

LMM-Det: Make Large Multimodal Models Excel in Object Detection

Jincheng Li^{1*} Chunyu Xie^{2,1*} Ji Ao¹ Dawei Leng^{1†} Yuhui Yin¹

¹360 AI Research ²Beihang University

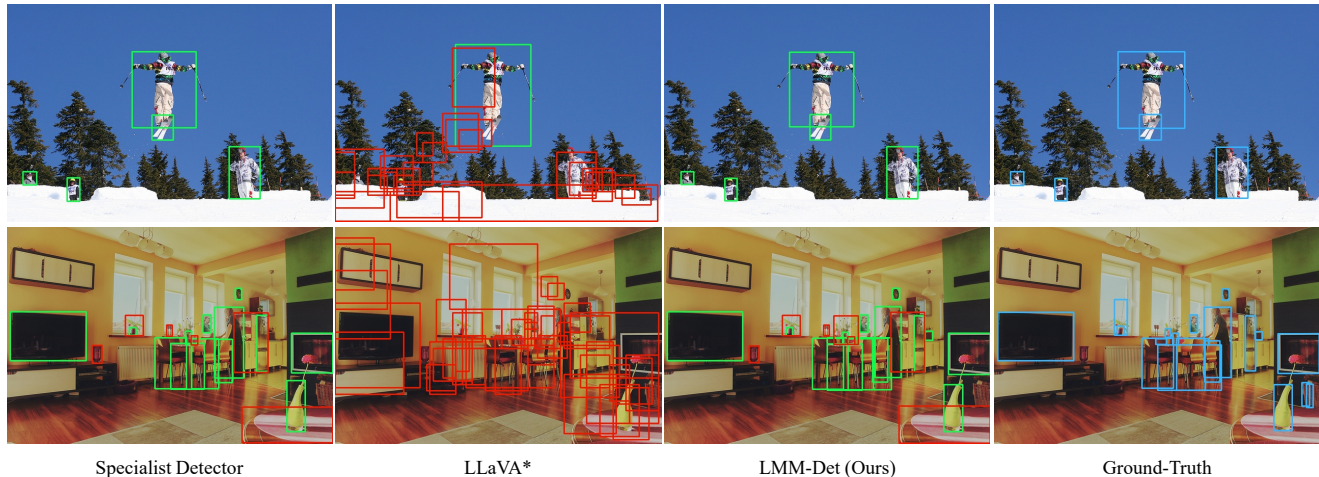


Figure 1. Visualizations of a specialist detector and large multimodal models for object detection on the validation set of COCO. For clarity, the green bounding boxes indicate the correctly predicted boxes, the red boxes represent the prediction errors, and the blue boxes denote the ground truth labels. Specifically, we employ Saliency-DETR [13] as the specialist detector. Following the prompt style of RefCOCO, we query the trained LLaVA-7B [23] with a sequence of questions, each targeting a single class category, which is called LLaVA*. In addition, a predicted box is deemed correct if it meets the following criteria: (1) The intersection over union (IoU) between the predicted box and the ground truth (GT) box is greater than 0.5; (2) The predicted class label matches the GT class label. As illustrated in the figure above, our approach (LMM-Det) enables a large multimodal model to perform object detection, achieving comparable results with state-of-the-art specialist detection methods. It is noteworthy that LMM-Det exhibits outstanding object detection capabilities without relying on additional specialist modules within its architecture. The detailed predicted labels of all images can be referred to the Figure A in Appendix.

Abstract

Large multimodal models (LMMs) have garnered widespread attention and interest within the artificial intelligence research and industrial communities, owing to their remarkable capability in multimodal understanding, reasoning, and in-context learning, among others. While LMMs have demonstrated promising results in tackling multimodal tasks like image captioning, visual question answering, and visual grounding, the object detection capabilities of LMMs exhibit a significant gap compared to specialist detectors. To bridge the gap, we depart from the conventional methods of integrating heavy detectors with LMMs and propose

LMM-Det, a simple yet effective approach that leverages a Large Multimodal Model for vanilla object Detection without relying on specialized detection modules. Specifically, we conduct a comprehensive exploratory analysis when a large multimodal model meets with object detection, revealing that the recall rate degrades significantly compared with specialist detection models. To mitigate this, we propose to increase the recall rate by introducing data distribution adjustment and inference optimization tailored for object detection. We re-organize the instruction conversations to enhance the object detection capabilities of large multimodal models. We claim that a large multimodal model possesses detection capability without any extra detection modules. Extensive experiments support our claim and show the effectiveness of the versatile LMM-Det. The datasets, models, and codes are available at <https://github.com/360CVGroup/LMM-Det>.

* Equal Contribution

† Corresponding Author, E-mail: lengdawei@360.cn

1. Introduction

Large multimodal models (LMMs) [1, 8, 18, 23, 38, 47, 53] have recently attracted substantial attention due to their exceptional capabilities in comprehending and processing multiple forms of data, such as texts, images, speech, and videos. LMMs have shown notable effectiveness in bridging the gap between different modalities, enabling more efficient human-machine interaction. In particular, LMMs have exhibited strong performance in various multimodal tasks like image captioning [19], visual question answering (VQA) [56], and visual grounding [7]. In the case of image captioning, a GPT assistant generates detailed textual descriptions for user-input images. Moreover, the goal of VQA and visual grounding is to produce detailed and contextually appropriate responses that meet the user’s requirements. These tasks demonstrate that large multimodal models are capable of aligning images and text at a fine-grained level, which holds significant potential in fine-grained image-text understanding tasks like object detection and segmentation.

Generally, object detection is a fundamental task in computer vision, serving as an extension of classification and a foundation for segmentation. Furthermore, it is essential for effective visual understanding to acquire the location and category of objects. Despite the impressive performance of large multimodal models in various multimodal tasks, their object detection capabilities remain underexplored in recent research. Compared to traditional state-of-the-art detection methods, there is a significant gap in LMMs.

To bridge this gap, existing methods [27, 42] attempt to integrate additional modules into LMMs, such as a specialized detection model [25, 34] or a region proposal network (RPN) [33]. While these approaches can achieve detection capabilities in user dialogues, they are limited by the performance of the extra modules and introduce additional latency during inference. More importantly, they do not fully explore the potential of LMMs to perform object detection tasks independently. On the other hand, some studies [7, 21, 30, 32, 49, 52] make LMMs directly output object categories and bounding box coordinates, demonstrating detection abilities in visual grounding tasks such as the referring expression comprehension (REC) and phase grounding. For instance, KOSMOS-2 [30] integrates grounding capabilities into downstream applications and provides the corresponding bounding boxes. Shikra [7] designs a simple architecture without the need for pre-/post-detection modules and external plug-in models to unlock LMMs’ grounding abilities. However, it is non-trivial for these works to finish vanilla object detection, which requires localizing and classifying all objects within an image.

The observation that large multimodal models exhibit fine-grained image-text alignment capability and possess preliminary detection capabilities (*e.g.*, REC), motivates us to further investigate this phenomenon. Our goal is to explore

the capability of LMMs in performing object detection tasks without relying on additional dedicated detection modules. To this end, we highlight the following inherent challenges when large multimodal models meet with object detection.

Most existing large multimodal models [7, 27, 30] evaluate their detection capabilities on RefCOCO [16] instead of specialist detection benchmarks like COCO [22], despite using large-scale object detection datasets (*e.g.*, Object365 [35], LVIS [10], OpenImages [17] or COCO) for pre-training. Moreover, customized instruction data for vanilla object detection is often ignored during the instruction tuning stage, leading to poor localization and identification of objects within images. This hinders practical applications in fields such as robotics, autonomous driving, and augmented reality [27].

On the other hand, large multimodal models typically generate a small number of proposal bounding boxes, resulting in poor object detection performance. For instance, LLaVA [23] shows suboptimal performance in detecting all objects as illustrated in Figure 1, using the following prompt: “Provide the bounding box coordinate of the region this sentence describes if region exists in the image: <category>”. For each query, LLaVA generates only a few bounding boxes, and most of them are inaccurate, resulting in a low recall rate for the object detection task.

In this paper, a comprehensive analysis is conducted to investigate the performance gap between specialist object detection models and LMMs. Specifically, we evaluate the large multimodal models on COCO, a specialist object detection benchmark. Through sample visualization and distribution comparison, we claim that the root cause of underperformance in the object detection task is the low recall rate. To address this, we present LMM-Det, a simple yet effective approach to make large multimodal models excel in object detection. We propose to increase the recall rate by adjusting the training data distribution since we observe that the trained model tends to approximate this distribution. Moreover, we introduce inference optimization and present an instruction-tuning dataset. To summarize, our main contributions are as follows:

- We focus on exploring the detection ability of large multimodal models to unlock their full potential in practical applications. To this end, we provide a comprehensive analysis to facilitate the detection performance improvement for large multimodal models.
- We propose a simple yet effective approach named LMM-Det without any extra detection modules. We introduce data distribution adjustment and inference optimization to increase the inherent recall rate when large multimodal models meet with object detection.
- Extensive experiments demonstrate that LMM-Det not only exhibits detection capability but also preserves inherent multimodal capabilities like captioning and VQA.

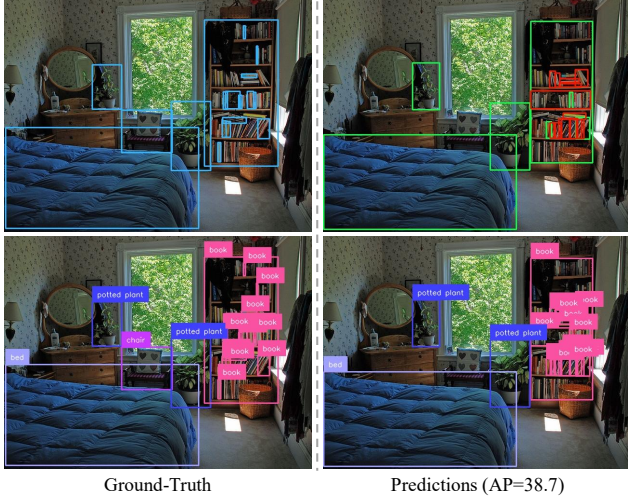


Figure 2. Visualizations of ground-truth and predicted bounding boxes, generated by the model in the 5th row of Table 1.

2. Examination of LMMs in Object Detection

To evaluate and analyze the detection performance of large multimodal models (LMMs), we first design a series of experiments from the perspective of data scale and image resolution. We then provide an in-depth analysis to facilitate better adaptation of LMMs to object detection. In this section, we use the standard large multimodal model LLaVA-7B [23] for all experiments. We choose RT-DETR [55] as the representative specialist detector for comparison.

2.1. Exploratory Experiments

This section conducts exploratory experiments to assess the detection capabilities of LMMs and the implementation details are provided in Section A of Appendix.

Zero-shot on COCO. A straightforward way to examine the detection performance of LMMs is through zero-shot evaluation. Thus, we first perform zero-shot experiments on COCO using LLaVA [23]. The qualitative results are shown in Figure 1, while the quantitative results are provided in the first row of Table 1. The experimental results reveal that LLaVA exhibits poor detection performance on the COCO dataset in a zero-shot setting, likely due to the absence of COCO data during its training.

Supervised Fine-tuning with COCO. To verify the above claim, we then leverage COCO for supervised fine-tuning. As illustrated in the second and third rows of Table 1, the incorporation of detection data (*i.e.*, COCO) yields a notable enhancement in AP, although there remains scope for further improvement. Inspired by the fact that expanding the data scale in conventional object detection methods is an effective technique to improve detection accuracy, we attempt to augment with additional detection data to make large multimodal models excel in object detection.

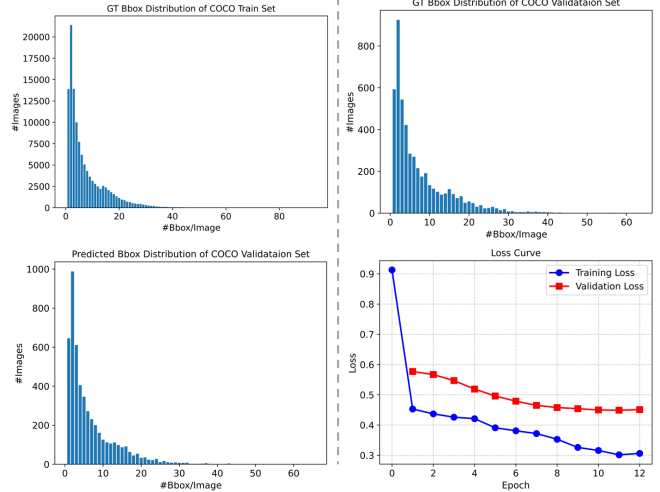


Figure 3. Bounding box (Bbox) distribution of ground-truth and predictions. We employ the model (AP=38.7) in the 5th row of Table 1 to generate boxes. In addition, we provide the loss curve to confirm that the model is not overfitting.

Method	SFT data	COCO	Object365	Res.	AP	AP ₅₀	AP ₇₅	AR@100
LLaVA	✓	×	×	336	0.2	0.6	0.2	11.2
	✓	✓	×	336	14.5	29.5	12.7	24.3
	×	✓	×	336	16.0	31.2	14.4	26.7
	×	✓	✓	336	15.6	28.5	15.1	21.5
	×	✓	×	644	17.7	30.3	17.5	26.9
	×	✓	✓	644	38.7	55.8	41.3	50.5
Specialist Detector	×	✓	✓	640	49.2	66.6	53.5	70.9

Table 1. Examinations of LLaVA-7B on the validation set of COCO. “SFT data” denotes the 665K instruction data during supervised fine-tuning (SFT) in LLaVA. “Res.” is image resolution.

Expand Detection Data. As aforementioned, we expand the detection data by incorporating Object365 during training. From the fourth row of Table 1, the addition of more data fails to provide a substantial boost in performance. The possible reasons are: (1) the low input image resolution is insufficient to perform object detection; (2) the instruction organization of COCO data needs to be improved.

Upsample Image Resolution. To ensure a fair comparison with the specialist detector, we employ interpolation to increase the image resolution from 336 to 644, thus aligning it with the detector’s native resolution. As shown in rows 5-6 of Table 1, the integration of Object365 using a higher input resolution of 644×644 indeed leads to a noticeable improvement in detection performance. However, LLaVA’s performance remains significantly inferior to that of the traditional specialist model, despite training on a similar resolution and an identical scale of detection data.

2.2. Phenomenon Analysis

To narrow the performance gap between LMMs and specialist detection models, we investigate the underlying causes.

Visualizations of Predictions and GT. We visualize the generated bounding boxes with the corresponding categories and the ground-truth (GT) in the bottom row of Figure 2. The visualization results are surprisingly satisfactory despite being 10.5% lower than the specialist model on AP. According to the criteria specified in Figure 1, we plot boxes of different colors and provide the results in the top row of Figure 2. We observe that the red bounding boxes, deemed as false positives, are actually with correct predicted labels. Meanwhile, the labels for ground truth (GT) are not fully annotated, *e.g.*, “book”. Furthermore, our statistical results (as shown in Table A in Appendix) indicate that the average number of bounding boxes per image in both the training and validation sets of COCO is approximately 7, which also matches the average number of boxes produced by the trained model. One possible reason is that the trained model has adapted to the data distribution of the COCO dataset.

Distribution Comparison. We provide the box distribution for predictions and ground truth to further analyze the above observations. From Figure 3, the predicted box distribution approximates the training set of COCO. However, incomplete ground-truth annotations lead to premature truncation of predictions, resulting in a small number of generated bounding boxes. In addition, the current simple auto-regressive training recipe results in predicting fewer bounding boxes than specialist detection models.

Actually, traditional object detection methods maintain an appropriate recall rate on proposals (candidate regions), such as 300 proposals in Faster RCNN [33] or 900 proposals in H-Deformable-DETR [15] to balance detection performance and computational costs. In this case, an insufficient recall rate can substantially degrade the detection performance. However, it is challenging for a large multimodal model to function as a Region Proposal Network (RPN) [33] and generate a large number of high-quality proposals due to the inherent limitations of LMMs’ next-token prediction loss under incomplete GT annotations. Therefore, increasing the recall rate is crucial for improving the overall detection performance of large multimodal models.

3. LMM-Det

This section introduces LMM-Det, which enhances LMMs’ detection capabilities by increasing the recall rate. We first introduce the model architecture in Section 3.1. The advancements of LMM-Det consist of data distribution adjustment (Section 3.2) and inference optimization (Section 3.3).

3.1. Preliminary: Model Architecture

As illustrated in Figure 4 (c), LMM-Det consists of a visual encoder, a projector, and a large language model.

Visual Encoder. LMM-Det employs the OWLv2-ViT model [28] as the visual encoder, which supports high-resolution image input (1008×1008) and captures detailed

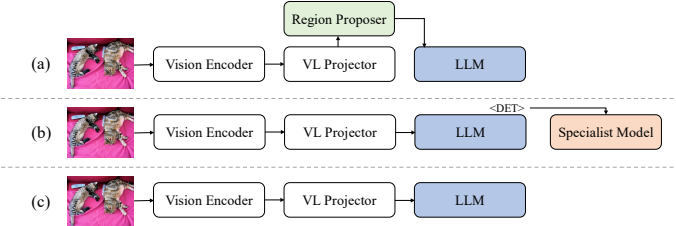


Figure 4. Overview of the proposed LMM-Det. Compared to the other LMMs with extra modules like region proposal generators (a) or specialist detection models (b), LMM-Det (c) enables LMMs to unlock the detection capability in a straightforward manner.

information for object detection. We do not compress the sequence of visual tokens and feed them directly into LMMs.

Projector. LLaVA [23] demonstrates that a linear projector is sufficient for vision language alignment. Additionally, Honeybee [4] shows that a linear projector can retain all the local context of visual features through a one-to-one projection without loss. Therefore, we use the linear projector to map visual features into the embedding space of text tokens.

Large Language Model. We utilize Vicuna-1.5-7B with a max sequence length of 16,000 as the large language model. LMM-Det is trained to perform token prediction using a paradigm of language modeling loss. Formally, given an image and the well-organized instruction text tokens, we maximize the following optimization problem:

$$\max_{\theta} \sum_{i=1}^L \log p_{\theta}(\tilde{y}_i | \mathbf{x}_v, \mathbf{x}_t, \mathbf{y}_{1:i-1}), \quad (1)$$

where θ is the trainable parameters, \mathbf{x}_v is the visual tokens from the visual encoder and the projector, \mathbf{x}_t is the instruction text tokens from the tokenizer of the large language model, $\mathbf{y}_{1:i-1}$ is the answer tokens in the instruction data before the current prediction token \tilde{y}_i , and L is the length of the output sequence.

3.2. Data Distribution Adjustment

As aforementioned in Section 2, the incomplete ground-truth annotations in the training set may lead to premature truncation of predictions in large multimodal models. To mitigate this issue, we adjust the data distribution of the training set since the trained large multimodal model tends to approximate this distribution. Moreover, the process of data distribution adjustment should achieve the goal of increasing the recall rate to enhance the detection performance of large multimodal models. To this end, a possible way is to use a self-training strategy, *i.e.*, generating pseudo labels by large multimodal models and training them with ground-truth labels iteratively. However, the generation of high-quality pseudo labels may face challenges in large multimodal models that do not possess strong object detection abilities. In

this work, we use a semi-supervised learning strategy to adjust the training data distribution, increase the recall rate, and thus improve LMMs’ detection accuracy.

Specifically, the data distribution adjustment of LMM-Det comprises three steps: 1) Pseudo Label Generation. This step aims to prepare more high-quality data for incomplete ground-truth annotations in the training set. 2) Data Merging and Design. The crafted pseudo labels and original ground-truth labels are merged and then designed to apply object detection to LMMs. 3) Token Representation Verification. We investigate different token representations to represent coordinates and confidence scores.

Pseudo Label Generation. We augment the ground-truth labels with high-quality pseudo-labels, thereby enhancing the annotation diversity. Concretely, we craft several pseudo-labels using a pre-trained specialist detector (*i.e.*, Saliency-DETR [13]). In this sense, we develop substantial proposals to improve the overall detection performance of large multimodal models. Notably, our modification only involves scaling the data, while the architecture of our LLM-Det does not rely on additional detection modules throughout both the pre-training and supervised fine-tuning stages.

Data Merging and Design. To remove the redundant boxes and further improve the data quality, we merge the generated pseudo-labels and the original ground-truth labels by performing Non-Maximum Suppression (NMS). Moreover, the existing data organization is non-trivial to compute average precision on COCO due to the absence of the confidence score of each predicted box. In practice, we empirically observe that the confidence score computed by the average probability of each coordinate token (after softmax) is not distinguishable. To address this, we make a large multimodal model output both the coordinates and the corresponding confidence score instead of only the coordinates. We assign a confidence score of 1 to the ground-truth labels, whereas the confidence scores of the pseudo-labels are determined by the pre-trained specialist detector.

Token Representation Verification. We examine how to represent the coordinates and their corresponding confidence scores. We explore two alternatives: directly outputting token predictions and expanding the vocabulary with extra tokens. The first approach increases the length of the model’s output tokens without requiring additional training of vocabulary embeddings, whereas the second approach achieves the opposite, necessitating extra training of vocabulary embeddings but resulting in shorter output tokens. Experiments in Section 4.5 show that the first approach achieves better detection accuracy. Thus, it is adopted for all experiments.

By adjusting the training data distribution, LMM-Det can generate more accurate and numerous bounding boxes, which can increase the recall rate and then improve the overall detection performance. In addition, we empirically show that LMM-Det can not only approximate the detection data

distribution but also maintain the original capabilities to approximate other data distributions like image captioning and visual question answering, as shown in Section 4.4.

3.3. Inference Optimization Tailored for Detection

At the inference stage, a straightforward approach for detecting all objects within an input image in large multimodal models is to predict all bounding boxes simultaneously, such as InternVL 2.5 [8]. However, we empirically observe that this approach is non-trivial to generate enough proposals with a fixed sampling strategy of LMMs. One possible reason is that current LMMs are hard to process all fine-grained image details in a single prediction step. We attempt to optimize this solution to make LMM-Det output more proposals to increase the recall rate at inference.

Specifically, we compromise on computational cost to obtain better detection accuracy of LMM-Det. Instead of outputting all bounding boxes in a single step, we make LMM-Det output all objects that belong to only one category in a single step and repeat this step several times. That is, LMM-Det predicts the bounding boxes for each category independently. To maintain consistency between inference and training, we re-organize the instruction conversations to adopt a class-specific prediction strategy. We illustrate the re-organized data in Figure B of Appendix.

4. Experiments

4.1. Implementation Details

Training Recipe. We train LMM-Det in three consecutive stages. Table B in Appendix summarizes the hyperparameters of all stages. We use 595K image-text pairs and 1.86M images to train LMM-Det in total. Training LMM-Det takes 176 hours on a cluster of 6 nodes, each equipped with 8 Nvidia H800 GPUs. Below are the details of each stage.

Stage I. We align the vision and language modules by pre-training the projector while freezing the visual encoder and the large language model. Specifically, we leverage the 595K image-text pairs employed in LLaVA [23].

Stage II. We utilize a large-scale object detection dataset (*i.e.*, Object365 [35]) to pre-train the proposed LMM-Det. In this stage, we train the projector and the large language model while freezing the visual encoder.

Stage III. We re-organize an object detection instruction dataset (as shown in Figure B of Appendix) built upon the COCO dataset, a widely-used benchmark for the object detection task. Then we fine-tune the projector and the large language model using this instruction data to further improve the detection capability of LMM-Det.

(Optional) Stage IV. In this stage, the projector and large language model are trained with the visual encoder frozen, using the 665k LLaVA [23] dataset alongside our proposed

Method	Visual Backbone	LLM	w. Specialist	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L	AR@100
LLaVA [23]	CLIP-L	Vicuna-7B	×	0.2	0.6	0.2	0.0	0.1	0.7	11.2
Shikra [7]	CLIP-L	Vicuna-7B	×	0.4	0.8	0.4	0.0	0.4	1.0	18.7
KOSMOS-2 [30]	CLIP-L	MAGNETO-24-layers	×	7.6	13.7	7.3	0.8	6.7	15.8	18.2
InternVL-2.5 [8]	InternViT-300M	Internlm2.5-7B	×	11.8	18.4	12.0	3.6	13.0	23.4	27.5
Groma [27]	DINOv2	Vicuna-7B	✓	12.8	17.0	13.8	3.3	12.8	24.2	22.5
LMM-Det (Ours)	OWLv2-L	Vicuna-7B	×	24.5	34.7	26.3	15.4	27.4	37.3	46.6

Table 2. Zero-shot results on COCO compared with state-of-the-art large multimodal models. We do not report the zero-shot results of VisionLLM v2 [42] since they do not release the pre-trained checkpoints. “w. Specialist” represents “with specialist detection model”.

Method	Visual Backbone	LLM	w. Specialist	AP	AP ₅₀	AP ₇₅	AP _S	AP _M	AP _L	AR@100
Faster R-CNN [33]	RN50-FPN	-	✓	40.2	61.0	43.8	24.2	43.5	52.0	54.0
Cascade R-CNN [3]	RN101-FPN	-	✓	42.7	61.6	46.6	23.8	46.2	57.4	-
Deformable-DETR [57]	RN50	-	✓	43.7	63.0	47.6	26.7	47.0	58.0	63.2
RT-DETR [55]	RN50	-	✓	55.3	73.4	60.0	38.0	59.9	71.6	74.4
H-Deformable-DETR [15]	Swin-L	-	✓	56.1	75.1	61.3	39.2	60.4	72.4	73.1
Saliency-DETR [13]	FocalNet-L	-	✓	57.3	75.5	62.4	40.9	61.8	74.5	75.4
Groma [27]	CLIP-L	Vicuna-7B	✓	32.4	-	-	-	-	-	-
Groma [27]	DINOv2	Vicuna-7B	✓	43.6	-	-	-	-	-	-
VisionLLM v2 [42]	Swin-T	Vicuna-7B	✓	56.3 ↓	74.3	61.6	-	-	-	-
Grounding DINO [25]	Swin-T	-	✓	57.2	-	-	-	-	-	-
Griffon-13B [51]	CLIP-L	LLaMA2-13B	×	24.8	40.6	25.1	5.9	25.5	48.7	-
Griffon v2 [50]	EVA2-CLIP-L	LLaMA2-13B	×	38.5	54.3	41.2	19.4	43.2	57.6	-
LLaVA* [23]	CLIP-L	Vicuna-7B	×	38.7	55.8	41.3	20.1	43.6	57.3	50.5
LMM-Det (Ours)	OWLv2-L	Vicuna-7B	×	47.5	66.5	51.1	34.7	51.8	60.3	63.6
LMM-Det[†] (Ours)	OWLv2-L	Vicuna-7B	×	47.1	66.2	50.5	35.0	51.6	60.1	63.1

Table 3. Fine-tuned results on COCO compared with traditional state-of-the-art detection models and large multimodal models relying on extra specialist detection models (w. Specialist). VisionLLM v2 employs Grounding DINO as the additional specialist module for object detection, yet this integration compromises the original detection performance of Grounding DINO. LLaVA* denotes that we retrain LLaVA with Object365 and COCO. LMM-Det[†] means that we apply the optional Stage IV to obtain a versatile LMM-Det.

re-organized instruction data and adopting the same hyper-parameters as Stage III. We denote it as LMM-Det[†].

Inference and Evaluation. We evaluate the detection capability of LMM-Det on the validation set of COCO. For each category of the validation set, we construct the corresponding prompt for questioning. We gather all the predicted outputs and regard them as the final proposals. We evaluate the mean average precision (mAP) to obtain quantitative results. For simplicity, we denote AP as mAP. We take AP, AP₅₀, AP₇₅, AP_S, AP_M, AP_L, and AR@100 as metric.

4.2. Zero-shot Experiments

In this section, we compare LMM-Det with state-of-the-art LMMs in a zero-shot setting. We adopt their official trained models and test them in a zero-shot way. The implementation details are given in the Appendix. LMM-Det only employs two stages as described in Section 4.1 and discards Stage III.

From Table 2, LMM-Det achieves the best results on COCO in a zero-shot manner, demonstrating its potential for detection after training on detection data. LLaVA, without training on detection data, only achieves 0.2 AP on COCO. Models like InternVL-2.5, which utilizes a large amount

of detection data, show better performance than those with limited data. Groma, which incorporates a specialist detection model, also performs well. Without extensive detection data and the specialist model, LMM-Det significantly outperforms other LMMs, validating our phenomenon analysis and the effectiveness of our proposed approach.

4.3. Fine-tuned Experiments

We further fine-tune LMM-Det on COCO. We compare LMM-Det with traditional detection models and LMMs relying on extra detection experts that can perform the object detection task (*i.e.*, Groma [27] and VisionLLMv2 [42]). Specifically, VisionLLMv2 uses Grounding DINO as the extra specialist detection model. For a comprehensive comparison, we also report the results of Grounding DINO.

As shown in Table 3, traditional state-of-the-art detection models outperform Groma, even when Groma is augmented with specialist detection models. Similarly, VisionLLMv2, despite incorporating Grounding DINO, exhibits a performance decline. We also retrain LLaVA using both COCO and Object365 datasets. The experimental results highlight a significant performance gap between traditional state-of-the-



Figure 5. Visualizations of LMM-Det on COCO.

Model	COCO AP	Image Captioning				VQAv2 Accuracy
		BLEU@4	METEOR	CIDEr	SPICE	
LLaVA	0.2	29.4	29.3	108.9	23.6	78.5
LMM-Det [†]	47.1	26.7	28.1	99.0	22.4	74.1

Table 4. Quantitative results for versatile LMM-Det[†].

art detection models and large multimodal models. LMM-Det, however, narrows this gap, supporting the argument that LMMs can inherently perform object detection without the need for additional specialist detection models.

4.4. Versatile LMM-Det

We provide visualizations of LMM-Det on the COCO validation set in Figure 5, showcasing its detection performance without extra specialist detectors. Furthermore, the qualitative and quantitative results (as given in Figure 6 and Table 4, respectively) demonstrate LMM-Det’s versatility, supporting that LMM-Det[†] not only unlocks the detection capabilities but also maintains high performance in image captioning and VQA. We provide more results in Table D in Appendix.

4.5. Ablation Study

Visual encoder. We first replace the CLIP-ViT [31] with OWLv2-ViT [28] to further increase the input resolution in ablation study. As shown in the first row of Table 5, this strategy achieves a 3.4% gain in AP. We further investigate the effectiveness of DINOv2 [29] in Table E of Appendix.

Data Distribution Adjustment. This approach aims to address the challenge of the insufficient recall rate. In Table 5, it effectively improves the recall rate and enhances

OWLv2-ViT	DDA	INO	AP	AP ₅₀	AP ₇₅	AR@100
×	×	×	38.7	55.8	41.3	50.5
✓	×	×	42.1	57.8	45.8	51.3
✓	✓	×	44.2	61.3	47.5	56.0
✓	✓	✓	47.5	66.5	51.1	63.6

Table 5. Ablation Study. The baseline experiment in the first row is LLaVA* [23], which uses the CLIP-ViT as the visual encoder. DDA: data distribution adjustment; INO: inference optimization.

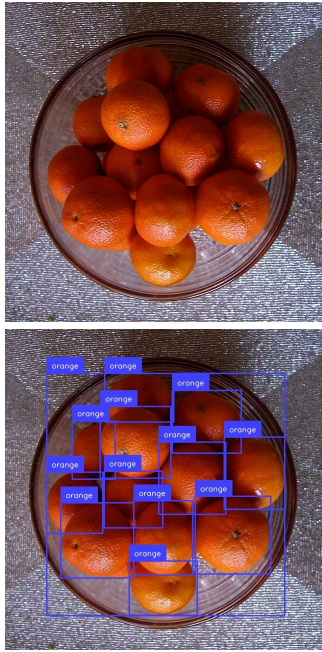
overall detection performance. We also draw the adjusted bounding box distribution in Figure C of Appendix.

Inference Optimization. We further investigate the effectiveness of the inference optimization. From Table 5, this strategy results in a significant increase in AP from 44.2% to 47.5% and AR@100 from 56.0% to 63.6%. To better visualize the enhancement of recall rate, we extend Figure 2 and show the qualitative results in Figure D of Appendix.

Effectiveness of extra vocabularies. We investigate different token representations to represent coordinates and confidence scores. From Table F in Appendix, directly outputting token achieves better performance than extra vocabularies, which aligns with the findings of Shikra [7].

Effectiveness of sampling strategy and inference cost. We conduct ablation studies using greedy decoding, beam search, and top-p sampling during inference. As shown in Table G in Appendix, beam search with beam=2 achieves the best results but requires more inference time. In particular, LMM-Det requires a computational time of approximately 4.0 seconds to process a single image.

Visual Input Example



User: Detect all the objects in the image that belong to the category set {orange}.

LMM-Det: [0.388, 0.722, 0.601, 0.891], 0.773; [0.298, 0.465, 0.322, 0.539], 0.183; [0.130, 0.446, 0.306, 0.637], 0.758; [0.310, 0.443, 0.493, 0.620], 0.751; [0.130, 0.446, 0.402, 0.779], 0.147; [0.311, 0.138, 0.527, 0.290], 0.771; [0.130, 0.446, 0.602, 0.893], 0.136; [0.130, 0.138, 0.528, 0.638], 0.179; [0.130, 0.138, 0.874, 0.893], 1.000; [0.298, 0.242, 0.518, 0.465], 0.791; [0.594, 0.521, 0.831, 0.763], 0.848; [0.480, 0.352, 0.690, 0.578], 0.781; [0.210, 0.286, 0.344, 0.468], 0.729; [0.402, 0.535, 0.587, 0.767], 0.791; [0.522, 0.192, 0.739, 0.386], 0.780; [0.298, 0.242, 0.690, 0.580], 0.137; [0.685, 0.337, 0.874, 0.562], 0.838; [0.174, 0.542, 0.402, 0.775], 0.790; [0.298, 0.138, 0.739, 0.466], 0.167; [0.300, 0.242, 0.874, 0.766], 0.156; [0.497, 0.281, 0.554, 0.373], 0.324

User: Detect all the objects in the image that belong to the category set {apple}.

LMM-Det: There are no objects in the image that belong to the required category set.

User: Describe the image in detail.

LMM-Det: The image features a clear glass bowl filled with a variety of oranges. There are at least 13 oranges in the bowl, with some of them being closer to the edge of the bowl and others being more centrally located. The oranges are arranged in a way that allows for a clear view of their vibrant colors and textures. The bowl is placed on a table, which is covered with a grey tablecloth. The oranges are arranged in a way that showcases their unique shapes and sizes, creating an appealing and visually pleasing display.

User: What color are oranges in the image?

LMM-Det: The oranges in the image are orange.

Figure 6. Chat example demonstrates that LMM-Det not only enables detection but also retains its intrinsic multimodal capabilities. The left-upper image represents the input, while the left-lower image illustrates the visualization of the output coordinates. For better visualization, we apply post-processing for the output coordinates by employing NMS with a threshold of 0.5.

5. Related Work

5.1. Large Multimodal Models

Pioneering works [5, 6, 9, 11, 19, 26, 36, 44–46, 48, 56] about large multimodal models (LMMs) primarily focus on vision-language alignment and support fundamental multimodal tasks such as image captioning and visual question answering. In this case, LMMs have demonstrated their considerable potential to perform multimodal tasks.

To further demonstrate the comprehensive capabilities of vision-language tasks, subsequent large multimodal models support resolution-sensitive tasks such as visual grounding and optical character recognition (OCR). These models leverage meticulously curated pre-training and instruction datasets, showcasing an effective strategy that highlights their robust capabilities. This type of LMMs encompass a broad range of models [2, 7, 8, 12, 14, 18, 20, 24, 30, 37, 39–41, 43, 47, 54], such as Shikra, KOSMOS-2, Ferret v2, MiniCPM-V 2.6, DeepSeek-VL2, InternVL-2.5 and Qwen2.5-VL. Compared to the coarse-grained visual question answering, the object localization capabilities of these LMMs are a crucial prerequisite for the effective deployment of LMMs in real-world applications such as agent, robotics, autonomous driving, and security surveillance. In this paper, we examine the detection potential of LMMs.

5.2. Detection Capability Exploration for LMMs

While LMMs show remarkable success on most vision-language benchmarks, they struggle to achieve comparable performance on object detection tasks. Existing approaches [8, 21, 27, 42, 50, 51] attempt to achieve object detection over large multimodal models. For example, VisionLLM-v2 [42] introduces Grounding-DINO [25] into LMMs and uses a special token <DET> to perform object detection. Groma [27] employs a region proposer to discover regions of interest to unlock the detection capabilities of LMMs. Instead, this paper aims to investigate the object detection capability of large multimodal models without the help of specialist detection models or proposal networks.

6. Conclusion and Limitations

LMM-Det has addressed the detection performance gap between LMMs and specialist detectors without relying on extra modules. Specifically, we provide a comprehensive exploration analysis and summarize that the key challenge is the insufficient recall rate. We enable LMM-Det to excel in object detection while preserving its inherent capabilities like caption generation and VQA. Besides, LMM-Det suffers from non-negligible inference time latency, making it less competitive with traditional real-time detections. In the future, we attempt to reduce the inference time of LMM-Det.

7. Acknowledgements

We sincerely thank Shuaicheng Niu from Nanyang Technological University and Xiaole Zhu from 360 AI Research for their valuable discussions and feedback. This study was carried out with the financial support from the National Key Research and Development Program of China (Nos. 2022YFB3904303).

References

- [1] 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*, pages 23716–23736, 2022. 2
- [2] 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. 8
- [3] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, 2018. 6
- [4] Junbum Cha, Wooyoung Kang, Jonghwan Mun, and Byungseok Roh. Honeybee: Locality-enhanced projector for multimodal llm. In *CVPR*, pages 13817–13827, 2024. 4
- [5] Feilong Chen, Minglun Han, Haozhi Zhao, Qingyang Zhang, Jing Shi, Shuang Xu, and Bo Xu. X-llm: Bootstrapping advanced large language models by treating multi-modalities as foreign languages. *arXiv preprint arXiv:2305.04160*, 2023. 8
- [6] Guiming Hardy Chen, Shunian Chen, Ruifei Zhang, Junying Chen, Xiangbo Wu, Zhiyi Zhang, Zhihong Chen, Jianquan Li, Xiang Wan, and Benyou Wang. Allava: Harnessing gpt4v-synthesized data for a lite vision-language model, 2024. 8
- [7] 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. 2, 6, 7, 8
- [8] Zhe Chen, Weiyun Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Erfei Cui, Jinguo Zhu, Shenglong Ye, Hao Tian, Zhaoyang Liu, et al. Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. *arXiv preprint arXiv:2412.05271*, 2024. 2, 5, 6, 8
- [9] Zhengxiao Du, Yujie Qian, Xiao Liu, Ming Ding, Jiezhong Qiu, Zhilin Yang, and Jie Tang. Glm: General language model pretraining with autoregressive blank infilling. In *ACL*, pages 320–335, 2022. 8
- [10] Agrim Gupta, Piotr Dollar, and Ross Girshick. LVIS: A dataset for large vocabulary instance segmentation. In *CVPR*, 2019. 2
- [11] Muyang He, Yexin Liu, Boya Wu, Jianhao Yuan, Yueze Wang, Tiejun Huang, and Bo Zhao. Efficient multimodal learning from data-centric perspective. *arXiv preprint arXiv:2402.11530*, 2024. 8
- [12] Wenyi Hong, Weihang Wang, Ming Ding, Wenmeng Yu, Qingsong Lv, Yan Wang, Yean Cheng, Shiyu Huang, Junhui Ji, Zhao Xue, et al. Cogvlm2: Visual language models for image and video understanding. *arXiv preprint arXiv:2408.16500*, 2024. 8
- [13] Xiuquan Hou, Meiqin Liu, Senlin Zhang, Ping Wei, and Badong Chen. Saliency detr: Enhancing detection transformer with hierarchical saliency filtering refinement. In *CVPR*, pages 17574–17583, 2024. 1, 5, 6
- [14] Anwen Hu, Haiyang Xu, Jiabo Ye, Ming Yan, Liang Zhang, Bo Zhang, Chen Li, Ji Zhang, Qin Jin, Fei Huang, et al. mplug-docowl 1.5: Unified structure learning for ocr-free document understanding. *EMNLP*, 2024. 8
- [15] Ding Jia, Yuhui Yuan, Haodi He, Xiaopei Wu, Haojun Yu, Weihong Lin, Lei Sun, Chao Zhang, and Han Hu. Detrs with hybrid matching. *CVPR*, 2023. 4, 6
- [16] Sahar Kazemzadeh, Vicente Ordonez, Mark Matten, and Tamara Berg. Referitgame: Referring to objects in photographs of natural scenes. In *EMNLP*, pages 787–798, 2014. 2
- [17] Ivan Krasin, Tom Duerig, Neil Alldrin, Vittorio Ferrari, Sami Abu-El-Haija, Alina Kuznetsova, Hassan Rom, Jasper Uijlings, Stefan Popov, Andreas Veit, Serge Belongie, Victor Gomes, Abhinav Gupta, Chen Sun, Gal Chechik, David Cai, Zheyun Feng, Dhyanes Narayanan, and Kevin Murphy. Openimages: A public dataset for large-scale multi-label and multi-class image classification. *Dataset available from <https://github.com/openimages>*, 2017. 2
- [18] Hugo Laurençon, Andrés Marafioti, Victor Sanh, and Léo Tronchon. Building and better understanding vision-language models: insights and future directions. *arXiv preprint arXiv:2408.12637*, 2024. 2, 8
- [19] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *ICML*, pages 19730–19742. PMLR, 2023. 2, 8
- [20] Zhang Li, Biao Yang, Qiang Liu, Zhiyin Ma, Shuo Zhang, Jingxu Yang, Yabo Sun, Yuliang Liu, and Xiang Bai. Monkey: Image resolution and text label are important things for large multi-modal models. In *CVPR*, 2024. 8
- [21] Chuang Lin, Yi Jiang, Lizhen Qu, Zehuan Yuan, and Jianfei Cai. Generative region-language pretraining for open-ended object detection. In *CVPR*, pages 13958–13968, 2024. 2, 8
- [22] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014. 2
- [23] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *NeurIPS*, 2023. 1, 2, 3, 4, 5, 6, 7
- [24] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, 2024. 8
- [25] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 2, 6, 8
- [26] Junyu Lu, Dixiang Zhang, Xiaojun Wu, Xinyu Gao, Ruyi Gan, Jiaying Zhang, Yan Song, and Pingjian Zhang. Ziya-

- visual: Bilingual large vision-language model via multi-task instruction tuning. *arXiv preprint arXiv:2310.08166*, 2023. 8
- [27] Chuofan Ma, Yi Jiang, Jiannan Wu, Zehuan Yuan, and Xiaojuan Qi. Groma: Localized visual tokenization for grounding multimodal large language models. *arXiv preprint arXiv:2404.13013*, 2024. 2, 6, 8
- [28] Matthias Minderer, Alexey Gritsenko, and Neil Houlsby. Scaling open-vocabulary object detection. *NeurIPS*, 36, 2024. 4, 7
- [29] Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran, Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Armand Joulin, and Piotr Bojanowski. Dino2: Learning robust visual features without supervision. *ICLR*, 2025. 7
- [30] 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 preprint arXiv:2306.14824*, 2023. 2, 6, 8
- [31] 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 *ICML*, pages 8748–8763, 2021. 7
- [32] Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M. Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S. Khan. Glamm: Pixel grounding large multimodal model. *CVPR*, 2024. 2
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *TPAMI*, 39(6):1137–1149, 2016. 2, 4, 6
- [34] Tianhe Ren, Qing Jiang, Shilong Liu, Zhaoyang Zeng, Wenlong Liu, Han Gao, Hongjie Huang, Zhengyu Ma, Xiaoke Jiang, Yihao Chen, Yuda Xiong, Hao Zhang, Feng Li, Peijun Tang, Kent Yu, and Lei Zhang. Grounding dino 1.5: Advance the "edge" of open-set object detection, 2024. 2
- [35] Shuai Shao, Zeming Li, Tianyuan Zhang, Chao Peng, Gang Yu, Xiangyu Zhang, Jing Li, and Jian Sun. Objects365: A large-scale, high-quality dataset for object detection. In *ICCV*, pages 8430–8439, 2019. 2, 5
- [36] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms, 2024. 8
- [37] Bin Wang, Chunyu Xie, Dawei Leng, and Yuhui Yin. Iaa: Inner-adaptor architecture empowers frozen large language model with multimodal capabilities. In *AAAI*, pages 21035–21043, 2025. 8
- [38] Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024. 2
- [39] Haoran Wei, Lingyu Kong, Jinyue Chen, Liang Zhao, Zheng Ge, Jinrong Yang, Jianjian Sun, Chunrui Han, and Xiangyu Zhang. Vary: Scaling up the vision vocabulary for large vision-language models. *arXiv preprint arXiv:2312.06109*, 2023. 8
- [40] Haoran Wei, Chenglong Liu, Jinyue Chen, Jia Wang, Lingyu Kong, Yanming Xu, Zheng Ge, Liang Zhao, Jianjian Sun, Yuang Peng, et al. General ocr theory: Towards ocr-2.0 via a unified end-to-end model. *arXiv preprint arXiv:2409.01704*, 2024.
- [41] Tianwen Wei, Liang Zhao, Lichang Zhang, Bo Zhu, Lijie Wang, Haihua Yang, Biye Li, Cheng Cheng, Weiwei Lü, Rui Hu, Chenxia Li, Liu Yang, Xilin Luo, Xuejie Wu, Lunan Liu, Wenjun Cheng, Peng Cheng, Jianhao Zhang, Xiaoyu Zhang, Lei Lin, Xiaokun Wang, Yutuan Ma, Chuanhai Dong, Yanqi Sun, Yifu Chen, Yongyi Peng, Xiaojuan Liang, Shuicheng Yan, Han Fang, and Yahui Zhou. Skywork: A more open bilingual foundation model. *arXiv preprint arXiv:2310.19341*, 2023. 8
- [42] Jiannan Wu, Muyan Zhong, Sen Xing, Zeqiang Lai, Zhaoyang Liu, Wenhai Wang, Zhe Chen, Xizhou Zhu, Lewei Lu, Tong Lu, et al. Visionllm v2: An end-to-end generalist multimodal large language model for hundreds of vision-language tasks. *NeurIPS*, 2024. 2, 6, 8
- [43] Zhiyu Wu, Xiaokang Chen, Zizheng Pan, Xingchao Liu, Wen Liu, Damai Dai, Huazuo Gao, Yiyang Ma, Chengyue Wu, Bingxuan Wang, Zhenda Xie, Yu Wu, Kai Hu, Jiawei Wang, Yaofeng Sun, Yukun Li, Yishi Piao, Kang Guan, Aixin Liu, Xin Xie, Yuxiang You, Kai Dong, Xingkai Yu, Haowei Zhang, Liang Zhao, Yisong Wang, and Chong Ruan. Deepseek-vl2: Mixture-of-experts vision-language models for advanced multimodal understanding, 2024. 8
- [44] Chunyu Xie, Heng Cai, Jincheng Li, Fanjing Kong, Xiaoyu Wu, Jianfei Song, Henrique Morimitsu, Lin Yao, Dexin Wang, Xiangzheng Zhang, et al. Ccmb: A large-scale chinese cross-modal benchmark. In *ACM MM*, pages 4219–4227, 2023. 8
- [45] Chunyu Xie, Bin Wang, Fanjing Kong, Jincheng Li, Dawei Liang, Gengshen Zhang, Dawei Leng, and Yuhui Yin. Fg-clip: Fine-grained visual and textual alignment. *arXiv preprint arXiv:2505.05071*, 2025.
- [46] Guowei Xu, Peng Jin, Hao Li, Yibing Song, Lichao Sun, and Li Yuan. Llava-cot: Let vision language models reason step-by-step, 2024. 8
- [47] Yuan Yao, Tianyu Yu, Ao Zhang, Chongyi Wang, Junbo Cui, Hongji Zhu, Tianchi Cai, Haoyu Li, Weilin Zhao, Zhihui He, et al. Minicpm-v: A gpt-4v level mllm on your phone. *arXiv preprint arXiv:2408.01800*, 2024. 2, 8
- [48] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, Chaoya Jiang, Chenliang Li, Yuanhong Xu, Hehong Chen, Junfeng Tian, Qian Qi, Ji Zhang, and Fei Huang. mplug-owl: Modularization empowers large language models with multimodality, 2023. *arXiv preprint arXiv:2304.14178*. 8

- [49] Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023. 2
- [50] Yufei Zhan, Yousong Zhu, Hongyin Zhao, Fan Yang, Ming Tang, and Jinqiao Wang. Griffon v2: Advancing multimodal perception with high-resolution scaling and visual-language co-referring, 2024. 6, 8, 1
- [51] Yufei Zhan, Yousong Zhu, Zhiyang Chen, Fan Yang, Ming Tang, and Jinqiao Wang. Griffon: Spelling out all object locations at any granularity with large language models. In *ECCV*, pages 405–422. Springer, 2025. 6, 8, 1
- [52] Hao Zhang, Hongyang Li, Feng Li, Tianhe Ren, Xueyan Zou, Shilong Liu, Shijia Huang, Jianfeng Gao, Lei Zhang, Chunyuan Li, et al. Llava-grounding: Grounded visual chat with large multimodal models. *arXiv preprint arXiv:2312.02949*, 2023. 2
- [53] Hang Zhang, Xin Li, and Lidong Bing. Video-llama: An instruction-tuned audio-visual language model for video understanding. *arXiv preprint arXiv:2306.02858*, 2023. 2
- [54] Haotian Zhang, Haoxuan You, Philipp Dufter, Bowen Zhang, Chen Chen, Hong-You Chen, Tsu-Jui Fu, William Yang Wang, Shih-Fu Chang, Zhe Gan, et al. Ferret-v2: An improved baseline for referring and grounding with large language models. *arXiv preprint arXiv:2404.07973*, 2024. 8
- [55] Yian Zhao, Wenyu Lv, Shangliang Xu, Jinman Wei, Guanzhong Wang, Qingqing Dang, Yi Liu, and Jie Chen. Detsr beat yolos on real-time object detection, 2024. CVPR. 3, 6
- [56] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. *arXiv preprint arXiv:2304.10592*, 2023. 2, 8
- [57] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. *ICLR*, 2021. 6