SpatialThinker: Reinforcing Scene Graph-Grounded Spatial Reasoning via Dense Rewards

Anonymous authors
Paper under double-blind review

Figure 1: Current MLLMs struggle with fine-grained spatial reasoning, inferring affordances, or interpreting real-world spatial relationships. SpatialThinker addresses this by integrating scene graph generation within reasoning, enforcing structured grounding.

Abstract

Multimodal large language models (MLLMs) have achieved remarkable progress in vision–language tasks, but continue to struggle with spatial reasoning. Existing spatial MLLMs rely on large-scale datasets, explicit 3D inputs, architecture-specific modifications, or sparse Reinforcement Learning (RL) methods that provide insufficient guidance for spatially-grounded reasoning. We introduce SpatialThinker. To our knowledge, it is the first MLLM unifying Scene Graph Generation (SGG) and visual reasoning in a single pass via online RL. The model simulates human-like spatial perception by constructing a mental scene graph of task-relevant objects and relations, and reasoning toward an answer via dense spatial rewards. Our contributions are threefold: (1) SGG-grounded reasoning: integrating SGG directly within the reasoning chain rather than as a disjoint preprocessing step; (2) STVQA-7K:

a high-quality spatial VQA training dataset via a scalable synthesis pipeline; and (3) a dense spatial reward design that enforces structured grounding during RL and generalizes to improve broad visual perception. SpatialThinker-7B achieves 3.6× larger gains over SFT and 1.7× better in- and out-of-distribution generalization than sparse RL. Trained on only 7K samples, SpatialThinker-7B matches GPT-5 and outperforms GPT-4o, while SpatialThinker-30B surpasses both GPT-5 and Claude 4 Sonnet on average across 14 spatial and real-world benchmarks, demonstrating that structured spatial grounding with reward-aligned reasoning enables robust spatial understanding with limited data.

1 Introduction

Spatial reasoning is central to human intelligence, enabling us to perceive, localize, and manipulate objects in complex environments, a capability crucial for embodied AI tasks like robotic manipulation Intelligence et al. (2025); Gao et al. (2023); Nasiriany et al. (2024), navigation Huang et al. (2022), and augmented reality Konenkov et al. (2024), where precise spatial awareness underpins real-world deployment Driess et al. (2023); Team et al. (2025). While MLLMs have advanced rapidly in vision-language tasks Hurst et al. (2024); Lin et al. (2024); Deitke et al. (2025); Bai et al. (2025b); Du et al. (2025); Liu et al. (2023); Google (2025), they continue to struggle with spatial understanding, especially in 3D space, which requires capturing geometry, structure, and relations beyond 2D projections Chen et al. (2024a); Tong et al. (2024b); Kamath et al. (2023); Yang et al. (2025a); Tong et al. (2024a); Ma et al. (2024b).

Existing approaches address this through large-scale data synthesis from 3D scene graphs Chen et al. (2024a); Ma et al. (2025b); Daxberger et al. (2025); Cheng et al. (2024), auxiliary spatial tokens or architectural modifications Hong et al. (2023b); Ma et al. (2025b), or explicit 3D inputs like depth maps and point clouds Hong et al. (2023c); Cheng et al. (2024); Cai et al. (2024). A common pattern is the use of scene graphs as an offline preprocessing tool for data curation: SpatialRGPT builds 3D scene graphs from point clouds to generate 700K training samples Cheng et al. (2024), while SpatialVLM and SpatialLLM rely on scene graph-derived annotations to supervise training on 2B and 1M samples respectively Chen et al. (2024a); Ma et al. (2025b), resulting in data-intensive pipelines that require either massive scale or architecture-specific modifications.

Recently, RLVR has shown superior generalization over SFT by learning diverse reasoning strategies rather than static patterns DeepSeek-AI et al. (2025); Shen et al. (2025b); Gandhi et al. (2025). However, existing RLVR approaches for spatial reasoning employ simple accuracy-only rewards, providing insufficient guidance for visually-grounded reasoning Shen et al. (2025a); Xiao et al. (2025); Ma et al. (2025a); Wang & Ling (2025); Xia et al. (2025); Zhu et al. (2025). Meanwhile, scene graphs offer natural structure to guide visual reasoning Hildebrandt et al. (2020); Wald et al. (2020), yet existing methods treat them either as external preprocessing pipelines for data curation Kim et al. (2024); Chen et al. (2023); Li et al. (2024c) or as isolated generation targets decoupled from downstream reasoning Chen et al. (2025c); Li et al. (2025), never integrating them end-to-end within the reasoning process itself.

We address these with SpatialThinker, the first MLLM to our knowledge that unifies Scene Graph Generation (SGG) and visual reasoning via online RL. Rather than treating SGG as a disjoint preprocessing step, SpatialThinker integrates scene graph construction directly within the reasoning chain, constructing region-of-interest scene graphs that capture task-relevant objects, spatial relations, and localized coordinates, and reasons over these structured representations toward an answer. The training leverages a multi-objective reward with lexicographic ordering: format rewards enforce structured reasoning; count penalties regulate regional focus; accuracy rewards prioritize correctness; and spatial rewards encourage precise localization. This promotes human-like reasoning following *observe, localize, think, answer*, where 2D grounding tells the model "where objects are" while 3D relational predicates tell it "how objects sit in the world", mirroring the mental sketches humans form when perceiving a scene.

SpatialThinker-7B, trained on only 7K samples from our synthesized STVQA-7K dataset, outperforms SFT (+5.5%) and conventional RL baselines (+3.2%) across fourteen benchmarks, surpassing GPT-4o (+4.7% avg.), Claude 3.5 Sonnet (+9.6% avg.), and Claude 4 Sonnet (+1.8% avg.) Hurst et al. (2024); Anthropic

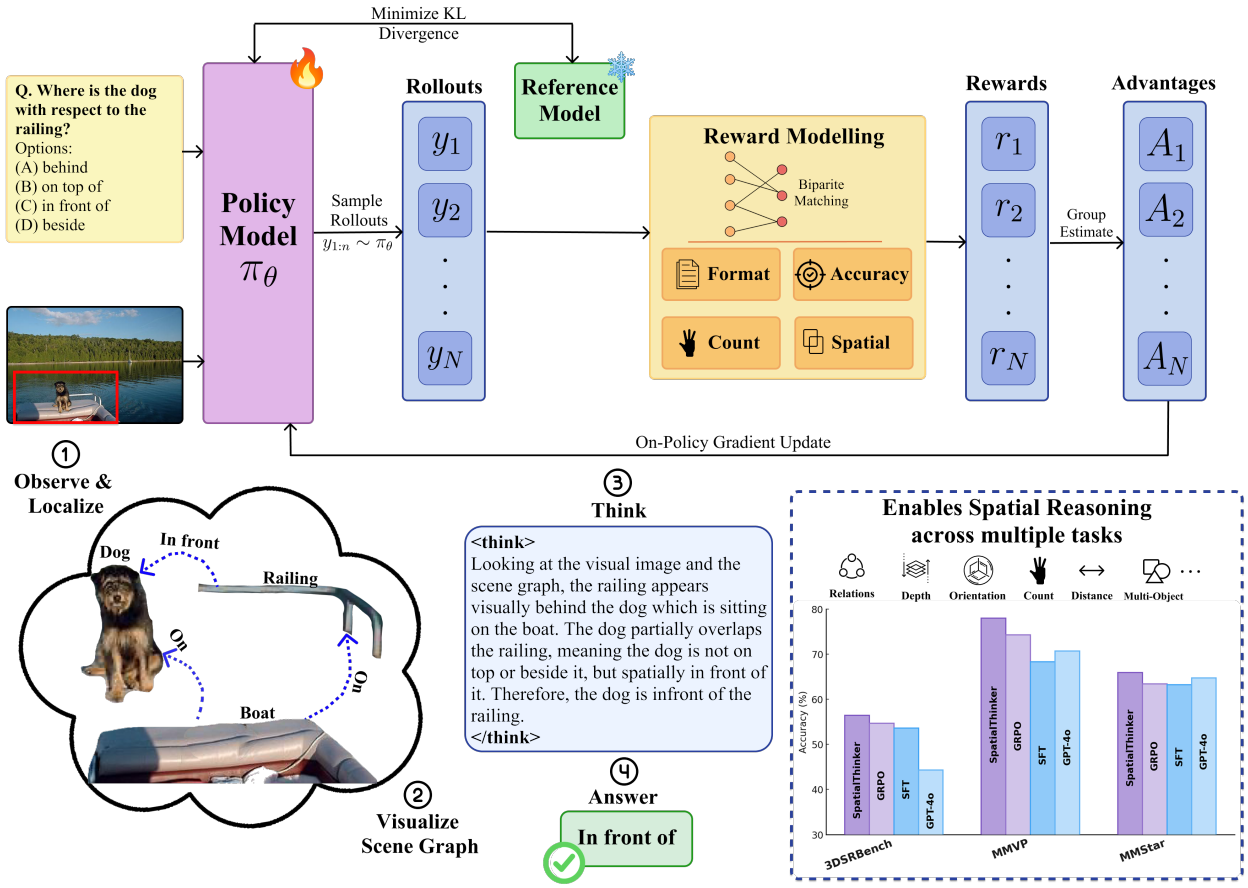

Figure 2: Method overview of SpatialThinker. Our framework unifies scene graph generation and visual reasoning via multi-objective dense RL training, enabling structured spatial grounding as part of the reasoning chain leading to enhanced spatial understanding in MLLMs across tasks and general visual perception.

(2024; 2025) while matching GPT-5 (-0.9% avg.). While vanilla sparse RL improves the base model by $+4.4\%$, our dense spatial rewards achieve $+7.7\%$, nearly doubling ($\times 1.7$) the RL benefit. Crucially, sparse RL merely matches SFT on OOD real-world VQA ($+2.7\%$ vs. $+2.9\%$), whereas SpatialThinker achieves $+5.2\%$, confirming that accuracy-only rewards fail to improve generalization to generic visual perception tasks. Scaling the same recipe to 30B, SpatialThinker-30B attains the best overall average, surpassing GPT-5 by $+3.0\%$ and Claude 4 Sonnet by $+5.8\%$, and reaching 93.6% on CV-Bench 3D Tong et al. (2024a) ($+3.3\%$ over GPT-5). This demonstrates that structured spatial grounding via dense rewards surpasses static SFT patterns learned from far larger datasets Chen et al. (2024a); Ma et al. (2024a).

Our main contributions are:

- We propose SpatialThinker, the first MLLM to our knowledge that unifies SGG and visual reasoning generation in a single pass via online RL, enabling the model to jointly perceive spatial structure and reason, achieving strong performance with only 7K training samples vs. hundreds of thousands to billions used by existing Spatial MLLMs.
- We introduce STVQA-7K, a high-quality spatial VQA training dataset grounded in scene graphs annotations, along with a scalable data generation pipeline supporting up to 108K samples with dual-LLM verification for quality control.
- We design a dense, lexicographically gated multi-objective reward that guides regionally focused spatial reasoning, and generalizes beyond spatial tasks to improve broad visual perception, achieving

superior in- and out-of-distribution generalization across spatial, generic VQA, and real-world benchmarks, and outperforming conventional RL and SFT baselines, open-sourced generalist and spatial MLLMs, and proprietary models. We release SpatialThinker at 3B, 7B, and 30B scales, with SpatialThinker-30B surpassing state-of-the-art proprietary models including GPT-5 and Claude 4 Sonnet on average across 14 benchmarks.

2 Preliminaries

Scene Graph Generation (SGG). A scene graph provides a structured representation of an image I as a directed graph $G = (V, E)$. Each node $v_i \in V$ denotes an object with a category label c_i and a 2D bounding box $b_i = (x_1, y_1, x_2, y_2)$; each edge $e_{ij} \in E$ is a relationship triplet $\langle v_i, r_{ij}, v_j \rangle$ consisting of subject v_i, predicate r_{ij}, and object v_j that capture spatial or interactive relations (e.g., *left of*, *on*, *under*) (Hildebrandt et al., 2020; Wald et al., 2020). Classical SGG decomposes prediction into object detection and relation recognition (Carion et al., 2020; Cong et al., 2023), while open-vocabulary methods leverage language or vision priors to generalize beyond fixed ontologies (Chen et al., 2024b; Li et al., 2023). We refer to *question-focused scene subgraphs* as $G_q = (V_q, E_q) \subseteq G$ that retain objects and relations relevant to a given query q.

Reasoning in Multimodal Large Language Models. MLLMs aim to solve reasoning tasks defined over a dataset $\mathcal{D}$ of multimodal instances $(\mathbf{x}_{\text{img}}, \mathbf{x}_{\text{text}}, \mathbf{y}^*)$, where $\mathbf{x}_{\text{img}}$ is a visual input, $\mathbf{x}_{\text{text}}$ is a natural language query, and $\mathbf{y}^*$ is the ground-truth answer. We model the MLLM as an autoregressive policy π_θ that outputs a trajectory $\mathbf{y} = (s_1, \ldots, s_T, a)$ consisting of reasoning steps s_t and a final answer a. The policy factorizes as:

$$\pi_\theta(\mathbf{y} \mid \mathbf{x}_{\text{img}}, \mathbf{x}_{\text{text}}) = \left(\prod_{t=1}^{T} \pi_\theta(s_t \mid \mathbf{x}_{\text{img}}, \mathbf{x}_{\text{text}}, s_{<t}) \right) \cdot \pi_\theta(a \mid \mathbf{x}_{\text{img}}, \mathbf{x}_{\text{text}}, s_{\leq T}) \tag{1}$$

Supervised fine-tuning enables imitation of reference reasoning traces but often struggles with out-of-distribution generalization. Reinforcement learning (RL) instead optimizes reasoning trajectories with explicit reward signals, improving robustness (Gandhi et al., 2025; DeepSeek-AI et al., 2025; Huang et al., 2025). The RL objective is given by: $\max_\theta \ \mathbb{E}_{(\mathbf{x}_{\text{img}}, \mathbf{x}_{\text{text}}, \mathbf{y}^*) \sim \mathcal{D},\, \mathbf{y} \sim \pi_\theta} \big[R(\mathbf{y}) \big]$, where $R(\mathbf{y})$ evaluates the trajectory based on format adherence, object counting, answer correctness, and spatial localization.

3 SpatialThinker: Scene Graph-Grounded Spatial Reasoning MLLMs

Task Formulation We cast spatial reasoning in MLLMs as the task of producing a visually grounded response $\mathbf{y}$ to a query $Q = \{\mathbf{x}_{\text{img}}, \mathbf{x}_{\text{text}}\}$. Unlike generic reasoning, our formulation explicitly requires constructing question-focused scene subgraphs G_q and reasoning over objects, bounding boxes, and relations. The policy π_θ is trained on spatially grounded VQA samples from our STVQA-7K dataset (3.3) using our multi-objective spatial reward design R (3.1) verified based on the ground-truth scene graphs from the dataset. This enforces structural validity, count fidelity, answer accuracy, and precise spatial grounding.

3.1 Multi-Objective Reward Design

SpatialThinker is trained with a fine-grained, multi-objective reward function that guides spatial reasoning via explicit visual grounding. Unlike prior RLVR methods that use sparse final-answer rewards (Peng et al., 2025; Zhu et al., 2025; Shen et al., 2025b), our dense reward design combines lexicographic gating with four components—format, count, accuracy, and spatial rewards. We further discuss our reward design process, including ablations and our rationale in Appendix C.

Format Reward. We enforce a visually-grounded and structured reasoning template: `<observe>` for scene description, `<scene>` for regional scene graphs with objects, bounding boxes, and relations, `<think>` for explicit reasoning, and `<answer>` for the final output. Beyond tag presence, the format reward validates the JSON inside `<scene>`, ensuring (1) it is parseable, (2) each object includes required fields (ID and bounding box), and (3) all relations are valid subject–predicate–object triplets. This encourages sequential grounding: perceive $\rightarrow$ localize $\rightarrow$ reason $\rightarrow$ answer. The reward $R_f \in 0, 1$ is weighted at $w_{\text{format}} = 0.1$.

Accuracy Reward. To prioritize task performance, we define the accuracy reward R_a as a binary score based on exact string match between the model's predicted answer and the ground-truth answer, enabled by our multiple-choice format. This component carries the highest weight ($w_{\text{accuracy}} = 0.5$), directly incentivizing correct final predictions, while the other rewards shape how the model arrives at correct answers.

Count Reward. The count reward encourages the model to predict the appropriate number of objects and relations relevant to the query, penalizing both under- and over-generation based on the deviation between predicted and ground-truth counts for both:

$$R_{\text{count}} = w_{\text{count}} \cdot \left(\lambda_{\text{obj}} \cdot \max\left(0,\ 1 - \frac{|N_{\text{obj}}^{\text{pred}} - N_{\text{obj}}^{\text{gt}}|}{\max(N_{\text{obj}}^{\text{gt}}, 1)}\right) + \lambda_{\text{rel}} \cdot \max\left(0,\ 1 - \frac{|N_{\text{rel}}^{\text{pred}} - N_{\text{rel}}^{\text{gt}}|}{\max(N_{\text{rel}}^{\text{gt}}, 1)}\right) \right)$$

where N^{pred} and N^{gt} denote predicted and ground truth counts respectively, $w_{\text{count}} = 0.2$ is the overall count reward weight, and λ_{obj} and λ_{rel} are set to 0.7 and 0.3 respectively. This guides the model to stay focused on question-relevant regions. Without it, we found the models tend to game the spatial reward by generating excessive objects and relations to maximize random matches—a form of reward hacking.

Spatial Reward. To supervise object localization, we compute the spatial reward only when the final answer is correct. Predicted and ground-truth objects are matched using the Hungarian algorithm for bipartite matching with a cost function that combines Complete IoU (CIoU) and semantic similarity:

$$C(o_i^{\text{pred}}, o_j^{\text{gt}}) = \lambda_{\text{spatial}} \left(1 - \text{IoU}(b_i, b_j)\right) + \lambda_{\text{semantic}} \left(1 - \text{sim}(l_i, l_j)\right) \tag{2}$$

where b and l denote bounding boxes and labels, respectively, $\lambda_{spatial} = 1.0$, and $\lambda_{semantic} = 2.0$. The reward is then computed as the average CIoU across matched pairs: $R_{\text{spatial}} = w_{\text{spatial}} \cdot \left(\frac{1}{|\mathcal{M}|} \sum_{(i,j)\in\mathcal{M}} \text{CIoU}(b_i^{\text{pred}}, b_j^{\text{gt}}) \right)$, where $w_{\text{spatial}} = 0.2$. CIoU offers dense supervision over IoU, even for non-overlapping boxes by incorporating distance and aspect ratio terms (Zheng et al., 2020). Although computed in 2D, these grounding signals combined with 3D-linked relations promote 3D-consistent spatial reasoning.

Lexicographic Gating. We apply lexicographic ordering with conditional gating (Skalse et al., 2022), prioritizing format $\succ$ {count, accuracy} $\succ$ spatial. The model must first satisfy formatting, then jointly optimize count and accuracy, and receives spatial reward only when the answer is correct. This ensures spatial grounding reinforces valid reasoning and makes the model robust to imperfect intermediate scene graphs. Without accuracy gating, we find that models tend to over-optimize intermediate spatial rewards at the expense of final answer correctness. The final reward is computed as follows, where $\mathbb{I}[\cdot]$ is the indicator function:

$$R_{\text{total}} = \mathbb{I}[R_{\text{format}} = 1] \cdot \left(w_{\text{format}} R_f + w_{\text{count}} R_c + w_{\text{accuracy}} R_a \mathbb{I}[R_{\text{accuracy}} = 1]\, w_{\text{spatial}} R_s \right) \tag{3}$$

3.2 Online RL Policy Optimization

To train SpatialThinker with dense, lexicographically gated rewards, we adopt Group-Relative Policy Optimization (GRPO) (DeepSeek-AI et al., 2025; Shao et al., 2024), an online RL method that avoids critic networks by estimating advantages through intra-group comparisons. Given an input $\mathbf{x}$, we sample N trajectories $\{y^{(1)}, \ldots, y^{(N)}\}$ from the current policy $\pi_{\theta_{\text{old}}}$. Each response is scored via our dense spatial reward function (3.1), and advantages are computed using group-normalized scores: $A^{(i)} = \frac{r^{(i)} - \mu}{\sigma + \varepsilon}$, where μ and σ are the group mean and standard deviation, and $\varepsilon = 10^{-6}$. We then update the policy using a PPO-style clipped loss with KL regularization:

$$\mathcal{L}_{\text{RL}}(\theta) = -\frac{1}{G} \sum_{i=1}^{G} \frac{1}{|y^{(i)}|} \sum_{t=1}^{|y^{(i)}|} \left[\min\left(r^{i,t} A^{(i)},\ \text{clip}(r^{i,t},\, 1-\epsilon_l,\, 1+\epsilon_h)\, A^{(i)}\right) - \beta\, D_{\text{KL}}^{i,t} \right] \tag{4}$$

where $r^{i,t} = \frac{\pi_\theta(y_t^{(i)} | \mathbf{x}, y_{<t}^{(i)})}{\pi_{\theta_{\text{old}}}(y_t^{(i)} | \mathbf{x}, y_{<t}^{(i)})}$ is the importance ratio between new and old policies, and $D_{\text{KL}}^{i,t}$ is the token-level KL divergence against a reference model. We set $\epsilon_l = 0.2$, $\epsilon_h = 0.3$, and $\beta = 10^{-2}$. This objective balances learning from dense spatial rewards while constraining policy divergence to ensure stability and generalization.

3.3 STVQA-7K: Dataset Construction

To facilitate reward-aligned spatial reasoning, we construct STVQA-7K, a synthetic visual question answering (VQA) dataset built from human-annotated scene graphs in Visual Genome (Krishna et al., 2017). STVQA-7K comprises 7,587 spatially grounded multiple-choice VQA pairs spanning both 2D and 3D spatial understanding, covering nine core reasoning types including relations, size, orientation, distance, depth, reach, location, count, and existence.

We augment the original VG150 predicate set with 34 additional spatial relations—covering distance (e.g., near, far), size (e.g., bigger, taller), orientation (e.g., facing away), and containment (e.g., inside, beneath)—to enrich the relational vocabulary beyond the standard 50 predicates. From these scene graphs, we generate 3D-reasoning VQA, e.g., depth: "*Which is closer to the camera?*" and orientation: "*From the person's perspective, which direction is the dog?*".

Each QA pair is generated from a scene graph using Claude Sonnet 4 (Anthropic, 2025), and rated by difficulty and quality. To mitigate potential biases from LLM-generated questions, we apply a consistency-based dual-LLM filtering pipeline for quality control: Claude Sonnet 4 generates QA pairs, then GPT-4o (Hurst et al., 2024) validates them via pass@2 agreement. This cross-model verification yields a +13% accuracy improvement (Table 5 , confirming that dual-LLM validation retains high-quality samples. From an initial pool of 56,224 questions, we retain the top 7,587 samples based on rating, difficulty, and verification.

To enable region-specific reasoning, we extract relevant objects and relations per question via lemmatized keyword matching, constructing question-aligned scene subgraphs as localized supervision. This localized supervision helps the model learn where to focus within complex scenes. Bounding box coordinates are retained in absolute pixel space to preserve real-world scale for CIoU-based reward training.

Importantly, our pipeline is scalable and can be extended to generate up to ∼108K samples, the maximum supported by Visual Genome, enabling future large-scale post-training or RL fine-tuning. 3 shows the distribution of QA categories. Full dataset construction details and examples are provided in Appendix A.

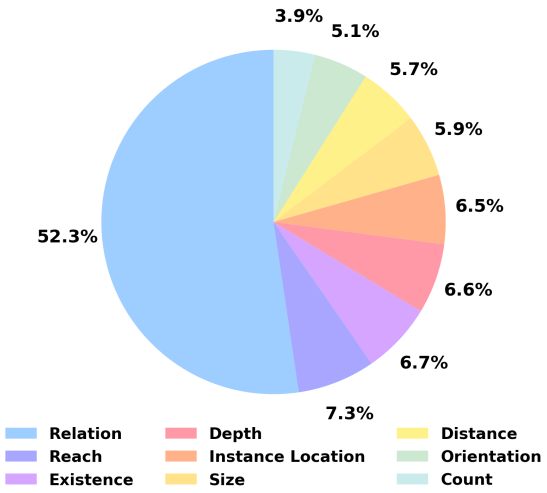

Figure 3: Distribution of STVQA-7K QA types. The dataset covers diverse spatial reasoning aspects, including spatial relations, localization, existence, reach, depth, distance, size, count & orientation.

3.4 Training Details

We build SpatialThinker upon three strong open-source multimodal base models: Qwen2.5-VL-3B, Qwen2.5-VL-7B (Bai et al., 2025b), and Qwen3-VL-30B (Bai et al., 2025a), using them as backbones for policy optimization with RL. No SFT is performed prior to RL training on our STVQA-7K dataset (3.3). We employ GRPO (Shao et al., 2024) as the advantage estimator as described in 3.2, using a rollout size of 8 samples per query and a sampling temperature of 1.0. The models are trained with a maximum context length of 16,384 tokens. The rollout batch size is set to 512, and the global batch size is 128. We train for 75 training steps i.e., 5 training episodes) on 4 $\times$ NVIDIA H100 80GB GPUs. Training time totals $\sim$ 13 hours for the 3B model and $\sim$ 15 hours for the 7B model. To assess scaling, we further apply the same dense-reward GRPO objective to a larger Qwen3-VL-30B-A3B-Instruct backbone Bai et al. (2025a), trained with LoRA (rank 64) via the Tinker API, yielding SpatialThinker-30B (details in Appendix B).

The models are trained on high-resolution image inputs ranging from 512×512 to 2048×2048 pixels, to preserve fine-grained spatial information. All model parameters, including the vision encoder, are updated

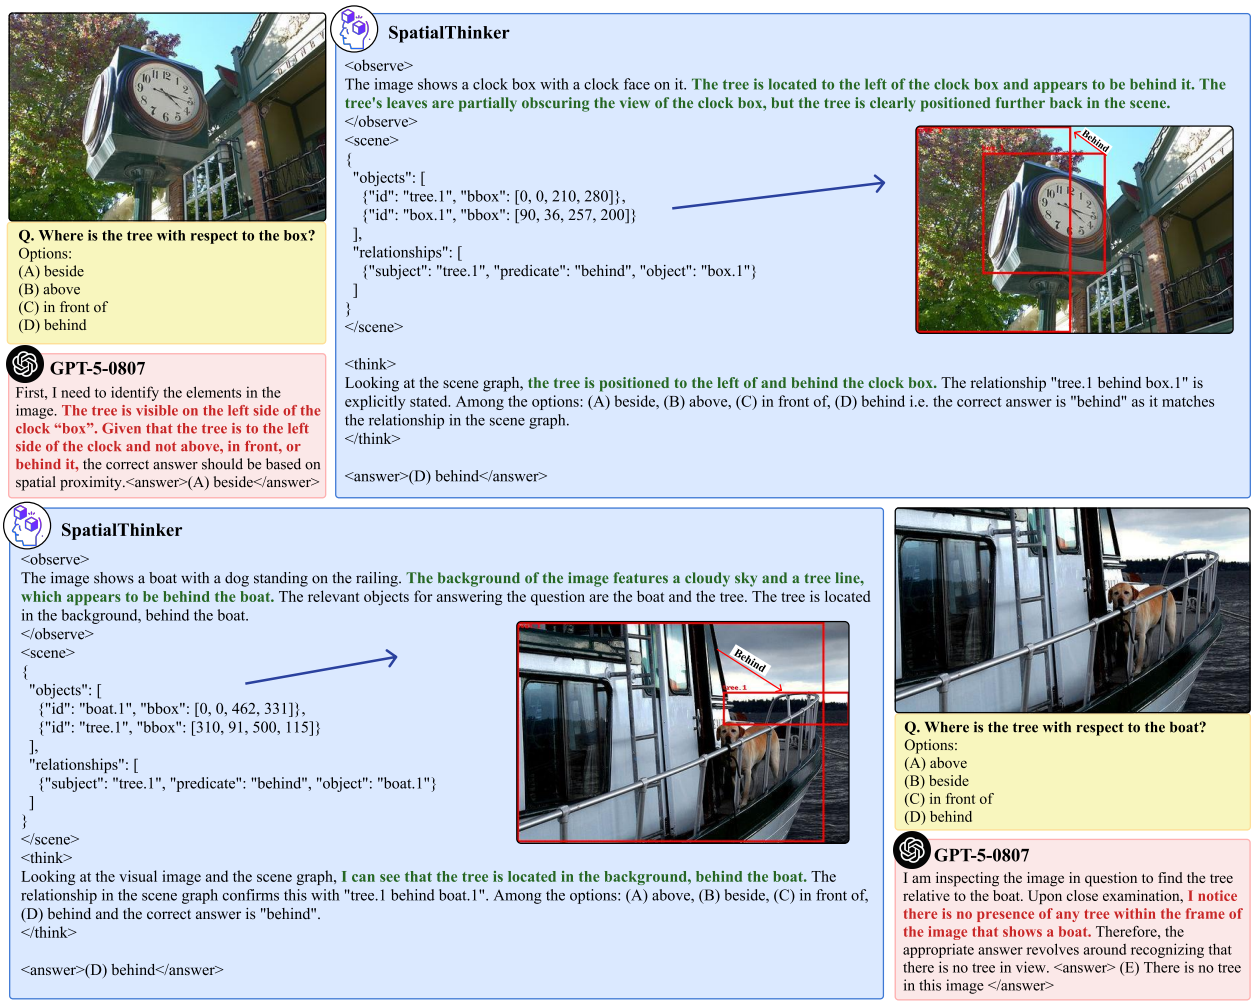

Figure 4: **Qualitative comparison between GPT-5-0807 and SpatialThinker-7B.** GPT-5-0807 often fails to distinguish objects at a *3D relational level*—for example, confusing spatial relations such as *beside*, *behind*, and *in front of*, or missing fine-grained object details.

during training. We use the AdamW optimizer with `bf16` precision, a learning rate of 1×10^{-6}, and weight decay of 1×10^{-2}. The KL penalty coefficient is set to 10^{-2} (Ablation in Appendix E). STVQA-7K is partitioned with a 90/10 train–validation split. Further details on prompts, SFT, and RL training setups, are provided in Appendices B.3, B.4, and B.5, respectively. Finally, Appendix B.5.1 illustrates how each reward component improves steadily under our multi-objective spatial reward, reflecting stable learning dynamics.

Inference Overhead. As SpatialThinker generates question-focused scene *sub*graphs rather than exhaustive scene descriptions, enforced by the count penalty and RoI-filtered supervision (3.1), the structured reasoning adds only modest overhead: on average ∼120 additional tokens for the scene graph component.

4 Experiments

We evaluate SpatialThinker across 14 diverse spatial understanding, and real-world and generic VQA benchmarks, encompassing both 2D and 3D reasoning tasks. Our experiments are guided by two core questions: (Q1) Does our spatial VQA generation pipeline, combined with dense reward RL, improve general spatial and visual reasoning in MLLMs? (Q2) Can MLLMs learn strong spatial capabilities from just 7K synthetic training samples, and how does this compare to models trained on orders-of-magnitude more data?

Table 1: Performance over 2D & 3D Spatial Understanding Benchmarks across different model types. Top-1 & Top-2 accuracies are represented using **bold text**, and underlines.

Model	3DSRBench	CV-Bench		Avg.	BLINK$_{val}$		Avg.
		2D	3D		Spatial Relation	Relative Depth	
Proprietary Models							
GPT-5-0807	**68.2**	**81.4**	90.3	85.8	**90.9**	**81.4**	**86.1**
GPT-4o-0513	44.3	75.8	83.0	79.4	82.5	78.2	80.4
Claude-4-Sonnet-0514	61.9	73.3	84.2	78.7	79.0	78.2	78.6
Claude-3.5-Sonnet-0620	48.2	60.2	71.5	65.9	58.7	67.7	63.2
Open-Source General MLLMs							
Qwen2.5-VL-3B	44.0	59.9	60.2	60.0	66.4	54.0	60.2
Qwen2.5-VL-7B	48.4	69.1	68.0	68.6	84.0	52.4	68.2
Qwen3-VL-30B	60.4	79.0	89.6	84.3	86.0	75.8	80.9
VLAA-Thinker-Qwen2.5-VL-7B	52.2	60.8	60.3	60.6	81.2	71.0	76.1
LLaVA-NeXT-8B	48.4	62.2	65.3	63.8	-	-	-
Cambrian-1-8B	42.2	72.3	72.0	72.2	69.9	73.4	71.7
Open-Source Spatial MLLMs							
RoboPoint-13B	-	-	61.2	-	60.8	61.3	61.1
SpatialBot-3B	41.1	-	69.1	-	67.8	67.7	67.8
SpaceLLaVA-13B	42.0	-	68.5	-	72.7	62.9	67.8
SATORI-R1	48.0	54.6	69.4	62.0	77.0	58.9	68.0
Spatial-RGPT-7B w/ depth	48.4	-	60.7	-	65.7	72.3	69.0
SpaceThinker	51.1	65.1	65.9	65.5	73.4	59.9	66.7
SpaceOm	52.2	72.1	69.3	70.7	81.1	65.3	73.2
Method Comparison (Trained on STVQA-7K)							
Qwen2.5-VL-3B + SFT	50.8	53.9	68.4	61.1	65.0	66.9	66.0
Qwen2.5-VL-3B + Vanilla GRPO	50.1	70.6	66.6	68.6	73.4	55.6	64.5
SpatialThinker-3B (Ours)	52.9	71.0	76.3	73.6	81.8	66.9	74.4
Qwen2.5-VL-7B + SFT	53.6	56.1	71.3	63.7	75.5	64.5	70.0
Qwen2.5-VL-7B + Vanilla GRPO	54.7	68.9	76.5	72.7	80.4	75.0	77.7
SpatialThinker-7B (Ours)	56.4	77.7	78.7	78.2	86.0	72.6	79.3
SpatialThinker-30B (Ours)	62.1	80.3	**93.6**	**87.0**	88.1	79.8	84.0

Benchmarks. We evaluate across 14 benchmarks: eight spatial and six real-world VQA. The spatial suite comprises eight benchmarks including CV-Bench 2D and 3D (Tong et al., 2024a), BLINK Spatial Relations and Relative Depth (Fu et al., 2024), 3DSRBench (Ma et al., 2024b), MMVP (Tong et al., 2024b), SpatialBench (Cai et al., 2024), and SpatialReasonerEval (Ma et al., 2025a); the multi-view MindCube-tiny (Wang et al., 2026); and the held-out validation split of our released STVQA-7K test set, which spans nine spatial reasoning categories: spatial relations, reach and interaction, comparative size, orientation, instance location, depth ordering, distance comparison, object counting, and existence. Together these cover spatial relations, depth, distance, counting, size, orientation, and egocentric multi-view 3D reasoning. To test generalization in real-world, embodied, and generalist VQA contexts, we use six further benchmarks: VStarBench (Wu & Xie, 2023), RealWorldQA (xAI, 2024), MME-RealWorld (Zhang et al., 2024), RoboSpatial-Home (Song et al., 2025) (Configuration and Compatibility only), MM-Star (Chen et al., 2024c), and HallusionBench (Guan et al., 2023), collectively spanning fine-grained visual perception, embodied spatial affordances, structured document and chart understanding, real-world domains (autonomous driving, remote sensing, surveillance), general visual reasoning, and hallucination robustness.

Baselines. We compare against proprietary MLLMs including GPT-5 (gpt-5-0807) (OpenAI, 2025), GPT-4o (gpt-4o-0513) (Hurst et al., 2024), Claude 4 Sonnet (claude-4-sonnet-0514) (Anthropic, 2025), and Claude 3.5 Sonnet (claude-3.5-sonnet-0620) (Anthropic, 2024), open-source generalist models like Qwen2.5-VL (Bai et al., 2025b), Qwen3-VL-30B (Bai et al., 2025a), LLaVA-NeXT (Li et al., 2024b), Cambrian-1 (Tong et al., 2024a), and VLAA-Thinker (Chen et al., 2025a), and spatially-tuned open-source MLLMs such as SpaceLLaVA (AI & Mayorquin, 2025a; Chen et al., 2024a), SpatialRGPT (Cheng et al., 2024), RoboPoint (Yuan et al., 2024), SpaceThinker (AI & Mayorquin, 2025c), SpaceOm (AI & Mayorquin, 2025b), SpatialReasoner (Ma et al., 2025a), SpatialBot (Cai et al., 2024), Visionary-R1 (Xia et al., 2025), and SATORI-R1 (Shen et al., 2025a). We evaluate ablations on variants of our model trained with the STVQA-7K dataset: a supervised fine-tuning (SFT) baseline, and a sparse-reward RL baseline that optimizes only format and accuracy rewards, each weighted equally at 0.5, to isolate the effect of our dense spatial reward.

Evaluation Setting. All models are evaluated in a zero-shot setting using greedy decoding (temp = 0.0). Models default prompting format is used where applicable (e.g., for VLAA-Thinker, SpaceOm, SpaceThinker). SpatialRGPT is evaluated with depth inputs; all other models use RGB. Accuracy is the primary evaluation

Table 2: Performance on additional spatial benchmarks. Top-1 & Top-2 accuracies are represented using **bold text**, and underlines. STVQA-7K$_{\text{val}}$ is the held-out in-distribution validation split.

Model	MMVP	SpatialReasonerEval	SpatialBench	STVQA-7K$_{\text{val}}$	MindCube-tiny
Proprietary Models					
GPT-5-0807	61.7	90.5	62.6	89.7	42.5
GPT-4o-0513	70.7	85.8	67.0	77.0	34.9
Claude-4-Sonnet-0514	71.3	85.7	60.9	80.5	44.8
Claude-3.5-Sonnet-0620	71.3	84.1	63.2	-	-
Open-Source General & Spatial MLLMs					
Qwen2.5-VL-3B	67.0	68.0	49.9	74.3	36.4
Qwen2.5-VL-7B	72.3	70.6	62.5	77.5	35.3
Qwen3-VL-30B	77.9	88.2	68.1	83.7	39.0
VLAA-Thinker-7B	75.3	61.2	66.2	76.2	36.7
SpaceThinker	63.0	69.6	57.9	75.4	36.0
SpaceOm	66.3	68.9	58.6	66.0	33.5
SpatialReasoner	64.0	76.4	59.2	74.0	34.6
SATORI-R1	67.7	70.5	60.3	51.2	36.0
Visionary-R1	70.3	72.9	59.8	74.4	36.0
Method Comparison (Trained on STVQA-7K)					
Qwen2.5-VL-3B + SFT	62.7	67.5	56.3	85.6	35.9
Qwen2.5-VL-3B + Vanilla GRPO	68.3	69.3	56.9	86.7	38.8
SpatialThinker-3B (Ours)	69.0	76.5	61.5	92.5	40.5
Qwen2.5-VL-7B + SFT	68.3	70.8	63.5	84.5	42.2
Qwen2.5-VL-7B + Vanilla GRPO	74.3	79.6	64.2	87.1	44.0
SpatialThinker-7B (Ours)	78.0	82.7	66.4	92.8	45.1
SpatialThinker-30B (Ours)	**79.7**	**92.6**	**69.5**	**93.0**	**45.4**

metric. Full benchmark descriptions, baseline details, and additional implementation specifics are provided in Appendix B.

4.1 Results

We evaluate SpatialThinker across eight spatial and six real-world VQA benchmarks to assess its effectiveness in learning spatial understanding and real-world VQA from limited training data through dense reward supervision.

Performance across Spatial Benchmarks. We evaluate SpatialThinker across eight spatial reasoning benchmarks including 2D relational understanding, 3D spatial alignment, multi-view reasoning, counting, depth ordering, and distance comparison. As shown in Tables 1 and 2, SpatialThinker achieves strong and consistent performance across all spatial tasks. SpatialThinker-30B attains the best results on nearly every benchmark, and SpatialThinker-7B remains the strongest small open-source model. On CV-Bench, SpatialThinker-7B attains an average accuracy of 78.2% across 2D and 3D tasks, approaching GPT-4o's 79.4% and leading all comparably-sized open models, while SpatialThinker-30B reaches 87.0%, exceeding GPT-5 (85.86%) and Claude 4 Sonnet (78.73%); on the 3D split specifically, SpatialThinker-30B attains a best-in-class 93.6% (+3.3% over GPT-5). On 3DSRBench, which requires orientation and multi-object reasoning, SpatialThinker-7B achieves 56.4% (surpassing GPT-4o by +12.1%) and SpatialThinker-30B reaches 62.1%. On BLINK's spatial relation and relative depth tasks, SpatialThinker-7B achieves 86.0% and 72.6% (a 79.3% average, near GPT-4o's 80.4%), and SpatialThinker-30B reaches an 84.0% average, outperforming other spatial MLLMs like Spatial-RGPT-7B (69.0%), which uses depth inputs and 700K training samples.

Despite being trained on just 7K synthetic samples and using only RGB inputs, SpatialThinker-7B consistently outperforms open-source baselines, including VLAA-Thinker-7B, Cambrian-1-8B, Spatial-RGPT, SpaceLLaVA, and RoboPoint-13B, all of which are trained on orders of magnitude more data. Notably, it exceeds specialized spatial models as well: on CV-Bench 3D, it outperforms SpaceLLaVA-13B (78.7% vs. 68.5%), and on BLINK tasks, it surpasses Spatial-RGPT-7B by +10.3%, and SpatialBot by +11.5% despite their reliance on depth information.

On the additional spatial benchmarks (Table 2), SpatialThinker-30B is the strongest model overall, topping MMVP (79.7%), SpatialReasonerEval (92.6%), and SpatialBench (69.5%) and surpassing both GPT-5 and Claude 4 Sonnet on each. SpatialThinker-7B is the best small model, with its 78.0% on MMVP exceeding every baseline, including GPT-4o, Claude 4 Sonnet, and GPT-5 (61.7%). These results highlight the effectiveness of our dense reward design in enabling generalizable spatial reasoning without the need for explicit geometric inputs or large-scale pretraining. Notably, even though training uses only RGB images

Table 3: Performance on VQA and Real-World benchmarks. Top-1 & Top-2 accuracies are represented using **bold text**, and underlines.

Model	MM-Star	VStarBench	RealWorldQA	MME-RealWorld-Lite	RoboSpatial-Home	HallusionBench
Proprietary and Open-Source MLLMs						
GPT-5-0807	58.9	73.3	**78.7**	**57.0**	71.5	73.8
GPT-4o-0513	64.7	66.0	75.4	51.6	68.4	55.0
Claude-4-Sonnet-0514	64.4	60.7	64.0	46.9	69.7	71.2
Claude-3.5-Sonnet-0620	65.1	51.8	60.1	45.2	57.0	55.5
Qwen2.5-VL-3B	55.9	74.9	58.2	41.9	58.7	46.3
Qwen2.5-VL-7B	63.9	75.9	68.4	44.1	70.6	52.9
Qwen3-VL-30B	64.3	81.2	64.8	45.8	53.1	61.5
VLAA-Thinker-7B	63.8	58.1	66.4	44.6	68.9	68.9
SpaceThinker	54.5	56.5	61.6	-	52.6	65.4
SpaceOm	57.7	56.5	53.3	-	68.9	62.9
Method Comparison (Trained on STVQA-7K)						
Qwen2.5-VL-3B + SFT	53.9	73.3	64.8	43.0	69.8	58.9
Qwen2.5-VL-3B + Vanilla GRPO	56.7	74.3	64.4	46.7	64.0	59.0
SpatialThinker-3B (Ours)	57.6	78.0	66.3	46.5	70.6	62.5
Qwen2.5-VL-7B + SFT	63.2	78.0	65.4	47.4	72.4	66.2
Qwen2.5-VL-7B + Vanilla GRPO	63.4	73.9	66.6	46.3	76.2	60.7
SpatialThinker-7B (Ours)	65.9	81.7	69.2	48.3	76.3	66.4
SpatialThinker-30B (Ours)	**66.9**	**85.9**	74.9	49.2	**78.1**	**75.2**

and 2D rewards, the relational scene-graph supervision encodes cues linked to depth and orientation. The RL objective trains the model to maintain geometric–relational consistency, which results in emergent 3D reasoning capabilities. Qualitative examples of model outputs are shown in 4, with additional comparisons in Appendix J. SpatialThinker demonstrates stronger 3D spatial grounding and fine-grained object distinction.

Performance across Real-World and General VQA Benchmarks. We further assess generalization to real-world visual question answering using six diverse benchmarks (Table 3). Compared to the base model, SpatialThinker-7B achieves 65.9% on MM-Star (+2.0%), 81.7% on VStarBench (+5.8% over base, and +15.7% over GPT-4o), and 76.3% on RoboSpatial-Home (+5.7% over base, and +7.9% over GPT-4o), surpassing all open-source and proprietary baselines. On hallucination-sensitive and real-world benchmarks, it scores 66.4% on HallusionBench (+13.5% over base, and 11.4% over GPT-4o), 69.2% on RealWorldQA, and 48.3% on MME-RealWorld-Lite (+4.2%). These results demonstrate that dense spatial rewards improve broad visual perception and not just spatial understanding, enhancing visual reasoning in the wild. We attribute these gains to spatial grounding forcing the model to attend to specific image regions and their geometric properties, which reduces hallucination by anchoring reasoning to visual evidence (+13.5% on HallusionBench), sharpens fine-grained visual distinction (+5.8% on VStarBench), and improves embodied scene understanding (+5.7% on RoboSpatial-Home). Scaling to 30B extends these gains further: SpatialThinker-30B attains the best result among all models on MM-Star (66.9%), VStarBench (85.9%), RoboSpatial-Home (78.1%), and HallusionBench (75.2%, exceeding GPT-5's 73.82%), improving over its Qwen3-VL-30B base by +25.0% on RoboSpatial-Home and +13.7% on HallusionBench.

RL Training with Dense Rewards Enables Superior Generalization. To isolate the contributions of our dense spatial reward design, we compare against two ablation variants: supervised fine-tuning (SFT) and reinforcement learning with sparse rewards (Vanilla GRPO) using only format and answer accuracy, trained on STVQA-7K. As shown in Table 4, SpatialThinker-7B achieves an average accuracy of 70.5% across all 14 benchmarks, exceeding the gains from SFT by +5.5% and the sparse GRPO variant by +3.2%, while essentially matching GPT-5 (−0.9% avg.) and beating GPT-4o, Claude 4 Sonnet, and every open baseline. Scaling the same recipe to 30B, SpatialThinker-30B reaches 74.5%, surpassing GPT-5 by +3.0% and Claude 4 Sonnet by +5.8%. These gains are consistent at 3B, where

Table 4: Avg. accuracy across 14 benchmarks (eight spatial and six real-world with relative gains (Δ).

Model	Avg. Acc. (14)	Δ_{Base}	$\Delta_{\text{GPT-5}}$	$\Delta_{\text{GPT-4o}}$	$\Delta_{\text{Claude 4}}$
Proprietary and Base MLLMs					
GPT-5-0807	71.5	-	-	-	-
GPT-4o-0513	65.8	-	-	-	-
Claude-4-Sonnet-0514	68.7	-	-	-	-
Claude-3.5-Sonnet-0620	60.9	-	-	-	-
Qwen2.5-VL-3B	56.8	-	-	-	-
Qwen2.5-VL-7B	62.8	-	-	-	-
Qwen3-VL-30B	68.1	-	-	-	-
Method Comparison (Trained on STVQA-7K)					
Qwen2.5-VL-3B + SFT	60.7	+3.9	-10.8	-5.1	-8.0
Qwen2.5-VL-3B + Vanilla GRPO	62.0	+5.2	-9.5	-3.8	-6.7
SpatialThinker-3B (Ours)	65.9	+9.1	-5.6	+0.1	-2.8
Qwen2.5-VL-7B + SFT	64.9	+2.1	-6.6	-0.9	-3.8
Qwen2.5-VL-7B + Vanilla GRPO	67.2	+4.4	-4.3	+1.4	-1.5
SpatialThinker-7B (Ours)	70.5	+7.7	-1.0	+4.7	+1.8
SpatialThinker-30B (Ours)	**74.5**	**+6.4**	**+3.0**	**+8.7**	**+5.8**

SpatialThinker-3B outperforms its SFT and GRPO counterparts by +5.2% and +3.9% respectively. Notably, Vanilla GRPO provides only modest improvements over the base model (+4.4 for 7B, +5.2 for 3B), whereas our dense spatial reward raises this to +7.7% and +9.1% (∼1.7×), underscoring the complementary learning signal provided by count and spatial objectives, along with lexicographic reward gating. Overall, these results affirm that structured reinforcement learning with dense spatial supervision significantly enhances the spatial and generic VQA capabilities of multimodal LLMs, with a small fraction of high-quality data.

Reward Design Ablation. To validate our reward formulation, we conduct a controlled ablation study on the STVQA-7K$_{\text{val}}$ set, progressively introducing each component and constraint as shown in Table 5. Naively adding spatial rewards causes a reward hacking behavior (23.7%), as models overgenerate cluttered boxes to exploit the CIoU reward. Introducing the count reward mitigates this (+38% relative gain), regularizing the object and relations count within the scene graph to match the ground-truth quantities. However, rewarding all scene objects biases the model toward exhaustive descriptions. To address this, we shift to local supervision—rewarding only Regions of Interest (RoIs) tied to question-relevant entities—and apply lexicographic gating to ensure spatial rewards are granted only when the final answer is correct, preventing the model from over-optimizing intermediate process rewards at the expense of outcome accuracy. These adjustments recover and slightly exceed baseline performance (76.3%), while additional dataset filtering via pass@2 correctness verification with GPT-4o to retain only high-quality, validated samples (7K) yields a further boost to 87.9%. This staged reward shaping process proves essential for stabilizing optimization and grounding learning in verifiable spatial reasoning. We discuss the full reward design process details in C.

Table 5: Reward ablation on STVQA-7K$_{\text{val}}$.

Reward Components	**STVQA-7K$_{\text{val}}$**
Format + Accuracy	74.9
+ Spatial	23.7
+ Count	61.7
+ Lexicographic Gating & RoI Filtering	76.3
+ Filtered Dataset (pass@2)	**87.9**$_{\text{(+13.0)}}$

Out-of-Distribution Generalization: Dense Rewards Enable Stronger Transfer. While both SFT and sparse-reward GRPO improve spatial reasoning over base models, their ability to generalize to out-of-distribution (OOD) real-world tasks is limited, when compared to SpatialThinker models. As shown in 6, sparse-reward GRPO provides solid spatial gains over its respective base model (+4.4% for 3B, +5.9% for 7B), but offers only marginal improvements on real-world benchmarks (+4.9% and +1.9% respectively), nearly matching or underperforming SFT (+4.6% for 3B, +2.8% for 7B). In contrast, SpatialThinker, trained with dense spatial and count rewards, achieves significantly stronger OOD generalization: +7.6% for 3B and +5.3% for 7B, outperforming baselines at both scales, and SpatialThinker-7B delivers nearly 3× the real-world gains of sparse-reward GRPO (+5.3% vs. +1.9%), highlighting the robustness of our reward design. The combination of structured reasoning formats and lexicographically gated rewards encourages models to internalize spatial priors and compositional patterns that transfer effectively to OOD tasks. Appendix F further demonstrates generalization to abstract and multi-view reasoning tasks.

Table 6: Average accuracy gains (Δ) over respective base models on all 13 held-out benchmarks (seven spatial and six real-world VQA), i.e., the full suite minus the in-domain STVQA-7K$_{\text{val}}$ split.

Model Variant	**Spatial VQA Δ_{Base}**	**Real-World VQA Δ_{Base}**
Qwen2.5-VL-3B + SFT	+2.1	+4.6
Qwen2.5-VL-3B + GRPO	+4.4	+4.9
SpatialThinker-3B	**+9.0**	**+7.6**
Qwen2.5-VL-7B + SFT	+0.9	+2.8
Qwen2.5-VL-7B + GRPO	+5.9	+1.9
SpatialThinker-7B	**+8.6**	**+5.3**

5 Related Work

3D Spatial Reasoning in MLLMs. While MLLMs have advanced core visual tasks (Hurst et al., 2024; Lin et al., 2024; Deitke et al., 2025; Bai et al., 2025b; Du et al., 2025; Li et al., 2024b), their spatial reasoning abilities remain limited (Mirzaee et al., 2021; Tong et al., 2024b; Kamath et al., 2023; Yamada et al., 2023; Li et al., 2024a; Yang et al., 2025a; Ma et al., 2024b). Recent works integrate 3D signals via point clouds or multi-view reconstructions (Hong et al., 2023c;a), or world models with physical priors (Wang et al., 2023; 2024). Large-scale efforts like SpatialVLM (Chen et al., 2024a), SpatialPIN (Ma et al., 2024a), SpatialBot,

(Cai et al., 2024) and SpatialRGPT (Cheng et al., 2024) use millions of 3D-augmented samples or RGB-D scene graphs. Others like MM-Spatial (Daxberger et al., 2025), SpatialLLM (Ma et al., 2025b), and SpaRE (Ogezi & Shi, 2025) similarly scale synthetic or reconstructed 3D data. However, these methods are data-intensive, rely on specialized inputs, or fall short on structured relational modeling. SpatialThinker attains robust relational, and regional reasoning using just 7K VQA samples trained with RL with dense spatial rewards.

Structured Visual Grounding in MLLMs. Scene graphs offer structured object–relation representations and have long supported visual reasoning (Hildebrandt et al., 2020; Wald et al., 2020; Gu et al., 2023; Carion et al., 2020; Cong et al., 2023). Recent LLM-based methods like LLM4SGG and GPT4SGG extract structured graphs from captions (Kim et al., 2024; Chen et al., 2023), while open-vocabulary SGG approaches use MLLMs to generalize beyond fixed ontologies (Chen et al., 2024b; Li et al., 2023). RL-trained SGG models like R1-SGG directly generate scene graphs via dense structural or cognitive rewards (Chen et al., 2025c), emphasizing the value of structured supervision. In parallel, region-aware MLLMs including KOSMOS-2 (Peng et al., 2023), GLaMM (Rasheed et al., 2024), and Ferret (You et al., 2023), enhance spatial grounding via bounding boxes and region-text alignment. SpatialThinker extends these ideas by grounding reasoning in scene subgraphs focused on the question's region of interest, combining structured understanding with reward-guided spatial reasoning.

Multimodal Reinforcement Learning. RL has been increasingly applied to enhance reasoning in MLLMs with verifiable rewards across tasks like math reasoning (Yang et al., 2025b; Meng et al., 2025), classification and grounding (Liu et al., 2025b), semantic segmentation (Liu et al., 2025a), regional understanding (Shen et al., 2025a), and open-vocabulary detection or referring expression comprehension (Pinto et al., 2023; Shen et al., 2025b). Spatial RL has also emerged (Wang & Ling, 2025; Shen et al., 2025b; Ma et al., 2025a), but remain limited to sparse signals like final accuracy or coarse location cues, offering limited support for fine-grained spatial reasoning. SpatialThinker introduces a dense, multi-objective reward framework encompassing regional subgraph construction, object localization, relational grounding, object counting, and final correctness.

6 Conclusion

We introduce SpatialThinker, a MLLM that achieves strong spatial reasoning by combining scene graph grounding with spatial rewards through RL. Trained on just 7K samples, it surpasses proprietary and open-sourced MLLMs on spatial, real-world, and generic VQA benchmarks with improved OOD generalization, outperforming models trained on orders of magnitude more data; SpatialThinker-7B matches GPT-5 and beats GPT-4o on average, while SpatialThinker-30B surpasses both GPT-5 and Claude 4 Sonnet. Dense spatial rewards nearly double the gains of standard RL via GRPO, underscoring the value of rich supervision signals. Our results show that 2D grounding combined with 3D-linked relational supervision is sufficient to induce robust 3D spatial priors. While our approach relies on explicit scene graphs, future work may explore implicit spatial reasoning in latent representations, extend our reward design to spatio-temporal and real-world tasks such as web navigation, and develop unified multi-objective policies and environments across visual reasoning domains.

References

Remyx AI and Salma Mayorquin. Spacellava models. *Hugging Face*, March 2025a. URL `https://huggingface.co/remyxai/SpaceLLaVA`.

Remyx AI and Salma Mayorquin. Spaceom models. *Hugging Face*, 2025b. URL `https://huggingface.co/remyxai/SpaceOm`.

Remyx AI and Salma Mayorquin. Spacethinker models. *Hugging Face*, April 2025c. URL `https://huggingface.co/remyxai/SpaceThinker-Qwen2.5VL-3B`.

Anthropic. Model card addendum: Claude 3.5 haiku and upgraded claude 3.5 sonnet. *Anthropic*, October 2024.

Anthropic. System card: Claude opus 4 & claude sonnet 4. *Anthropic System Cards*, May 2025.

Shuai Bai, Yuxuan Cai, Ruizhe Chen, Keqin Chen, Xionghui Chen, Zesen Cheng, Lianghao Deng, Wei Ding, Chang Gao, Chunjiang Ge, Wenbin Ge, Zhifang Guo, Qidong Huang, Jie Huang, Fei Huang, Binyuan Hui, Shutong Jiang, Zhaohai Li, Mingsheng Li, Mei Li, Kaixin Li, Zicheng Lin, Junyang Lin, Xuejing Liu, Jiawei Liu, Chenglong Liu, Yang Liu, Dayiheng Liu, Shixuan Liu, Dunjie Lu, Ruilin Luo, Chenxu Lv, Rui Men, Lingchen Meng, Xuancheng Ren, Xingzhang Ren, Sibo Song, Yuchong Sun, Jun Tang, Jianhong Tu, Jianqiang Wan, Peng Wang, Pengfei Wang, Qiuyue Wang, Yuxuan Wang, Tianbao Xie, Yiheng Xu, Haiyang Xu, Jin Xu, Zhibo Yang, Mingkun Yang, Jianxin Yang, An Yang, Bowen Yu, Fei Zhang, Hang Zhang, Xi Zhang, Bo Zheng, Humen Zhong, Jingren Zhou, Fan Zhou, Jing Zhou, Yuanzhi Zhu, and Ke Zhu. Qwen3-vl technical report, 2025a. URL `https://arxiv.org/abs/2511.21631`.

Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, Humen Zhong, Yuanzhi Zhu, Mingkun Yang, Zhaohai Li, Jianqiang Wan, Pengfei Wang, Wei Ding, Zheren Fu, Yiheng Xu, Jiabo Ye, Xi Zhang, Tianbao Xie, Zesen Cheng, Hang Zhang, Zhibo Yang, Haiyang Xu, and Junyang Lin. Qwen2.5-vl technical report. *ArXiv*, abs/2502.13923, 2025b.

Wenxiao Cai, Iaroslav Ponomarenko, Jianhao Yuan, Xiaoqi Li, Wankou Yang, Hao Dong, and Bo Zhao. Spatialbot: Precise spatial understanding with vision language models. *arXiv preprint arXiv:2406.13642*, 2024.

Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pp. 213–229. Springer, 2020.

Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14455–14465, 2024a.

Guikun Chen, Jin Li, and Wenguan Wang. Scene graph generation with role-playing large language models. *ArXiv*, abs/2410.15364, 2024b.

Hardy Chen, Haoqin Tu, Fali Wang, Hui Liu, Xianfeng Tang, Xinya Du, Yuyin Zhou, and Cihang Xie. Sft or rl? an early investigation into training r1-like reasoning large vision-language models. *ArXiv*, abs/2504.11468, 2025a.

Kaiyuan Chen, Shuangyu Xie, Zehan Ma, and Ken Goldberg. Robo2vlm: Visual question answering from large-scale in-the-wild robot manipulation datasets, 2025b.

Lin Chen, Jinsong Li, Xiao wen Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, and Feng Zhao. Are we on the right way for evaluating large vision-language models? *ArXiv*, abs/2403.20330, 2024c.

Zuyao Chen, Jinlin Wu, Zhen Lei, Zhaoxiang Zhang, and Changwen Chen. Gpt4sgg: Synthesizing scene graphs from holistic and region-specific narratives. *arXiv preprint arXiv:2312.04314*, 2023.

Zuyao Chen, Jinlin Wu, Zhen Lei, Marc Pollefeys, and Chang Wen Chen. Compile scene graphs with reinforcement learning. *arXiv preprint arXiv:2504.13617*, 2025c.

An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. Spatialrgpt: Grounded spatial reasoning in vision-language models. *Advances in Neural Information Processing Systems*, 37:135062–135093, 2024.

Yuren Cong, Michael Ying Yang, and Bodo Rosenhahn. Reltr: Relation transformer for scene graph generation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(9):11169–11183, 2023.

Erik Daxberger, Nina Wenzel, David Griffiths, Haiming Gang, Justin Lazarow, Gefen Kohavi, Kai Kang, Marcin Eichner, Yinfei Yang, Afshin Dehghan, et al. Mm-spatial: Exploring 3d spatial understanding in multimodal llms. *arXiv preprint arXiv:2503.13111*, 2025.

DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Jun-Mei Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiaoling Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bing-Li Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dong-Li Ji, Erhang Li, Fangyun Lin, Fucong Dai, Fuli Luo, Guangbo Hao, Guanting Chen, Guowei Li, H. Zhang, Han Bao, Hanwei Xu, Haocheng Wang, Honghui Ding, Huajian Xin, Huazuo Gao, Hui Qu, Hui Li, Jianzhong Guo, Jiashi Li, Jiawei Wang, Jingchang Chen, Jingyang Yuan, Junjie Qiu, Junlong Li, Jiong Cai, Jiaqi Ni, Jian Liang, Jin Chen, Kai Dong, Kai Hu, Kaige Gao, Kang Guan, Kexin Huang, Kuai Yu, Lean Wang, Lecong Zhang, Liang Zhao, Litong Wang, Liyue Zhang, Lei Xu, Leyi Xia, Mingchuan Zhang, Minghua Zhang, M. Tang, Meng Li, Miaojun Wang, Mingming Li, Ning Tian, Panpan Huang, Peng Zhang, Qiancheng Wang, Qinyu Chen, Qiushi Du, Ruiqi Ge, Ruisong Zhang, Ruizhe Pan, Runji Wang, R. J. Chen, Ruiqi Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shao-Kang Wu, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wen-Xia Yu, Wentao Zhang, Wangding Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xi aokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyu Jin, Xi-Cheng Shen, Xiaosha Chen, Xiaowen Sun, Xiaoxiang Wang, Xinnan Song, Xinyi Zhou, Xianzu Wang, Xinxia Shan, Y. K. Li, Y. Q. Wang, Y. X. Wei, Yang Zhang, Yanhong Xu, Yao Li, Yao Zhao, Yaofeng Sun, Yaohui Wang, Yi Yu, Yichao Zhang, Yifan Shi, Yi Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yu-Jing Zou, Yujia He, Yunfan Xiong, Yu-Wei Luo, Yu mei You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanping Huang, Yao Li, Yi Zheng, Yuchen Zhu, Yunxiang Ma, Ying Tang, Yukun Zha, Yuting Yan, Zehui Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhen guo Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zi-An Li, Ziwei Xie, Ziyang Song, Zizheng Pan, Zhen Huang, Zhipeng Xu, Zhongyu Zhang, and Zhen Zhang. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *ArXiv*, abs/2501.12948, 2025.

Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, 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.

Yihe Deng, Hritik Bansal, Fan Yin, Nanyun Peng, Wei Wang, and Kai-Wei Chang. Openvlthinker: Complex vision-language reasoning via iterative sft-rl cycles. 2025.

Danny Driess, F. Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Ho Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Peter R. Florence. Palm-e: An embodied multimodal language model. In *International Conference on Machine Learning*, 2023.

Kimi Team Angang Du, Bohong Yin, Bowei Xing, Bowen Qu, Bowen Wang, Cheng Chen, Chenlin Zhang, Chenzhuang Du, Chu Wei, Congcong Wang, Dehao Zhang, Dikang Du, Dongliang Wang, Enming Yuan, Enzhe Lu, Fang Li, Flood Sung, Guangda Wei, Guokun Lai, Han Zhu, Hao Ding, Hao-Xing Hu, Hao Yang, Hao Zhang, Haoning Wu, Haotian Yao, Haoyu Lu, Heng Wang, Hongcheng Gao, Huabin Zheng, Jiaming Li, Jianling Su, Jianzhou Wang, Jiaqi Deng, Jiezhong Qiu, Jin Xie, Jinhong Wang, Jingyuan Liu, Junjie Yan, Kun Ouyang, Liang Chen, Lin Sui, Long Yu, Mengfan Dong, Meng Dong, Nuo Xu, Peng Cheng, Qizheng Gu, Runjie Zhou, Shaowei Liu, Sihan Cao, Tao Yu, Tianhui Song, Tongtong Bai, Wei Song, Weiran He, Weixiao Huang, Weixin Xu, Xiao feng Yuan, Xingcheng Yao, Xingzhe Wu, Xinxing Zu, Xinyu Zhou, Xinyuan Wang, Y. Charles, Yan-Qing Zhong, Yang Li, Yan-Ni Hu, Yanru Chen, Ye-Jia Wang, Yibo Liu, Yibo Miao, Yidao Qin, Yimin Chen, Yiping Bao, Yiqin Wang, Yongsheng Kang, Yuan-Qing Liu, Yulun Du, Yuxin Wu, Yuzhi Wang, Yuzi Yan, Zaida Zhou, Zhaowei Li, Zhejun Jiang, Zheng Zhang, Zhilin Yang, Zhiqi Huang, Zihao Huang, Zijia Zhao, and Ziwei Chen. Kimi-vl technical report. *ArXiv*, abs/2504.07491, 2025.

Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A. Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but not perceive. *ArXiv*, abs/2404.12390, 2024.

Kanishk Gandhi, Ayush Chakravarthy, Anikait Singh, nathan lile, and Noah D. Goodman. Cognitive behaviors that enable self-improving reasoners, or, four habits of highly effective stars. *ArXiv*, abs/2503.01307, 2025.

Jensen Gao, Bidipta Sarkar, F. Xia, Ted Xiao, Jiajun Wu, Brian Ichter, Anirudha Majumdar, and Dorsa Sadigh. Physically grounded vision-language models for robotic manipulation. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 12462–12469, 2023.

Google. Gemini 2.0 flash: Model card. *Technical Report*, April 2025. Published April 15, 2025.

Qiao Gu, Ali Kuwajerwala, Sacha Morin, Krishna Murthy Jatavallabhula, Bipasha Sen, Aditya Agarwal, Corban Rivera, William Paul, Kirsty Ellis, Ramalingam Chellappa, Chuang Gan, Celso de Melo, Joshua B. Tenenbaum, Antonio Torralba, Florian Shkurti, and Liam Paull. Conceptgraphs: Open-vocabulary 3d scene graphs for perception and planning. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 5021–5028, 2023.

Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, Dinesh Manocha, and Tianyi Zhou. Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14375–14385, 2023.

Marcel Hildebrandt, Hang Li, Rajat Koner, Volker Tresp, and Stephan Günnemann. Scene graph reasoning for visual question answering. *ArXiv*, abs/2007.01072, 2020.

Yining Hong, Chunru Lin, Yilun Du, Zhenfang Chen, Joshua B. Tenenbaum, and Chuang Gan. 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, 2023a.

Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36:20482–20494, 2023b.

Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36:20482–20494, 2023c.

Audrey Huang, Wenhao Zhan, Tengyang Xie, Jason D. Lee, Wen Sun, Akshay Krishnamurthy, and Dylan J. Foster. Correcting the mythos of kl-regularization: Direct alignment without overoptimization via chi-squared preference optimization. *ArXiv*, abs/2407.13399, 2024.

Chen Huang, Oier Mees, Andy Zeng, and Wolfram Burgard. Visual language maps for robot navigation. *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 10608–10615, 2022.

Wenxuan Huang, Bohan Jia, Zijie Zhai, Shaoshen Cao, Zheyu Ye, Fei Zhao, Zhe Xu, Yao Hu, and Shaohui Lin. Vision-r1: Incentivizing reasoning capability in multimodal large language models. *ArXiv*, abs/2503.06749, 2025.

Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. Gpt-4o system card. *arXiv preprint arXiv:2410.21276*, 2024.

Physical Intelligence, Kevin Black, Noah Brown, James Darpinian, Karan Dhabalia, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Manuel Y. Galliker, Dibya Ghosh, Lachy Groom, Karol Hausman, Brian Ichter, Szymon Jakubczak, Tim Jones, Liyiming Ke, Devin LeBlanc, Sergey Levine, Adrian Li-Bell, Mohith Mothukuri, Suraj Nair, Karl Pertsch, Allen Z. Ren, Lucy Xiaoyang Shi, Laura Smith, Jost Tobias Springenberg, Kyle Stachowicz, James Tanner, Quan Vuong, Homer Rich Walke, Anna

Walling, Haohuan Wang, Lili Yu, and Ury Zhilinsky. $\pi 0.5$: a vision-language-action model with open-world generalization. *ArXiv*, abs/2504.16054, 2025.

Amita Kamath, Jack Hessel, and Kai-Wei Chang. What's" up" with vision-language models? investigating their struggle with spatial reasoning. *arXiv preprint arXiv:2310.19785*, 2023.

Kibum Kim, Kanghoon Yoon, Jaehyeong Jeon, Yeonjun In, Jinyoung Moon, Donghyun Kim, and Chanyoung Park. Llm4sgg: Large language models for weakly supervised scene graph generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28306–28316, 2024.

Mikhail Konenkov, Artem Lykov, Daria Trinitatova, and Dzmitry Tsetserukou. Vr-gpt: Visual language model for intelligent virtual reality applications. *ArXiv*, abs/2405.11537, 2024.

Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123(1):32–73, 2017.

Chengzu Li, Caiqi Zhang, Han Zhou, Nigel Collier, Anna Korhonen, and Ivan Vulić. Topviewrs: Vision-language models as top-view spatial reasoners. *arXiv preprint arXiv:2406.02537*, 2024a.

Feng Li, Renrui Zhang, Hao Zhang, Yuanhan Zhang, Bo Li, Wei Li, Zejun Ma, and Chunyuan Li. Llava-next-interleave: Tackling multi-image, video, and 3d in large multimodal models. *ArXiv*, abs/2407.07895, 2024b.

Lin Li, Jun Xiao, Guikun Chen, Jian Shao, Yueting Zhuang, and Long Chen. Zero-shot visual relation detection via composite visual cues from large language models. *ArXiv*, abs/2305.12476, 2023.

Lin Li, Wei Chen, Jiahui Li, Kwang-Ting Cheng, and Long Chen. Relation-r1: Progressively cognitive chain-of-thought guided reinforcement learning for unified relation comprehension. *arXiv preprint arXiv:2504.14642*, 2025.

Rongjie Li, Songyang Zhang, Dahua Lin, Kai Chen, and Xuming He. From pixels to graphs: Open-vocabulary scene graph generation with vision-language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 28076–28086, 2024c.

Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. Vila: On pre-training for visual language models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 26689–26699, 2024.

Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *ArXiv*, abs/2304.08485, 2023.

Yuqi Liu, Bohao Peng, Zhisheng Zhong, Zihao Yue, Fanbin Lu, Bei Yu, and Jiaya Jia. Seg-zero: Reasoning-chain guided segmentation via cognitive reinforcement. *ArXiv*, abs/2503.06520, 2025a.

Ziyu Liu, Zeyi Sun, Yuhang Zang, Xiao wen Dong, Yuhang Cao, Haodong Duan, Dahua Lin, and Jiaqi Wang. Visual-rft: Visual reinforcement fine-tuning. *ArXiv*, abs/2503.01785, 2025b.

Chenyang Ma, Kai Lu, Ta-Ying Cheng, Niki Trigoni, and Andrew Markham. Spatialpin: Enhancing spatial reasoning capabilities of vision-language models through prompting and interacting 3d priors. *Advances in neural information processing systems*, 37:68803–68832, 2024a.

Wufei Ma, Haoyu Chen, Guofeng Zhang, Celso M de Melo, Alan L. Yuille, and Jieneng Chen. 3dsrbench: A comprehensive 3d spatial reasoning benchmark. *ArXiv*, abs/2412.07825, 2024b.

Wufei Ma, Yu cheng Chou, Qihao Liu, Xingrui Wang, Celso M de Melo, Jieneng Chen, Jianwen Xie, and Alan L. Yuille. Spatialreasoner: Towards explicit and generalizable 3d spatial reasoning. *ArXiv*, abs/2504.20024, 2025a.

Wufei Ma, Luoxin Ye, Celso M de Melo, Alan Yuille, and Jieneng Chen. Spatialllm: A compound 3d-informed design towards spatially-intelligent large multimodal models. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 17249–17260, 2025b.

Fanqing Meng, Lingxiao Du, Zongkai Liu, Zhixiang Zhou, Quanfeng Lu, Daocheng Fu, Tiancheng Han, Botian Shi, Wenhai Wang, Junjun He, Kaipeng Zhang, Ping Luo, Yu Qiao, Qiaosheng Zhang, and Wenqi Shao. Mm-eureka: Exploring the frontiers of multimodal reasoning with rule-based reinforcement learning. 2025.

Roshanak Mirzaee, Hossein Rajaby Faghihi, Qiang Ning, and Parisa Kordjmashidi. Spartqa:: A textual question answering benchmark for spatial reasoning. *arXiv preprint arXiv:2104.05832*, 2021.

Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, Quan Ho Vuong, Tingnan Zhang, Tsang-Wei Edward Lee, Kuang-Huei Lee, Peng Xu, Sean Kirmani, Yuke Zhu, Andy Zeng, Karol Hausman, Nicolas Manfred Otto Heess, Chelsea Finn, Sergey Levine, and Brian Ichter. Pivot: Iterative visual prompting elicits actionable knowledge for vlms. *ArXiv*, abs/2402.07872, 2024.

Michael Ogezi and Freda Shi. Spare: Enhancing spatial reasoning in vision-language models with synthetic data. *arXiv preprint arXiv:2504.20648*, 2025.

OpenAI. Gpt-5 system card. *OpenAI System Cards*, August 2025.

Yi Peng, Gongrui Zhang, Miaosen Zhang, Zhiyuan You, Jie Liu, Qipeng Zhu, Kai Yang, Xingzhong Xu, Xin Geng, and Xu Yang. Lmm-r1: Empowering 3b lmms with strong reasoning abilities through two-stage rule-based rl. *ArXiv*, abs/2503.07536, 2025.

Zhiliang Peng, Wenhui Wang, Li Dong, Yaru Hao, Shaohan Huang, Shuming Ma, and Furu Wei. Kosmos-2: Grounding multimodal large language models to the world. *arXiv preprint arXiv:2306.14824*, 2023.

André Susano Pinto, Alexander Kolesnikov, Yuge Shi, Lucas Beyer, and Xiaohua Zhai. Tuning computer vision models with task rewards. *ArXiv*, abs/2302.08242, 2023.

Hanoona Rasheed, Muhammad Maaz, Sahal Shaji, Abdelrahman Shaker, Salman Khan, Hisham Cholakkal, Rao M Anwer, Eric Xing, Ming-Hsuan Yang, and Fahad S Khan. Glamm: Pixel grounding large multimodal model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13009–13018, 2024.

Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Jun-Mei Song, Mingchuan Zhang, Y. K. Li, Yu Wu, and Daya Guo. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *ArXiv*, abs/2402.03300, 2024.

Chuming Shen, Wei Wei, Xiaoye Qu, and Yu Cheng. Satori-r1: Incentivizing multimodal reasoning with spatial grounding and verifiable rewards. *ArXiv*, abs/2505.19094, 2025a.

Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun Zhang, Kangjia Zhao, Qianqian Zhang, Ruochen Xu, and Tiancheng Zhao. Vlm-r1: A stable and generalizable r1-style large vision-language model. *ArXiv*, abs/2504.07615, 2025b.

Joar Skalse, Lewis Hammond, Charlie Griffin, and Alessandro Abate. Lexicographic multi-objective reinforcement learning. *ArXiv*, abs/2212.13769, 2022.

Chan Hee Song, Valts Blukis, Jonathan Tremblay, Stephen Tyree, Yu Su, and Stan Birchfield. Robospatial: Teaching spatial understanding to 2d and 3d vision-language models for robotics. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 15768–15780, 2025.

Kexian Tang, Junyao Gao, Yanhong Zeng, Haodong Duan, Yanan Sun, Zhening Xing, Wenran Liu, Kaifeng Lyu, and Kai Chen. Lego-puzzles: How good are mllms at multi-step spatial reasoning? *ArXiv*, abs/2503.19990, 2025.

Gemini Robotics Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montse Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauzá, Michiel Blokzijl, Steven Bohez, Konstantinos Bousmalis, Anthony Brohan, Thomas Buschmann, Arunkumar Byravan, Serkan Cabi, Ken Caluwaerts, Federico Casarini, Os car Chang, José Enrique Chen, Xi Chen, Hao-Tien Lewis Chiang, Krzysztof Choromanski, David D'Ambrosio, Sudeep Dasari, Todor Davchev, Coline Devin, Norman Di Palo, Tianli Ding, Adil Dostmohamed, Danny Driess, Yilun Du, Debidatta Dwibedi, Michael Elabd, Claudio Fantacci, Cody Fong, Erik Frey, Chuyuan Kelly Fu, Marissa Giustina, Keerthana Gopalakrishnan, Laura Graesser, Leonard Hasenclever, Nico-351 las Heess, Brandon Hernaez, Alex Herzog, R. Alex Hofer, Jan Humplik, Atil Iscen, Mithun George Jacob, Deepali Jain, Ryan C. Julian, Dmitry Kalashnikov, Mustafa Emre Karagozler, Stefani Karp, Chase Kew, Jerad Kirkland, Sean Kirmani, Yuheng Kuang, Thomas Lampe, Antoine Laurens, Isabel Leal, Alex X. Lee, Tsang-Wei Edward Lee, Jacky Liang, Yixin Lin, Sharath Maddineni, Anirudha Majumdar, Assaf Hurwitz Michaely, Robert Moreno, Michael Neunert, Francesco Nori, Carolina Parada, Emilio Parisotto, Peter Pastor, Acorn Pooley, Kanishka Rao, Krista Reymann, Dorsa Sadigh, Stefano Saliceti, Pannag R. Sanketi, Pierre Sermanet, Dhruv Shah, Mohit Sharma, Kathryn Shea, Charles Shu, Vikas Sindhwani, Sumeet Singh, Radu Soricut, Jost Tobias Springenberg, Rachel Sterneck, Razvan Surdulescu, Jie Tan, Jonathan Tompson, Vincent Vanhoucke, Jake Varley, Grace Vesom, Giulia Vezzani, Oriol Vinyals, Ayzaan Wahid, Stefan Welker, Paul Wohlhart, Fei Xia, Ted Xiao, Annie Xie, Jinyu Xie, Peng Xu, Sichun Xu, Ying Xu, Zhuo Xu, Yuxiang Yang, Rui Yao, Sergey Yaroshenko, Wenhao Yu, Wentao Yuan, Jingwei Zhang, Tingnan Zhang, Allan Zhou, and Yuxiang Zhou. Gemini robotics: Bringing ai into the physical world. *ArXiv*, abs/2503.20020, 2025.

Shengbang Tong, Ellis L Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *ArXiv*, abs/2406.16860, 2024a.

Shengbang Tong, Zhuang Liu, Yuexiang Zhai, Yi Ma, Yann LeCun, and Saining Xie. Eyes wide shut? exploring the visual shortcomings of multimodal llms. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9568–9578, 2024b.

Jean Vassoyan, Nathanael Beau, and Roman Plaud. Ignore the kl penalty! boosting exploration on critical tokens to enhance rl fine-tuning. *ArXiv*, abs/2502.06533, 2025.

Johanna Wald, Helisa Dhamo, Nassir Navab, and Federico Tombari. Learning 3d semantic scene graphs from 3d indoor reconstructions. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3960–3969, 2020.

Peiyao Wang and Haibin Ling. Svqa-r1: Reinforcing spatial reasoning in mllms via view-consistent reward optimization. *ArXiv*, abs/2506.01371, 2025.

Qineng Wang, Baiqiao Yin, Pingyue Zhang, Jianshu Zhang, Kangrui Wang, Zihan Wang, Jieyu Zhang, Keshigeyan Chandrasegaran, Han Liu, Ranjay Krishna, Saining Xie, Jiajun Wu, Li Fei-Fei, and Manling Li. Mindcube: Spatial mental modeling from limited views, 2026. URL `https://arxiv.org/abs/2506.21458`.

Xingrui Wang, Wufei Ma, Zhuowan Li, Adam Kortylewski, and Alan L Yuille. 3d-aware visual question answering about parts, poses and occlusions. *Advances in Neural Information Processing Systems*, 36: 58717–58735, 2023.

Xingrui Wang, Wufei Ma, Angtian Wang, Shuo Chen, Adam Kortylewski, and Alan Yuille. Compositional 4d dynamic scenes understanding with physics priors for video question answering. *arXiv preprint arXiv:2406.00622*, 2024.

Penghao Wu and Saining Xie. V*: Guided visual search as a core mechanism in multimodal llms. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 13084–13094, 2023.

xAI. Grok-1.5 vision preview. *xAI Blog*, apr 2024.

Jiaer Xia, Y.-F. Zang, Peng Gao, Yixuan Li, and Kaiyang Zhou. Visionary-r1: Mitigating shortcuts in visual reasoning with reinforcement learning. *ArXiv*, abs/2505.14677, 2025.

Tong Xiao, Xin Xu, Zhenya Huang, Hongyu Gao, Quan Liu, Qi Liu, and Enhong Chen. Advancing multimodal reasoning capabilities of multimodal large language models via visual perception reward. *ArXiv*, abs/2506.07218, 2025.

Yutaro Yamada, Yihan Bao, Andrew K Lampinen, Jungo Kasai, and Ilker Yildirim. Evaluating spatial understanding of large language models. *arXiv preprint arXiv:2310.14540*, 2023.

Jihan Yang, Shusheng Yang, Anjali W Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in space: How multimodal large language models see, remember, and recall spaces. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 10632–10643, 2025a.

Yi Yang, Xiaoxuan He, Hongkun Pan, Xiyan Jiang, Yan Deng, Xingtao Yang, Haoyu Lu, Dacheng Yin, Fengyun Rao, Minfeng Zhu, Bo Zhang, and Wei Chen. R1-onevision: Advancing generalized multimodal reasoning through cross-modal formalization. *ArXiv*, abs/2503.10615, 2025b.

Haoxuan You, Haotian Zhang, Zhe Gan, Xianzhi Du, Bowen Zhang, Zirui Wang, Liangliang Cao, Shih-Fu Chang, and Yinfei Yang. Ferret: Refer and ground anything anywhere at any granularity. *arXiv preprint arXiv:2310.07704*, 2023.

Qiying Yu, Zheng Zhang, Ruofei Zhu, Yufeng Yuan, Xiaochen Zuo, Yu Yue, Tiantian Fan, Gaohong Liu, Lingjun Liu, Xin Liu, Haibin Lin, Zhiqi Lin, Bole Ma, Guangming Sheng, Yuxuan Tong, Chi Zhang, Mofan Zhang, Wang Zhang, Hang Zhu, Jinhua Zhu, Jiaze Chen, Jiangjie Chen, Chengyi Wang, Honglin Yu, Weinan Dai, Yuxuan Song, Xiang Wei, Haodong Zhou, Jingjing Liu, Wei Ma, Ya-Qin Zhang, Lin Yan, Mu Qiao, Yong-Xu Wu, and Mingxuan Wang. Dapo: An open-source llm reinforcement learning system at scale. *ArXiv*, abs/2503.14476, 2025.

Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial affordance prediction for robotics. *ArXiv*, abs/2406.10721, 2024.

Yi-Fan Zhang, Huanyu Zhang, Haochen Tian, Chaoyou Fu, Shuangqing Zhang, Jun Wu, Feng Li, Kun Wang, Qingsong Wen, Zhang Zhang, Liang Wang, Rong Jin, and Tien-Ping Tan. Mme-realworld: Could your multimodal llm challenge high-resolution real-world scenarios that are difficult for humans? *ArXiv*, abs/2408.13257, 2024.

Yaowei Zheng, Richong Zhang, Junhao Zhang, Yanhan Ye, Zheyan Luo, and Yongqiang Ma. Llamafactory: Unified efficient fine-tuning of 100+ language models. *ArXiv*, abs/2403.13372, 2024.

Yaowei Zheng, Junting Lu, Shenzhi Wang, Zhangchi Feng, Dongdong Kuang, and Yuwen Xiong. Easyr1: An efficient, scalable, multi-modality rl training framework. *arXiv preprint arXiv:2501.12345*, 2025.

Zhaohui Zheng, Ping Wang, Dongwei Ren, Wei Liu, Rongguang Ye, Qinghua Hu, and Wangmeng Zuo. Enhancing geometric factors in model learning and inference for object detection and instance segmentation. *IEEE Transactions on Cybernetics*, 52:8574–8586, 2020.

Hengguang Zhou, Xirui Li, Ruochen Wang, Minhao Cheng, Tianyi Zhou, and Cho-Jui Hsieh. R1-zero's "aha moment" in visual reasoning on a 2b non-sft model. *ArXiv*, abs/2503.05132, 2025.

Fangrui Zhu, Hanhui Wang, Yiming Xie, Jing Gu, Tianye Ding, Jianwei Yang, and Huaizu Jiang. Struct2d: A perception-guided framework for spatial reasoning in large multimodal models. *ArXiv*, abs/2506.04220, 2025.

Appendix

A STVQA-7K: Dataset Construction

High-quality spatial VQA datasets remain scarce, as most existing benchmarks either lack grounded scene-graph annotations (i.e., explicit spatial coordinates for objects and relations) or fail to comprehensively cover both 2D and 3D spatial reasoning categories. Visual Genome (Krishna et al., 2017) provides dense, human-annotated scene graphs that support strict grounding of both question generation and answer verification within a unified representational framework. Using Visual Genome, we synthetically constructed a spatial visual question answering dataset called SpatialThinker Visual Question Answering dataset i.e., STVQA-7K comprising 7,587 samples, fully grounded in human-annotated scene graphs (Krishna et al., 2017), which we employed for post-training the SpatialThinker models. Importantly, our pipeline is scalable and can be extended to generate up to 108K samples, the maximum supported by Visual Genome, enabling future large-scale post-training or RL fine-tuning.

The original VG150 predicate set is limited to 50 relations, missing several important categories such as positional relations (e.g., left, right, beside), distance-based relations (e.g., near, far, next to), comparative size (e.g., smaller, taller, bigger), orientation (e.g., facing towards/away), and containment (e.g., inside, beneath). To address this gap, we extended the scene graph relation space with an additional 34 predicates, ensuring richer spatial coverage in both 2D and 3D reasoning. Bounding box coordinates are retained in absolute pixel space, rather than normalized values, to preserve real-world scale and spatial alignment, to enable both improved spatial reasoning and effective use of CIoU-based supervision during reward optimization. From these scene graphs, we generate 3D-reasoning VQA, e.g., depth: "*Which is closer to the camera?*" and orientation: "*From the person's perspective, which direction is the dog?*". The dataset construction pipeline proceeds in three stages: (1) synthetic question generation from ground-truth scene graphs, (2) automated quality filtering with external verification, and (3) scene graph adaptation for regional alignment with individual questions.

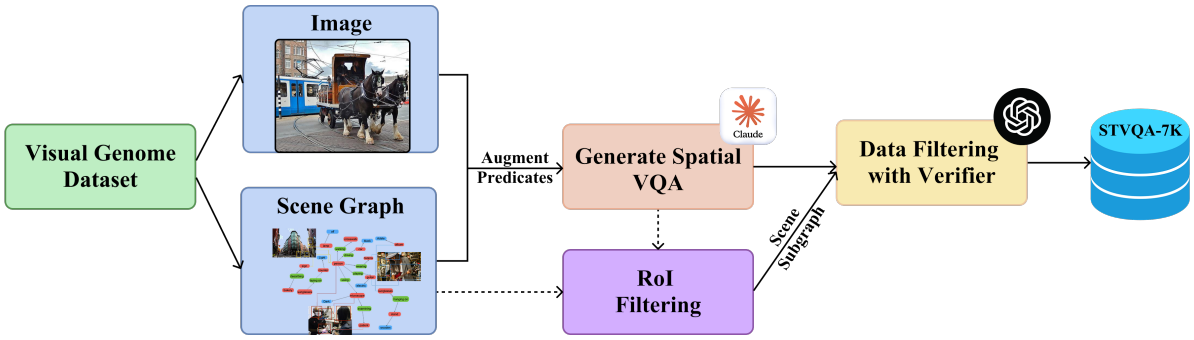

Figure 5: STVQA-7K dataset construction pipeline.

Synthetic Question Generation. Visual Genome scene graphs serve as our foundational ground truth, providing object categories, bounding boxes, and relational triplets for over 150,000 images. We synthetically generate question-answer pairs for a given scene graph data using Claude Sonnet 4 (Anthropic, 2025), synthesizing multiple-choice questions based on the salient objects and meaningful spatial relations explicitly present in each graph. Each question-answer pair is accompanied with a rating generated out of 10 and the difficulty level. Our question generation encompasses nine distinct spatial reasoning categories: spatial relations (above, behind, near, etc.), physical reach and interaction (holding, touching), comparative size, orientation from specific viewpoints, instance location within image frames, depth ordering relative to the camera, distance comparisons to reference objects, object counting, and existence verification. This comprehensive taxonomy spans both 2D and 3D spatial understanding, providing a broad coverage of visual-spatial reasoning capabilities. To promote robust perception, we also include questions involving objects that are partially visible or occluded in the scene, encouraging the model to reason about spatial arrangements and fine-grained details. For each question, we generate a rating out of 10.

Figure 6: Examples of generated QA pairs across the nine spatial reasoning categories in STVQA-7K. Each category highlights distinct reasoning skills, ranging from relative spatial relations and depth ordering to distance, size, orientation, reach, location, count and existence.

Quality Filtering and Validation. To ensure semantic correctness at scale, we implement a consistency-based verification procedure using GPT-4o (Hurst et al., 2024) as an external validation model. To mitigate potential biases from LLM-generated questions, we apply a consistency-based dual-LLM filtering pipeline for quality control: Claude Sonnet 4 generates QA pairs, then GPT-4o (Hurst et al., 2024) validates the ground truth label with pass@2 agreement. Questions that fail this initial consistency check undergo additional evaluation with two supplementary model responses. Items for which all four collected responses disagree with the generated label are discarded as potentially incorrect or ambiguous. This cross-model verification yields a +13% accuracy improvement (Table 5, confirming that dual-LLM validation retains high-quality samples. This filtering process begins with 56,224 initially generated questions by Claude Sonnet 4 (Anthropic, 2025). We select the 10,000 highest-rated samples based on the questions complexity and rating towards its contribution to enhance spatial intelligence as judged by Claude Sonnet 4. Following consistency filtering, we retain 6,895 training samples and 692 validation samples (75%), indicating high label reliability. The final set consists of 50% samples from the relation category, and the remaining 50% distributed across the eight other categories. To prevent positional bias, answers are uniformly distributed across options A, B, C, and D. Figure Figure 3 illustrates the distribution of QA types in STVQA-7K, highlighting the emphasis on spatial relations while maintaining balanced coverage across the remaining reasoning categories. Representative examples of generated QA pairs across the nine spatial reasoning categories are shown in Figure 6, illustrating the diversity of question types in STVQA-7K.

Scene Graph Adaptation. Since each question focuses on specific objects and relationships within the broader scene, we derive question-aligned scene subgraphs that capture only the relevant spatial context. For each question, we extract content words through tokenization and lemmatization to obtain both singular and plural word forms. We then filter the original scene graph to retain only object nodes whose labels appear in the extracted question vocabulary. Relational triplets are preserved when both the subject and object entities are retained and the predicate appears in the question context. The resulting focused scene graph representations enable training the model to generate question-aligned region-of-interest subgraphs, encouraging it to localize attention, ground reasoning in relevant entities and relations, and ultimately learn where to focus within complex visual scenes.

B Experimental Setup Details

This section presents comprehensive evaluations of SpatialThinker across multiple spatial reasoning benchmarks, demonstrating the effectiveness of our multi-objective dense reward design and data-efficient training approach.

B.1 Implementation Details

We build SpatialThinker upon two strong open-source multimodal base models: Qwen2.5-VL-3B and Qwen2.5-VL-7B Bai et al. (2025b), using them as backbones for policy optimization with reinforcement learning. No supervised fine-tuning is performed prior to RL training on our STVQA-7K dataset (section 3.3). We employ GRPO Shao et al. (2024) as the advantage estimator as described in section 3.2, using a rollout size of 8 samples per query and a sampling temperature of 1.0. The models are trained with a maximum context length of 16,384 tokens. The rollout batch size is set to 512, and the global batch size is 128. We train for 75 training steps i.e., 5 training episodes) on 4 $\times$ NVIDIA H100 80GB GPUs. Training time totals around 13 hours for the 3B model and 15 hours for the 7B model.

The models are trained on high-resolution image inputs ranging from 512×512 to 2048×2048 pixels, to preserve fine-grained spatial information. All model parameters, including the vision encoder, are updated during training. We use the AdamW optimizer with `bf16` precision, a learning rate of 1×10^{-6}, and a weight decay of 1×10^{-2}. The KL penalty coefficient is set to 10^{-2}. STVQA-7K is partitioned with a 90/10 train–validation split.

SpatialThinker-30B is trained on the Qwen3-VL-30B-A3B-Instruct backbone Bai et al. (2025a) with the same dense-reward GRPO objective and STVQA-7K data, but using LoRA adapters (rank 64) via the Tinker training API rather than full-parameter updates. It otherwise follows the 7B recipe: AdamW with learning

rate 1×10^{-6}, a rollout group size of 8 samples per query for group-relative advantages, a completion-length cap of 2048 tokens, and a KL penalty (β=10^{-2}) against the frozen base policy.

B.2 Experimental Setup

We evaluate SpatialThinker across a diverse suite of 14 benchmarks: eight spatial (six external sets, the multi-view MindCube-tiny, and our in-distribution STVQA-7K$_{\text{val}}$ split) and six real-world VQA, covering both 2D and 3D understanding aspects to assess fine-grained spatial reasoning capabilities and real-world generalization. We compare against both proprietary and open-source baselines, including models specifically trained for spatial reasoning tasks. Our experiments address two key questions: (Q1) Does our spatial VQA data generation pipeline, combined with dense reward RL, improve MLLMs' general spatial reasoning capabilities? (Q2) How effectively can MLLMs learn spatial understanding from just 7K synthetic training samples, and how does this compare to models trained on orders-of-magnitude larger datasets?

Benchmarks. We evaluate models across eight spatial benchmarks (six external sets, the multi-view MindCube-tiny, and our in-distribution STVQA-7K$_{\text{val}}$ split), and six general-purpose VQA and real-world understanding datasets. The six external spatial benchmarks include CV-Bench Tong et al. (2024a) that measures 2D spatial relations, object counting, depth ordering, and distance reasoning. BLINK's Spatial Relations and Relative Depth tasks Fu et al. (2024) test directional and positional understanding, and fine-grained point-level depth perception, particularly challenging as SpatialThinker receives no explicit point-level supervision during training. 3DSRBench Ma et al. (2024b) assesses egocentric 3D spatial reasoning via relational and multi-object comparisons. MMVP Tong et al. (2024b) examines visual pattern recognition across attributes such as orientation, positional relations, existence, viewpoint, and size. SpatialBench Cai et al. (2024) assesses general spatial comprehension across counting, existence, positional relationships, physical interactions such as reach, and size comparisons. SpatialReasonerEval Ma et al. (2025a) emphasizes depth and distance reasoning within 3D spatial tasks. We further include the multi-view MindCube-tiny Wang et al. (2026), an egocentric benchmark probing spatial reasoning across multiple images of a scene, and the held-out validation split of our released STVQA-7K test set, which measures in-domain spatial mastery across nine reasoning categories (spatial relations, reach and interaction, comparative size, orientation, instance location, depth ordering, distance comparison, object counting, and existence).

To assess broader generalization, we further evaluate models on six diverse real-world benchmarks. VStarBench Wu & Xie (2023) measures accurate localization and recognition of key objects in complex natural scenes. RealWorldQA xAI (2024) requires integrating visual inputs with commonsense and multi-step reasoning for real-world understanding. MME-RealWorld Zhang et al. (2024) spans five challenging domains including optical character recognition in the wild, remote sensing, diagram and table interpretation, autonomous driving, and scene monitoring. RoboSpatial-Home Song et al. (2025) simulates embodied spatial reasoning tasks involving object-object relationships, compatibility, and reference-frame switching (ego-centric, object-centric, and world-centric). We only use Configuration and Compatibility subsets of RoboSpatial-Home. MM-Star Chen et al. (2024c) provides a holistic benchmark covering math, logical reasoning, instance recognition, and fine/coarse visual perception. HallusionBench Guan et al. (2023) evaluates hallucination resistance in multimodal models, requiring accurate visual grounding to counteract entangled linguistic or perceptual illusions. Together, these benchmarks allow us to probe spatial and perceptual reasoning across synthetic, embodied, and naturalistic settings.

Closed-Source MLLM Baselines. Among proprietary models, we evaluate GPT-5 (gpt-5-0807, queried at its default *medium* reasoning effort) OpenAI (2025), GPT-4o (gpt-4o-0513) Hurst et al. (2024), Claude 4 Sonnet (claude-4-sonnet-0514) Anthropic (2025), and Claude 3.5 Sonnet (claude-3.5-sonnet-0620) Anthropic (2024), which represent the current state-of-the-art in commercial multimodal reasoning. These serve as upper bounds for spatial generalization under non-public training regimes.

Open-Source Generalist MLLM Baselines. We compare against generalist open-source MLLMs including Qwen2.5-VL 3B and 7B models Bai et al. (2025b), Qwen3-VL-30B Bai et al. (2025a), LLaVA-NeXT Li et al. (2024b), Cambrian-1 Tong et al. (2024a), and VLAA-Thinker (3B and 7B) Chen et al. (2025a). These models represent state-of-the-art vision-language architectures, offering strong general visual reasoning but without specific spatial tuning.

Open-Source Spatial MLLM Baselines. We benchmark against specialized open-source models designed for spatial reasoning: SpaceLLaVA-13B AI & Mayorquin (2025a); Chen et al. (2024a) – a public re-implementation of SpatialVLM, SpatialRGPT-7B Cheng et al. (2024) incorporates region-level supervision and explicit depth maps into training, RoboPoint-13B Yuan et al. (2024) which instruction-tunes an MLLM to predict image key-point affordances for robotics and spatial affordance tasks, SpaceThinker AI & Mayorquin (2025c), a fine-tuned VLAA-Thinker model for spatial reasoning, and its improved successor SpaceOm AI & Mayorquin (2025b), which incorporates deeper chain-of-thought traces and Robo2VLM data Chen et al. (2025b). Other baselines include SpatialReasoner Ma et al. (2025a) trained with RL and explicit 3D representations, SpatialBot Cai et al. (2024), which integrates RGB and depth inputs for robust spatial perception, Visionary-R1 Xia et al. (2025) which mitigates shortcut learning in visual reasoning by enforcing captioning before reasoning without reliance on chain-of-thought data, and SATORI-R1 Shen et al. (2025a) which decomposes visual question answering into verifiable stages with explicit rewards for improved spatial grounding and reasoning accuracy.

In addition to the above, we compare against our training variants including supervised fine-tuning (SFT) baselines and vanilla GRPO trained with sparse rewards (accuracy and format only) to isolate the contribution of our dense spatial reward framework.

Evaluation Setting. We report accuracy as the primary evaluation metric across all benchmarks. All models are evaluated under zero-shot settings, using greedy decoding (temperature = 0.0, max_new_tokens = 2048) to ensure deterministic and reproducible outputs. For models with specific reasoning templates such as VLAA-Thinker, SpaceThinker, and SpaceOm, we utilize their corresponding structured prompts. In line with their original training setup, SpatialRGPT receives depth inputs, while all other models are evaluated using RGB images alone. Our evaluation pipeline builds upon OpenVLThinker's evaluation framework Deng et al. (2025), adapted to support our new benchmark and dataset formats.

B.3 SpatialThinker Prompt Format

SpatialThinker Prompt

You FIRST observe the image in <observe> </observe> tags, then visualise the relevant scene graph in <scene> </scene> tags, followed by thinking about the reasoning process as an internal monologue within <think> </think> tags and then provide the final answer. The final answer MUST BE put within <answer> </answer> tags, and only return the final choice including the correct option and answer within the answer tags, e.g., <answer> (C) The red cube is left of the green sphere </answer>. Image size: {Width} × {Height}

We use a structured prompt to guide the model through a four-stage reasoning process, explicitly separated using the tags `<observe>`, `<scene>`, `<think>`, and `<answer>`. This format is enforced during training via a binary format reward $R_f \in \{0, 1\}$, with weight $w_{\text{format}} = 0.1$, which verifies the presence, ordering, and validity of all required tags. The `<scene>` section must contain a JSON-encoded subgraph with object IDs, bounding boxes, and relational triplets, while the final answer must be clearly placed within the `<answer>` tags.

Each prompt also includes the input image dimensions in the form `Image size: {Width} × {Height}`, which are dynamically replaced with actual values. Including this information helps the model constrain predicted bounding box coordinates within image bounds, enabling better spatial localization. These coordinates are directly evaluated using IoU-based spatial rewards such as Complete IoU (CIoU), making dimension-aware prediction essential for optimizing structured spatial grounding.

B.4 Details on SFT Training

To establish a comprehensive baseline for comparison with our reinforcement learning approach, we conduct supervised fine-tuning (SFT) experiments using the same base models (Qwen2.5-VL-3B and Qwen2.5-VL-7B) and training dataset (STVQA-7K). Due to compute bottleneck, we do not train the SFT baseline for Qwen3-

VL-30B, and only train it with our spatial reward design. The SFT implementation utilizes LLaMA-Factory framework Zheng et al. (2024) with Low-Rank Adaptation (LoRA) for parameter-efficient fine-tuning.

The training configuration employs LoRA with rank 8 applied to all available modules within the model architecture, enabling comprehensive adaptation while maintaining computational efficiency. Models are trained for 3 epochs totaling 645 training steps, using a context window length of 2048 tokens. We adopt BF16 mixed precision training with a learning rate of 1×10^{-4}, following a cosine learning rate schedule with a warmup ratio of 0.1.

For the SFT experiments, we train models directly on question-answer pairs without intermediate reasoning traces or chain-of-thought prompting. This design choice reflects the practical constraint that generating ground-truth reasoning traces would require additional dataset processing, annotation, and API credits budget. In contrast, reinforcement learning approaches with verifiable rewards (RLVR) naturally enables training with answer supervision alone, as the model learns to generate its own reasoning strategies through environmental feedback rather than imitating pre-specified reasoning patterns.

The SFT baseline serves a critical role in our experimental evaluation, providing direct evidence of the generalization advantages offered by reinforcement learning with dense spatial rewards compared to traditional supervised learning on the same dataset.

B.5 Details on RL Training

We implement reinforcement learning training using the EasyR1 framework Zheng et al. (2025), building upon Qwen2.5-VL-3B, Qwen2.5-VL-7B, and Qwen3-VL-30B as base models without any prior supervised fine-tuning. This direct application of RL to the base models enables us to isolate the effects of reward-driven learning from potential confounding factors introduced by intermediate training stages. Additionally, performing an SFT stage prior to RL would require generating ground-truth reasoning traces, which is limited by API budget. Moreover, explicit reasoning supervision is not strictly necessary—our multi-objective dense spatial rewards encourage the model to acquire structured reasoning and self-reflection abilities directly during RL training.

The training employs Group Relative Policy Optimization (GRPO) Shao et al. (2024) as the advantage estimation method, configured with a rollout size of 8 samples per query at a sampling temperature of 1.0. This configuration balances exploration diversity with computational efficiency, allowing the model to discover multiple reasoning strategies while maintaining stable convergence. The training process utilizes a rollout batch size of 512 and a global batch size of 128, processing data through 75 training steps (approximately 5 training episodes) to achieve convergence. The entire training pipeline runs on 4 $\times$ NVIDIA H100 80GB GPUs, requiring approximately $\sim$ 13 hours for the 3B model and $\sim$ 15 hours for the 7B variant. To assess scaling, we further apply the same dense-reward GRPO objective to a larger Qwen3-VL-30B-A3B-Instruct backbone Bai et al. (2025a), trained with LoRA (rank 64) via the Tinker API, yielding SpatialThinker-30B

To preserve fine-grained spatial information critical for accurate object localization and spatial reasoning, models process high-resolution image inputs ranging from 512×512 to 2048×2048 pixels. The training configuration updates all model parameters including the vision encoder, enabling comprehensive adaptation to spatial reasoning tasks. Optimization employs AdamW with BF16 mixed precision, a conservative learning rate of 1×10^{-6}, and weight decay of 1×10^{-2}. The KL penalty coefficient is set to 10^{-2} to prevent excessive divergence from the base model distribution while allowing sufficient exploration for spatial reasoning strategies. The training utilizes a 90/10 train-validation split of the STVQA-7K dataset, with a maximum context length of 16,384 tokens to accommodate detailed scene descriptions and reasoning traces.

For baseline comparisons, we train vanilla GRPO models (Qwen2.5-VL-3B + Vanilla GRPO and Qwen2.5-VL-7B + Vanilla GRPO) using a simplified reward structure consisting solely of accuracy ($w_{acc} = 0.5$) and format rewards ($w_{format} = 0.5$), with the model trained to output both a reasoning trace and a final answer, following standard chain-of-thought GRPO setups. This configuration represents standard RLVR approaches that rely on sparse final-answer supervision DeepSeek-AI et al. (2025); Shen et al. (2025b); Chen et al. (2025a). Note that we do not train the baseline RL model for Qwen3-VL-30B due to compute bottlenecks, and only train it with our final spatial reward design. The full multi-objective reward design employed for

SpatialThinker training, incorporating format, count, accuracy, and spatial rewards with lexicographic gating, is detailed in section 3.1. The substantial performance improvements of SpatialThinker over vanilla GRPO baselines demonstrate the critical importance of dense spatial supervision in teaching models to perform visually-grounded reasoning.

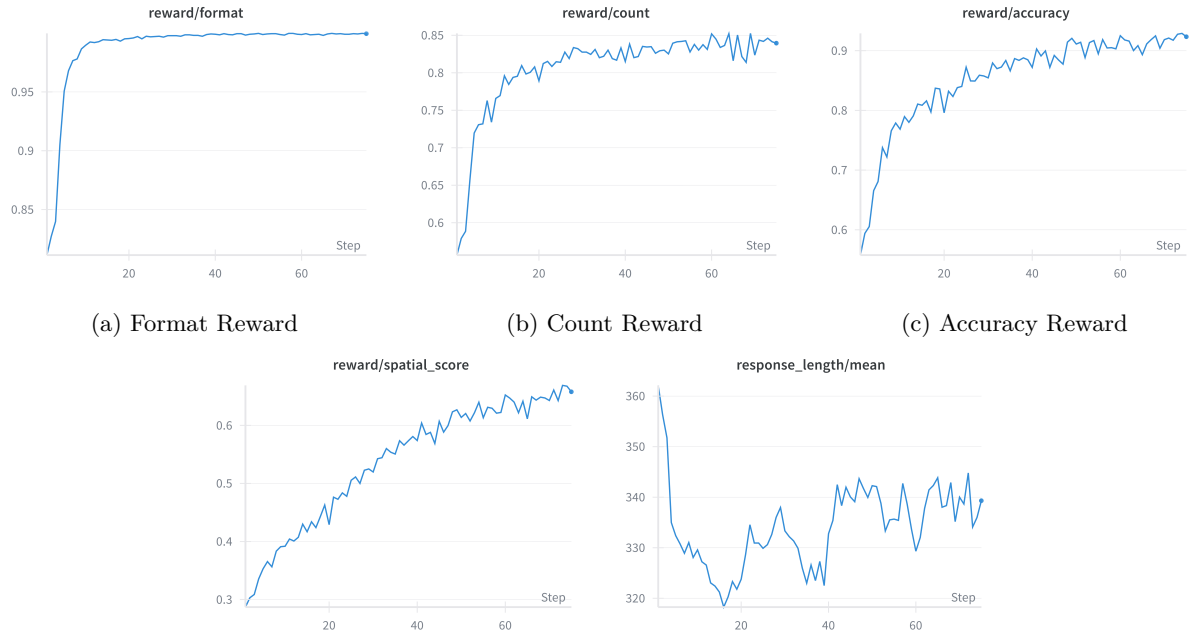

(a) Format Reward (b) Count Reward (c) Accuracy Reward

(d) Spatial Reward (e) Response Length

Figure 7: RL training dynamics of SpatialThinker. All reward components (a–d) improve consistently, reflecting stable optimization. Response length (e) shows a non-monotonic trend, indicating emergent reasoning strategies.

B.5.1 SpatialThinker RL Training Curves

Throughout reinforcement learning, all four reward components: format, accuracy, count, and spatial; demonstrate consistent and interpretable improvement, reflecting stable learning under our lexicographically gated, multi-objective reward structure. The format reward quickly converges early in training, indicating the model learns to produce structurally valid outputs that adhere to the required scene-grounded reasoning format. Accuracy steadily improves across steps, highlighting the model's increasing ability to provide correct answers. Count reward rises consistently, showing that the model learns to focus on predicting only question-relevant objects and relations, rather than describing the entire scene. The spatial reward also improves gradually, indicating better object localization and grounding, as the model increasingly aligns predicted bounding boxes with ground truth annotations. Together, these trends reflect how each reward component scaffolds a different stage of the reasoning process, enforcing structure, correctness, focus, and grounding in tandem.

Response length initially declines, then rises again as it begins producing more deliberate, structured reasoning, signaling an "aha moment" where the model starts to produce more deliberate reasoning traces DeepSeek-AI et al. (2025); Zhou et al. (2025). This emergent behavior suggests the development of internal problem-solving strategies, as the model learns to spend more "thinking time" before answering, consistent with the emergence of self-reflection and structured planning in its spatial reasoning process.

C Reward Design Process

This section details our approach to designing a robust reward system that guides models toward genuine spatial reasoning while preventing degenerate solutions. Our reward design emerged from iterative refinement to address systematic reward hacking behaviors observed during training as shown in Table 5. Early experiments revealed that models readily exploit loopholes in reward functions—particularly when spatial localization rewards were provided without proper constraints. To empirically motivate our design choices, we first present an ablation over successive reward components on the $\text{STVQA}_{\text{val}}$ split, then an ablation over the object-count/relation-count weights inside the count reward, and finally discuss the observed reward-hacking behaviors and the rationale behind each component.

C.1 Reward Design Ablation

To empirically validate our design choices, we conduct a controlled ablation study on the $\text{STVQA-7K}_{\text{val}}$ set, progressively introducing each reward component and constraint. The ablation results support our design rationale by highlighting how each component mitigates specific failure modes. Adding spatial rewards naively without generation constraints, causes performance to collapse by over 50% (from 74.9% to 23.7%), as models exploit the reward by generating cluttered bounding boxes to game the CIoU metric. Introducing the count reward addresses this issue, improving accuracy by 38% relative (to 61.7%), as it constrains overgeneration and forces models to focus on question-relevant elements. However, residual overfitting persists because rewarding spatial alignment across all scene objects biases the model toward exhaustive global descriptions. To address this, we shift from global to local spatial supervision—rewarding only Regions of Interest (RoIs) derived from question-relevant objects and relations—thereby training the model to attend selectively to meaningful spatial cues rather than densely describing the entire scene. Lexicographic gating further ensures that spatial rewards are only applied when the final answer is correct, preventing the model from over-optimizing intermediate process rewards at the expense of outcome accuracy.. Together, these interventions restore and slightly surpass the original performance (76.3%), demonstrating the importance of grounding rewards in both correctness and relevance. Finally, dataset filtering using pass@2 correctness verification amplifies these effects, yielding a substantial gain and culminating in the best validation accuracy of 87.9%. This step ensures that only high-quality, verifiable supervision signals contribute to training, reinforcing the alignment between spatial grounding and task success.

C.2 Count Reward Weight Ablation

Holding the high-level reward weights (w_f, w_c, w_a, w_s)=$(0.1, 0.2, 0.5, 0.2)$ and the rest of the training pipeline fixed (SpatialThinker-3B, lexicographic gating, KL 0.01, 75 steps on STVQA-7K), we ablate the object-count/relation-count split $(\lambda_{\text{obj}}, \lambda_{\text{rel}})$ inside the count reward $R_c = \lambda_{\text{obj}} R_c^{\text{obj}} + \lambda_{\text{rel}} R_c^{\text{rel}}$ (section 3.1). The count reward checks whether the model predicts the same number of objects and the same number of relations as the question-relevant ground truth, with a penalty on both under-generation (too few predicted) and over-generation (too many predicted). This dual penalty also helps localize the output to the right objects and relations: under-generation misses required entities, and over-generation is suppressed by the over-prediction penalty. The $(\lambda_{\text{obj}}, \lambda_{\text{rel}})$ split sets how much of this pressure is applied to the object count versus the relation count.

We compare four splits: 0.7/0.3 (ours), 0.8/0.2, 1.0/0.0, and 0.5/0.5. Putting more weight on the object count helps: our 0.7/0.3 is best at 76.02%, 0.8/0.2 follows at 75.24%, and both 1.0/0.0 and 0.5/0.5 drop to 74.85% (table 7). Object counts are the cleaner signal, since relations are constrained by the objects present, so the object-count penalty curbs object over-generation first and the relation-count penalty then refines the remaining relation set. Both extremes underperform: 1.0/0.0 removes the relation correction entirely, while 0.5/0.5 underweights the dominant object signal.

C.3 Reward Design Rationale

Mitigating Spatial Reward Hacking. Our initial reward formulation, which directly rewarded spatial localization quality, led to unexpected model behavior. Without constraints on generation quantity, models

Table 7: **Count reward object-count/relation-count weight ablation.** High-level weights (w_f, w_c, w_a, w_s)=$(0.1, 0.2, 0.5, 0.2)$, lexicographic gating, and the training/validation pipeline are held fixed (SpatialThinker-3B, STVQA-7K$_{\text{val}}$, n=692). Our 0.7/0.3 split is best.

λ_{obj}	λ_{rel}	Acc.	Format	Count	Spatial
0.7	**0.3**	**76.02**	99.90	75.35	49.48
0.8	0.2	75.24	99.90	74.98	48.61
1.0	0.0	74.85	99.90	71.97	48.53
0.5	0.5	74.85	99.71	77.65	48.54

discovered they could maximize spatial rewards by generating numerous bounding boxes with varying coordinates. Through Hungarian matching that selects the best-matching boxes, even random predictions would occasionally yield high Complete IoU (CIoU) scores. This reward hacking manifested as models producing excessive, hallucinated objects while achieving poor task accuracy—the spatial reward was inflated despite the clutter of irrelevant predictions degrading actual performance. To address this exploitation, we introduced the Count Reward that penalizes deviations from expected object and relation counts. This reward serves dual purposes: (1) preventing reward hacking by constraining the generation space, and (2) encouraging models to focus on question-relevant scene elements rather than exhaustively describing the entire image. The count reward formulation provides a linear penalty proportional to relative deviations from ground truth RoI counts, normalized to prevent domination by scenes with many objects.

Scene Graph Filtering. Another form of overfitting emerged when training with complete Visual Genome scene graphs. Models would memorize exhaustive scene descriptions, including irrelevant background objects, leading to poor generalization. We addressed this by filtering ground truth scene graphs to retain only objects and relations relevant to the given question, focusing supervision on task-critical information.

CIoU over IoU for Spatial Reward. For spatial localization, we adopt Complete IoU (CIoU) instead of standard IoU to compute the spatial reward. Unlike IoU, which returns zero when predicted and ground-truth boxes do not overlap, CIoU provides meaningful gradients by incorporating center distance, aspect ratio, and overlap Zheng et al. (2020). This makes CIoU a denser and more robust supervisory signal during training.

Balancing Supervision with Exploration. Our experiments reveal a crucial insight: models learn simple reward functions significantly faster than complex ones. Tasks with straightforward rewards (e.g., format compliance) show rapid improvements, while multi-component rewards require careful balancing. However, counterintuitively, highly detailed reward functions that attempt to supervise every aspect often degrade performance. Models overfit to maximize minute reward components, converging to template-style answers that score well on individual metrics while losing flexibility. We observed accuracy drops mid-training when rewards became too prescriptive, as models focused on reward optimization rather than genuine task understanding. Effective reinforcement learning requires providing guidance while preserving exploration space. Our final design addresses this by providing soft signals through format checks, count constraints, and accuracy rewards, with spatial localization rewards activated only for correct answers. This maintains the delicate balance between guidance and exploration necessary for robust learning.

Sequential Optimization via Lexicographic Gating. To prevent models from gaming individual reward components at the expense of task accuracy, we implement lexicographic gating Skalse et al. (2022). Rewards are applied in a strict hierarchy: format $\succ$ {count, accuracy} $\succ$ spatial. This forces models to first master output formatting, then simultaneously learn to control generation scope and achieve correctness, before optimizing spatial grounding:

$$R_{\text{total}} = \mathbb{I}[R_{\text{format}} = 1] \cdot \Big(w_{\text{format}} \cdot R_f + w_{\text{count}} \cdot R_c + w_{\text{accuracy}} \cdot R_a + \mathbb{I}[R_{\text{accuracy}} = 1] \cdot w_{\text{spatial}} \cdot R_s \Big)$$

where $\mathbb{I}[\cdot]$ is the indicator function, with weights $w_{\text{format}} = 0.1$, $w_{\text{count}} = 0.2$, $w_{\text{accuracy}} = 0.5$, $w_{\text{spatial}} = 0.2$. This gated design ensures spatial rewards are only applied when the final answer is correct, aligning grounding

Table 8: **SGG correctness vs. final-answer correctness on STVQA-7K$_{\mathbf{val}}$** (n=692, SpatialThinker-7B). *SGG correct* is defined as $J_{\text{trip}}\geq 0.5$ against the ground-truth scene graph. Cells report count (row-conditional %). Even when the predicted scene graph misses most ground-truth triples, the model still answers correctly 79.35% of the time, indicating that scene-graph errors have a measurable but limited downstream impact (Δ=+15.91 pp; Phi ϕ=+0.209; Spearman ρ=+0.221, $p<10^{-8}$).

	Answer		
SGG	✓	×	**Row**
✓ ($J_{\text{trip}}\geq 0.5$)	221 (95.26%)	11 (4.74%)	232
× ($J_{\text{trip}}<0.5$)	365 (79.35%)	95 (20.65%)	460
Col.	586 (84.68%)	106 (15.32%)	692

quality with task success and preventing scenarios where models achieve high spatial scores through precise but irrelevant localizations.

D Scene-Graph Error Propagation Impact

SpatialThinker generates a question-focused scene graph as an intermediate step within its reasoning chain, so errors at this stage could in principle propagate through the reasoning process to the final answer. We measure how strongly scene-graph quality couples to answer correctness on the held-out validation split of STVQA-7K (n=692; section 3.3), which provides ground-truth Visual Genome Krishna et al. (2017) scene graphs alongside each VQA instance; all numbers are computed on SpatialThinker-7B. For each sample we parse the predicted scene graph from the `<scene>...</scene>` block and compute two complementary set-similarity scores against the ground truth: an *object Jaccard* on the set of normalized object classes O (numeric identifiers stripped, e.g. `mountain.1` $\rightarrow$ `mountain`), and a *triple Jaccard* on the set of (subject, predicate, object) relations T under the same normalization,

$$J_{\text{obj}} = \frac{|O_{\text{pred}} \cap O_{\text{gt}}|}{|O_{\text{pred}} \cup O_{\text{gt}}|}, \qquad J_{\text{trip}} = \frac{|T_{\text{pred}} \cap T_{\text{gt}}|}{|T_{\text{pred}} \cup T_{\text{gt}}|}.$$

Object identification saturates ($\bar{J}_{\text{obj}}$=0.957, with 663/692 samples scoring $J_{\text{obj}}\geq 0.5$) and is therefore uninformative as a predictor of downstream accuracy (ϕ=−0.009); the relation-level term carries the discriminative signal, so we adopt $J_{\text{trip}}\geq 0.5$ as our definition of *SGG correct*.

table 8 reports the contingency between SGG correctness and final-answer correctness. Triple-level scene-graph recovery is difficult ($\bar{J}_{\text{trip}}$=0.319 overall, with only 232/692 (33.5%) samples scoring $J_{\text{trip}}\geq 0.5$), and the two outcomes reveal complementary effects. When the predicted scene graph is correct, conditional answer accuracy rises to 95.26%, indicating that a faithful intermediate improves the final answer; when it is incorrect, answer accuracy still holds at 79.35%, for a gap of +15.91 pp (Phi coefficient ϕ=+0.209; Spearman ρ=+0.221 on continuous J_{trip}, $p<10^{-8}$). This signal is consistent across stricter and looser triple-based cuts ($\Delta \in [+15.9, +17.2]$ pp). Scene-graph errors therefore have a measurable but limited downstream impact: a correct scene graph reliably improves the final answer, while the trained reasoning policy can recover from an imperfect intermediate by reasoning over the partial scaffold rather than treating it as a strict dependency.

Taken together, scene-graph generation acts as useful but non-load-bearing scaffolding for SpatialThinker's reasoning chain: when the intermediate is accurate it lifts final-answer correctness by roughly 16 pp, and when it is not the trained policy still answers correctly on about 80% of samples. The +15.91 pp gap further indicates that explicit scene-graph generation contributes positively to correct spatial reasoning. Continued investment in the intermediate, through longer RL training, stronger scene-graph supervision, or richer scene-graph annotations, is thus a promising direction for converting improvements in triple-level recovery into higher final-answer accuracy.

Table 9: Ablation on divergence constraints for SpatialThinker-3B on CV-Bench tasks. KL-regularization with $\beta = 0.01$ yields the highest overall average and strongest 3D reasoning performance.

Model Variant	Count	Relation	Depth	Distance	CV-Bench 2D	CV-Bench 3D	CV-Bench Avg.
SpatialThinker-3B + No KL Penalty	65.5	**76.8**	74.8	70.2	**71.2**	72.5	71.9
SpatialThinker-3B + Chi^2 (0.01)	64.5	73.7	71.2	66.2	69.1	68.7	68.9
SpatialThinker-3B + KL (0.01)	**68.5**	73.5	**79.7**	**72.8**	71.0	**76.3**	**73.7**

Table 10: Results on abstract reasoning benchmarks. Lego Puzzles measure compositional reasoning over object arrangements, while BLINK Multi-View requires integrating multi-view spatial cues.

Model	Lego Puzzles	BLINK Multi-View
Proprietary and Open-Source MLLMs		
GPT-4o-0513	**57.7**	**54.1**
Claude-3.5-Sonnet-0620	53.6	51.9
Qwen2.5-VL-3B	29.9	42.9
Qwen2.5-VL-7B	35.8	44.4
VLAA-Thinker-7B	33.4	51.1
SpaceThinker	31.5	50.4
SpaceOm	32.0	48.9
Method Comparison (Trained on SpatialThinkerVQA)		
Qwen2.5-VL-3B + SFT	34.7	42.1
Qwen2.5-VL-3B + Vanilla GRPO	27.0	45.9
SpatialThinker-3B (Ours)	33.9	45.1
Qwen2.5-VL-7B + SFT	36.6	44.4
Qwen2.5-VL-7B + Vanilla GRPO	29.7	51.9
SpatialThinker-7B (Ours)	37.7	52.6

E Ablation on Divergence Constraints

Recent works such as DAPO Yu et al. (2025); Vassoyan et al. (2025) argue that KL regularization can unnecessarily constrain policy updates and recommend removing the KL penalty entirely to allow freer exploration. In contrast, Huang et al. (2024) revisit divergence regularization and propose using a chi-squared penalty to better control overoptimization. Motivated by these findings, we ablate the effect of different divergence constraints in our reinforcement learning setup for spatial reasoning.

Table table 9 reports results on CV-Bench 2D and 3D tasks Tong et al. (2024a) for three variants of SpatialThinker-3B: (i) no KL penalty, (ii) chi-squared divergence penalty with a coefficient of 0.01, and (iii) our default KL divergence penalty with a coefficient of 0.01. Removing the KL penalty leads to a noticeable drop in performance, particularly on 3D tasks. Using a chi-squared divergence penalty underperforms both the no-penalty and KL variants on several subtasks, especially depth and distance reasoning. The KL-regularized model achieves the best overall performance, yielding a CV-Bench average of 73.7% and providing the strongest results on 3D reasoning tasks.

These findings suggest that a modest KL penalty stabilizes policy updates and prevents reward overoptimization in our spatial reasoning setting, leading to more reliable improvements. While recent language-only alignment work has advocated for removing divergence constraints, our results indicate that retaining a small KL term remains beneficial for multimodal reasoning tasks where stability and coherent spatial grounding are crucial.

F Additional Results: Abstract Reasoning

To further evaluate the generalization capacity of SpatialThinker, we examine its performance on two abstract reasoning benchmarks: **Lego Puzzles** Tang et al. (2025), which test compositional object reasoning and multi-step spatial reasoning, and **BLINK Multi-View** Fu et al. (2024), which requires integrating spatial cues across multiple viewpoints, including visual-spatial understanding and perspective understanding. These tasks are not part of the training distribution and measure the ability of models to extrapolate structured reasoning skills to abstract domains.

Across both tasks, SpatialThinker-7B achieves the highest open-source performance improving over generalist and spatial MLLMs, and scoring 37.7% on Lego Puzzles and 52.6% on BLINK Multi-View, closely approaching GPT-4o and surpassing Claude 3.5 Sonnet on the latter. Interestingly, we observe that vanilla GRPO provides competitive performance on BLINK Multi-View but underperforms on Lego Puzzles, suggesting that dense spatial rewards offer complementary signals that better support compositional reasoning. These results demonstrate that the spatial grounding learned through reinforcement learning transfers to more abstract domains that require compositional and multi-view integration skills.

G Detailed Results: CV-Bench

Table 11: Detailed breakdown of CV-Bench Tong et al. (2024a) results across Count, Relation, Depth, and Distance subtasks.

Model	CV-Bench Tasks				CV-Bench		Avg.
	Count	Relation	Depth	Distance	2D	3D	
Proprietary Models							
GPT-5-0807	68.3	94.5	92.2	88.6	81.4	90.4	85.9
GPT-4o-0513	65.9	85.7	87.8	78.2	75.8	83.0	79.4
Gemini-1.5-Pro	70.4	85.2	82.4	72.8	77.8	77.6	77.7
Claude-4-Sonnet-0514	63.3	83.3	83.0	85.4	73.3	84.2	78.8
Claude 3.7 Sonnet	-	74.2	85.8	84.2	-	85.0	-
Open-Source General MLLMs							
Qwen2-VL-2B	54.7	22.6	16.7	31.7	38.7	24.2	31.5
Qwen2.5-VL-3B	61.5	58.3	67.3	53.0	59.9	60.2	60.1
Qwen2.5-VL-7B	55.9	82.2	70.0	66.0	69.1	68.0	68.6
Qwen3-VL-30B	65.9	92.1	92.0	87.2	79.0	89.6	84.3
VLAA-Thinker-3B	61.6	83.5	53.0	46.8	72.6	49.9	61.3
VLAA-Thinker-7B	47.0	74.6	61.3	59.2	60.8	60.3	60.6
LLaVA-NeXT-34B	-	-	-	-	73.0	74.8	73.9
Mini-Gemini-HD-34B	-	-	-	-	71.5	79.2	75.4
Cambrian-1-34B	-	-	-	-	74.0	79.7	76.9
Open-Source Spatial MLLMs							
Spatial-LLaVA-7B	-	-	57.3	52.2	-	54.8	-
SATORI-R1	42.8	66.3	65.5	73.3	54.6	69.4	62
VisualThinker-R1-2B	59.6	66.8	54.2	56.7	63.2	55.45	59.3
Spatial-RGPT-7B w/ depth	-	-	62.3	59.0	-	60.7	-
RoboPoint-13B	-	75.6	77.8	44.5	-	61.15	-
SpaceThinker-3B	61.0	69.2	70.5	61.3	65.1	65.9	65.5
SpaceLLaVA-13B	-	63.7	66.8	70.2	-	68.5	-
SpatialBot-3B	-	69.4	77.3	60.8	-	69.05	-
Method Comparison (Trained on STVQA-7K)							
Qwen2.5-VL-3B + SFT	30.2	77.5	61.2	75.5	53.9	68.4	61.2
Qwen2.5-VL-3B + Vanilla GRPO	67.5	73.7	64.0	69.2	70.6	66.6	68.6
SpatialThinker-3B (Ours)	68.5	73.5	79.7	72.8	71.0	76.3	73.7
Qwen2.5-VL-7B + SFT	33.3	78.9	64.8	77.7	56.1	71.3	63.7
Qwen2.5-VL-7B + Vanilla GRPO	58.9	78.8	79.3	73.7	68.9	76.5	72.7
SpatialThinker-7B (Ours)	68.7	86.7	81.2	76.2	77.7	78.7	78.2
SpatialThinker-30B (Ours)	67.4	93.2	96.1	91.1	80.3	93.6	87.0

H Detailed Results: 3DSRBench

Table 12: Detailed Breakdown of 3DSRBench Ma et al. (2024b) Height, Location, Orientation, and Multi-Object tasks.

Model	**3DSRBench Tasks**				**Avg.**
	Height	Location	Orientation	Multi-Object	
Proprietary Models					
GPT-5-0807	72.9	79.5	59.0	60.6	68.2
GPT-4o-0513	53.2	59.6	21.6	39.0	44.3
Claude-4-Sonnet-0514	68.6	72.7	48.4	56.8	62.0
Claude-3.5-Sonnet-0620	53.5	63.1	31.4	41.3	48.2
Gemini 2.0 Flash	49.7	68.9	32.2	41.5	49.9
Gemini 2.0 Flash (thinking)	53.0	67.1	35.8	43.6	51.1
Open-Source MLLMs					
Qwen2.5-VL-3B	45.2	56.8	35.7	35.7	44.0
Qwen2.5-VL-7B	44.1	62.7	40.6	40.5	48.4
Qwen3-VL-30B	59.1	75.4	51.0	54.6	60.4
Qwen2.5-VL-72B	53.3	71.0	43.1	46.6	54.9
Cambrian-1-8B	23.2	53.9	35.9	41.9	42.2
LLaVA-NeXT-8B	50.6	59.9	36.1	43.4	48.4
VLAA-Thinker-7B	54.0	60.2	42.9	49.1	52.2
Open-Source Spatial MLLMs					
SpatialBot-3B	40.4	54.4	31.9	33.5	41.1
SpaceLLaVA-13B	49.3	54.4	27.6	35.4	42.0
SpatialLLM-8B	45.8	61.6	30.0	36.7	44.9
SATORI-R1	50.9	52.3	41.9	46.3	48.0
SpatialRGPT-7B w/ depth	55.9	60.0	34.2	42.3	48.4
SpaceThinker-3B	53.1	57.3	41.9	49.6	51.1
Method Comparison (Trained on STVQA-7K)					
Qwen2.5-VL-3B + SFT	51.1	58.3	42.7	48.1	50.8
Qwen2.5-VL-3B + Vanilla GRPO	48.9	57.9	42.5	47.2	50.1
SpatialThinker-3B (Ours)	52.6	61.8	43.4	49.8	52.9
Qwen2.5-VL-7B + SFT	50.6	66.3	43.8	47.9	53.6
Qwen2.5-VL-7B + Vanilla GRPO	54.3	64.7	45.5	50.4	54.7
SpatialThinker-7B (Ours)	52.0	70.3	45.5	50.9	56.4
SpatialThinker-30B (Ours)	62.6	74.9	50.5	54.3	62.1

I Detailed Results: STVQA-7K$_{\text{val}}$

Table 13: Per-category accuracy on the held-out STVQA-7K$_{\text{val}}$ split across the nine spatial reasoning categories. SpatialThinker models lead across nearly all categories, with the largest gains on relational, size, and existence reasoning.

Model	Relation	Reach	Size	Orient.	Location	Depth	Distance	Count	Existence	Overall
Proprietary Models										
GPT-5-0807	89.4	94.0	89.5	88.4	86.9	88.0	**100.0**	87.5	87.1	89.7
GPT-4o-0513	75.8	82.1	64.7	84.7	78.4	72.1	90.1	84.5	80.8	77.0
Claude-4-Sonnet-0514	79.2	78.0	93.8	69.2	82.6	74.0	86.7	81.3	90.3	80.5
Open-Source General & Spatial MLLMs										
Qwen2.5-VL-3B	72.6	80.0	79.2	73.1	78.3	72.0	73.4	78.1	74.2	74.3
Qwen2.5-VL-7B	77.0	80.2	85.6	81.0	80.6	78.2	73.5	65.8	74.4	77.5
Qwen3-VL-30B	84.5	86.0	81.3	80.8	82.6	80.0	83.4	87.5	80.7	83.7
VLAA-Thinker-7B	75.5	70.0	81.3	88.5	80.5	74.0	73.4	71.9	80.7	76.2
SpaceThinker	74.1	77.5	80.3	71.0	85.6	63.5	66.2	75.1	96.6	75.4
SpaceOm	64.3	70.3	73.6	65.4	68.4	64.3	65.6	67.3	66.5	66.0
SpatialReasoner	75.9	72.9	72.1	81.4	66.8	72.9	66.3	71.1	72.0	74.0
SATORI-R1	49.4	59.8	54.9	41.2	54.0	49.8	55.8	57.0	50.0	51.2
Visionary-R1	73.6	80.0	77.1	73.1	78.2	72.0	66.6	78.1	74.2	74.4
Method Comparison (Trained on STVQA-7K)										
Qwen2.5-VL-3B + SFT	85.3	86.3	89.9	85.0	87.3	76.3	87.0	87.8	90.7	85.6
Qwen2.5-VL-3B + Vanilla GRPO	86.0	86.0	89.6	84.6	89.1	82.0	86.7	87.5	**96.8**	86.7
SpatialThinker-3B (Ours)	91.4	93.3	**95.1**	95.4	97.1	**93.3**	92.6	86.8	96.1	92.5
Qwen2.5-VL-7B + SFT	84.5	80.1	85.5	88.6	84.9	78.1	86.8	**93.9**	84.0	84.5
Qwen2.5-VL-7B + Vanilla GRPO	87.3	82.0	87.5	80.7	93.4	82.0	86.6	90.6	93.5	87.1
SpatialThinker-7B (Ours)	93.2	94.0	93.8	88.5	97.9	90.0	90.0	87.5	93.6	92.8
SpatialThinker-30B (Ours)	**93.4**	**94.1**	93.4	**98.5**	**99.2**	80.1	93.4	91.3	93.9	**93.0**

J Additional Qualitative Results

Q. Where is the tower with respect to the train?
Options:
(A) beside
(B) on top of
(C) in front of
(D) behind

GPT-5-0807
In the image, there is a train and a tower is visible. I need to determine the relative position of the tower with respect to the train. The tower is clearly visible **to the right of the train. Based on the options given, the best answer describing the tower's position relative to the train, considering the orientation in the image, is 'in front of the train'.**
<answer> (C) in front of
</answer>

SpatialThinker
<observe>
The image shows a train at a station platform. **The train is on the left side of the image, and there is a tower visible in the background, behind the train. The tower is located further back in the scene, and it is not directly in front of or beside the train.** The tower is also not on top of the train, as it is a separate structure.
</observe>

<scene>
{
"objects": [
{"id": "train.1", "bbox": [0, 72, 474, 204]},
{"id": "tower.1", "bbox": [355, 25, 380, 90]}
],
"relationships": [
{"subject": "tower.1", "predicate": "behind", "object": "train.1"}
]
}
</scene>

<think>
Looking at the visual image and the scene graph, **I can see the tower is positioned further back in the scene, behind the train.** The relationship in the scene graph confirms this with "tower.1 behind train.1". Among the options: (A) beside, (B) on top of, (C) in front of, (D) behind. The correct answer is "behind".
</think>

<answer>(D) behind</answer>

Q. Where is the pole with respect to the building?
Options:
(A) on top of
(B) beside
(C) in front of
(D) behind

GPT-5-0807
To locate the pole with respect to the building, I observe the picture provided. The pole is clearly visible and positioned **at the right side of the building. Considering the options provided, the best description of its position in this context is that the pole is "beside" the building.**
<answer>(B) beside

SpatialThinker
<observe>
The image shows a street scene with a vintage orange tram in the foreground. The tram is on a track, and there is a building in the background. The building has large windows and a modern design. There is a pole visible in the image, which appears to be a streetlight or utility pole, **located in front of the building. The pole is positioned to the right of the tram and is closer to the viewer than the building.** The relevant objects for answering the question are the pole and the building.
</observe>

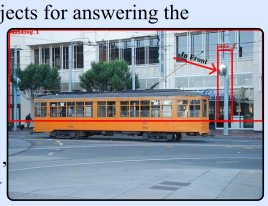

<scene>
{
"objects": [
{"id": "pole.1", "bbox": [410, 40, 440, 181]},
{"id": "building.1", "bbox": [1, 0, 500, 180]}
],
"relationships": [
{"subject": "pole.1", "predicate": "in front of", "object": "building.1"}
]
}</scene>
<think>
Looking at the visual image and the scene graph, **I can see that the pole is positioned in front of the building.** The relationship in the scene graph confirms this with "pole.1 in front of building.1". Among the options: (A) on top of, (B) beside, (C) in front of, (D) behind. Therefore, the correct answer is "in front of".
</think>
<answer>(C) in front of</answer>

Figure 8: **Additional qualitative comparisons between GPT-5-0807 and SpatialThinker.** SpatialThinker shows stronger spatial grounding and fine-grained object distinction, accurately identifying 3D relations that GPT-5-0807 often confuses.