88.25	77.21	85.58	85.48
	Qwen-VL-7B-Chat	88.55	92.27	84.51	82.82	88.59	76.79	85.96	86.32
	COMM-7B (Ours)	91.73	94.06	88.85	87.21	91.74	81.39	87.32	88.33
Specialist SOTAs (Specialist/Finetuned)	G-DINO-L	90.56	93.19	88.24	82.75	88.95	75.92	86.13	87.02
	UNINEXT-H	92.64	94.33	91.46	85.24	89.63	79.79	88.73	89.37
	ONE-PEACE	92.58	94.18	89.26	88.77	92.21	83.23	89.22	89.27

Table 3: Results on standard referring expression generation (REG) task in CIDEr score. We reproduce the results of Shikra-7B using its officially released checkpoint. SLR is a finetuned listener-speaker model with an added reward-based module (SLR).

Model	RefCOCO			RefCOCO+			RefCOCog	
	val	test-A	test-B	val	test-A	test-B	val-u	test-u
SLR Yu et al. (2017)	-	69.7	132.3	-	49.4	70.9	59.2	-
SLR+Rerank Yu et al. (2017)	-	77.5	132.0	-	52.0	73.5	66.2	-
Shikra	75.61	44.26	104.83	56.42	40.98	68.25	62.71	65.58
Kosmos-2	-	-	-	-	-	-	62.3	-
COMM (Ours)	93.35	54.95	131.13	70.00	52.27	79.05	79.22	77.96

fine-grained perception ability. In both stages, we freeze the visual encoder and tune all parameters in LLMs, alignment layer and multi-level feature fusion module. We adopt AdamW (Loshchilov & Hutter, 2019) as the optimizer and cosine annealing scheduler (Loshchilov & Hutter, 2017) as learning rate scheduler with an initial learning rate of $2e-5$ and global batch size of 64. All training runs on 8 NVIDIA A800 GPUs. It takes around 100h for stage one training and 20h for stage two.

5.1 REFERRING EXPRESSION COMPREHENSION

To evaluate the fine-grained understanding and positioning capability of our model, we investigate the referring expression comprehension task on benchmarks as RefCOCO (Kazemzadeh et al., 2014), RefCOCO+ (Mao et al., 2016) and RefCOCog (Mao et al., 2016), where models are asked to localize the object described with an expression. As shown in Table 2, compared with generalist VL models and previous SOTA MLLMs, **COMM** achieves significant performance gain on all benchmarks, *i.e.*, **COMM-7B** outperforms Shikra-13B and Qwen-VL-7B-Chat by 4.87% and 3.10% accuracy on average, respectively. With more powerful visual capabilities of our proposed fusion model, we can evidently surpass recent SOTA MLLMs in a more efficient way, *e.g.*, using a smaller LLM than Shikra (7B vs. 13B) and less training data than Qwen (3.6M vs. 1.4B). Besides, our generalist model even achieves comparable results with specialist SOTA methods, showing the superior grounding ability of our MLLMs.

5.2 REFERRING EXPRESSION GENERATION

Moreover, we evaluate the ability to understand image regions or objects referred via inputting bounding boxes. Instead of referring image regions or objects via detailed text descriptions, directly referring to image regions via its bounding boxes is more effective and can reduce the ambiguity. The experiments are conducted on the referring expression generation task with RefCOCO, RefCOCO+ and RefCOCog, aiming to generate text descriptions of specific regions in the bounding box. Table 3 shows that our model outperforms Shikra and Kosmos-2 by a considerable margin of 16.51 CIDEr

Table 4: Object hallucination benchmark using POPE evaluation pipeline (Li et al., 2023d). The results of Shikra-7B are taken from its paper. Except for Shikra-7B, the other results are obtained from Li et al., 2023d.

Datasets	COMM	Shikra	InstructBLIP	MiniGPT-4	LLaVA	mPLUG-Owl
Random	87.29	86.90	88.57	79.67	50.37	53.97
Popular	86.50	83.97	82.77	69.73	49.87	50.90
Adversarial	84.50	83.10	72.10	65.17	49.70	50.67

Table 5: Results on visual question answering (VQA) and image captioning. For VQA, we evaluate SOTA generalist models and our **COMM** on VQAv2 and OK-VQA following the normalization rules. Shikra is based on the 13B variant. For image captioning, we evaluate them on COCO and Flickr30k in CIDEr score. We call Flamingo as FM for short.

Datasets		COMM	Qwen	Shikra	FM-80B	BLIP-2	Unified-IO	VPGTrans	VisionLLM
VQA	VQAv2 ^{val}	79.05	-	75.33	-	65.2	-	65.2	-
	VQAv2 ^{dev}	81.04	79.5	77.36	56.3	65.0	77.9	-	-
	VQAv2 ^{std}	81.17	-	77.51	-	-	-	-	-
	OK-VQA	59.18	58.6	47.16	50.6	45.9	54.0	45.0	-
Caption	Flickr30k	88.2	85.8	73.9	67.2	-	-	-	-
	COCO	132.7	-	117.5	84.3	-	122.3	-	114.2

and 16.92 CIDEr gain on RefCOCOg, demonstrating the effectiveness of our model for fine-grained understanding. Besides, **COMM** even outperforms finetuned SLR on RefCOCO+ and RefCOCOg.

5.3 OBJECT HALLUCINATION BENCHMARK

We compare our model against the baseline models on the hallucination evaluation dataset recently introduced by POPE (Li et al., 2023d), which randomly selects 500 images from COCO (Caesar et al., 2018). Table 4 shows that **COMM** surpasses recent popular MLLMs with 1.44% and 4.95% higher accuracy on average than Shikra and InstructBLIP, respectively. By enhancing the fine-grained visual capabilities, **COMM** can effectively alleviate the object hallucination problem.

5.4 VISUAL QUESTION ANSWERING AND IMAGE CAPTIONING

We evaluate **COMM** on conventional VL tasks of VQA and Image Captioning. Specifically, image captioning requires the model to generate description for the given image and VQA asks the model to generate answer for the given image-question pair. For image captioning, we choose COCO (Chen et al., 2015) and Flickr30K (Plummer et al., 2015) as benchmarks and report the CIDEr score. For VQA task, we experiment on VQAv2 (Antol et al., 2015) and OK-VQA (Marino et al., 2019). As shown in Table 5, **COMM** achieves state-of-the-art performance on image captioning task, *i.e.*, 88.2 CIDEr score on Flickr30K and 132.7 CIDEr score on COCO, even outperforms previous SOTA models with much more parameters (*e.g.*, Shikra-13B with 13B parameters) or much more training data (*e.g.*, Qwen with 1.4B data). For VQA task, our model also shows significant advantages compared to other MLLMs. On VQAv2 val, dev and std, our model achieves 79.05, 81.04 and 81.17 accuracy respectively, which surpasses recent proposed Shikra with the same training data and procedure by a large margin, demonstrating the effectiveness of merging visual embeddings of DINOv2 and CLIP for enhancing visual capabilities. Besides, our **COMM** model outperforms Qwen with 1.54 and 0.58 accuracy gain on VQAv2 dev and OK-VQA respectively with less VQA training data, *i.e.*, we use 0.6M and Qwen with 3.6M. Training with more VQA data might further improve performance and we leave it as future work.

5.5 DEMONSTRATIONS

As shown in Fig. 5, our **COMM** model exhibits a multitude of promising capabilities including visual grounding, fine-grained region understanding and robustness to object hallucination. The first example showcases our strong fine-grained perception ability, which identifies implicit strawberries in a blender. The second example exhibits our strong visual grounding ability to successfully locates

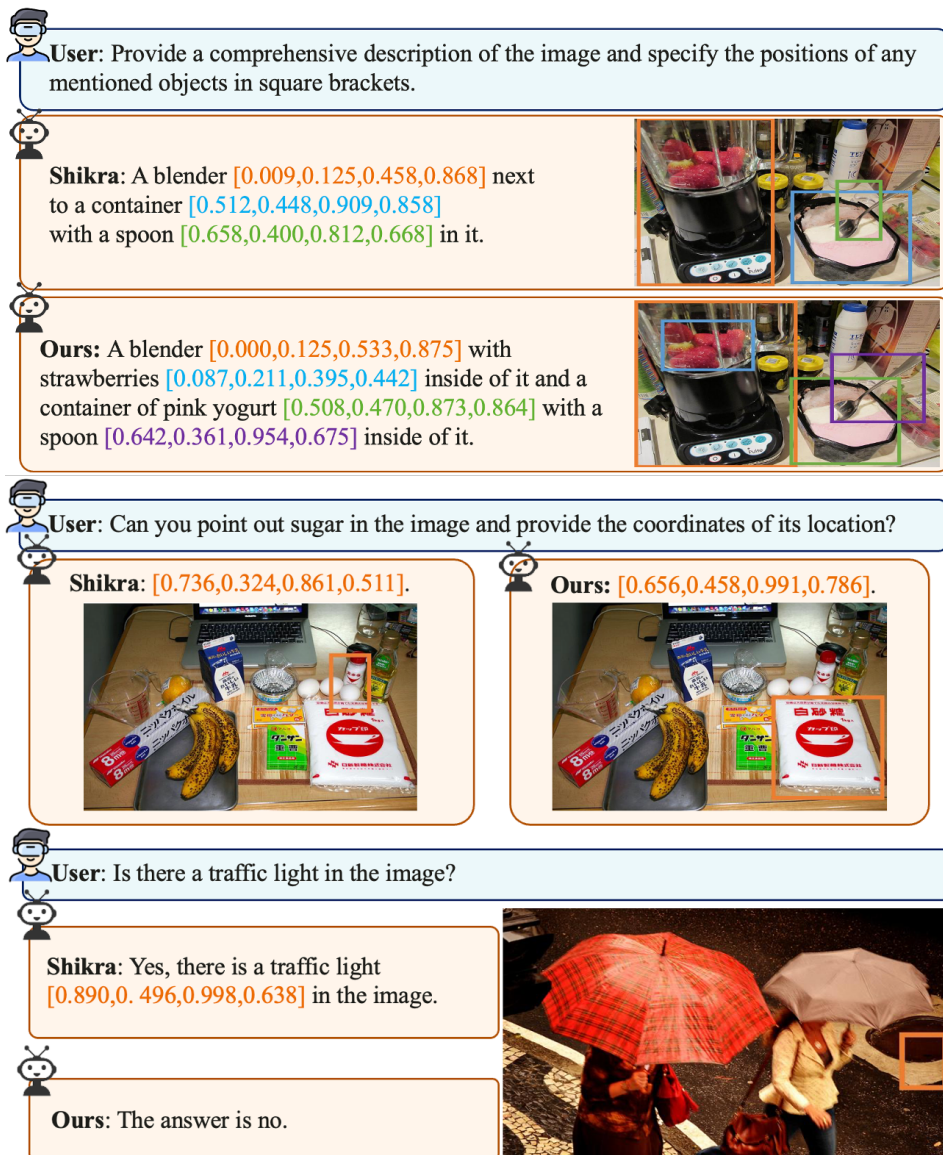


Figure 5: Qualitative comparison between Shikra with its official checkpoint and our COMM.

the object of sugar. The third case demonstrates our robustness to object hallucination. In contrast, Shikra fails on these challenging cases, showing the superior capabilities of our model.

6 CONCLUSION

This paper presented an extensive investigation into the efficacy of different visual models when employed as the visual branch in MLLMs. Through a systematic analysis, we highlight the significance of shallow layer features, which capture low-level details that prove beneficial for grounding and positioning tasks. Furthermore, we recognize the potential of the vision-only model DINOv2, which leverages its inherent fine-grained pixel-level information for enhanced fine-grained perception in MLLMs when combined with an MLP layer for alignment purposes. Motivated by our analysis, we introduce a fusion approach to combine the visual features obtained from CLIP and DINOv2, thereby further augmenting the visual capabilities and performance of MLLMs. Through qualitative analysis and extensive quantitative experiments, we demonstrate the effectiveness of our proposed method, surpassing the performance of existing MLLM models across diverse benchmark datasets. Looking ahead, we encourage future research to explore the integration of more powerful vision models to enhance the capabilities of visual branches in MLLMs. We believe that this avenue of investigation holds the key to unlocking the potential of the next generation of MLLMs.

REFERENCES

- 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. *NeurIPS*, 2022.
- Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *Proceedings of the IEEE international conference on computer vision*, pp. 2425–2433, 2015.
- 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.
- Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv: Computer Vision and Pattern Recognition*, 2021.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- Holger Caesar, Jasper Uijlings, and Vittorio Ferrari. Coco-stuff: Thing and stuff classes in context. In *CVPR*, 2018.
- Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm’s referential dialogue magic. *arXiv preprint arXiv:2306.15195*, 2023.
- Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv preprint arXiv:1504.00325*, 2015.
- WL Chiang, Z Li, Z Lin, Y Sheng, Z Wu, H Zhang, L Zheng, S Zhuang, Y Zhuang, JE Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, mar. 2023.
- 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. *arXiv preprint arXiv:2305.06500*, 2023.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc V Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. *arXiv preprint arXiv:1901.02860*, 2019.
- Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. PaLM-E: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023.
- Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 16000–16009, 2022.
- Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, 2014.
- Jing Yu Koh, Ruslan Salakhutdinov, and Daniel Fried. Grounding language models to images for multimodal inputs and outputs. 2023.
- Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual Genome: Connecting language and vision using crowdsourced dense image annotations. 123:32–73, 2017.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *arXiv preprint arXiv:2306.00890*, 2023a.

- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv:2301.12597*, 2023b.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023c.
- 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*, 2023d.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023a.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv:2304.08485*, 2023b.
- Jingyu Liu, Liang Wang, and Ming-Hsuan Yang. Referring expression generation and comprehension via attributes. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 4856–4864, 2017.
- Ilya Loshchilov and Frank Hutter. SGDR: stochastic gradient descent with warm restarts. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL <https://openreview.net/forum?id=Skq89Scxx>.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL <https://openreview.net/forum?id=Bkg6RiCqY7>.
- Arjun Mani, Nobline Yoo, Will Hinthorn, and Olga Russakovsky. Point and ask: Incorporating pointing into visual question answering. *arXiv preprint arXiv:2011.13681*, 2020.
- Junhua Mao, Jonathan Huang, Alexander Toshev, Oana Camburu, Alan L Yuille, and Kevin Murphy. Generation and comprehension of unambiguous object descriptions. pp. 11–20, 2016.
- Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pp. 3195–3204, 2019.
- OpenAI. Chatgpt: A language model for conversational ai. Technical report, OpenAI, 2023a. URL <https://www.openai.com/research/chatgpt>.
- OpenAI. Gpt-4 technical report, 2023b.
- Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023.
- Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv:2306.14824*, 2023.
- Bryan A Plummer, Liwei Wang, Chris M Cervantes, Juan C Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. In *Proceedings of the IEEE international conference on computer vision*, pp. 2641–2649, 2015.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.
- Yixuan Su, Tian Lan, Huayang Li, Jialu Xu, Yan Wang, and Deng Cai. Pandagpt: One model to instruction-follow them all. *arXiv preprint arXiv:2305.16355*, 2023.

- Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Stanford alpaca: An instruction-following llama model, 2023.
- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International conference on machine learning*, pp. 10347–10357. PMLR, 2021.
- Hugo Touvron, Matthieu Cord, and Herve Jegou. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv:2302.13971*, 2023a.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurelien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *arXiv:2307.09288*, 2023b.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language model with self generated instructions. *arXiv preprint arXiv:2212.10560*, 2022.
- Jason Wei, Maarten Bosma, Vincent Y Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. Finetuned language models are zero-shot learners. *arXiv preprint arXiv:2109.01652*, 2021.
- Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv:2304.14178*, 2023.
- Licheng Yu, Hao Tan, Mohit Bansal, and Tamara L Berg. A joint speaker-listener-reinforcer model for referring expressions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7282–7290, 2017.
- Rowan Zellers, Yonatan Bisk, Ali Farhadi, and Yejin Choi. From recognition to cognition: Visual commonsense reasoning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6720–6731, 2019.
- Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023a.
- Shilong Zhang, Peize Sun, Shoufa Chen, Min Xiao, Wenqi Shao, Wenwei Zhang, Kai Chen, and Ping Luo. Gpt4roi: Instruction tuning large language model on region-of-interest. *arXiv:2307.03601*, 2023b.
- Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan Yuille, and Tao Kong. ibot: Image bert pre-training with online tokenizer. *arXiv preprint arXiv:2111.07832*, 2021.
- Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv:2304.10592*, 2023.

APPENDIX

A ANALYSIS OF DEiT AS THE VISUAL BRANCH

We conduct experiment with recent supervised learning model DeiT III (Touvron et al., 2022). Since the visual model does not align with text, we employ a non-linear MLP layer as the alignment network. Here we use the ViT-Large variant as the visual branch and employ Vicuna-7B as the language branch of MLLMs. The architecture and training process follow Shikra (Chen et al., 2023) but with fewer training iterations (9400 iterations, batch size 16 on 4 A800) to save the computation cost. As shown in Table 6, DeiT suffers from evident performance degradation as the visual branch of MLLMs. We speculate that maybe the supervised training is so strong that it learns specialized visual space, making it difficult to align with the word embedding space.

Table 6: Comparison of the visual model using CLIP, DINOv2 with our proposed multi-level feature merging (MFM), MAE and DeiT. Experiments are conducted on referring expression comprehension and object hallucination benchmark on Random (R), Adversarial (A), and Popular (P). MAE-20 denotes using the features output by the 20-th layer of MAE. DeiT-20 denotes using the features output by the 20-th layer of DeiT.

Visual Model	RefCOCO+			RefCOCOg		RefCOCO			POPE
	test-A	test-B	val	test-u	val-u	test-A	test-B	val	A/P/R
CLIP w/ MFM	73.7	53.8	64.3	69.1	70.3	83.8	68.4	76.4	80.7/84.2/85.8
DINOv2 w/ MFM	75.3	59.3	67.0	73.0	71.8	84.4	74.1	79.6	80.3/84.2/85.5
MAE-20	64.7	49.4	56.8	63.7	62.8	77.9	68.6	73.6	66.8/71.1/76.7
MAE-22	65.9	50.0	58.5	64.2	63.2	79.3	69.8	74.9	68.0/71.2/77.5
MAE-24	67.7	51.5	60.2	65.4	64.3	80.5	69.9	75.8	69.8/76.1/80.8
DeiT-20	18.4	13.0	15.9	17.0	16.2	29.0	21.6	25.7	66.2/69.6/77.9
DeiT-22	25.3	15.4	19.4	22.6	21.8	36.9	25.3	32.0	67.9/71.6/78.7

B ABLATION ON THE MLP OF DINOv2

We conduct ablation study on the MLP module of DINOv2 for aligning visual and text embedding space. Here we use the ViT-Large variant as the visual branch and employ Vicuna-7B as the language branch of MLLMs. The architecture and training process follow Shikra Chen et al. (2023) but with fewer training iterations (9400 iterations, batch size 16 on 4 A800) to save the computation cost. We ablate on the number and the expanding ratio of MLP module. Table 7 shows that increasing the number of MLP to 2 can evidently improve performance, demonstrate the effectiveness of using a more powerful network to align the vision only model DINOv2 to the word embedding space. However, increasing the number beyond 2 suffers the degraded performance. For the expanding ratio, increasing to 8 can improve performance, while increasing to 16 does not achieve significant performance gain. Moreover, we experiment with one linear layer, which suffers severe performance degradation. Thus, non-linear MLP is necessary for aligning the features of vision-only DINOv2 to the word embedding space.

C EVALUATION BENCHMARK FOR MULTIMODAL LARGE LANGUAGE MODELS

We evaluate our final model on the recently proposed MME benchmark to further demonstrate its strong generalizability to follow a diverse range of instructions. MME benchmark measures both perception and cognition abilities on a total of 14 subtasks. We report the average results of perception tasks and cognition tasks, respectively. As shown in Table 8, the performance of our **COMM** model outperforms Shikra by a large margin of 176.59 and 39.95 score for perception and cognition abilities on MME benchmark, demonstrating the effectiveness of enhancing visual capabilities with

Table 7: Ablation study on the number and expanding ratio of MLP module. Experiments are conducted on referring expression comprehension and object hallucination benchmark on Random (R), Adversarial (A), and Popular (P).

Visual Model	RefCOCO+			RefCOCOg		RefCOCO			POPE
	test-A	test-B	val	test-u	val-u	test-A	test-B	val	A/P/R
DINOv2 w/ MLP Ratio 4	75.3	59.3	67.0	73.0	71.8	84.4	74.1	79.6	80.3/84.2/85.5
DINOv2 w/ 2MLP Ratio 4	77.5	60.3	69.2	74.6	74.7	86.5	75.3	81.4	82.4/84.5/86.2
DINOv2 w/ 4MLP Ratio 4	53.7	34.4	45.3	49.0	48.8	65.4	48.0	57.9	79.2/82.9/84.6
DINOv2 w/ 8MLP Ratio 4	8.2	6.5	7.4	6.8	6.7	14.8	12.9	14.9	56.0/55.3/59.0
DINOv2 w/ MLP Ratio 8	77.4	59.9	69.7	73.7	73.3	85.7	74.1	80.9	81.5/85.8/86.7
DINOv2 w/ MLP Ratio 16	76.2	60.2	69.7	74.5	74.6	85.7	75.5	81.5	80.4/83.7/85.7
DINOv2 w/ Linear	61.8	48.8	55.1	64.1	62.9	76.5	67.0	71.9	75.6/79.3/83.7

Table 8: Zero-shot evaluation of perception and cognition abilities on MME benchmark.

	COMM-13B	COMM-7B	Shikra-7B [†]	Qwen	InstructBLIP	LLaVA	MiniGPT-4	mPLUG-Owl
Perception	1421.00	1387.00	1210.41	1487.58	1212.82	502.82	866.57	967.34
Cognition	315.00	314.60	274.65	360.71	291.79	214.64	292.14	276.07

our fused visual embeddings. Compared with Qwen, our model suffers a litter performance degradation, and we speculate it is due to less training data than Qwen (3.6M vs. 1.4B) and we will try to train with more data for evaluation in the future work.

D MORE DEMONSTRATIONS

We provide additional demonstrations of our COMM model in this section to demonstrate a multitude of promising capabilities including visual grounding, fine-grained region understanding and robustness to object hallucination. For instance, we showcase Referring Expression Comprehension in Fig. 6 and Spot Captioning in Fig. 7. We also include cases of Referring Expression Generation in Fig. 8 and Object Hallucination in Fig. 9.

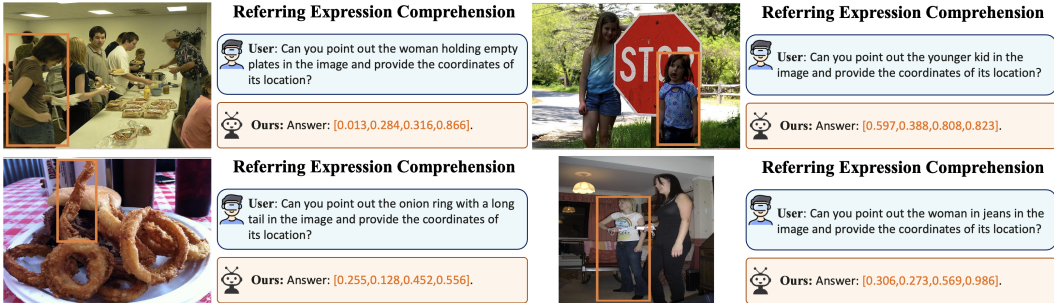


Figure 6: Referring Expression Comprehension (REC) using our COMM-7B. The task intends to localize a target object in an image described by a referring expression.

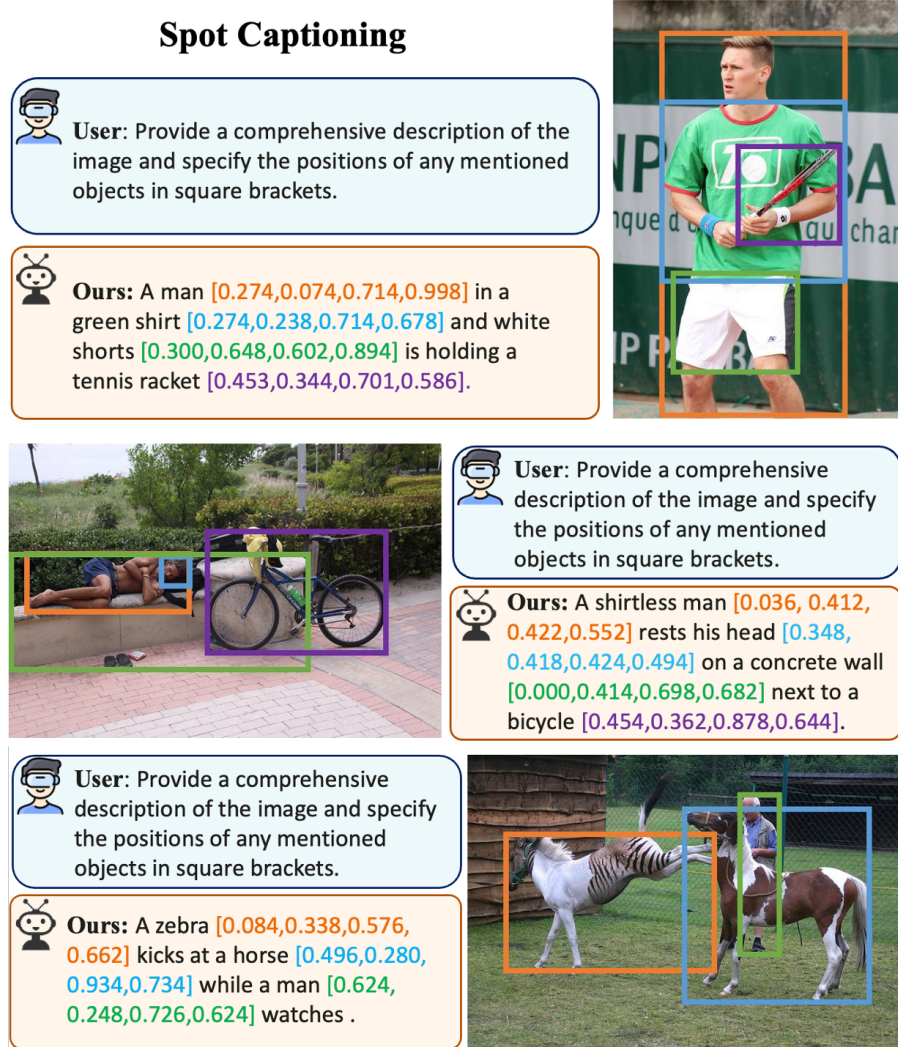


Figure 7: Spotting Captioning using our COMM-7B. The task requires the model to describe the image and spots the mentioned objects or regions using points or boxes.

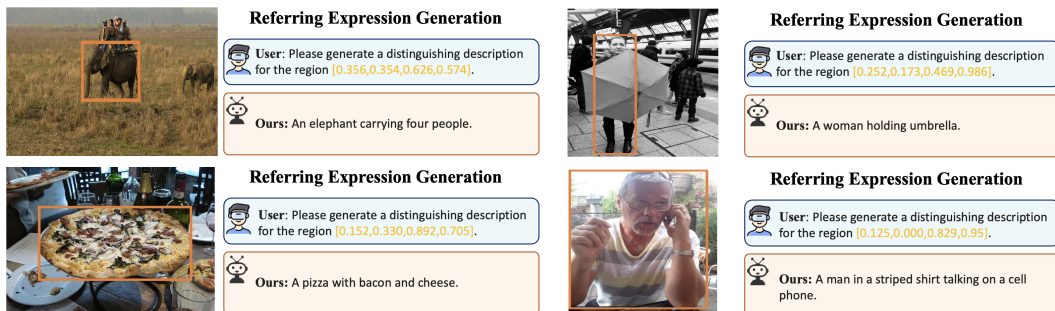


Figure 8: Referring Expression Generation (REG) using our COMM-7B. This task aims to generate a unique description for a specified location.




Figure 9: Object hallucination using our **COMM-7B**. This task aims to evaluate the robustness to object hallucination, *i.e.*, answer yes or no for the existence of questioned object.