

VIEW2SPACE: Studying Multi-View Visual Reasoning from Sparse Observations

Fucai Ke¹, Zhixi Cai¹, Boying Li¹, Long Chen¹, Beibei Lin¹, Weiqing Wang¹, Pari Delir Haghighi¹, Gholamreza Haffari¹, and Hamid Rezaatofghi¹

¹ Monash University, Australia

² National University of Singapore, Singapore

Abstract. Multi-view visual reasoning is essential for intelligent systems that must understand complex environments from sparse and discrete viewpoints, yet existing research has largely focused on single-image or temporally dense video settings. In real-world scenarios, reasoning across views requires integrating partial observations without explicit guidance, while collecting large-scale multi-view data with accurate geometric and semantic annotations remains challenging. To address this gap, we leverage physically grounded simulation to construct diverse, high-fidelity 3D scenes with precise per-view metadata, enabling scalable data generation that remains transferable to real-world settings. Based on this engine, we introduce VIEW2SPACE, a multi-dimensional benchmark for sparse multi-view reasoning, together with a scalable, disjoint training split supporting millions of grounded question-answer pairs. Using this benchmark, a comprehensive evaluation of state-of-the-art vision-language and spatial models reveals that multi-view reasoning remains largely unsolved, with most models performing only marginally above random guessing. We further investigate whether training can bridge this gap. Our proposed Grounded Chain-of-Thought with Visual Evidence substantially improves performance under moderate difficulty, and generalizes to real-world data, outperforming existing approaches in cross-dataset evaluation. We further conduct difficulty-aware scaling analyses across model size, data scale, reasoning depth, and visibility constraints, indicating that while geometric perception can benefit from scaling under sufficient visibility, deep compositional reasoning across sparse views remains a fundamental challenge. Code and benchmark are available online³.

Keywords: Spatial Reasoning · Multi-view Visual Reasoning · Cross-view Reasoning · 3D spatial understanding · Spatial Intelligence · Embodied AI · Vision-Language Models

1 Introduction

Visual reasoning in the real world rarely comes from a single glance. When observing complex environments, humans often draw on information gathered from

³ github.com/pokerme7777/VIEW2SPACE



Fig. 1: A single glance is often insufficient for completing real-world tasks. Humans naturally integrate observations across sparse viewpoints to form a shared spatial understanding. Given heterogeneous views (e.g., robotic dog and drone), one can efficiently align what different agents observe and reason across sparse viewpoints.

previous viewpoints or actively seek additional perspectives to support judgment and decision making. Even though each individual view may reveal only partial and biased information, people can still infer scene layout, reason about occluded objects, and understand spatial relationships that are never simultaneously visible [22, 23, 55, 56].

Such multi-view visual reasoning is a foundational capability for intelligent systems operating in complex environments. It enables spatial understanding from sparse, discrete viewpoints, whether acquired by a single observer across locations or collaboratively by multiple observers sharing observations. This setting arises in applications ranging from individual spatial reasoning to collaborative systems such as multi-robot coordination, multi-agent interaction, large-scale multi-camera environments, and autonomous driving, where effective decision making relies on integrating partial visual observations [16, 20, 21, 34–37, 42, 51, 57], as illustrated in Figure 1.

Despite recent progress in vision-language models (VLMs) [2, 10, 12, 14, 26, 29–31], robust visual reasoning has yet to be consistently achieved even in single-image and densely sampled video settings, which have received the most research attention. Multi-view visual reasoning is therefore even less understood, despite its importance for comprehensive spatial understanding and downstream tasks (e.g., robotic planning and autonomous driving). A few recent works (e.g., [15, 55]) touch on this setting, but they are typically evaluated under non-sparse or human-centric viewpoints, relatively small-scale scenes, and tasks that tend to emphasize perception rather than deeper, multi-hop reasoning across views. We attribute these limitations to both the difficulty of collecting large-scale real-world multi-view data with accurate per-view annotations and the tendency of existing research to emphasize spatial reconstruction and basic post-hoc relational queries, rather than deeper, multi-hop reasoning over sparse views.

To meaningfully study sparse multi-view reasoning, it is necessary to establish a controlled evaluation setting with large-scale, high-fidelity scenes and objects, together with precise cross-view geometric and semantic metadata. Such

a setting should enable systematic variation of viewpoint sparsity, reasoning depth, and task complexity, factors that are difficult to reliably manipulate using real-world data alone, while still supporting transfer to real-world scenarios. We therefore adopt physically grounded simulation as a controllable testbed, designed to minimize the synthetic-to-real gap while preserving geometric and semantic fidelity. This enables scalable generation of high-quality supervision for both training and evaluation.

Within this controlled framework, we build a scalable data engine that generates 3 million QA pairs from 2,000 diverse 3D scenes with precise cross-view geometric and semantic metadata for systematic evaluation and large-scale training. Based on this engine, we introduce VIEW2SPACE, a multi-dimensional benchmark for studying multi-view reasoning under sparse viewpoints.

To assess the current state of multi-view reasoning, we conduct a comprehensive evaluation of state-of-the-art vision-language and spatial models. The results show that multi-view reasoning remains largely unsolved, with most models performing only marginally above random guessing, even on relatively simple reasoning questions. This gap motivates an investigation into whether different training strategies can bridge it, comparing direct-answer, standard CoT, and our proposed Grounded CoT with Visual Evidence. Training consistently improves performance, and grounded CoT further amplifies these gains, achieving up to +52% mIoU on challenging visual grounding tasks in VIEW2SPACE, measured on a test set covering varying reasoning difficulty levels. Remarkably, the resulting trained model also generalizes strongly to real-world data: when evaluated on the MINDCUBE [55] without further fine-tuning, it surpasses the officially released trained models by over 9% accuracy.

While these results suggest that multi-view reasoning can be improved under training, a deeper analysis reveals that the problem is far from solved. Scaling across model size, data volume, and task difficulty indicates that although scale improves perceptual performance under sufficient visibility, it exhibits poor scaling efficiency as reasoning complexity increases, suggesting inherent structural limitations in current approaches. In summary, our contributions are fourfold:

1. **Scalable Data Engine and Benchmark.** We introduce a scalable data engine and VIEW2SPACE, a large-scale benchmark for multi-view spatial reasoning with diverse 3D scenes, geometrically varied viewpoints, rich metadata, graded difficulty, and progressively stricter visual-evidence constraints.
2. **Empirical Analysis of Model Behaviour.** We provide a systematic empirical analysis of SOTA models on VIEW2SPACE, revealing performance gaps under increasing evidence constraints, shortcut-based reasoning patterns, and the limitations of explicit thinking in sparse multi-view settings.
3. **Visually Grounded Chain-of-Thought Dataset and Method.** We build a step-wise, visually grounded chain-of-thought training dataset and propose Grounded Chain-of-Thought with Visual Evidence, achieving significant gains and strong synthetic-to-real generalization.
4. **Scaling Analysis of Multi-View Reasoning.** We reveal, through controlled scaling analysis, persistent limitations in deep multi-view reasoning

despite increased data and model capacity, suggesting the need for more compositional visual reasoning approaches.

2 Literature Review

Recent VLMs such as InternVL [10], Molmo [14], Qwen3VL [1], and GPTs [47] have shown promising spatial reasoning abilities through large-scale training. However, their evaluation settings are largely limited to single images, unrelated image collections, or temporally continuous videos, and rarely cover the common scenario of reasoning over sparse, discrete viewpoints of the same scene. Similarly, spatially focused or robotic models (e.g., RoboBrain [27, 49], OpenView [5], SpatialMLLM [52]) aim to enhance spatial understanding, but are typically constrained to single-view inputs or rely on explicit spatial supervision.

Existing spatial reasoning benchmarks [6, 17, 18, 33, 41, 45, 46, 48, 53, 58] predominantly focus on single-view or temporally dense video settings. Only a few recent works [15, 55] explore multi-view scenarios. However, these settings are generally not sparse, often adopt human-centric viewpoints, and primarily emphasize perceptual or low-hop reasoning. Moreover, systematic investigation of model shortcut behaviors or guessability under varying evidence constraints remains limited. In contrast, real-world spatial reasoning requires inferring cross-view relationships directly from visual observations across sparse and heterogeneous viewpoints beyond fixed human perspectives, often involving compositional and multi-hop inference. Our work targets this underexplored regime and shows that although perceptual-level multi-view reasoning can be learned with sufficient data, deeper reasoning across discrete and diverse viewpoints remains challenging (Sec. 4). Detailed benchmark comparisons are provided in the appendix.

Recent studies have explored incorporating explicit 3D information into LLMs and VLMs to enhance spatial reasoning. One line of work directly introduces structured 3D inputs, such as point clouds [7, 24] and depth maps [59], by projecting 3D geometric features into a shared vision–language embedding space [8, 38]. Building on this paradigm, methods such as 3D-R1 [25] and SpatialThinker [3] further employ explicit chain-of-thought supervision and reinforcement learning with spatial rewards to improve multi-step 3D reasoning, often coupled with view selection or scene graph representations.

However, many of these methods [3, 7, 24, 25, 59] either rely on input modalities beyond RGB that are uncommon in everyday settings, or primarily focus on 3D reconstruction rather than semantic understanding. As a result, their applicability to semantic tasks such as visual question answering or visual grounding under realistic RGB-only conditions remains limited. In contrast, we focus on multi-view spatial reasoning directly from readily available RGB observations.

3 Multi-view Data Engine

Our multi-view data engine is designed to produce high-fidelity, physically grounded 3D scenes with precise and verifiable cross-view annotations. It constructs diverse




Fig. 2: Overview of the multi-view data engine. Stage 1: Assets are collected, categorized, and size-normalized, and annotated using OSD-Tag with verification. Stage 2: Given theme configurations defining asset categories and spatial relations, scenes are synthesized by sampling layouts under compositional constraints. Stage 3: Multi-view images are rendered together with physics-based metadata. Stage 4: Spatial relations are selected to automatically generate grounded question-answer pairs.

environments from over *1,000 high-quality photorealistic assets* with multi-level semantic labels, supporting flexible composition across more than *40 everyday themes*. Built on a physically grounded 3D engine [13], it enables scalable scene generation with geometry-consistent supervision across views. In the following, we describe the overall design and generation pipeline (Figure 2).

3D Asset Library and Automatic Annotation Pipeline. To ensure high realism and physical fidelity, we manually curate a large collection of 3D assets, prioritizing high-quality, scan-based objects that closely reflect real-world geometry and appearance. Assets expected to co-occur within the same scene theme are organized into category-level clusters to maintain semantic coherence during scene composition and annotation. To guarantee metric accuracy and spatial validity, asset scales are normalized to real-world proportions. This calibration integrates prior knowledge from LLMs with manual verification, ensuring object dimensions remain physically consistent. All assets are further assigned semantic categories based on functional role and appearance, forming the foundation for structured scene generation and hierarchical annotation.

Overview-driven Semantic Differentiation and Tagging (OSD-Tag). Given the need for a scalable and largely automated yet reliable annotation pipeline, we introduce Overview-driven Semantic Differentiation and Tagging. Using aggregated asset previews within each category as an overview, OSD-Tag leverages a VLM (*i.e.*, GPT-4o [26]) to identify key visual and functional dif-

ferences and construct a hierarchical semantic tag library. These tags are then automatically assigned to each asset, followed by *manual verification* to ensure correctness. The resulting hierarchical annotations support controlled scene composition, question generation, and diagnostic analysis. Implementation details are provided in the appendix.

Thematic Scene Design and Layout Generation. To ensure realistic scene composition while maintaining layout diversity, each scene theme is manually designed with explicit semantic and spatial constraints. Designers define theme-level configurations that govern object categories, placements, and inter-category relationships, enabling coherent yet varied scene generation. These constraints are encoded in theme-specific configuration files and used to sample valid layouts with controlled randomness. Scene generation is implemented in Blender [13], enabling precise geometric control and adjustable density, thereby producing diverse but structurally consistent scenes within each theme.

Multi-View Rendering and Metadata Extraction. To support the diverse requirements of benchmarking, diagnostic analysis, and model training, we design a set of four canonical viewpoint configurations for multi-view rendering, including drone views, bird-eye views, human-like egocentric views, and fixed surveillance camera views. Within each viewpoint category, users can flexibly control camera placement and orientation (*e.g.*, viewing the scene from the center or from peripheral locations, and from different azimuthal directions), as well as camera coverage and optics, such as the proportion of the scene captured and the use of wide-angle versus narrow field-of-view settings. These viewpoints capture scenes from complementary perspectives and introduce systematic variation in viewpoint geometry and visibility.

Using the underlying physics and rendering engines, the system automatically computes rich view-dependent metadata, including object bounding boxes, occlusion ratios, and camera-relative visibility statistics (see figure in the appendix). We further record precise geometric information for each object, such as 3D pose, orientation, and spatial relationships, which are stored alongside rendered images. This comprehensive metadata enables fine-grained evaluation, controlled diagnostic studies, and the generation of structured supervision for multi-view visual reasoning.

Grounded Multi-View Question and Supervision Generation. We construct question templates spanning multiple reasoning difficulty levels to promote diversity while maintaining answer reliability. Answers are then deterministically derived through rule-based mapping from scene geometry, avoiding heuristic or generative inference. Using controllable 3D scenes and engine-provided ground-truth metadata, including object identities, 3D poses, camera parameters, view-dependent bounding boxes, and occlusion statistics, we compute spatial relations in both object-centric and camera-centric coordinate frames. This process produces strictly grounded multi-view question-answer pairs and structured supervision such as instance-level grounding, explicit inter-view relations, and geometry-derived reasoning traces. Additional implementation details are provided in the appendix.




Fig. 3: Overview of VIEW2SPACE task design. VIEW2SPACE comprises three task types: MCQ, counting, and visual grounding detection. Reasoning difficulty ranges from one-hop perceptual queries to multi-hop cross-view spatial reasoning.

4 VIEW2SPACE Benchmark and Evaluation

VIEW2SPACE is designed to study multi-view spatial reasoning under sparse and geometrically diverse viewpoints, where no single view is sufficient to resolve the query. Beyond simply providing multiple images, the benchmark covers a graded spectrum of reasoning difficulty, from basic cross-view alignment to high-dimensional, multi-hop compositional reasoning, and includes task formats with increasingly strict visual-evidence constraints (MCQ, counting, and detection).

Empirically, most models perform only marginally above random guessing even on multiple-choice questions, which are the most permissive and tolerant to informed guessing. When explicit visual grounding is required, performance drops further, with detection tasks yielding near-zero results. This pattern indicates that current models struggle to integrate information across sparse viewpoints and fail under stricter grounding constraints.

4.1 VIEW2SPACE Benchmark

Overview We introduce VIEW2SPACE, a benchmark for evaluating multi-view spatial reasoning under sparse viewpoints. The benchmark tests whether models can integrate incomplete cross-view observations into a coherent spatial representation as shown in Figure 3. VIEW2SPACE comprises groups of images capturing the same static scene from different viewpoints, paired with reasoning questions whose answers require cross-view inference. In contrast to prior benchmarks [4, 11, 39, 44, 53, 55], which either focus on single-image understanding or continuous video streams, or involve human-centric viewpoints with limited layout diversity and predominantly low-hop perceptual queries, VIEW2SPACE is designed to evaluate deep compositional reasoning across sparse and geometrically diverse viewpoints.

Task Types and Evaluation Metric. To our knowledge, VIEW2SPACE is the first benchmark that systematically spans tasks from highly permissive answer spaces to strictly constrained grounding-based detection. It includes three task types: multiple-choice questioning (MCQ), counting, and visual grounding detection, that progressively restrict the feasible answer space and reduce the likelihood of correct predictions obtained through guessing or unsupported reasoning. Rather than differing only in output structure, these tasks reflect increasing resistance to correct answers obtained through linguistic bias or hallucinated reasoning. MCQ is the most permissive, allowing correctness via partial evidence or informed guessing, and is evaluated using accuracy (ACC). Counting requires preliminary grounding of relevant instances before aggregation, narrowing plausible guesses, and is evaluated using ACC and mean absolute error (MAE). Visual grounding detection is the most restrictive, demanding explicit localization of target objects and is evaluated using mIoU and F1-score (IoU > 50%), making unsupported yet correct answers least likely.

Reasoning Difficulty. Within each task type, questions are organized along a graded reasoning difficulty spectrum, ranging from basic perceptual alignment across views to compositional, multi-hop spatial reasoning. Difficulty is defined by the minimal number of relational hops required in a per-view object–relation graph, where objects are nodes and spatial relations are edges. A single perceptual grounding step counts as one hop, and complexity further increases with the logarithm of the number of key objects involved. Detailed reasoning difficulty settings are provided in the Appendix.

Object Visibility Difficulty. We define it based on the proportion of an object’s projected area on the image plane that is not directly visible to the camera. Using the underlying Blender rendering engine, we perform ray casting from object surface points toward the camera to determine whether each point is directly observable. The fraction of rays that fail to reach the camera due to occlusion is taken as the visibility difficulty score. For each task, we define *key-object visibility difficulty* as the average visibility difficulty of all task-relevant objects aggregated across views.

Viewpoint Configurations. VIEW2SPACE systematically varies viewpoint pairings to probe spatial reasoning under heterogeneous perspectives, as shown in the *Different Viewpoints* block in Figure 2. These pairings include combinations of human-like egocentric views, drone views, bird’s-eye views, and fixed surveillance camera views. By combining viewpoints with distinct geometric properties, the benchmark creates controlled variations in viewpoint geometry and visibility, enabling evaluation of reasoning under diverse cross-view conditions.

Statistics of VIEW2SPACE. We introduce VIEW2SPACE-v1, a curated evaluation split comprising 3,591 questions (1,400 MCQ, 591 Counting, and 1,600 Detection), with balanced answer distributions across 40 themes and 1,096 primary 3D assets spanning graded difficulty levels. As sparse multi-view reasoning is already challenging for current VLMs, even lower-difficulty cases remain non-trivial; thus, we maintain a graded distribution rather than concentrating on extreme difficulty. Although VIEW2SPACE scales to **3 million sam-**

Table 1: Comparison on VIEW2SPACE-v1

Method	Multiple Choice Answering ACC (%)	Visual Counting		Visual Grounding	
		MAE	ACC (%)	mIoU (%)	F1-score (%)
Baseline					
Random (chance)	28.59	11.37	3.39	0.22	0.0
Random (frequency)	28.59	2.46	13.05	2.01	0.27
Open-source MLLMs					
Intervl2.5-8B [9]	30.93	2.04	22.50	0.79	0.0
Mantis-8B (SigLip) [28]	30.07	2.94	17.60	1.24	0.23
DeepSeek-VL2-Small [40]	29.43	2.20	24.37	2.86	0.19
Gemma-3-12B-it [50]	31.14	2.08	24.53	2.49	0.24
Idefics2-8B [32]	26.64	3.92	11.34	0.39	0.0
Qwen2.5-VL-7B [2]	30.29	2.70	13.03	2.68	0.32
Qwen3-VL-2B [1]	31.43	2.63	16.34	11.20	5.07
Qwen3-VL-4B [1]	35.19	2.26	21.15	16.53	18.72
Qwen3-VL-8B [1]	37.41	2.10	24.92	10.05	7.37
Open-source Spatial Models					
Molmo2-4B [12]	31.36	2.21	22.17	1.17	0.13
Molmo2-8B [12]	34.57	5.93	21.66	1.51	0.0
RoBoBrain2.0 [49]	30.79	2.22	21.49	0.0	0.0
Spatial-MLLM [52]	27.93	2.10	9.98	0.0	0.0
Openview [5]	30.72	2.57	19.80	1.08	0.30
MINDCUBE-Aug [55]	29.00	4.52	18.10	0.0	0.0
MINDCUBE-Plain [55]	30.21	4.77	19.12	0.12	0.0
Closed-source MLLMs					
GPT-4o [26]	33.93	1.51	29.03	3.60	0.56
GPT-5-nano [47]	33.36	1.73	24.28	2.89	0.87
GPT-5-mini [47]	49.00	1.20	38.71	7.42	2.77
GPT-5 [47]	59.86	1.31	38.10	8.18	3.43
GPT-5.2 [43]	37.07	1.28	36.50	8.74	2.84
Fine-tuned Models (Qwen3VL-4B)					
Ours (Instruct-tuning)	36.50	0.71	50.08	47.58	50.07
Ours (Grounded CoT, Direct Answer at Test)	33.95	10.75	29.10	16.35	11.50
Ours (CoT)	59.57	0.68	51.43	49.77	52.34
Ours (Grounded CoT)	64.93	0.58	54.99	69.34	70.92
$\Delta(\text{abs})$ / $\Delta(\%)$	$\uparrow 5.07$ / 8.47	$\downarrow 0.62$ / 51.67	$\uparrow 16.28$ / 42.06	$\uparrow 52.81$ / 319.48	$\uparrow 52.2$ / 278.85

ples, large-scale evaluation, particularly on closed-source models, incurs substantial computational and API costs. To balance evaluation cost (approximately \$100 per closed-source model) and reproducibility, we benchmark the 3k-sample VIEW2SPACE-v1 in this work. Detailed statistics are provided in the Appendix.

A Scalable Training Split for VIEW2SPACE. We construct a scalable training dataset with a size of 300K with scenes disjoint from the evaluation split, supporting expansion to millions of question-answer pairs. Each instance is derived from engine-level geometric and physical ground truth, including bounding boxes, occlusion ratios, and precise cross-view spatial relations. Via a symbolic pipeline, we generate deterministic, visually grounded chain-of-thought supervision with verifiable intermediate steps, eliminating the stochastic noise and inconsistency of MLLMs generated reasoning traces and yielding clean, geometry-consistent training signals. Additional details are provided in the appendix.

4.2 Evaluation on VIEW2SPACE

We evaluate models on VIEW2SPACE-v1, a difficulty-balanced split with a relatively higher proportion of lower-difficulty cases. Since sparse multi-view reasoning remains underexplored and most models struggle even on basic multi-view and grounding tasks, this design avoids over-concentrating evaluation on only the hardest instances. The evaluated models include open-source MLLMs [1, 2, 9, 28, 32, 40, 50], specialized spatial reasoning models [5, 12, 27, 52, 55] with explicit fine-tuning on spatial reasoning dataset, as well as closed-source MLLMs [26, 43, 47].

An overview of all evaluated models is provided in Table 1. As reference, we include two baselines: Random (chance), which samples uniformly, and Random (frequency), which follows the empirical answer distribution, estimating lower-bound performance under different prior knowledge assumptions.

Performance Across Tasks with Shrinking Answer Space. To our knowledge, VIEW2SPACE is the first benchmark that spans tasks ranging from those more tolerant to hallucinated reasoning to those least tolerant, progressively enforcing stricter visual-evidence constraints and tighter solution spaces.

Model performance consistently declines as visual-evidence constraints tighten, closely tracking the reduction in answer space from left to right in Table 1. Under strict grounding constraints, most models perform near random baselines, indicating limited grounding capability and reflecting the current capability profile of VLMs, which remain stronger in text-centric generation and selection than in precise localization. Importantly, this does not imply that grounding is unattainable. Recent efforts increasingly incorporate grounding-oriented training into VLMs (*i.e.*, [1,9,12]). Notably, Qwen3VL [1] performs measurably above baseline despite low absolute accuracy; consistent with its reported grounding-focused training, this suggests that grounding robustness is achievable with explicit supervision. The GPTs [26,43,47] likewise show steady improvements in grounding performance across versions, aligning with their reported enhancements. In contrast, although Molmo2 [12] also emphasizes grounding, its performance remains limited under sparse multi-view constraints, indicating that grounding learned in conventional settings may not readily transfer to multi-view reasoning.

Comparative Performance of State-of-the-Art Models. Table 1 shows that most models retain substantial headroom even on the most hallucination-tolerant MCQ tasks. We do not interpret this as an inability to perform multi-view reasoning *per se*, but rather as a limited transfer of spatial skills learned in conventional settings to sparse multi-view scenarios without dedicated training.

Closed-source GPTs [26,47] exhibit relatively stable performance; for instance, GPT-5 [47] achieves 59.86% accuracy on MCQ. This may relate to larger model scale and exposure to diverse perceptual and reasoning data, potentially including 3D-derived or perception-heavy VQA-style supervision. Interestingly, GPT-5.2 [43] does not outperform GPT-5 [47] or GPT-5-mini, and performs closer to GPT-4o [26] on MCQ. While GPT-5.2 [43] reports improvements in grounding and an increased emphasis on coding, these enhancements do not translate into clear gains in our multi-view setting, suggesting that fine-tuning objectives may not strongly target sparse multi-view spatial alignment.

Open-source spatial models do not show consistent advantages over general open-source MLLMs, and generally underperform the Qwen3-VL series [1] as shown in Table 1. This suggests that spatial capabilities developed through robotic spatial training or static single-view supervision may have limited transfer to discrete multi-view reasoning. While single-view spatial understanding remains fundamental, multi-view reasoning additionally requires cross-view alignment and geometric calibration, which likely accounts for the observed gap.

Finally, as reasoning difficulty progressively increases within the dataset, overall performance trends also reflect the structured escalation of task complexity.

Thinking vs. Non-Thinking Inference. In the main experiments, models were prompted without restricting reasoning style. We additionally compare thinking (CoT) and direct-answer inference for open-source MLLMs (Appendix). As grounding performance is near zero for most of these models, we focus on counting and MCQ tasks. Explicit thinking does not improve performance and may degrade accuracy. Manual inspection suggests that in sparse multi-view settings, direct answers may exploit shortcut cues. In contrast, enforced thinking often produces repetitive or weakly grounded reasoning, extending the reasoning trace without effectively integrating cross-view evidence and ultimately leading to incorrect conclusions. This indicates that explicit reasoning without sufficient multi-view training may introduce instability rather than gains.

These findings raise a central question: **can sufficient data and appropriate training enable models to acquire spatial reasoning for sparse multi-view settings?** We investigate this by training on VIEW2SPACE.

5 From Data to Multi-View Reasoning: What Can Training Achieve?

In this section, we investigate how different training strategies affect multi-view visual reasoning. We compare standard instruction tuning, language-only chain-of-thought training, and our proposed Grounded Chain-of-Thought with Visual Evidence to understand the role of explicit visual grounding. We further examine synthetic-to-real generalization and analyze how performance scales with model size, data volume, reasoning difficulty, and visibility constraints. Together, these experiments aim to clarify what training can and cannot achieve for sparse multi-view reasoning.

5.1 Supervised Fine-Tuning for Multi-View Reasoning

In this subsection, we study how supervised fine-tuning enables multi-view visual reasoning, examining its effectiveness under different supervision strategies.

Training Setup. We adopt Qwen3-VL-4B-Instruct [1] as the base model and perform full-parameter supervised fine-tuning under three strategies: instruction tuning, standard CoT, and our proposed grounded CoT with visual evidence. Training samples at the 300K scale are generated by our data engine. Additional implementation details are provided in the appendix.

Limited Gains from Instruction Tuning. We evaluate the effectiveness of monolithic instruction tuning for multi-view visual reasoning as shown in Table 1. While instruction tuning consistently outperforms random guessing on VIEW2SPACE-v1, the overall gains remain limited. In particular, performance on MCQ tasks (36.50% ACC) falls well below that of GPT-5 [47] (59.86% ACC), indicating that instruction tuning alone is insufficient for learning deeper cross-view reasoning patterns.

Instruction tuning yields relatively stronger improvements on counting and visual grounding tasks, where supervision is directly anchored to visual evidence. Consistent with our analysis in Sec. 4.2, we attribute this to task familiarity: these tasks align more closely with the types of visually grounded supervision that the community has increasingly emphasized, and instruction tuning provides missing multi-view training signals that help models develop initial competence in this unfamiliar setting. These results suggest that instruction tuning alone is insufficient for learning structured cross-view relational reasoning, motivating the exploration of more explicitly grounded training strategies in the following subsection.

CoT Improves MCQ but Not Grounding-Intensive Tasks. As shown in Table 1, standard CoT reasoning yields substantial improvements on MCQ tasks, increasing accuracy from 36.50% under instruction tuning to 59.57%, approaching the performance of GPT-5 [47] on the same setting. However, compared to instruction tuning, improvements on counting and visual grounding tasks remain limited, despite their stronger reliance on localization ability. This suggests that language-driven reasoning alone may not sufficiently compensate for missing visual grounding capability, indicating a bottleneck when grounding is required.

Grounded Chain-of-Thought Enables Broad Improvements. Beyond standard instruction tuning, we propose a grounded CoT training strategy that explicitly ties intermediate reasoning steps to visual evidence (explicitly output key object bounding box in the corresponding view). Instead of relying on free-form textual reasoning, the model is encouraged to reference visually grounded entities and spatial relations when forming intermediate conclusions.

This design reduces ungrounded linguistic reasoning and promotes reasoning anchored in observable visual cues across multiple views. Empirically, grounded CoT leads to consistent and substantial improvements across tasks. On MCQ, it achieves over 77% relative accuracy improvement compared to standard instruction tuning. More notably, on hallucination-sensitive visual grounding tasks, grounded CoT yields over 300% relative improvement in detection mIoU, indicating significantly stronger evidence-grounded reasoning. These results demonstrate that explicitly grounding intermediate reasoning steps in visual evidence is critical for robust multi-view reasoning beyond surface-level pattern matching.

These results suggest that explicitly incorporating visual evidence into intermediate reasoning steps is essential for robust multi-view reasoning, especially in settings where correct answers cannot be obtained through linguistic shortcuts.

5.2 Synthetic-to-Real Generalization

Beyond in-domain evaluation, we assess transfer to real-world data. We evaluate on MindCube-Tiny [55], a real-world multi-view dataset, following the official evaluation protocol. Compared to our benchmark, MindCube frequently encodes inter-view relations through language descriptions, employs a relatively limited set of viewpoint configurations, and involves simpler object-relation structures.

Table 2: Comparison on MINDCUBE-Tiny. The second column indicates whether models are fine-tuned on MINDCUBE (w Tuning vs. w/o Tuning).

Method	Tuning on MINDCUBE	Overall	Rotation Among Around		
Baseline					
Random (chance)	-	32.35	36.36	32.29	30.66
Random (frequency)	-	33.02	38.30	32.66	35.79
Open-source MLLMs					
Internvl2.5-8B [9]	w/o Tuning	18.68	36.45	47.11	26.91
Molmo2-4B [12]	w/o Tuning	37.33	36.00	38.83	35.75
Molmo2-8B [12]	w/o Tuning	34.33	27.50	36.50	34.50
Mantis-8B (SigLip)	w/o Tuning	41.05	37.65	40.23	50.99
DeepSeek-VL2-Small	w/o Tuning	47.62	37.00	50.38	26.91
Gemma-3-12B-it	w/o Tuning	46.67	38.39	48.38	34.63
Idedics2-8B	w/o Tuning	35.86	35.15	35.94	35.49
Qwen2.5-VL-7B [2]	w/o Tuning	29.26	38.76	29.50	21.35
Qwen3-VL-2B [1]	w/o Tuning	32.58	26.50	36.33	30.00
Qwen3-VL-4B [1]	w/o Tuning	36.67	41.50	34.67	37.25
Qwen3-VL-8B [1]	w/o Tuning	37.42	46.00	35.83	35.50
RoBoBrain [27]	w/o Tuning	37.38	35.80	38.28	29.53
Spatial-MLLM [52]	w/o Tuning	32.06	38.39	20.92	32.82
Openview [5]	w/o Tuning	40.08	24.50	45.50	39.75
Closed-source MLLMs					
Claude-4-Sonnet	w/o Tuning	44.75	48.42	44.21	47.62
GPT-4o	w/o Tuning	38.81	32.65	40.17	29.16
GPT-5-nano	w/o Tuning	38.58	44.00	38.67	35.75
GPT-5-mini	w/o Tuning	46.64	89.00	44.67	51.75
MINDCUBE-Checkpoints					
MINDCUBE-Aug [55]	w Tuning	55.24	49.50	52.50	66.40
MINDCUBE-Plain [55]	w Tuning	60.76	47.50	62.33	67.60
VIEW2SPACE-Trained Models					
Ours 2B (Grounded-CoT)	w/o Tuning	59.92	49.00	56.50	70.50
Ours 4B (Grounded-CoT)	w/o Tuning	70.00	71.50	66.33	74.75
$\Delta(\text{abs})$ / $\Delta(\%)$		$\uparrow$ 9.24 / 15.21			

Moreover, it contains only MCQ without grounding-based evaluation. Results are shown in Table 2.

To assess cross-domain generalization, we compare models trained solely on VIEW2SPACE with models fine-tuned directly on the MindCube training set. Despite not being trained on any MindCube data, VIEW2SPACE-trained models outperform the official MindCube checkpoints by 9.24% absolute accuracy and 15.21% relative gain. These results demonstrate that the supervision provided by VIEW2SPACE generalizes effectively across domains, even when transferring from synthetic sparse-view settings to real-world data.

5.3 Limits of Scaling in Multi-View Reasoning

Despite strong overall performance and encouraging synthetic-to-real transfer, multi-view reasoning remains far from solved. VIEW2SPACE-v1 spans a range of difficulty from perceptual alignment to multi-hop reasoning; given the novelty of sparse multi-view tasks, a substantial portion remains perceptual or low-hop. As a result, aggregate metrics average across difficulty levels and can obscure model performance on genuinely complex cases. To probe the limits of training and assess the effects of dataset and model scale on multi-view reasoning, we conduct controlled experiments on a balanced 3K-sample analysis set. We fine-tune Qwen3-VL-2B/4B [1] with increasing training data (1K–300K) and evaluate




Fig. 4: Scaling Behaviour across Reasoning and Key-Object Visibility Difficulty (see Sec. 4.1). Each subplot plots mIoU versus task difficulty. Each curve corresponds to a specific dataset size, fitted over the averaged performance at each difficulty level (aggregated across multiple samples and weighted accordingly). The dense empirical dot plots and detailed results are provided in the appendix.

performance across reasoning and key-object visibility difficulty (Figure 4). To isolate factors, reasoning analysis restricts visibility difficulty below 0.5, while visibility analysis restricts reasoning difficulty below 6. Additional details are provided in the appendix.

Our analysis shows that while scaling dataset and model size improves multi-view perceptual capability and yields strong performance under sufficient visibility, it exhibits poor scaling efficiency as reasoning complexity increases. As illustrated in Figure 4, training improvements span roughly 70% of the visibility spectrum shown, whereas diminishing returns already emerge within only a tiny portion of the effectively unbounded reasoning complexity spectrum. This suggests that scaling may raise performance ceilings without removing underlying structural constraints.

Visibility difficulty reflects a learnable perceptual regime with an observability limit. Performance declines gradually across a broad range of increasing visibility difficulty, indicating that current methods can learn effective geometric and perceptual representations when visual evidence remains sufficient, as shown in the left region of Figure 4(c,d). However, once visibility crosses a critical threshold around 0.7, accuracy drops sharply across settings. This abrupt collapse is consistent with intuition: when key visual evidence is largely unobservable, answering correctly becomes inherently more difficult.

Reasoning difficulty drives sustained degradation and poor scaling efficiency. Despite potential noise in our difficulty metric, the trend is clear: accuracy declines steadily as reasoning difficulty increases, and gains diminish

beyond moderate data scale, especially beyond 50K. Because compositional complexity grows rapidly through simple combinations and permutations and can extend far beyond the range shown here, scaling data and model size quickly becomes inefficient. One possible explanation is that current fine-tuning strengthens reasoning chains but struggles with deeper compositional search. Multi-hop reasoning across sparse viewpoints resembles structured graph search that requires exploring alternative relational hypotheses and maintaining cross-view consistency, whereas CoT follows a single linear trajectory that limits exploration and accumulates early errors. These findings suggest that high-difficulty multi-view reasoning may require mechanisms beyond linear CoT, such as structured or tree-based search (*e.g.* [54]) to better approximate graph-like relational inference (*e.g.* [19]).

6 Conclusion

This work studies multi-view visual reasoning in collaborative settings, where no single observation is sufficient and models must integrate information across discrete viewpoints to form a shared spatial understanding. To support systematic evaluation and scalable training, we develop an automated multi-view data engine with rich geometric and view-consistent annotations, and introduce the VIEW2SPACE benchmark spanning multiple task types, viewpoint configurations, and graded reasoning difficulties. Empirical results show that, while current vision-language models can learn perceptual-level cross-view skills with sufficient data, performance remains limited for high-difficulty multi-hop reasoning and high-occlusion cases, indicating that compositional multi-view reasoning remains an open challenge.

Acknowledgements

This work was supported by the ONR Global X-Challenge Grant (N62909-25-1-2067), and partially supported by DARPA ANSR (FA8750-23-2-1016) and DECRA (DE250100032).

References

1. Bai, S., Cai, Y., Chen, R., Chen, K., Chen, X., Cheng, Z., Deng, L., Ding, W., Gao, C., Ge, C., Ge, W., Guo, Z., Huang, Q., Huang, J., Huang, F., Hui, B., Jiang, S., Li, Z., Li, M., Li, M., Li, K., Lin, Z., Lin, J., Liu, X., Liu, J., Liu, C., Liu, Y., Liu, D., Liu, S., Lu, D., Luo, R., Lv, C., Men, R., Meng, L., Ren, X., Ren, X., Song, S., Sun, Y., Tang, J., Tu, J., Wan, J., Wang, P., Wang, P., Wang, Q., Wang, Y., Xie, T., Xu, Y., Xu, H., Xu, J., Yang, Z., Yang, M., Yang, J., Yang, A., Yu, B., Zhang, F., Zhang, H., Zhang, X., Zheng, B., Zhong, H., Zhou, J., Zhou, F., Zhou, J., Zhu, Y., Zhu, K.: Qwen3-vl technical report. arXiv preprint arXiv:2511.21631 (2025)
2. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2. 5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)

3. Batra, H., Tu, H., Chen, H., Lin, Y., Xie, C., Clark, R.: Spatialthinker: Reinforcing 3d reasoning in multimodal llms via spatial rewards. arXiv preprint arXiv:2511.07403 (2025)
4. Chen, B., Xu, Z., Kirmani, S., Ichter, B., Sadigh, D., Guibas, L., Xia, F.: Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14455–14465 (2024)
5. Chen, Q., Zhang, C., Fu, C.W., Ye, J., Cai, J.: Openview: Empowering mllms with out-of-view vqa. arXiv preprint arXiv:2512.18563 (2025)
6. Chen, S., Zhu, T., Zhou, R., Zhang, J., Gao, S., Niebles, J.C., Geva, M., He, J., Wu, J., Li, M.: Why is spatial reasoning hard for vlms? an attention mechanism perspective on focus areas. arXiv preprint arXiv:2503.01773 (2025)
7. Chen, S., Chen, X., Zhang, C., Li, M., Yu, G., Fei, H., Zhu, H., Fan, J., Chen, T.: Ll3da: Visual interactive instruction tuning for omni-3d understanding reasoning and planning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 26428–26438 (2024)
8. Chen, Z., Zhang, M., Yu, X., Luo, X., Sun, M., Pan, Z., An, X., Feng, Y., Pei, P., Cai, X., Huang, R.: Think with 3d: Geometric imagination grounded spatial reasoning from limited views. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2613–2624 (June 2026)
9. Chen, Z., Wang, W., Cao, Y., Liu, Y., Gao, Z., Cui, E., Zhu, J., Ye, S., Tian, H., Liu, Z., et al.: Expanding performance boundaries of open-source multimodal models with model, data, and test-time scaling. arXiv preprint arXiv:2412.05271 (2024)
10. Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al.: Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 24185–24198 (2024)
11. Cheng, A.C., Yin, H., Fu, Y., Guo, Q., Yang, R., Kautz, J., Wang, X., Liu, S.: Spatialrgpt: Grounded spatial reasoning in vision-language models. *Advances in Neural Information Processing Systems* **37**, 135062–135093 (2024)
12. Clark, C., Zhang, J., Ma, Z., Park, J.S., Salehi, M., Tripathi, R., Lee, S., Ren, Z., Kim, C.D., Yang, Y., et al.: Molmo2: Open weights and data for vision-language models with video understanding and grounding. arXiv preprint arXiv:2601.10611 (2026)
13. Community, B.O.: Blender-a 3D modelling and rendering package. Blender Foundation, Blender Institute, Amsterdam (2018)
14. Deitke, M., Clark, C., Lee, S., Tripathi, R., Yang, Y., Park, J.S., Salehi, M., Muenighoff, N., Lo, K., Soldaini, L., et al.: Molmo and pixmo: Open weights and open data for state-of-the-art vision-language models. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 91–104 (2025)
15. Deng, N., Gu, L., Ye, S., He, Y., Chen, Z., Li, S., Wang, H., Wei, X., Yang, T., Dou, M., et al.: Internspatial: A comprehensive dataset for spatial reasoning in vision-language models. arXiv preprint arXiv:2506.18385 (2025)
16. Du, W., Yuan, Z., Chen, T., Ke, F., Lin, B., Zhang, S.: Weatherreasonseg: A benchmark for weather-aware reasoning segmentation in visual language models. arXiv preprint arXiv:2603.17680 (2026)
17. Fu, X., Hu, Y., Li, B., Feng, Y., Wang, H., Lin, X., Roth, D., Smith, N.A., Ma, W.C., Krishna, R.: Blink: Multimodal large language models can see but not perceive. In: European Conference on Computer Vision. pp. 148–166. Springer (2024)

18. Gao, Y., Li, H., Liu, Y., Ji, X., Gong, Y., Liao, Y., Liu, F., Zhang, M., Yang, Y., Xu, D., et al.: Holi-spatial: Evolving video streams into holistic 3d spatial intelligence. arXiv preprint arXiv:2603.07660 (2026)
19. Gardenfors, P.: Conceptual spaces: The geometry of thought. MIT press (2004)
20. Gawel, A., Del Don, C., Siegwart, R., Nieto, J., Cadena, C.: X-view: Graph-based semantic multi-view localization. *IEEE Robotics and Automation Letters* **3**(3), 1687–1694 (2018)
21. Ge, H., Gao, D., Sun, L., Hou, Y., Yu, C., Wang, Y., Tan, G.: Multi-agent transfer reinforcement learning with multi-view encoder for adaptive traffic signal control. *IEEE Transactions on Intelligent Transportation Systems* **23**(8), 12572–12587 (2021)
22. Gholami, M., Rezaei, A., Weimin, Z., Mao, S., Zhou, S., Zhang, Y., Akbari, M.: Spatial reasoning with vision-language models in ego-centric multi-view scenes. arXiv preprint arXiv:2509.06266 (2025)
23. Hong, Y., Lin, C., Du, Y., Chen, Z., Tenenbaum, J.B., Gan, C.: 3d concept learning and reasoning from multi-view images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 9202–9212 (2023)
24. Huang, J., Yong, S., Ma, X., Linghu, X., Li, P., Wang, Y., Li, Q., Zhu, S.C., Jia, B., Huang, S.: An embodied generalist agent in 3d world. arXiv preprint arXiv:2311.12871 (2023)
25. Huang, T., Zhang, Z., Tang, H.: 3d-r1: Enhancing reasoning in 3d vlms for unified scene understanding. arXiv preprint arXiv:2507.23478 (2025)
26. Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. arXiv preprint arXiv:2410.21276 (2024)
27. Ji, Y., Tan, H., Shi, J., Hao, X., Zhang, Y., Zhang, H., Wang, P., Zhao, M., Mu, Y., An, P., et al.: Robobrain: A unified brain model for robotic manipulation from abstract to concrete. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 1724–1734 (2025)
28. Jiang, D., He, X., Zeng, H., Wei, C., Ku, M., Liu, Q., Chen, W.: Mantis: Interleaved multi-image instruction tuning. *Transactions on Machine Learning Research* (2024)
29. Ke, F., Cai, Z., Jahangard, S., Wang, W., Haghighi, P.D., Rezatofghi, H.: Hydra: A hyper agent for dynamic compositional visual reasoning. In: *European Conference on Computer Vision*. pp. 132–149. Springer (2024)
30. Ke, F., G, V.K.B., Leng, X., Cai, Z., Khan, Z., Wang, W., Haghighi, P.D., Rezatofghi, H., Chandraker, M.: Dwim: Towards tool-aware visual reasoning via discrepancy-aware workflow generation & instruct-masking tuning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 3378–3389 (October 2025)
31. Ke, F., Hsu, J., Cai, Z., Ma, Z., Zheng, X., Wu, X., Huang, S., Wang, W., Haghighi, P.D., Haffari, G., et al.: Explain before you answer: A survey on compositional visual reasoning. arXiv preprint arXiv:2508.17298 (2025)
32. Laurençon, H., Tronchon, L., Cord, M., Sanh, V.: What matters when building vision-language models? *Advances in Neural Information Processing Systems* **37**, 87874–87907 (2024)
33. Lee, P.Y., Je, J., Park, C., Uy, M.A., Guibas, L., Sung, M.: Perspective-aware reasoning in vision-language models via mental imagery simulation. arXiv preprint arXiv:2504.17207 (2025)
34. Li, B., Cai, Z., Li, Y.F., Reid, I., Rezatofghi, H.: Hier-slam: Scaling-up semantics in slam with a hierarchically categorical gaussian splatting. In: *2025 IEEE Inter-*

- national Conference on Robotics and Automation (ICRA). pp. 9748–9754. IEEE (2025)
35. Li, B., Huang, Y., Liu, Z., Zou, D., Yu, W.: Structdepth: Leveraging the structural regularities for self-supervised indoor depth estimation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12663–12673 (2021)
 36. Li, B., Zou, D., Huang, Y., Niu, X., Pei, L., Yu, W.: Textslam: Visual slam with semantic planar text features. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **46**(1), 593–610 (2023)
 37. Li, B., Zou, D., Sartori, D., Pei, L., Yu, W.: Textslam: Visual slam with planar text features. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 2102–2108. IEEE (2020)
 38. Li, H., Cao, Q., Tang, T., Xiang, K., Guo, Z., Han, J., Xu, H., Bian, J., Liang, X.: Thinking with geometry: Active geometry integration for spatial reasoning. *arXiv preprint arXiv:2602.06037* (2026)
 39. Liu, F., Emerson, G., Collier, N.: Visual spatial reasoning. *Transactions of the Association for Computational Linguistics* **11**, 635–651 (2023)
 40. Lu, H., Liu, W., Zhang, B., Wang, B., Dong, K., Liu, B., Sun, J., Ren, T., Li, Z., Yang, H., et al.: Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525* (2024)
 41. Ma, W., Chen, H., Zhang, G., Chou, Y.C., Chen, J., de Melo, C., Yuille, A.: 3dsr-bench: A comprehensive 3d spatial reasoning benchmark. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6924–6934 (2025)
 42. Ma, Y., Ye, W., Cui, C., Zhang, H., Xing, S., Ke, F., Wang, J., Miao, C., Chen, J., Rezaatofghi, H., et al.: Position: Prospective of autonomous driving-multimodal llms world models embodied intelligence ai alignment and mamba. In: Proceedings of the Winter Conference on Applications of Computer Vision. pp. 1010–1026 (2025)
 43. OpenAI: Introducing gpt-5.2. <https://openai.com/index/introducing-gpt-5-2/> (Dec 2025), accessed: 2025-12-30
 44. Pothiraj, A., Stengel-Eskin, E., Cho, J., Bansal, M.: Capture: Evaluating spatial reasoning in vision language models via occluded object counting. *arXiv preprint arXiv:2504.15485* (2025)
 45. Qi, J., Liu, J., Tang, H., Zhu, Z.: Beyond semantics: Rediscovering spatial awareness in vision-language models. *arXiv preprint arXiv:2503.17349* (2025)
 46. Ramakrishnan, S.K., Wijmans, E., Kraehenbuehl, P., Koltun, V.: Does spatial cognition emerge in frontier models? *arXiv preprint arXiv:2410.06468* (2024)
 47. Singh, A., Fry, A., Perelman, A., Tart, A., Ganesh, A., El-Kishky, A., McLaughlin, A., Low, A., Ostrow, A., Ananthram, A., et al.: Openai gpt-5 system card. *arXiv preprint arXiv:2601.03267* (2025)
 48. Tang, Y., Qu, A., Wang, Z., Zhuang, D., Wu, Z., Ma, W., Wang, S., Zheng, Y., Zhao, Z., Zhao, J.: Sparkle: Mastering basic spatial capabilities in vision language models elicits generalization to spatial reasoning. In: The First Workshop on Multimodal Knowledge and Language Modeling (2025)
 49. Team, B.R., Cao, M., Tan, H., Ji, Y., Chen, X., Lin, M., Li, Z., Cao, Z., Wang, P., Zhou, E., et al.: Robobrain 2.0 technical report. *arXiv preprint arXiv:2507.02029* (2025)
 50. Team, G., Kamath, A., Ferret, J., Pathak, S., Vieillard, N., Merhej, R., Perrin, S., Matejovicova, T., Ramé, A., Rivière, M., et al.: Gemma 3 technical report. *arXiv preprint arXiv:2503.19786* (2025)

51. Wang, W., Zhang, B., Chen, F., Wang, Z., Li, W., Li, C., Lei, Y.: Agentic collaborative cognition for zero-shot 3d understanding. arXiv preprint arXiv:2606.24649 (2026)
52. Wu, D., Liu, F., Hung, Y.H., Duan, Y.: Spatial-mlm: Boosting mllm capabilities in visual-based spatial intelligence. arXiv preprint arXiv:2505.23747 (2025)
53. Yang, J., Yang, S., Gupta, A.W., Han, R., Fei-Fei, L., Xie, S.: Thinking in space: How multimodal large language models see, remember, and recall spaces. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10632–10643 (June 2025)
54. Yao, S., Yu, D., Zhao, J., Shafran, I., Griffiths, T., Cao, Y., Narasimhan, K.: Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems* **36**, 11809–11822 (2023)
55. Yin, B., Wang, Q., Zhang, P., Zhang, J., Wang, K., Wang, Z., Zhang, J., Chandrasegaran, K., Liu, H., Krishna, R., et al.: Spatial mental modeling from limited views. In: Structural Priors for Vision Workshop at ICCV’25 (2025)
56. Yu, S., Lee, S., Kim, N., Shin, J., Park, J., Ryu, W., Jung, R., Shim, H.: Waymoqa: A multi-view visual question answering dataset for safety-critical reasoning in autonomous driving. arXiv preprint arXiv:2511.20022 (2025)
57. Zhang, C., Li, B., Wei, M., Cao, Y.P., Gambardella, C.C., Phung, D., Cai, J.: Unified camera positional encoding for controlled video generation. arXiv preprint arXiv:2512.07237 (2025)
58. Zhang, W., Zhou, Z., Zheng, Z., Gao, C., Cui, J., Li, Y., Chen, X., Zhang, X.P.: Open3dvqa: A benchmark for comprehensive spatial reasoning with multimodal large language model in open space. arXiv preprint arXiv:2503.11094 (2025)
59. Zhu, C., Wang, T., Zhang, W., Pang, J., Liu, X.: Llava-3d: A simple yet effective pathway to empowering llms with 3d-awareness. arXiv preprint arXiv:2409.18125 (2024)