

SPHERE: Unveiling Spatial Blind Spots in Vision-Language Models Through Hierarchical Evaluation

Wenyu Zhang¹, Wei En Ng², Lixin Ma^{3*}, Yuwen Wang^{2*}, Junqi Zhao^{4*}, Allison Koencke⁵
Boyang Li⁴, Lu Wang¹

¹Institute for Infocomm Research (I²R), Agency for Science, Technology and Research (A*STAR)

²National University of Singapore (NUS), ³Tongji University

⁴Nanyang Technological University (NTU), ⁵Cornell University

Abstract

Current vision-language models may grasp basic spatial cues and simple directions (e.g. left, right, front, back), but struggle with the multi-dimensional spatial reasoning necessary for human-like understanding and real-world applications. To address this gap, we develop SPHERE (Spatial Perception and Hierarchical Evaluation of REasoning), a hierarchical evaluation framework supported by a new human-annotated dataset. SPHERE systematically probes models across increasing levels of complexity, from fundamental skills to multi-skill integration and high-level reasoning that combines spatial, visual, and logical understanding. Benchmark evaluation of state-of-the-art models reveals significant deficiencies, especially in reasoning about distance and proximity, understanding both egocentric and allocentric perspectives, and applying spatial logic in physical contexts. These findings expose critical blind spots in existing models and underscore the need for more advanced spatial reasoning techniques, driving the development of vision-language models that align more closely with human spatial cognition. The SPHERE benchmark is available at [this repository](#).

1 Introduction

Spatial perception and reasoning play an essential role in how vision-language models (VLMs) understand complex, context-rich environments (Du et al., 2024; Liu et al., 2023a; Zeng et al., 2024). These capabilities are crucial for the deployment of VLMs in physical-world applications such as robotics, embodied AI, or human-assistive systems (Nair et al., 2022; Brohan et al., 2023; Stone et al., 2023; Leal et al., 2023).

Current research efforts to enhance spatial sensing in VLMs have incorporated features such as

* Contributed equally to this work; authors are listed in alphabetical order.

Model	Single-skill	Multi-skill	Reasoning	Overall
Human baseline	95.4	92.5	89.0	93.0
Random baseline*	50.0	44.3	50.0	49.1
Best model score	78.2	58.6	64.7	67.9
Rank 1	Gemini 2.0 Flash	GPT-4o	GPT-4o	GPT-4o
Rank 2	GPT-4o	Qwen2.5-VL	SpatialBot-RGB	Qwen2.5-VL
Rank 3	Qwen2.5-VL	LLaVA-OneVision	Llama-3.2-Vision	Gemini 2.0 Flash

Table 1: VLM rankings on SPHERE tasks. * Random baseline scores excludes open-ended counting tasks.

depth information, object bounding boxes, and simple spatial directions (e.g. left/right, inside/outside, front/back) into the models (Chen et al., 2024a; Zhao et al., 2023; AI, 2024a; Cai et al., 2024; Cheng et al., 2024). However, the enhancements often focus on isolated, simplistic spatial cues, and therefore limit the models’ ability to handle the complexities of real-world scenarios. Practical applications demand more sophisticated spatial perception and reasoning capabilities, akin to human-like understanding. Figure 1 demonstrates failure cases of GPT-4o. In question 2, GPT-4o has difficulty counting while interpreting object proximity. In question 3, GPT-4o understands that the laptop is at the edge of the desk, but fails to reason that the laptop cannot be moved further back.

In this work, we propose a new benchmarking framework called SPHERE (Spatial Perception and Hierarchical Evaluation of REasoning), to systematically disentangle and measure spatial skills, in order to pinpoint model strengths and weaknesses and guide future research. We are motivated by prior research findings (West et al., 2023; Tiong et al., 2024) which reveal fundamental differences between machine and human intelligence. While LLMs excel in generating coherent and contextually relevant responses for advanced topics, they often exhibit surprising errors in basic understanding (West et al., 2023). We cannot assume good performance in one task to always transfer to a

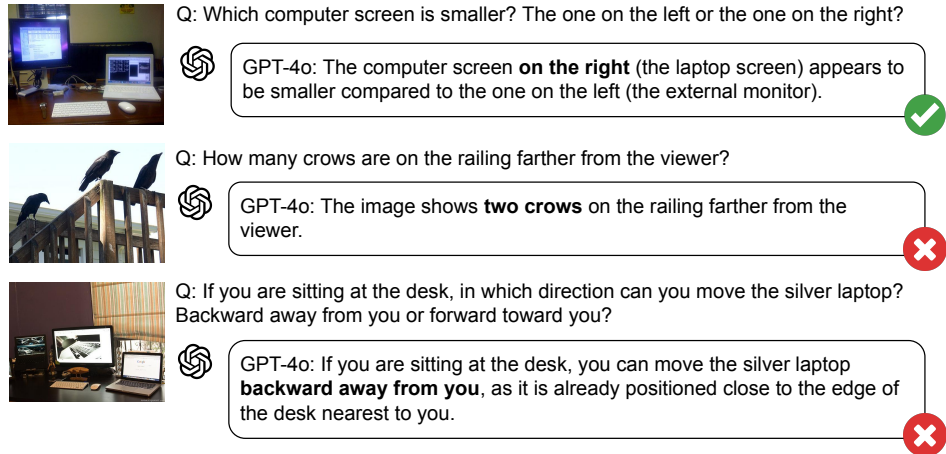


Figure 1: State-of-the-art models such as GPT-4o still has difficulty on questions that require multiple spatial, visual and reasoning skills. GPT-4o itself bolded the words in the second and third examples.

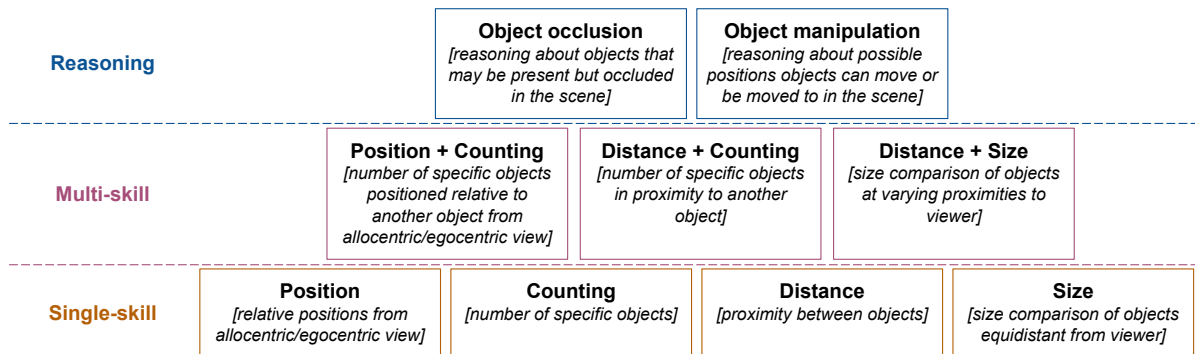


Figure 2: The proposed SPHERE framework evaluates vision-language models on a hierarchy of tasks, advancing from single-skill tasks to multi-skill tasks, and ultimately to complex reasoning tasks that require the integration of multiple spatial and visual cues with logical reasoning abilities.

seemingly correlated or even more basic task.

SPHERE, illustrated in Figure 2, is a hierarchical evaluation framework that covers a spectrum of basic and advanced tasks. We begin with single-skill tasks (i.e. object positions, distances, sizes and counts), then expand to multi-skill tasks that require multiple spatial and visual skills (e.g. counting objects at specific positions, counting objects at certain distances, etc.). We also design two novel reasoning tasks (i.e. reasoning about occluded objects and object manipulation) that require advanced understanding of a scene as a 3-dimensional environment and the objects as physical entities that can be interacted with. We measure reasoning ability as the ability to infer about scenes or objects not explicitly observed but physically plausible in the input images. These challenging tasks allow us to assess the alignment of VLMs with human-like spatial reasoning capabilities. Our contributions are summarized as follows:

- We design a hierarchical evaluation framework to study the spatial perception and rea-

soning capabilities of models, progressing from single-skill, multi-skill, to complex reasoning tasks. The reasoning tasks are further decomposed into simpler sub-tasks to disentangle spatial perception and reasoning.

- We manually annotate and curate an evaluation dataset using real-world images from MS COCO based on the designed tasks.
- We evaluate state-of-the-art VLMs and reveal significant shortcomings in their capabilities. These models especially lack the ability to understand distance and proximity, to reason from both allocentric and egocentric viewpoints, and to perform complex reasoning in a physical world context.

2 Related works

2.1 Spatial Capabilities in Multimodal LLMs

The majority of literature on multimodal LLMs focus on their conversational ability to respond flu-

ently to the queries of human users through natural language (Dai et al., 2023; Liu et al., 2023b, 2024; Lyu et al., 2023) or other modalities such as audio and image (Wu et al., 2023; Tang et al., 2024). More recently, works have started to explore the integration of spatial capabilities in LLMs to better understand the physical world. Zheng et al. (2024) developed BAT based on text and audio by synthesizing data with a spatial audio simulator, and Devnani et al. (2024) trained a spatially-aware audio and text embedding model. SpatialBot (Cai et al., 2024) directly includes pairs of RGB and depth images in existing datasets for training. A popular approach for vision-language models is to use existing algorithms and models (e.g. object detection, segmentation, depth estimation) to extract additional information from available data, and then incorporate these enhanced data for training or fine-tuning (Chen et al., 2024a; Zhao et al., 2023; AI, 2024a; Cai et al., 2024; Cheng et al., 2024). The resulting models have improved capabilities to infer spatial relationships between objects and to estimate object dimensions from natural images.

2.2 Benchmark Evaluations

Spatial reasoning is necessary for VLMs to understand and interact with the physical world. This has spurred the development of numerous benchmarks that vary in scope, tasks, and the types of spatial skills they focus on. Some benchmarks (Li et al., 2024b; Yu et al., 2024) involve spatial capabilities in only a small subset of tasks. EmbSpatial-Bench (Du et al., 2024) evaluates only a limited number of positional relationships, namely ‘above’, ‘below’, ‘left’, ‘right’, ‘far’ and ‘close’. Zeng et al. (2024) tests the effects of spatial relations on visual grounding tasks. VSR (Liu et al., 2023a) includes proximity and topological concepts. SpatialBench (Cai et al., 2024) consists of a small dataset of 120 images to evaluate spatial relationships, counting and size comparison between objects. Q-Spatial Bench (Liao et al., 2024) and SpatialRGPT-Bench (Cheng et al., 2024) focuses on model ability to recognize quantitative information such as metric distances and object sizes. These evaluations and datasets do not disentangle the effects of multiple spatial and/or visual skills. CLEVR (Johnson et al., 2016) tests compositional reasoning with multiple spatial and visual skills, but the dataset is synthetic and the objects comprise only simple 3D shapes such as spheres, cubes and cylinders. Kamath et al. (2023) isolates the effect of spatial attributes by

varying the spatial relations of objects while fixing object identities, but the collected dataset focuses only on positional relationships. We study the interaction effect of a wider range of skills, and introduce a novel study of model ability to infer about scenes or objects not explicitly observed but physically plausible in the input images.

3 The SPHERE Benchmark

3.1 Overview

The SPHERE evaluation framework is designed to systematically disentangle essential spatial and visual capabilities and to assess VLM performance in tasks that can guide physical world actions. The hierarchy of tasks is illustrated in Figure 2. The first two levels emphasize image understanding, while the final level builds upon image understanding and focuses on analysis and reasoning.

SPHERE begins at the first level with 4 single-skill tasks: Position, Counting, Distance and Size. We choose these skills because they are foundational building blocks for higher-order tasks, and are of interest to the research community as demonstrated through existing efforts to endow VLMs with these skills (Cai et al., 2024; Chen et al., 2024a; Du et al., 2024). At the second level, the first level skills are integrated into 3 multi-skill tasks: Position + Counting, Distance + Counting, and Distance + Size. We choose to combine 2 skills at a time to minimize confounding factors. Tasks in this category are more challenging. Specifically, Distance + Size require models to understand the concept of size constancy, instead of simply measuring size by the number of pixels occupied in an image. At the third level, 2 reasoning tasks assess models on their ability to reason about object occlusion and object manipulation. These tasks reflect scenarios commonly encountered in the physical world, such as inferring the presence or position of hidden objects or determining how objects can be moved or rearranged. Logical reasoning together with multiple spatial and visual understanding abilities may be required to successfully answer the questions in this category.

3.2 Benchmark Dataset

We manually curate and annotate a question-answering dataset for SPHERE, as automatic generation with VLMs remain unreliable due to accuracy limitations, as demonstrated in Table 1 and GPT-4o failure cases in Figure 1. We use images from

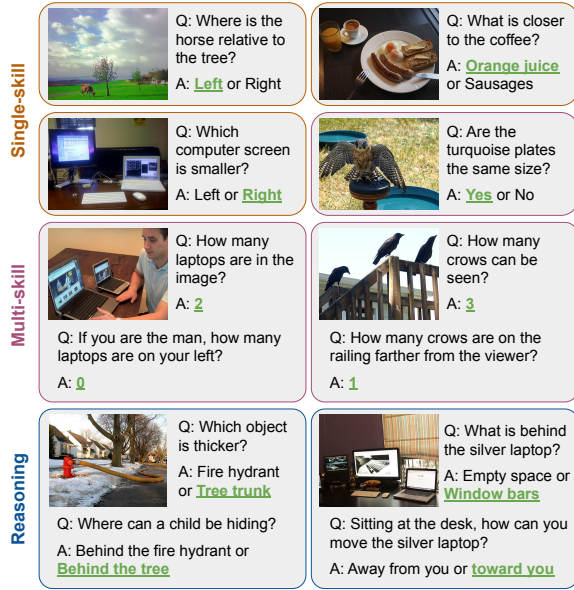


Figure 3: Examples of open-ended questions for counting-related tasks and multiple-choice questions for other tasks in the annotated SPHERE dataset. Ground-truth answers are in green.

the test split of MS COCO-2017 (Lin et al., 2014) due to its rich diversity in scenes and object categories. Since COCO-2017 is a widely-used and publicly available dataset, its images are likely to be in-distribution with respect to training datasets used by VLMs, ensuring that the evaluation aligns with their learned representations. This alignment reduces potential biases stemming from dataset shifts.

Annotations for each task are conducted and cross-verified by at least two authors for consistency and accuracy. Questions that did not elicit the same responses are discarded and replaced. The general guidelines for annotation are:

- Ensure non-ambiguity and diversity of question and answers.
- Omit images that may be irrelevant or ambiguous for the task.
- Counting-related tasks require open-ended numerical responses, while other tasks use multiple-choice questions (MCQs) with two options by default and three options when additional clarity is needed.

We use diverse types of options for MCQs, including boolean (e.g. yes vs. no), position descriptor (e.g. left vs. right), and name of object. Figure 3 shows examples of annotated data.

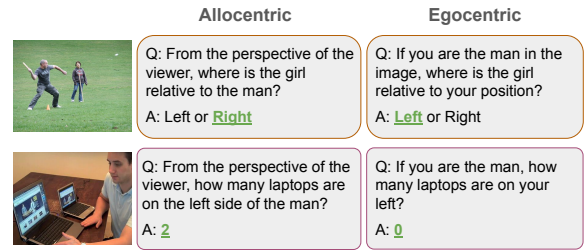


Figure 4: Examples of position-related questions asked from allocentric and egocentric viewpoints.

Task level	Task type	Sample size
Single-skill	Position	Ego.: 172 Allo.: 185
	Counting	201
	Distance	202
	Size	198
Multi-skill	Position + Counting	Ego.: 64 Allo.: 105
	Distance + Counting	158
	Distance + Size	199
Reasoning	Object occlusion	Intermed.: 202 Final: 200
	Object manipulation	Intermed.: 199 Final: 200

Table 2: The annotated dataset has a total of 2,285 question-answer pairs. For position-related tasks, we annotate for both egocentric (ego.) and allocentric (allo.) viewpoints. For reasoning tasks, we additionally annotate for intermediate (intermed.) questions.

3.2.1 Dataset Design and Statistics

We annotate a total of 2,285 question-answer pairs. The sample size for each task is delineated in Table 2. Detailed dataset statistics and analysis of dataset diversity are in Appendix A.

Single-skill tasks. We annotate for 4 single-tasks. Basic object recognition capabilities are assumed.

- **Position:** We pose questions on relative positional relationships (e.g. left vs. right, in front vs. behind, on top vs. below) between objects in the scene. For entities with clearly defined front and back orientations (e.g. humans, animals, vehicles), we include questions for both *egocentric* and *allocentric* viewpoints. Egocentric viewpoint questions assess the model’s ability to understand spatial relationships from a specified entity’s perspective, such as identifying objects to the entity’s left or right relative to its orientation. Allocentric viewpoint questions evaluate the model’s ability to describe spatial relationships from the external




Figure 5: Distribution of ground-truth answers for the counting task.




Figure 6: Distribution of base spatial and visual skills required for questions in reasoning tasks.

point of view of the camera or observer. Example questions are shown in Figure 4.

- **Counting:** These questions require the model to count the number of specified objects visible in the image. To evaluate the model’s susceptibility to object hallucinations, we included ‘trick’ questions where the correct answer is zero. Figure 5 illustrates the distribution of ground-truth answers.
- **Distance:** We ask questions on the relative proximity (e.g. closer vs. farther) between objects, and between objects and the viewer.
- **Size:** We ask questions on the relative size (e.g. smaller vs. bigger, shorter vs. taller) of objects. We focus on objects positioned approximately equidistant from the camera, where object size can be inferred directly from the number of pixels occupied by the object on the image.

Multi-skill tasks: We integrate the first level skills into 3 multi-skill tasks.

- **Position + Counting:** Questions require the model to count the number of specified objects that have a specified spatial relationship with another object in the image. We include both egocentric and allocentric viewpoints. Example questions are shown in Figure 4.
- **Distance + Counting:** Questions require the

model to count the number of specified objects that are at a specified proximity to another object in the image or the observer. An example of a question in Figure 3 is “How many crows are on the railing farther from the viewer?”

- **Distance + Size:** We pose questions about the relative size of objects located at different distances from the camera. To succeed at this task, the model must demonstrate an understanding of size constancy, or the ability to recognize that an object’s actual size remains consistent despite changes in its apparent size due to distance or perspective.

Reasoning tasks: We design 2 tasks to evaluate the model’s ability to logically reason about the 3D physical world based on its understanding of the input 2D image. In addition to the final reasoning-based questions, we include understanding-based questions as an intermediate step to check the model’s comprehension of the 2D image. In our design, “intermediate” questions focus on assessing the perception and understanding of visible objects and scenes. These tasks align with cognitive theories like Piaget’s stages of cognitive development, which suggest that perceptual skills are crucial for building more abstract reasoning abilities assessed through the “final” reasoning questions. The “final” reasoning questions require the ability to infer the existence, position, or state of partially or non-visible objects and scenes that are physically plausible. These inferences rely on understanding spatial relationships and object permanence, as the missing information could only be revealed by changing viewpoints or interacting with the scene, such as moving objects to uncover hidden ones.

- **Object occlusion:** This task evaluates a model’s ability to infer the presence or properties of objects that may be fully occluded from the perspective of observer. From Figure 3, an example of an intermediate understanding-based question is “Which object is thicker? Fire hydrant or tree trunk?” The final reasoning-based question is “Where can a child be hiding? Behind the fire hydrant or behind the tree?” To answer this question correctly, the model needs to reason that a child is more likely to be obscured by the thicker tree.
- **Object manipulation:** This task tests a model’s ability to reason about potential or hypothetical positions that objects can move to

or be moved to within a scene, based on physical constraints. From Figure 3, an example of an intermediate understanding-based question is “What is behind the silver laptop? Empty space or window bars?” The final reasoning-based question is “Sitting at the desk, how can you move the silver laptop? Away from you or toward you?” To answer this question correctly, the model needs to reason that a person at the desk cannot move the laptop further back as the laptop is already positioned against the window bars.

Successfully solving reasoning tasks requires multiple spatial and visual skills. We categorize these foundational skills as counting, distance, position, shape, size and other visual skills, such as perceiving color and transparency. We annotate these foundational skills and illustrate the distribution of their occurrences in Figure 6.

4 Evaluations

4.1 Setup

We describe the experimental setup used to evaluate the performance of state-of-the-art VLMs on SPHERE tasks. All experiments were performed on NVIDIA GeForce RTX 3090 GPUs.

Models: Our evaluation encompasses two model categories: general-purpose VLMs and VLMs with enhanced spatial capabilities. The general-purpose VLMs evaluated are Phi-3.5-Vision (Abdin et al., 2024), LLaVA-NeXT (Liu et al., 2024), LLaVA-OneVision (Li et al., 2024a), InstructBLIP (Dai et al., 2023), Idefics2 (HuggingFaceM4, 2024), InternVL2.5 (Chen et al., 2024b), Janus-Pro (Chen et al., 2025), Qwen-VL (Bai et al., 2023), Qwen2-VL (Wang et al., 2024), Qwen2.5-VL (Bai et al., 2025), Llama-3.2-Vision (Meta, 2024), Gemini 2.0 Flash (Google, 2024), and GPT-4o (OpenAI, 2024). The VLMs specifically trained for spatial understanding and reasoning are SpatialBot (Cai et al., 2024), SpaceMantis (AI, 2024b), and SpatialRGPT (Cheng et al., 2024). For SpatialRGPT, we use the RGB-only version as the model requires highly accurate mask proposals to make use of depth maps. All model implementations follow their released versions on Hugging Face. We use the default prompt templates and hyperparameters in the demos of the respective models.

The text instructions given take the form:

“<question> Answer the question directly.”

For example, for the first image in Figure 4, the text instruction given is *“From the perspective of the viewer, where is the girl relative to the man? Left or Right? Answer the question directly.”* We added the instruction “Answer the question directly” to reduce the refusal-to-answer rate and to encourage concise answers that align with the options provided.

Metrics: We computed both the *validity* and *accuracy* of the VLM responses. Since the VLM generative responses are open-ended, we deem a response as valid if it is a relevant answer to the question, that is, a number is given for a counting question and the response contains one of the options for a multiple-choice question (MCQ). Accuracy measures the correctness of the model responses, and invalid responses are treated as incorrect. For MCQs, we set the baseline accuracy for each question as $1/\#options$, since this is the probability of selecting the correct answer purely by random guessing.

To account for the probabilistic nature of VLM generations, we evaluate each question 5 times using different random seeds and average the validity and accuracy scores across these runs. For MCQs, we introduce additional randomness by shuffling the order of the answer options at each seed. This prevents the models from exploiting positional biases in the option sequence and ensures that their performance reflects true understanding rather than reliance on specific patterns. Through these considerations, we increase the robustness of our evaluation, providing a more reliable measure of the models’ generalization and consistency.

4.2 Results

Table 3 shows model accuracy on each of the SPHERE tasks. The results highlight significant room for improvement, especially for multi-skill and reasoning tasks.

The proprietary model GPT-4o outperforms the open-weights models evaluated in almost all tasks and has an overall accuracy of 67.9%. Qwen2.5-VL (72B) has the second highest overall accuracy of 64.3%. Interestingly, amongst open-weights models, we observe that the smaller model LLaVA-OneVision (7B) can outperform larger models in terms of overall accuracy, and performs on par with Gemini 2.0 Flash.

Model	Single-skill					Multi-skill				Reasoning			Overall
	Pos.	Count.	Dist.	Size	Avg	P + C	D + C	D + S	Avg	Occl.	Manip.	Avg	
Human baseline	90.4	98.0	98.0	95.0	95.4	95.6	91.0	91.0	92.5	87.0	91.0	89.0	93.0
Random baseline*	50.0	-	50.0	50.0	50.0	-	-	44.3	44.3	50.0	50.0	50.0	49.1
General purpose VLMs													
Phi-3.5-Vision (4B)	61.8	58.7	56.7	76.0	63.3	44.4	36.1	47.2	42.6	54.4	55.5	54.9	54.3
LLaVA-NeXT (7B)	53.6	68.7	54.5	70.9	61.9	42.0	35.4	36.1	37.8	48.7	49.1	48.9	50.6
LLaVA-OneVision (7B)	60.2	76.1	64.2	84.0	71.1	53.3	53.8	52.9	53.3	53.0	59.6	56.3	61.5
Qwen2-VL (7B)	63.2	74.1	53.0	78.6	67.2	36.1	33.5	50.8	40.1	54.0	55.2	54.6	55.0
Qwen2.5-VL (7B)	59.3	82.1	58.9	73.6	68.5	44.7	31.6	37.9	38.1	48.7	54.1	51.4	54.1
Janus-Pro (7B)	57.1	70.6	59.3	76.5	65.9	37.3	33.5	46.8	39.2	53.5	54.3	53.9	54.0
InstructBLIP (8B)	45.6	64.7	50.4	56.2	54.2	32.0	28.5	38.1	32.8	54.8	52.7	53.7	47.0
Idefics2 (8B)	50.1	45.3	49.1	58.5	50.7	22.5	27.8	40.1	30.1	51.0	53.3	52.1	44.2
InternVL2.5 (8B)	62.2	72.6	64.6	78.3	69.4	50.2	43.8	41.8	45.3	52.7	53.6	53.1	57.3
Qwen-VL (10B)	55.9	72.9	59.2	72.0	65.0	37.2	29.9	34.7	33.9	54.3	54.3	54.3	52.0
Llama-3.2-Vision (11B)	58.4	53.5	52.9	67.2	58.0	29.8	27.7	40.8	32.8	60.7	56.1	58.4	49.7
Qwen2-VL (72B)	62.8	80.6	60.8	85.5	72.4	37.3	38.0	69.3	48.2	53.9	55.0	54.4	59.8
Qwen2.5-VL (72B)	63.7	81.1	69.9	89.3	76.0	51.5	55.7	66.4	57.9	50.3	56.4	53.3	64.3
Llama-3.2-Vision (90B)	66.0	59.2	65.2	75.8	66.5	40.1	31.8	52.8	41.6	61.6	55.3	58.4	56.2
Gemini 2.0 Flash	68.6	82.1	75.2	86.9	78.2	50.3	52.5	49.7	50.9	48.0	49.2	48.6	61.7
GPT-4o	71.7	76.6	71.3	89.4	77.3	54.4	62.7	58.8	58.6	66.0	63.3	64.7	67.9
Spatial VLMs													
SpatialBot-RGB (3B)	55.9	67.7	52.8	73.0	62.3	38.5	39.2	38.7	38.8	57.6	60.3	59.0	53.6
SpatialBot (3B)	55.5	72.1	51.9	75.4	63.7	39.1	36.1	38.9	38.0	56.5	54.6	55.5	53.1
SpaceMantis (8B)	52.7	52.7	57.0	63.1	56.4	28.4	43.7	45.6	39.2	56.6	54.2	55.4	50.4
SpatialRGPT-RGB (8B)	59.3	70.1	59.2	74.6	65.8	42.6	46.2	40.9	43.2	54.5	53.2	53.8	55.3

Table 3: Average accuracy (%) of VLMs on SPHERE tasks. Values highlighted in gold, silver and bronze denote the first, second and third place for each task. The random baseline accuracy for multiple-choice questions is the probability of selecting the correct answer from given options purely by random guessing. * Average scores for random baseline excludes open-ended counting tasks.

For the single-skill tasks, Gemini 2.0 Flash has the highest average accuracy (78.2%). The best performances on Counting and Size are above 80%. The best performances on Position and Distance are between 70% and 80%, indicating that current VLMs still have difficulty in these two skills. In Section 4.3, we analyze how viewpoint affects model performance on understanding positions.

For the multi-skill tasks, GPT-4o has the highest average accuracy (58.6%). The multi-skill tasks are markedly more challenging than the single-skill tasks. In particular, we observe that half of the models perform worse than baseline random guessing on Distance + Size. Qwen2-VL (72B) obtains the highest 69.3% on Distance + Size, and 85.5% on Size. This gap reflects the reliance of the VLMs on the 2D representation of objects in an image, and measures object size as being directly proportional to the number of pixels occupied. The failure to incorporate the principle of size constancy may cause VLMs to struggle in 3D spaces.

For the reasoning tasks, GPT-4o has the highest

average accuracy (64.7%). Most models perform close to random guessing, suggesting that their training datasets may lack sufficient examples emphasizing reasoning skills. Notably, Gemini 2.0 Flash has the worst average accuracy (48.6%), below the random guessing baseline.

In general, we note that the spatial VLMs evaluated performed comparably with or better than most similarly-sized general-purpose VLMs, except for LLaVA-OneVision. In particular, SpatialBot (3B) performed comparably with or better than larger models with up to 11B parameters. This shows that the techniques in spatial VLMs can benefit spatial capabilities. Possible reasons on why the spatial VLMs do not excel on the SPHERE tasks may be the lack of training on the specific skills emphasized by the SPHERE tasks, and the lack of model generalization to the SPHERE tasks.

We refer readers to Table 7 in the Appendix B for detailed validity scores of the models. The validity scores reflect the instruction-following capabilities of the VLMs. Idefics2 has above 95% validity, and

Model	Position		Position + Counting	
	Allo.	Ego.	Allo.	Ego.
General purpose VLMs				
Phi-3.5-Vision (4B)	77.9	44.5	53.3	29.7
LLaVA-NeXT (7B)	69.3	36.7	47.6	32.8
LLaVA-OneVision (7B)	73.7	45.7	62.9	37.5
Qwen2-VL (7B)	69.1	56.0	37.1	37.5
Qwen2.5-VL (7B)	69.0	48.8	49.1	37.5
Janus-Pro (7B)	65.0	48.6	37.1	37.5
InstructBLIP (8B)	43.0	48.4	31.4	32.8
Idefics2 (8B)	48.5	51.9	20.0	26.6
InternVL2.5 (8B)	75.7	47.7	59.8	34.4
Qwen-VL (10B)	65.6	45.5	41.5	30.0
Llama-3.2-Vision (11B)	60.5	56.0	31.6	26.9
Qwen2-VL (72B)	70.3	55.6	36.2	35.9
Qwen2.5-VL (72B)	78.7	47.6	58.1	40.6
Llama-3.2-Vision (90B)	72.9	58.6	41.1	38.4
Gemini 2.0 Flash	83.2	52.9	55.2	42.2
GPT-4o	79.5	63.4	62.9	40.6
Spatial VLMs				
SpatialBot-RGB (3B)	68.6	42.2	47.6	23.4
SpatialBot (3B)	69.7	40.2	46.7	26.6
SpaceMantis (8B)	56.2	48.8	31.4	23.4
SpatialRGPT-RGB (8B)	73.2	44.4	48.6	32.8

Table 4: Average accuracy (%) of VLMs on position-related questions asked from allocentric (allo.) and egocentric (ego.) viewpoints. Values highlighted in gold, silver and bronze denote the first, second and third place.

all other models have above 98% validity.

4.3 Further Analysis

Bias toward allocentric or egocentric viewpoint.

We further analyze model performance on position-related tasks by distinguishing between allocentric and egocentric viewpoints. This analysis provides insights into whether VLMs can comprehend object orientations and reason from both relative (object-to-object) and absolute (viewer-to-object) perspectives. Example questions are shown in Figure 4. From the results in Table 4, we observe that most VLMs demonstrate a significant bias toward the allocentric viewpoint, with performance gaps as large as 33.4% for Phi-3.5-Vision in the Position task, and 25.4% for LLaVA-OneVision in the Position + Counting task.

Effectiveness of additional information for reasoning. Since the reasoning tasks first require image understanding of the factual information presented in the images and then logical reasoning to conduct analysis and inference, we further disentangle the two types of abilities. We measure the strength of model reasoning capabilities alone by directly providing understanding-based informa-

tion in the input prompts. That is, we prefix the final step question with the intermediate step question and ground-truth answer as:

“Given that for the question: <intermediate step question> The answer is: <intermediate step answer>. <final step question> Answer the question directly.”

For the example in Figure 3, the prompt will be *“Given that for the question: Which object is thicker? Fire hydrant or tree trunk? The answer is: Tree trunk. Where can a child be hiding? Behind the fire hydrant or behind the tree? Answer the question directly.”*

From Table 5, we observe that incorporating additional information improves final question accuracy across all models for the object occlusion task, but yields mixed results for the object manipulation task. Although the additional information can introduce relevant facts that the models might otherwise overlook, the models sometimes struggle to establish a clear connection between these facts and the specific reasoning they need to conduct.

Skill composition. The different performance of the evaluated models across tasks reflect that they have different compositions of skills, and that proficiency in one task does not necessarily translate to proficiency in related tasks. For instance, relatively weaker basic skills do not necessarily imply correspondingly weaker higher-order skills. In Table 3, SpaceMantis has 7.3% lower average accuracy than SpatialBot on the single-skill tasks, but has 1.2% higher average accuracy than SpatialBot on the multi-skill tasks. In Table 5, on the object manipulation task, Idefics2 and Llama-3.2-Vision (11B) have attained higher accuracy on the final reasoning questions (53.3% and 56.1%) than on the supposedly easier intermediate perception or understanding questions (51.1% and 53.7%).

5 Conclusion

This work introduces SPHERE, a comprehensive evaluation framework designed to address the critical gaps in spatial understanding and reasoning exhibited by current vision-language models. Through a hierarchy of tasks and a meticulously curated dataset, SPHERE highlights significant limitations in the ability of state-of-the-art models, especially in processing spatial cues such as distance, reasoning from both allocentric and egocentric viewpoints, and analyzing the depicted scene

Model	Object occlusion			Object manipulation		
	Intermed.	Final	Final (w/ gt intermed.)	Intermed.	Final	Final (w/ gt intermed.)
General purpose VLMs						
Phi-3.5-Vision (4B)	60.6	54.4	67.0 (↑ 12.6)	59.8	55.5	57.6 (↑ 2.1)
LLaVA-NeXT (7B)	60.4	48.7	54.4 (↑ 5.7)	58.5	49.1	48.8 (↓ 0.3)
LLaVA-OneVision (7B)	63.9	53.0	59.8 (↑ 6.8)	64.2	59.6	57.1 (↓ 2.5)
Qwen2-VL (7B)	63.4	54.0	60.8 (↑ 6.8)	55.1	55.2	55.4 (↑ 0.2)
Qwen2.5-VL (7B)	65.2	48.7	51.1 (↑ 2.4)	61.1	54.1	55.4 (↑ 1.3)
Janus-Pro (7B)	58.4	53.5	55.0 (↑ 1.5)	62.6	54.3	53.0 (↓ 1.3)
InstructBLIP (8B)	48.7	54.8	57.8 (↑ 3.0)	52.9	52.7	45.7 (↓ 7.0)
Idefics2 (8B)	56.5	51.0	60.0 (↑ 9.0)	51.1	53.3	60.7 (↑ 7.4)
InternVL2.5 (8B)	63.7	52.7	58.9 (↑ 6.2)	62.9	53.6	57.2 (↑ 3.6)
Qwen-VL (10B)	57.8	54.3	57.3 (↑ 3.0)	54.5	54.3	49.0 (↓ 5.3)
Llama-3.2-Vision (11B)	52.3	60.7	62.5 (↑ 1.8)	53.7	56.1	57.5 (↑ 1.4)
Qwen2-VL (72B)	65.6	53.9	62.4 (↑ 8.5)	68.5	55.0	76.9 (↑ 21.9)
Qwen2.5-VL (72B)	67.3	50.3	57.6 (↑ 7.3)	64.2	56.4	66.9 (↑ 10.5)
Llama-3.2-Vision (90B)	62.7	61.6	68.0 (↑ 6.4)	59.8	55.3	74.0 (↑ 18.7)
Gemini 2.0 Flash	62.4	48.0	61.0 (↑ 13.0)	65.5	49.2	58.3 (↑ 9.1)
GPT-4o	67.3	66.0	72.0 (↑ 6.0)	68.5	63.3	76.4 (↑ 13.1)
Spatial VLMs						
SpatialBot-RGB (3B)	59.9	57.6	65.9 (↑ 8.3)	60.8	60.3	54.0 (↓ 6.3)
SpatialBot (3B)	58.6	56.5	63.2 (↑ 6.7)	59.2	54.6	56.1 (↑ 1.5)
SpaceMantis (8B)	63.1	56.6	62.4 (↑ 5.8)	60.1	54.2	57.6 (↑ 3.4)
SpatialRGPT-RGB (8B)	61.1	54.5	63.2 (↑ 8.7)	60.1	53.2	58.5 (↑ 5.3)

Table 5: Average accuracy (%) of VLMs on the intermediate (intermed.) and final questions in the reasoning tasks. The approach ‘final (w/ gt intermed.)’ denotes that the intermediate questions and ground-truth answers are used as additional information to assist the answering of the final questions.

in a physical world context. Through this work, we seek to lay the groundwork for advancing vision-language models toward more human-like spatial perception and reasoning, and to set a benchmark for future research in this domain.

6 Limitations

As we seek to ensure the quality of the SPHERE dataset through manual annotations, the dataset may be limited in scope. The framework and the curated tasks may not encompass all possible spatial understanding and reasoning challenges encountered in practice. In addition, the dataset focuses on static images, which may not fully capture the challenges of dynamic spatial understanding and reasoning in real-world scenarios. As VLMs advance in spatial perception and reasoning, automatic annotation methods may become more reliable and warrant exploration in future work.

7 Ethics Statement

All real-world images used in our annotated dataset are sourced from the MS COCO-2017 dataset (Lin et al., 2014), a widely used and publicly available

resource. Our annotations focus solely on the spatial properties of the entities in the images and do not involve any personally identifiable information. The authors have unanimously agreed that the annotated data will be used exclusively for academic research purposes, in compliance with the dataset’s terms of use (COCO Consortium, 2024).

Acknowledgements

We thank Chanwoo Hwang from Hanyang University for valuable discussions. This research is partially supported by the Nanyang Associate Professorship and the National Research Foundation Fellowship (NRF-NRFF13-2021-0006), Singapore.

References

- Marah Abdin, Sam Ade Jacobs, Ammar Ahmad Awan, Jyoti Aneja, Ahmed Awadallah, and et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv*.
- Remyx AI. 2024a. SpaceLLaVA. <https://huggingface.co/remyxai/SpaceLLaVA>.
- Remyx AI. 2024b. SpaceMantis. <https://huggingface.co/remyxai/SpaceMantis>.

- Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. 2023. Qwen-VL: A versatile vision-language model for understanding, localization, text reading, and beyond. *arXiv*.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. 2025. Qwen2.5-VL technical report. *arXiv*.
- Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Xi Chen, Krzysztof Choromanski, Tianli Ding, Danny Driess, Avinava Dubey, Chelsea Finn, Pete Florence, Chuyuan Fu, Montse Arenas, Keerthana Gopalakrishnan, Kehang Han, Karol Hausman, Alexander Herzog, Jasmine Hsu, Brian Ichter, and Brianna Zitkovich. 2023. RT-2: Vision-language-action models transfer web knowledge to robotic control. *arXiv*.
- Wenxiao Cai, Yaroslav Ponomarenko, Jianhao Yuan, Xiaoqi Li, Wankou Yang, Hao Dong, and Bo Zhao. 2024. SpatialBot: Precise spatial understanding with vision language models. *arXiv*.
- Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. 2024a. SpatialVLM: Endowing vision-language models with spatial reasoning capabilities. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Xiaokang Chen, Zhiyu Wu, Xingchao Liu, Zizheng Pan, Wen Liu, Zhenda Xie, Xingkai Yu, and Chong Ruan. 2025. Janus-Pro: Unified multimodal understanding and generation with data and model scaling. *arXiv*.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. 2024b. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. 2024. SpatialRGPT: Grounded spatial reasoning in vision-language models. *arXiv*.
- COCO Consortium. 2024. COCO Terms of Use. <https://cocodataset.org/#termsofuse>.
- Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Albert Li, Pascale Fung, and Steven C. H. Hoi. 2023. InstructBLIP: Towards general-purpose vision-language models with instruction tuning. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Bhavika Suresh Devnani, Skyler Seto, Zakaria Aldeneh, Alessandro Toso, Yelena Menyaylenko, Barry-John Theobald, Jonathan Sheaffer, and Miguel Sarabia. 2024. Learning spatially-aware language and audio embeddings. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Mengfei Du, Binhao Wu, Zejun Li, Xuanjing Huang, and Zhongyu Wei. 2024. EmbSpatial-bench: Benchmarking spatial understanding for embodied tasks with large vision-language models. In *Annual Meeting of the Association for Computational Linguistics*.
- Google. 2024. Introducing Gemini 2.0: our new AI model for the agentic era. <https://blog.google/technology/google-deepmind/google-gemini-ai-update-december-2024>.
- HuggingFaceM4. 2024. idefics2-8B. <https://huggingface.co/HuggingFaceM4/idefics2-8b>.
- Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross B. Girshick. 2016. CLEVR: A diagnostic dataset for compositional language and elementary visual reasoning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Amita Kamath, Jack Hessel, and Kai-Wei Chang. 2023. What’s “up” with vision-language models? Investigating their struggle with spatial reasoning. In *EMNLP*.
- Isabel Leal, Krzysztof Choromanski, Deepali Jain, Kumar Avinava Dubey, Jake Varley, Michael Ryoo, Yao Lu, Frederick Liu, Vikas Sindhwani, Quan Ho Vuong, Tamás Sarlós, Kenneth Oslund, Karol Hausman, and Kanishka Rao. 2023. SARA-RT: Scaling up robotics transformers with self-adaptive robust attention. In *International Conference on Robotics and Automation (ICRA)*.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. 2024a. LLaVA-OneVision: Easy visual task transfer. *arXiv*.
- Bohao Li, Yuying Ge, Yixiao Ge, Guangzhi Wang, Rui Wang, Ruimao Zhang, and Ying Shan. 2024b. Seed-bench: Benchmarking multimodal large language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yuan-Hong Liao, Rafid Mahmood, Sanja Fidler, and David Acuna. 2024. Reasoning paths with reference objects elicit quantitative spatial reasoning in large vision-language models. *arXiv*.
- Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. 2014. Microsoft COCO: Common objects in context. In *European Conference on Computer Vision (ECCV)*.
- Fangyu Liu, Guy Emerson, and Nigel Collier. 2023a. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 11:635–651.

- Haotian Liu, Chunyuan Li, Yuheng Li, Bo Li, Yuanhan Zhang, Sheng Shen, and Yong Jae Lee. 2024. *LLaVA-NeXT: Improved reasoning, OCR, and world knowledge*.
- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. 2023b. Visual instruction tuning. In *Conference on Neural Information Processing Systems (NeurIPS)*.
- Chenyang Lyu, Minghao Wu, Longyue Wang, Xinting Huang, Bingshuai Liu, Zefeng Du, Shuming Shi, and Zhaopeng Tu. 2023. Macaw-LLM: Multi-modal language modeling with image, audio, video, and text integration. *arXiv*.
- Meta. 2024. Llama 3.2: Revolutionizing edge AI and vision with open, customizable models. <https://ai.meta.com/blog/llama-3-2-connect-2024-vision-edge-mobile-devices/>.
- Suraj Nair, Aravind Rajeswaran, Vikash Kumar, Chelsea Finn, and Abhi Gupta. 2022. R3M: A universal visual representation for robot manipulation. In *Conference on Robot Learning*.
- OpenAI. 2024. Gpt-4o system card. *arXiv*.
- Austin Stone, Ted Xiao, Yao Lu, Keerthana Gopalakrishnan, Kuang-Huei Lee, Quan Vuong, Paul Wohlhart, Sean Kirmani, Brianna Zitkovich, Fei Xia, Chelsea Finn, and Karol Hausman. 2023. Open-world object manipulation using pre-trained vision-language models. *arXiv*.
- Zineng Tang, Ziyi Yang, Mahmoud Khademi, Yang Liu, Chenguang Zhu, and Mohit Bansal. 2024. CoDi-2: In-context, interleaved, and interactive any-to-any generation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Anthony Meng Huat Tiong, Junqi Zhao, Boyang Albert Li, Junnan Li, Steven C.H. Hoi, and Caiming Xiong. 2024. What are we measuring when we evaluate large vision-language models? an analysis of latent factors and biases. In *North American Chapter of the Association for Computational Linguistics*.
- Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei Du, Xuancheng Ren, Rui Men, Dayiheng Liu, Chang Zhou, Jingren Zhou, and Junyang Lin. 2024. Qwen2-VL: Enhancing vision-language model’s perception of the world at any resolution. *arXiv*.
- Peter West, Ximing Lu, Nouha Dziri, Faeze Brahman, Linjie Li, Jena D. Hwang, Liwei Jiang, Jillian R. Fisher, Abhilasha Ravichander, Khyathi Raghavi Chandu, Benjamin Newman, Pang Wei Koh, Allyson Ettinger, and Yejin Choi. 2023. The generative ai paradox: "what it can create, it may not understand". *arXiv*.
- Shengqiong Wu, Hao Fei, Leigang Qu, Wei Ji, and Tat-Seng Chua. 2023. NExT-GPT: Any-to-any multi-modal LLM. *arXiv*.
- Weihao Yu, Zhengyuan Yang, Linjie Li, Jianfeng Wang, Kevin Lin, Zicheng Liu, Xinchao Wang, and Lijuan Wang. 2024. MM-Vet: Evaluating large multimodal models for integrated capabilities. In *International Conference on Machine Learning*. PMLR.
- Yunan Zeng, Yan Huang, Jinjin Zhang, Zequn Jie, Zhenhua Chai, and Liang Wang. 2024. Investigating compositional challenges in vision-language models for visual grounding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yongqiang Zhao, Zhenyu Li, Zhi Jin, Feng Zhang, Haiyan Zhao, Chengfeng Dou, Zhengwei Tao, Xinhai Xu, and Donghong Liu. 2023. Enhancing the spatial awareness capability of multi-modal large language model. *arXiv*.
- Zhisheng Zheng, Puyuan Peng, Ziyang Ma, Xie Chen, Eunsol Choi, and David Harwath. 2024. BAT: Learning to reason about spatial sounds with large language models. *arXiv*.

A Detailed Dataset Statistics

We provide additional details of our annotated SPHERE dataset.

Number of options. For non-counting tasks, questions are posed as multiple-choice questions (MCQs). Figure 7 shows the distribution of number of options across all tasks. All non-counting tasks except Distance + Size only have two-option MCQs, while around a third of MCQs in Distance + Size have three options. When comparing the size of two objects and asking the models to identify the smaller/bigger object, we provide options of the form ‘the closer object’, ‘the farther object’ and ‘the two objects are of similar size’. Such three-option MCQs are designed to evaluate if the model has an understanding of size constancy when similar objects may appear to be of different sizes in the image due to their positioning in the 3D scene.

Numerical answers. Counting tasks are phrased as open-ended numerical questions, designed to evaluate the models’ ability to recognize objects matching the required criterion while distinguishing such objects as distinct entities. Figure 8 shows the distribution of ground-truth answers across all counting-related tasks. We limit the number of objects to be counted to a maximum of 9 to avoid including overly small objects.

We check the count distributions by fitting a negative binomial model to each distribution using the `goodfit` function from R’s `vcd` package. The results of Pearson’s χ^2 tests, shown in Table 6, indicate that all p -values exceeded 0.10. This suggests that the count distributions align well with the negative binomial model and are not excessively skewed toward specific values.

Formats. Figure 9a shows the distribution of the following option formats across all tasks:

- **Boolean:** True/false or other similar phrasings, e.g. “yes” vs. “no”.
- **Name:** Name of objects or their attributes, e.g. “empty space” vs. “window pane”, “black laptop” vs. “white laptop”, “curved” vs “straight”.
- **Numeric:** For counting tasks, a non-negative integer indicating the number of objects matching the required criterion. For non-counting tasks, MCQs may contain numbers in the provided options, e.g. “1/2” vs. “less than 1/4”, or require numerical estimation, e.g.




Figure 7: Distribution of number of options for MCQs across all non-counting tasks.

Task	χ^2	df	$P(> \chi^2)$
Counting	6.2420	7	0.51179
Distance + Counting	7.8268	6	0.25107
Position + Counting	7.8879	5	0.16252

Table 6: Goodness-of-fit tests on ground-truth answers for counting tasks fitted to a negative-binomial model.

“less than the length of the painting” vs. “more than the length of the painting”.

- **Position:** Position of objects, e.g. “left” vs. “right”, “on top” vs. “below”.

Viewpoints. Figure 9b shows the distribution of egocentric vs. allocentric viewpoints across all tasks. Position, Position + Counting, Distance + Counting and Object manipulation comprise both allocentric and egocentric questions, while all other tasks comprise only allocentric questions.

Object concepts. We assess the diversity of object concepts in the annotations by counting the occurrences of unique nouns. Nouns were extracted from each question using the `en_core_web_sm` model from `spaCy`, then manually filtered to exclude propositions and commonly referenced objects in question phrasings, such as ‘camera’, ‘viewer’, ‘photo’, ‘image’. The dataset contains a total of more than 600 object concepts. Figure 10 shows the distribution of object concepts across the dataset for the top 30 concepts.

B Detailed Evaluation Results

Table 7 shows the validity scores of the models across all tasks. These scores reflect the instruction-following capabilities of the VLMs. Most models




Figure 8: Distribution of ground-truth answers for all counting tasks.




Figure 9: Distribution of option formats and viewpoints across all tasks.

effectively follow instructions and provide valid responses, such as numerical answers for counting-related questions and selecting an option from the provided list for multiple-choice questions. Idefics2 has above 95% validity, and all other models have above 98% validity. Common types of invalid responses include providing an image caption instead of addressing the question, refusing to answer, or offering an answer that is not among the given options.

We provide example questions and answers for each task in Table 8, 9 and 10.










Figure 10: Distribution of object concepts across the dataset.

Model	Single-skill					Multi-skill				Reasoning			Overall
	Pos.	Count.	Dist.	Size	Avg	P + C	D + C	D + S	Avg	Occl.	Manip.	Avg	
General purpose VLMs													
Phi-3.5-Vision (4B)	100.0	100.0	99.7	100.0	99.9	100.0	100.0	99.0	99.7	100.0	100.0	100.0	99.9
LLaVA-NeXT (7B)	100.0	100.0	100.0	100.0	100.0	100.0	100.0	98.8	99.6	100.0	100.0	100.0	99.9
LLaVA-OneVision (7B)	99.7	100.0	99.8	98.4	99.5	100.0	100.0	97.3	99.1	99.7	100.0	99.8	99.4
Qwen2-VL (7B)	97.4	100.0	98.3	99.1	98.7	100.0	100.0	98.5	99.5	98.5	99.5	99.0	99.0
Qwen2.5-VL (7B)	94.0	100.0	98.8	98.8	97.9	100.0	99.4	96.8	98.7	98.4	97.3	97.9	98.2
Janus-Pro (7B)	100.0	100.0	100.0	100.0	100.0	100.0	100.0	99.8	99.9	99.8	100.0	99.9	100.0
InstructBLIP (8B)	97.1	100.0	100.0	99.7	99.2	99.4	100.0	99.9	99.8	99.8	100.0	99.9	99.6
Idefics2 (8B)	92.9	93.0	96.5	91.9	93.6	99.4	96.8	91.4	95.9	98.7	96.0	97.3	95.3
InternVL2.5 (8B)	98.2	100.0	99.4	98.9	99.1	100.0	100.0	99.3	99.8	99.8	100.0	99.9	99.5
Qwen-VL (10B)	99.0	100.0	99.5	100.0	99.6	100.0	100.0	98.4	99.5	96.8	98.8	97.8	99.1
Llama-3.2-Vision (11B)	99.0	100.0	99.0	98.3	99.1	100.0	100.0	98.6	99.5	100.0	99.9	99.9	99.4
Qwen2-VL (72B)	97.6	100.0	99.5	99.0	99.0	100.0	100.0	99.3	99.8	99.6	100.0	99.8	99.5
Qwen2.5-VL (72B)	99.2	100.0	100.0	100.0	99.8	100.0	99.2	99.2	99.5	99.6	99.9	99.8	99.7
Llama-3.2-Vision (90B)	95.9	100.0	99.5	98.3	98.4	100.0	100.0	98.3	99.4	100.0	98.8	99.4	99.0
Gemini 2.0 Flash	99.4	100.0	100.0	100.0	99.9	100.0	100.0	98.5	99.5	100.0	100.0	100.0	99.8
GPT-4o	99.2	99.5	98.5	97.0	98.5	100.0	100.0	94.0	98.0	98.5	100.0	99.2	98.5
Spatial VLMs													
SpatialBot-RGB (3B)	99.2	99.0	99.8	98.4	99.1	100.0	100.0	99.0	99.7	100.0	99.5	99.7	99.4
SpatialBot (3B)	99.2	100.0	100.0	100.0	99.8	100.0	99.4	99.3	99.6	100.0	100.0	100.0	99.8
SpaceMantis (8B)	99.4	100.0	100.0	99.7	99.8	100.0	100.0	98.5	99.5	100.0	100.0	100.0	99.7
SpatialRGPT-RGB (8B)	98.9	100.0	100.0	99.1	99.5	100.0	100.0	99.2	99.7	97.9	95.9	96.9	98.9




Table 7: Average validity (%) of VLM responses on SPHERE tasks.

Image	Question	Answer
	If you are the man in the image, where is the girl relative to your position? Left or Right?	Ground-truth: Left GPT-4o: Right ✗ Gemini 2.0 Flash: Right ✗
	From the perspective of the viewer, where is the girl relative to the man? Left or Right?	Ground-truth: Right GPT-4o: Right ✓ Gemini 2.0 Flash: Right ✓
	If you are the man pouring wine, where is the man with glasses relative to your position? Left or Right?	Ground-truth: Left GPT-4o: Right ✗ Gemini 2.0 Flash: Right ✗

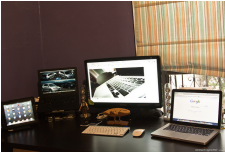


(a) Task: Position

Image	Question	Answer
	How many buses are visible?	Ground-truth: 0 GPT-4o: 0 ✓ Gemini 2.0 Flash: 0 ✓
	How many trees are visible?	Ground-truth: 4 GPT-4o: 2 ✗ Gemini 2.0 Flash: 3 ✗
	How many crows can be seen?	Ground-truth: 3 GPT-4o: 3 ✓ Gemini 2.0 Flash: 3 ✓

(b) Task: Counting



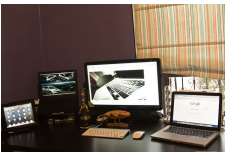
Image	Question	Answer
	Which object is more distant from the plant? Chair or Table?	Ground-truth: Chair GPT-4o: Chair ✓ Gemini 2.0 Flash: Chair ✓
	Which group of people is farther away from the motorcycle? People sitting on the left or People standing on the right?	Ground-truth: People sitting on the left GPT-4o: People standing on the right ✗ Gemini 2.0 Flash: People sitting on the left ✓
	Which boat has a shorter distance to the bank? Boat with a cleaning sign or Boat with passengers?	Ground-truth: Boat with passengers GPT-4o: Boat with a cleaning sign ✗ Gemini 2.0 Flash: Boat with a cleaning sign ✗

(c) Task: Distance




Image	Question	Answer
	Which computer is the smallest? The leftmost one or The rightmost one?	Ground-truth: The leftmost one GPT-4o: The leftmost one ✓ Gemini 2.0 Flash: The leftmost one ✓
	Is the fish icon on the ground bigger than the star icon in the picture? Yes or No?	Ground-truth: Yes GPT-4o: Yes ✓ Gemini 2.0 Flash: Yes ✓
	Which cup is higher? The left one or The right one?	Ground-truth: The right one GPT-4o: The left one ✗ Gemini 2.0 Flash: The left one ✗

(d) Task: Size




Table 8: Examples from single-skill tasks.

Image	Question	Answer
	From the perspective of the viewer, how many laptops are on the left side of the man?	Ground-truth: 2 GPT-4o: 2 ✓ Gemini 2.0 Flash: 1 ✗
	If you are the man in the image, how many laptops are on your left?	Ground-truth: 0 GPT-4o: 2 ✗ Gemini 2.0 Flash: 2 ✗
	How many displays are to the right of the keyboard?	Ground-truth: 1 GPT-4o: 2 ✗ Gemini 2.0 Flash: 3 ✗

(a) Task: Position + Counting




Image	Question	Answer
	From the perspective of the person, how many horses are farther from him than the fence?	Ground-truth: 3 GPT-4o: 3 ✓ Gemini 2.0 Flash: 2 ✗
	From the perspective of the person, how many horses are closer to him than the fence?	Ground-truth: 0 GPT-4o: 1 ✗ Gemini 2.0 Flash: 1 ✗
	How many crows are on the railing farther from the viewer?	Ground-truth: 1 GPT-4o: 3 ✗ Gemini 2.0 Flash: 2 ✗

(b) Task: Distance + Counting


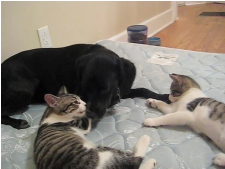

Image	Question	Answer
	Are the turquoise plates likely the same size? Yes or No?	Ground-truth: Yes GPT-4o: Yes ✓ Gemini 2.0 Flash: Yes ✓
	The chairs on the top half of the picture are differently sized? Yes or No?	Ground-truth: No GPT-4o: Yes ✗ Gemini 2.0 Flash: Yes ✗
	Which pillow is bigger? The left one or The right one or They are similar in size?	Ground-truth: They are similar in size GPT-4o: The left one ✗ Gemini 2.0 Flash: The right one ✗

(c) Task: Distance + Size

Table 9: Examples from multi-skill tasks.

Image	Question	Answer
	Where can the pepper shaker most likely be found? Behind the ketchup bottle or Behind the margarita glass?	Ground-truth: Behind the ketchup bottle GPT-4o: Behind the ketchup bottle ✓ Gemini 2.0 Flash: Behind the ketchup bottle ✓
	Where can another teddy bear be found? Inside the suitcase on the white chair or Under the blue dress?	Ground-truth: Inside the suitcase on the white chair GPT-4o: Under the blue dress ✗ Gemini 2.0 Flash: Under the blue dress ✗
	Is it likely that more glasses are hidden behind the pizza boxes? Yes or No?	Ground-truth: No GPT-4o: No ✓ Gemini 2.0 Flash: Yes ✗

(a) Task: Object occlusion

Image	Question	Answer
	From the perspective of the woman in pink, can the man wearing glasses stand on her left side? Yes or No?	Ground-truth: Yes GPT-4o: Yes ✓ Gemini 2.0 Flash: Yes ✓
	Can the dog lie between the two cats? Yes or No?	Ground-truth: No GPT-4o: Yes ✗ Gemini 2.0 Flash: No ✓
	From the perspective of the viewer, can we pour all the wine from the nearer glass into the one behind? Yes or No?	Ground-truth: Yes GPT-4o: Yes ✓ Gemini 2.0 Flash: No ✗

(b) Task: Object manipulation

Table 10: Examples from reasoning tasks.