

DetToolChain: A New Prompting Paradigm to Unleash Detection Ability of MLLM

Yixuan Wu^{1,2,5*}, Yizhou Wang^{2*}, Shixiang Tang^{3(✉)}, Wenhao Wu⁴, Tong He²,
Wanli Ouyang^{2,3}, Philip Torr⁵, and Jian Wu^{1(✉)}

¹ Zhejiang University

² Shanghai Artificial Intelligence Laboratory

³ The Chinese University of Hong Kong

⁴ The University of Sydney

⁵ University of Oxford

wyx_chloe@zju.edu.cn

Abstract. We present DetToolChain, a novel prompting paradigm, to unleash the zero-shot object detection ability of multimodal large language models (MLLMs), such as GPT-4V and Gemini. Our approach consists of a detection prompting toolkit inspired by high-precision detection priors and a new Chain-of-Thought to implement these prompts. Specifically, the prompts in the toolkit are designed to guide the MLLM to focus on regional information (e.g., zooming in), read coordinates according to measure standards (e.g., overlaying rulers and compasses), and infer from the contextual information (e.g., overlaying scene graphs). Building upon these tools, the new detection chain-of-thought can automatically decompose the task into simple subtasks, diagnose the predictions, and plan for progressive box refinements. The effectiveness of our framework is demonstrated across a spectrum of detection tasks, especially hard cases. Compared to existing state-of-the-art methods, GPT-4V with our DetToolChain improves state-of-the-art object detectors by **+21.5%** AP_{50} on MS COCO `Novel class` set for open-vocabulary detection, **+24.23%** Acc on RefCOCO val set for zero-shot referring expression comprehension, **+14.5%** AP on D-cube describe object detection FULL setting. The codes shall be released upon acceptance.

Keywords: Multitmodal Large Language Model · Prompting · Detection

1 Introduction

Large language models (LLMs), *e.g.*, GPT3 [6], Gemini [45], InternLM [46], and Qwen [2], have shown unprecedented capabilities in understanding human languages and solving practical problems such as scientific question answering and code generation. When integrated with visual encoders, large language models can be upgraded to multimodal large language models (MLLMs), which can

* Equal Contributions.

achieve the ability similar to human visual intelligence and tackle some visual understanding tasks, such as image captioning. Despite these advances, the potential of MLLMs in detection tasks is still underestimated among common vision tasks [18, 48, 57, 65, 70, 74]. When required to ask precise coordinates in complicated object detection tasks, *e.g.*, detecting highly occluded objects, rotational objects, or small objects in the scene images, MLLMs often miss the target objects or answer inaccurate bounding boxes [65]. The poor performance on object detection significantly limits the applications of MLLMs in the real world, *e.g.*, defect detection [8, 17, 30, 77], clinical diagnosis [16, 60] and sports analysis [7, 47, 49, 53, 54].

To enhance the detection capabilities of MLLMs, prior efforts can be categorized into two classes: (1) Finetuning MLLMs with high-quality question-answer instructions with abundant location information [3, 12, 29, 29, 37] in the answer. Despite the considerable improvements achieved, preparing high-quality question-answer pairs requires great manual efforts and finetuning multimodal large language models suffers from large computational costs. Furthermore, since current state-of-the-art MLLMs [1, 3, 45] are closed-sourced and their performances have been significantly superior to open-sourced models, the “finetuning” method can not be implemented on the most powerful MLLMs at the moment (and most likely in the future), which significantly limits its potential to continuously improve the emerging state-of-the-art MLLMs. (2) Designing textual or visual prompting with location information to advance the localization ability of MLLMs. While intuition-based prompting methods have greatly advanced the performance of regional comprehension tasks such as compositional reasoning [36] and spatial understanding [10], their effectiveness on detection tasks remains underexplored.

This work explores how the detection ability of multimodal large language models can be unlocked by a new chain of thoughts on detection prompting toolkits (dubbed as DetToolchain). The new DetToolchain is motivated by three ideas. First, visual prompts are identified as a crucial component of the detection prompting toolkits. They offer a more direct and intuitive approach to enhancing the spatial comprehension of Multimodal Large Language Models (MLLMs) compared to language prompts. This is because current MLLMs still struggle with accurately translating textual coordinates and descriptions into precise regions and visual information. Instead, visual prompts directly drawn in the image can significantly narrow the gap between visual and textual information and ultimately contribute to the improved detection ability of MLLMs. Second, detecting challenging instances, such as occluded and small objects, can be more efficiently tackled by breaking them down into smaller, simpler subtasks. Third, the detection results should be modified step by step using Chain-of-Thought, similar to the progressive refinement of bounding boxes in current state-of-the-art object detection algorithms such as DETR [9], SparseRCNN [44], and DiffusionDet [14]. Based on these ideas, our proposed DetToolChain consists of a detection prompts toolkit, including *visual processing* prompts and *detection reasoning* prompts, and a multimodal Chain-of-Thought to properly apply these

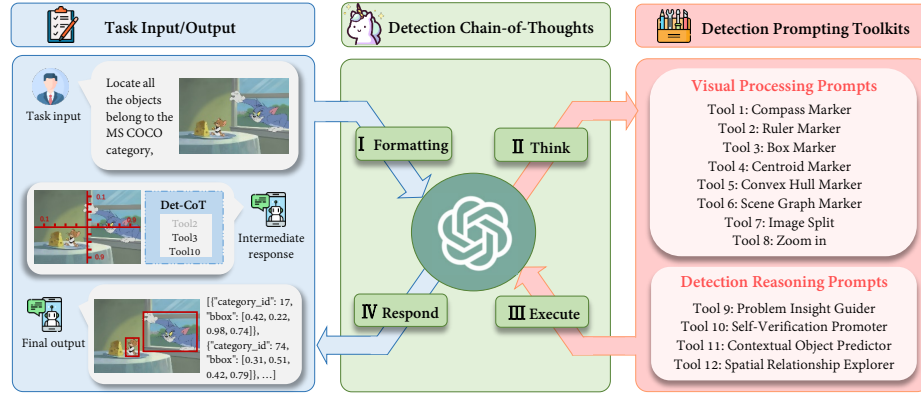


Fig. 1: The overall framework of DetToolChain for unleashing the detection potential of MLLMs. *e.g.*, GPT-4V. Given a query image, the MLLM is instructed to: I. format the task raw input into a suitable instruction template, II. decompose the specific detection task into simpler subtasks and choose effective prompts from Detection Prompting Toolkits, III. execute the specific prompt sequentially, IV. after verifying the predictions, return the response.

detection prompts to unleash the detection ability of MLLMs. Their critical designs and insights involve:

(1) **A comprehensive set of *visual processing* prompts that support a wide range of detection tasks.** The *visual processing* prompts process the given images to facilitate the detection performance of MLLMs, according to some well-accepted prior knowledge and techniques in effective detectors. Specifically, the *visual processing* prompts can be divided into three categories, *i.e.*, the regional amplifier, the spatial measurement standard, and the scene image parser. These prompts facilitate different key factors for better detectors. First, the regional amplifier consists of image splitting and zooming in, capable of highlighting the region of interest in detection tasks. Second, the spatial measurement standard includes rulers and compass with linear graduations, which provide transnational references and rotational references in object detection (particularly in rotated object detection), respectively. Similar to human intelligence, these spatial measurement standard can help MLLMs to locate the objects and read their coordinates out. Third, the scene parser marks the predicted positions or spatial relations of objects in the images by convex hull, bounding boxes of objects and scene graphs of the image. The markers by the scene parser can facilitate the detection capability of MLLMs by encouraging them to reason from contextual information in the scene images.

(2) **A comprehensive set of *detection reasoning* prompts that help MLLMs to diagnose the detection results and reason the next *visual processing* prompts to be applied.** Different from *visual processing* prompts, the *detection reasoning* prompts do not process the images but are dedicated to

evaluating the predicted bounding boxes and diagnosing the inaccurate predictions even using the *visual processing* prompts. By analyzing the relationship and reasoning the co-occurrence of detected objects in the scene image with the commonsense knowledge in MLLM, these prompts help the MLLM to prevent hallucination on detection tasks.

(3) A multimodal detection Chain-of-Thought (Det-CoT) that enables the MLLM to manage the whole process of detecting targets. With these powerful detection prompting toolkits, the detection Chain-of-Thought helps the MLLM comprehend spatial detection tasks and produce reliable bounding boxes. As illustrated in Fig. 1, the multimodal detection Chain-of-Thought are instructed to (1) format the raw input with a suitable instruction template, (2) decompose the complex detection task into smaller sub-tasks and select the corresponding detection prompting toolkit, (3) execute detection prompting toolkit iteratively, and (4) apply its own reasoning and critical thinking to oversee the whole detection process and return the final response. Given a detection task, our proposed Det-CoT can automatically decompose it into multimodal subtasks and refine the predicted bounding boxes progressively.

DetToolChain allows MLLMs to support various detection tasks without instruction tuning. Concretely, it significantly improves baseline GPT-4V and Gemini by **20%-50%** on main detection metrics of open-vocabulary detection, described object detection, referring expression comprehension, and oriented object detection. Compared to existing state-of-the-art methods, GPT-4V with our DetToolChain improves state-of-the-art object detectors by **+21.5% AP_{50}** on MS COCO Novel class set for open-vocabulary detection, **+24.23% Acc** on RefCOCO val set for zero-shot referring expression comprehension, **+14.5% AP** on D-cube describe object detection FULL setting. Notably, our method without trained on COCO train2017 set even achieves better AP_{50} on COCO val2017 than the DETR trained on COCO train2017 set. To summarize, our contributions are two-fold: (1) We propose a new prompting paradigm to instruct the MLLM to manage object detection by applying tools in the detection prompting toolkit with a multimodal detection Chain-of-Thought. (2) We propose detection prompting toolkits including *visual processing* prompts and *detection reasoning* prompts to facilitate MLLMs on detection tasks. We showcase that our method achieves remarkable performance on a range of detection tasks.

2 Related Work

Multimodal Large Language Models on Detection Tasks. The implementation of Multimodal Large Language Models (MLLMs) in detection tasks has been a recent focus. The widely used strategy is to finetune MLLMs using a high-quality image-text instruction tuning dataset consisting of detection problems. To promote the detection performance, BuboGPT [80], Kosmos-2 [37], Shikra [12], Qwen-VL [3], SPHINX [29], and Ferret [71] constructed instruction datasets with high-quality question-and-answer pairs about detection. Consequently, they consumed great manual efforts and heavy computing costs and

suffered from unsatisfied generalization ability to detect objects unseen in the instruction dataset. Furthermore, since most state-of-the-art MLLMs are closed-sourced [1, 3, 45], it is infeasible to apply instruction tuning on these MLLMs. In this paper, we believe the current state-of-the-art MLLM is potentially a zero-shot detector with proper promptings, and therefore design a Chain-of-Thought with new detection prompts to unleash their potential on detection tasks.

Visual Prompting. Visual prompting, inspired by textural prompting in Natural Language Processing [6, 38], manipulates images to improve the visual perception ability of MLLMs, such as classification and spatial reasoning. For example, RedCircle [42] used a circle marker to guide the model to focus on the specific region for fine-grained classification while FGVP [64], SCAFFOLD [27] and SOM [63] explored prompts for spatial reasoning with dot matrices or pre-trained models such as SAM [25]. Although CPT [69], SOM [63] are claimed to perform well on the visual grounding task and referring expression comprehension task, they just choose one of the pre-extracted bounding boxes and segmentation masks by high-quality detectors [72] and segmentors [25] for the correct answer, which is essential not the ability of object detection. To unleash the detection potential of MLLMs, we propose a general detection prompting toolkit that incorporates some well-adopted prior knowledge and strategies but does not include any pretrained model for detection or segmentation.

Multimodal Chain-of-thought. Chain-of-Thought [55, 78] methods and their variants [5, 26, 67, 68] have greatly boosted the reasoning ability of large language models (LLMs). By encouraging LLMs to follow a human-like, step-by-step reasoning process, these methods decompose complex tasks into easier subtasks and solve them sequentially to get the final answer. However, the original CoT shows inferior performance [36] on vision-language tasks, *e.g.*, compositional reasoning. Therefore, recent researches pivoted towards developing multimodal Chain-of-Thought methods, *e.g.*, Multimodal CoT [79], DDCoT [81], HoT [66] for multimodal reasoning, Chameleon [33], Compositional CoT [36] for compositional reasoning, and Spatial CoT [10] for spatial comprehension. Nevertheless, there remain multimodal Chain-of-Thought methods tailored for the detection ability of MLLMs unexplored. Our research fills this gap by introducing a detection-tailored Chain-of-Thought approach, coupled with a detection prompting toolkit. Our method differs from other practices by in visual promptings that use detection priors to progressively refine detection results through the CoT process.

3 Method

While state-of-the-art MLLMs, *e.g.*, GPT-4V and Gemini, exhibit commendable reasoning and recognition capabilities, their detection proficiency is still **unleashed**. To unlock the potential of MLLMs in object detection, we introduce a comprehensive detection prompts toolkit (Sec. 3.2) including *visual processing* prompts and *detection reasoning* prompts, and a detection Chain-of-Thought (Det-CoT) to reason the sequential implementation of the detection prompts in the toolkit (Sec. 3.1).

3.1 Detection Chain-of-Thought for Object Detection

Let \mathbb{L} and \mathbb{I} denote the set of finite language strings and images, respectively. Given a test-time query $x = \{x_l, x_i | x_l \in \mathbb{L}, x_i \in \mathbb{I}\}$, which can be a specific detection task described in images and natural languages, we aim to obtain corresponding textual outputs with the help of a frozen MLLM. This model, like GPT-4V [1], also maintains a prompt history that may include a list of previous messages, symbolized by \mathcal{H} , and produces a corresponding textual response y . The textual response y includes the detection outputs, the diagnosis of current outputs and the suggestions of next prompts in the toolkit $\text{TK} = \{T_1, \dots, T_n\}$ to be called, where T_i is the i -th tool pretrained in the toolkit. Finally, we define string extractors e , which is designed to retrieve a substring that is enclosed within specific delimiters.

Algorithmic Procedure. As illustrated in Algorithm 1, we provide a conceptual overview of the procedure below:

1. **Formatting the Test-time Query:** Using the transformation function f , the raw query is replaced with a suitable template for the MLLM.
2. **Loop Iteration in each prompting step:**
 - (a) **Proposing a thought based on the response y_s :** The current message history, named \mathcal{H}_s , guides the MLLM to propose a thought - either to end it here and return the final outputs, or to continue the algorithm and engage another visual prompting tool.
 - (b) **Engaging detection prompting tool:** If the MLLM does not return the result, it can select the appropriate visual prompting tool and construct it with corresponding instructions. The constructed prompting will be appended with the message history \mathcal{H}_s and be sent to the next iteration.
3. **Returning the Final Response:** If the string extractor e discovers the final detection with special markers in the MLLM’s response y_s , the algorithm will be ended and return the extracted bounding boxes.

3.2 Detection Prompting Toolkits

Compared with language promptings, visual promptings are more intuitive and can effectively teach MLLMs prior knowledge for detection. Following this insight, we design a comprehensive set of visual prompting toolkits to enhance various detection tasks. The toolkits are summarized in Fig. 2 and Tab. 1 and described as follows. Detailed descriptions and usage instructions can be found in the *supplementary materials*.

Visual Processing Prompts. We introduce *visual processing* prompts to pre-process the input image to enhance the detection performance of MLLMs. Inspired by prior techniques for detection, we design our *visual processing* prompts (see Fig. 2) with three types focusing on different aspects, *i.e.*, better visibility on details, more accurate spatial reference, and better contextual comprehension of the given images.

Algorithm 1 Detection Chain-of-Thought (Det-CoT)

Input: $x = \{x_l, x_i | x_l \in \mathbb{L}, x_i \in \mathbb{I}\}; \text{thought} \in \mathbb{L}; \text{TK} = \{T_1, \dots, T_n\}; \text{MLLM} : \mathbb{L}, \mathbb{I} \rightarrow \mathbb{L}; f, e : \mathbb{L} \rightarrow \mathbb{L}; \text{CoT step } s; \text{string concatenation } \oplus$

```

1:  $\mathcal{H}_1 \leftarrow f(x)$  ▷ Format the test-time query to templated input
2: while true do
3:    $y_s \leftarrow \text{MLLM}(\mathcal{H}_s)$ 
4:    $\text{thought}_s \leftarrow e(y_s)$  ▷ Propose a thought in current step  $s$ 
5:   if  $\text{thought}_s$  is not ending in this step then
6:      $T_n \leftarrow \text{thought}_s \times \text{TK}$  ▷ Select the appropriate visual prompting tool
7:      $\mathcal{H}_{s+1} \leftarrow T_n(\mathcal{H}_s) \oplus \mathcal{H}_s$  ▷ Execute visual prompting tool
8:   else
9:     return  $e(y_s)$  ▷ Respond final output until all conditions are satisfied
10:  end if
11: end while

```

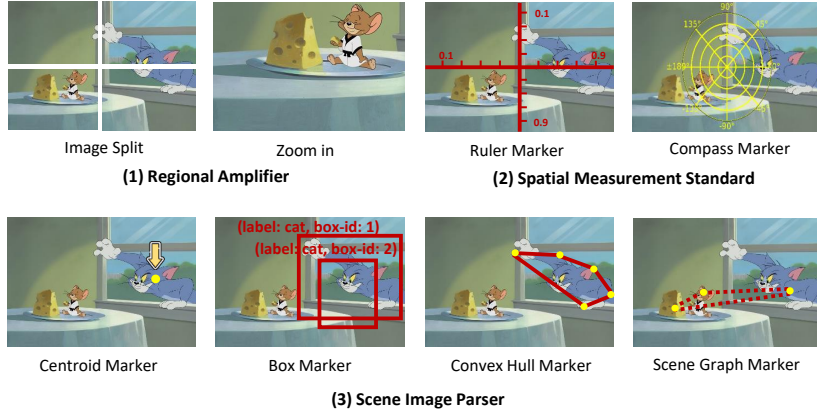


Fig. 2: Illustration of visual processing prompts. We design Regional Amplifier, Spatial Measurement Standard, and Scene Image Parser to facilitate the detection ability of MLLMs from different perspectives.

1. **Regional Amplifier** aims at enhancing the visibility of the region of interest for MLLMs. Specifically, **Image Split** crops the image into disparate parts, and **Zoom in** enables close-up inspection of specific regions in the image.
2. **Spatial Measurement Standard** provides a more explicit reference for object detection by overlaying rulers and compasses with linear graduations on the original image, as depicted in Fig. 2(2). The auxiliary ruler and compass enable MLLMs to output accurate coordinates and angles with the help of translational and rotational references overlaid in the image. Essentially, this standard simplifies the detection task, allowing MLLMs to read out the coordinates of the objects instead of directly predicting them.
3. **Scene Image Parser** marks the predicted location or relations of objects, enabling a more localized understanding of images using spatial and contextual information. The proposed scene image parser can be divided into

Table 1: Detection Reasoning prompts and their effects.

Name	Effect
Problem Insight Guider	Find potential detection problems in the image and provide advice.
Spatial Relationship Explorer	Leverage spatial reasoning ability to analyze the relationship of objects.
Contextual Object Predictor	Based on common sense, inquire the model what objects are likely to co-occur in the image.
Self-Verification Promoter	Verify the detection results itself to ensure consistent predictions and decide when to end the algorithm.

two categories. First, we mark the predicted objects with centroids, convex hulls and bounding boxes with label names and box indices. These markers represent the object location information in different formats, enabling the MLLM to detect diverse objects with different shapes and backgrounds, especially with irregular shapes or significant occlusions. For example, **Convex Hull Marker** marks the boundary points of objects and connects them as the convex hull to enhance the detection performance when objects have very irregular shapes. Second, we mark the scene graph by connecting the center of different objects with **Scene Graph Marker** to highlight the relationships of objects in the image. Based on the scene graph, the MLLM can leverage its contextual reasoning abilities to refine its predicted boxes and avoid hallucinations. For instance, as shown in Fig. 2(3), Jerry Mouse is going to eat cheese, so their bounding boxes should be pretty close to each other.

Detection Reasoning Prompts. To improve the reliability of predicted boxes, we conduct *detection reasoning* prompts (illustrated in Tab. 1) to examine the predictions and diagnose the potential problem that may exist. Based on these intuitions, we first propose **Problem Insight Guider**, which highlights difficult issues and offers effective detection suggestions for query images. For example, given Fig. 3, the prompt will point out the issue of small object detection and suggest solving it by zooming in the region of the surfboard. Second, to leverage the inherent spatial and contextual ability of MLLMs, we design **Spatial Relation Analyzer** and **Contextual Object Predictor** to ensure the detection results accord with common sense. As shown in Fig. 3, surfboards may co-occur with the ocean (contextual knowledge), and the surfing person should have a surfboard very close to his feet (spatial knowledge). Furthermore, we apply **Self-Verification Promoter** to enhance the consistency of multi-round responses. To further promote MLLMs’ reasoning ability, we engage in widely adopted prompting methods, *e.g.*, debating [20, 83] and self-debugging [15, 40]. Detailed descriptions are listed in the *supplementary materials*.



Fig. 3: Detection reasoning prompts can help diagnose small object detection problems, *e.g.*, using common sense to locate the surfboards under people’s feet and encouraging the model to detect surfboards in the ocean.

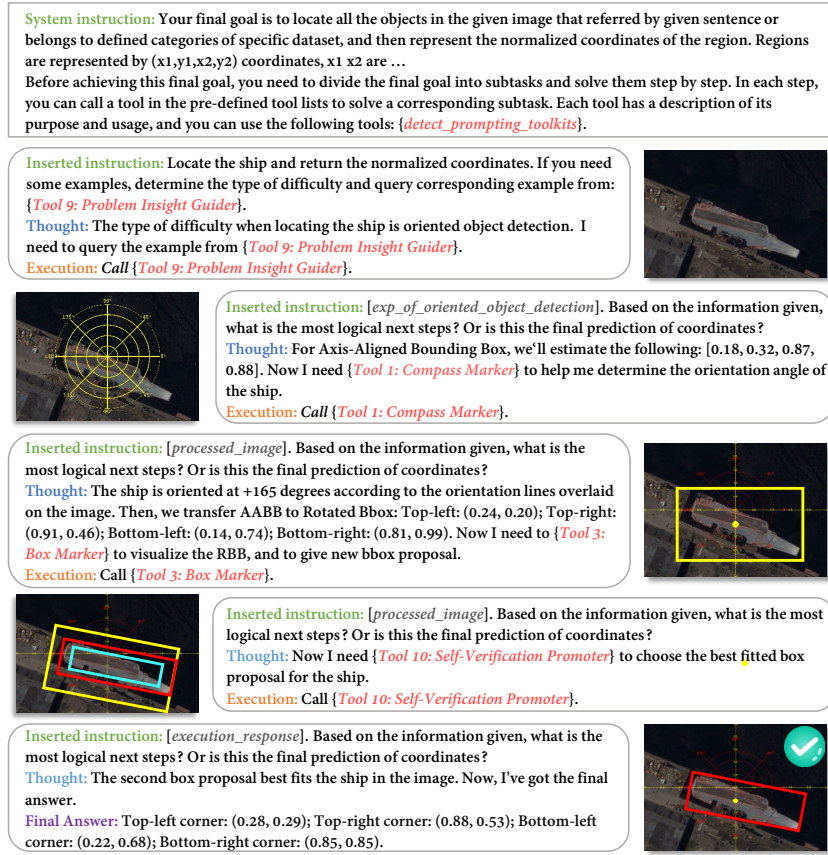


Fig. 4: An example of DetToolChain for oriented object detection on HRSC2016.

4 Experiments

4.1 Experimental Setup

Benchmarks. To assess the effectiveness of our proposed methodology, we undertook a comprehensive evaluation across a diverse array of tasks and datasets related to object detection. We also extend our method to Referring expression comprehension which also needs to localize objects referred to in the expression. These evaluations encompass:

- **Object detection** utilize the MS COCO 2017 dataset [28], with results reported on the validation subset by default.
- **Described object detection** on the D-cube dataset [59, 61], aiming to confirm the presence of objects described in arbitrary open-set expressions and to localize them accordingly.

- **Referring expression comprehension**, a multimodal task that involves grounding a referent based on a given expression. This was tested on three representative benchmarks: RefCOCO [73], RefCOCO+ [73], and RefCOCOg [34].
- **Oriented object detection** on the widely recognized HRSC2016 dataset [32], aiming to identify and locate objects of arbitrary orientation within images.

Models. We utilized the following Multimodal Large Language Models (MLLMs) in our experiments to explore the performance of our proposed DetToolChain:

- GPT-4 [1], specifically the `gpt-4-vision-preview` model.
- Gemini [45], specifically the `gemini-pro-vision` variant.

Implementations. As shown in Fig. 4, the system instruction is set to ensure MLLM to comprehend its ultimate task objective and the preceding steps. Subsequently, the task description and input image are provided as the initial round of inserted instruction. Upon reaching the final answer, an automatic function extracts coordinate results, streamlining the prediction collection process. To accommodate varying image sizes across datasets, we performed a transformation from absolute to normalized coordinate axes. The returned normalized coordinates are ultimately converted back to absolute ones for standard evaluation. For comparison, we implement baselines of GPT-4V and Gemini. We provide foundational prompt to enable successful result return (detailed prompt information available in the *supplementary materials*). Across all experiments, we uniformly applied identical parameters and instructions, setting the maximum token count to 1024, with other settings at default values.

4.2 Experimental Results

Open Vocabulary Object Detection As shown in Tab. 2, we evaluate our method on open vocabulary detection (OVD), where the AP_{50} results on 17 novel classes, 48 base classes, and all classes of the COCO OVD benchmark [4] are reported. With our DetToolChain, the performance of GPT-4V and Gemini are both remarkably enhanced. GPT-4V+DetToolChain significantly outperforms the state-of-the-art method CORA [58] by **22.7** AP_{50} , **4.1** AP_{50} , and **9.0** AP_{50} on Novel, Base and All classes, respectively. Benefiting from general comprehension abilities, MLLMs can originally recognize most objects in the image, including those belonging to novel classes. However, they fall short in predicting precise textural coordinates. When integrated with our DetToolChain, the MLLM is instructed to locate the position with overlaying measure standards. This enhancement allows the MLLM to recognize and accurately locate objects across novel classes, effectively transforming it into a super open-vocabulary detector.

To position our method in previous state-of-the-art object detectors, we further compare our zero-shot detection performance with their detection performance by the model trained on COCO `train2017` in Tab. 3. While the baseline performance of GPT-4V and Gemini are only AP is 0.2 and 0.3 respectively, there is a significant performance improvement after employing our DetToolChain, *i.e.*, Gemini’s AP , AP_{50} , and AP_{75} increasing to 30.6, 58.0, and

Table 2: Main results on the COCO OVD benchmark. We report AP_{50} .

Methods	Novel	Base	All
OCR-CNN [76]	22.8	46.0	39.9
ViLD [21]	27.6	59.9	51.3
OV-DETR [75]	29.4	61.0	52.7
MEDet [13]	32.6	54.0	49.4
CFM-ViT [24]	34.1	-	46.0
OADP [50]	35.6	55.8	50.5
RegionCLIP [82]	39.3	61.6	55.7
BARON [56]	42.7	54.9	51.7
CORA [58]	43.1	60.9	56.2
Gemini (baseline) [45]	2.2	0.9	1.2
Gemini + DetToolChain	63.6	58.5	59.8
GPT-4V (baseline) [1]	2.3	1.2	1.5
GPT-4V + DetToolChain	65.8	65.0	65.2

Table 3: Performance comparisons on COCO val2017 for object detection.

Methods	AP	AP_{50}	AP_{75}
<i>Supervised methods</i>			
Faster R-CNN-R50 [39]	36.4	58.4	39.1
Mask R-CNN-R50 [23]	37.3	59.0	40.2
DETR-R50 [9]	42.0	62.4	44.2
VisionLLM-R50 [52]	44.6	64.0	48.1
UNINEXT-R50 [62]	51.3	68.4	56.2
Co-DETR-L [84]	60.7	78.5	66.7
<i>Zero-Shot methods</i>			
Gemini (baseline) [45]	0.2	1.1	0.0
Gemini+DetToolChain	30.6	58.0	27.9
GPT-4V* (baseline) [1]	0.3	1.4	0.0
GPT-4V+DetToolChain	34.5	64.8	31.5

27.9, and GPT-4V’s AP , AP_{50} and AP_{75} increasing to 34.5, 64.8, and 31.5. Notably, GPT-4V with DetToolChain surpasses Faster R-CNN on AP_{50} by **6.4**, and even achieves better performance than recent detectors, *e.g.*, DETR-R50 [9] and VisionLLM-R50 [52], demonstrating its potential to be a general object detector in the real world. The higher AP_{50} but lower AP_{75} than state-of-the-art detectors, *e.g.*, Mask R-CNN and DETR, indicates that our method enables MLLMs to detect most objects in a zero-shot manner but with less accurate box boundaries because we consider accurate boundaries involves annotation priors in each dataset and can only be regressed after learning the model on the training set.

Described Object Detection

Described Object Detection [61] is a new task that can be regarded as a superset of open-vocabulary detection and referring expression comprehension. The dataset D-cube [59, 61] in the task presents two primary challenges: (1) The presence of negative samples, where the object referred by sentences does not necessarily exist within the given image. This demands the model’s capability to understand

semantically incongruent concepts. (2) The referring sentences are exceedingly complex, often containing the use of negation, such as "*the person riding a horse who is **not** wearing a hat*", evaluating the model’s comprehension of complex and lengthy sentences. The experiment evaluates the model’s ability across three categories: FULL, PRES, and ABS, which respectively mean the full descriptions

Table 4: Performance comparisons on the described object detection task.

Methods	FULL	PRES	ABS
OFA-L [51]	4.2	4.1	4.6
OWL-ViT-L [35]	9.6	10.7	6.4
G-DINO-B [31]	20.7	20.1	22.5
UNINEXT-H [62]	20.0	20.6	18.1
SPHINX-7B [29]	10.6	11.4	7.9
OFA-DOD-B [61]	21.6	23.7	15.4
FIBER-B [19]	22.7	21.5	26.0
Gemini (baseline) [45]	14.3	14.2	14.2
Gemini + DetToolChain	35.9	35.9	35.8
GPT-4V (baseline) [1]	16.3	16.3	16.2
GPT-4V + DetToolChain	37.2	37.5	36.2

Table 5: Performance comparisons on the referring expression comprehension task. We report Top-1 Accuracy@0.5 (%) as the evaluation metric.

Methods	RefCOCO			RefCOCO+			RefCOCOg	
	val	test-A	test-B	val	test-A	test-B	val-u	test-u
<i>Supervised sota</i>								
UNINEXT [62]	92.64	94.33	91.46	85.24	89.63	79.79	88.73	89.37
G-DINO-L [31]	90.56	93.19	88.24	82.75	88.95	75.92	86.13	87.02
MiniGPT-v2-7B [11]	88.69	91.65	85.33	79.97	85.12	74.45	84.44	84.66
Shikra-13B [12]	87.83	91.11	81.81	82.89	87.79	74.41	82.64	83.16
SPHINX [29]	89.15	91.37	85.13	82.77	87.29	76.85	84.87	83.65
Qwen-VL-7B [3]	89.36	92.26	85.34	83.12	88.25	77.21	85.58	85.48
<i>Zero-Shot methods</i>								
CPT-Seg [69]	32.20	36.10	30.30	31.90	35.20	28.80	36.70	36.50
CPT-adapted [43]	23.79	22.87	26.03	23.46	21.73	26.32	21.77	22.78
GradCAM [41]	42.85	51.07	35.21	47.83	56.92	37.70	50.86	49.70
ReCLIP [43]	45.78	46.10	47.07	47.87	50.10	45.10	59.33	59.01
StructureSimi [22]	49.37	47.76	51.68	48.89	50.02	46.86	60.95	59.99
GPT-4V (baseline)	25.48	26.22	24.39	10.59	18.23	8.87	14.26	15.42
GPT-4V + DetToolChain	70.01	72.33	67.24	57.32	62.74	52.28	63.48	64.54

(422 categories), presence descriptions (316 categories), and absence descriptions (106 categories).

As shown in Tab. 4, with our DetToolChain method, both GPT-4V and Gemini significantly outperform existing approaches, *e.g.*, GPT-4V + DetToolChain improve the best method, FIBER-B [19], by 14.5, 16.0, and 10.2 *AP* in the FULL, PRES, and ABS settings, respectively. Such improvements come from two aspects. First, Gemini and GPT-4V have stronger ability to comprehend detection-related descriptions, especially the negative concepts, and get rid of unrelated concepts in both PRES and ABS settings than other models, *e.g.*, OFA [51], OWL-ViT [35]. Second, our proposed DetToolChain significantly unleash the ability of MLLMs to detect the described objects by prompts in the detection toolkit and the new detection Chain-of-Thought.

Referring Expression Comprehension To demonstrate the effectiveness of our method on referring expression comprehension, we compare our approach with other zero-shot methods on the RefCOCO, RefCOCO+, and RefCOCOg datasets (Tab. 5). First, our DetToolChain improves GPT-4V baseline by 44.53%, 46.11%, and 24.85% on val, test-A, and test-B, respectively, which exhibits the best zero-shot referring expression comprehension performance on RefCOCO. Since the referring sentences in RefCOCO contain a substantial number of directional words, such as "*right*" and "*left*", our designed *visual processing* prompts, *e.g.*, **Ruler Marker**, can help MLLMs to more effectively understand the spatial rationale and detect target objects. Furthermore, while other zero-shot methods [22, 43] introduce an out-of-the-shelf object detector [72] to generate proposals, our DetToolChain instructs MLLM to predict and refine the detection boxes with their inherent spatial and contextual ability, which is more simple and adaptable in real deployments. In addition, our method significantly narrows

Table 6: Ablation studies of our proposed visual prompting tools and comparison with other Chain-of-Thought methods. **Veri.**, **SG**, and **Co-Obj.** denote **Self-Verification Promoter**, **Scene Graph Marker**, and **Contextual Object Predictor**, respectively.

Methods	MS COCO val2017 subsets										HRSC2016 test set	
	Single object		Multi-instances		Multi-objects		Occlusion object		Small object		Oriented Object	
	mIoUp	mIoUr	mIoUp	mIoUr	mIoUp	mIoUr	mIoUp	mIoUr	mIoUp	mIoUr	mIoUp	mIoUr
(a) Gemini (baseline)	0.19	0.19	0.25	0.13	0.17	0.09	0.22	0.17	0.13	0.07	0.16	0.09
(b) GPT-4V (baseline)	0.23	0.23	0.25	0.15	0.18	0.11	0.24	0.18	0.16	0.10	0.13	0.08
(c) GPT-4V + Ruler	0.52	0.52	0.40	0.29	0.30	0.21	0.43	0.25	0.22	0.18	0.32	0.23
(d) GPT-4V + Compass	0.22	0.22	0.27	0.16	0.19	0.10	0.28	0.17	0.14	0.06	0.47	0.33
(e) GPT-4V + Split + Zoom	0.43	0.43	0.52	0.43	0.45	0.37	0.43	0.33	0.46	0.38	0.34	0.28
(f) GPT-4V + Ruler + Split + Zoom	0.64	0.64	0.61	0.54	0.52	0.47	0.55	0.49	0.54	0.47	0.49	0.41
(g) GPT-4V + Ruler + Convex + Veri.	0.69	0.69	0.59	0.54	0.49	0.43	0.70	0.61	0.48	0.42	0.51	0.44
(h) GPT-4V + Ruler + SG + Relation	0.60	0.60	0.60	0.54	0.53	0.48	0.64	0.59	0.49	0.44	0.47	0.40
(i) GPT-4V + Ruler + Co-Obj. + Veri.	0.66	0.66	0.54	0.50	0.51	0.48	0.59	0.53	0.52	0.48	0.41	0.33
(j) GPT-4V + ZS-CoT	0.25	0.25	0.27	0.18	0.19	0.11	0.33	0.23	0.17	0.12	0.16	0.09
(k) GPT-4V + FS-CoT	0.37	0.37	0.28	0.20	0.22	0.10	0.38	0.22	0.18	0.13	0.34	0.28
(l) GPT-4V + MM-CoT	0.29	0.29	0.38	0.29	0.32	0.24	0.45	0.33	0.22	0.19	0.29	0.24
(m) Gemini + DetToolChain	0.75	0.75	0.67	0.61	0.56	0.51	0.77	0.70	0.52	0.44	0.62	0.56
(n) GPT-4V + DetToolChain	0.78	0.78	0.75	0.69	0.64	0.60	0.78	0.72	0.61	0.57	0.63	0.59

the gap between zero-shot methods and instruction-tuned MLLMs [3, 11, 12, 29], indicating the promising direction of designing new multimodal promptings.

Oriented Object Detection To investigate the efficacy of our DetToolChain in detecting rotated objects, we conducted experiments on the HRSC2016 test set. The images in the HRSC2016 dataset contain several ships of varying orientations. We employ $mIoU_{precision}$ and $mIoU_{recall}$ as performance metrics. $mIoU_p$ reflects the precision in bounding box coordinates prediction for objects that have been identified and $mIoU_r$ evaluate recall for the model’s ability to recognize all targeted objects (definitions detailed in *supplementary materials*).

As shown in the last two columns of Tab. 6, our DetToolChain significantly improves the baseline results of Gemini and GPT-4 by 0.46 $mIoUp$ and 0.47 $mIoUr$ on Gemini, and 0.50 $mIoUp$ and 0.51 $mIoUr$ on GPT-4V. Furthermore, compared with other combinations of detection prompts in Tab. 6, **Compass Marker** brings the most substantial enhancements on oriented object detection, 0.34 $mIoUp$ and 0.25 $mIoUr$ performance gains over the GPT-4 baseline, which is consistent with our expectations. With a compass drawn by **Compass Marker** on the query image, it is easier for MLLMs to read the angle out, which will ultimately improve the accuracy of ship angle predictions.

4.3 Ablation Study and Further Analysis

To further analyze the individual improvement of our proposed visual prompting tools, we classify the samples in COCO val2017 into several detection scenarios:

- Single: Images containing only a single object.
- Multi-instances: Images containing multiple instances of a single category.
- Multi-objects: Images containing multiple objects of various categories.
- Occluded objects: Images with objects significantly occluded (occlusion area greater than 50%).

- Small objects: Images with objects smaller than 32x32 pixels.

The comparison with other CoTs? In Tab. 6, we compare our DetToolChain with other Chain-of-Thought approaches, including: (1) Zero-shot CoT (ZS-CoT), initially proposed by adding "*Let's think step by step*" to the original prompt, is adapted for our detection task as "*Let's detect and localize target objects step by step*". (2) Few-shot CoT (FS-CoT), which provides demonstrations with intermediate reasoning steps in the prompt. (3) Multimodal CoT (MM-CoT), which introduces a two-stage process: generating a rationale based on the input of image and text, and performing inference based on the generated rationale. Among them, MM-CoT works better in scenarios with multiple instances, multiple objects, and occluded objects. While MM-CoT has managed to grasp the spatial relationships within images to some extent, it has not been directly applied to the localization of each target. However, by considering the distinct characteristics of each image, our DetToolChain marks spatial relationships directly on the images and explicitly uses them to detect every target, thereby achieving superior performance.

Which tool is the best to use? Tab. 6 shows the effectiveness of various detection prompting tools. It is observed in exp-c and exp-d that **Ruler Marker** and **Compass Marker** enable a more intuitive grasp of translational and rotational scale for MLLM, and therefore benefit general detection and oriented object detection, respectively. For small object detection, we have two effective tools: **Image Split + Zoom in** to help MLLM select the region of interest and focus on fine-grained information in this region (exp-e); and **Contextual Object Predictor** to reduce the missing rate of detection by the commonsense in LMMs (e.g., books highly exist in the library) (exp-i). For multi-instance and multi-object detections, **Scene Graph Marker** and **Spatial Relationship Explorer** can improve the detection performance by using the spatial rationale among multiple objects (exp-h). For occluded targets, leveraging **Convex Hull Marker** and **Self-Verification Promoter** can contour the irregular visible parts of objects, thereby referring accurate bounding boxes (exp-g). Moreover, our DetToolChain can achieve better performance by strategically applying a combination of different detection prompting tools (exp-m & exp-n).

5 Conclusion

We have presented DetToolChain, a groundbreaking prompting paradigm, to unleash the potential of multimodal large language models (MLLMs), like GPT-4V and Gemini, as zero-shot detectors. By employing a multimodal detection Chain-of-Thought (Det-CoT) to manage the proposed *visual processing* prompts and *detection reasoning* prompts, MLLMs are instructed to perceive objects by regional focus, prediction progressive refinement, and contextual inference. Our approach not only improves interpretability and precision in detection tasks across a variety of challenging scenarios, including occluded, small, and oriented objects, but also sets new records in open-vocabulary detection, described object detection, and referring expression comprehension without instruction tuning.

Acknowledgments.

This research was partially supported by National Natural Science Foundation of China under grants No. 62176231, No. 62106218, No. 82202984, No. 92259202 and No. 62132017, Zhejiang Key R&D Program of China under grant No. 2023C03053. This work is Supported by Shanghai Artificial Intelligence Laboratory. This work was supported by the JC STEM Lab of AI for Science and Engineering, funded by The Hong Kong Jockey Club Charities Trust.

References

1. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
2. Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., Fan, Y., Ge, W., Han, Y., Huang, F., et al.: Qwen technical report. arXiv preprint arXiv:2309.16609 (2023)
3. Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., Zhou, J.: Qwen-vl: A frontier large vision-language model with versatile abilities. ArXiv **abs/2308.12966** (2023)
4. Bansal, A., Sikka, K., Sharma, G., Chellappa, R., Divakaran, A.: Zero-shot object detection. In: Proceedings of the European conference on computer vision (ECCV). pp. 384–400 (2018)
5. Besta, M., Blach, N., Kubicek, A., Gerstenberger, R., Gianinazzi, L., Gajda, J., Lehmann, T., Podstawski, M., Niewiadomski, H., Nyczyk, P., et al.: Graph of thoughts: Solving elaborate problems with large language models. arXiv preprint arXiv:2308.09687 (2023)
6. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33**, 1877–1901 (2020)
7. Burić, M., Pobar, M., Ivašić-Kos, M.: Object detection in sports videos. In: 2018 41st International Convention on Information and Communication Technology, Electronics and Microelectronics (MIPRO). pp. 1034–1039. IEEE (2018)
8. Cao, W., Liu, Q., He, Z.: Review of pavement defect detection methods. *IEEE Access* **8**, 14531–14544 (2020). <https://doi.org/10.1109/ACCESS.2020.2966881>
9. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: European conference on computer vision. pp. 213–229. Springer (2020)
10. Chen, B., Xu, Z., Kirmani, S., Ichter, B., Driess, D., Florence, P., Sadigh, D., Guibas, L., Xia, F.: Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. arXiv preprint arXiv:2401.12168 (2024)
11. Chen, J., Li, D.Z.X.S.X., Zhang, Z.L.P., Xiong, R.K.V.C.Y., Elhoseiny, M.: Minigpt-v2: Large language model as a unified interface for vision-language multi-task learning. arXiv preprint arXiv:2310.09478 (2023)
12. Chen, K., Zhang, Z., Zeng, W., Zhang, R., Zhu, F., Zhao, R.: Shikra: Unleashing multimodal llm’s referential dialogue magic. arXiv preprint arXiv:2306.15195 (2023)
13. Chen, P., Sheng, K., Zhang, M., Lin, M., Shen, Y., Lin, S., Ren, B., Li, K.: Open vocabulary object detection with proposal mining and prediction equalization. arXiv preprint arXiv:2206.11134 (2022)

14. Chen, S., Sun, P., Song, Y., Luo, P.: Diffusiondet: Diffusion model for object detection. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 19830–19843 (2023)
15. Chen, X., Lin, M., Schärli, N., Zhou, D.: Teaching large language models to self-debug. *arXiv preprint arXiv:2304.05128* (2023)
16. Chou, Y.C., Li, B., Fan, D.P., Yuille, A., Zhou, Z.: Acquiring weak annotations for tumor localization in temporal and volumetric data. *Machine Intelligence Research* **21**(2), 318–330 (2024)
17. Czimmermann, T., Ciuti, G., Milazzo, M., Chiurazzi, M., Roccella, S., Oddo, C.M., Dario, P.: Visual-based defect detection and classification approaches for industrial applications—a survey. *Sensors* **20**(5), 1459 (2020)
18. Dai, G., Shu, X., Wu, W.: Gpt4ego: Unleashing the potential of pre-trained models for zero-shot egocentric action recognition. *arXiv preprint arXiv:2401.10039* (2024)
19. Dou, Z.Y., Kamath, A., Gan, Z., Zhang, P., Wang, J., Li, L., Liu, Z., Liu, C., LeCun, Y., Peng, N., et al.: Coarse-to-fine vision-language pre-training with fusion in the backbone. *Advances in neural information processing systems* **35**, 32942–32956 (2022)
20. Du, Y., Li, S., Torralba, A., Tenenbaum, J.B., Mordatch, I.: Improving factuality and reasoning in language models through multiagent debate. *arXiv preprint arXiv:2305.14325* (2023)
21. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. *arXiv preprint arXiv:2104.13921* (2021)
22. Han, Z., Zhu, F., Lao, Q., Jiang, H.: Zero-shot referring expression comprehension via structural similarity between images and captions. *arXiv preprint arXiv:2311.17048* (2023)
23. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: *Proceedings of the IEEE international conference on computer vision*. pp. 2961–2969 (2017)
24. Kim, D., Angelova, A., Kuo, W.: Contrastive feature masking open-vocabulary vision transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 15602–15612 (2023)
25. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. *arXiv preprint arXiv:2304.02643* (2023)
26. Lei, B., Liao, C., Ding, C., et al.: Boosting logical reasoning in large language models through a new framework: The graph of thought. *arXiv preprint arXiv:2308.08614* (2023)
27. Lei, X., Yang, Z., Chen, X., Li, P., Liu, Y.: Scaffolding coordinates to promote vision-language coordination in large multi-modal models. *arXiv preprint arXiv:2402.12058* (2024)
28. Lin, T.Y., Maire, M., Belongie, S., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft coco: Common objects in context. In: *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*. pp. 740–755. Springer (2014)
29. Lin, Z., Liu, C., Zhang, R., Gao, P., Qiu, L., Xiao, H., Qiu, H., Lin, C., Shao, W., Chen, K., et al.: Sphinx: The joint mixing of weights, tasks, and visual embeddings for multi-modal large language models. *arXiv preprint arXiv:2311.07575* (2023)
30. Liu, J., Xie, G., Wang, J., Li, S., Wang, C., Zheng, F., Jin, Y.: Deep industrial image anomaly detection: A survey. *Machine Intelligence Research* **21**(1), 104–135 (2024)

31. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Yue Li, C., Yang, J., Su, H., Zhu, J.J., Zhang, L.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. ArXiv **abs/2303.05499** (2023)
32. Liu, Z., Wang, H., Weng, L., Yang, Y.: Ship rotated bounding box space for ship extraction from high-resolution optical satellite images with complex backgrounds. *IEEE geoscience and remote sensing letters* **13**(8), 1074–1078 (2016)
33. Lu, P., Peng, B., Cheng, H., Galley, M., Chang, K.W., Wu, Y.N., Zhu, S.C., Gao, J.: Chameleon: Plug-and-play compositional reasoning with large language models. *Advances in Neural Information Processing Systems* **36** (2024)
34. Mao, J., Huang, J., Toshev, A., Camburu, O., Yuille, A.L., Murphy, K.: Generation and comprehension of unambiguous object descriptions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 11–20 (2016)
35. Minderer, M., Gritsenko, A., Stone, A., Neumann, M., Weissenborn, D., Dosovitskiy, A., Mahendran, A., Arnab, A., Dehghani, M., Shen, Z., et al.: Simple open-vocabulary object detection. In: *European Conference on Computer Vision*. pp. 728–755. Springer (2022)
36. Mitra, C., Huang, B., Darrell, T., Herzig, R.: Compositional chain-of-thought prompting for large multimodal models. arXiv preprint arXiv:2311.17076 (2023)
37. Peng, Z., Wang, W., Dong, L., Hao, Y., Huang, S., Ma, S., Wei, F.: Kosmos-2: Grounding multimodal large language models to the world. arXiv preprint arXiv:2306.14824 (2023)
38. Raffel, C., Shazeer, N., Roberts, A., Lee, K., Narang, S., Matena, M., Zhou, Y., Li, W., Liu, P.J.: Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research* **21**(1), 5485–5551 (2020)
39. Ren, S., He, K., Girshick, R., Sun, J.: Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* **28** (2015)
40. Schick, T., Dwivedi-Yu, J., Jiang, Z., Petroni, F., Lewis, P., Izacard, G., You, Q., Nalmpantis, C., Grave, E., Riedel, S.: Peer: A collaborative language model. arXiv preprint arXiv:2208.11663 (2022)
41. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: *2017 IEEE International Conference on Computer Vision (ICCV)*. pp. 618–626 (2017). <https://doi.org/10.1109/ICCV.2017.74>
42. Shtedritski, A., Rupprecht, C., Vedaldi, A.: What does clip know about a red circle? visual prompt engineering for vlms. arXiv preprint arXiv:2304.06712 (2023)
43. Subramanian, S., Merrill, W., Darrell, T., Gardner, M., Singh, S., Rohrbach, A.: Reclip: A strong zero-shot baseline for referring expression comprehension. arXiv preprint arXiv:2204.05991 (2022)
44. Sun, P., Zhang, R., Jiang, Y., Kong, T., Xu, C., Zhan, W., Tomizuka, M., Li, L., Yuan, Z., Wang, C., et al.: Sparse r-cnn: End-to-end object detection with learnable proposals. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 14454–14463 (2021)
45. Team, G., Anil, R., Borgeaud, S., Wu, Y., Alayrac, J.B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., et al.: Gemini: a family of highly capable multimodal models. arXiv preprint arXiv:2312.11805 (2023)
46. Team, I.: Internlm: A multilingual language model with progressively enhanced capabilities (2023)
47. Thomas, G., Gade, R., Moeslund, T.B., Carr, P., Hilton, A.: Computer vision for sports: Current applications and research topics. *Computer Vision and Image Understanding* **159**, 3–18 (2017)

48. Tong, S., Liu, Z., Zhai, Y., Ma, Y., LeCun, Y., Xie, S.: Eyes wide shut? exploring the visual shortcomings of multimodal llms. arXiv preprint arXiv:2401.06209 (2024)
49. Vandeghen, R., Cioppa, A., Van Droogenbroeck, M.: Semi-supervised training to improve player and ball detection in soccer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3481–3490 (2022)
50. Wang, L., Liu, Y., Du, P., Ding, Z., Liao, Y., Qi, Q., Chen, B., Liu, S.: Object-aware distillation pyramid for open-vocabulary object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11186–11196 (2023)
51. Wang, P., Yang, A., Men, R., Lin, J., Bai, S., Li, Z., Ma, J., Zhou, C., Zhou, J., Yang, H.: Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In: International Conference on Machine Learning. pp. 23318–23340. PMLR (2022)
52. Wang, W., Chen, Z., Chen, X., Wu, J., Zhu, X., Zeng, G., Luo, P., Lu, T., Zhou, J., Qiao, Y., et al.: Visionllm: Large language model is also an open-ended decoder for vision-centric tasks. arXiv preprint arXiv:2305.11175 (2023)
53. Wang, Y., Chen, M., Tang, S., Zhu, F., Yang, H., Bai, L., Zhao, R., Yan, Y., Qi, D., Ouyang, W.: Unsupervised object detection pretraining with joint object priors generation and detector learning. *Advances in neural information processing systems* **35**, 12435–12448 (2022)
54. Wang, Y., Wu, Y., Tang, S., He, W., Guo, X., Zhu, F., Bai, L., Zhao, R., Wu, J., He, T., et al.: Hulk: A universal knowledge translator for human-centric tasks. arXiv preprint arXiv:2312.01697 **8** (2023)
55. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems* **35**, 24824–24837 (2022)
56. Wu, S., Zhang, W., Jin, S., Liu, W., Loy, C.C.: Aligning bag of regions for open-vocabulary object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15254–15264 (2023)
57. Wu, W., Yao, H., Zhang, M., Song, Y., Ouyang, W., Wang, J.: Gpt4vis: what can gpt-4 do for zero-shot visual recognition? arXiv preprint arXiv:2311.15732 (2023)
58. Wu, X., Zhu, F., Zhao, R., Li, H.: Cora: Adapting clip for open-vocabulary detection with region prompting and anchor pre-matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7031–7040 (2023)
59. Wu, Y., Zhang, Z., Xie, C., Zhu, F., Zhao, R.: Advancing referring expression segmentation beyond single image. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2628–2638 (2023)
60. Wu, Y., Zheng, B., Chen, J., Chen, D.Z., Wu, J.: Self-learning and one-shot learning based single-slice annotation for 3d medical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 244–254. Springer (2022)
61. Xie, C., Zhang, Z., Wu, Y., Zhu, F., Zhao, R., Liang, S.: Described object detection: Liberating object detection with flexible expressions. *Advances in Neural Information Processing Systems* **36** (2024)
62. Yan, B., Jiang, Y., Wu, J., Wang, D., Luo, P., Yuan, Z., Lu, H.: Universal instance perception as object discovery and retrieval. 2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 15325–15336 (2023)
63. Yang, J., Zhang, H., Li, F., Zou, X., Li, C., Gao, J.: Set-of-mark prompting unleashes extraordinary visual grounding in gpt-4v. arXiv preprint arXiv:2310.11441 (2023)

64. Yang, L., Wang, Y., Li, X., Wang, X., Yang, J.: Fine-grained visual prompting. *Advances in Neural Information Processing Systems* **36** (2024)
65. Yang, Z., Li, L., Lin, K., Wang, J., Lin, C.C., Liu, Z., Wang, L.: The dawn of lmms: Preliminary explorations with gpt-4v (ision). *arXiv preprint arXiv:2309.17421* **9**(1), 1 (2023)
66. Yao, F., Tian, C., Liu, J., Zhang, Z., Liu, Q., Jin, L., Li, S., Li, X., Sun, X.: Thinking like an expert: Multimodal hypergraph-of-thought (hot) reasoning to boost foundation modals. *arXiv preprint arXiv:2308.06207* (2023)
67. Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T., Cao, Y., Narasimhan, K.: Tree of thoughts: Deliberate problem solving with large language models. *Advances in Neural Information Processing Systems* **36** (2024)
68. Yao, Y., Li, Z., Zhao, H.: Beyond chain-of-thought, effective graph-of-thought reasoning in large language models. *arXiv preprint arXiv:2305.16582* (2023)
69. Yao, Y., Zhang, A., Zhang, Z., Liu, Z., Chua, T.S., Sun, M.: Cpt: Colorful prompt tuning for pre-trained vision-language models. *arXiv preprint arXiv:2109.11797* (2021)
70. Yin, Z., Wang, J., Cao, J., Shi, Z., Liu, D., Li, M., Huang, X., Wang, Z., Sheng, L., Bai, L., et al.: Lamm: Language-assisted multi-modal instruction-tuning dataset, framework, and benchmark. *Advances in Neural Information Processing Systems* **36** (2024)
71. You, H., Zhang, H., Gan, Z., Du, X., Zhang, B., Wang, Z., Cao, L., Chang, S.F., Yang, Y.: Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704* (2023)
72. Yu, L., Lin, Z., Shen, X., Yang, J., Lu, X., Bansal, M., Berg, T.L.: Mattnet: Modular attention network for referring expression comprehension. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1307–1315 (2018)
73. Yu, L., Poirson, P., Yang, S., Berg, A.C., Berg, T.L.: Modeling context in referring expressions. In: *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part II* 14. pp. 69–85. Springer (2016)
74. Zang, Y., Li, W., Han, J., Zhou, K., Loy, C.C.: Contextual object detection with multimodal large language models. *arXiv preprint arXiv:2305.18279* (2023)
75. Zang, Y., Li, W., Zhou, K., Huang, C., Loy, C.C.: Open-vocabulary detr with conditional matching. In: *European Conference on Computer Vision*. pp. 106–122. Springer (2022)
76. Zareian, A., Rosa, K.D., Hu, D.H., Chang, S.F.: Open-vocabulary object detection using captions. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 14393–14402 (2021)
77. Zeng, N., Wu, P., Wang, Z., Li, H., Liu, W., Liu, X.: A small-sized object detection oriented multi-scale feature fusion approach with application to defect detection. *IEEE Transactions on Instrumentation and Measurement* **71**, 1–14 (2022)
78. Zhang, Z., Zhang, A., Li, M., Smola, A.: Automatic chain of thought prompting in large language models. *arXiv preprint arXiv:2210.03493* (2022)
79. Zhang, Z., Zhang, A., Li, M., Zhao, H., Karypis, G., Smola, A.: Multimodal chain-of-thought reasoning in language models. *arXiv preprint arXiv:2302.00923* (2023)
80. Zhao, Y., Lin, Z., Zhou, D., Huang, Z., Feng, J., Kang, B.: Bubogpt: Enabling visual grounding in multi-modal llms. *arXiv preprint arXiv:2307.08581* (2023)
81. Zheng, G., Yang, B., Tang, J., Zhou, H.Y., Yang, S.: Ddcot: Duty-distinct chain-of-thought prompting for multimodal reasoning in language models. *Advances in Neural Information Processing Systems* **36** (2024)

82. Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., et al.: Regionclip: Region-based language-image pretraining. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16793–16803 (2022)
83. Zhuge, M., Liu, H., Faccio, F., Ashley, D.R., Csordás, R., Gopalakrishnan, A., Hamdi, A., Hammoud, H.A.A.K., Herrmann, V., Irie, K., et al.: Mindstorms in natural language-based societies of mind. arXiv preprint arXiv:2305.17066 (2023)
84. Zong, Z., Song, G., Liu, Y.: Detrs with collaborative hybrid assignments training. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6748–6758 (2023)