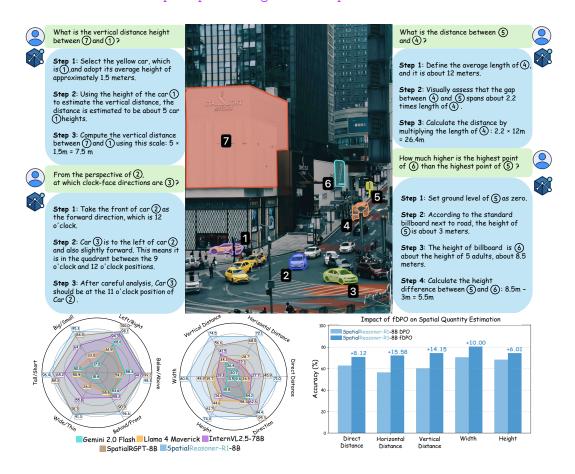


Fine-Grained Preference Optimization Improves Spatial Reasoning in VLMs

Yifan Shen¹, Yuanzhe Liu², Jingyuan Zhu², Xu Cao¹, Xiaofeng Zhang³, Yixiao He¹, Wenming Ye⁴, James M. Rehg¹, Ismini Lourentzou¹

¹University of Illinois Urbana-Champaign ²University of Pennsylvania ³Shanghai Jiao Tong University ⁴Google {yifan26,lourent2}@illinois.edu

https://plan-lab.github.io/spatialreasoner



Abstract

Current Vision-Language Models (VLMs) struggle with fine-grained spatial reasoning, particularly when multi-step logic and precise spatial alignment are required. In this work, we introduce **SpatialReasoner-R1**, a vision-language reasoning model designed to address these limitations. To construct high-quality supervision for spatial reasoning, we design a Multi-Model Monte Carlo Tree Search (M3CTS) method that generates diverse, logically consistent Long Chain-of-Thought (Long-CoT) reasoning trajectories. In addition, we propose a fine-grained Direct Preference Optimization (fDPO) method that introduces segment-specific preference

granularity for descriptive grounding and logical reasoning, guided by a spatial reward mechanism that evaluates candidate responses based on visual consistency, spatial grounding, and logical coherence. Experimental results demonstrate that fDPO achieves relative performance gains of 4.1% and 9.0% over standard DPO on spatial quality and spatial quantity tasks, respectively. SpatialReasoner-R1, trained with fDPO, sets a new SoTA on SPATIALRGPT-BENCH, outperforming the strongest baseline by 9.8% in average accuracy, while maintaining competitive performance on general vision-language tasks.

1 Introduction

Vision-Language Models (VLMs) have demonstrated significant advancements in multimodal understanding tasks, such as image captioning, visual question answering, object detection, and video interpretation [2, 31, 54, 63, 96]. However, their ability to perform spatial reasoning remains limited, especially in scenarios involving complex object arrangements and occlusions [7, 13, 18, 82]. This gap poses a significant challenge for applications such as robotics, autonomous driving, and augmented reality, where robust spatial understanding is essential for effective decision-making [51].

Historically, early VLMs predominantly employed direct-response paradigms [2, 54], *i.e.*, producing immediate answers without explicit reasoning, which often leads to shallow understanding. Recent advances in Chain-of-Thought (CoT) prompting have introduced step-by-step reasoning [73], but standard CoT traces are often too brief or abstract to capture fine-grained spatial logic. In contrast, Long Chain-of-Thought (LongCoT) prompting produces richer, more interpretable reasoning paths that better support comprehension [10, 43, 74]. Still, such prompting must go beyond simple depth estimation, as accurate spatial reasoning requires understanding occlusions, relative orientations, and positional ambiguity, all of which are difficult to capture without structured, fine-grained supervision.

To address these challenges, we introduce **SpatialReasoner-R1**, a novel VLM designed to perform spatial reasoning directly from 2D images. SpatialReasoner-R1 employs structured, interpretable LongCoT reasoning to systematically parse and solve spatial queries without relying on additional modalities or external sensor data. To optimize the training process for multi-step reasoning, we introduce a new **fine-grained Direct Preference Optimization (fDPO)** method that applies differentiated learning updates tailored to two semantically distinct components, descriptive grounding and logical reasoning. Unlike traditional DPO, fDPO introduces segment-specific preference granularity, allowing SpatialReasoner-R1 to adjust its optimization for each generation phase, emphasizing spatial localization during descriptive grounding and enhancing multi-step logical inferences during reasoning.

To curate diverse high-quality spatial reasoning data for training, we propose a **Multi-Model Monte Carlo Tree Search** (M3CTS) that generates high-quality LongCOT responses by leveraging collaborative exploration across multiple VLMs, and a **fine-grained spatial reward mechanism** that evaluates candidate responses across three dimensions: descriptive accuracy, spatial grounding precision, and logical coherence, which are then used to construct positive and negative sample pairs for DPO and fDPO training. Empirical results across several challenging spatial reasoning tasks demonstrate that SpatialReasoner-R1 achieves state-of-the-art performance, significantly outperforming existing VLMs and CoT-based methods, particularly on complex, multi-step spatial reasoning. Specifically, SpatialReasoner-R1 surpasses the best baseline by 9.8% in average accuracy on spatial understanding. Our fDPO improves by 4.1% and 9.0% on average over standard DPO on spatial quality and spatial quantity tasks, respectively. Our contributions are as follows:

- (1) We introduce SpatialReasoner-R1, a LongCoT spatial reasoning VLM that effectively generates interpretable, step-by-step explanations directly from 2D images. SpatialReasoner-R1 establishes a new SoTA in spatial understanding, while maintaining robust performance on general vision-language benchmarks.
- (2) To enhance training stability and precision, we propose a new fine-grained Direct Preference Optimization (fDPO) method that employs segment-specific learning updates tailored explicitly for descriptive grounding and logical reasoning.
- (3) To address the scarcity of high-quality spatial reasoning data, we introduce a data generation pipeline that combines Multi-Model Monte Carlo Tree Search (M3CTS) with fine-grained spatial rewards, enabling the creation of diverse, logically consistent LongCoT trajectories for fine-grained preference training.

2 Related Work

Vision Language Models and Spatial Reasoning. Recent advances in VLMs have significantly enhanced the ability of multimodal models to understand and generate descriptive text grounded in visual contexts [31, 39, 40, 49, 67, 96]. Models such as Flamingo [1], BLIP-2 [32], and Qwen-VL [39] use high-capacity vision encoders [54] paired with LLMs [5, 65] to achieve state-of-the-art performance in various multimodal tasks, such as visual question answering, image captioning, and instruction following [2, 15, 37, 66, 77, 101]. Current trends involve scaling models to improve general understanding [12, 25, 63] and using large-scale instruction tuning datasets [40, 57, 94]. Both proprietary [21, 26, 50] and open-source VLMs [12, 17, 90] have shown impressive results.

While VLMs show promise in visual understanding, accurately perceiving and reasoning about spatial arrangements remains a challenge [13]. Recent efforts to improve spatial understanding include fine-tuning VLMs on spatial VQA datasets [7, 8, 13, 41, 76, 28, 56], and zero-shot frameworks that leverage external 3D foundation models for geometric priors [44]. Region-aware models have also been proposed for better grounding and finer spatial queries [24, 88, 92]. These advances extend to scenarios such as video understanding [82] and 3D generation [46, 51]. To track progress, specialized benchmarks like Q-Spatial Bench [36], SpatialRGPT-Bench [13], VSI-Bench [82], and 3DSRBench [45] have been introduced to assess spatial skills. However, current models still struggle with complex, multi-step spatial reasoning. SpatialReasoner-R1 addresses this gap by introducing fine-grained preference optimization and multi-level reward mechanisms.

Aligning VLMs using Preference Optimization. Preference-based learning methods, particularly DPO [55], have become standard techniques for aligning models with human intentions. These methods bypass the need for explicit reward model training and have often demonstrated strong performance compared to earlier Reinforcement Learning with Human Feedback (RLHF) approaches [3, 19, 48, 99]. In the multimodal domain, DPO and its variants have been adapted to address specific challenges such as reducing hallucinations and improving visual grounding [71, 79, 89]. The adaptability of DPO is further highlighted by its recent application in aligning generative models beyond language, such as text-to-image diffusion models [22, 33, 68, 83, 91]. Adaptation methods often involve constructing preference pairs based on human corrections, AI feedback, or contrasting inputs to guide the model towards desired behaviors [11, 14, 18, 20, 62, 69, 75, 78, 80, 86].

Standard DPO methods treat the reasoning process as a single structure. To address this, preference granularity in DPO has been explored at the token [38, 58, 95, 98, 100], step [29, 97], sentence [52, 55, 59], and turn [60, 61, 81] levels. While effective in certain domains, these approaches overlook the semantic roles of different segments in LongCoT, where descriptive grounding and logical reasoning require distinct optimization. In contrast, our proposed fDPO introduces functional-level preference granularity.

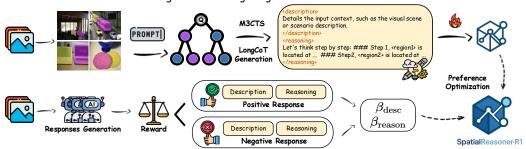
Multi-LLM Guided Reasoning Recent work has explored leveraging multiple LLMs to collaboratively solve complex reasoning tasks, often integrated with Monte Carlo Tree Search (MCTS). Methods such as MoA [70], MoSA [85], AlphaLLM-CPL [72], and LE-MCTS [53] enhance multiagent text-based reasoning using ensemble methods and stepwise search. CoMCTS (Mulberry) [87] extends multi-LLM MCTS to multimodal reasoning, primarily targeting collaborative reflection and error correction. In contrast, our method, M3CTS, addresses the challenge of spatial reasoning in VLMs, introducing fine-grained preference learning and multi-level spatial rewards that allow for coherent, visually-grounded reasoning paths across multimodal data.

3 Method

3.1 Spatial Reasoning from Images

Spatial reasoning is a core vision-language challenge, requiring models to understand visual layouts and perform logical inference over spatial relationships. We define spatial reasoning as a multimodal understanding problem where the goal is to generate accurate reasoning paths based on visual and textual inputs. Formally, a spatial reasoning instance can be represented as a tuple $\mathcal{T} = (I, Q, P) \xrightarrow{\pi_{\theta}} \mathbf{R}$ where I represents the input image containing the visual content, Q is the textual query specifying the spatial reasoning task, P denotes the visual prompt tokens pointing to a specific object or region in the image, and \mathbf{R} the textual response, providing the answer or step-by-step reasoning path. The

Stage 1. Generating LongCoT based on M3CTS



Stage 2. Preference Pair Construction

Stage 3. Fine-grained Direct Preference Optimization

Figure 1: **Method Overview.** To train SpatialReasoner-R1, we (1) generate reasoning paths using M3CTS, (2) construct fine-grained preference pairs via reward-based selection, and (3) train with fine-grained DPO (fDPO) to optimize descriptive and logical reasoning separately.

primary objective of a spatial reasoning model, denoted as π_{θ} , is to map the multimodal input \mathcal{T} to a logically sound and spatially grounded response \mathbf{R} .

Unlike typical direct-response VQA tasks, **SpatialReasoner-R1** is designed to output Long-CoT reasoning traces that decompose spatial reasoning into clear, verifiable steps. To train the model, we introduce a finegrained preference objective that optimizes descriptive and reasoning responses separately (§3.2) and a spatial reward mechanism that evaluates candidate reasoning paths based on spatial and logical understanding (§3.3). Fi-

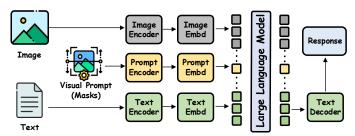


Figure 2: **Architecture Overview.** SpatialReasoner-R1 is a VLM that takes as input a text instruction, visual prompts, and an image, and generates LongCoT reasoning responses.

nally, to address the lack of LongCoT supervision for spatial reasoning, we propose a multi-model collaborative tree search method that generates diverse, reward-aligned reasoning trajectories to enable preference-based training (§3.4). An overview of the proposed training framework is depicted in Figure 1 and the SpatialReasoner-R1 architecture is shown in Figure 2.

3.2 Fine-grained Direct Preference Optimization (fDPO)

We propose fDPO as a novel fine-grained off-policy preference learning algorithm to optimize LongCoT spatial reasoning. Traditional DPO methods apply a single global trade-off parameter β uniformly across all reasoning steps [69, 75, 86], implicitly treating all response segments as equally learnable. However, this can lead to degenerate solutions, as the model may overfit to simpler descriptive responses while under-optimizing the more complex reasoning paths. This observation motivates the design of our fine-grained preference mechanism, which introduces segment-level preference granularity.

To facilitate fine-grained preference optimization, we first segment each LongCoT response R into its constituent description $R_{\rm desc}$ and reasoning $R_{\rm reason}$ components, represented as $R = [R_{\rm desc}, R_{\rm reason}]$. We then quantify the preference signal for each segment by calculating the score difference between the corresponding segments derived from the positive R^p and negative responses R^l , yielding segment-wise preference differentials:

$$\Delta \mathbf{R}_{\text{desc}} = \text{score}(\mathbf{R}_{\text{desc}}^p) - \text{score}(\mathbf{R}_{\text{desc}}^l), \quad \Delta \mathbf{R}_{\text{reason}} = \text{score}(\mathbf{R}_{\text{reason}}^p) - \text{score}(\mathbf{R}_{\text{reason}}^l), \quad (1)$$

where the differentials $\Delta \mathbf{R}_{\text{desc}}$ and $\Delta \mathbf{R}_{\text{reason}}$ quantify the preference margin for description and reasoning segments based on the preference pair, and $score(\cdot)$ composite scores are introduced in Section 3.3. The design of fDPO is guided by two key principles:

Principle 1: Preference optimization strength should be dynamically balanced according to the intrinsic complexity and quality disparity between description and reasoning components.

Our analysis of the fine-grained reward signals reveals that descriptive segments (R_{desc}) are easier to optimize while models struggle with reasoning segments (R_{reason}) that are typically longer and require multi-hop logic. Thus, a unified optimization parameter β may lead to reasoning under-optimization. To address this, fDPO introduces separate, adaptively-tuned trade-off parameters, β_{desc} and β_{reason} , which dynamically control the learning signals for each segment independently, to allow the model to prioritize deeper logical inference while maintaining visual and attribute accuracy.

Principle 2: The choice of segment-specific optimization parameters (β_{desc} and β_{reason}) should prioritize the component that exhibits a larger preference differential, such that learning focuses on harder-to-learn segments.

We further empirically observe that the preference score differential for descriptive components $\Delta R_{\rm desc}$ is consistently smaller than $\Delta R_{\rm reason}$. To account for this, fDPO computes dynamic segment weights $w_{\rm desc}$ and $w_{\rm reason}$ to adaptively adjust the learning signals for the description and reasoning components, respectively:

$$w_s = \frac{\exp(\lambda \cdot \Delta \mathbf{R}_s)}{\exp(\lambda \cdot \Delta \mathbf{R}_{\text{desc}}) + \exp(\lambda \cdot \Delta \mathbf{R}_{\text{reason}})}, \quad s \in \{\text{desc}, \text{reason}\},$$
(2)

where $\Delta \mathbf{R}_s$ is the preference score differential for segment s (either description or reasoning), $\lambda > 0$ controls the sensitivity of weights, and $\{w_{\text{desc}}, w_{\text{reason}}\}$ reflect the relative importance of each segment.

These weights are then mapped to adjustment factors centered around 1 and applied to the base optimization parameter β to yield segment-specific trade-off parameters $\beta_{\rm desc}$ and $\beta_{\rm reason}$ for description and reasoning, respectively:

$$\beta_s = \beta \times f(w_s) = \beta \left[1 + \alpha \left(2w_s - 1 \right) \right], \quad s \in \{ \text{desc, reason} \}, \tag{3}$$

where $w_s \in [0,1]$ is the respective segment-specific weight (either description or reasoning), α is a hyperparameter that controls the maximum scaling amplitude, and β is the base hyperparameter value. This design implements a dynamic learning strategy: segments with larger preference differentials (higher relative importance w_s) receive a higher effective β_s , amplifying the learning signal and prioritizing those components. Conversely, smaller preference differentials yield lower β_s , enabling finer-grained updates. This adaptive mechanism allows fDPO to balance optimization based on segment-specific learning difficulty, promoting better alignment for complex reasoning steps while preserving descriptive accuracy. The final optimization objective for each segment is defined as

$$\mathcal{F}_{s}(\mathcal{T}, \mathbf{R}^{p}, \mathbf{R}^{l}) = \log \frac{\pi_{\theta}(\mathbf{R}_{s}^{p} \mid x)}{\pi_{\text{ref}}(\mathbf{R}_{s}^{p} \mid x)} - \log \frac{\pi_{\theta}(\mathbf{R}_{s}^{l} \mid x)}{\pi_{\text{ref}}(\mathbf{R}_{s}^{l} \mid x)}, \quad s \in \{\text{desc, reason}\}.$$
(4)

Here, \mathcal{T} represents the multimodal input (image, text query, visual prompt), π_{θ} is the model's learned policy, and \mathcal{F}_s measures the segment-specific preference margin in log-likelihood ratios relative to a reference policy π_{ref} . The overall optimization objective for fDPO is

$$\mathcal{L}_{\text{fDPO}}(\theta) = -\mathbb{E}_{(\mathcal{T}, \mathbf{R}^p, \mathbf{R}^l) \sim \mathcal{D}} \left[\log \sigma \Big(\beta_{\text{desc}} \, \mathcal{F}_{\text{desc}}(\mathcal{T}, \mathbf{R}^p, \mathbf{R}^l) + \beta_{\text{reason}} \, \mathcal{F}_{\text{reason}}(\mathcal{T}, \mathbf{R}^p, \mathbf{R}^l) \Big) \right], \quad (5)$$

where $\mathcal{D} = \{(x^{(i)}, y_{\mathrm{desc}}^{(p,i)}, y_{\mathrm{reason}}^{(p,i)}, y_{\mathrm{desc}}^{(l,i)}, y_{\mathrm{reason}}^{(l,i)})\}_{i=1}^N$ denotes the dataset of segment-level preference pairs and $\sigma(\cdot)$ is a sigmoid activation function.

3.3 Fine-Grained Spatial Rewards

To optimize spatial reasoning paths effectively, we introduce a fine-grained reward mechanism that evaluates candidate reasoning paths across visual, spatial, and logical dimensions. Rewards capture alignment with image content, spatial relationships, and logical inference. Figure 3 illustrates the proposed fine-grained spatial rewards for fDPO. Specifically, we define four scalar rewards; details about their formulations and rationale behind each reward are provided in Appendix A.

 \diamond Visual Consistency Reward (\mathcal{R}_{vc}) evaluates the description R_{desc} to ensure spatial grounding and fidelity. The reward verifies key aspects of the quality and alignment of the description with

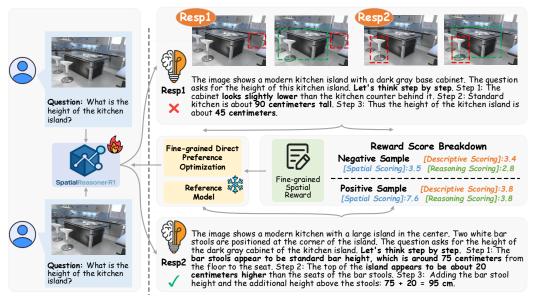


Figure 3: **Fine-Grained Spatial Rewards.** Candidate reasoning paths are decomposed into three aspects, *descriptive*, *spatial*, and *reasoning*, scored separately; the higher value in each row is marked by \checkmark and the lower by \times . **Explanation of Scoring:** *Descriptive:* Negative response omits the two bar-stools and uses generic "modern kitchen" wording, whereas the positive response lists every salient object; *Spatial:* Negative response wrongly claims the island is *lower* than the rear counter and ignores the 20cm offset revealed by the stool reference, whereas the positive response provides its estimate to the 75cm stool height plus that offset; *Reasoning:* Negative response uses an illogical "half-height" heuristic $90\text{cm} \rightarrow 45\text{cm}$ without intermediate computation, whereas the positive response explicitly adds reference height and gap (75cm + 20cm = 95cm). These per-category deficits yield lower composite reward, designating the upper response as negative sample.

the visual scene, such as whether all referenced objects are present and identifiable, whether the stated properties (such as color, size, and shape) match the visual content, whether the description includes all necessary details prompted by the query, and whether it remains contextually appropriate and free from extraneous information.

- \diamond **Depth-Guided Spatial Reward** (\mathcal{R}_{sp}) measures fine-grained spatial understanding by leveraging depth information. This reward is independently computed for the description R_{desc} and reasoning R_{reason} components, with two adaptive weighting mechanisms: an uncertainty weight that adjusts the score for spatial expressions with qualifiers (*e.g.*, "approximately","possibly") to account for reduced confidence, and a context-aware weight that emphasizes spatial relations directly relevant to the query. The final spatial rewards are computed as the uncertainty and context-aware weighted correctness scores across all spatial assertions in the corresponding description and reasoning components, validated against both the RGB image and its corresponding depth map. This ensures that more confident and contextually aligned relations have stronger influence on the reward.
- \diamond Logical Coherence Reward (\mathcal{R}_{lc}) evaluates the reasoning R_{reason} for structural integrity and logical correctness. This reward captures multi-hop inference and factual alignment by verifying that premises are consistent with the image, depth map, and preceding descriptions, reasoning steps maintain spatial and causal logic, the application of physical, spatial, and logical principles remains coherent throughout, and the conclusion is fully supported by the reasoning chain.

The preference differentials in fDPO are computed from composite rewards that aggregate fine-grained evaluations across segments: $score(\mathbf{R}_{desc}) = \mathcal{R}_{vc} + \mathcal{R}_{sp,desc}$ and $score(\mathbf{R}_{reason}) = \mathcal{R}_{lc} + \mathcal{R}_{sp,reason}$.

3.4 Multi-Model MCTS (M3CTS)

We introduce a Multi-Model Monte Carlo Tree Search (M3CTS) framework for generating high-quality LongCoT data $\mathcal{D} = \{(\mathcal{T}, \mathbf{R}^p, \mathbf{R}^l)\}_{i=1}^N$ tailored to spatial reasoning. Inspired by DeepSeek-R1-Zero [16] and prior multimodal MCTS methods [87, 53, 85], M3CTS explores diverse reasoning

trajectories across multiple large language models (LLMs) to effectively search for logical, spatially-consistent explanations that satisfy the query. Formally, the reasoning process is defined as a sequence of reasoning states: $S = \{s_0, \ldots, s_t, \ldots, s_T\}$, where $s_t \in S$ represents a partial reasoning state, s_0 is the initial state derived from T, and s_T is a terminal state corresponding to a fully reasoned path. M3CTS operates through four key stages: *Expand, Simulate, Backprop*, and *Select*.

Expand. At each step t, M3CTS expands the current state s_t by generating diverse candidate reasoning states S_c using multiple VLMs $\{\pi_k\}_{k=1}^K$ concurrently, *i.e.*,

$$S_{c} = \bigcup_{k=1}^{K} \pi_{k} \Big(s_{t} \mid \mathcal{T}, \operatorname{Parent}(s_{t}) \Big), \tag{6}$$

where π_k is the k-th VLM, \mathcal{T} multimodal input, and Parent(s_t) ancestor reasoning states of s_t . To ensure consistency, we enforce a structured output format across all VLMs (Appendix B).

Simulate. Each candidate $s_{k,t} \in \mathcal{S}_c$ generated during expansion is evaluated based on three distinct criteria: (i) visual description accuracy against the original image, (ii) spatial correctness of inferred spatial relationships utilizing both original and depth-derived images, and (iii) logical coherence of the textual reasoning steps. The evaluation score $R(s_{k,t})$ is computed as:

$$R(s_{k,t}) = \frac{1}{M} \sum_{m=1}^{M} \left[\mathbb{I}_{\text{visual}}^{(m)}(s_{k,t}) + \mathbb{I}_{\text{spatial}}^{(m)}(s_{k,t}) + \mathbb{I}_{\text{logical}}^{(m)}(s_{k,t}) \right], \tag{7}$$

where M is the number of evaluation models, and each $\mathbb{I}_{\text{eval}}^{(m)}(s_{k,t})$ indicator function defined as:

$$\mathbb{I}_{\text{eval}}^{(m)}(s_{k,t}) = \{+1 \text{ (fully accurate)}, 0 \text{ (neutral)}, -1 \text{ (inaccurate)}\}.$$
(8)

We preserve high-quality paths by pruning the candidate set according to the evaluation score, *i.e.*, $S_c^* = \{s_{k,t} \mid R(s_{k,t}) \geq 0\}$. Appendix C provides detailed descriptions of the evaluation.

Backprop. To perform credit assignment, scores from the simulation phase are recursively propagated upwards through the search tree. The objective is to update the value estimates $V(s_{k,t})$ and visit counts $N(s_{k,t})$ for each parent node $s_{k,t}$ based on the performance of its children $\mathcal{S}_{c}^{*} = \text{Child}(s_{k,t})$,

$$V(s_{k,t}) \leftarrow \frac{N(s_{k,t})V(s_{k,t}) + \sum_{s_c \in \mathcal{S}_c^*} N(s_c)R(s_c)}{N(s_{k,t}) + \sum_{s_c \in \mathcal{S}_c^*} N(s_c)}, \quad N(s_{k,t}) \leftarrow N(s_{k,t}) + \sum_{s_c \in \mathcal{S}_c^*} N(s_c). \quad (9)$$

Select. This phase is responsible for choosing the most promising candidate state for further exploration in the next iteration of tree expansion. We use the Upper Confidence Bound (UCB) strategy to select the next state $s_{k',t+1}^{\star}$ to traverse, based on updated values and visitation statistics. UCB ensures that high-value paths are prioritized, while also exploring less-visited nodes to discover new reasoning trajectories. The candidate selected maximizes the UCB objective, *i.e.*,

$$s_{k',t+1}^{\star} = \underset{s_c \in \mathcal{S}_c^{\star}}{\operatorname{arg\,max}} \left[V(s_c) + \alpha \sqrt{\frac{\log N(s_{k,t})}{1 + N(s_c)}} \right],$$
 (10)

where $V(s_c)$ is the value estimate of the candidate state s_c , $N(s_c)$ its visit count, and $\alpha > 0$ is a hyperparameter that balances exploration vs. exploitation.

4 Experiments

4.1 Experimental Setup

We evaluate SpatialReasoner-R1 across diverse spatial reasoning and general vision-language established benchmarks to assess the model's fine-grained spatial understanding and logical reasoning capabilities. Implementation details are provided in Appendix D.

Spatial Reasoning Benchmarks. Our primary benchmark is SPATIALRGPT-BENCH [13], comprising image-based spatial reasoning questions and their corresponding ground truth answers. Detailed descriptions of benchmarks and evaluation protocols are provided in Appendix E.

Table 1: **Spatial Reasoning Success Rates** (†) **on SPATIALRGPT-BENCH**. Classification (top) and numeric distance/direction (bottom). are General Large VLMs, are Customized VLMs, are **SpatialReasoner-R1** variants. "/" indicates the model refuses to provide a response for that metric.

	Below/ Above	Left/ Right	Big/ Small	Tall/ Short	Wide/ Thin	Behind/ Front	Qual. Acc.
Gemini 2.0 Flash [21]	58.33	68.57	16.98	50.00	15.38	53.63	44.29
Llama 4 Maverick [47]	54.17	61.90	33.02	50.89	25.96	55.45	47.18
Gemini 1.5 Pro [64]	85.83	56.19	58.49	71.42	55.76	60.00	65.14
ChatGPT-4o [50]	87.50	80.00	53.77	63.39	51.92	60.90	66.67
SpatialBot-3B [6]	52.50	62.86	57.54	49.11	49.04	62.73	55.56
SpaceThinker Qwen2.5VL-3B [4]	89.16	63.81	76.41	56.25	56.73	70.91	69.25
InternVL2.5-78B [12]	94.16	94.28	64.15	65.17	55.76	58.18	72.29
Sa2VA 4B [90]	22.50	25.71	25.47	16.07	27.88	30.91	24.65
Sa2VA 8B [90]	50.00	39.04	45.28	26.78	45.19	53.63	43.37
SpatialRGPT-8B [13]	99.17	100.00	84.90	89.28	91.34	90.90	92.69
SpatialReasoner-R1 SFT 4B	79.16	78.09	55.66	66.96	59.61	75.45	69.41
SpatialReasoner-R1 SFT 8B	81.66	81.90	75.47	75.89	79.80	83.63	79.75
SpatialReasoner-R1 DPO 4B	91.66	91.42	69.81	65.17	71.15	85.45	79.29
SpatialReasoner-R1 DPO 8B	94.16	93.33	89.62	90.18	88.64	92.27	91.48
SpatialReasoner-R1 fDPO 4B	95.83	93.33	83.96	74.10	87.50	89.09	87.37
SpatialReasoner-R1 fDPO 8B	98.33	<u>98.10</u>	95.28	96.43	91.34	93.64	95.59

	Direct Distance	Horizontal Distance	Vertical Distance	Width	Height	Direction	Quan. Acc.
Gemini 2.0 Flash [21]	9.45	10.65	26.41	10.52	30.82	54.20	22.43
Llama 4 Maverick [47]	24.48	28.68	34.28	35.71	44.61	58.09	36.72
Gemini 1.5 Pro [64]	14.18	17.21	14.15	19.54	36.09	30.84	21.90
ChatGPT-4o [50]	/	/	/	/	/	60.75	/
SpatialBot-3B [6]	6.00	15.51	8.00	10.52	18.75	39.00	15.62
SpaceThinker Qwen2.5VL-3B [4]	24.32	17.21	59.43	23.27	23.62	32.35	28.97
InternVL2.5-78B [12]	27.70	22.13	41.50	29.32	34.58	62.61	35.25
Sa2VA 4B [90]	13.51	15.57	19.81	13.53	12.03	10.28	14.02
Sa2VA 8B [90]	14.18	14.75	9.43	14.28	19.54	14.18	14.55
SpatialRGPT-8B [13]	45.90	<u>68.00</u>	56.60	48.90	61.70	95.30	61.42
SpatialReasoner-R1 SFT 4B	22.29	27.86	31.13	25.56	33.80	47.66	30.71
SpatialReasoner-R1 SFT 8B	28.43	20.49	44.05	33.59	51.63	46.72	37.12
SpatialReasoner-R1 DPO 4B	47.97	46.72	60.37	45.11	55.63	91.58	56.61
SpatialReasoner-R1 DPO 8B	62.83	56.55	60.37	70.45	68.42	93.45	68.22
SpatialReasoner-R1 fDPO 4B	60.13	59.01	<u>71.70</u>	65.41	57.89	92.52	66.76
SpatialReasoner-R1 fDPO 8B	70.95	72.13	74.52	80.45	74.43	<u>94.39</u>	77.30

General Vision-Language Benchmarks. To validate the robustness of SpatialReasoner-R1 beyond purely spatial tasks, we evaluate on broader vision-language datasets such as MME, POPE, SEED-BENCH, AI2D, SQA-TEST, MMMUV, MMSTAR, and HALLUSIONBENCH [35, 34, 30, 27, 42, 93, 9, 23]. These datasets cover fundamental vision-language tasks such as object grounding, hierarchical scene parsing, multimodal understanding, and multi-turn reasoning in diverse multimodal contexts.

Baselines. We benchmark SpatialReasoner-R1 against two categories of baseline models:

General Large VLMs. This includes powerful, widely-accessible models such as Gemini 2.0 Flash [21], Llama 4 Maverick [47], Gemini 1.5 Pro [64], and ChatGPT-4o [50]. These are evaluated in zero-shot/few-shot settings as reference of standard VLM capabilities w/o task-specific fine-tuning. Specialized VLMs. This baseline set comprises models specifically developed, adapted, or fine-tuned for spatial understanding tasks, allowing us to assess our contributions relative to other specialized approaches. The models included are: SpatialBot-3B [6], SpaceThinker Qwen2.5VL-3B [4], InternVL2.5-78B [12], Sa2VA (4B, 8B) [90], and SpatialRGPT-8B [13].

SpatialReasoner-R1 VLMs. We also include SpatialReasoner-R1 4B and 8B variants with different training strategies, such as SpatialReasoner-R1 SFT, SpatialReasoner-R1 DPO, trained with standard DPO, and SpatialReasoner-R1 fDPO trained with the proposed fine-grained DPO method.

Table 2: General Vision-Language Understanding Results. Best performance in bold.

Models	MME	POPE	SEED-BENCH	AI2D	SQATEST	MMMUv	MMSTAR	HALLUSIONBENCH
SpatialRGPT-8b [13]	1667 /348	85.50	67.00	67.42	81.81	41.40	43.98	40.80
SpatialReasoner-R1 fDPO 8B	1667/503	89.71	76.21	78.85	93.85	48.11	55.43	51.10

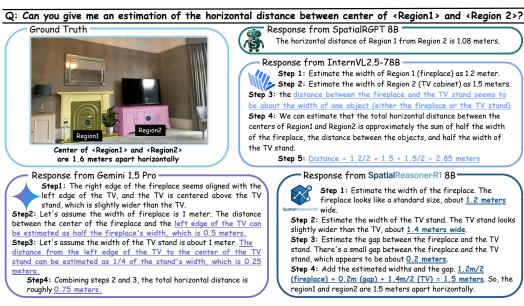


Figure 4: **Qualitative Examples of Spatial Reasoning Across Models.** SpatialReasoner-R1 demonstrates a coherent, multi-step logical chain that closely matches the ground truth, while other models exhibit less precise or less interpretable reasoning paths.

4.2 Experimental Results

Spatial Reasoning. As shown in Table 1, SpatialReasoner-R1 models achieve substantial improvements over both general-purpose and spatial-specialized VLMs across all spatial tasks. Notably, SpatialReasoner-R1 fDPO 8B sets a new benchmark for average accuracy with **2.9**% and **15.8**% gains over SpatialReasoner-R1 fDPO 4B outperforms larger models like InternVL2.5-78B, highlighting the effectiveness of our fine-tuning strategy. Finally, when compared to its predecessor DPO 8B, our optimized variant fDPO 8B boosts average accuracy by **4.1**% across quality tasks and by **9.0**% in quantity tasks.

General Vision-Language Understanding. Beyond achieving state-of-the-art performance in spatial reasoning tasks, our SpatialReasoner-R1 fDPO 8B also demonstrates significant gains in general vision-language benchmarks compared to SpatialRGPT-8B, as presented in Table 2.

4.3 Qualitative Examples

Figure 4 provides qualitative examples that demonstrate SpatialReasoner-R1's advanced capability for coherent, multi-step spatial reasoning. **SpatialReasoner-R1** first estimates the fireplace and TV-stand widths at 1.2m and 1.4m, then computes $\frac{1.2}{2} + 0.2$ (gap) $+ \frac{1.4}{2} = 1.5$ m, a value that nearly matches the reference while transparently tying each term to an observed feature. In contrast, **InternVL2.5-78B** adopts similar width guesses (1.2m, 1.5m) but assumes "the distance between the fireplace and the TV stand seems to be about the width of one object". This assumption is inconsistent with what is shown in the image. **Gemini 1.5 Pro** aligns the fireplace's right edge with the TV's left edge, assigns both objects a 1m width, and combines only half a fireplace width (0.5m) with one quarter of the stand width (0.25m). These two estimates are not accurate and ignore the gap between the two regions. **SpatialRGPT-8B** yields a more accurate estimate than Gemini 1.5 Pro and InternVL2.5-78B. However, since it is not designed as a reasoning model, it cannot generate step-by-step reasoning traces, *i.e.*, does not explicitly reveal the logical chain of spatial deductions or intermediate calculations

Table 3: **Effect of Alpha** (α).

Metric	10%	20%	30%	40%
Direct Dist.	53.38	56.76	60.13	58.11
Horiz. Dist.	52.46	55.74	59.01	56.55
Vert. Dist.	65.09	67.92	71.75	69.81
Width	51.88	57.89	65.41	63.16
Height	56.39	57.14	57.89	58.64

92.23

92.52

94.39

91.58

Table 4: **Effect of Lambda** (λ) at α =30%.

Metric	λ =0.2	λ =0.4	λ=0.6	λ =0.8
Direct Dist.	54.05	57.38	60.13	59.45
Horiz. Dist.	53.27	57.37	59.01	58.19
Vert. Dist.	65.09	68.86	71.75	69.81
Width	53.38	60.15	65.41	54.67
Height	56.39	57.14	57.89	57.89
Direction	91.58	91.58	92.53	93.45

leading to that estimate. These qualitative examples show that SpatialReasoner-R1 has more accurate spatial awareness. Additional examples can be found in Appendix F.

4.4 Ablations

Direction

Table 3 illustrates the impact of varying the Alpha (α) parameter, which modulates the magnitude of segment-specific learning adjustments during fDPO optimization. When α is set too high, the model may overly focus on the reasoning part at the expense of the other, introducing instability and degraded performance, as observed when α reaches 40%. Conversely, if α is too low, both description and reasoning segments are optimized equally. A moderate value of $\alpha=30\%$ allows the model to effectively amplify learning signals for fine-grained spatial distinctions, leading to substantial improvements across all spatial metrics. Furthermore, Table 4 presents the impact of varying the Lambda (λ) parameter, which modulates the sensitivity of segment-specific weights to preference differentials, controlling how responsively the model shifts learning focus based on the observed preference margins. As λ increases, the model becomes more sensitive to segment-specific preference differences, leading to noticeable changes in performance across spatial metrics. We observe that while a moderate value of $\lambda=0.6$ achieves the best overall results, setting λ too high can introduce slight performance degradation in some spatial metrics, likely due to overly aggressive re-weighting.

5 Conclusion

In this work, we introduce **SpatialReasoner-R1**, a novel VLM with state-of-the-art spatial reasoning capabilities, trained with a proposed fine-grained DPO (fDPO) method that decomposes LongCoT paths into description and reasoning components, allowing for targeted preference-based learning and enhanced logical reasoning. fDPO is guided by a set of comprehensive rewards that evaluate reasoning paths across visual consistency, spatial alignment, logical coherence, and depth-based verification. Additionally, we propose a Multi-Model Monte Carlo Tree Search (M3CTS) strategy that leverages multiple LLMs to generate high-quality, diverse LongCoT data. Our comprehensive evaluations demonstrate SpatialReasoner-R1 achieves state-of-the-art performance, outperforming significantly larger models. Moving forward, we plan to evaluate fDPO on additional VLM tasks such as GUI navigation and reasoning segmentation.

Acknowledgments

This research was supported by a gift from Google AI, the Google TPU Research Cloud (TRC) program, and the U.S. Defense Advanced Research Projects Agency (DARPA) under award numbers HR00112390062 and HR001125C0303. We gratefully acknowledge the cloud TPU credits from the Google TPU Research Cloud (TRC) program and the Google Tunix (Tune-in-JAX) team for their feedback. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of Google, DARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

References

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. In *Advances in Neural Information Processing Systems* (NeurIPS), 2022.
- [2] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C Lawrence Zitnick, and Devi Parikh. Vqa: Visual question answering. In *International Conference on Computer Vision (ICCV)*, 2015.
- [3] Mohammad Gheshlaghi Azar, Zhaohan Daniel Guo, Bilal Piot, Remi Munos, Mark Rowland, Michal Valko, and Daniele Calandriello. A general theoretical paradigm to understand learning from human preferences. In *Annual Conference on Artificial Intelligence and Statistics* (AISTATS), 2024.
- [4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibo Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2.5-vl technical report. *arXiv:2502.13923*, 2025.
- [5] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In Advances in Neural Information Processing Systems (NeurIPS), 2020.
- [6] Wenxiao Cai, Iaroslav Ponomarenko, Jianhao Yuan, Xiaoqi Li, Wankou Yang, Hao Dong, and Bo Zhao. Spatialbot: Precise spatial understanding with vision language models. In *International Conference on Robotics and Automation (ICRA)*, 2025.
- [7] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brian Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [8] Keqin Chen, Zhao Zhang, Weili Zeng, Richong Zhang, Feng Zhu, and Rui Zhao. Shikra: Unleashing multimodal llm's referential dialogue magic. *arXiv*:2306.15195, 2023.
- [9] Lin Chen, Jinsong Li, Xiaoyi Dong, Pan Zhang, Yuhang Zang, Zehui Chen, Haodong Duan, Jiaqi Wang, Yu Qiao, Dahua Lin, et al. Are we on the right way for evaluating large vision-language models? In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [10] Qiguang Chen, Libo Qin, Jinhao Liu, Dengyun Peng, Jiannan Guan, Peng Wang, Mengkang Hu, Yuhang Zhou, Te Gao, and Wanxiang Che. Towards reasoning era: A survey of long chain-of-thought for reasoning large language models. *arXiv:2503.09567*, 2025.
- [11] Zhangquan Chen, Xufang Luo, and Dongsheng Li. Visrl: Intention-driven visual perception via reinforced reasoning. In *International Conference on Computer Vision (ICCV)*, 2025.
- [12] Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [13] 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. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [14] Chenhang Cui, An Zhang, Yiyang Zhou, Zhaorun Chen, Gelei Deng, Huaxiu Yao, and Tat-Seng Chua. Fine-grained verifiers: Preference modeling as next-token prediction in vision-language alignment. In *International Conference on Learning Representations (ICLR)*, 2024.
- [15] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale N Fung, and Steven Hoi. Instructblip: towards general-purpose vision-language models with instruction tuning. In *Advances in Neural Information Processing Systems* (NeurIPS), 2023.

- [16] DeepSeek-AI, Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, Xiaokang Zhang, Xingkai Yu, Yu Wu, Z. F. Wu, Zhibin Gou, Zhihong Shao, Zhuoshu Li, Ziyi Gao, Aixin Liu, Bing Xue, Bingxuan Wang, Bochao Wu, Bei Feng, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, Damai Dai, Deli Chen, Dongjie 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, J. L. 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, Minghui 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, R. L. Jin, Ruyi Chen, Shanghao Lu, Shangyan Zhou, Shanhuang Chen, Shengfeng Ye, Shiyu Wang, Shuiping Yu, Shunfeng Zhou, Shuting Pan, S. S. Li, Shuang Zhou, Shaoqing Wu, Shengfeng Ye, Tao Yun, Tian Pei, Tianyu Sun, T. Wang, Wangding Zeng, Wanjia Zhao, Wen Liu, Wenfeng Liang, Wenjun Gao, Wenqin Yu, Wentao Zhang, W. L. Xiao, Wei An, Xiaodong Liu, Xiaohan Wang, Xiaokang Chen, Xiaotao Nie, Xin Cheng, Xin Liu, Xin Xie, Xingchao Liu, Xinyu Yang, Xinyuan Li, Xuecheng Su, Xuheng Lin, X. Q. Li, Xiangyue Jin, Xiaojin 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, Yiliang Xiong, Ying He, Yishi Piao, Yisong Wang, Yixuan Tan, Yiyang Ma, Yiyuan Liu, Yongqiang Guo, Yuan Ou, Yuduan Wang, Yue Gong, Yuheng Zou, Yujia He, Yunfan Xiong, Yuxiang Luo, Yuxiang You, Yuxuan Liu, Yuyang Zhou, Y. X. Zhu, Yanhong Xu, Yanping Huang, Yaohui Li, Yi Zheng, Yuchen Zhu, Yunxian Ma, Ying Tang, Yukun Zha, Yuting Yan, Z. Z. Ren, Zehui Ren, Zhangli Sha, Zhe Fu, Zhean Xu, Zhenda Xie, Zhengyan Zhang, Zhewen Hao, Zhicheng Ma, Zhigang Yan, Zhiyu Wu, Zihui Gu, Zijia Zhu, Zijun Liu, Zilin 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 2501.12948, 2025.
- [17] 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 multimodal models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [18] Yuhao Dong, Zuyan Liu, Hai-Long Sun, Jingkang Yang, Winston Hu, Yongming Rao, and Ziwei Liu. Insight-v: Exploring long-chain visual reasoning with multimodal large language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [19] Kawin Ethayarajh, Winnie Xu, Niklas Muennighoff, Dan Jurafsky, and Douwe Kiela. Model alignment as prospect theoretic optimization. In *International Conference on Machine Learning (ICML)*, 2024.
- [20] Yuhan Fu, Ruobing Xie, Xingwu Sun, Zhanhui Kang, and Xirong Li. Mitigating hallucination in multimodal large language model via hallucination-targeted direct preference optimization. In *Findings of the Association for Computational Linguistics (ACL)*, 2025.
- [21] Google Deepmind. Gemini 2.0 is now available to everyone. https://blog.google/technology/google-deepmind/gemini-model-updates-february-2025/, 2025.
- [22] Yi Gu, Zhendong Wang, Yueqin Yin, Yujia Xie, and Mingyuan Zhou. Diffusion-rpo: Aligning diffusion models through relative preference optimization. *arXiv:2406.06382*, 2024.
- [23] Tianrui Guan, Fuxiao Liu, Xiyang Wu, Ruiqi Xian, Zongxia Li, Xiaoyu Liu, Xijun Wang, Lichang Chen, Furong Huang, Yaser Yacoob, et al. Hallusionbench: An advanced diagnostic suite for entangled language hallucination and visual illusion in large vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [24] Qiushan Guo, Shalini De Mello, Hongxu Yin, Wonmin Byeon, Ka Chun Cheung, Yizhou Yu, Ping Luo, and Sifei Liu. Regiongpt: Towards region understanding vision language model. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.

- [25] 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:2410.21276, 2024.
- [26] Amazon Artificial General Intelligence. The amazon nova family of models: Technical report and model card. 2024.
- [27] Aniruddha Kembhavi, Mike Salvato, Eric Kolve, Minjoon Seo, Hannaneh Hajishirzi, and Ali Farhadi. A diagram is worth a dozen images. In *European Conference on Computer Vision (ECCV)*, 2016.
- [28] Yangzhe Kong, Daeun Song, Jing Liang, Dinesh Manocha, Ziyu Yao, and Xuesu Xiao. Autospatial: Visual-language reasoning for social robot navigation through efficient spatial reasoning learning. *arXiv:2503.07557*, 2025.
- [29] Xin Lai, Zhuotao Tian, Yukang Chen, Senqiao Yang, Xiangru Peng, and Jiaya Jia. Step-dpo: Step-wise preference optimization for long-chain reasoning of llms. *arXiv*:2406.18629, 2024.
- [30] Bohao Li, Rui Wang, Guangzhi Wang, Yuying Ge, Yixiao Ge, and Ying Shan. Seed-bench: Benchmarking multimodal llms with generative comprehension. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [31] 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. In *International Conference on Learning Representations (ICLR)*, 2025.
- [32] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International Conference on Machine Learning (ICML)*, 2023.
- [33] Shufan Li, Konstantinos Kallidromitis, Akash Gokul, Yusuke Kato, and Kazuki Kozuka. Aligning diffusion models by optimizing human utility. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [34] Yifan Li, Yifan Du, Kun Zhou, Jinpeng Wang, Wayne Xin Zhao, and Ji-Rong Wen. Evaluating object hallucination in large vision-language models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023.
- [35] Zijing Liang, Yanjie Xu, Yifan Hong, Penghui Shang, Qi Wang, Qiang Fu, and Ke Liu. A survey of multimodel large language models. In *International Conference on Computer, Artificial Intelligence and Control Engineering*, 2024.
- [36] Yuan-Hong Liao, Rafid Mahmood, Sanja Fidler, and David Acuna. Reasoning paths with reference objects elicit quantitative spatial reasoning in large vision-language models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [37] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 2014.
- [38] Aiwei Liu, Haoping Bai, Zhiyun Lu, Yanchao Sun, Xiang Kong, Simon Wang, Jiulong Shan, Albin Madappally Jose, Xiaojiang Liu, Lijie Wen, et al. Tis-dpo: Token-level importance sampling for direct preference optimization with estimated weights. In *International Conference on Learning Representations (ICLR)*, 2024.
- [39] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [40] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (CVPR), 2024.

- [41] Yuecheng Liu, Dafeng Chi, Shiguang Wu, Zhanguang Zhang, Yaochen Hu, Lingfeng Zhang, Yingxue Zhang, Shuang Wu, Tongtong Cao, Guowei Huang, et al. Spatialcot: Advancing spatial reasoning through coordinate alignment and chain-of-thought for embodied task planning. *arXiv*:2501.10074, 2025.
- [42] Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kai-Wei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. Learn to explain: Multimodal reasoning via thought chains for science question answering. In *Advances in Neural Information Processing Systems* (NeurIPS), 2022.
- [43] Haotian Luo, Haiying He, Yibo Wang, Jinluan Yang, Rui Liu, Naiqiang Tan, Xiaochun Cao, Dacheng Tao, and Li Shen. Ada-r1: From long-cot to hybrid-cot via bi-level adaptive reasoning optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [44] 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. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [45] Wufei Ma, Haoyu Chen, Guofeng Zhang, Celso M de Melo, Jieneng Chen, and Alan Yuille. 3dsrbench: A comprehensive 3d spatial reasoning benchmark. In *International Conference on Computer Vision (ICCV)*, 2025.
- [46] Wufei Ma, Yu-Cheng Chou, Qihao Liu, Xingrui Wang, Celso de Melo, Jieneng Chen, Jianwen Xie, and Alan Yuille. Spatialreasoner: Towards explicit and generalizable 3d spatial reasoning. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
- [47] Meta. The llama 4 herd: The beginning of a new era of natively multimodal ai innovation. https://ai.meta.com/blog/llama-4-multimodal-intelligence/, April 2025.
- [48] Remi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Côme Fiegel, et al. Nash learning from human feedback. In *International Conference on Machine Learning (ICML)*, 2024.
- [49] Kiet A Nguyen, Adheesh Juvekar, Tianjiao Yu, Muntasir Wahed, and Ismini Lourentzou. Calico: Part-focused semantic co-segmentation with large vision-language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [50] OpenAI. Gpt-4o system card. https://openai.com/index/gpt-4o-system-card/, 2024.
- [51] Zhenyu Pan and Han Liu. Metaspatial: Reinforcing 3d spatial reasoning in vlms for the metaverse. *arXiv:2503.18470*, 2025.
- [52] Richard Yuanzhe Pang, Weizhe Yuan, He He, Kyunghyun Cho, Sainbayar Sukhbaatar, and Jason Weston. Iterative reasoning preference optimization. In Advances in Neural Information Processing Systems (NeurIPS), 2024.
- [53] Sungjin Park, Xiao Liu, Yeyun Gong, and Edward Choi. Ensembling large language models with process reward-guided tree search for better complex reasoning. In *Nations of the Americas Chapter of the Association for Computational Linguistics (NAACL)*, 2025.
- [54] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, 2021.
- [55] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023.
- [56] Kanchana Ranasinghe, Satya Narayan Shukla, Omid Poursaeed, Michael S Ryoo, and Tsung-Yu Lin. Learning to localize objects improves spatial reasoning in visual-llms. In 2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 12977–12987. IEEE Computer Society, 2024.

- [57] 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [58] Ruichen Shao, Bei Li, Gangao Liu, Yang Chen, Xiang Zhou, Jingang Wang, Xunliang Cai, and Peng Li. Earlier tokens contribute more: Learning direct preference optimization from temporal decay perspective. In *International Conference on Learning Representations (ICLR)*, 2025.
- [59] Shuaijie She, Wei Zou, Shujian Huang, Wenhao Zhu, Xiang Liu, Xiang Geng, and Jiajun Chen. Mapo: Advancing multilingual reasoning through multilingual-alignment-as-preference optimization. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [60] Wentao Shi, Mengqi Yuan, Junkang Wu, Qifan Wang, and Fuli Feng. Direct multi-turn preference optimization for language agents. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [61] Yifan Song, Da Yin, Xiang Yue, Jie Huang, Sujian Li, and Bill Yuchen Lin. Trial and error: Exploration-based trajectory optimization of llm agents. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [62] Zhiqing Sun, Sheng Shen, Shengcao Cao, Haotian Liu, Chunyuan Li, Yikang Shen, Chuang Gan, Liang-Yan Gui, Yu-Xiong Wang, Yiming Yang, et al. Aligning large multimodal models with factually augmented rlhf. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2023.
- [63] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv*:2312.11805, 2023.
- [64] Gemini Team, Petko Georgiev, Ving Ian Lei, Ryan Burnell, Libin Bai, Anmol Gulati, Garrett Tanzer, Damien Vincent, Zhufeng Pan, Shibo Wang, et al. Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context. *arXiv:2403.05530*, 2024.
- [65] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv*:2302.13971, 2023.
- [66] Shangqing Tu, Yucheng Wang, Daniel Zhang-Li, Yushi Bai, Jifan Yu, Yuhao Wu, Lei Hou, Huiqin Liu, Zhiyuan Liu, Bin Xu, et al. Longwriter-v: Enabling ultra-long and high-fidelity generation in vision-language models. In ACM International Conference on Multimedia (ACM MM), 2025.
- [67] Muntasir Wahed, Kiet A Nguyen, Adheesh Sunil Juvekar, Xinzhuo Li, Xiaona Zhou, Vedant Shah, Tianjiao Yu, Pinar Yanardag, and Ismini Lourentzou. Prima: Multi-image vision-language models for reasoning segmentation. *arXiv:2412.15209*, 2024.
- [68] Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. Diffusion model alignment using direct preference optimization. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [69] Fei Wang, Wenxuan Zhou, James Y Huang, Nan Xu, Sheng Zhang, Hoifung Poon, and Muhao Chen. mdpo: Conditional preference optimization for multimodal large language models. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [70] Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. Mixture-of-agents enhances large language model capabilities. In *International Conference on Learning Representations (ICLR)*, 2025.

- [71] Weiyun Wang, Zhe Chen, Wenhai Wang, Yue Cao, Yangzhou Liu, Zhangwei Gao, Jinguo Zhu, Xizhou Zhu, Lewei Lu, Yu Qiao, et al. Enhancing the reasoning ability of multimodal large language models via mixed preference optimization. *arXiv*:2411.10442, 2024.
- [72] Xiyao Wang, Linfeng Song, Ye Tian, Dian Yu, Baolin Peng, Haitao Mi, Furong Huang, and Dong Yu. Towards self-improvement of llms via mcts: Leveraging stepwise knowledge with curriculum preference learning. *arXiv:2410.06508*, 2024.
- [73] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [74] Liang Wen, Yunke Cai, Fenrui Xiao, Xin He, Qi An, Zhenyu Duan, Yimin Du, Junchen Liu, Lifu Tang, Xiaowei Lv, et al. Light-r1: Curriculum sft, dpo and rl for long cot from scratch and beyond. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2025.
- [75] Junkang Wu, Yuexiang Xie, Zhengyi Yang, Jiancan Wu, Jinyang Gao, Bolin Ding, Xiang Wang, and Xiangnan He. beta-dpo: Direct preference optimization with dynamic beta. In Advances in Neural Information Processing Systems (NeurIPS), 2024.
- [76] Wenshan Wu, Shaoguang Mao, Yadong Zhang, Yan Xia, Li Dong, Lei Cui, and Furu Wei. Mind's eye of llms: Visualization-of-thought elicits spatial reasoning in large language models. In Advances in Neural Information Processing Systems (NeurIPS), 2024.
- [77] Xun Wu, Shaohan Huang, Guolong Wang, Jing Xiong, and Furu Wei. Multimodal large language models make text-to-image generative models align better. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [78] Yuxi Xie, Guanzhen Li, Xiao Xu, and Min-Yen Kan. V-dpo: Mitigating hallucination in large vision language models via vision-guided direct preference optimization. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [79] Shuo Xing, Yuping Wang, Peiran Li, Ruizheng Bai, Yueqi Wang, Chengxuan Qian, Huaxiu Yao, and Zhengzhong Tu. Re-align: Aligning vision language models via retrieval-augmented direct preference optimization. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2025.
- [80] Tianyi Xiong, Xiyao Wang, Dong Guo, Qinghao Ye, Haoqi Fan, Quanquan Gu, Heng Huang, and Chunyuan Li. Llava-critic: Learning to evaluate multimodal models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [81] Wei Xiong, Chengshuai Shi, Jiaming Shen, Aviv Rosenberg, Zhen Qin, Daniele Calandriello, Misha Khalman, Rishabh Joshi, Bilal Piot, Mohammad Saleh, et al. Building math agents with multi-turn iterative preference learning. In *International Conference on Learning Representations (ICLR)*, 2025.
- [82] 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. *arXiv:2412.14171*, 2024.
- [83] Kai Yang, Jian Tao, Jiafei Lyu, Chunjiang Ge, Jiaxin Chen, Weihan Shen, Xiaolong Zhu, and Xiu Li. Using human feedback to fine-tune diffusion models without any reward model. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8941–8951, 2024.
- [84] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. In *Advances in Neural Information Processing Systems* (NeurIPS), 2024.
- [85] Sen Yang, Yafu Li, Wai Lam, and Yu Cheng. Multi-llm collaborative search for complex problem solving. *arXiv*:2502.18873, 2025.

- [86] Zhihe Yang, Xufang Luo, Dongqi Han, Yunjian Xu, and Dongsheng Li. Mitigating hallucinations in large vision-language models via dpo: On-policy data hold the key. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [87] Huanjin Yao, Jiaxing Huang, Wenhao Wu, Jingyi Zhang, Yibo Wang, Shunyu Liu, Yingjie Wang, Yuxin Song, Haocheng Feng, Li Shen, et al. Mulberry: Empowering mllm with o1-like reasoning and reflection via collective monte carlo tree search. *arXiv:2412.18319*, 2024.
- [88] 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. In *International Conference on Learning Representations (ICLR)*, 2023.
- [89] Tianyu Yu, Yuan Yao, Haoye Zhang, Taiwen He, Yifeng Han, Ganqu Cui, Jinyi Hu, Zhiyuan Liu, Hai-Tao Zheng, Maosong Sun, et al. Rlhf-v: Towards trustworthy mllms via behavior alignment from fine-grained correctional human feedback. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [90] Haobo Yuan, Xiangtai Li, Tao Zhang, Zilong Huang, Shilin Xu, Shunping Ji, Yunhai Tong, Lu Qi, Jiashi Feng, and Ming-Hsuan Yang. Sa2va: Marrying sam2 with llava for dense grounded understanding of images and videos. arXiv:2501.04001, 2025.
- [91] Huizhuo Yuan, Zixiang Chen, Kaixuan Ji, and Quanquan Gu. Self-play fine-tuning of diffusion models for text-to-image generation. In *Advances in Neural Information Processing Systems* (*NeurIPS*), 2024.
- [92] Yuqian Yuan, Wentong Li, Jian Liu, Dongqi Tang, Xinjie Luo, Chi Qin, Lei Zhang, and Jianke Zhu. Osprey: Pixel understanding with visual instruction tuning. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [93] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multi-modal understanding and reasoning benchmark for expert agi. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [94] Xiang Yue, Tianyu Zheng, Ge Zhang, and Wenhu Chen. Mammoth2: Scaling instructions from the web. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [95] Yongcheng Zeng, Guoqing Liu, Weiyu Ma, Ning Yang, Haifeng Zhang, and Jun Wang. Token-level direct preference optimization. In *International Conference on Machine Learning (ICML)*, 2024.
- [96] Duzhen Zhang, Yahan Yu, Jiahua Dong, Chenxing Li, Dan Su, Chenhui Chu, and Dong Yu. Mm-llms: Recent advances in multimodal large language models. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2024.
- [97] Xuan Zhang, Chao Du, Tianyu Pang, Qian Liu, Wei Gao, and Min Lin. Chain of preference optimization: Improving chain-of-thought reasoning in llms. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
- [98] Qi Zhao, Haotian Fu, Chen Sun, and George Konidaris. Epo: Hierarchical Ilm agents with environment preference optimization. In *Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2024.
- [99] Yao Zhao, Rishabh Joshi, Tianqi Liu, Misha Khalman, Mohammad Saleh, and Peter J Liu. Slic-hf: Sequence likelihood calibration with human feedback. *arXiv:2305.10425*, 2023.
- [100] Han Zhong, Zikang Shan, Guhao Feng, Wei Xiong, Xinle Cheng, Li Zhao, Di He, Jiang Bian, and Liwei Wang. Dpo meets ppo: Reinforced token optimization for rlhf. In ICML 2024 Workshop on Models of Human Feedback for AI Alignment, 2024.
- [101] Yiyang Zhou, Zhiyuan Fan, Dongjie Cheng, Sihan Yang, Zhaorun Chen, Chenhang Cui, Xiyao Wang, Yun Li, Linjun Zhang, and Huaxiu Yao. Calibrated self-rewarding vision language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.

A Fine-Grained Spatial Reward Details

This appendix details design and hyperparameter choices for the fine-grained spatial rewards introduced in Section 3.3. The prompt template for estimating rewards is shown in Figure 10.

A.1 Visual Consistency Reward (\mathcal{R}_{vc})

The Visual Consistency Reward quantifies alignment between the generated description and the visual scene across four continuous criteria: Existence, Attribute accuracy, Completeness, and Appropriateness. Each component yields a score in the range [0.0, 1.0], where its continuous range enables fine-grained assessment, permitting fractional scores when descriptions partially satisfy evaluation criteria. Scores near 0.0 indicate misalignment, while scores near 1.0 denote perfect alignment. Intermediate values reflect varying degrees of partial correctness or uncertainty. The total reward, $\mathcal{R}_{vc} \in [0, 4.0]$, distinguishes varying degrees of alignment across responses.

A.2 Depth-Guided Spatial Reward (\mathcal{R}_{sp})

We introduce a depth-guided reward to evaluate the spatial accuracy of model outputs using ground-truth depth maps. The reward is computed independently for the description R_{desc} and reasoning R_{reason} components, yielding two sub-scores: $\mathcal{R}_{\text{sp,desc}}$ and $\mathcal{R}_{\text{sp,reason}}$, each ranging from 0 to 4. These scores capture the alignment of spatial expressions with geometric cues in the image. The final spatial reward is given by $\mathcal{R}_{\text{sp}} = \mathcal{R}_{\text{sp,desc}} + \mathcal{R}_{\text{sp,reason}}$.

Uncertainty Weight (W_u) . Spatial expressions in model outputs often include uncertain qualifiers. W_u ranges from 0.8 to 1.0, with 1.0 indicating complete certainty in spatial assertions, and the lower bound of 0.8 representing cautious but plausible uncertainty. Setting the lower bound at 0.8 balances cautious language (e.g., "approximately," "possibly") without overly penalizing reasonable uncertainty. Lower values (below 0.8) would overly penalize reasonable, conservative predictions and discourage the model from producing cautious but informative reasoning.

Context-aware Weight (W_c) . The context-aware weight $W_c \in [0.8, 1.0]$ reflects the relevance of spatial statements to the question. Explicitly asked spatial relationships are assigned $W_c = 1.0$, while auxiliary or indirect spatial references are assigned $W_c = 0.8$. This distinction prioritizes primary spatial relations explicitly required by the query, ensuring the model emphasizes essential spatial assertions more significantly. Scores below 0.8 would disproportionately underemphasize auxiliary information, degrading the model's ability to handle broader contextually relevant details.

Given a response, we extract all spatial relationship statements from the description and reasoning response sections. Each statement is then evaluated using GPT-40 by comparing the original image and its corresponding depth image, which is generated using Depth Anything [84], to obtain a correctness score $r_i \in [0,1]$. Every statement is also assigned an associated $W_u^{(i)}$ and $W_c^{(i)}$. The spatial reward scores are computed as

$$\mathcal{R}_{\text{sp,desc}} = \frac{1}{n} \sum_{i=1}^{n} W_u^{(i)} \cdot W_c^{(i)} \cdot r_i, \qquad \mathcal{R}_{\text{sp,reason}} = \frac{1}{m} \sum_{i=1}^{m} W_u^{(i)} \cdot W_c^{(i)} \cdot r_i, \tag{11}$$

where n and m denote the number of spatial statements in the description and reasoning components, respectively, and r_i represents the correctness of each spatial relationship, validated against both the RGB image and its corresponding depth map.

A.3 Logical Coherence Reward (\mathcal{R}_{lc})

This reward quantifies the logical robustness of a response by aggregating four components: Factual Consistency, Logical Coherence, Correct Rule Application, and Conclusion Validity. Each component is scored in the range [0.0, 1.0], with fractional values capturing partial correctness, *e.g.*, from minor gaps in logical sequences to partial inaccuracies in applying physical, spatial, or logical rules. Scores of 0.0 and 1.0 indicate complete logical coherence failure or perfect logical chains, respectively. The final reward, $\mathcal{R}lc \in [0, 4.0]$, reflects the overall logical quality of the reasoning.

B Structured Output Format Specification for M3CTS

To enable reliable parsing and downstream analysis, M3CTS requires reasoning paths from LLMs to follow a standardized structured format. This format uses Markdown-style headings to clearly segment key components of the reasoning trace. Each section begins with a line prefixed by ###, followed by a descriptive heading. The defined sections are:

Description: Details the input context, such as the visual scene or scenario description.

Rationale: Summarizes the overall reasoning strategy or justification.

Let's think step by step: An optional phrase before detailed step-by-step breakdown.

Step \mathbb{N} (e.g., ### Step 1, ### Step 2, ...): Enumerates the sequential steps involved in the reasoning procedure. Multiple steps are typically present.

In Conclusion: States the final derived conclusion of the reasoning process.

Figure 5 shows the reasoning tree example produced by M3CTS.

C Node Evaluation Protocol for M3CTS

To ensure the semantic and visual quality of each candidate reasoning step $s_{k,t}$ within the M3CTS framework, we employ a structured multi-criteria evaluation. Individual steps that form a reasoning path are segmented by ###. Each candidate $s_{k,t} \in \mathcal{S}_c$ is independently evaluated by two multimodal models, Gemini 1.5 Pro and Qwen2.5VL-72B, along key distinct dimensions:

- \diamond Visual Description Accuracy: Assesses whether the entities, attributes, and contextual cues described in $s_{k,t}$ correctly reflect the visual content of the input image. This includes references to objects, colors, spatial layouts, and contextual cues.
- \diamond Spatial Consistency: Evaluates whether the spatial relations expressed in $s_{k,t}$ (e.g., "above," "to the left of," "behind") are consistent with both the RGB image and depth map generated via the Depth Anything model [84]. Errors such as inversion of relations (e.g., stating "behind" instead of "in front") are penalized.
- ♦ Logical Reasoning Coherence: For steps within the "think step-by-step" chain-of-thought reasoning block, this component checks whether the logical flow of inferences is coherent and justified. This includes identifying unsupported jumps in logic or contradictions.

Each criterion $\mathbb{I}_{\text{eval}}^{(m)}(s_{k,t})$ is rated as follows:

$$\mathbb{I}^{(m)}_{\mathrm{eval}}(s_{k,t}) = \begin{cases} +1, & \text{if the content is entirely accurate according to model } m; \\ \\ 0, & \text{if the content is ambiguous or partially accurate;} \\ \\ -1, & \text{if there is any clear inaccuracy.} \end{cases}$$

We preserve high-quality paths by pruning the candidate set. Specifically, we retain any node $s_{k,t}$ whose aggregated score across all evaluators and criteria is non-negative, i.e., $\mathcal{S}_{c}^{*} = \{s_{k,t} \mid R(s_{k,t}) \geq 0\}$. This threshold is chosen empirically to balance filtering out incorrect steps while maintaining adequate reasoning diversity.

D Training Details

D.1 Implementation Details

SpatialReasoner-R1 is built upon the Sa2VA architecture [90], which is based on InternVL2.5. We train the 8B-parameter model in two stages on two NVIDIA H100 GPUs, each stage taking approximately 2.5 days. For supervised fine-tuning, we employ AdamW optimizer with a learning rate of 4×10^{-5} , weight decay of 0.05, and a 5% linear warm-up schedule, using a batch size of 2 per device with gradient accumulation over 4 steps. For Direct Preference Optimization, we similarly use AdamW with learning rate of 1×10^{-7} , weight decay of 0.05, and a 5% warm-up, training with a batch size of 1 per device.

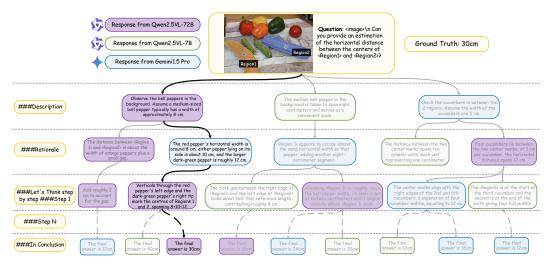


Figure 5: Example Reasoning Tree from the M3CTS Data Generation Pipeline. Diverse candidate reasoning paths are sampled from multiple models. Each path follows a structured LongCoT format with markdown-style section headers that decompose the answer into interpretable reasoning stages.

D.2 Training Data

For SFT, we convert samples from the OPEN SPATIAL dataset [13] to reasoning chains using the M3CTS pipeline. While the original OPEN SPATIAL dataset provides single-sentence answers, we transform 400K samples, grounded in distinct images, into structured LongCoT reasoning chains, where examples are used to teach the model to generate high-quality, step-by-step spatial reasoning responses. For Direct Preference Optimization (DPO) training, the goal is to train the model to distinguish high-quality spatial reasoning from suboptimal or subtly flawed alternatives. To this end, we utilize our OPEN SPATIAL REASONING dataset, described below, that consists of spatial reasoning preference pairs.

An additional set of 100K challenging negative pairs is meticulously crafted by perturbing only the conclusion keywords of high-quality positive samples. Each original response represents a coherent and accurate reasoning path with a factually correct outcome. To create the corresponding negative sample, we retain the exact description and reasoning segments and alter only the final conclusion value. This yields tightly controlled preference pairs that isolate correctness at the conclusion level. For example, a positive sample may assert "The distance between region1 and region2 is 11 meters.", while its negative perturbed counterpart is "The distance between region1 and region2 is 10 meters."

Our method adopts a data-centric strategy that emphasizes high-quality supervision and reasoning diversity. Instead of collecting large volumes of weakly aligned or noisy data, we curate training examples using the M3CTS sampling strategy guided by structured reward evaluations. By applying reward-based filtering, we reduce noise and enforce a consistent output structure. In parallel, using multiple LLMs during generation introduces variation in reasoning styles, improving coverage of diverse spatial patterns and edge cases. The effectiveness of this approach is evident in the substantial performance gains of DPO-trained models over their simpler SFT counterparts (Table 1) and the reasoning improvements and diversity depicted in Figure 5.

D.3 Open Spatial Reasoning Dataset

We curate the OPEN SPATIAL REASONING dataset, a collection of 400K Vision Question Answering (VQA) preference pairs (y_p, y_l) , to support training of preference-based spatial reasoning models. This dataset is derived from the OPEN SPATIAL dataset [13], which provides image-based spatial questions paired with ground-truth answers and offers 10 question variations per image-grounding scenario. To construct each preference pair, we randomly sample a question instance from the source dataset, and generate a diverse pool of eight candidate answers using four distinct sources: our M3CTS pipeline, Gemini 1.5 Pro, GPT-40, and our SpatialReasoner-R1 Supervised Fine-Tuned

Chosen Response: The image shows a street intersection with reegion2× marked on a vehicle to the left and region1× marked on a person to the right. region2× is located on the left side of the image, while region1× is on the right side of the image, while region1× is on the right side of region2× on the left side of the street. Step 2: Locate region1× on the right side of the street, and further down the road. Step 3: Compare this direction to a clock face, aligning it with the 4 o'clock direction. The find answer is region [2] will find region [1] around the 4 o'clock direction.



Rejected Response: The image shows a street scene with buildings on both sides. There are two labeled regions: region2-, which is a car parked on the left side of the street, and region12, which is a person walking on the sidewalk on the right side of the street. Let's think step by step. Step 1. Imagine standing at the location of region2-, which is the car on the left side of the street. Step 2. Look around for region2-, which is the person walking on the sidewalk. Step 3. The person is at the right of car. The final answer is region [2] will find region [1] around the 3 o clock direction.

Between Region [1] and Region [2], which one has less height?

Chosen Response: The image shows a collection of bottles on a table. Region[1] is a green bottle on the right, and Region[2] is a blue bottle on the left. The bottles' heights can be compared by their relative positions and surrounding context. Let's think step by step. Step 1: The middle bottles have a similar height to Region[2]. Step 2: Region[1] is taller than the middle bottles. The final answer is Region[2].



Rejected Response: The image shows a collection of bottles on a table. Region[1] is a green bottle on the right, and Region[2] is a blue bottle on the left. Region[1] appears shorter than Region[2]. Let's think step by step. Step 1: Compare the heights visually. Region[1] is noticeably shorter than the surrounding bottles. The final answer is Region[1].

What is the distance between <region1> and <region2>?

Chosen Response: The image shows two taxis in front of Trump Tower. Region[1] is a yellow taxi on the right, and Region[2] is a blue taxi on the left. Typical road width is about 3 meters, with each taxi approximately 1.5 meters wide. Let's think step by step. Step 1: The taxis are aligned with their fronts and backs on the same line. Step 2: With each taxi being 1.5 meters wide and the road 3 meters wide, the space between them is about 1.5 meters. The final answer is: 1.5 meters.



Rejected Response: The image shows two taxis in front of Trump Tower. Region[1] is a yellow taxi on the right, and Region[2] is a blue taxi on the left. The taxis are parked closely side by side. Let's think step by step. Step 1: Region[1] and Region[2] are adjacent with minimal visible space. Step 2: Given typical urban parking, assume a standard spacing. The final answer is: Approximately 0.5 meters.

Is <region1> positioned at right of <region2>?

Chosen Response: The image shows a side mirror of a vehicle reflecting a road scene. Region[1] is the highlighted road, and region[2] is the blue car. Let's think step by step. Step 1: Identify the positions of Region[1] and Region[2] in the mirror. Step2: Region[1] is on the left side. Since the image is shown in mirror, so Region[1] is to the right of Region[2]. The final answer is Yes region1> positioned at right region2>.



Rejected Response: The image shows a vehicle reflecting a road scene. Region[1] is the highlighted road, and Region[2] is the blue car. Let's think step by step. Step 1: Identify the positions of Region[1] and Region[2]. Region[1] is on the left side. Step2: Region[2] is on the right side. The final answer is No (region12) positioned at left (region2).

Figure 6: **Example DPO Pairs of our OPEN SPATIAL REASONING Dataset,** constructed from M3CTS-generated reasoning trajectories. Each pair consists of a preferred and a rejected response to the same spatial question. The examples highlight differences in descriptive accuracy, spatial alignment, and reasoning coherence, which guide preference optimization during training.

(SFT) model, with each method contributing two response variants. All eight candidate responses are independently evaluated by our fine-grained spatial reward mechanism (Appendix A). The highest-scoring response is selected as the preferred answer (y_p) , while the response with the lowest score is designated as the less-preferred (y_l) , ensuring that each preference pair is anchored in meaningful fine-grained spatial reasoning quality. Figure 6 shows dataset examples.

E Evaluation Details

We evaluate on SPATIALRGPT-BENCH [13], a benchmark specifically designed to assess the 3D spatial reasoning abilities of VLMs, featuring 657 qualitative and 749 quantitative VQA pairs, covering 88 object classes across diverse environments. We employ the same GPT-4 evaluation proposed in SpatialRGPT-Bench [13] for evaluating the free-form responses generated by the models. For *qualitative* questions, GPT-40 assesses the semantic alignment between the model's response and the ground-truth answer, assigning a binary score (1 for correct, 0 for incorrect). For *quantitative* questions (*e.g.*, distance, size), GPT-40 first extracts numerical values from both the prediction and the ground truth, standardizing them to a common unit (meters). We then compute accuracy (*e.g.*, success rate defined as predictions within $\pm 25\%$ of the ground truth).

We also evaluate on several general vision-language benchmarks to provide a comprehensive assessment of SpatialReasoner-R1's capabilities. Specifically, we use MME [34] to assess multimodal models on perception and cognition tasks across a wide range of domains. POPE [34] is employed

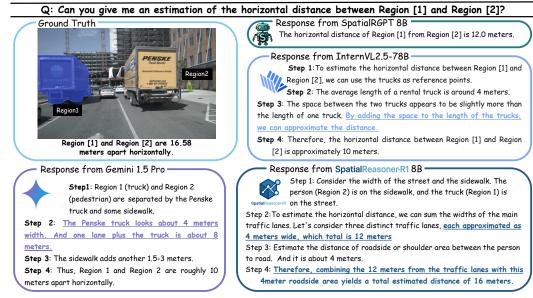


Figure 7: Qualitative Examples of Spatial Reasoning Across Models. SpatialReasoner-R1 demonstrates coherent, step-by-step spatial reasoning that closely aligns with ground truth estimates. In contrast, baseline models produce less precise or partially incorrect reasoning steps, often neglecting key visual cues or misestimating spatial references.

to evaluate object hallucination in testing the ability of VLMs to ground responses to visual content, while SEED-BENCH [30] offers a multi-dimensional evaluation, covering aspects from image understanding to complex reasoning across various modalities and tasks. We further utilize AI2D [27], a benchmark focusing on diagram understanding and reasoning, which requires parsing visual elements and their relationships within schematic representations. SQA [42] is used to measure the model's ability to answer science-related questions based on visual context, often requiring domain-specific knowledge and reasoning. MMMU [93] evaluates massive multi-disciplinary multimodal understanding and reasoning across diverse college-level subjects. Moreover, MMSTAR [9] provides a challenging benchmark with meticulously curated, multimodal instances that require advanced reasoning, low hallucination, and resistance to leading questions. Finally, HALLUSIONBENCH [23] is specifically designed to quantitatively measure and analyze the hallucination phenomena in VLMs, probing for both object-level and attribute-level inconsistencies.

F Qualitative Experiment Examples

In this section, we provide additional qualitative experiment examples. Figure 7 shows a question that requires estimation of the horizontal distance between a truck and a pedestrian. **SpatialReasoner-R1** demonstrates a clear advantage by decomposing the scene into semantically meaningful components, explicitly reasoning over the widths of multiple traffic lanes, the roadside, and the sidewalk. This results in an estimated distance that closely matches the ground truth and provides full transparency into the model's stepwise deductions. In contrast, **InternVL2.5-78B** bases its answer primarily on the width of the trucks and the space between them, omitting the crucial step of accounting for the distance from the pedestrian to the roadway, which leads to significant underestimation. **Gemini1.5Pro** correctly recognizes that the separation includes the truck, traffic lane, and sidewalk, but substantially underestimates the width of the sidewalk, causing a notable error in its final answer. Meanwhile, **SpatialRGPT-8B** provides a more accurate estimate than Gemini or InternVL2.5-78B, but still has a gap compared to the ground truth. Most importantly, it cannot generate step-by-step reasoning traces.

Figure 8 presents another illustrative example evaluating spatial reasoning capabilities of various models, specifically focusing on size comparison between two highlighted image regions. The question is whether Region 1 (a computer monitor) appears smaller than Region 2 (a computer tower). **SpatialReasoner-R1** accurately identifies Region 2 as a computer tower and explicitly reasons by comparing Region 1 with the closest computer tower positioned adjacent to the monitor. This

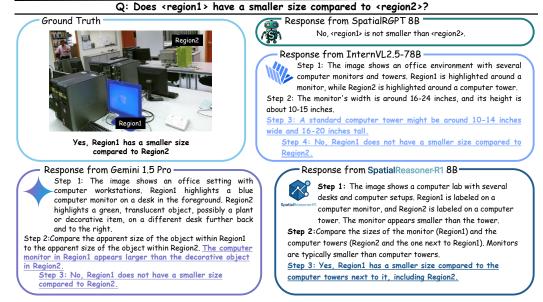


Figure 8: Qualitative Examples of Spatial Reasoning Across Models. SpatialReasoner-R1 correctly recognizes Region2 as a computer tower and compares it clearly with the nearby monitor, reaching an accurate conclusion. InternVL2.5-78B relies on general object size knowledge but provides incorrect reasoning, Gemini1.5Pro fails to identify Region2 clearly and draws incorrect visual conclusions, while SpatialRGPT-8B directly provides a wrong answer.

systematic visual grounding and clear comparative reasoning enable SpatialReasoner-R1 to correctly conclude that the monitor is indeed smaller than the tower. By contrast, the baseline models exhibit varying degrees of errors and reasoning inadequacies. **InternVL2.5-78B** relies significantly on prior general knowledge about typical object dimensions and incorrectly concludes the monitor is not smaller, without effectively validating this against the visual evidence provided. textbfGemini1.5Pro fails entirely to recognize what object Region 2 represents, causing it to inaccurately rely purely on the objects' visual proximity and perspective, leading to an incorrect conclusion. Lastly, the **SpatialRGPT-8B** model directly presents an incorrect judgment ("Region1 is not smaller") without providing any interpretable reasoning steps or visual grounding.

We also provide a failure example in this section. Figure 9 illustrates a representative failure case on vertical size estimation in an indoor setting. The query is "How tall is Region 1?". Here, Region 1 corresponds to the dresser mirror on the right, adjacent to a sleigh-style bed headboard on the left. SpatialReasoner-R1 produces an estimate of approximately 2.0m, while the ground truth is closer to 1.5m. Our model's reasoning proceeds as follows: (1) segments Region 1 (the mirror) and searches for a nearby object of familiar scale, (2) identifies the bed headboard and assumes a typical headboard height of 1.5m, further estimating that the mirror extends about 0.5m above the headboard, and (3) sums these values to obtain about 2.0m. This error arises from overreliance on default furniture priors rather than fully grounding the estimate in image evidence, such as the mirror's vertical extent relative to the floor plane and its contact points with the dresser. To mitigate this, training incorporates fine-grained reward signals that explicitly reward consistency between predicted measurements and image-derived cues, encouraging verification of intermediate steps (e.g., floor contact, vanishing-line alignment) before finalizing a measurement.

G Broader Impacts

This work aims to improve the spatial reasoning capabilities of vision-language models through fine-grained preference optimization. Accurate spatial understanding is critical for downstream applications such as robotics, autonomous navigation, assistive technologies, and visual analytics. By introducing more interpretable and structured reasoning mechanisms, our method can contribute to building AI systems that are safer, more transparent, and more aligned with human expectations in

Response from SpatialReasoner-R1 8B Step 1: The image shows a furnished bedroom with a dresser mirror on the right (Region1) and a sleigh-style spatialReasoner-R1 bed headboard on the left. Region1 is labeled on the dresser mirror, which appears adjacent to the bed. Step 2:Compare the height of the mirror (Region1) with the bed headboard nearby. Bed headboards typically have a standard height of around 1.5 m. The mirror appears to extend approximately 0.5 m taller than the headboard. Step 3: Region1 is approximately 1.5 m headboard height + 0.5 m additional height = 2.0 m total.

Figure 9: Failure case on height estimation in a furnished bedroom.

spatially grounded tasks. However, as with other vision-language systems, potential risks remain. If deployed in safety-critical domains, incorrect spatial inferences, especially in edge cases, could lead to unintended consequences. Additionally, reward scoring and generation rely on foundation models that may encode hidden biases, which can propagate through the training pipeline. Although we attempt to mitigate these risks via multi-source sampling and structured evaluation, future work should explore robustness to distribution shifts, adversarial spatial prompts, and the inclusion of human-in-the-loop verification for high-stakes use cases.

H Limitations

While our work demonstrates strong improvements in spatial reasoning, a limitation of our approach is its reliance on explicit region representations provided as input to disambiguate object references within the spatial queries. Enabling the model to implicitly ground entities solely based on natural language descriptions remains an avenue for future investigation, which would enhance the model's flexibility in real-world scenarios. Future work could focus on integrating implicit linguistic context understanding to alleviate this constraint. Finally, our focus is limited to 2D spatial reasoning; extending this framework to 3D or embodied contexts would require structural adjustments left for future work.

System Prompt for LongCOT Reward Evaluation

The following is a spatial reasoning task, and this is the question: question and the ground truth is: ground_truth. The response is divided into different sections. There are 4 dimensions to evaluate, and I will provide you with the corresponding image and text for reference. You will need to evaluate the response based on the following criteria:

The first task: Descriptive Scoring (Total 0-4.0 points)

Evaluate the "Description" section based on:

- Existence: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Does the description correctly identify objects that actually appear in the image?
- Attribute Accuracy: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Are the object's attributes (color, shape, size, etc.) described accurately?
- Completeness: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Does the description include all key objects and necessary details relevant to the question?
- Appropriateness: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Does the description focus on the core aspects of the question?

Clearly state the score for each sub-category and sum them to obtain the final descriptive score. You need to give the score with the following format: {"task1_score": your score}

The second task: Depth-Guided Spatial Relationship Scoring – Description (Total 0–4.0 points)

Evaluate all spatial statements within the "Description" section using the provided depth image as ground truth.

- For each spatial claim in the description:
 - Correctness score: Assign 1 if the spatial claim is correct based on the depth image, and 0 if not correct.
 - Uncertainty score: For claims expressed with uncertainty (using words like "approximately", "roughly", "possibly"), assign a score from 0.8 to 1.0, where 1.0 means the statement is expressed with high certainty.
 - Relationship score: Assign a weight from 0.8 to 1.0 based on whether the relationship is explicitly
 emphasized by the question (1.0) or is extra/irrelevant information (0.8).

Provide a detailed breakdown for each spatial claim. Calculate the final score as: (Sum of (Correctness score \times Uncertainty score \times Relationship score)) / (Number of claims), then scale to 4.0.

You need to give the score with the following format:

{"task2_claim_score": [Correctness score, Uncertainty score, Relationship score]}

The third task: Depth-Guided Spatial Relationship Scoring – Reasoning (Total 0–4.0 points)

Apply the same evaluation method as in Task 2 to the spatial statements within the "Reasoning" section.

- · For each spatial claim in the reasoning:
 - Correctness score: Assign 1 if the spatial claim is correct based on the depth image, and 0 if not correct.
 - Uncertainty score: For claims expressed with uncertainty, assign a score from 0.8 to 1.0, where 1.0
 means the statement is expressed with high certainty.
 - Relationship score: Assign a weight from 0.8 to 1.0 based on whether the relationship is explicitly emphasized by the question (1.0) or is extra/irrelevant information (0.8).

Provide a detailed breakdown for each spatial claim. Calculate the final score as: (Sum of (Correctness score \times Uncertainty score \times Relationship score)) / (Number of claims), then scale to 4.0.

You need to give the score with the following format:

{"task3_claim_score": [Correctness score, Uncertainty score, Relationship score]}

The fourth task: Reasoning Scoring (Total 0-4.0 points)

Evaluate the "Reasoning" section (the chain-of-thought) based on:

- Factual Consistency: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Are the claims consistent with the image, depth image, and the earlier description?
- Logical Coherence: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Do the reasoning steps flow logically without gaps or contradictions?
- Correct Application of Rules: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Are physical, spatial, and logical rules applied correctly?
- Conclusion Validity: Assign scores from 0 to 1.0, where 1.0 means mostly confidently correct and 0 means mostly confidently incorrect. Does the reasoning properly support the final answer?

Clearly state the score for each sub-category and sum them to obtain the final reasoning score. You need to give the score with the following format: {"task4_score": your score}

Figure 10: **System Prompt for Evaluating LongCoT Spatial Reasoning** w.r.t. descriptive accuracy, spatial alignment, and logical consistency of reasoning steps.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We claim to design a new DPO algorithm for LongCOT optimization, and a reward method, and we verify improvements with extensive experiments.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: Appendix H is about the limitations.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: We do not include theoretical assumptions and proofs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and crossreferenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: Yes, everything is reported, and we will release our data. We provide implementation details in Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We will open source our code at a later date.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: In Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental
 material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: The results are provided as accuracy on the test set. And we have explained how they are calculated in Appendix E.

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).

- It should be clear whether the error bar is the standard deviation or the standard error
 of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: In Appendix D.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

Answer: [Yes]

Justification: We have reviewed the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: In Appendix G.

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: This paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with
 necessary safeguards to allow for controlled use of the model, for example by requiring
 that users adhere to usage guidelines or restrictions to access the model or implementing
 safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do
 not require this, but we encourage authors to take this into account and make a best
 faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Yes. We acknowledge the codebases.

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

• If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- · Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and research with human subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional review board (IRB) approvals or equivalent for research with human subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: The paper does not involve crowdsourcing nor research with human subjects. Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- · For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. Declaration of LLM usage

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this paper does not involve LLMs as any important, original, or non-standard components.

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.