

Spatial457: A Diagnostic Benchmark for 6D Spatial Reasoning of Large Multimodal Models

Xingrui Wang¹, Wufei Ma¹, Tiezheng Zhang¹, Celso M de Melo², Jieneng Chen^{1†}, Alan Yuille^{1†}

¹Johns Hopkins University ²DEVCOM Army Research Laboratory

Abstract

Although large multimodal models (LMMs) have demonstrated remarkable capabilities in visual scene interpretation and reasoning, their capacity for complex and precise 3-dimensional spatial reasoning remains uncertain. Existing benchmarks focus predominantly on 2D spatial understanding and lack a framework to comprehensively evaluate 6D spatial reasoning across varying complexities. To address this limitation, we present Spatial457, a scalable and unbiased synthetic dataset designed with 4 key capability for spatial reasoning: multi-object recognition, 2D location, 3D location, and 3D orientation. We develop a cascading evaluation structure, constructing 7 question types across 5 difficulty levels that range from basic single object recognition to our new proposed complex 6D spatial reasoning tasks. We evaluated various large multimodal models (LMMs) on Spatial457, observing a general decline in performance as task complexity increases, particularly in 3D reasoning and 6D spatial tasks. To quantify these challenges, we introduce the Relative Performance Dropping Rate (RPDR), highlighting key weaknesses in 3D reasoning capabilities. Leveraging the unbiased attribute design of our dataset, we also uncover prediction biases across different attributes, with similar patterns observed in real-world image settings.¹ The code is released in <https://github.com/XingruiWang/Spatial457>.

1. Introduction

In recent years, large multimodal models (LMMs) have demonstrated remarkable reasoning capabilities, enabling them to interpret visual scenes from images and articulate their understanding through natural language [1, 3, 11, 26, 27, 38, 39, 48]. These advancements have been evaluated using a variety of benchmarks that primarily emphasize semantic features—such as object categories and appearances, spatial relationships within the 2D image plane [15,

17, 18] (e.g. left-right positioning from the camera view), and limited 2.5D features such as camera depth [40]. Tasks like object detection [13, 22], image captioning [21, 42, 47] and visual question answering [18, 50, 51] have significantly benefited from these evaluations, driving LMMs to new heights in visual understanding.

However, a critical gap remains in assessing the ability of LMMs to comprehend and reason about the full six degrees of freedom (6D) of objects in spatial environments, which includes both 3D positions and orientations [43, 46, 53]. Despite its significance in many real-world applications, such as robotics [2], autonomous navigation [14], and augmented reality [5], there are no available benchmarks for LMMs that evaluate 6D spatial reasoning capabilities.

Building a scalable diagnostic dataset for 6D vision tasks poses considerable challenges due to the extensive annotations required across vast image collections. One strategy is to develop 6D vision question-answering benchmarks on top of existing image datasets [4, 6, 7, 36]. However, real-world image collections often exhibit significant biases in 3D locations and poses, influenced by camera perspective limitations and the prevalence of common object orientations in everyday scenes. For instance, statistics from the SUN-RGBD [36] and nuScenes [7] datasets reveal that over 70% of objects cluster into a single predominant direction (see Fig. 2), limiting the diversity required for comprehensive 6D reasoning evaluation.

To overcome these challenges, we introduce an unbiased and scalable synthetic dataset rendered in a realistic style, specifically designed to diagnose the 6D spatial reasoning abilities of LMMs. To comprehensively evaluate diverse aspects of 6D spatial reasoning, we focus on four core capabilities: **multi-object recognition**, **2D locations**, **3D locations**, and **3D orientation**. Multi-object recognition serves as the foundation of scene understanding, while 2D location represents the simplest way to depict the spatial relationship between objects within the image plane [18, 19]. The 3D location factor extends understanding into three-dimensional space, capturing the depth and distance of objects, which is essential for tasks such as reasoning about occlusions [9, 10, 44]. The 3D orientation introduces object

¹Approved for public release: distribution is unlimited.

†: Equal advising.

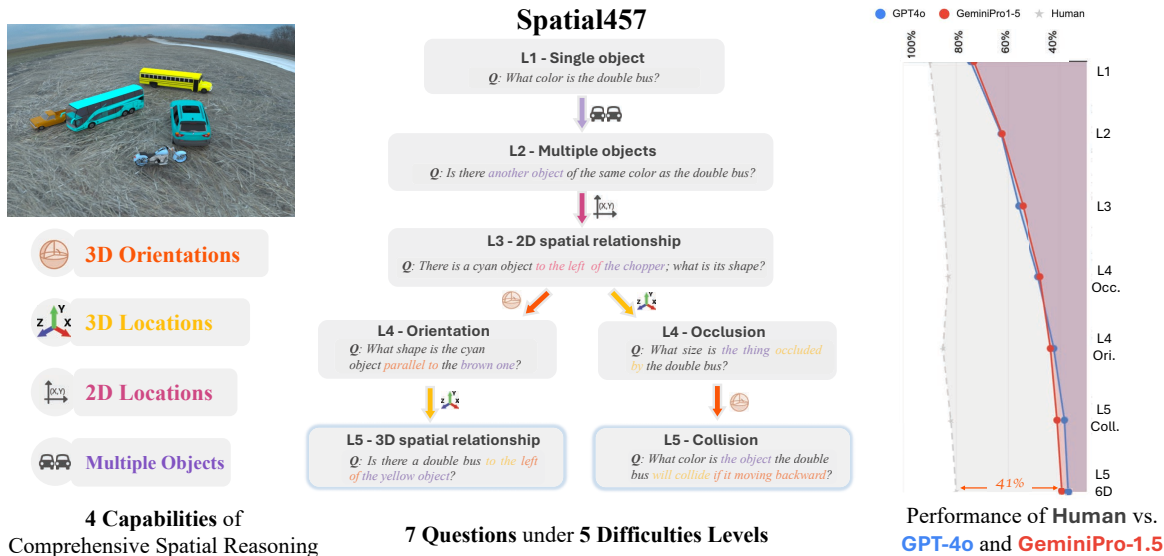


Figure 1. Overview of the Spatial457 benchmark. We define four core capabilities for spatial reasoning. By incorporating these capabilities step-by-step, the benchmark assesses models across five progressive difficulty levels and seven question types, ranging from single-object recognition to advanced 6D spatial relationships and collision prediction. Our evaluation demonstrates a significant performance drop for more advanced questions and highlights the gap between state-of-the-art (SoTA) models and human performance in complex spatial reasoning VQA tasks.




Figure 2. The distribution of 3D poses in representative 6D datasets (a&b) highlights an **imbalance** in object orientations. This inspired us to develop a **balanced** benchmark encompassing all degrees of 3D orientation.

rotation [44, 46], enabling models to reason about precise 3D poses of objects within a scene.

By progressively incorporating the four core capabilities, we establish five levels of difficulty, resulting in seven distinct question types. These questions span a spectrum of spatial reasoning complexity, ranging from basic semantic identification and 2D spatial relationships to advanced 3D positional understanding and orientation-based queries, culminating in comprehensive 6D spatial reasoning tasks.

Our evaluation measures performance across varied question types and examines how it is affected by different combinations of the four capabilities. This tiered design provides comprehensive benchmarks for spatial reasoning capabilities, highlighting each model’s strengths and areas of vulnerability. To the best of our knowledge, this work presents the first comprehensive benchmark designed to evaluate the 6D spatial reasoning capabilities of vision-

language models.

Additionally, our dataset’s unbiased attribute design allows us to analyze potential biases in model predictions. By systematically exploring model performance under different attributes, we identify tendencies or limitations that may affect real-world applications. Altogether, our benchmark not only provides a foundational tool for evaluating LMMs but also serves as a guide for developing future benchmarks that support models with advanced 6D spatial reasoning.

Our main contributions are threefold:

- 1. An unbiased and scalable synthetic dataset:** We present a realistic, unbiased synthetic dataset designed to assess 6D spatial reasoning in LMMs. We define four core capabilities and propose five difficulty levels for spatial reasoning VQA with seven question types.
- 2. Comprehensive benchmarking of 6D capabilities:** We benchmark various LMMs across all difficulty levels. The performance variations across settings reveal potential limitations or strengths, providing valuable insights into which components should be prioritized to improve the spatial reasoning of LMMs.
- 3. Bias analysis in model predictions:** Leveraging the unbiased nature of our dataset, we conduct an in-depth bias analysis to explore how LMM predictions may be unbalanced across different attributes, especially 3D-related ones. This reveals potential weaknesses in model predictions, offering insights into limitations that may also appear in real-world scenarios.

2. Related Work

2.1. Synthetic Datasets in VQA

The use of synthetic data in visual question answering (VQA) has gained traction as a means to circumvent the limitations and biases inherent in real-world datasets. Synthetic environments [16, 30, 34] allow for controlled manipulation of scene elements, enabling researchers to probe specific reasoning abilities in VQA models without the confounding factors often present in natural imagery. Previous works [19, 24, 25] have utilized synthetic environments to evaluate multimodal models across a range of visual reasoning tasks, typically focusing on 2D spatial relationships and object features. For instance, CLEVR [19], a synthetic dataset widely used in the VQA community, evaluates compositional reasoning by introducing artificially generated scenes with diverse object arrangements. Super-CLEVR-3D [45] represents a significant step forward in evaluating 3D locations, it lacks an assessment of 6D spatial relationships that consider both 3D location and 3D orientation. Our work extends this tradition by creating a synthetic dataset tailored for 6D spatial reasoning, filling a gap in current synthetic datasets that do not address the complete spatial understanding of objects in 3D space.

2.2. LMM Benchmarks

Existing benchmarks [15, 18, 35, 51] for large multimodal models (LMMs) predominantly assess their performance on tasks involving 2D visual question answering and multimodal reasoning. For instance, datasets like VQAv2 [15] and GQA [18] focus on answering questions based on visual content. However, these benchmarks leave a gap in evaluating models' capabilities in spatial contexts. Recent benchmarks, such as CV-Bench [40], assess spatial understanding, including depth ordering for 2.5D evaluation. More recent work [9, 10, 32] introduces questions regarding either aspect of 3D locations or rotations. Our proposed benchmark goes beyond traditional 2D, 2.5D, and 3D location or rotation assessments by introducing tasks that challenge LMMs to understand 3D positions and orientations together. We also evaluate models in a cascading manner across five difficulty levels and summarize their performance in four reasoning abilities, along with their output biases, providing a deeper analysis of the weaknesses in spatial reasoning within multimodal AI.

3. Spatial457: Comprehensive Spatial Reasoning Benchmark

To comprehensively evaluate the spatial reasoning capabilities of vision-language models, we introduce a synthetic dataset named **Spatial457**. We define four core capabilities for comprehensive spatial reasoning (Sec. 3.1) and generate scenes with full 6D annotations (Sec. 3.2). The full combination of these capabilities form the new 6D spatial

relationship VQA tasks (Sec. 3.3).

To thoroughly assess model capabilities, we define five levels of difficulty in vision question answering (VQA) tasks, converting seven question subsets that range from single-object recognition to the newly proposed 6D spatial reasoning tasks (Sec. 3.4). This structured progression in question difficulty is achieved by sequentially introducing new spatial reasoning capabilities, as illustrated in Fig. 1.

3.1. Four Capabilities for Spatial Reasoning

We define 4 capabilities for comprehensive spatial understanding in a 3D scene: *multiple objects*, *2D locations*, *3D locations*, and *3D orientations*.

(1) Multiple objects. Objects are fundamental to defining scene structure [52]. Large vision-language models show strong results in aligning semantic features with text [23, 31], handling multiple objects from scene becomes the first toward scene understanding. Models must recognize individual attributes and interpret relationships among them. Our dataset evaluates multi-object reasoning, such as comparison and counting. For example, in Fig. 1, the image shows two cyan objects, requiring the model to verify color similarity and count cyan objects accurately.

(2) 2D locations. 2D spatial relationships are essential in prior image-based VQA tasks [18, 19]. In traditional multimodal models [20, 28, 49], these relationships can be learned through object detection training [8, 33]. However, in large-scale vision-language models, explicit location information is often absent from the latent representations. Following Li et al. [24], we examine four spatial relationships (*left*, *right*, *front*, *back*) between objects from 2D camera views as the first tasks of spatial reasoning. Questions use these relations to filter targets, such as "What is the shape of the object to the left of the chopper?" (see Fig. 1).

(3) 3D locations. Beyond 2D relationships, 3D locations can incorporate depth, capturing spatial relationships that 2D alone cannot represent. This is crucial for tasks involving occlusions and complex spatial arrangements. Our dataset assesses the 3D location understanding abilities of model through questions about occlusions query as in [45]. For example, as shown in Fig. 1, questions may ask about an object's position relative to others in 3D space or whether one object is occluded, challenging models to judge distances and spatial positions accurately.

(4) 3D orientations. The 3D orientation of objects introduces an additional layer of complexity by incorporating rotational attributes. In current computer vision tasks, 3D orientation is often derived by pose estimation methods. However, large-scale vision-language models generally lack explicit 3D orientation information in their latent representations, potentially making it challenging to accurately infer the object's orientations. In our dataset, we eval-

uate models’ abilities to understand 3D orientations by focusing on object poses, including their facing direction and alignments with others. For example, as depicted in Fig. 1, the model may be asked to identify whether a cyan object is facing left, oriented parallel to a brown object, or positioned in a specific spatial relationship — i.e. left, right, front, or behind— from the perspective of the objects themselves.

3.2. Scene Rendering

We build upon the generation pipeline from prior VQA datasets [19, 24] and uniformly select 21 object models across 5 categories to construct the dataset as Li et al. [24]. To improve the realism of our rendered dataset, we use real image style environment maps and refined object textures. For benchmarking new 6D spatial understanding ability, we add new 6D spatial relationship annotations among objects.

Background. To create realistic scene backgrounds, we use HDRI images² with environment maps and natural lightings. We selected 224 different outdoor scenes with normal brightness levels without high light contrast by their metadata to maintain visual consistency across images.

Objects. We reconstruct 3D mesh models [24] with realistic textures, introducing more variation and realistic appearances for objects. Similar to Li et al. [24], each object in our dataset is defined by distinct attributes and precise spatial locations. The original attributes include shape, color, and size. For benchmarking 6D spatial reasoning, each object also includes a vector specifying its 2D location (x, y), 3D world coordinates (X, Y, Z), and pose direction derived from its 3D orientation [45].

Relationships. As shown in Fig. 3 (a.1 to a.3), we define three types of spatial relationships: 1. *2D Spatial Relationships*: These describe the relative positioning of objects within the 2D camera view, using terms such as *left*, *right*, *front*, and *behind*. 2. *6D Spatial Relationships*: Extending spatial descriptions into 3D space, these relationships incorporate both the 3D position and orientation of objects to define relative positions from a target object’s perspective, including *left*, *right*, *front*, and *behind*. 3. *Collision Relationships*: This relationship describes potential collisions if the target object moves forward or backward. Collision prediction involves assessing the orientation and 3D location of target objects relative to others to determine possible collisions.

3.3. New Questions for 6D Spatial Reasoning

Based on the four capabilities defined in Sec. 3.1, we construct new question types for **6D spatial reasoning** and **collision prediction**. These questions require an understanding of both 3D location and 3D orientation across multiple objects to generate accurate answers. This level of complexity represents the L5 difficulty level shown in Fig. 1.

²<https://polyhaven.com/hdri>

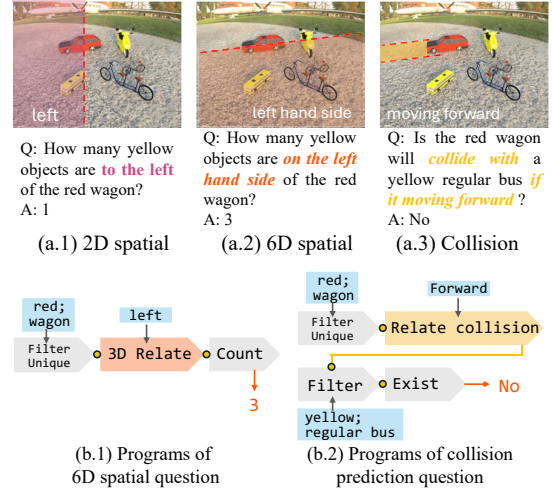


Figure 3. (a.1-3) Example of the 2D spatial, 3D spatial and collision with images and questions. (b.1-2) shows the new operation programs we generated for the new 6D spatial reasoning questions.

The 6D spatial relationship questions challenge the model to understand the spatial relationship involving both 3D position and orientation from a target object’s perspective (shown in Fig. 3 (a.2)). Compared with prior VQA benchmarks focusing on 2D [19, 24], we introduce new 3D spatial reasoning operations, *3D Relate*(\cdot), to determine relative positions like *left*, *right*, *front*, and *behind* within a 3D coordinate space (see Fig. 3 (b.1)). For example, a question might ask, “How many objects are to the right of the red wagon?” requiring the model to filter and count objects based on their orientation and relative 3D position, which would yield different answers compared to similar 2D left and right questions (Fig. 3 (a.1)). This task is particularly challenging as it requires models to integrate 3D spatial relationships while managing occlusions and depth perception—capabilities not typically covered in simpler 2D spatial tasks.

The collision prediction questions, as shown in Fig. 3 (a.3), extend 6D spatial reasoning by introducing motion-based inference. These questions assess the model’s ability to anticipate future states and potential interactions (collisions) based on accurate estimation of 3D location and direction, denoted as *forward* or *backward*. For instance, a question might ask, “Will the red wagon collide with a yellow bus if it moves forward?”. This requires the model to consider the orientation and direction of movement relative to surrounding objects to predict possible outcomes. Such tasks test the model’s advanced spatial reasoning about 3D orientation and position, and its capacity to predict accurate trajectories and directions within the future scene.

3.4. Questions under Other 4 Difficulty Levels

The L5 questions above represent the most challenging cases for 6D spatial reasoning, testing all four capabilities

in Sec. 3.1. However, the purpose of this benchmark is to comprehensively diagnose each factor in spatial reasoning. Thus, we design a structured roadmap to evaluate the model’s performance across different levels of spatial reasoning, assessing its capabilities in each specific factor.

By composing the four factors of comprehensive spatial reasoning, we define an additional four levels of question difficulty based on the generated scenes, encompassing a total of seven tasks in our benchmark. As the complexity of spatial reasoning increases with the introduction of additional factors, we organize these questions into a structured roadmap, as illustrated in Fig. 1. Although many questions are reimplementations from previous benchmarks [24, 45] on our new images, we aim to integrate these into a unified benchmark to provide deeper insights into the model’s spatial reasoning capabilities. The composition of the four capabilities at each level, along with comparisons to previous VQA benchmarks, is shown in Tab. 1.

Level 1 - single object questions [19, 24]. We begin with the questions with single objects. These questions are generated using a filtered set of templates from Li et al. [24]. We select five templates involving only one object in the question, such as “What color is the double bus?” in Fig. 1. This question type provides a base score for the tested vision-language model. The comparison with L1 helps to determine whether the decrease of performance are due to challenge of spatial reasoning or object recognition.

Level 2 - multiple objects questions [19, 24]. Level 2 questions extend Level 1 by involving multiple objects. These questions involves understand the relationship between objects and counting. For example, “Is there another object of the same color as the double bus?” as shown in Fig. 1, where the question templates follows Li et al. [24]. We evaluate the model’s capacity for reasoning in scenes with multiple entities as a critical step toward understanding spatial relationships in next levels.

Level 3 - 2D spatial questions [19, 24]. Level 3 questions further introduce 2D spatial relationships among multiple objects, requiring the model to understand the position of objects and their relationships from the 2D camera’s perspective. For example, “What shape is the object to the left of the chopper?” as depicted in Fig. 1. The 2D relationship is the first step of spatial reasoning.

Level 4 - 3D pose questions [45]. Level 4 questions extend 3D orientation reasoning beyond the 2D spatial relationships in Level 3. Beyond appearance attributes (*shape*, *color*, *size*) or 2D locations, these questions require understanding 3D orientation, including the rotation angle or facing direction of an object from the camera’s perspective. For example, “What shape is the object that is parallel to the brown object?” as shown in Fig. 1. 3D pose questions test the model’s ability to interpret scenes with 3D awareness, involving rotation and alignment with other objects.

Level 4 - occlusion questions [45]. In addition to 3D orientation, incorporating 3D location presents another challenge beyond the Level 3 questions on 2D spatial relationships. A unique aspect of 3D spatial reasoning from images is occlusion. The occlusion questions in Level 4 require the model to comprehend depth perception and the relative positions of objects in 3D space, determining which object is closer to the camera within overlapping areas. For example, an occlusion question might ask, “Is the red wagon occluded by the yellow bus?” as shown in Fig. 1. This level evaluates the model’s ability to reason about occlusions and depth perception, which are essential for understanding complex spatial relationships in 3D scenes.

Question Level	Multi-object	2D location	3D location	3D Orientation	VQA Benchmarks
L1					spatial-agnostic bench (e.g., TextVQA [35])
L2	✓				CLEVR-Math [25], CLEVR-Count [19]
L3	✓	✓			VQA [15], GQA [18], CLEVR [19], Super-CLEVR [24]
L4 (a)	✓	✓	✓		CV-Bench (DepthOrder) [40], Super-CLEVR-3D [45]
L4 (b)	✓	✓		✓	Super-CLEVR-3D [45]
L5	✓	✓	✓	✓	Spatial457

Table 1. Comparison of VQA Benchmarks by Factors and Question Levels. Compared with existing VQA benchmarks, Spatial457 includes all four spatial reasoning factors, providing a comprehensive evaluation of 6D spatial understanding.

3.5. Evaluation Protocols

To better evaluate the model’s spatial reasoning capabilities, we propose a structured evaluation protocol that measures the relative capability of models on each factor for spatial reasoning. We define this evaluation metric as the Relative Performance Dropping Rate (RPDR), which assesses the decline in performance after introducing a new factor in the questions. Based on the roadmap in Fig. 1, the RPDR for each factor is defined as follows:

$$RPDR_{\text{Multi-Obj.}} = Acc_{\text{Level 1}} \Rightarrow Acc_{\text{Level 2}}$$

$$RPDR_{\text{2D Loc.}} = Acc_{\text{Level 2}} \Rightarrow Acc_{\text{Level 3}}$$

$$RPDR_{\text{3D Loc.}} = \text{Avg} \left(\begin{array}{l} Acc_{\text{Level 3}} \Rightarrow Acc_{\text{Level 4 (Occ.)}}, \\ Acc_{\text{Level 4 (3D Pose)}} \Rightarrow Acc_{\text{Level 5 (6D)}} \end{array} \right)$$

$$RPDR_{\text{3D Ori.}} = \text{Avg} \left(\begin{array}{l} Acc_{\text{Level 3}} \Rightarrow Acc_{\text{Level 4 (3D Pose)}}, \\ Acc_{\text{Level 4 (Occ.)}} \Rightarrow Acc_{\text{Level 5 (Col.)}} \end{array} \right)$$

where Acc_{Level} denotes the model’s accuracy on the corresponding questions and $(Acc_{Q1} \Rightarrow Acc_{Q2}) = \frac{Acc_{Q2} - Acc_{Q1}}{Acc_{Q1}}$ denotes the relative performance dropping rate from $Q1$ to $Q2$. The results of RPDR will be discussed in Sec. 4.3.

4. Experiments

4.1. Experimental Setup

We generate 1,000 images to test the performance of large vision-language models. For the seven types of questions, we generate 1000 single-object questions (*L1-Single*), 5,000

Difficulty Level	Level 1	Level 2	Level 3	Level 4		Level 5	
Questions type	Single Object	Multi-Obj.	2D Spatial	Occlusion	3D Pose	Collisions	6D Spatial
Random	33.05	32.77	33.47	22.04	18.99	21.02	19.41
GPT-4o [29]	74.46	62.88	56.14	48.40	42.41	38.41	37.01
GeminiPro 1.5 [37]	73.26	62.54	54.49	47.65	43.67	41.19	39.36
Claude 3.5 Sonnet [3]	68.24	57.40	54.19	30.84	38.40	35.34	33.48
Qwen2-VL-7B-Inst. [39]	71.96	61.44	55.34	27.87	34.29	36.58	33.75
InternVL2-8B [38]	58.11	58.76	57.40	33.25	32.32	34.30	34.30
LLaVA-v1.5-7B [26]	44.87	44.72	42.20	24.34	24.55	23.63	23.86
LLaVA-NeXT-vicuna-7B [27]	50.72	49.47	46.01	29.68	29.35	31.95	31.95
LLaVA-NeXT-llama3-8B [41]	52.15	49.73	45.92	30.31	29.77	32.12	32.12
PO3D-VQA [45]	86.46	82.55	80.64	70.49	81.40	68.12	71.06
Human	89.97	86.83	84.95	82.76	84.95	81.82	79.94

Table 2. Performance comparison across all 7 question types under 5 difficulty levels. As task complexity increases from single-object questions to 6D spatial relationships, performance generally decreases in most cases, highlighting the challenges in handling multi-object interactions, 3D orientations, and predictive spatial reasoning. The highest score among all models for each question type is marked in red, and the second-highest scores are underlined. The best performance among open-source models is marked in blue. PO3D-VQA is the state-of-the-art (SOTA) 3D-aware neural symbolic reasoning model as proposed in [45], and is listed as the potential upper bound of evaluated methods. The comparison with human performance also shows an increasing gap for LLMs as task difficulty increases.

multiple-object questions (*L2-Multi-obj.*), 4,995 2D spatial relationship questions (*L3-2D-Spatial*), 3,534 occlusion questions (*L4-Occlusion*), 4,555 3D pose questions (*L4-3D-Pose*), 4,357 collision questions (*L5-Collision*), and 4,995 6D spatial relationship questions (*L5-6D-Spatial*).

We evaluate two series of large vision-language models on Spatial457: (1) API-based large language models, including GPT-4o [29], Gemini-1.5-Pro [37], and Claude-3.5-Sonnet [3]; and (2) open-source large vision-language models, including Qwen2-VL-7B-Instruct [39], InternVL2-8B [38], llava-v1.5-7b [26], llava-next-vicuna-7b [27], and llava-next-llama3-8b [41]. The evaluation follows the format of Vlmekit [12]. The details of the model implementation including prompts are listed in Appendix.

4.2. Evaluation Results

In Tab. 2, we present the performance of various large vision-language models across different levels of difficulty on Spatial457. The levels range from basic single-object tasks to complex 6D spatial reasoning, with increasing difficulty as more factors are introduced. We compare both API-based models (e.g., GPT-4o, GeminiPro1.5, Claude3-5V) and open-source models (e.g., Qwen2-VL-7B-Instruct, InternVL2-8B, LLaVA series) to assess their spatial reasoning capabilities.

Comparison across difficulty levels. We evaluate model performance across different levels of difficulty in Spatial457, ranging from basic single-object tasks to complex 6D spatial reasoning. As shown in Tab. 2, performance generally decreases as the complexity of spatial relationships increases, highlighting the challenges models face

when handling multi-object interactions, 3D orientations, and predictive spatial reasoning tasks.

Comparison between API models and open-source models. From Tab. 2, we observe that API models generally outperform open-source models across all difficulty levels. In 2D spatial relationship questions, InternVL2-8B achieves slightly higher results than the API models, while GPT-4o is only 1% lower. Qwen2-VL-7B-Instruct achieves comparable results, especially in simpler tasks like single-object and multiple-object reasoning. This analysis highlights the strengths of API models in complex spatial reasoning tasks.

Comparison among API models. API-based models, *i.e.* GPT-4o, GeminiPro1.5, and Claude3-5V, demonstrate strengths across all difficulty levels. For instance, GPT-4o excels in 2D and 3D spatial tasks, while GeminiPro1.5 shows robust performance in collision prediction tasks. This comparison provides insights into the unique capabilities and limitations of each API model in handling spatial reasoning across different levels of complexity.

4.3. RPDR Analysis for 4 Capabilities

As outlined in Sec. 3.5, we use the Relative Performance Dropping Rate (RPDR) to analyze model performance across the four core factors of Spatial457: multi-object reasoning, 2D location understanding, 3D location comprehension, and 3D orientation interpretation. The RPDR analysis provides a structured approach to quantify the decline in accuracy with the introduction of each new factor, as detailed in Tab. 3. Our findings reveal that all models exhibit weaknesses in 3D reasoning. Specifically, GPT-4o and GeminiPro-1.5 perform the worst in 3D orientation tasks,

while the other two struggle most with 3D location.

Model	Multi.	2D Loc.	3D Ori.	3D Loc.
GPT-4o [29]	84.45	89.27	77.46	86.74
GeminiPro 1.5 [37]	85.37	87.13	83.28	88.79
Claude 3.5 S. [3]	84.12	94.41	85.43	72.05
Qwen2-VL-7B [39]	85.39	90.06	80.99	74.40

Table 3. RPDR (%) analysis on the four spatial reasoning capabilities, where higher values indicate better performance. For each row, we highlight the lowest RPDR score among all factors.

4.4. Bias Analysis

In this section, we analyze prediction biases in large language models (LLMs) regarding their understanding of key attributes, specifically shape, color, size, and pose. Leveraging the unbiased generation of dataset images and questions in Spatial457, we examine how these models may demonstrate attribute-specific biases in their predictions.

We first compute the confusion matrices and distribution statistics for both ground truth and model predictions across each attribute (shape, color, size, and pose). For example, the confusion matrix plots for color and pose attributes of GPT-4o and GeminiPro1.5, together with the distribution of prediction and ground truth label are shown in Fig. 4. We find both GPT-4o and GeminiPro1.5 exhibit bias in both color and pose predictions. These results highlight areas for improvement to reduce bias and enhance predictive reliability across models.

To quantify prediction variability, we calculate the Coefficient of Variation (CV) for each attribute as follows:

$$CV = \frac{\text{Std}(p_1, \dots, p_k)}{\text{Mean}(p_1, \dots, p_k)}, \quad (1)$$

where $\text{Std}(\cdot)$ and $\text{Mean}(\cdot)$ are the standard deviation and mean of the probability distribution p_i for each attribute label i (e.g., "left," "right" for pose).

Tab. 4 shows the CV values for shape, color, size, and pose attributes across different models. Lower CV values indicate more consistent predictions with less bias, which reveals significant biases in model predictions across attributes. Higher CV values, especially in color and pose, indicate consistency challenges.

4.5. Comparison to the 3D Symbolic Model

As neural symbolic models achieve state-of-the-art (SOTA) performance on many prior synthetic datasets, we implement the 3D-aware neural symbolic model PO3D-VQA [45] on our new image domain and all levels of questions. We train PO3D-VQA on a separate set of 20,000 images and questions and report its performance on the same test dataset in Tab. 2. A direct comparison between the 3D neural symbolic model and large language models is not entirely fair, as each module of the neural symbolic model (6D pose estimation, attribute classification, and language reasoning) is trained independently as a sub-task. However, we




Figure 4. The distribution of color and pose attributes in the L4-3D-Pose task for GPT-4o and Gemini-Pro1.5. Although the ground truth labels are well balanced, the predicted colors and poses for both models are imbalanced, with “yellow” being the most frequent color and “front” the most frequent pose.

Models	Avg. ↓	Shape	Color	Size	Pose
GPT4o [29]	0.485	0.623	0.521	0.043	0.752
GeminiPro1.5 [37]	0.498	0.567	0.366	0.297	0.760
Claude3.5V-S [3]	0.720	0.777	0.465	0.574	1.063
Qwen2-VL-7B [39]	0.545	0.868	0.558	0.067	0.688
Human	0.093	0.110	0.082	0.055	0.126

Table 4. Coefficient of Variation (CV) for predictions on each attribute in the L4-3D-Pose task. We evaluate three API models and one open-source model, calculating the CV across four attributes (shape, color, size, 3D pose) based on the distribution of predicted labels. Lower CV values indicate less bias. The first column shows the average CV across attributes, and GPT4-o has the lowest bias.

include this model to illustrate the upper bound that a data-driven machine learning model can achieve on our dataset.

4.6. Qualitative Results

Fig. 5 presents an example of GPT-4o’s performance across various spatial reasoning tasks on the same image. GPT-4o correctly answers multi-object, 2D spatial, and occlusion tasks but fails on 3D pose and 6D spatial questions, indicating challenges in 3D orientation understanding in this case. As we prompt the models to output the answer as well as the reasoning steps, the models correctly detect the existence of objects (e.g. in Fig. 5 the yellow minivan in L4-Pose or the fighter plane and road bike in L5 questions, but has the wrong understand to the orientations or 3D locations.

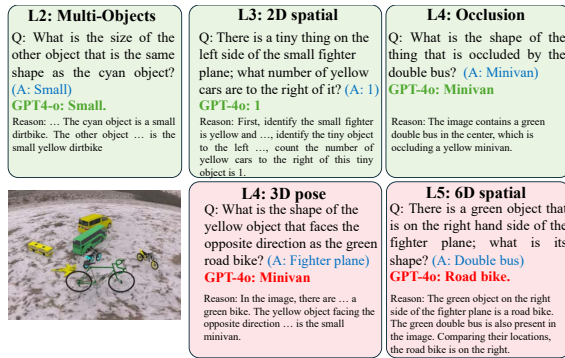


Figure 5. Example of GPT-4o across different spatial reasoning tasks on the same image with the ground truth answer marked in blue. GPT-4o answers multi-object, 2D spatial, and occlusion tasks correctly (in the green block) but fails on 3D pose and 6D spatial questions (the red block).

5. Extension to Real Dataset

We extend the setting of 3D questions to real-world image data. Due to limitations in fully annotated attributes such as category, color, 3D orientation, position, and occlusion, we focus on questions from the *L4-Pose* category. Using the scene structure from SUN-RGBD [36], we construct a new subset, *L4-Pose-Real*, which involves only the category and pose attributes. To reduce ambiguity in object localization, we add a red bounding box around target objects, ensuring clarity in identifying categories within the scene.

We provide the question along with the image as input to the LLM. In this example in Fig. 6, GPT-4o provides the correct answer, while Gemini-Pro 1.5 is incorrect. However, the reasoning process reveals that both models rely on common sense knowledge or 2D visibility cues to interpret the 3D orientation.

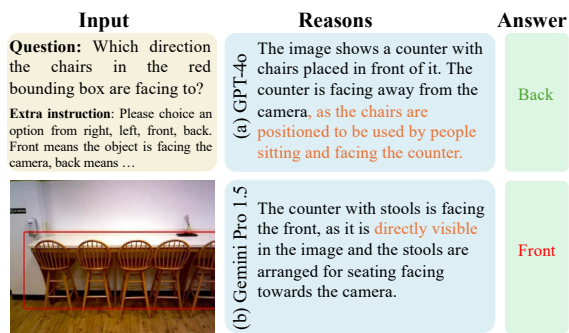


Figure 6. Example of real-world 3D pose reasoning tasks *L4-Pose-Real*. Here we shows the reasoning and answers provided by LLMs where GPT-4o is correct and Gemini-Pro 1.5 is wrong. But the reasoning process reveals these two LLMs are both reasoning 3D orientation from the command sense or 2D visibility.

As shown in Tab. 5, Gemini Pro 1.5 achieves the highest accuracy (40.50%) in *L4-Pose-Real*, and the gap between the three API-based models and the open-source

model Qwen2-VL-7B-Instruct (26.50%) is larger. However, all these models perform significantly lower than humans, which have accuracy 92.64%. We further analyze the bias of predictions in this real-world question set. Among the models, Claude3-5V Sonnet exhibits the lowest CV (0.374), indicating relatively stable predictions, while Qwen2-VL-7B-Instruct has the highest CV (0.934). The distribution of ground truth and predicted pose and the confusion matrix are shown in Fig. 7.

Models	Accuracy (%) \uparrow	CV \downarrow
GPT-4o [29]	38.75	0.690
GeminiPro1-5 [37]	40.50	0.762
Claude3-5V Sonnet [3]	38.75	0.374
Qwen2-VL-7B-Instruct [39]	26.50	0.934
Human	92.64	0.197

Table 5. Accuracy and Coefficient of Variation (CV) for *L4-Pose-Real* task: Comparison of accuracy and CV for GPT-4o, GeminiPro1.5, Claude3-5V, and Qwen2-VL-7B-Instruct on real-world 3D pose tasks. Human performance is included as a reference.

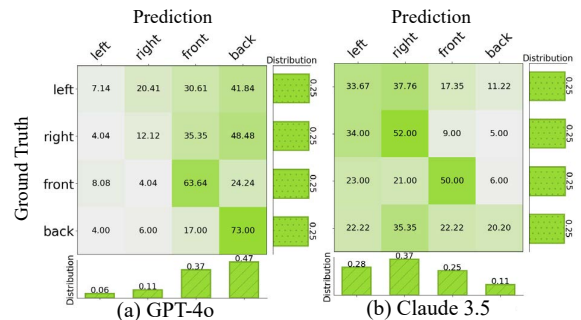


Figure 7. Distribution of 3D poses for in *L4-Pose-Real* task.

6. Conclusion

In this work, we introduce Spatial457, a scalable and unbiased benchmark designed to comprehensively assess 6D spatial reasoning capabilities in LMMs. Spatial457 extends traditional 2D and 3D VQA benchmarks by incorporating complex 6D spatial reasoning tasks, structured across five levels of difficulty. Through evaluations, we observe that current LMMs demonstrate strong performance in basic object recognition and 2D spatial relationships, but struggle significantly with 3D spatial understanding and collision prediction tasks. Additionally, our analysis of prediction biases uncovers attribute-specific inconsistencies that could affect model reliability in real-world applications. **Limitations.** To real images settings, current evaluation limites to L3 pose questions. We are actively working on building the additional required annotations for these real images to expand the domain of the proposed benchmark.

Acknowledgement

This work is supported by ONR with award N00014-23-1-2641, from ARL Army Research Laboratory with award W911NF2320008.

References

- [1] Alibaba DAMO Academy. Gemini: Vision-language model by alibaba. *Alibaba DAMO*, 2023.
- [2] Boshi An, Yiran Geng, Kai Chen, Xiaoqi Li, Qi Dou, and Hao Dong. Rgbmanip: Monocular image-based robotic manipulation through active object pose estimation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7748–7755. IEEE, 2024.
- [3] Anthropic. Claude: An anthropic language model. *Anthropic*, 2023.
- [4] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *arXiv preprint arXiv:2111.08897*, 2021.
- [5] Mark Billinghurst, Adrian Clark, and Gun Lee. A survey of augmented reality. *Foundations and Trends in Human-Computer Interaction*, 8(2-3):73–272, 2015.
- [6] Garrick Brazil, Abhinav Kumar, Julian Straub, Nikhila Ravi, Justin Johnson, and Georgia Gkioxari. Omni3d: A large benchmark and model for 3d object detection in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13154–13164, 2023.
- [7] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multi-modal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020.
- [8] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020.
- [9] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14455–14465, 2024.
- [10] An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. Spatialrgpt: Grounded spatial reasoning in vision language model. *arXiv preprint arXiv:2406.01584*, 2024.
- [11] Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality, march 2023. URL <https://lmsys.org/blog/2023-03-30-vicuna>, 3(5), 2023.
- [12] Haodong Duan, Junming Yang, Yuxuan Qiao, Xinyu Fang, Lin Chen, Yuan Liu, Xiaoyi Dong, Yuhang Zang, Pan Zhang, Jiaqi Wang, et al. Vlmevalkit: An open-source toolkit for evaluating large multi-modality models. In *Proceedings of the 32nd ACM international conference on multimedia*, pages 11198–11201, 2024.
- [13] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.
- [14] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3354–3361, 2012.
- [15] 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.
- [16] Roland Hess. *Blender foundations: The essential guide to learning blender 2.5*. Routledge, 2013.
- [17] Yushi Hu, Hang Hua, Zhengyuan Yang, Weijia Shi, Noah A Smith, and Jiebo Luo. Promptcap: Prompt-guided task-aware image captioning. *arXiv preprint arXiv:2211.09699*, 2022.
- [18] 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.
- [19] Justin Johnson, Bharath Hariharan, Laurens Van Der Maaten, Li Fei-Fei, C Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2901–2910, 2017.
- [20] Aishwarya Kamath, Mannat Singh, Yann LeCun, Gabriel Synnaeve, Ishan Misra, and Nicolas Carion. Mdetrm: Modulated detection for end-to-end multi-modal understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1780–1790, 2021.
- [21] Dong-Jin Kim, Jinsoo Choi, Tae-Hyun Oh, and In So Kweon. Dense relational captioning: Triple-stream networks for relationship-based captioning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6271–6280, 2019.
- [22] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ilya Krasin, Jordi Pont-Tuset, Gleb Muryshkin, Sergey Popov, Marwa Elhachem, Pedro de Carvalho, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *International Journal of Computer Vision*, 128(7):1956–1981, 2020.
- [23] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023.

- [24] Zhuowan Li, Xingrui Wang, Elias Stengel-Eskin, Adam Kortylewski, Wufei Ma, Benjamin Van Durme, and Alan L Yuille. Super-clevr: A virtual benchmark to diagnose domain robustness in visual reasoning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14963–14973, 2023.
- [25] Adam Dahlgren Lindström and Savitha Sam Abraham. Clevr-math: A dataset for compositional language, visual and mathematical reasoning. *arXiv preprint arXiv:2208.05358*, 2022.
- [26] Haotian Liu et al. Llava: Large language and vision assistant. *LLaVA Project*, 2023.
- [27] Haotian Liu et al. Llava-next: Extended large language and vision assistant. *LLaVA Project*, 2023.
- [28] Jiayuan Mao, Chuang Gan, Pushmeet Kohli, Joshua B Tenenbaum, and Jiajun Wu. The neuro-symbolic concept learner: Interpreting scenes, words, and sentences from natural supervision. *arXiv preprint arXiv:1904.12584*, 2019.
- [29] OpenAI. Gpt-4 technical report. *OpenAI*, 2023.
- [30] Weichao Qiu and Alan Yuille. Unrealcv: Connecting computer vision to unreal engine. In *Computer Vision—ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part III 14*, pages 909–916. Springer, 2016.
- [31] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [32] Arijit Ray, Jiafei Duan, Reuben Tan, Dina Bashkirova, Rose Hendrix, Kiana Ehsani, Aniruddha Kembhavi, Bryan A Plummer, Ranjay Krishna, Kuo-Hao Zeng, et al. Sat: Spatial aptitude training for multimodal language models. *arXiv preprint arXiv:2412.07755*, 2024.
- [33] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *IEEE transactions on pattern analysis and machine intelligence*, 39(6):1137–1149, 2016.
- [34] Andrew Sanders. *An introduction to Unreal engine 4*. AK Peters/CRC Press, 2016.
- [35] 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.
- [36] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015.
- [37] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv preprint arXiv:2403.05530*, 2024.
- [38] Intern Team. Internvl 2.0: A vision-language model. *Intern Foundation*, 2023.
- [39] Qwen Team. Qwen-vl: A large multimodal model. *Alibaba Group*, 2023.
- [40] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, et al. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *arXiv preprint arXiv:2406.16860*, 2024.
- [41] Hugo Touvron et al. Llama: Open and efficient foundation language models. *Meta AI*, 2023.
- [42] Oriol Vinyals, Alexander Toshev, Samy Bengio, and Dumitru Erhan. Show and tell: A neural image caption generator. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3156–3164, 2015.
- [43] He Wang, Srinath Sridhar, Jingwei Huang, Julien Valentin, Shuran Song, and Leonidas J Guibas. Normalized object coordinate space for category-level 6d object pose and size estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2642–2651, 2019.
- [44] Xingrui Wang, Wufei Ma, Zhuowan Li, Adam Kortylewski, and Alan Yuille. 3d-aware visual question answering about parts, poses and occlusions. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023.
- [45] Xingrui Wang, Wufei Ma, Zhuowan Li, Adam Kortylewski, and Alan L Yuille. 3d-aware visual question answering about parts, poses and occlusions. *Advances in Neural Information Processing Systems*, 36, 2024.
- [46] Yu Xiang, Roozbeh Mottaghi, and Silvio Savarese. Beyond pascal: A benchmark for 3d object detection in the wild. In *IEEE winter conference on applications of computer vision*, pages 75–82. IEEE, 2014.
- [47] Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pages 2048–2057. PMLR, 2015.
- [48] Lin Xu, Yilin Zhao, Daquan Zhou, Zhijie Lin, See Kiong Ng, and Jiashi Feng. Pllava: Parameter-free llava extension from images to videos for video dense captioning. *arXiv preprint arXiv:2404.16994*, 2024.
- [49] Kexin Yi, Jiajun Wu, Chuang Gan, Antonio Torralba, Pushmeet Kohli, and Josh Tenenbaum. Neural-symbolic vqa: Disentangling reasoning from vision and language understanding. *Advances in neural information processing systems*, 31, 2018.
- [50] Kexin Yi, Chuang Gan, Yunzhu Li, Pushmeet Kohli, Jiajun Wu, Antonio Torralba, and Joshua B Tenenbaum. Clevrer: Collision events for video representation and reasoning. *arXiv preprint arXiv:1910.01442*, 2019.
- [51] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567, 2024.

- [52] Alan Yuille and Daniel Kersten. Vision as bayesian inference: analysis by synthesis? *Trends in cognitive sciences*, 10(7):301–308, 2006.
- [53] Jiyao Zhang, Weiyao Huang, Bo Peng, Mingdong Wu, Fei Hu, Zijian Chen, Bo Zhao, and Hao Dong. Omni6dpose: A benchmark and model for universal 6d object pose estimation and tracking. In *European Conference on Computer Vision*, pages 199–216. Springer, 2024.