

Video Spatial Reasoning with Object-Centric 3D Rollout

Haoran Tang¹, Meng Cao^{3,1*}, Ruyang Liu¹, Xiaoxi Liang¹, Linglong Li¹, Ge Li^{1†}, Xiaodan Liang^{2,3}

¹ School of Electronic and Computer Engineering, Shenzhen Graduate School, Peking University

² Sun Yat-sen University

³ Mohamed bin Zayed University of Artificial Intelligence

Abstract

Recent advances in Multi-modal Large Language Models (MLLMs) have showcased remarkable capabilities in vision-language understanding. However, enabling robust video spatial reasoning—the ability to comprehend object locations, orientations, and inter-object relationships in dynamic 3D scenes—remains a key unsolved challenge. Existing approaches primarily rely on spatially grounded supervised fine-tuning or reinforcement learning, yet we observe that such models often exhibit *query-locked reasoning*, focusing narrowly on objects explicitly mentioned in the prompt while ignoring critical contextual cues. To address this limitation, we propose Object-Centric 3D Rollout (OCR), a novel strategy that introduces structured perturbations to the 3D geometry of selected objects during training. By degrading object-specific visual cues and projecting the altered geometry into 2D space, OCR compels the model to reason holistically across the entire scene. We further design a rollout-based training pipeline that jointly leverages vanilla and region-noisy videos to optimize spatial reasoning trajectories. Experiments demonstrate state-of-the-art performance: our 3B-parameter model achieves 47.5% accuracy on VSI-Bench, outperforming several 7B baselines. Ablations confirm OCR’s superiority over prior rollout strategies (e.g., T-GRPO, NoisyRollout).

Introduction

Human cognition effortlessly navigates and interprets the complex 3D spatial relationships inherent in dynamic visual scenes. We intuitively understand not just what objects are present, but where they are located relative to each other and how they are oriented in space. However, empowering Multi-modal Large Language Models (MLLMs) with analogous *video spatial reasoning* (Chen et al. 2024; Cheng et al. 2024; Cai et al. 2024) capabilities remains a significant challenge, although they have achieved significant progress.

To enhance MLLMs’ spatial reasoning capabilities, recent works have proposed constructing various large-scale datasets through spatially grounded Supervised Fine-Tuning (SFT) or Reinforcement Learning (RL). For instance, VSI-100K (Liao et al. 2025) and SpaceR-151K (Ouyang et al.

2025) leverage ScanNet (Dai et al. 2017) videos and 3D assets to generate synthetic yet richly annotated spatial reasoning tasks. While these datasets have led to performance gains on benchmark evaluations, our analysis reveals a critical limitation in the reasoning patterns learned by such models, **query-locked reasoning**. Specifically, MLLMs tend to focus narrowly on query-mentioned objects, neglecting the broader context of object-to-object spatial interactions. As illustrated in Figure 1, when asked a question about the relative direction between objects, VideoR1 (Feng et al. 2025) concentrates exclusively on the `pillow`, `laptop`, and `radiator`—objects directly mentioned in the query—while ignoring other spatially relevant items such as the `bed` and `desk`. These omitted entities provide important global cues necessary for accurate spatial inference, but are overlooked due to the model’s over-reliance on query-salient tokens.

To overcome this limitation and encourage models to consider object-wise correlations, we propose an **Object-Centric 3D Rollout (OCR)** strategy. The key idea is to perturb the spatial integrity of selected objects during training, thereby forcing MLLMs to reason through contextual relationships and leverage a more holistic understanding of the scene. Specifically, we randomly select 3D object assets from each video scene, inject structured noise into their 3D representations, and project the altered geometry back into the 2D space to generate region-noisy videos. Then we feed both the original (vanilla) and region-noisy videos into the GRPO rollout module to produce two parallel sets of reasoning trajectories. Both types of rollouts are aggregated to compute the final reward signals, encouraging the model to learn robust reasoning policies. This process selectively degrades visual cues associated with certain objects, compelling the model to rely on surrounding context and inter-object relationships to maintain accurate spatial reasoning.

To facilitate our object-centric 3D rollout training, we construct a suite of SFT and RL datasets (dubbed as OCR-SFT and OCR-RL, respectively), which offers two key advantages over existing spatial reasoning datasets: 1) it encompasses a more diverse range of spatial perception types, and 2) it includes high-quality Chain-of-Thought (CoT) annotations as code-start data. Based on this dataset, our method achieves 47.5% average results on the VSI-Bench benchmark at the cost of only 3B parameters, which is even

*Equal Contribution

†Corresponding Author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

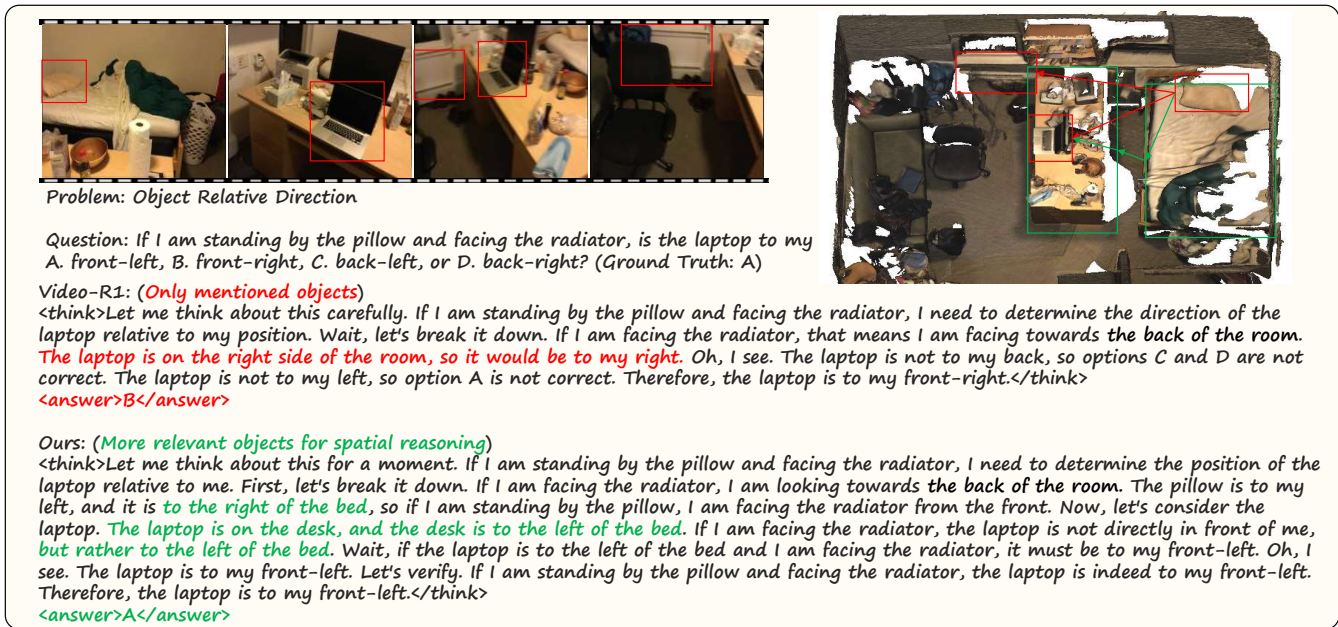


Figure 1: **Motivation of our proposed OCR.** The previous method exhibits *query-locked reasoning*, focusing narrowly on objects explicitly mentioned in the prompt while ignoring critical contextual cues. Our method encourages MLLMs to rely on surrounding context and inter-object relationships for comprehensive spatial reasoning.

higher than several existing 7B models. In addition, we conduct extensive ablation studies to demonstrate the superiority of our OCR to previous rollout strategies, including T-GRPO, NoisyRollout, *etc.*

Our contributions are summarized as follows:

- We propose a novel object-centric rollout strategy that introduces structured noise to 3D object assets to enforce reasoning based on global scene context rather than memorizing salient tokens.
- A suite of SFT and RL datasets is proposed to facilitate our object-centric 3D rollout training with broader category coverage and high-quality CoT annotations.
- Extensive experimental and ablative results have demonstrated the superior performance of our proposed OCR training strategy.

Related Works

Video Spatial Reasoning. Recent advancements in large language models (LLMs) (Hurst et al. 2024; Hui et al. 2024; Guo et al. 2025) have motivated the development of multimodal LLMs (MLLMs) that extend reasoning capabilities to spatial understanding, particularly in embodied AI scenarios. Visual spatial reasoning has gained increasing attention, with models such as SpatialVLM (Chen et al. 2024), SpatialRGPT (Cheng et al. 2024), and SpatialBot (Cai et al. 2024) utilizing grounding and depth estimation on 2D images to answer spatial questions. VSI-Bench (Yang et al. 2025b) introduced the first benchmark for video-based spatial reasoning, facilitating evaluation on real-world video inputs. Building upon this setting, VG-LLM (Zheng et al.

2025a) and Spatial-MLLM (Wu et al. 2025) incorporate 3D scene reconstruction via VGGT (Wang et al. 2025b) to extract spatial embeddings and apply supervised fine-tuning (SFT). In contrast to SFT-based methods, reinforcement learning (RL) approaches (Feng et al. 2025; Ouyang et al. 2025; Liao et al. 2025)—particularly those leveraging GRPO (Shao et al. 2024)—have shown superior adaptability and performance in spatial reasoning tasks. For example, Video-R1 (Feng et al. 2025) introduces T-GRPO, employing shuffled video sequences for reward estimation, while SpaceR (Ouyang et al. 2025) incorporates 2D grid layouts as auxiliary supervision. Despite the progress, the reasoning process in 3D scenes always focuses on query-mentioned objects while neglecting the broader contexts. To bridge this gap, we propose an object-centric 3D rollout strategy that encourages more diverse and spatially grounded learning in complex environments by introducing noise to regions of the specific 3D objects.

GRPO Rollout Modification. Reinforcement learning has become a standard training paradigm for enhancing the reasoning capabilities of large language models (LLMs), as demonstrated in GPT-4o (Hurst et al. 2024), Qwen-3 (Yang et al. 2025a), and DeepSeek-R1 (Guo et al. 2025). Among various algorithms, Group Relative Policy Optimization (GRPO) (Shao et al. 2024) has proven particularly effective in mathematical and code generation tasks by replacing the value model in PPO (Schulman et al. 2017) with group-based rollout. In the vision-language domain, GRPO has been adapted with modified rollout strategies. For instance, NoisyRollout (Liu et al. 2025) injects visual noise into image inputs to improve the robustness of advantage

estimation and encourage exploration of diverse reasoning paths. A series of methods (Cao et al. 2025; Zheng et al. 2025b; Zhang et al. 2025b) incorporate image regional features into the CoT process, enabling the transformation of textual reasoning into visual-textual interleaved reasoning.

However, in the context of video spatial reasoning, such adaptations remain limited. Existing GRPO-based methods often struggle with overfitting due to the scarcity of diverse, high-quality 3D indoor video data, leading to repeated memorization of scene layouts. This highlights a critical gap and the need for task-specific rollout strategies that better accommodate the geometric and temporal complexity of video-based spatial reasoning tasks.

Methodology

Preliminary

Group Relative Policy Optimization. (GRPO) (Shao et al. 2024) was originally developed and applied in DeepSeek-R1 (Guo et al. 2025) to enhance mathematical reasoning in natural language processing. However, it has also been effectively adapted to improve video-based spatial reasoning in recent advancements (Feng et al. 2025; Ouyang et al. 2025) with multimodal large language models (MLLMs). Given an input pair (V, \mathbf{q}) sampled from the training distribution, where V represents a video sequence derived from a 3D scene and \mathbf{q} denotes the corresponding textual query, the model generates outputs denoted as \mathbf{o} . During the training stage, the model produces n response rollouts based on the previous policy $\pi_{\theta_{\text{old}}}$ denoted as $\{(\mathbf{o}_i | V, \mathbf{q})\}_{i=1}^n$. A rule-based reward function $r(V, \mathbf{q}, \mathbf{o})$ is employed to prevent reward hacking. Specifically, the reward function assigns $r(V, \mathbf{q}, \mathbf{o}) = 1$ if the generated response \mathbf{o} is correct and adheres to the expected format. Otherwise, $r(V, \mathbf{q}, \mathbf{o}) = 0$. The rewards for all rollouts are represented as $\mathbf{r} = \{r_i\}_{i=1}^n = \{r(V, \mathbf{q}, \mathbf{o}_i)\}_{i=1}^n$. GRPO computes the advantage of each rollout based on the mean reward of all generated rollouts, denoted as $\text{mean}(\mathbf{r})$. The normalized advantage of the i -th rollout is then defined as:

$$A_i = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})}. \quad (1)$$

Following Proximal Policy Optimization (PPO) (Schulman et al. 2017) formulation, the objective function of GRPO is defined as:

$$\begin{aligned} \mathcal{J}_{\text{GRPO}}(\theta) = & \mathbb{E}_{(V, \mathbf{q}), \mathbf{o} \sim \pi_{\theta_{\text{old}}}(\cdot | V, \mathbf{q})} \\ & \left[\frac{1}{n} \sum_{i=1}^n \min \left(\frac{\pi_{\theta}(\mathbf{o}_i | V, \mathbf{q})}{\pi_{\theta_{\text{old}}}(\mathbf{o}_i | V, \mathbf{q})} A_i, \text{clip} \left(\frac{\pi_{\theta}(\mathbf{o}_i | V, \mathbf{q})}{\pi_{\theta_{\text{old}}}(\mathbf{o}_i | V, \mathbf{q})}, 1 \right. \right. \right. \\ & \left. \left. \left. - \epsilon, 1 + \epsilon \right) A_i \right) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right]. \quad (2) \end{aligned}$$

where π_{θ} is the current policy, $\epsilon > 0$ sets the clipping range.

Object-Centric 3D Rollout

While directly applying GRPO to video spatial reasoning has shown advantages over supervised fine-tuning (SFT), MLLMs still face significant limitations, as discussed in

the Introduction. Specifically, it remains challenging for the model to perform complex spatial reasoning based solely on limited video content. To address this, we propose a method that incorporates an object-centric 3D rollout policy to enhance the GRPO training process. Our design includes object-guided spatial perturbation and step-wise annealed object selector.

Object-Guided Spatial Perturbation

The key idea of object-aware spatial noisy injection is to create more challenging training examples, encouraging the model to reason about a broader set of objects beyond those mentioned in the prompt. To this end, we utilize the 3D bounding boxes of the input scene as augmented data. Specifically, for a scene S represented by video V containing k input frames $\{f_1, \dots, f_k\}$, we first extract the 3D bounding boxes of all visible objects, denoted as $\{\mathbf{B}_1, \dots, \mathbf{B}_M\}$. Each bounding box \mathbf{B}_j consists of an axis-aligned 3D box defined by its center coordinates and dimensions. For each \mathbf{B}_j , we project all enclosed 3D points $(x, y, z) \in \mathbf{B}_j$ into 2D coordinates across all frames, forming the set $\{(x, y)_{f_1}, \dots, (x, y)_{f_k}\}$. We denote this projected region across frames as $R_j = \{b_{j, f_1}, \dots, b_{j, f_k}\}$, where b_{j, f_k} denotes the selected region of object j on frame k . Given m selected bounding boxes $\{\mathbf{B}_1, \dots, \mathbf{B}_m\}$, we define the perturbed region of input frames as $\hat{R} = \bigcup_{j=1}^m R_j$. The resulting video input with noise is denoted as \hat{V} . For each input pair (V, \mathbf{q}) , Gaussian noise is added to the region \hat{R} , forming a perturbed input pair (\hat{V}, \mathbf{q}) . To increase GRPO training diversity, for every n rollouts generated from (V, \mathbf{q}) using policy $\pi_{\theta_{\text{old}}}$, we simultaneously generate n spatially perturbed rollouts from (\hat{V}, \mathbf{q}) , denoted as $\{(\hat{\mathbf{o}}_i | \hat{V}, \mathbf{q})\}_{i=n+1}^{2n}$, as illustrated in Figure 2. The combined rewards are represented as:

$$\hat{\mathbf{r}} = \{\hat{r}_i\}_{i=1}^{2n} = \{r(V, \mathbf{q}, \mathbf{o}_i)\}_{i=1}^n \cup \{r(\hat{V}, \mathbf{q}, \hat{\mathbf{o}}_i)\}_{i=n+1}^{2n} \quad (3)$$

The normalized advantage for each rollout is then calculated as:

$$\hat{A}_i = \frac{\hat{r}_i - \text{mean}(\hat{\mathbf{r}})}{\text{std}(\hat{\mathbf{r}})}. \quad (4)$$

Step-wise Annealed Object Selector

An essential question is how to choose which objects to perturb. In practice, we adopt a simple yet effective strategy: randomly selecting objects from the scene. However, since models typically learn from easy to hard cases during training, we propose a step-wise annealing strategy to encourage more complex spatial reasoning over time. This strategy involves gradually decreasing the number of objects used for perturbation, while simultaneously reducing the intensity of the applied Gaussian noise. We implement this using a scheduler $T_{\delta_t}(\cdot)$, where δ_t is a decay factor computed from the current training step α and total training steps β . Let M be the total number of objects in scene S . We select m' objects as:

$$\{\mathbf{B}_j\}_{j=1}^{m'} = T_{\delta_t}(\{\mathbf{B}_j\}_{j=1}^M). \quad (5)$$

In our experiments, δ_t is annealed linearly from 0.5 to 0 throughout training. Additional ablation studies are provided in ablative analysis.

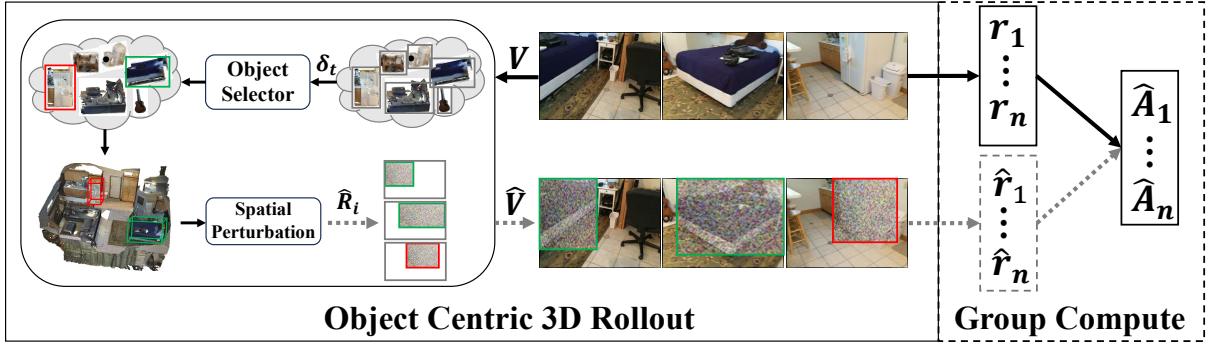


Figure 2: **Overview of our Object-Centric 3D Rollout (OCR) strategy.** Given a video scene, we generate additional rollouts by injecting spatial noise into selected object regions using a step-wise object selector. These perturbed rollouts are used for advantage estimation, while only the clean rollouts (outlined in black) contribute to gradient updates.

Finally, the training objective of our proposed methods is represented as:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{(V, \mathbf{q}), \{\mathbf{o}_j\}_1^n \sim \pi_{\theta_{\text{old}}}(\cdot | V, \mathbf{q}), \{\hat{\delta}_j\}_{n+1}^{2n} \sim \pi_{\theta_{\text{old}}}(\cdot | \hat{V}, \mathbf{q})} \left[\frac{1}{2n} \sum_{i=1}^{2n} \min \left(\frac{\pi_{\theta}(\mathbf{o}_i | V, \mathbf{q})}{\pi_{\theta_{\text{old}}}(\mathbf{o}_i | V, \mathbf{q})} \hat{A}_i, \text{clip} \left(\frac{\pi_{\theta}(\mathbf{o}_i | V, \mathbf{q})}{\pi_{\theta_{\text{old}}}(\mathbf{o}_i | V, \mathbf{q})}, 1 - \epsilon, 1 + \epsilon \right) \hat{A}_i \right) - \beta \mathbb{D}_{\text{KL}}(\pi_{\theta} \| \pi_{\text{ref}}) \right]. \quad (6)$$

Note that only rollouts from the vanilla GRPO input pair are used to compute gradients and update the model parameters.

Experiment

Num.	Config	Training Data		Acc. (%)
		Correct	Wrong	
①	\mathcal{F}^2	24,391	73,712	24.86
②	\mathcal{F}^{16}	25,287	72,816	25.78
③	$(\mathcal{F}^{16})_{bev}$	26,380	71,723	26.89
④	$(\mathcal{F}^{16})_{grpo}$	41,485	56,618	42.29

Table 1: **Statistics of model-based data filtering.** This table summarizes the selection results and model configurations used to curate high-quality training samples.

Data Construction

As video spatial reasoning lacks high-quality training datasets in VSI-Bench (Yang et al. 2025b), we construct a comprehensive training dataset containing 98K samples, covering all video spatial categories listed in VSI-Bench. We name this dataset **OCR-RL**. Additionally, to enhance the model’s spatial reasoning ability, we curate a smaller, high-quality subset with detailed spatial reasoning annotations for the cold start SFT stage, referred to as **OCR-SFT**. Details of the construction process are provided below.

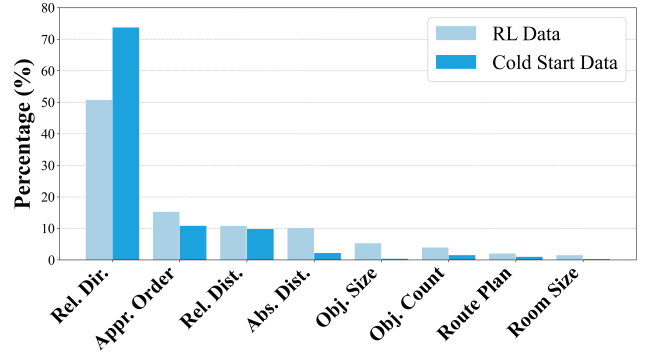


Figure 3: **Distribution of the constructed dataset.** The chart illustrates the category-wise distribution of our data. Light blue denotes the training set, while blue indicates the curated cold-start subset. Categories are sorted in descending order based on their frequency in the RL data.

Training Data Construction. For the training dataset, we follow the benchmark data generation practices outlined in VSI-Bench, excluding the categories of *route planning* and *appearance order*, as these require extensive manual annotation. To construct the dataset, for these two categories, we use the annotated data from SpaceR (Ouyang et al. 2025) and VLM-3R (Fan et al. 2025). Low-quality annotations are filtered using GPT-4o (Hurst et al. 2024), resulting in a final training set of 98K samples. The data distribution of our training dataset is shown in Figure 3. Compared with VSI-100K and SpaceR-151K, our dataset covers all categories in VSI-Bench, making the training process unbiased.

Cold Start Data Construction. The necessity of a cold start dataset arises from two key challenges: (1) the model must learn the reasoning and thinking format, and (2) modern models cannot acquire high-quality spatial reasoning capabilities solely through GRPO rollout policies. To address these challenges, we extract a subset from the training set that highlights the spatial reasoning deficiencies of the model. This subset is then used to generate high-quality spatial chain-of-thought (CoT) annotations using proprietary

Model	#Params	Avg.	Obj. Count	Abs. Dist.	Obj. Size	Room Size	Rel. Dist.	Rel. Dir.	Route Plan	Appr. Order
			Numerical Answer				Multiple-Choice Answer			
<i>Proprietary Models (API)</i>										
GPT-4o	-	34.0	46.2	5.3	43.8	38.2	37.0	41.3	31.5	28.5
Gemini-1.5-Flash	-	42.1	49.8	30.8	53.5	54.4	37.7	41.0	31.5	37.8
Gemini-1.5-Pro	-	45.4	56.2	30.9	64.1	43.6	51.3	46.3	36.0	34.6
<i>Open-source Models</i>										
InternVL2-8B	8B	34.6	23.1	28.7	48.2	39.8	36.7	30.7	29.9	39.6
InternVL2-40B	40B	36.0	34.9	26.9	46.5	31.8	42.1	32.2	34.0	39.6
LLaVA-NeXT-Video-7B	7B	35.6	48.5	14.0	47.8	24.2	43.5	42.4	34.0	30.6
LLaVA-NeXT-Video-72B	72B	40.9	48.9	22.8	57.4	35.3	42.4	36.7	35.0	48.6
LLaVA-OneVision-7B	7B	32.4	47.7	20.2	47.4	12.3	42.5	35.2	29.4	24.4
LLaVA-OneVision-72B	72B	40.2	43.5	23.9	57.6	37.5	42.5	39.9	32.5	44.6
Qwen2.5-VL-7B-Instruct	7B	33.0	40.9	14.8	43.4	10.7	38.6	38.5	33.0	29.8
Qwen2.5-VL-72B-Instruct	72B	37.0	25.1	29.3	54.5	38.8	38.2	37.0	34.0	28.9
<i>Spatial-SFT Models</i>										
SPAR	8B	41.1	-	-	-	-	-	-	-	-
SpaceR(SFT)	7B	41.6	-	-	-	-	-	-	-	-
VG-LLM	4B	46.1	66.4	36.6	55.2	56.3	40.8	43.4	30.4	39.5
<i>Spatial-RL Models</i>										
Video-R1	7B	37.1	-	-	-	-	-	-	-	-
vsGRPO-V	7B	40.7	59.9	29.6	50.8	48.3	35.4	35.6	34.0	31.5
SpaceR(SG-RLVR)	7B	45.6	-	-	-	-	-	-	-	-
Ours	3B	47.5	63.2	34.1	57.4	46.7	39.6	45.5	44.3	49.8

Table 2: Comparison with state-of-the-art models on VSI-Bench. *Spatial-SFT Models* are trained using supervised fine-tuning, while *Spatial-RL Models* are trained using reinforcement learning approaches.

Type	ColdStart Data		Rollout Method		VSI-Bench
	Video-R1	Ours	Vanilla	GRPO	Avg.
1	✓				35.1
2		✓			40.4
3	✓		✓		42.6
4		✓	✓		43.9
5	✓			✓	45.6
6		✓		✓	47.5

Table 3: Ablations on the main component of our method.

models such as GPT-4o (Hurst et al. 2024). To identify hard spatial reasoning questions, we employ Qwen2.5-VL-7B-Instruct (Bai et al. 2025) as a filtering model. The model is evaluated on the training set under four different settings, as shown in Table 1. The settings include: ① 2-frame input; ② 16-frame input; ③ 16-frame input with a global BEV (bird’s-eye view) map; ④ 16-frame input with the model fine-tuned for 2000 steps using the GRPO algorithm.

The accuracy results in Table 1 demonstrate that increasing the number of frames improves performance. Additionally, the inclusion of a BEV map provides global spatial information about the input scene, highlighting that modern models struggle with local spatial reasoning. Although GRPO training improves the model’s ability to answer questions correctly, a significant portion of errors still stem from improper spatial reasoning. To construct the cold start dataset, we focus on hard spatial problems by filtering samples based on the following criteria: (1) Questions where the model transitions from incorrect to correct answers when provided with more frames (② - ①); (2) Questions where

Method	VSI-Bench (Avg.)
Baseline	35.1
+ Vanilla GRPO	43.9
+ T-GRPO	44.1
+ Noisy Rollout	45.6
+ Downsample Rollout	44.2
+ Ours	47.5

Table 4: Ablations on rollout policy.

Schedule	VSI-Bench (Avg.)
<i>Fix</i>	45.0
<i>Linear</i>	47.5
<i>Exp</i>	46.0
<i>Cos</i>	46.3

Table 5: Ablations on object schedule.

the model transitions from incorrect to correct answers when provided with a global BEV map (③ - ①); (3) Questions that remain incorrectly answered even after GRPO training (¬④). To collect representative hard spatial problems, we filter samples from the intersections (② - ①) \wedge ¬④ and (③ - ①) \wedge ¬④, selecting 1000 samples for each category. For annotation generation, as shown in Table 1, the BEV map, which includes global spatial information, proves useful for generating high-quality data. Therefore, we annotate the dataset

using a BEV map extracted by (Wang et al. 2025a) as the first image and 32 frames as the video input for GPT-4o annotations.

Figure 3 illustrates the data distribution of our RL dataset and Cold Start SFT dataset. It can be observed that, through our data filtering process, more questions are selected from the categories of *relative direction* and *relative distance*, while fewer questions are selected from *object size*, *object count*, and *room size*, as these are not the core challenges of hard spatial reasoning.

Experimental Settings

Implementation Details. All experiments are conducted on the VSI-Bench (Yang et al. 2025b), which comprises 5130 testing samples curated from ScanNet (Dai et al. 2017), ScanNet++ (Yeshwanth et al. 2023), and ARK-itScenes (Baruch et al. 2021). We evaluate our approach using the Qwen2.5-VL-3B-Instruct (Bai et al. 2025) model, with 16-frame inputs for training and 32-frame inputs for inference. For model training, we first perform supervised fine-tuning (SFT) on OCR-SFT for 2 epochs. Subsequently, GRPO training is conducted with 2000 steps on OCR-RL. The number of clean rollouts and perturbed rollouts is set to 4 and 4. Other configurations are set following Video-R1 (Feng et al. 2025). All experiments are executed on a cluster of 8 NVIDIA A100 GPUs.

Performance Comparison

We compare our Qwen2.5-VL-3B model trained with the proposed Object-Centric 3D Rollout (OCR) strategy against recent state-of-the-art models on the VSI-Bench. As shown in Table 2, our 3B model achieves an average score of 47.5, surpassing several larger-scale models, including proprietary systems such as GPT-4o (Hurst et al. 2024) and Gemini-1.5-Pro (Comanici et al. 2025), as well as open-source models like LLaVA-Next-Video-72B (Zhang et al. 2024), LLaVA-OneVision (Li et al. 2024), and Qwen2.5-VL-72B (Bai et al. 2025). This demonstrates the strong effectiveness of our OCR training pipeline and high-quality curated dataset. To further validate the impact of our rollout strategy, we compare against both *Spatial-SFT* and *Spatial-RL* baselines. Among the *Spatial-SFT* models (Zhang et al. 2025a; Ouyang et al. 2025; Zheng et al. 2025a), VG-LLM, built on the same Qwen2.5-VL-3B backbone, shows competitive performance—especially in object counting and absolute distance reasoning—benefiting from the powerful 3D encoding capabilities of VGGT (Wang et al. 2025b). However, this comes at the cost of increased computational overhead and slower inference. In the *Spatial-RL* category (Zheng, Huang, and Wang 2025; Liao et al. 2025; Ouyang et al. 2025), our method consistently outperforms previous 7B-scale models that employ GRPO variants, such as Video-R1, vsGRPO-V, and SpaceR, highlighting the unique advantages of our object-centric rollout design. These results collectively confirm that incorporating structured object-level perturbations during training significantly enhances a model’s ability to reason spatially in complex 3D environments.

Ablation Analysis

We conduct comprehensive ablation experiments to analyze the contributions of each component in our proposed Object-Centric 3D Rollout (OCR) framework. As shown in Table 3, we first compare our curated cold-start dataset (OCR-SFT) against the one adopted in Video-R1 (Feng et al. 2025). Comparing setups (1 vs. 2), (3 vs. 4), and (5 vs. 6), we observe consistent improvements when using our dataset, both in the initial cold-start phase and during full GRPO training. These results highlight two key benefits: (1) our annotated data offers high-quality Chain-of-Thought(CoT) guidance, directly supporting spatial reasoning, and (2) the model-based filtering strategy effectively selects challenging samples that require deeper reasoning, thereby enriching the training distribution. To isolate the effect of our proposed rollout mechanism, we compare vanilla GRPO rollouts with our object-centric 3D rollout (comparing 3 vs. 5 and 4 vs. 6). The consistent performance gain across these pairs demonstrates the effectiveness of injecting structured spatial perturbations into the rollout policy.

Further analysis of rollout variants is presented in Table 4. We benchmark our method against T-GRPO (Feng et al. 2025), NoisyRollout (Liu et al. 2025), and a simple frame-downsampling strategy. For NoisyRollout, we extend image-based noise augmentation to all video frames. The downsampling variant randomly removes frames following the same scheduling function used in OCR. While NoisyRollout improves robustness in spatial tasks, our object-centric perturbation achieves superior performance by explicitly encoding spatial object interactions.

Finally, we evaluate different scheduling strategies used in the object selector. We evaluate four perturbation schedules: *Fix* (a constant 25% object perturbation), *Linear* (linearly decreasing from 50% to 0%), *Exp*, and *Cos* (both using exponential and cosine decay, respectively, with the same range as *Linear*). We find that the *Linear* schedule yields the best performance under GRPO training. This schedule encourages the model to learn from easier cases initially and gradually adapt to harder scenarios, thus discovering more diverse reasoning paths. All experiments are conducted using the Qwen2.5-VL-3B (Bai et al. 2025) backbone to ensure consistent comparisons.

Visualization Results

We provide additional qualitative results on VSI-Bench in Figure 4 to demonstrate the effectiveness of our proposed GRPO training strategy. The visualizations reveal two key advantages of our method. **(1) Reasoning with Extra Objects.** As highlighted in green in Figure 4(a) and (b), our model frequently leverages additional objects not explicitly mentioned in the question. In Figure 4(a), the model infers the relative positions of the *refrigerator*, *table*, and *dishwasher* by incorporating the *sink*, which provides essential spatial cues absent from the original query and frames. Similarly, in Figure 4(b), the model utilizes contextual objects such as the *shower*, *shelves*, and *cabinets* to identify the room type and estimate its size based on object arrangement. These examples demonstrate



Problem: Object Relative Direction

Question: If I am standing by the refrigerator and facing the table, is the dishwasher to my A. front-left, B. front-right, C. back-left, or D. back-right? (Ground Truth: A)

Ours: (More relevant objects for spatial reasoning)

<think>Let me think about the layout of the kitchen based on the images provided. The refrigerator is positioned along one side of the kitchen, and the table is in the center. The dishwasher is located near the sink, which is adjacent to the refrigerator. Hmm, let's break it down. The dishwasher is located near the sink, which is on the same side as the refrigerator. Since the sink is on the same side as the refrigerator, the dishwasher would be to my left if I am facing the table. Wait, let's verify this. If the table is directly in front of me, and the refrigerator is directly in front of me, then the dishwasher, which is near the sink, would be to my left. This aligns with the layout where the dishwasher is positioned to the left of the refrigerator when facing the table. Oh, I see. The dishwasher is indeed to my front-left when standing by the refrigerator and facing the table.</think>

<answer>A</answer>

(a)



Problem: Room Size Estimation

Question: What is the size of this room (in square meters)? If multiple rooms are shown, estimate the size of the combined space.

(Ground Truth: 3.3)

Ours: (Global layout for spatial reasoning)

<think>Let me think about this. The images show a sequence of a bathroom, which includes a toilet, a sink, a shower, and some storage areas. The images provide a sense of the layout and the size of the room. First, let's break down the images. The bathroom appears to be compact, with a toilet, sink, and shower all in close proximity. The images show a small space with tiled walls and floors, which suggests a compact layout typical of a small bathroom. The images also show some storage areas, like shelves and cabinets, which are often used to store toiletries and other items. These storage areas are relatively small, indicating that the overall space is not very large. Considering the typical dimensions of a bathroom, a compact space like this might be around 3 to 4 square meters. A small bathroom like this could be approximately 3.5 square meters. Hmm, let's verify this. If the bathroom is about 3.5 square meters, and there are no other rooms shown, then the size of the combined space is simply the size of the bathroom itself. Oh, I see. The size of the room, based on the images and typical bathroom layouts, is approximately 3.5 square meters.</think>

<answer>3.5</answer>

(b)

Figure 4: **Visualization samples of our method.** Text in green indicates instances where the model incorporates additional relevant objects beyond those mentioned in the query, enhancing spatial reasoning. Text in blue reflects the model's ability to reason from a global scene perspective rather than relying solely on local visual cues.

the model's ability to introduce relevant auxiliary objects to enhance spatial reasoning. **(2) Global Spatial Perspective.** While existing spatial reasoning models tend to rely on local visual cues, our model exhibits a broader understanding by first establishing a global layout of the scene before addressing the specific question. In Figure 4(a), the model outlines a holistic scene layout and then selects the *sink* to support finer-grained spatial inference. In Figure 4(b), the model first composes a global view of the bathroom and then accurately estimates its size using the observed object configuration. These findings further demonstrate that our object-centric 3D rollout policy encourages the model to reason beyond explicitly mentioned entities and promotes spatially coherent reasoning grounded in 3D scene structure.

Conclusion

We present Object-Centric 3D Rollout (OCR), a novel training strategy for enhancing video spatial reasoning in MLLMs. By introducing structured 3D perturbations to object regions and jointly leveraging vanilla and region-noisy rollouts, OCR encourages holistic scene understanding beyond query-mentioned objects. Trained on a curated dataset with diverse spatial scenarios and high-quality CoT annotations, our 3B model achieves state-of-the-art results on VSI-Bench, outperforming several 7B baselines. This work underscores the value of geometry-aware rollouts and paves the way for more context-grounded spatial reasoning in multimodal systems.

Acknowledgments

This work is supported by the National Key Research and Development Program of China (2024YFE0203100).

References

- Bai, S.; Chen, K.; Liu, X.; Wang, J.; Ge, W.; Song, S.; Dang, K.; Wang, P.; Wang, S.; Tang, J.; et al. 2025. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*.
- Baruch, G.; Chen, Z.; Dehghan, A.; Dimry, T.; Feigin, Y.; Fu, P.; Gebauer, T.; Joffe, B.; Kurz, D.; Schwartz, A.; et al. 2021. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *arXiv preprint arXiv:2111.08897*.
- Cai, W.; Ponomarenko, I.; Yuan, J.; Li, X.; Yang, W.; Dong, H.; and Zhao, B. 2024. Spatialbot: Precise spatial understanding with vision language models. *arXiv preprint arXiv:2406.13642*.
- Cao, M.; Zhao, H.; Zhang, C.; Chang, X.; Reid, I.; and Liang, X. 2025. Ground-R1: Incentivizing Grounded Visual Reasoning via Reinforcement Learning. *arXiv preprint arXiv:2505.20272*.
- Chen, B.; Xu, Z.; Kirmani, S.; Ichter, B.; Sadigh, D.; Guibas, L.; and Xia, F. 2024. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14455–14465.
- Cheng, A.-C.; Yin, H.; Fu, Y.; Guo, Q.; Yang, R.; Kautz, J.; Wang, X.; and Liu, S. 2024. Spatialrgpt: Grounded spatial reasoning in vision-language models. *Advances in Neural Information Processing Systems*, 37: 135062–135093.
- Comanici, G.; Bieber, E.; Schaekermann, M.; Pasupat, I.; Sachdeva, N.; Dhillon, I.; Blistein, M.; Ram, O.; Zhang, D.; Rosen, E.; et al. 2025. Gemini 2.5: Pushing the frontier with advanced reasoning, multimodality, long context, and next generation agentic capabilities. *arXiv preprint arXiv:2507.06261*.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5828–5839.
- Fan, Z.; Zhang, J.; Li, R.; Zhang, J.; Chen, R.; Hu, H.; Wang, K.; Qu, H.; Wang, D.; Yan, Z.; et al. 2025. VLM-3R: Vision-Language Models Augmented with Instruction-Aligned 3D Reconstruction. *arXiv preprint arXiv:2505.20279*.
- Feng, K.; Gong, K.; Li, B.; Guo, Z.; Wang, Y.; Peng, T.; Wu, J.; Zhang, X.; Wang, B.; and Yue, X. 2025. Video-r1: Reinforcing video reasoning in mllms. *arXiv preprint arXiv:2503.21776*.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Hui, B.; Yang, J.; Cui, Z.; Yang, J.; Liu, D.; Zhang, L.; Liu, T.; Zhang, J.; Yu, B.; Lu, K.; et al. 2024. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186*.
- Hurst, A.; Lerer, A.; Goucher, A. P.; Perelman, A.; Ramesh, A.; Clark, A.; Ostrow, A.; Welihinda, A.; Hayes, A.; Radford, A.; et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*.
- Li, B.; Zhang, Y.; Guo, D.; Zhang, R.; Li, F.; Zhang, H.; Zhang, K.; Zhang, P.; Li, Y.; Liu, Z.; et al. 2024. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*.
- Liao, Z.; Xie, Q.; Zhang, Y.; Kong, Z.; Lu, H.; Yang, Z.; and Deng, Z. 2025. Improved visual-spatial reasoning via r1-zero-like training. *arXiv preprint arXiv:2504.00883*.
- Liu, X.; Ni, J.; Wu, Z.; Du, C.; Dou, L.; Wang, H.; Pang, T.; and Shieh, M. Q. 2025. Noisyrollout: Reinforcing visual reasoning with data augmentation. *arXiv preprint arXiv:2504.13055*.
- Ouyang, K.; Liu, Y.; Wu, H.; Liu, Y.; Zhou, H.; Zhou, J.; Meng, F.; and Sun, X. 2025. SpaceR: Reinforcing MLLMs in Video Spatial Reasoning. *arXiv preprint arXiv:2504.01805*.
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal Policy Optimization Algorithms. *arXiv:1707.06347*.
- Shao, Z.; Wang, P.; Zhu, Q.; Xu, R.; Song, J.; Bi, X.; Zhang, H.; Zhang, M.; Li, Y.; Wu, Y.; et al. 2024. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*.
- Wang, H.; Zhao, Y.; Wang, T.; Fan, H.; Zhang, X.; and Zhang, Z. 2025a. Ross3d: Reconstructive visual instruction tuning with 3d-awareness. *arXiv preprint arXiv:2504.01901*.
- Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Rupprecht, C.; and Novotny, D. 2025b. Vggt: Visual geometry grounded transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5294–5306.
- Wu, D.; Liu, F.; Hung, Y.-H.; and Duan, Y. 2025. Spatial-mllm: Boosting mllm capabilities in visual-based spatial intelligence. *arXiv preprint arXiv:2505.23747*.
- Yang, A.; Li, A.; Yang, B.; Zhang, B.; Hui, B.; Zheng, B.; Yu, B.; Gao, C.; Huang, C.; Lv, C.; et al. 2025a. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- Yang, J.; Yang, S.; Gupta, A. W.; Han, R.; Fei-Fei, L.; and Xie, S. 2025b. Thinking in space: How multimodal large language models see, remember, and recall spaces. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 10632–10643.
- Yeshwanth, C.; Liu, Y.-C.; Nießner, M.; and Dai, A. 2023. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12–22.
- Zhang, J.; Chen, Y.; Zhou, Y.; Xu, Y.; Huang, Z.; Mei, J.; Chen, J.; Yuan, Y.-J.; Cai, X.; Huang, G.; et al. 2025a. From flatland to space: Teaching vision-language models to perceive and reason in 3d. *arXiv preprint arXiv:2503.22976*.
- Zhang, X.; Gao, Z.; Zhang, B.; Li, P.; Zhang, X.; Liu, Y.; Yuan, T.; Wu, Y.; Jia, Y.; Zhu, S.-C.; et al. 2025b. Chain-of-Focus: Adaptive Visual Search and Zooming for Multimodal Reasoning via RL. *arXiv preprint arXiv:2505.15436*.

Zhang, Y.; Wu, J.; Li, W.; Li, B.; Ma, Z.; Liu, Z.; and Li, C. 2024. Video instruction tuning with synthetic data. *arXiv preprint arXiv:2410.02713*.

Zheng, D.; Huang, S.; Li, Y.; and Wang, L. 2025a. Learning from Videos for 3D World: Enhancing MLLMs with 3D Vision Geometry Priors. *arXiv preprint arXiv:2505.24625*.

Zheng, D.; Huang, S.; and Wang, L. 2025. Video-3d llm: Learning position-aware video representation for 3d scene understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8995–9006.

Zheng, Z.; Yang, M.; Hong, J.; Zhao, C.; Xu, G.; Yang, L.; Shen, C.; and Yu, X. 2025b. DeepEyes: Incentivizing” Thinking with Images” via Reinforcement Learning. *arXiv preprint arXiv:2505.14362*.