
Prism: A Framework for Decoupling and Assessing the Capabilities of VLMs

Yuxuan Qiao^{2*}, Haodong Duan^{1†}, Xinyu Fang⁴, Junming Yang⁵, Lin Chen⁶,
Songyang Zhang¹, Jiaqi Wang¹, Dahua Lin^{1,3,7}, Kai Chen^{1†}

¹Shanghai AI Laboratory ²Nanjing University ³The Chinese University of Hong Kong

⁴Tongji University ⁵Nanjing University of Posts and Telecommunications

⁶University of Science and Technology of China ⁷CPII under InnoHK

yuxuanqiao@smail.nju.edu.cn

duanhaodong@pjlab.org.cn

Abstract

Vision Language Models (VLMs) demonstrate remarkable proficiency in addressing a wide array of visual questions, which requires strong perception and reasoning faculties. Assessing these two competencies independently is crucial for model refinement, despite the inherent difficulty due to the intertwined nature of seeing and reasoning in existing VLMs. To tackle this issue, we present **Prism**, an innovative framework designed to disentangle the perception and reasoning processes involved in visual question solving. Prism comprises two distinct stages: a **perception stage** that utilizes a VLM to extract and articulate visual information in textual form, and a **reasoning stage** that formulates responses based on the extracted visual information using a Large Language Model (LLM). This modular design enables the systematic comparison and assessment of both proprietary and open-source VLM for their perception and reasoning strengths. Our analytical framework provides several valuable insights, underscoring Prism’s potential as a cost-effective solution for vision-language tasks. By combining a streamlined VLM focused on perception with a powerful LLM tailored for reasoning, Prism achieves superior results in general vision-language tasks while substantially cutting down on training and operational expenses. Quantitative evaluations show that Prism, when configured with a vanilla 2B LLaVA and freely accessible GPT-3.5, delivers performance on par with VLMs 10× larger on the rigorous multimodal benchmark MMStar.

1 Introduction

With the rapid development of Large Language Models (LLMs) [45, 59, 60, 56, 5, 19, 4], Vision Language Models (VLMs) [46, 55, 6, 36, 16, 18] have also experienced significant advancements. As an end-to-end approach, VLMs trained on large-scale multimodal data [36, 11, 51, 8] exhibit superior performance on a variety of tasks. These tasks range from basic ones such as object localization [31] and optical character recognition [44, 38] to more complex challenges like document or diagram comprehension [26, 43, 42] and solving geometric problems [40]. For VLMs to solve a general visual question, two essential capabilities are required: 1) **Perception**: extracting necessary information from the image; 2) **Reasoning**: generating answers based on the extracted information and contextual understanding. Limitations in either capability can impede the overall performance of a VLM.

A systematic evaluation of the perception and reasoning capabilities is crucial to provide valuable insights for future model optimization. However, seeing and reasoning are mostly **entangled** in

* The work was done during an internship at Shanghai AI Laboratory.

† Corresponding Author.

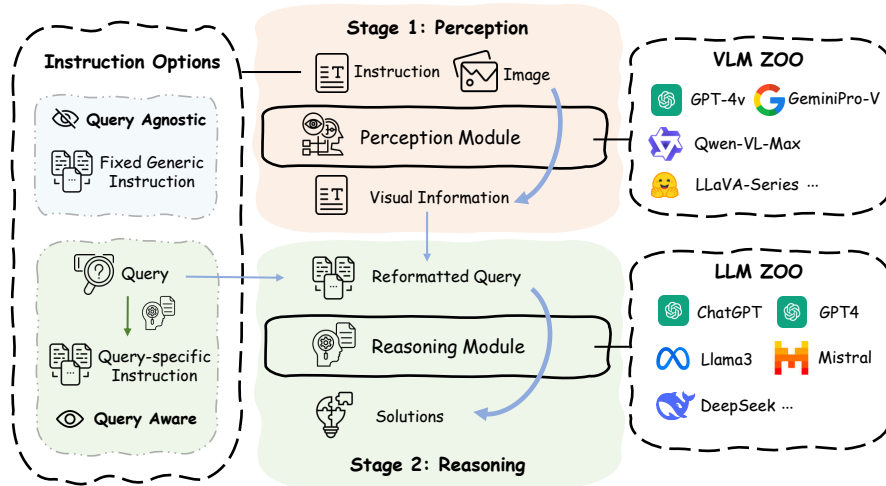


Figure 1: **Prism Framework Architecture.** Prism framework takes image-query pairs as input. An instruction (can be query-agnostic or query-aware) and the image are first fed into the VLM to extract visual information. Then, an LLM is used to generate the answer based on the reformatted query which combines the original question and visual information in textual form.

existing VLMs. Proprietary VLMs, such as GPT-4v are opaque systems that only provide the final answer to visual questions. Meanwhile, open-source VLMs [36] commonly utilize vision encoders to extract visual embeddings, which are often difficult to interpret, from the image and employ adapted LLMs to generate answers based on these visual and linguistic embeddings. In light of this challenge, we introduce **Prism**, a framework designed to disentangle the perception and reasoning processes of any given VLM. It can serve as a proxy for assessing the real perception capabilities of VLMs.

Prism decomposes the process of solving general visual questions into two distinct stages: 1) a **perception stage** that concentrates on extracting visual information from the image using VLMs and articulating this information in textual form; and 2) a **reasoning stage** that utilizes LLMs to answer the question based on the extracted visual information. Prism facilitates the analysis of VLMs’ capabilities through two approaches. To assess the true perception capabilities, one can use a constant LLM in the reasoning stage while testing various VLMs in the perception stage. Conversely, by keeping the VLM fixed and varying the LLM in the reasoning stage, one can determine whether a VLM’s performance is limited by its reasoning capabilities. Through this analysis, we have uncovered several key insights: 1) Proprietary VLMs, such as GPT-4o and GPT-4v, take the lead in the perception capabilities competition; 2) For open-source VLMs, perception capabilities remain relatively consistent regardless of the language model’s size; and 3) The overall performance of open-source VLMs, particularly those with smaller-scale language models like 7B variants, is often constrained by the limited reasoning capabilities.

Beyond its role as an evaluation framework, Prism also excels as an efficient general Vision-Language Model (VLM). Building upon findings 2 and 3, we posit that integrating a small-scale VLM as a visual captioner with a powerful LLM as a reasoning engine offers a promising and efficient strategy for general vision-language tasks. By concentrating on visual information extraction, a lightweight VLM can achieve decent performance on par with much larger VLMs. When paired with a powerful yet economical LLM, thanks to advancements in deployment techniques, one can achieve a robust solution for visual-language information processing, requiring significantly fewer hardware resources for training and deployment. In our experiments, we trained an approximately 2B-parameter vanilla LLaVA to extract visual information and observed that it exhibits perception performance comparable to LLaVA-NeXT [35] that is equipped with a 34B powerful language model. Quantitative evaluations indicate that Prism, when instantiated with a streamlined visual captioner and the freely available ChatGPT-3.5, outperforms many open-source VLMs on multiple multimodal benchmarks including the stringent multimodal understanding benchmark MMStar [10]. Notably, the advantage is particularly pronounced on reasoning-related visual questions.

In summary, the contributions of this work are as follows:

1. We introduce **Prism**, a highly adaptable framework designed to explicitly disentangle the perception and reasoning processes. Prism enables the breakdown analysis of VLM capabilities and serves as a solution for vision-language tasks by integrating any given VLM and LLM.
2. Utilizing Prism, we conduct a **decoupled analysis** of the perception and reasoning capabilities of existing VLMs. Several intriguing findings emerge from the ablation study.
3. Drawing inspiration from these findings, we integrate a lightweight VLM focused on perception with a powerful LLM dedicated to reasoning within the Prism framework. Quantitative results demonstrate that this combination exhibits **outstanding performance and efficiency** across a range of vision-language tasks.

2 Methodology

2.1 Prism Architecture

Prism is characterized by a modular design that decomposes the process of solving visual questions into two stages, consisting of a perception module and a reasoning module, both of which can be flexibly replaced, as depicted in Fig. 1. The perception module, typically a VLM, initially follows the instruction to extract visual information from images and articulates this information in textual form. The instruction can be generic or query specific, *ie.*, written by the reasoning module given the question (text-only) as contextual information. Meanwhile, the reasoning module, usually an LLM, performs text-based reasoning on the textual information to generate answers to the questions. To assess the perception capabilities of various VLMs, we carefully select an appropriate benchmark corresponding to the principle we illustrate in Sec. 2.2. In Sec. 2.3, we describe how we utilize the Prism framework to assess the perception and reasoning capabilities of VLMs, respectively. To substantiate our belief in the potential of combining a small-scale VLM with a powerful LLM, we train an approximately 2B-parameter VLM based on the LLaVA architecture to serve as a visual captioner and integrate it into Prism as the perceptual module, as detailed in Sec. 2.4.

2.2 To Analyze with a Suitable Benchmark

Numerous multimodal benchmarks [37, 20, 28, 40, 66, 10] exist, each evaluating the capabilities of VLMs from various perspectives. To maximize the utility of the Prism framework, careful selection of the benchmark is imperative for the decoupling analysis. In summary, we adhere to the following principles for benchmark selection: 1) **Vision Indispensability**: The benchmark must require visual information for question solving, and questions that can be answered without utilizing visual information are excluded; 2) **Minimal Data Leakage**: The visual questions should not be part of the model’s training data; and 3) **Complexity**: The solving process of the visual question should involve both perception and reasoning components. Considering these three principals, we select MMStar [10] as the primary benchmark for the decoupling analysis and ablation. MMStar ensures vision indispensability and makes a concerted effort to minimize data leakage, while many other benchmarks are plagued by the two issues.

2.3 Prism as an Analytical Framework

Prism functions as an analytical framework for evaluating the perception and reasoning capabilities of a given VLM. In this section, we elaborate on the methodology for evaluating these capabilities.

To assess the perception capabilities of various VLMs, we employ ChatGPT (GPT-3.5-turbo-0125) as the reasoning module and standardize instructions for description within Prism. With the reasoning module and instructions fixed, the final accuracy of the visual questions is solely determined by the quality of the visual information extracted. Under the controlled setting, we consider the VQA accuracy as a proxy to measure the perception capability of VLMs.

Within the Prism framework, the instructions for visual information extraction are crucial as they are designed to elicit the fundamental perceptual capabilities of VLMs. We have adopted two types of instructions to assess the perceptual abilities of models:

1. **Generic Instruction**: a standardized, universal instruction aimed at extracting and describing the basic elements present in an image;
2. **Query-Specific Instruction**: a combination of the generic instruction and an incremental instruction that directs the VLM to provide a detailed account of the visual information relevant to the question. The incremental instruction is crafted by the reasoning module given the text-only question.

Employed with the generic instruction, Prism offers a straightforward decoupling pipeline. Meanwhile, using query-specific instructions is a more realistic setting and can better realize the full potential of the Prism framework. A Prism decoupled GPT-4o (GPT-4o for both perception and reasoning) achieves almost the same quantitative performance compared to the end-to-end GPT-4o. Analyzing the results, we find that with the reasoning module and instruction controlled, VLMs with language models of different sizes display a much narrowed gap in perception performance compared to the results under the end-to-end setting.

Besides evaluating the perception performance, Prism can also roughly measure the reasoning capabilities of VLMs. When used to solve general visual questions in an end-to-end manner, the VLM implicitly performs reasoning with its language encoder. Alternatively, Prism provides an explicit pipeline, in which an off-the-shelf LLM is separately used to predict the answer based on the visual information. For small-scale VLMs (7B-parameter, *etc.*), we find that Prism equipped with the VLM and ChatGPT-3.5 outperforms the end-to-end VLM in quantitative performance, especially for reasoning related VQA. The results reveal that for small-scale VLMs, the overall performance can be heavily constrained by the parameter size of the language model.

2.4 Prism as a Vision-Language Task Solver

Beyond serving as an evaluation framework, Prism can also function as an **efficient vision-language task solver**. The perception module in Prism can incorporate one or multiple VLMs to extract high-quality visual information. Concurrently, the reasoning module can be instantiated with a powerful LLM to harness its advanced reasoning capabilities.

In Sec. 2.3, we observed that VLMs paired with language models of varying sizes exhibit similar performance in perception capability. This suggests that it is promising to employ small-scale VLMs to generate informative visual descriptions, serving as an efficient perception module. To validate this concept, we conduct extensive experiments to train visual captioners, utilizing the widely adopted LLaVA architecture and only open-source datasets. Below, we elaborate on the specific settings considered for this ablation study in detail:

Instruction Tuning Data. We use ALLaVA-Caption-4V and Evol-Instruct-GPT4-Turbo-143K in ALLaVA [9] as our instruction tuning data. The former comprises 715K pairs of images and detailed visual captions produced by GPT-4v, whereas the latter contains 143K text instruction tuning data, generated by GPT-4-Turbo. Utilizing the descriptive data for instruction tuning better triggers the VLM’s ability to extract and articulate visual information more effectively, compared to instruction tuning data in QA formats [36, 32, 52, 42].

Model Architecture. To investigate the impact of vision encoder in LLaVA, we experimented with multiple encoders, including Clip ViT-L/14 [50], SigLip-S0400M [67], and InternViT-6B [13]. For the language encoder, we tested two lightweight variants of InternLM2 [7]: InternLM2-7B and InternLM2-1.8B. The inference of all combinations (excluding InternViT-6B) can be efficiently executed on consumer-level GPUs such as RTX 4090, *etc.* During instruction tuning, we maintain the vision encoder fixed and apply QLoRA [17] to the language encoder.

Regarding the vision backbone, SigLip exhibits relatively superior performance. Additionally, we observed that a larger language model only results in minor differences in perception performance. Our findings indicate that, within the Prism framework, a 2B vision captioner can achieve strong perception performance on par with LLaVA-NeXT [35] equipped with a 34B language backbone.

3 Evaluation Results

3.1 Implementation Details

Evaluation Details. We use Prism to evaluate the capabilities of various VLMs, which can be categorized into two major groups: (a) **Proprietary VLMs**, including GPT-4o (20240513) [46], GPT-4v (20231106) [46], GeminiPro-V [55], and Qwen-VL-Max [6]; (b) **Open-Source VLMs**, including LLaVA-v1.5 [34], InternLM-XComposer2 [18], mPLUG-Owl2 [65], LLaVA-NeXT [35], InternVL-Chat-v1.5 [12], DeepSeek-VL [39], MiniCPM-V-2 [47]. When integrating these VLMs as perception modules within Prism, we employ greedy decoding and limit the maximum number of output tokens to 512. The evaluation encompasses both generic and query-specific instructions.

Model	Generic Instruction							Query-Specific Instruction						
	CP	FP	IR	LR	Math	ST	Overall	CP	FP	IR	LR	Math	ST	Overall
<i>Proprietary VLMs</i>														
GPT-4o [46]	<u>64.0</u>	41.6	<u>54.4</u>	51.6	47.6	31.2	48.4	67.2	51.6	63.2	56.4	52.0	36.8	54.5
GPT-4v [46]	62.4	33.2	51.2	39.2	44.4	32.8	43.9	67.2	43.6	56.8	46.8	42.0	30.8	47.9
GeminiPro-V [55]	57.2	36.0	51.6	43.2	45.6	23.2	42.8	61.2	35.2	53.2	47.6	48.4	19.2	44.1
Qwen-VL-Max [5]	62.4	28.8	49.2	45.2	37.2	21.2	40.7	62.4	36.0	54.8	46.4	38.0	24.0	43.6
<i>Open-Source VLMs</i>														
InternVL-Chat-v1.5 [12]	59.2	33.2	51.2	<u>42.8</u>	40.4	<u>30.4</u>	<u>42.9</u>	<u>65.6</u>	<u>42.4</u>	56.4	<u>50.8</u>	<u>46.8</u>	<u>28.0</u>	<u>48.3</u>
InternLM-XComposer2 [18]	56.0	32.0	54.8	33.6	37.2	20.8	39.1	60.8	39.2	<u>58.4</u>	47.6	42.4	19.6	44.7
LLaVA-NeXT (Yi-34B) [35]	64.4	<u>37.2</u>	50.0	34.0	39.2	24.8	41.6	60.0	41.6	52.4	44.0	40.8	25.2	44.0
DeepSeek-VL-7B [39]	57.2	31.6	49.6	42.0	<u>43.6</u>	24.0	41.3	61.2	34.8	54.4	42.8	45.2	25.2	43.9
LLaVA-NeXT (Mistral-7B) [35]	60.8	30.8	49.6	37.6	36.8	20.4	39.3	62.8	36.4	53.6	42.0	35.2	26.4	42.7
MiniCPM-V-2 [47]	55.6	28.4	49.2	39.6	38.4	19.2	38.4	61.2	30.4	52.0	44.4	40.8	22.8	41.9
LLaVA-NeXT (Vicuna-13B) [35]	60.8	38.0	48.8	35.2	<u>43.6</u>	20.8	41.2	63.2	38.8	49.2	38.8	37.6	23.2	41.8
LLaVA-NeXT (Vicuna-7B) [35]	63.2	28.8	46.4	40.8	34.0	22.8	39.3	56.8	37.6	47.2	42.4	36.4	21.6	40.3
mPLUG-Owl2 [65]	47.2	24.0	45.2	33.6	28.8	18.8	32.9	53.2	34.4	45.2	38.8	38.8	23.6	39.0
LLaVA-v1.5-13B [34]	45.2	28.4	45.2	35.6	30.0	17.2	33.6	54.8	28.8	48.8	37.6	32.0	20.0	37.0
LLaVA-v1.5-7B [34]	48.8	29.6	48.0	34.4	27.6	17.2	34.3	50.8	31.6	50.4	38.8	32.0	25.2	38.1

Table 1: **Detailed Perception Performance on MMStar under Prism Evaluation Framework.** Reasoning module: ChatGPT. Abbreviations: Coarse Perception (CP), Fine-grained Perception (FP); Instance Reasoning (IR); Logical Reasoning (LR); Science&Technology (ST). Overall best scores are marked as **bold**, and intra-category best scores are marked as underline.

Unless otherwise specified, GPT-3.5-Turbo-0125 is adopted as the reasoning module. All evaluations are conducted using VLMEvalKit [14]. Further details on the evaluation are provided in Appx. B.1.

Post-Processing. Within the Prism framework, the LLM (particularly proprietary APIs) in the reasoning module often declines to answer questions due to insufficient clues in the extracted visual information. When Prism is employed as an evaluation framework, we classify this as a failure of perception for the visual question and refrain from any post-processing. For fair comparisons, whether Prism is contrasted with other VLMs as a vision-language task solver or different LLMs are evaluated as reasoning modules with varying rejection rates, a random choice is utilized as a fallback when option matching fails.

3.2 Main Results

We present the primary evaluation results of various VLMs’ perception performance in Tab. 1. Results are categorized under two instruction types: *Generic Instruction* and *Query-Specific Instruction*.

Generic Instruction Results. When using generic instructions, GPT-4o exhibits exceptional performance across a variety of tasks. InternVL-Chat-v1.5 achieves the highest overall perception performance among open-source VLMs, nearly on par with the proprietary GPT-4v. LLaVA-NeXT (Yi-34B) demonstrates strong coarse and fine-grained perceptual abilities but lags behind GPT-4v in recognizing abstract elements in LR and ST visual questions. Smaller open-source VLMs, such as mPLUG-Owl2 and LLaVA-v1.5-7B, struggle with both coarse and fine-grained perception and encounter difficulties in recognizing essential elements in reasoning-related VQA.

Query-Specific Instruction Results. GPT-4o demonstrates significantly superior performance compared to all other models in the extraction and expression of visual information across all dimensions. GPT-4v is equally adept at handling coarse perception contents. Among open-source VLMs, InternVL-Chat-v1.5 excels in perception tasks across all dimensions excluding LR. Its performance not only surpasses that of other open-source VLMs but is also slightly ahead of GPT-4v. Occasionally, a decline in performance can be observed for specific VLM-task pairs when using query-specific instructions instead of generic ones, which can stem from difficulty in understanding query-specific instructions for specific domains.

3.3 Detailed Analysis

Are Proprietary Models Getting Ahead of the Game? When employing both generic and query-specific instructions, proprietary VLMs, particularly GPT-4o, significantly surpass other models in perceptual capabilities and can adeptly manage a wide range of tasks, as demonstrated in Tab. 1. Certain open-source models, such as InternVL-Chat-v1.5 and LLaVA-NeXT (Yi-34B), have achieved

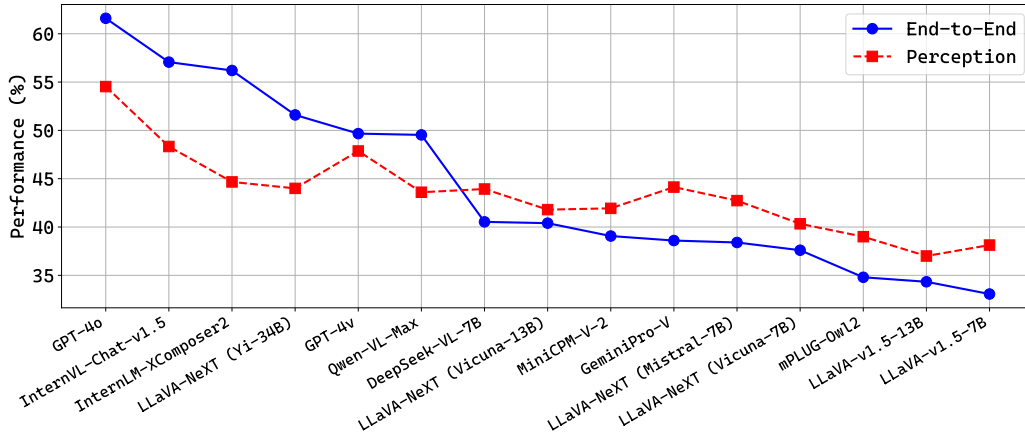


Figure 2: **Comparing End-to-End Performance and Perception Capability on MMStar.** We display model accuracies in end-to-end VQA and the Prism perception test with query-specific instructions. Most small-scale (7B, 13B, etc.) VLMs achieve better performance within Prism.

notable performance, approaching the capabilities of proprietary VLMs like GPT-4v and GeminiPro-V. Other open-source models, due to their limited perceptual abilities, generally perform slightly worse in Math and ST assessments. Notably, MiniCPM-V-2, a lightweight VLM with ~ 3 B parameters, displays better perceptual performance compared to some 7B VLMs.

The Gap between Perception Ability and End-to-End Performance. In addition to solving visual questions in an end-to-end manner, Prism provides an alternative pipeline where the VLM is solely utilized for perception. The distinction between these two methods lies in the reasoning process: the former conducts reasoning internally within VLMs, whereas the latter performs reasoning based on VLM extracted information using an external LLM (ChatGPT). The comparison between these two approaches on MMStar is depicted in Fig. 2. For state-of-the-art large-scale VLMs such as GPT-4o and InternVL-Chat-v1.5, which are expected to possess excellent reasoning capabilities, employing an external ChatGPT for reasoning may diminish overall performance. Conversely, for most small-scale VLMs, using ChatGPT for reasoning significantly improves their performance, particularly in reasoning-related VQA, as shown in Fig. 5. This phenomenon indicates that the overall performance of small-scale VLMs can be heavily constrained by the size of the language model. To investigate whether the reasoning ability of ChatGPT constrains state-of-the-art VLMs, we implemented a Prism pipeline that decouples GPT-4o by using it as both the perception and reasoning module. The result reveals that this Prism pipeline, with post-processing, achieves an overall accuracy of 61%, nearly identical to the end-to-end GPT-4o performance of 61.6%.

How does Language Model Size Affect Perception Ability? During evaluation, we observe that the LLaVA-v1.5 series shows no significant improvement when using a larger language model (Vicuna-13B instead of Vicuna-7B, etc.). This suggests that perception performance may be independent of the language model size when using a relatively low-resolution vision backbone. However, the situation appears to differ with LLaVA-NeXT. Quantitative results for the LLaVA-NeXT series tell that scaling up the language model slightly enhances model perception, particularly when using query-specify instructions. Through a detailed qualitative analysis, we identified the primary factors contributing to the superior performance of larger LLaVA-NeXT models over smaller ones as follows: (1) **More Elaborate Expression:** Models equipped with a larger language encoder exhibit enhanced ability to articulate visual information. More detailed and organized narratives make it easier for the reasoning module to answer the question; (2) **More Adaptive to Instruction:** Larger language backbones entitle the model with a better understanding of instructions, yielding more suitable textual visual information for reasoning, particularly in response to query-specific instructions. In Fig. 3, we provide some qualitative results about the two typical modes.

3.4 Error Analysis

Error Mode. We conduct a thorough analysis and categorize errors in perception and reasoning into the following major modes:

1. **Major Error Modes in Perception:** (1) **Factual Errors.** VLMs may describe images with inaccuracies, such as stating that prominent elements are "not visible"; (2) **Incomplete Details.**



Figure 3: **The Effect of Language Model Size on Perception Ability.** We compare visual information extracted from different LLaVA-NeXT models. **Left:** LLaVA-NeXT (Yi-34B) tells the spatial arrangement in a more detailed way; **Right:** LLaVA-NeXT (Vicuna-7B) dismisses the query on the man’s hair while LLaVA-NeXT (Yi-34B) tells all contents elaborately following the instruction.

VLM	Agreement Rate (VLM)	Error Rate (VLM)
GPT-4o	68.6	29.9
GPT-4v	61.5	38.0
LLaVA-NeXT (Yi-34B)	60.7	43.4
LLaVA-v1.5-7B	50.6	56.5

Table 2: **Agreement Rates and Error Rates of Different VLMs.**

LLM	Alignment Rate (LLM)
GPT-3.5-Turbo-0125	87.1
DeepSeek-v2-Chat	91.3
Llama-3-70B-Instruct	92.8
GPT-4-Turbo-0125	90.7

Table 3: **Alignment Rates of Different LLMs.**

Even in the absence of factual errors, VLMs may lack detailed content, resulting in insufficient information for reasoning; (3) **Instruction Following:** VLMs sometimes fail to follow instructions when providing corresponding descriptions.

2. **Major Error Modes in Reasoning:** (1) **Logical Errors.** LLMs may produce incorrect conclusions or reasoning processes due to limited reasoning abilities; (2) **Lack of Knowledge.** The absence of relevant domain knowledge prevents LLM from solving corresponding problems, especially in specific fields; (3) **Misunderstanding of the Query.** In rare cases, the query-specific part generated by LLMs deviates from the original question, misleading the perception of VLMs.

Hallucination. To reveal the source of hallucinations, we employ various LLMs to reason about descriptions generated by the given VLM and analyze the patterns in their results. We conduct experiments on various VLMs with a set of powerful language models, including GPT-3.5-Turbo-0125, GPT-4-Turbo-0125, Llama-3-70B-Instruct, and DeepSeek-v2-Chat. For each question, we gather results from four LLMs and consider cases where the same choice appears three times or more as "**agreement**". In agreement cases, we define the corresponding choice as "**voted choice**" and focus on the following cases:

Case 1: With descriptive texts generated by a specific VLM, the voted choice is wrong.

Case 2: With descriptive texts generated by a specific VLM, the voted choice is correct.

Case 3: For a specific LLM, it makes the same choice as the voted choice by all LLMs, while the voted choice is also the correct one.

For each VLM, we analyze the rate of case 1, which indicates the potential hallucinations of VLMs. For each LLM, we correspondingly calculate the rate of case 3 to case 2 to observe the alignment of its predictions with voted choices. The higher alignment rate means more robust reasoning since voted options are considered better. The notions are as follows.

$$\text{Agreement Rate (VLM)} = \frac{\text{number of agreement cases}}{\text{number of all cases}}$$

$$\text{Error Rate (VLM)} = \frac{\text{number of case 1}}{\text{number of agreement cases}}$$

$$\text{Alignment Rate (LLM)} = \frac{\text{number of case 3}}{\text{number of case 2}}$$

As presented in Tabs. 2 and 3, stronger VLMs show lower error rates, indicating that more capable VLMs experience fewer hallucination issues. By delving deep into the specific cases, we find that VLMs are prone to hallucinations in spatial awareness and fine-grained perception. All LLMs show relatively good alignment rates, demonstrating relatively robust reasoning performance. GPT-3.5-Turbo-0125 is more unstable by comparison. The cases where LLM’s prediction does not align with the voted choice may stem both from reasoning ability issues and hallucinations.

Instructions	GPT-4o	LLaVA-NeXT (Yi-34B)
Human 2	47.7	40.4
GPT Synthesize 1	47.1	41.3
GPT Synthesize 2	48.1	40.1
CoT	47.7	41.4
Decompose	47.3	41.4
Human 1	48.4	41.6

Table 4: Ablation on Different Generic Instructions.

Model	GPT-4o	GPT-4v	LLaVA-NeXT (Yi-34B)	LLaVA-v1.5-7B
GPT-3.5-Turbo-0125	54.7	48.5	44.1	38.5
GPT-4-Turbo-0125	56.9	48.7	48.4	38.3
Llama-3-70B-Instruct	59.3	51.3	47.7	38.5
DeepSeek-v2-Chat	57.8	49.7	45.8	35

Table 5: Ablation on Using Different LLMs as the Reasoning Module.

Different Generic Instructions
Human 1: Describe the fine-grained content of the image, including scenes, objects, relationships, instance location, and any text present.
Human 2: Describe the fine-grained content of the image, including scenes, objects, relationships, instance location, background and any text present. Please skip generating statements for non-existent contents and describe all you see.
GPT Synthesize 1: Given the image below, please provide a detailed description of what you see.
GPT Synthesize 2: Analyze the image below and describe the main elements and their relationship.
CoT: Describe the fine-grained content of the image, including scenes, objects, relationships, instance location, and any text present. Let’s think step by step.
Decompose: Decompose the image into several parts and describe the fine-grained content of the image part by part, including scenes, objects, relationships, instance location, and any text present.

Figure 4: Different Generic Instructions we adopted in the Ablation Study.

4 Ablation Study

Ablation on Generic Instructions. Within Prism, the generic instructions for visual information extraction are crucial. We experimented with a variety of instructions to elicit the fundamental perceptual capabilities of VLMs, including human-written instructions, GPT-4 generated instructions, and those incorporating chain-of-thought [61] or explicit decomposition, as shown in Fig. 4. We conducted an ablation study on MMStar using the state-of-the-art VLM GPT-4o and LLaVA-NeXT (Yi-34B). As illustrated in Tab. 4, Human 1 outperforms others in eliciting the fundamental perceptual capabilities of the models, while the differences among various instructions are not significant. Therefore, we adopt Human 1 as the generic instruction for all evaluations.

Ablation on the Reasoning Module. The reasoning module is critical for accurately determining the correct answer based on the visual information. To evaluate the impact of the reasoning module on overall performance, we select four LLMs: GPT-3.5-Turbo-0125, GPT-4-Turbo-0125, Llama-3-70B-Instruct, and DeepSeek-v2-Chat, and assess these models across four VLMs with varying capacities. The comparative results are presented in Tab. 5. As a freely available model, GPT-3.5 demonstrates good reasoning performance under our framework, and the more advanced GPT-4 shows improved performance in line with other benchmarks. Notably, Llama3-70B-Instruct, as a representative of open-source models, exhibits competitive capabilities compared to GPT-4-Turbo-0125 under different perceptual conditions. This suggests that open-source models could be valuable for further exploration in visual reasoning.

Ablation on the Vision Backbone. To investigate the impact of the vision encoder on perception ability within the LLaVA architecture, we conduct an ablation study on three pre-trained visual backbones, including CLIP ViT-L/14, SigLip-S0400M, and InternViT-6B. We use InternLM2-7B as the fixed language encoder in LLaVA and train the VLMs for one epoch with the vision encoder fixed. The results in Tab. 6 show that SigLip-S0400M achieves better performance compared to CLIP ViT-L/14 and InternViT-6B on MMStar.

5 PrismCaptioner

Within Prism, we explore the use of small-scale Vision-Language Models (VLMs) as a perception module. We use SigLip as the vision encoder, InternLM2- [1.8B/7B] as the language encoder to develop two visual captioners at different scales, referred as PrismCaptioner- [2B/7B].

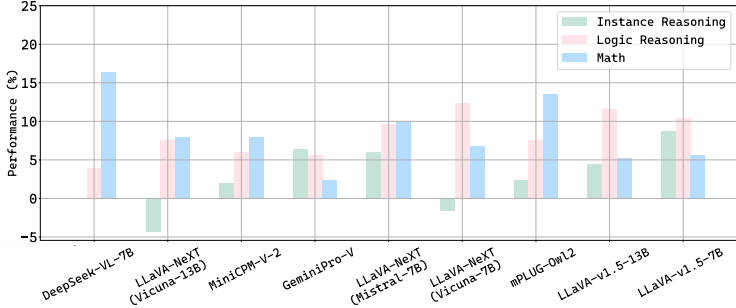


Figure 5: The Performance Changes of Using an External LLM (ChatGPT) for Reasoning of Small Scale VLMs.

Vision Backbone	Overall
SigLip-SO400M	43.5
InternViT-6B	42.7
CLIP ViT-L/14	41.7

Table 6: Ablation on the Vision Backbone.

5.1 Training Details

We perform one-stage training for two epochs using ZeRO2 with XTuner [15] on 8 NVIDIA A800-80GB GPUs, and the training lasts less than a day. The training data include of one copy of ALLaVA-Caption-4V and two copies of Evol-Instruct-GPT4-Turbo-143K. The batch size is set to 16 per device for PrismCaptioner-2B and 8 for PrismCaptioner-7B. We utilized the AdamW optimizer and $2e-4$ learning rate, with the warm-up ratio set to 0.03 and (β_1, β_2) set to (0.9, 0.999). No weight decay was applied and a maximum norm value of 1 is applied for gradient clipping. Full details about QLoRA are presented in Appx. B.2.

5.2 The Performance of PrismCaptioner

We conduct a thorough evaluation of PrismCaptioners across multiple benchmarks, employing GPT-3.5-Turbo-0125 and Llama-3-70B-Instruct as the reasoning module. We utilize MMStar as our primary benchmark to assess comprehensive multimodal capabilities. For domain-specific evaluations, we choose AI2D to gauge diagram comprehension, MMMU for expert knowledge assessment, and MathVista to test mathematical proficiency. The results are presented in Tab. 7. In line with the benchmark selection principles outlined in Sec. 2.2, we apply a consistent filtering strategy across AI2D, MMMU, and MathVista, mirroring that of MMStar, to retain only vision-essential and uncontaminated questions, denoted by the suffix (F). In addition to comparisons with existing open-source and proprietary VLMs, we also assess two baseline models, LLaVA-InternLM2- [1.8B/7B], which are PrismCaptioners trained on LLaVA-v1.5 instruction tuning data. We also compare PrismCaptioners and ShareCaptioner, an open-source VLM designed for generating informative image captions, under identical Prism framework configurations.

As depicted in Tab. 7, each PrismCaptioner, with an external potent LLM for reasoning, markedly surpasses its corresponding end-to-end baseline. PrismCaptioners also outperform ShareCaptioner across all multimodal benchmarks. For the 7B variant, the integration of Llama3 results in a substantial enhancement, positioning PrismCaptioner-7B as a highly competitive vision-language solver, particularly on MMStar and MMMU. For PrismCaptioner-2B, employing ChatGPT yields superior results, outperforming nearly all 7B VLMs in general aptitude, expert knowledge, and mathematical skills. Remarkably, it achieves performance levels on par with some ten times larger VLMs, such as LLaVA-InternLM2-20B, Yi-VL-34B, and Emu2-Chat. This demonstrates that Prism enables the creation of a robust yet efficient vision-language solver, exemplified by PrismCaptioner-2B with ChatGPT, which delivers impressive results.

6 Discussion

Prism’s Value as an Evaluation Framework. Prism’s value as an Evaluation Framework lies in its ability to disentangle and measure the perception and reasoning capabilities of VLMs across various data sources. There exist specialized multimodal benchmarks designed to assess VLMs’ perception and reasoning capabilities, yet they often focus on specific domains. For instance, RealWorldQA [62] evaluates real-world perception with high-resolution images, OCRVQA [44] assesses text recognition in publications, and POPE [30] determines object existence in images. However, many interested domains (geometry, medical images, GUIs, etc.) are not covered by those perception benchmarks. Prism fills this gap by enabling the measurement and comparison of VLMs’ perception capabilities on

Model	MMStar	MMMU	MMMU (F)	MathVista	MathVista (F)	AI2D	AI2D (F)
<i>Proprietary VLMs</i>							
GPT-4o [46]	61.6	62.8	45.9	<u>56.5</u>	<u>46.5</u>	82.2	59.7
GPT-4v [46]	49.7	53.8	42.0	48.7	32.0	75.9	45.7
<i>Open-Source VLMs</i>							
InternVL-Chat-v1.5 [12]	<u>57.1</u>	46.8	<u>33.7</u>	54.7	47.5	80.6	55.0
InternLM-XComposer2 [18]	56.2	41.4	21.2	59.5	49.0	<u>81.2</u>	<u>57.5</u>
LLaVA-NeXT (Yi-34B) [35]	51.6	<u>48.8</u>	26.7	40.4	29.8	78.9	51.6
LLaVA-NeXT (Vicuna-13B) [35]	40.4	37.3	16.5	34.1	19.6	72.2	36.2
LLaVA-NeXT (Mistral-7B) [35]	38.4	37.0	19.2	34.6	21.3	69.0	32.3
LLaVA-NeXT (Vicuna-7B) [35]	37.6	37.6	18.0	31.5	17.1	67.0	29.3
Yi-VL-34B [1]	40.5	45.1	21.2	31.5	12.0	65.9	26.4
Emu2-Chat [53]	40.7	35.0	27.8	30.7	14.0	49.7	22.7
LLaVA-InternLM2-20B [15]	41.9	39.4	18.0	25.3	9.9	65.4	28.4
DeepSeek-VL-7B [39]	40.5	38.3	19.6	36.9	20.7	65.3	36.9
MiniCPM-V-2 [47]	39.1	38.2	21.2	39.8	25.4	62.9	26.9
LLaVA-v1.5-13B [34]	34.3	37.0	15.3	27.7	10.3	61.1	23.0
LLaVA-v1.5-7B [34]	33.1	35.7	15.7	25.6	8.5	55.5	17.8
mPLUG-Owl2 [65]	34.8	34.7	19.6	25.4	8.5	55.7	20.5
<i>Open-Source VLMs (E2E Baselines)</i>							
LLaVA-InternLM2-1.8B [15]	34.5	30.2	18.4	<u>26.3</u>	<u>9.1</u>	43.6	23.5
LLaVA-InternLM2-7B [15]	<u>38.3</u>	<u>40.1</u>	21.9	26.0	7.4	<u>63.6</u>	<u>26.6</u>
<i>Prism Models</i>							
ShareCaptioner-ChatGPT	38.7	45.2	30.2	30.2	12.6	59.6	22.2
PrismCaptioner-2B-ChatGPT	43.3	46.6	32.0	33.5	18.2	62.0	27.4
PrismCaptioner-2B-Llama3	42.0	46.7	32.4	33.5	19.8	59.8	30.3
PrismCaptioner-7B-ChatGPT	43.7	47.3	29.8	35.1	24.6	65.4	30.6
PrismCaptioner-7B-Llama3	<u>45.9</u>	<u>53.3</u>	<u>35.6</u>	<u>39.0</u>	<u>27.3</u>	<u>68.1</u>	<u>37.2</u>

Table 7: **Detail Results of Models under Prism Framework.**¹(F) represents the sub-dataset filtered by our strategy in order to ensure vision indispensability and avoid data leakage. The suffix of Prism models (ChatGPT or Llama3) indicates the reasoning module adopted.

general VQA datasets in these domains. Additionally, existing ‘reasoning’ benchmarks [40, 23] are compositional, requiring VLM to recognize key elements and before reasoning. Comparing an end-to-end VLM with Prism equipped with the same VLM and an external LLM, like ChatGPT, provides insights into the VLM’s intrinsic reasoning capabilities and potential performance constraints.

Prism’s Value as a Vision-Language Solver.

By integrating small-scale VLMs with readily available external LLMs within Prism, we achieve superior performance compared to the standard end-to-end capabilities of the standalone VLM. This approach also renders it practical to address vision-language tasks using a 2B

parameter VLM, as reasoning is effectively outsourced to external LLMs. When implemented with an LLM API, Prism’s inference process (without quantization) only consumes several gigabytes of GPU memory. Furthermore, Prism allows for the flexible incorporation of multiple VLMs to enhance perception. For instance, the straightforward concatenation of outputs from GPT-4v and GeminiPro-V has demonstrated substantial improvements across the majority of metrics on the MMStar benchmark, as substantiated by the data presented in Table 8.

VLM	CP	FP	IR	LR	Math	ST	Overall
GPT-4v	67.2	43.6	56.8	46.8	42.0	30.8	47.9
GeminiPro-V	61.2	35.2	53.2	47.6	48.4	19.2	44.1
Ensemble	66.4	47.2 \uparrow	58.8 \uparrow	50.8 \uparrow	46.4	32.0 \uparrow	50.3 \uparrow

Table 8: **Performance of Perception Module with Multiple VLMs on MMStar.**

¹For comparison, we use single image as input and limit the maximum numbers of output tokens to 512. Better performance of PrismCaptioners is presented in Appx. B.3

Acknowledgement

This project is supported by the National Key R&D Program of China (No.2022ZD0161600), the Shanghai Postdoctoral Excellence Program (No.2023023), China Postdoctoral Science Fund (No.2024M751559), and Shanghai Artificial intelligence Laboratory. This project is funded in part by the Centre for Perceptual and Interactive Intelligence (CPII) Ltd under the Innovation and Technology Commission (ITC)'s InnoHK. Dahua Lin is a PI of CPII under the InnoHK.

References

- [1] 01-ai. Yi-vl. <https://huggingface.co/01-ai/Yi-VL-34B>, 2023.
- [2] Harsh Agrawal, Karan Desai, Yufei Wang, Xinlei Chen, Rishabh Jain, Mark Johnson, Dhruv Batra, Devi Parikh, Stefan Lee, and Peter Anderson. Nocaps: Novel object captioning at scale. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8948–8957, 2019.
- [3] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in Neural Information Processing Systems*, 35:23716–23736, 2022.
- [4] Anthropic. The claude 3 model family: Opus, sonnet, haiku. 2024.
- [5] Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. Qwen technical report. *arXiv preprint arXiv:2309.16609*, 2023.
- [6] 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.
- [7] Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, Xiaoyi Dong, Haodong Duan, Qi Fan, Zhaoye Fei, Yang Gao, Jiaye Ge, Chenya Gu, Yuzhe Gu, Tao Gui, Aijia Guo, Qipeng Guo, Conghui He, Yingfan Hu, Ting Huang, Tao Jiang, Penglong Jiao, Zhenjiang Jin, Zhikai Lei, Jiaying Li, Jingwen Li, Linyang Li, Shuaibin Li, Wei Li, Yining Li, Hongwei Liu, Jiangning Liu, Jiawei Hong, Kaiwen Liu, Kuikun Liu, Xiaoran Liu, Chengqi Lv, Haijun Lv, Kai Lv, Li Ma, Runyuan Ma, Zerun Ma, Wenchang Ning, Linke Ouyang, Jiantao Qiu, Yuan Qu, Fukai Shang, Yunfan Shao, Demin Song, Zifan Song, Zhihao Sui, Peng Sun, Yu Sun, Huanze Tang, Bin Wang, Guoteng Wang, Jiaqi Wang, Jiayu Wang, Rui Wang, Yudong Wang, Ziyi Wang, Xingjian Wei, Qizhen Weng, Fan Wu, Yingtong Xiong, Chao Xu, Ruiliang Xu, Hang Yan, Yirong Yan, Xiaogui Yang, Haochen Ye, Huaiyuan Ying, Jia Yu, Jing Yu, Yuhang Zang, Chuyu Zhang, Li Zhang, Pan Zhang, Peng Zhang, Ruijie Zhang, Shuo Zhang, Songyang Zhang, Wenjian Zhang, Wenwei Zhang, Xingcheng Zhang, Xinyue Zhang, Hui Zhao, Qian Zhao, Xiaomeng Zhao, Fengzhe Zhou, Zaida Zhou, Jingming Zhuo, Yicheng Zou, Xipeng Qiu, Yu Qiao, and Dahua Lin. Internlm2 technical report, 2024.
- [8] Soravit Changpinyo, Piyush Sharma, Nan Ding, and Radu Soricut. Conceptual 12m: Pushing web-scale image-text pre-training to recognize long-tail visual concepts. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3558–3568, 2021.
- [9] 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.

- [10] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. Are we on the right way for evaluating large vision-language models? *arXiv preprint arXiv:2403.20330*, 2024.
- [11] Lin Chen, Jisong Li, Xiaoyi Dong, Pan Zhang, Conghui He, Jiaqi Wang, Feng Zhao, and Dahua Lin. Sharegpt4v: Improving large multi-modal models with better captions. *arXiv preprint arXiv:2311.12793*, 2023.
- [12] Zhe Chen, Weiyun Wang, Hao Tian, Shenglong Ye, Zhangwei Gao, Erfei Cui, Wenwen Tong, Kongzhi Hu, Jiapeng Luo, Zheng Ma, Ji Ma, Jiaqi Wang, Xiaoyi Dong, Hang Yan, Hewei Guo, Conghui He, Botian Shi, Zhenjiang Jin, Chao Xu, Bin Wang, Xingjian Wei, Wei Li, Wenjian Zhang, Bo Zhang, Pinlong Cai, Licheng Wen, Xiangchao Yan, Min Dou, Lewei Lu, Xizhou Zhu, Tong Lu, Dahua Lin, Yu Qiao, Jifeng Dai, and Wenhai Wang. How far are we to gpt-4v? closing the gap to commercial multimodal models with open-source suites, 2024.
- [13] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, Bin Li, Ping Luo, Tong Lu, Yu Qiao, and Jifeng Dai. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. *arXiv preprint arXiv:2312.14238*, 2023.
- [14] OpenCompass Contributors. Opencompass: A universal evaluation platform for foundation models. <https://github.com/open-compass/opencompass>, 2023.
- [15] XTuner Contributors. Xtuner: A toolkit for efficiently fine-tuning llm. <https://github.com/InternLM/xtuner>, 2023.
- [16] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. *arXiv preprint arXiv:2305.06500*, 2023.
- [17] Tim Dettmers, Artidoro Pagnoni, Ari Holtzman, and Luke Zettlemoyer. Qlora: Efficient finetuning of quantized llms, 2023.
- [18] Xiaoyi Dong, Pan Zhang, Yuhang Zang, Yuhang Cao, Bin Wang, Linke Ouyang, Xilin Wei, Songyang Zhang, Haodong Duan, Maosong Cao, Wenwei Zhang, Yining Li, Hang Yan, Yang Gao, Xinyue Zhang, Wei Li, Jingwen Li, Kai Chen, Conghui He, Xingcheng Zhang, Yu Qiao, Dahua Lin, and Jiaqi Wang. Internlm-xcomposer2: Mastering free-form text-image composition and comprehension in vision-language large model. *arXiv preprint arXiv:2401.16420*, 2024.
- [19] 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 *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 320–335, 2022.
- [20] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Zhenyu Qiu, Wei Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, and Rongrong Ji. Mme: A comprehensive evaluation benchmark for multimodal large language models. *ArXiv*, abs/2306.13394, 2023.
- [21] Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but not perceive. *arXiv preprint arXiv:2404.12390*, 2024.
- [22] Yash Goyal, Tejas Khot, Douglas Summers-Stay, Dhruv Batra, and Devi Parikh. Making the v in vqa matter: Elevating the role of image understanding in visual question answering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6904–6913, 2017.
- [23] Xiaotian Han, Quanzeng You, Yongfei Liu, Wentao Chen, Huangjie Zheng, Khalil Mrini, Xudong Lin, Yiqi Wang, Bohan Zhai, Jianbo Yuan, et al. Infimm-eval: Complex open-ended reasoning evaluation for multi-modal large language models. *arXiv e-prints*, pages arXiv–2311, 2023.

- [24] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 6700–6709, 2019.
- [25] Wang Jin, Dong Shichao, Zhu Yapeng, Yao Kelu, Zhao Weidong, Li Chao, and Luo Ping. Diagnosing the compositional knowledge of vision language models from a game-theoretic view. In *To appear in ICML*, 2024.
- [26] Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images. In *Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part IV 14*, pages 235–251. Springer, 2016.
- [27] Hugo Laurençon, Lucile Saulnier, Léo Tronchon, Stas Bekman, Amanpreet Singh, Anton Lozhkov, Thomas Wang, Siddharth Karamcheti, Alexander M. Rush, Douwe Kiela, Matthieu Cord, and Victor Sanh. Obelics: An open web-scale filtered dataset of interleaved image-text documents, 2023.
- [28] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. *arXiv preprint arXiv:2307.16125*, 2023.
- [29] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [30] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. *arXiv preprint arXiv:2305.10355*, 2023.
- [31] 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 *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014.
- [32] Fangyu Liu, Guy Edward Toh Emerson, and Nigel Collier. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 2023.
- [33] Fuxiao Liu, Tianrui Guan, Zongxia Li, Lichang Chen, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: You see what you think? or you think what you see? an image-context reasoning benchmark challenging for gpt-4v (ision), llava-1.5, and other multi-modality models. *arXiv preprint arXiv:2310.14566*, 2023.
- [34] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. *arXiv preprint arXiv:2310.03744*, 2023.
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. Llava-next: Improved reasoning, ocr, and world knowledge, January 2024.
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *arXiv preprint arXiv:2304.08485*, 2023.
- [37] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, et al. Mmbench: Is your multi-modal model an all-around player? *arXiv preprint arXiv:2307.06281*, 2023.
- [38] Yuliang Liu, Zhang Li, Hongliang Li, Wenwen Yu, Mingxin Huang, Dezhi Peng, Mingyu Liu, Mingrui Chen, Chunyuan Li, Lianwen Jin, et al. On the hidden mystery of ocr in large multimodal models. *arXiv preprint arXiv:2305.07895*, 2023.
- [39] Haoyu Lu, Wen Liu, Bo Zhang, Bingxuan Wang, Kai Dong, Bo Liu, Jingxiang Sun, Tongzheng Ren, Zhuoshu Li, Yaofeng Sun, et al. Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525*, 2024.

- [40] Pan Lu, Hritik Bansal, Tony Xia, Jiacheng Liu, Chunyuan Li, Hannaneh Hajishirzi, Hao Cheng, Kai-Wei Chang, Michel Galley, and Jianfeng Gao. Mathvista: Evaluating mathematical reasoning of foundation models in visual contexts. *arXiv preprint arXiv:2310.02255*, 2023.
- [41] Kenneth Marino, Mohammad Rastegari, Ali Farhadi, and Roozbeh Mottaghi. Ok-vqa: A visual question answering benchmark requiring external knowledge. In *Proceedings of the IEEE/cvf conference on computer vision and pattern recognition*, pages 3195–3204, 2019.
- [42] Ahmed Masry, Do Xuan Long, Jia Qing Tan, Shafiq Joty, and Enamul Hoque. Chartqa: A benchmark for question answering about charts with visual and logical reasoning. *arXiv preprint arXiv:2203.10244*, 2022.
- [43] Minesh Mathew, Dimosthenis Karatzas, and CV Jawahar. Docvqa: A dataset for vqa on document images. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 2200–2209, 2021.
- [44] Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. Ocr-vqa: Visual question answering by reading text in images. In *2019 international conference on document analysis and recognition (ICDAR)*, pages 947–952. IEEE, 2019.
- [45] OpenAI. Chatgpt. <https://openai.com/blog/chatgpt>, 2023.
- [46] OpenAI. Gpt-4 technical report. *ArXiv*, abs/2303.08774, 2023.
- [47] OpenBMB. Minicpm: Unveiling the potential of end-side large language models, 2024.
- [48] Aitor Ormazabal, Che Zheng, Cyprien de Masson d’Autume, Dani Yogatama, Deyu Fu, Donovan Ong, Eric Chen, Eugenie Lamprecht, Hai Pham, Isaac Ong, et al. Reka core, flash, and edge: A series of powerful multimodal language models. *arXiv preprint arXiv:2404.12387*, 2024.
- [49] 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.
- [50] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [51] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021.
- [52] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Dhruv Batra, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8317–8326, 2019.
- [53] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiying Yu, Zhengxiong Luo, Yueze Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, et al. Generative multimodal models are in-context learners. *arXiv preprint arXiv:2312.13286*, 2023.
- [54] Quan Sun, Yuxin Fang, Ledell Wu, Xinlong Wang, and Yue Cao. Eva-clip: Improved training techniques for clip at scale. *arXiv preprint arXiv:2303.15389*, 2023.
- [55] Gemini Team, Rohan Anil, Sebastian Borgeaud, Yonghui Wu, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023.
- [56] InternLM Team. Internlm: A multilingual language model with progressively enhanced capabilities. <https://github.com/InternLM/InternLM-techreport>, 2023.

- [57] Tristan Thrush, Ryan Jiang, Max Bartolo, Amanpreet Singh, Adina Williams, Douwe Kiela, and Candace Ross. Winoground: Probing vision and language models for visio-linguistic compositionality. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5228–5238, 2022.
- [58] Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide shut? exploring the visual shortcomings of multimodal llms. *arXiv preprint arXiv:2401.06209*, 2024.
- [59] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [60] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [61] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models, 2023.
- [62] XAI. Grok-1.5 vision preview. 2024.
- [63] Peng Xu, Wenqi Shao, Kaipeng Zhang, Peng Gao, Shuo Liu, Meng Lei, Fanqing Meng, Siyuan Huang, Yu Jiao Qiao, and Ping Luo. Lvlm-ehub: A comprehensive evaluation benchmark for large vision-language models. 2023.
- [64] Qinghao Ye, Haiyang Xu, Guohai Xu, Jiabo Ye, Ming Yan, Yiyang Zhou, Junyang Wang, Anwen Hu, Pengcheng Shi, Yaya Shi, et al. mplug-owl: Modularization empowers large language models with multimodality. *arXiv preprint arXiv:2304.14178*, 2023.
- [65] Qinghao Ye, Haiyang Xu, Jiabo Ye, Ming Yan, Anwen Hu, Haowei Liu, Qi Qian, Ji Zhang, Fei Huang, and Jingren Zhou. mplug-owl2: Revolutionizing multi-modal large language model with modality collaboration, 2023.
- [66] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. *arXiv preprint arXiv:2311.16502*, 2023.
- [67] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11975–11986, 2023.
- [68] Jianrui Zhang, Mu Cai, Tengyang Xie, and Yong Jae Lee. Countercurate: Enhancing physical and semantic visio-linguistic compositional reasoning via counterfactual examples. *Findings of the Association for Computational Linguistics: ACL 2024*, 2024.
- [69] Yizhe Zhang, He Bai, Ruixiang Zhang, Jiatao Gu, Shuangfei Zhai, Josh Susskind, and Navdeep Jaitly. How far are we from intelligent visual deductive reasoning? *arXiv preprint arXiv:2403.04732*, 2024.
- [70] 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.

A Related Work

A.1 Large Vision-Language Models (LVLMs)

The landscape of Large Language Models (LLMs) is continually evolving, with an expanding body of research focused on integrating multimodal capabilities to enhance their perceptual abilities in real-world contexts [49, 70, 29]. Early efforts in this direction, such as Flamingo [3], introduced gated dense blocks of cross-attention within pre-trained language encoder layers to fuse visual features. Subsequent models like BLIP2 [29] and InstructBLIP [16] utilized a Q-former to align features across different modalities, enabling tasks such as zero-shot visual question answering. More recent models, including LLaVA [36] and MiniGPT-4 [70], have simplified modality bridging through the use of MLP-based projection layers, offering a more straightforward approach compared to the Q-former. The architecture of LLaVA has been widely adopted in subsequent works [35, 1, 12, 39]. The choice of vision encoders and the language model is considered critical for the overall performance of VLMS. Most VLMs [36, 18, 35, 27, 70] employ CLIP-based Vision Transformer [50, 67, 54] as the vision encoder, owing to its good pre-training alignment of visual and textual modalities. There is a prevalent belief [37, 35] that the scale of the language model significantly impacts the performance of VLMs, though detailed analyses are lacking. In addition to the open-source VLMs developed by the academic community, numerous proprietary VLMs [46, 55, 6, 48] demonstrate robust performance across various multimodal benchmarks. This paper presents a breakdown capability analysis of both open-source and proprietary VLMs using the Prism framework.

A.2 LVLMs Capability Evaluation

Large-scale VLMs have demonstrated promising outcomes across a diverse range of multimodal tasks, as evidenced by extensive qualitative and quantitative evaluations. Early assessments of LVLMs often involved open-ended Visual Question Answering (VQA) [2, 24, 41, 22] and human-based subjective evaluations [64, 63]. However, these methods face limitations in accurately reflecting the true performance of VLMs. Open-ended VQA tasks typically demand an exact match between the model’s prediction and the ground truth, which can lead to a significant number of false positives. Conversely, subjective evaluations introduce biases and make the results challenging to reproduce. Subsequent research has shifted towards structuring visual questions in closed-ended formats, such as multiple-choice or Yes-or-No questions. Pioneering works like MMBench [37], MME [20], or SEED-Bench [28] have presented comprehensive evaluations of VLMs using closed-ended VQA, covering various perception and reasoning capabilities. Additionally, specialized multimodal benchmarks have emerged to assess VLMs from specific angles. For instance, MMMU [66] evaluates VLMs’ ability to handle multimodal examination questions, while POPE [30] and HallusionBench [33] scrutinize hallucination and illusion phenomena in VLMs. RealWorldQA [62] focuses on real-world perception with high-resolution images. Recently, MMStar [10] has addressed the issues of vision dispensability and data contamination in existing benchmarks by compiling high-quality, vision-indispensable questions from multiple sources, ensuring minimal data leakage and covering six core capabilities. In this study, Prism primarily utilizes MMStar for capability evaluation.

A.3 LVLMs Capability Breakdown Investigation

To offer insightful and detailed feedback for the future optimization of VLMs, some researchers have initiated efforts to dissect the abilities of VLMs, aiming to uncover strategies for enhancement. To more effectively explore the disparity between VLMs and human cognition, Zhang et al. [69] employ Raven’s Progressive Matrices (RPM) to examine the model’s deductive reasoning skills grounded in visual perception. Through error case analysis, the authors observe that VLMs often make compounding and confounding errors when articulating individual elements within RPM, which subsequently results in erroneous reasoning. Although this study provides qualitative insights, it does not encompass a systematic evaluation of perception and reasoning capabilities. Meanwhile, researchers have developed specialized benchmarks to assess specific capabilities. For example, InfiMM-Eval [23] and MathVista [40] have implemented rigorous, step-by-step evaluations of the model’s complex reasoning abilities on natural images and mathematical VQA problems, respectively. However, these reasoning benchmarks require a foundational perception capability to accurately identify key elements. Concurrently, there are perception benchmarks [58, 62, 30, 21] that exclusively focus on evaluating the perception skills of VLMs across various scenarios. Moreover, some

specific crucial capabilities of VLMs have received researchers’ attention, such as fine-grained understanding [68] and visio-linguistic compositional knowledge [57, 68, 25]. In this study, Prism introduces a general decoupling framework that facilitates a detailed analysis of perception and reasoning capabilities, applicable to any multimodal benchmark.

B Supplementary Details of Prism

B.1 Details of Prism Evaluation Framework

B.1.1 Query-Specific Instruction Details

Query-Specific Instruction is a combination of generic instruction and query-specific part, as depicted in Fig. 6. To ensure that the query-specific part generated by the reasoning module is closely related to the questions and options, we adopt few-shot learning to guide the LLM. For each visual question, we feed the request with multiple examples. This approach helps the reasoning module understand what “contents to observe” means in different contexts and allows it to make accurate inferences in response to specific questions, as illustrated by the prompts in Fig. 7.

Instructions

Generic Instruction: Describe the fine-grained content of the image, including scenes, objects, relationships, instance location, and any text present.

Query-Specific Instruction: Describe the fine-grained content of the image, including scenes, objects, relationships, instance location, and any text present. Especially, pay attention to <query-specific part>.

Figure 6: **Generic Instruction vs. Query-Specific Instruction.**

The Few-shot Prompt Template

Your task is to give a concise instruction about what basic elements are needed to be described based on the given question. Ensure that your instructions do not cover the raw question, options, or thought process of answering the question.

Question: In which period the number of full-time employees is the maximum?
Contents to observe: the number of full-time employees

Question: What is the value of the smallest bar?
Contents to observe: the heights of all bars and their values

Question: What is the main subject of the image?
Contents to observe: the central theme or object

Question: What is the position of the catcher relative to the home plate?
Contents to observe: the spatial arrangement of the objects

Question: What is the expected ratio of offspring with white spots to offspring with solid coloring? Choose the most likely ratio.
Contents to observe: the genetic information

Question: <question>
Contents to observe:

Figure 7: **The Prompt Template for the Reasoning Module to Generate the "Contents to Observe" Part.**

B.1.2 Inference Prompt Template of the Reasoning Module

After the perception module generates detailed visual information about the image, the query needs to be reformatted to enable the reasoning module to answer more accurately based on the information

Hyperparameter	Assignment
Image Resolution	384 * 384
Patch Size	14
Max Length	1296
QLoRA Quantization	4bit
LLM int8 Threshold	6.0
BnB 4bit Compute Dtype	torch.float16
BnB 4bit Double Quant	True
BnB 4bit Quant Dtype	nf4
LoRA r	512
LoRA α	256
LoRA Dropout	0.05

Table 9: PrismCaptioner Training Details.

and the question. The template for reformatting is shown in Fig. 8. The reformatting process involves simple splicing, making it intuitive for the reasoning module to respond directly. An example of this can be seen in Fig. 9.

Reformatting Template

You are an excellent text-based reasoning expert. You are required to answer the question based on the detailed description of the image.

Description: <description>

Question: <question>

Figure 8: The Template for Reformatting Query.

Reformatted Query Example

You are an excellent text-based reasoning expert. You are required to answer the question based on the detailed description of the image.

Description: The image presents a delightful celebration of Father’s Day. Dominating the center of the image is a blue tie, adorned with white stripes, symbolizing the essence of fatherhood. The tie is slightly tilted to the right, adding a touch of dynamism to the composition. On the left side of the tie, the phrase "Happy Father’s Day" is elegantly inscribed in a white cursive font, extending warm wishes to all dads. The text and the tie are set against a dark blue background, creating a striking contrast that draws attention to the main elements of the image. Adding a final touch of sophistication, a thin white border frames the entire image, encapsulating the joyous message of Father’s Day. The image, in its entirety, serves as a heartfelt tribute to all the wonderful fathers out there.

Question: Which special day is associated with this poster?
Options: A. Earth Day. B. National Reading Day. C. Father’s Day. D. Mother’s Day
Please select the correct answer from the options above.

Figure 9: An Example of the Reformatted Query.

B.2 More Details of PrismCaptioner Training

The details about QLoRA for training PrismCaptioner is demonstrated in Tab. 9

Model	MMStar	MMMU	MMMU (F)	MathVista	MathVista (F)	AI2D	AI2D (F)
<i>Prism Models (Multiple Image Inputs & Max Output Length 2048)</i>							
PrismCaptioner-2B-ChatGPT	43.0	46.1	31.4	34.8	16.7	62.1	27.9
PrismCaptioner-2B-Llama3	42.6	49.0	34.1	35.1	19.8	59.8	30.3
PrismCaptioner-7B-ChatGPT	43.4	47.9	33.7	36.5	25.0	65.4	30.8
PrismCaptioner-7B-Llama3	<u>46.2</u>	<u>56.5</u>	<u>42.0</u>	<u>39.8</u>	<u>26.2</u>	<u>67.9</u>	<u>37.7</u>

Table 10: **More Detailed Performance Results of PrismCaptioners**

B.3 More Performance Details of PrismCaptioner

In addition to the results of Tab. 7, we use multiple images as inputs if there are and set maximum output length to 2048. The performance is presented in Tab. 10

C Detailed Examples of Prism

C.1 End-to-end v.s. Prism Predictions

Some end-to-end models with small language models, often predict incorrect answers due to their limited reasoning capabilities. Utilizing an external LLM, Prism endows VLMs with the reasoning ability to solve the vision-language tasks, as presented in Tab. 11 and Tab. 12.

C.2 PrismCaptioner Performances

PrismCaptioner can professionally extract and express detailed visual information to solve coarse perception, fine-grained perception, instance reasoning, logical reasoning, science & technology, and math tasks, as presented in Tabs. 13 to 18.

C.3 Performances Comparison between PrismCaptioner and GPT-4o

Although PrismCaptioner can generate detailed descriptions of images, there is still room for improvement compared to GPT-4o.

- (1) GPT-4o generates descriptions that are more relevant to the questions. Due to limited training data, PrismCaptioner’s captions are less adaptive to query-specific parts compared to those of GPT-4o, as shown in Tab. 19.
- (2) GPT-4o’s expression is more detailed and specific, and the responses about spatiotemporal information are also more accurate, as shown in Tab. 20.

D Limitations and Broader Impacts

Limitations. In this study, we introduce the Prism framework and showcase its effectiveness as both an analytical tool and a versatile vision-language task solver. Prism relies on language descriptions while natural languages may struggle to express some obscure visual elements, especially some aesthetic contents like abstract artworks and surrealist paintings. It is difficult to obtain high-quality descriptions in these cases. Given budget constraints, our evaluation focuses on a select group of representative open-source and proprietary VLMs, which may not encompass all the most recent high-performing models. Our experiments with training visual captioners aim to illustrate Prism’s capability to deliver strong performance on vision-language tasks while minimizing costs. To fully realize the potential of Prism, additional experiments are recommended, particularly on domain-specific visual instruction tuning data, such as tables and diagrams, screens, and graphical user interfaces (GUIs), medical images, etc., to thoroughly assess Prism’s efficacy in these specialized contexts.

Broader Impacts. As an analytical framework, Prism offers detailed insights into the perception and reasoning capabilities of vision-language models (VLMs), providing valuable guidance for future model optimization. Training large-scale VLMs necessitates extensive multi-modal data and

Question: What is the primary scene depicted in the image?

Image:



Options:

- A. A group of people taking pictures of a statue in a plaza
- B. A large crowd taking pictures of each other in an outdoor event
- C. A group of people taking pictures of a Christmas tree in a plaza
- D. A group of people taking pictures of a famous building in New York City

Category: Coarse Perception

Model: LLaVA-NeXT (Vicuna-7B)

End-to-end Prediction: D. A group of people taking pictures of a famous building in New York City (✘)

Perception Description: The image captures a bustling scene in a city square, where a large crowd of people has gathered. The crowd is diverse, with individuals of various ages and attire, all seemingly engaged in the event taking place. The central focus of the image is a large, illuminated Christmas tree, which stands prominently in the square. The tree is adorned with a multitude of colorful lights and ornaments, creating a festive atmosphere.

In the background, the cityscape is visible with tall buildings, some of which are adorned with American flags, suggesting a location in the United States. The flags are flying at half-mast, adding a solemn tone to the scene.

The sky above is overcast, casting a soft light over the scene. The overall mood of the image is one of celebration and unity, as the crowd seems to be enjoying the holiday festivities. The image does not contain any text. The relative positions of the objects suggest a well-organized event, with the Christmas tree as the focal point and the crowd surrounding it. The flags in the background provide a sense of location and national pride. The overcast sky adds a touch of solemnity to the otherwise joyful scene.

Perception Prediction: C. (✔)

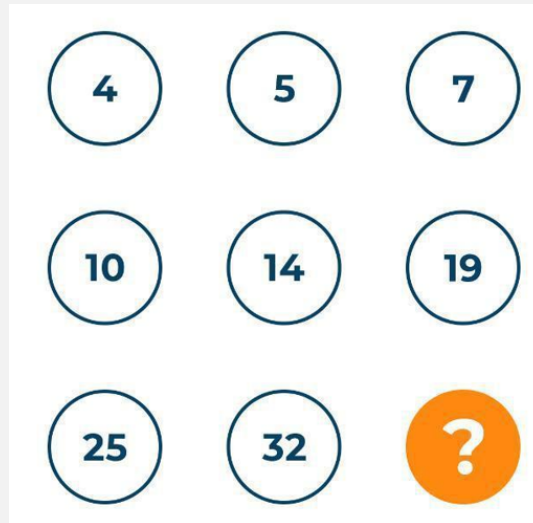
Table 11: An Example of End-to-End v.s. Prism LLaVA-NeXT (Vicuna-7B) Prediction.

computational resources. Prism mitigates this challenge by training small-scale VLMs that specialize in visual captioning tasks and perform reasoning with large language models (LLMs), which are now cost-effective.² When employed as a vision-language task solver, Prism can be trained and deployed at a significantly reduced cost, making it a promising approach for highly customized applications or tasks with limited training data. However, Prism also carries potential societal impacts, as it could lower the barrier to building multimodal applications, some of which may be harmful. There is a risk that Prism could be used to develop harmful multimodal AI systems. Additionally, data-driven methods often inherit biases, which can persist in downstream tasks. We urge users to thoughtfully consider the implications of these biases when implementing our model.

²Thanks to the development in deploying technologies, the inference of LLM is now cheap, provided by various corporations at a price as low as millions of tokens per US dollar.

Question: Which number is missing?

Image:



Options: A. 40 B. 10 C. 22 D. 34

Category: Logical Reasoning

Model: MiniCPM-V-2

End-to-end Prediction: B. 10 (✘)

Perception Description: The image presents a grid of nine circles, each containing a number from 4 to 25. The numbers are arranged in a 3x3 grid, with the numbers 4, 5, 7, 10, 14, 19, 25, 32, and a question mark (?) at the bottom right corner. **The sequence of numbers in the grid is 4, 5, 7, 10, 14, 19, 25, 32, with the question mark indicating the missing number.**

Perception Prediction: The pattern in the sequence is as follows: $5 - 4 = 1$, $7 - 5 = 2$, $10 - 7 = 3$, $14 - 10 = 4$, $19 - 14 = 5$, $25 - 19 = 6$, $32 - 25 = 7$. Therefore, the missing number should be $32 + 8 = 40$.

So, the correct answer is: A. 40 (✔)

Table 12: An Example of End-to-End v.s. Prism MiniCPM-V-2 Prediction.

Question: Which image is more brightful?

Image:



Options: A. The first image B. The second image

Category: Coarse Perception

Description: The image is a split composition with two distinct halves.

On the left side, there is a photograph of a **bright red sports car parked on the side of a road**. The car appears to be a modern coupe with a sleek design, featuring a long hood and a short rear deck, which are characteristic of high-performance vehicles. It has five-spoke alloy wheels with a silver finish and low-profile tires. The car's windows are tinted, and it has a visible antenna on the rear part of the roof. The background shows a blurred natural landscape with greenery, suggesting a rural or semi-rural setting.

On the right side of the image, there is a digital artwork depicting a **swirling, fractal-like pattern in shades of green and purple, set against a dark background**. The pattern has a hypnotic, almost hypnotic quality, with a sense of motion and energy emanating from it. The artwork has a hypnotic, almost hypnotizing effect, drawing the viewer's attention to the center of the swirl where the colors seem to converge.

Prediction: A. The first image (✓)

Table 13: PrismCaptioner Generates Description to Solve Coarse Perception Task.

Question: What color is the shirt of the man playing the guitar?
Image:



Options: A. Blue B. Brown C. Red D. Black

Category: Fine-grained Perception

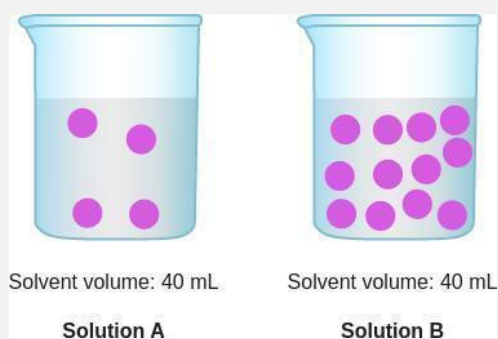
Description: The image depicts a live music performance by a band on stage. There are three members visible: a vocalist, a guitarist, and a keyboard player. The vocalist, positioned in the center, is a woman with long hair, wearing a red top and a black skirt. She is holding a microphone stand with her right hand and appears to be singing into it. To her left, a man is playing an electric guitar, dressed in a black shirt and jeans, focused on his instrument. To the right of the vocalist, another man is playing a keyboard, wearing a black shirt and dark pants. The stage is equipped with various musical instruments and equipment, including a drum set behind the vocalist, a keyboard in front of the keyboard player, and a microphone stand in front of the vocalist. The lighting is dim, with spotlights highlighting the performers, and the background is a draped curtain in shades of blue and red. The atmosphere suggests an intimate concert setting.

Prediction: D. Black (✓)

Table 14: PrismCaptioner Generates Description to Solve Fine-Grained Perception Task.

Question: Which solution has a higher concentration of pink particles?

Image:



Options: A. neither; their concentrations are the same B. Solution A C. Solution B

Category: Instance Reasoning

Description: The image displays two beaker-like containers, labeled Solution A and Solution B, each containing a clear liquid. Solution A has a volume of 40 mL, while Solution B has a volume of 40 mL. Inside Solution A, there are multiple small, round, purple particles dispersed uniformly throughout the liquid. In contrast, Solution B contains a similar number of purple particles, but they are clustered more densely, suggesting a higher concentration. The background is plain white, focusing attention on the containers and their contents. There are no additional background elements or context provided.

Prediction: C. Solution B (✓)

Table 15: PrismCaptioner Generates Description to Solve Instance Reasoning Task.

Question: In nature, what's the relationship between these two creatures?

Image:



Options:

- A. Predatory relationships
- B. Competitive relationships
- C. Parasitic relationships
- D. Symbiotic relationship

Category: Logical Reasoning

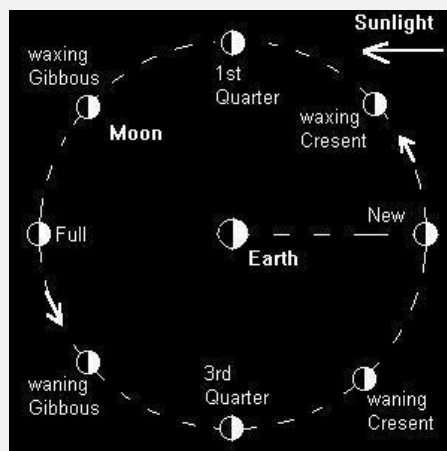
Description: The image depicts two antelopes engaged in what appears to be a **physical confrontation** on a grassy savanna. The antelope on the left is rearing up on its hind legs, with its front legs extended towards the other antelope, which is also on its hind legs, facing the first antelope. Both animals have prominent, curved horns, and their fur is a rich, reddish-brown color. The background is a blurred natural landscape, suggesting a wild, open environment typical of a savanna ecosystem. There are no human elements or artificial structures visible, reinforcing the natural setting. The colors are vibrant, with the green of the grass contrasting with the reddish-brown of the antelopes.

Prediction: B. Competitive relationships (✓)

Table 16: PrismCaptioner Generates Description to Solve Logical Reasoning Task.

Question: What comes after the full moon?

Image:



Options: A. waning crescent B. 3rd quarter C. waning Gibbous D. new moon

Category: Science & Technology

Description: The image depicts a simplified diagram illustrating the phases of the Moon in relation to the Earth and the Sun. It is a circular diagram with the Sun at the top right and the Earth at the bottom left. **The Moon is shown as a crescent in the top left, waxing crescent, first quarter, waxing gibbous, full, waning gibbous, third quarter, and waning crescent.** Arrows indicate the direction of the waxing and waning of the Moon's phases. The background is solid black, emphasizing the diagram's white and gray elements. There are no background elements or colors other than the white and gray of the Moon and Earth, and the black of the background.

Prediction: C. waning Gibbous. (✓)

Table 17: PrismCaptioner Generates Description to Solve Science & Technology Task.

Question: For case A accompanying table, answer the questions that follow. Calculate the future value of the annuity, assuming that it is an ordinary annuity.

Image:

Case	Amount of annuity	Interest rate	Deposit period (years)
A	\$ 2,500	8%	10
B	500	12	6
C	30,000	20	5
D	11,500	9	8
E	6,000	14	30

Options: A. \$126,827 B. \$223,248 C. \$36,216 D. \$4,058

Category: Math

Description: The image displays a table with five rows and four columns, excluding the header row. The header row is labeled from left to right as "Case," "Amount of annuity," "Interest rate," and "Deposit period (years)." Below the header, there are five cases labeled A, B, C, D, and E. The "Amount of annuity" column lists monetary values: \$2,500 for case A, \$500 for case B, \$30,000 for case C, \$11,500 for case D, and \$6,000 for case E. The "Interest rate" column lists percentage values: 8% for case A, 12% for case B, 20% for case C, 9% for case D, and 14% for case E. The "Deposit period (years)" column lists numerical values: 10 years for case A, 6 years for case B, 5 years for case C, 8 years for case D, and 30 years for case E. The table is set against a plain background with no additional elements or decorations.

Prediction: To calculate the future value of an ordinary annuity, you can use the formula:

$$FV = Pmt * [(1 + r)^n - 1]/r,$$

where FV = Future Value, Pmt = Amount of annuity (\$2,500), r = Interest rate (8% or 0.08), n = Deposit period (10 years).

Plugging in the values:

$$FV = \$2,500 * [(1 + 0.08)^{10} - 1]/0.08$$

$$FV = \$2,500 * [(1.08)^{10} - 1]/0.08$$

$$FV = \$2,500 * [2.158924 - 1]/0.08$$

$$FV = \$2,500 * 1.158924/0.08$$

$$FV = \$36,216.55$$

Therefore, the correct answer is: C. \$36,216 (✓)

Table 18: PrismCaptioner Generates Description to Solve Math Task.

Question

Where is the rug in the living room located?



Options

- A. Next to the door B. Under the table
C. In front of the window D. Under the couch

Models

GPT-4o

PrismCaptioner

Descriptions

The image depicts a cozy living room corner with a modern and inviting aesthetic. Here is a detailed description of the scene:

1. Furniture and Objects:

Sofa: A gray upholstered sofa is positioned on the left side of the image. It has a few decorative pillows in various patterns and colors, including white, green, and teal.

Side Table: Next to the sofa, there is a wooden side table with two drawers. On top of the table, there is a blue ceramic lamp with a white lampshade and a potted plant with large green leaves.

Rug: A gray rug with a white geometric pattern is placed on the floor. The rug is positioned in front of the sofa and extends partially under it, covering the wooden floor.

Curtains: Light gray curtains with a subtle pattern hang from a black curtain rod above a window on the left side of the image.

Wall Art: Above the side table, there is a framed piece of art featuring a colorful floral arrangement.

2. Placement and Relationships:

The rug is centrally placed in front of the sofa, extending slightly under it, and covers a portion of the wooden floor.

The side table is positioned between the sofa and the door, with the lamp and plant on top providing both functional and decorative elements.

The framed art is centered above the side table, adding a pop of color to the wall.

...

Overall, the arrangement of furniture and decor items creates a balanced and aesthetically pleasing living room corner. The rug's placement in relation to the sofa and side table helps define the seating area and adds warmth to the space.

The image shows a neatly arranged living room corner with a focus on a side table and a portion of a couch. The side table, which appears to be made of wood with a natural finish, has two drawers with metal handles. On top of the table, there is a large blue lamp with a white lampshade, a small white bowl, and a green potted plant with broad leaves. To the left of the side table, there is a couch with a grey upholstery, adorned with a teal throw pillow and a pink throw blanket casually draped over it. The floor is covered with a grey and white patterned area rug, and the walls are painted in a light grey tone. A window with white trim is partially visible, dressed with patterned curtains in shades of grey and white. On the wall, there is a framed artwork featuring a bouquet of flowers in a vase. The door in the background is a dark grey with a black doorknob, and there is a light switch plate visible on the wall to the right of the door. The overall aesthetic of the room is modern with a touch of warmth provided by the wooden furniture and the cozy textiles.

Predictions

D. Under the couch (✓)

C. In front of the window (✗)

Table 19: Performances between GPT-4o and PrismCaptioner.

Question **What is the relative position of the man and the woman sitting at the table?**



Options

- A. The man is on the left and the woman is on the right
- B. The man is on the right and the woman is on the left
- C. The man and the woman are facing each other
- D. The man and the woman are back-to-back

Models

Descriptions

GPT-4o

The image depicts a man and a woman sitting closely together at a table in what appears to be a social setting, possibly a bar or a restaurant. **The woman, on the left**, has blonde hair styled in loose curls and is wearing a white dress with a plunging neckline. She is leaning in towards the man, appearing to whisper something in his ear. Her left hand is resting on the table, and her right hand is not visible.

The man, on the right, is wearing a dark suit with a white shirt and a dark tie. He is holding a cigarette in his right hand, which is raised near his face. His left hand is resting on the table. **He is looking directly at the camera with a slight smile.**

On the table in front of them, there is a champagne glass near the woman and a glass with a dark liquid, possibly a cocktail, near the man. There are also other objects on the table, including what appears to be a hand of another person partially visible on the left side of the image. In the background, there is a colorful arrangement of red and white flowers on the left side. The backdrop includes a painted mural or artwork featuring abstract shapes and possibly a cityscape with buildings and a sailboat.

The spatial arrangement shows the man and woman sitting very close to each other, with the woman leaning in towards the man, indicating a sense of intimacy or a private conversation. **The man is positioned slightly turned towards the camera, while the woman is more focused on him.**

PrismCaptioner

The image depicts a man and a woman seated at a table in what appears to be a bar or restaurant setting. The man is dressed in a dark suit with a white shirt and a dark tie, and he is holding a cigarette between his fingers. The woman is wearing a white dress with a plunging neckline, and her blonde hair is styled in curls.

Both individuals are looking at each other, and the woman seems to be whispering into the man's ear. On the table, there are two glasses, one of which appears to contain a clear liquid, possibly water, and the other is not fully visible. There is also a small, dark object that could be an ashtray or a coaster.

In the background, there is a bouquet of flowers with red and white blooms, and the wall behind them features a large, abstract painting with blue and yellow tones. The lighting in the scene is warm, and the overall ambiance suggests a casual yet intimate setting.

Predictions

A. The first image (✓)

B. The second image (✗)

Table 20: **Performances between GPT-4o and PrismCaptioner.**

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: As denoted by the abstract and introduction, in this work, we present the Prism framework that can serve as an analytical framework and a general vision-language task solver. Qualitative and quantitative results in this paper can support the contributions.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: The limitations of Prism are discussed in Appx. D.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: The paper does not include theoretical results.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: In Secs. 3.1 and 5.1, we describe the hyper-parameters and settings of Prism inference and the training settings of VLMs trained in this work in detail. So it would be easy to reproduce the results we reported.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: Data, codes, and other resources required to reproduce the main results will be released later. We need some time to prepare a more tidy version of those resources.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: Those settings are described in In Secs. 3.1 and 5.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [No]

Justification: During the inference process in Prism, we perform greedy decoding for both VLMs and LLMs, so the quantitative results we reported will be stable. Besides, reporting error bars will be expensive and is not a common practice for research in large vision-language models.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.

- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: These information is included in Sec. 5.1.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: This research conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The potential impacts of this work are discussed in Appx. D.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This work will release the evaluation results and some weights of small-scale VLMs trained on publicly available dataset. We believe this paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: All assets used in this work are properly credited. Our evaluation codes are modified based on VLMEvalKit, which goes with Apache 2.0 License. Our self-trained vision captioner is built with SigLip and InternLM2, and trained on ALLaVA and Evol-Instruct-GPT-4-Turbo, all of those assets go with Apache 2.0 License. Other evaluation datasets adopt in this paper go with CC-BY 4.0 License.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.

- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: The paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.