

FoREST: Frame of Reference Evaluation in Spatial Reasoning Tasks

Anonymous ACL submission

Abstract

Spatial reasoning is a fundamental aspect of human intelligence. One key concept in spatial cognition is the Frame of Reference (FoR), which identifies the perspective of spatial expressions. Despite its significance, FoR has received limited attention in AI models that need spatial intelligence. There is a lack of dedicated benchmarks and in-depth evaluation of large language models (LLMs) in this area. To address this issue, we introduce the **Frame of Reference Evaluation in Spatial Reasoning Tasks (FoREST)** benchmark, designed to assess FoR comprehension in LLMs. We evaluate LLMs on answering questions that require FoR comprehension and layout generation in text-to-image models using FoREST. Our results reveal a notable performance gap across different FoR classes in various LLMs, affecting their ability to generate accurate layouts for text-to-image generation. This highlights critical shortcomings in FoR comprehension. To improve FoR understanding, we propose Spatial-Guided prompting, which improves LLMs' ability to extract essential spatial concepts. Our proposed method improves overall performance across spatial reasoning tasks.

1 Introduction

Spatial reasoning plays a significant role in human cognition and daily activities. It is also a crucial aspect in many AI problems, including language grounding (Zhang and Kordjamshidi, 2022; Yang et al., 2024), navigation (Yamada et al., 2024), computer vision (Liu et al., 2023; Chen et al., 2024), medical domain (Gong et al., 2024), and image generation (Gokhale et al., 2023). One key concept in spatial reasoning is the Frame of Reference (FoR), which identifies the perspective of spatial expressions. FoR has been studied extensively in cognitive linguistics (Edmonds-Wathen, 2012; Vukovic and Williams, 2015). Levinson (2003) initially defines three FoR classes: *relative*, based

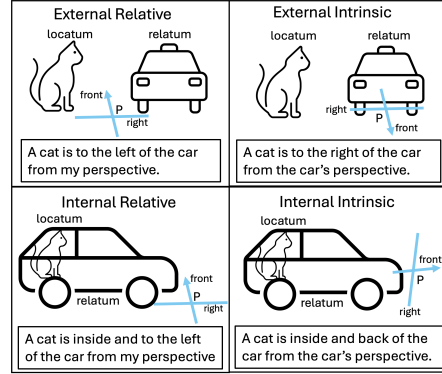


Figure 1: Illustration of FoR classes. The cat is the locatum, the car is the relatum, and the arrow denotes the perspective.

on the observer's perspective; *intrinsic*, based on an inherent feature of the reference object; and *absolute*, using environmental cues like cardinal directions -See Figure 1. This framework was expanded by Tenbrink (2011) to create a more comprehensive framework, serving as the basis of our work. Understanding FoR is important for many applications, especially in embodied AI. In such applications, an agent must simultaneously comprehend multiple perspectives, including the one from the instruction giver and from the instruction follower, to communicate and perform tasks effectively. However, recent spatial evaluation benchmarks have largely overlooked FoR. For example, the text-based benchmarks Shi et al. (2022); Mirzaee and Kordjamshidi (2022); Rizvi et al. (2024) and text-to-images benchmarks (Gokhale et al., 2023; Huang et al., 2023; Cho et al., 2023a,b) assume a fixed perspective for all spatial expressions. This inherent bias limits situated spatial reasoning, restricting adaptability in interactive environments where perspectives can change.

To systematically investigate the role of FoR in spatial understanding and create a new resource, that is, **Frame of Reference Evaluation in Spatial**

Reasoning Tasks (FoREST), FoREST is designed to evaluate models’ ability to comprehend FoR from textual descriptions and extend this evaluation to grounding and visualization. Our benchmark includes spatial expressions with FoR ambiguity—where multiple FoRs may apply to the described situation—and spatial expressions with only a single valid FoR. This design allows evaluation of the models’ understanding of FoR in both scenarios. We evaluate several LLMs in a QA setting that require FoR understanding and apply the FoR concept in text-to-image models. Our findings reveal performance differences across FoR classes and show that LLMs exhibit bias toward specific FoRs when handling ambiguous cases. This bias extends to layout-diffusion models, which rely on LLM-generated layouts in the image generation pipeline. To enhance FoR comprehension in LLMs, we propose Spatial-Guided prompting, which enables models to analyze and extract additional spatial information, including directional, topological, and distance relations. We demonstrate that incorporating spatial information improves question-answering and layout generation, ultimately enhancing text-to-image generation performance.

Our contribution¹ are summarized as follows, 1. We introduce the FoREST benchmark to evaluate LLMs’ FoR comprehension in a QA setting systematically. 2. We analyze the impact of FoR information on text-to-image generation using multiple diffusion models. 3. We propose a prompting approach that generates spatial information, which can be incorporated into QA and layout diffusion to enhance performance.

2 Spatial Primitives

We review three semantic aspects of spatial information expressed in language: Spatial Roles, Spatial Relations, and Frame of Reference.

Spatial Roles. We focus on two main spatial roles (Kordjamshidi et al., 2010) of *Locatum*, and *Relatum*. The locatum is the object described in the spatial expression, while the relatum is the other object used to describe the position of the locatum. An example is *a cat is to the left of a dog*, where the *cat* is the locatum, and the *dog* is the relatum.

Spatial Relations. When dealing with spatial knowledge representation and reasoning, three main relations are often considered, directional, topological, and distance (Hernández, 1994; Cohn

and Renz, 2008; Kordjamshidi et al., 2011). *Directional* describes an object’s direction based on specific coordinates, e.g., left and right. *Topological* describes the containment between two objects, such as inside. *Distance* describes qualitative and quantitative relations between entities. Examples of qualitative are far, and quantitative are 3km.

Spatial Frame of Reference. We use four frames of references investigated in the cognitive linguistic studies (Tenbrink, 2011). These are defined based on the concept of *Perspective*, which is the origin of a coordinate system to determine the direction. The four frames of reference are defined as follows.

1. *External Intrinsic* describes a spatial relation from the relatum’s perspective, where the relatum does not contain the locatum. The top-right image in Figure 1 illustrates this with the sentence, *A cat is to the right of the car from the car’s perspective*.
2. *External Relative* describes a spatial relation from the observer’s perspective. The top-left image in Figure 1 shows an example with the sentence, *A cat is to the left of a car from my perspective*.
3. *Internal Intrinsic* describes a spatial relation from the relatum’s perspective, where the relatum contains the locatum. The bottom-right image in Figure 1 shows this with the sentence, *A cat is inside and back of the car from the car’s perspective*.
4. *Internal Relative* describes a spatial relation from the observer’s perspective where the locatum is inside the relatum. The bottom-left image in Figure 1 shows this FoR with the sentence, *A cat is inside and to the left of the car from my perspective*.

3 FoREST Dataset Construction

To systematically evaluate LLM on the frame of reference (FoR) recognition, we introduce the **Frame of Reference Evaluation in Spatial Reasoning Tasks (FoREST)** benchmark. Each instance in FoREST consists of a spatial context (T), a set of corresponding FoR (FoR) which is a subset of $\{external\ relative, external\ intrinsic, internal\ intrinsic, internal\ relative\}$, a set of questions and answers ($\{Q, A\}$), and a set of visualizations ($\{I\}$). An example of T is *A cat is to the right of a dog. A dog is facing toward the camera*. The FoR of T is $\{external\ intrinsic, external\ relative\}$. A question-answer is $Q = Based\ on\ the\ camera’s\ perspective, where\ is\ the\ cat\ from\ the\ dog’s\ position?, A = \{left, right\}$. There is an ambiguity in the FoR for this expression. Thus, the answer will be *left* if the model assumes the external relative. Conversely, it will

¹code and dataset available at [Anonymous repository](#).

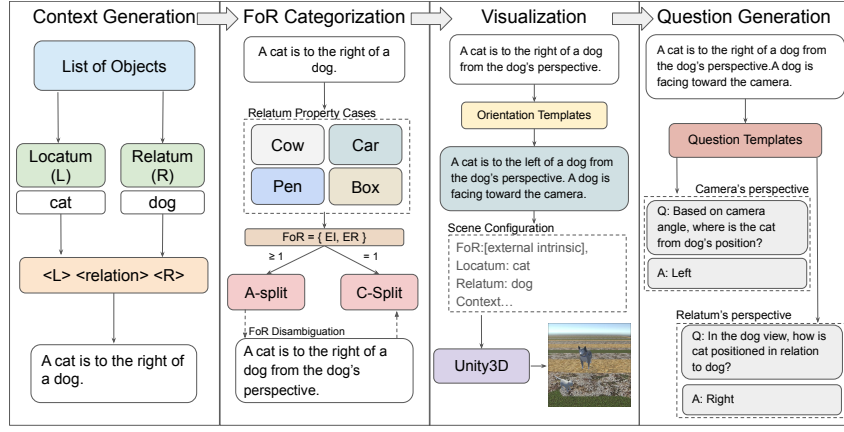


Figure 2: Pipeline for dataset creation, starting from selecting a locatum and relatum from available objects and then applying a spatial template to generate the spatial expression (T). FoRs are assigned based on relatum properties. T is categorized based on the number of FoRs. For example, *A cat is to the right of a dog* (with two possible FoRs: external intrinsic and external relative) belongs to the A-split. Then, its disambiguated version (*A cat is to the right of a dog from the dog’s perspective*) is added to the C-split. Next, if applicable, a relatum’s orientation is included for visualization and question generation. Finally, Unity3D generates scene configurations, and question-answer pairs are derived from T .

be *right* if the model assumes the external intrinsic. The visualization of this example is in Figure 2.

3.1 Context Generation

We select two distinct objects—a relatum (R) and a locatum (L)—from a set of 20 objects and apply them to a Spatial Relation template, $\langle L \rangle \langle spatial\ relation \rangle \langle R \rangle$ to generate the context T . FoRs for T are determined based on the properties of the selected objects. Depending on the number of possible FoRs, T is categorized as ambiguous (A-split), where multiple FoRs apply, or clear (C-split), where only one FoR is valid. We further augment the C-split with disambiguated spatial expressions derived from the A-split, as shown in Figure 2.

3.2 Categories based on Relatum Properties

Using the FoR classes in Section 2, we identify two key properties contributing to FoR ambiguity. The first property is the relatum’s intrinsic direction. It creates ambiguity between intrinsic and relative FoRs as spatial relations may originate from either relatum’s or observer’s perspectives. The second is the relatum’s affordance as a container. It introduces the ambiguity between internal and external FoR, as spatial relations may refer to either inside or outside of the relatum. Note that containment is defined as the ability of the relatum to contain the locatum, considering both objects’ sizes. Based on these properties, we define four distinct cases: *Cow*, *Box*, *Car*, and *Pen*.

Case 1: Cow Case. In this case, the selected re-

latum has intrinsic directions but cannot be the container for the locatum. An example is a cow. In such cases, the relatum potentially provides a perspective for spatial relations. The applicable FoR classes are $FoR = \{external\ intrinsic, external\ relative\}$. We augment the C-split with expressions of this case but include the perspective to resolve their ambiguity. To specify the perspective, we use predefined templates for augmenting clauses, such as *from {relatum}’s perspective* for *external intrinsic* or *from the camera’s perspective* for *external relative*. For example, if the context is *A cat is to the right of the cow*, in the A-split. The counterparts included in the C-split are *A cat is to the right of the cow from cow’s perspective*. for *external intrinsic* and *A cat is to the right of the cow from my perspective*. for *external intrinsic*.

Case 2: Box Case. The relatum in this case has the property of being a container but lacks intrinsic directions, making the internal FoR applicable. An example is a box. The applicable FoR classes are $FoR = \{external\ relative, internal\ relative\}$. To include their unambiguous counterparts in the C-split, we specify the topological relation by adding *inside* for *internal relative* and *outside* for *external relative*. For example, for the sentence *A cat is to the right of the box.*, the unambiguous T with *internal relative* FoR is *A cat is inside and to the right of the box*. The counterpart for *external relative* is *A cat is outside and to the right of the box*.

Case 3: Car Case. A relatum with an intrinsic direction and container affordance falls into this case,

allowing all FoR classes. An obvious example is a car that can be a container with intrinsic directions. The applicable FoR classes are $FoR = \{ external\ relative, external\ intrinsic, internal\ intrinsic, internal\ relative \}$. To augment C-split with this case’s disambiguated counterparts, we add perspective and topology information similar to the Cow and Box cases. An example expression for this case is *A person is in front of the car*. The four disambiguated counterparts to include in the C-split are *A person is outside and in front of the car from the car itself*. for *external intrinsic*, *A person is outside and in front of the car from the observer*. for *external relative*, *A person is inside and in front of the car from the car itself*. for *internal intrinsic*, and *A person is inside and in front of the car from the observer*. for *internal relative*.

Case 4: Pen Case. In this case, the relatum lacks both the intrinsic direction and the affordance as a container. An obvious example is a pen. Lacking these two properties, the created context has only one applicable FoR, $FoR = \{ external\ relative \}$. Therefore, we can categorize this case into both splits without any modification. An example of such a context is *The book is to the left of a pen*.

3.3 Context Visualization

In our visualization, complexity arises when the relatum has an intrinsic direction within the intrinsic FoR, as its orientation can complicate the spatial representation. For example, for visualization of *A cat is to the right of a dog from the dog’s view*, the cat can be placed in different coordinates based on the dog’s orientation. To address this issue, we add a template sentence for each direction, such as *<relatum> is facing toward the camera*, to specify the relatum’s orientation of all applicable T for visualization and QA. For instance, *A cat is to the left of a dog*. becomes *A cat is to the left of a dog. The dog is facing toward the camera*. To avoid occlusion issues, we generate visualizations only for external FoRs, as one object may become invisible in internal FoR classes. We use only expressions in C-split since those have a unique FoR interpretation for visualization. We then create a scene configuration by applying a predefined template, as illustrated in Figure 2. Images are generated using the Unity 3D simulator (Juliani et al., 2020), producing four variations per expression T with different backgrounds and object positions. Further details on the simulation process are in Appendix A.2.

3.4 Question-Answering Generation

We generate questions for all generated spatial expressions (T). Note that we include the relatum orientation for cases where the relatum has an intrinsic direction, as mentioned in the visualization. Our benchmark includes two types of questions. The first type asks for the spatial relation between two given objects from the camera’s perspective, following predefined templates such as, *Based on the camera’s perspective, where is the locatum relative to the relatum’s position?* Template variations are made based on GPT4o; details are provided in Appendix A.2.4. The second type of question queries the spatial relation from the relatum’s perspective. This question type follows the same templates but replaces the camera with the relatum. The first type of question is generated for all T , while the second type is only generated for T where the relatum has intrinsic direction and a perspective can be defined accordingly. Answers are determined based on the corresponding FoRs, the spatial relation in T , and the relatum’s orientation when applicable.

4 Models and Tasks

This paper focuses on Question-Answering and Text-to-Image using the FoREST benchmark to evaluate FoR in spatial reasoning comprehensively. FoREST also supports other tasks, such as FoR identification, detailed in Appendix H.

4.1 Question-Answering (QA)

Task. This QA task evaluates LLMs’ ability to adapt contextual perspectives across different FoRs. The input includes a spatial expression T , relatum orientation (if available), and a question Q querying the spatial relation from either the observer’s or the relatum’s perspective. The output is a spatial relation S , restricted to {left, right, front, back}. We use the context from both splits for this task.

Zero-shot baseline. We call the LLM with instructions, a spatial context, and a question expecting a spatial relation as the response. The prompt instructs the model to answer the question with one of the candidate spatial relations without explanation.

Few-shot baseline. We create four spatial expressions, each assigned to a single FoR class to prevent bias. Following the steps in Section 3.4, we generate a corresponding question and answer for each. These serve as examples in our few-shot prompting. The input to the model is the instruction, example, spatial context, and the question.

Chain-of-Thought (CoT) baseline (Wei et al., 2023). To create CoT examples, we modify the prompt to require reasoning before answering. We manually craft reasoning explanations with the necessary information for the few-shot examples. The input to the model is the instruction, CoT example, spatial context, and the question.

Human baseline. We conducted a small scope human study with three participants, paid as research assistants, to establish a human baseline for A-split. Each participant is shown 25 randomly sampled QA examples per relatum case, totaling 150 examples. Each includes a textual scene description and corresponding question. See Appendix I for details. We report average accuracy across participants to reflect the human baseline.

4.2 Text-To-Image (T2I)

Task. This task assesses the diffusion models’ ability to consider FoR by evaluating their generated images. The input is a spatial expression, T , and the output is a generated image (I). C and A splits with external FoRs are used for this task.

Stable Diffusion Models. We use the stable diffusion models as the baseline for the T2I task. This model only needs the scene description as input.

Layout Diffusion Models. This model operates in two phases: text-to-layout and layout-to-image. Given that LLMs can generate the bounding box layout (Cho et al., 2023b), we provide them with instructions and T to create the layout. The layout consists of bounding box coordinates for each object in the format of {object: $[x, y, w, h]$ }, where x and y denote the starting point and h and w denote the height and width. The bounding box coordinates and T are then passed to the layout-to-image model to produce the final image, I .

4.3 Spatial-Guide (SG) Prompting

We hypothesize that the spatial relation types and FoR classes defined in Section 2 can improve QA and layout generation. For example, the *external intrinsic* FoR emphasizes that spatial relations originate from the relatum’s perspective. To leverage this, we propose SG prompting, an additional step applied before QA or layout generation. This step extracts spatial information, including direction, topology, distance, and the FoR, from the spatial expression T , as supplementary input to guide LLMs in QA or layout generation. We manually craft four examples covering these aspects. First, we specify the perspective for *directional relations*, e.g.,

left relative to the observer, to distinguish intrinsic from relative FoR. Next, we indicate whether the locatum is inside or outside the relatum for *topological relations* to differentiate internal from external FoR. Lastly, we provide an estimated quantitative distance to support topological and directional relation identification, e.g., *far*. These examples are provided in a few-shot setting to guide the model in automatically extracting such information.

5 Experimental Results

5.1 Evaluation Metrics

QA. We report an accuracy (acc.) based on the correct answer defined as follows. Since the questions can have multiple correct answers, as explained in Section 3, the prediction is correct if it matches any valid answer. Additionally, we report the model’s bias distribution when FoR ambiguity exists. $I\%$ is the percentage of correct answers when assuming an intrinsic FoR, while $R\%$ is this percentage with a relative FoR assumption. Note that cases where both FoR assumptions lead to the same answer are excluded from the bias calculation.

T2I. We adopt *spatialEval* (Cho et al., 2023b) for evaluating T2I spatial ability. We modify it to account for FoR by converting relations to a camera perspective before passing them to *spatialEval*, which assumes this viewpoint. Accuracy is determined by comparing the bounding box and depth map of the relatum and locatum. For FoR ambiguity, a generated image is correct if it aligns with at least one valid FoR interpretation. We report results using $VISOR_{cond}$ and $VISOR_{uncond}$ (Gokhale et al., 2023) metrics. $VISOR_{cond}$ evaluates spatial relations only when both objects appear correctly, aligning with our focus on spatial reasoning. While, $VISOR_{uncond}$ evaluates the overall performance, including object creation errors.

5.2 Experimental Setting

QA. We use Llama3-70B (Llama, 2024), Qwen2-72B (Qwen Team, 2024), Qwen2VL-72B (Wang et al., 2024), and GPT-4o (OpenAI, 2024) as backbone models. All models are evaluated with *zero-shot*, *few-shot*, *CoT*, and our SG+CoT prompting under temperature 0 to ensure reproducibility.

T2I. We select Stable Diffusion SD-1.5 and SD-2.1 (Rombach et al., 2021) as our stable diffusion models and GLIGEN (Li et al., 2023) as the layout-to-image backbone. For translating spatial descriptions into bounding box information, we

Model	Camera perspective									Relatum perspective							
	Cow			Car			Box	Pen	Avg.	Cow			Car			Avg.	
	R%	I%	Acc.	R%	I%	Acc.	Acc.	Acc.	Acc.	R%	I%	Acc.	R%	I%	Acc.	Acc.	
Llama3-70B (1)	48.1	51.5	62.5	58.0	41.6	65.5	73.3	72.5	64.3	61.0	38.7	62.1	51.8	47.9	61.8	62.1	
Llama3-70B (2)	49.1	50.5	62.2	52.2	47.4	64.7	85.8	85.5	65.8	59.6	40.1	57.1	55.5	44.2	61.8	57.7	
Llama3-70B (3)	49.4	50.3	80.7	49.4	50.3	79.6	95.8	94.9	82.6	60.8	39.0	77.2	55.1	44.6	80.9	77.7	
Llama3-70B (4)	59.4	40.2	73.6	57.9	41.7	74.8	100.0	100.0	77.5	60.6	39.1	65.7	56.0	43.7	67.7	66.0	
Qwen2-72B (1)	96.6	2.9	95.6	95.9	3.6	95.0	100.0	100.0	96.1	8.8	90.6	79.3	7.8	91.7	83.6	79.9	
Qwen2-72B (2)	89.0	10.5	84.4	85.6	13.9	85.5	100.0	100.0	86.8	17.7	81.8	78.3	10.4	89.1	86.3	79.4	
Qwen2-72B (3)	67.2	32.4	88.6	62.0	37.6	83.4	100.0	100.0	89.6	21.3	78.3	85.5	22.7	76.9	83.6	85.2	
Qwen2-72B (4)	93.0	6.5	90.1	94.6	4.9	93.3	100.0	98.6	91.7	8.2	91.2	86.0	10.5	89.0	87.4	86.2	
Qwen2VL-72B (1)	49.5	50.5	78.3	51.6	48.4	80.0	98.3	96.9	81.2	41.0	59.0	55.4	44.7	55.3	59.0	56.0	
Qwen2VL-72B (2)	40.9	59.1	89.4	44.7	55.3	79.0	100.0	100.0	89.6	32.7	67.3	66.9	28.5	71.5	67.0	66.9	
Qwen2VL-72B (3)	63.5	36.5	84.0	72.2	27.8	84.5	100.0	100.0	86.3	51.0	49.0	77.7	54.6	45.4	82.8	78.4	
Qwen2VL-72B (4)	50.5	49.5	78.3	59.5	40.5	67.4	98.3	99.6	79.9	23.5	76.5	78.2	33.0	67.0	67.7	76.7	
GPT-4o (1)	84.3	15.3	94.5	88.5	11.0	97.3	99.2	99.8	95.6	21.6	78.0	91.6	16.1	83.5	90.5	91.4	
GPT-4o (2)	69.0	30.6	76.6	80.3	19.2	89.5	100.0	100.0	81.5	29.0	70.5	74.7	30.9	68.7	77.5	75.1	
GPT-4o (3)	41.5	58.3	92.3	38.2	61.6	91.0	100.0	99.8	93.2	33.9	65.8	93.9	32.0	67.6	93.9	93.9	
GPT-4o (4)	26.0	73.9	79.2	27.7	72.1	79.4	96.7	94.3	81.4	16.2	83.4	95.5	19.2	80.4	94.8	95.4	
Human-baseline	36.6	63.4	90.7	27.8	72.2	96.0	72.0	82.7	85.3	41.4	58.6	97.3	36.1	63.9	96.0	96.7	

Table 1: QA accuracy in the A-Split. R% and I% represent the percentage the model assumes relative or intrinsic FoR for ambiguous expression, explained in Section 5.1. Acc is the accuracy, and Avg is the micro-average of accuracy. (1): 0-shot, (2): 4-shot, (3): CoT, and (4): SG+CoT.

Model	Camera perspective					Relatum perspective				
	ER (CP)	EI (RP)	II (RP)	IR (CP)	Avg.	ER (CP)	EI (RP)	II (RP)	IR (CP)	Avg.
Llama3-70B (0-shot)	44.8	38.4	39.7	54.4	42.6	42.2	47.1	62.5	34.4	45.1
Llama3-70B (4-shot)	43.0	40.0	39.1	47.3	41.9	41.8	60.9	77.7	35.2	52.0
Llama3-70B (CoT)	57.8	46.1	44.7	46.0	51.5	55.5	56.8	71.5	49.0	56.6
Llama3-70B (SG+CoT)	47.6	42.9	50.0	35.6	45.0	55.4	64.5	75.0	47.1	60.1
Qwen2-72B (0-shot)	94.5	35.2	31.8	93.2	66.9	28.7	89.3	93.6	23.8	59.0
Qwen2-72B (4-shot)	90.2	39.5	39.1	68.5	65.3	33.5	92.1	94.0	29.5	62.7
Qwen2-72B (CoT)	81.4	57.4	58.6	62.5	69.1	39.5	83.7	85.2	37.7	61.6
Qwen2-72B (SG+CoT)	97.6	42.5	31.3	93.8	71.4	42.8	86.6	92.0	34.0	64.5
Qwen2VL-72B (0-shot)	68.3	40.2	42.4	75.3	56.0	37.4	60.1	82.2	35.5	50.0
Qwen2VL-72B (4-shot)	78.4	41.0	44.5	66.1	60.5	38.0	85.4	91.2	29.9	61.6
Qwen2VL-72B (CoT)	58.4	62.4	68.4	40.7	59.3	48.6	43.5	36.1	61.9	46.4
Qwen2VL-72B (SG+CoT)	99.0	50.0	47.5	93.0	75.9	27.3	37.3	37.1	34.0	32.7
GPT-4o (0-shot)	79.7	45.1	39.5	90.2	64.2	46.9	88.5	98.2	34.8	67.5
GPT-4o (4-shot)	68.0	52.6	60.7	74.1	61.8	44.9	98.2	100.0	37.5	71.2
GPT-4o (CoT)	81.7	76.1	82.4	71.5	78.8	53.0	91.1	90.6	50.8	71.9
GPT-4o (SG+CoT)	97.9	72.2	72.7	93.4	85.8	48.9	96.3	95.9	36.1	71.8

Table 2: QA accuracy in the C-Split across various LLMs. ER, EI, II, and IR denote external relative, external intrinsic, internal intrinsic, and internal relative FoRs. Avg represents the micro-average accuracy. CP refers to context with camera perspective, while RP denotes context with relatum perspective.

use Llama3-8B and Llama3-70B, as detailed in Section 4.2. The same LLMs are used to generate spatial information for SG prompting. We generate four images to compute the VISOR score following (Gokhale et al., 2023). Inference steps for all T2I models are set to 50. For the evaluation, we select grounding DINO (Liu et al., 2025) for object detection and DPT (Ranftl et al., 2021) for depth mapping, following VPEval (Cho et al., 2023b). The experiments were conducted on two A6000 GPUs, totaling approximately 300 GPU hours.

5.3 Results

RQ1. What is the LLM’s bias for the ambiguous FoR? Table 1 presents the QA results for the A-split. Ideally, a model extracting spatial relations without considering perspective should achieve 100% accuracy, as the context lacks a fixed perspective. However, our focus is not on this. We aim to assess LLMs’ bias by measuring how often

they adopt a specific perspective when answering. In the Box and Pen cases, relatum properties do not introduce FoR ambiguity in directional relations, making the task pure extraction rather than reasoning. Thus, we focus on the I% and R% of the Cow and Car cases, which best reflect LLMs’ bias. Qwen2 achieves 80% accuracy across all experiments by selecting spatial relations directly from context, suggesting it may disregard the question’s perspective. This is supported by attention analysis in Appendix C. The attention maps show that most correct Qwen2 responses assume a shared perspective between context and question and pay low attention to the perspective mentioned in the question. GPT-4o exhibits similar bias in 0-shot and 4-shot settings but shifts toward intrinsic interpretation with CoT. This shift reduces accuracy in camera-perspective questions where FoR adaptation is more significant than relation extraction. Llama3-70B shows no strong preference, slightly

Model	VISOR(%)							
	A-Split				C-Split			
	cond (I)			cond (R)	cond (avg)	cond (I)	cond (R)	cond (avg)
	EI FoR	ER FoR	all					
SD-1.5	51.11	21.61	72.72	48.95	68.72	53.92	53.77	53.83
SD-2.1	57.97	21.49	79.46	54.10	75.39	60.06	59.64	59.83
Llama3-8B + GLIGEN	53.67	25.78	79.45	66.08	77.38	57.51	65.98	62.12
Llama3-70B + GLIGEN	54.49	29.45	83.94	68.68	81.43	56.47	69.53	63.49
Llama3-8B + SG + GLIGEN (Our)	57.46	27.96	85.42	71.14	83.17	58.84	70.36	65.15
Llama3-70B + SG + GLIGEN (Our)	56.54	30.59	87.13	66.56	83.75	56.77	70.04	64.06

Table 3: VISOR_{cond} score explained in Section 5.1 where *I* refers to the Cow and Car cases where relatum has intrinsic directions, and *R* refers to the Box and Pen cases where relatum lacks intrinsic directions, *avg* is the micro-average of *I* and *R*. EI and ER FoR represent the generated image considered corrected by EI or ER FoR.

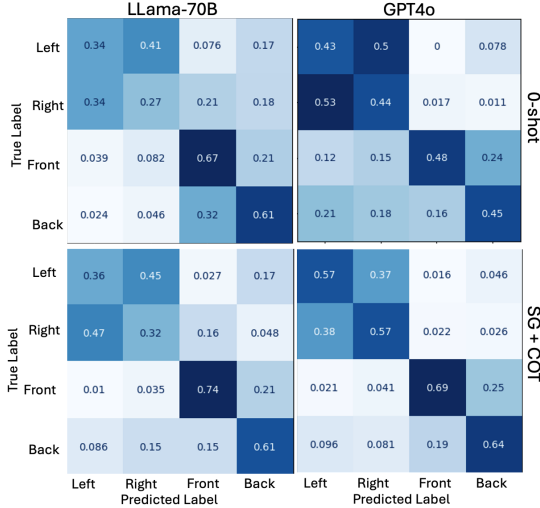


Figure 3: Confusion matrices of predicted spatial relations by Llama3 and GPT-4o when adapting to FoR is required. 0-shot and SG-CoT settings are reported here.

favoring relative FoR, but this balanced assumption leads to lower performance due to increased reasoning demands. Qwen2VL exhibits a similar pattern but achieves higher accuracy, likely due to visual training. GPT-o4-mini with high visual reasoning shows the same pattern, which is later discussed as an additional experiment in Appendix J. Moreover, our experiments with humans, presumed to have prior visual knowledge, show strong reasoning ability, achieving over 90% accuracy even in scenarios where perspective preference requires FoR adaptation. In our results, humans slightly prefer intrinsic FoR over relative FoR. Their cultural background and recently seen examples may also influence their preference as discussed in (Coventry et al., 2018). A more detailed analysis of human performance is provided in Appendix I. In summary, we conjecture that Qwen2 performs well by focusing on extraction without reasoning FoR, while other models attempt reasoning but often fail to reach correct conclusions, leading to lower accuracy.

RQ2. Can LLMs adapt FoR when answering

the questions? To address this RQ, we analyze QA results of C-split in Table 2, where context and question explicitly specify perspective. Results show that LLMs struggle with FoR conversion, achieving only up to 55.5% accuracy when the question uses the relatum’s perspective and the context uses the camera’s. We further examine how Llama3 and GPT-4o adapt FoR using the confusion matrix in Figure 3. Our findings reveal that pure-text LLM (Llama3) reverses left and right. This contrasts with humans when describing the spatial relation from a perspective in English, which reverses front and back while preserving lateral directions. This leads to poor adaptation to the camera perspective of Llama3. In contrast, large multimodal models like GPT-4o follow expected patterns, consistent with Zhang et al. 2025. While our GPT-4o results suggest some abilities to convert the relatum’s perspective into the camera’s with in-context learning (72% accuracy), the reverse transformation in the textual domain remains challenging (53% accuracy). A similar trend appears when comparing Qwen2 and Qwen2VL, as discussed in Appendix B. Qwen2VL performs better when shifting to the camera perspective but performs worse with reverse transformation, often failing to generate correct reasoning. This difficulty persists for several models when converting spatial relations from images to relatum’s perspective, as noted in Zhang et al. 2025. A possible reason is that the model often overlooks orientation tokens, critical for answering FoR questions, based on Attention analysis provided in Appendix C.

RQ3. How can an explicit FoR identification help spatial reasoning in QA? We compare CoT and CoT+SG results to assess the effect of FoR identification on LLMs’ spatial reasoning in QA. Based on C-Split results (Table 2), incorporating SG improves the model’s ability to identify the correct perspective from input expression ranging from 2.9% to 30% of cases where the context and

question share the same perspective. These cases are easier as the models do not need FoR adaptation. Two notable exceptions emerge. First, Llama3 performs poorly on camera-perspective questions, and FoR identification via SG fails to improve its performance. This may be due to Llama3’s lack of visual training, which we speculate limits its FoR understanding. Second, Qwen2VL struggles with relatum-perspective reasoning, showing negative gains even with CoT. SG is less effective when context and question differ in perspective; while it helps identify the correct FoR in context, it does not enhance reasoning across perspectives. This limitation is evident in A-Split results (Table 1), where models only improve significantly when SG aligns their preference with the question’s perspective, as seen in Qwen2-72B and GPT-4o. SG identification results are reported in the Appendix H. Still, FoR identification improves overall spatial reasoning (see Avg. column for SG+CoT in Table 3).

RQ4. How can explicit FoR identification help spatial reasoning in visualization? We evaluate SG layout diffusion to assess the impact of incorporating FoR in image generation. We focus on $VISOR_{cond}$ metric, which better reflects the model’s spatial understanding than the overall performance measured by $VISOR_{uncond}$, reported in Appendix F. Table 3 shows that adding spatial information and FoR classes (SG+GLIGEN) improves performance across all splits compared to the baseline models (GLIGEN). In particular, SG improved the model’s performance when expressions follow a relative FoR. This finding aligns with QA results in Table 1, where *Llama3 prefers relative FoR in camera-perspective scenarios*. In contrast, baseline diffusion models (SD-1.5 and SD-2.1) perform better for intrinsic FoR, even though GLIGEN is based on SD-2.1. This outcome may be due to GLIGEN’s reliance on bounding boxes for spatial configurations, which makes it struggle with intrinsic FoR due to the absence of object properties and orientation. Despite this bias, incorporating FoR information via SG-prompting improves performance across all FoR classes. We provide further analysis on SG for the layout generation in Appendix G.

6 Related Work

Frame of Reference in Cognitive Studies. The concept of the frame of reference in cognitive studies was introduced by Levinson 2003 and later expanded with more diverse spatial relations

(Tenbrink, 2011). Subsequent research investigated the human preferences for specific FoR classes (Vukovic and Williams, 2015; Shusterman and Li, 2016; Ruotolo et al., 2016; Coventry et al., 2018). For instance, Ruotolo et al. 2016 examined how FoR influences scene memorization. Their study found that participants performed better when spatial relations were based on their position rather than external objects, highlighting a distinction between relative and intrinsic FoR.

Frame of Reference in AI. Several benchmarks have been developed to evaluate the spatial understanding of AI models in multiple modalities; for instance, textual QA (Shi et al., 2022; Mirzaee and Kordjamshidi, 2022; Rizvi et al., 2024), and text-to-image (T2I) benchmarks (Gokhale et al., 2023; Cho et al., 2023a,b). However, most of these benchmarks overlook FoR, assuming a single FoR for all instances despite its significance in cognitive studies. Recent vision-language studies have begun addressing this gap (Liu et al., 2023; Zhang et al., 2025). For instance, Zhang et al. 2025 explores FoR ambiguity by evaluating spatial relations from camera-perspective images, with FoR specified in the question. In contrast, our work focuses on the spatial reasoning when dealing with multiple FoRs and when there are changes in perspective in explaining the context beyond the camera’s viewpoint. We further show that explicitly identifying FoR enhances spatial reasoning in QA and T2I.

7 Conclusion

Given the significance of spatial reasoning in AI, we introduce **Frame of Reference Evaluation in Spatial Reasoning Tasks** (FoREST) benchmark to evaluate Frame of Reference comprehension in textual spatial expressions through question-answering and grounding in visual modality via diffusion models. Using this benchmark, we identify notable differences in FoR comprehension across LLMs and their struggle with questions that require adaptation between multiple FoRs. Moreover, the bias in FoR interpretations impacts the layout generation with LLMs for text-to-image models. To improve FoR comprehension, we propose Spatial-Guided prompting, which extracts spatial relations’ topological, distal, and directional type information in addition to FoR and includes this information in downstream task prompting. Employing SG improves performance in both QA tasks requiring FoR understanding and text-to-image generation.

8 Limitations

While we analyze LLMs’ shortcomings, our benchmark only highlights areas for improvement. The trustworthiness and reliability of the LLMs are still a research challenge. Our analysis is confined to the spatial reasoning domain and does not account for biases related to gender or race. However, we acknowledge that linguistic and cultural variations in spatial expression are not considered, as our study focuses solely on English. Extending this work to multiple languages could reveal important differences in FoR adaptation. Our analysis is still limited to the synthetic environment. Future research should consider the broader implications of the frame of reference of spatial reasoning in real-world applications. Additionally, our experiments require substantial GPU resources, limiting the selection of LLMs and constraining the feasibility of testing larger models. The computational demands also pose accessibility challenges for researchers with limited resources. We find no ethical concerns in our methodology or results, as our study does not involve human subjects or sensitive data.

References

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

Jaemin Cho, Abhay Zala, and Mohit Bansal. 2023a. Dall-eval: Probing the reasoning skills and social biases of text-to-image generation models. In *ICCV*.

Jaemin Cho, Abhay Zala, and Mohit Bansal. 2023b. Visual programming for step-by-step text-to-image generation and evaluation. In *Advances in Neural Information Processing Systems*, volume 36, pages 6048–6069. Curran Associates, Inc.

Anthony G. Cohn and Jochen Renz. 2008. Chapter 13 qualitative spatial representation and reasoning. In Frank van Harmelen, Vladimir Lifschitz, and Bruce Porter, editors, *Handbook of Knowledge Representation*, volume 3 of *Foundations of Artificial Intelligence*, pages 551–596. Elsevier.

Kenny R. Coventry, Elena Andonova, Thora Tenbrink, Harmen B. Gudde, and Paul E. Engelhardt. 2018. Cued by what we see and hear: Spatial reference frame use in language. *Frontiers in Psychology*, Volume 9 - 2018.

Cris Edmonds-Wathen. 2012. False friends in the multilingual mathematics classroom. In *.*, pages 5857–5866.

Pierre-Pascal Forster, Katja Fiehler, and Harun Karim-pur. 2023. Egocentric cues influence the allocentric spatial memory of object configurations for memory-guided actions. *Journal of Neurophysiology*, 130(5):1142–1149. PMID: 37791381.

Tejas Gokhale, Hamid Palangi, Besmira Nushi, Vibhav Vineet, Eric Horvitz, Ece Kamar, Chitta Baral, and Yezhou Yang. 2023. Benchmarking spatial relationships in text-to-image generation. *Preprint*, arXiv:2212.10015.

Shizhan Gong, Yuan Zhong, Wenao Ma, Jinpeng Li, Zhao Wang, Jingyang Zhang, Pheng-Ann Heng, and Qi Dou. 2024. 3dsam-adapter: Holistic adaptation of sam from 2d to 3d for promptable tumor segmentation. *Medical Image Analysis*, 98:103324.

Daniel Hernández, editor. 1994. *Reasoning with qualitative representations*, pages 55–103. Springer Berlin Heidelberg, Berlin, Heidelberg.

Kaiyi Huang, Kaiyue Sun, Enze Xie, Zhenguo Li, and Xihui Liu. 2023. T2i-compbench: A comprehensive benchmark for open-world compositional text-to-image generation. *Preprint*, arXiv:2307.06350.

Arthur Juliani, Vincent-Pierre Berges, Ervin Teng, Andrew Cohen, Jonathan Harper, Chris Elion, Chris Goy, Yuan Gao, Hunter Henry, Marwan Mattar, and Danny Lange. 2020. Unity: A general platform for intelligent agents. *Preprint*, arXiv:1809.02627.

Parisa Kordjamshidi, Martijn Van Otterlo, and Marie-Francine Moens. 2010. Spatial role labeling: Task definition and annotation scheme. In *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC’10)*, Valletta, Malta. European Language Resources Association (ELRA).

Parisa Kordjamshidi, Martijn van Otterlo, and Marie-Francine Moens. 2011. Spatial role labeling: towards extraction of spatial relations from natural language. *ACM - Transactions on Speech and Language Processing*, 8:1–36.

Stephen C. Levinson. 2003. *Space in Language and Cognition: Explorations in Cognitive Diversity*. Language Culture and Cognition. Cambridge University Press.

Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. 2023. Gligen: Open-set grounded text-to-image generation. *CVPR*.

Fangyu Liu, Guy Emerson, and Nigel Collier. 2023. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 11:635–651.

Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, Jun Zhu, and Lei Zhang. 2025. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *Computer*

736	<i>Vision – ECCV 2024</i> , pages 38–55, Cham. Springer	Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhi-	788
737	Nature Switzerland.	hao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin	789
738	Llama. 2024. The llama 3 herd of models . <i>Preprint</i> ,	Wang, Wenbin Ge, Yang Fan, Kai Dang, Mengfei	790
739	arXiv:2407.21783.	Du, Xuancheng Ren, Rui Men, Dayiheng Liu,	791
740	Roshanak Mirzaee and Parisa Kordjamshidi. 2022.	Chang Zhou, Jingren Zhou, and Junyang Lin. 2024.	792
741	Transfer learning with synthetic corpora for spatial	Qwen2-vl: Enhancing vision-language model’s per-	793
742	role labeling and reasoning . In <i>Proceedings of the</i>	ception of the world at any resolution . <i>Preprint</i> ,	794
743	<i>2022 Conference on Empirical Methods in Natural</i>	arXiv:2409.12191.	795
744	<i>Language Processing</i> , pages 6148–6165, Abu	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten	796
745	Dhabi, United Arab Emirates. Association for Com-	Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and	797
746	putational Linguistics.	Denny Zhou. 2023. Chain-of-thought prompting elic-	798
747	OpenAI. 2024. Gpt-4 technical report . <i>Preprint</i> ,	its reasoning in large language models . <i>Preprint</i> ,	799
748	arXiv:2303.08774.	arXiv:2201.11903.	800
749	Alibaba Group Qwen Team. 2024. Qwen2 technical	Yutaro Yamada, Yihan Bao, Andrew Kyle Lampinen,	801
750	report . <i>Preprint</i> , arXiv:2407.10671.	Jungo Kasai, and Ilker Yildirim. 2024. Evaluat-	802
751	René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun.	ing spatial understanding of large language models .	803
752	2021. Vision transformers for dense prediction . In	<i>Transactions on Machine Learning Research</i> .	804
753	<i>2021 IEEE/CVF International Conference on Com-</i>	Jianing Yang, Xuweiyi Chen, Shengyi Qian, Nikhil	805
754	<i>puter Vision (ICCV)</i> , pages 12159–12168.	Madaan, Madhavan Iyengar, David F. Fouhey, and	806
755	Md Imbesat Rizvi, Xiaodan Zhu, and Iryna Gurevych.	Joyce Chai. 2024. Llm-grounder: Open-vocabulary	807
756	2024. SpaRC and SpaRP: Spatial reasoning char-	3d visual grounding with large language model as an	808
757	acterization and path generation for understanding	agent . In <i>2024 IEEE International Conference on</i>	809
758	spatial reasoning capability of large language models .	<i>Robotics and Automation (ICRA)</i> , pages 7694–7701.	810
759	In <i>Proceedings of the 62nd Annual Meeting of the</i>	Yue Zhang and Parisa Kordjamshidi. 2022. LOViS:	811
760	<i>Association for Computational Linguistics (Volume 1:</i>	Learning orientation and visual signals for vision	812
761	<i>Long Papers)</i> , pages 4750–4767, Bangkok, Thailand.	and language navigation . In <i>Proceedings of the 29th</i>	813
762	Association for Computational Linguistics.	<i>International Conference on Computational Linguis-</i>	814
763	Robin Rombach, A. Blattmann, Dominik Lorenz,	<i>tics</i> , pages 5745–5754, Gyeongju, Republic of Korea.	815
764	Patrick Esser, and Björn Ommer. 2021. High-	International Committee on Computational Linguis-	816
765	resolution image synthesis with latent diffusion mod-	<i>tics</i> .	817
766	els . <i>2022 IEEE/CVF Conference on Computer Vi-</i>	Zheyuan Zhang, Fengyuan Hu, Jayjun Lee, Freda Shi,	818
767	<i>sion and Pattern Recognition (CVPR)</i> , pages 10674–	Parisa Kordjamshidi, Joyce Chai, and Ziqiao Ma.	819
768	10685.	2025. Do vision-language models represent space	820
769	Francesco Ruotolo, Tina Iachini, Gennaro Ruggiero,	and how? evaluating spatial frame of reference under	821
770	Ineke J. M. van der Ham, and Albert Postma. 2016.	ambiguities . In <i>The Thirteenth International Confer-</i>	822
771	Frames of reference and categorical/coordinate spa-	<i>ence on Learning Representations</i> .	823
772	tial relations in a “what was where” task . <i>Experimen-</i>		
773	<i>tal Brain Research</i> , 234(9):2687–2696.		
774	Zhengxiang Shi, Qiang Zhang, and Aldo Lipani. 2022.		
775	Stepgame: A new benchmark for robust multi-hop		
776	spatial reasoning in texts . In <i>Proceedings of the AAAI</i>		
777	<i>Conference on Artificial Intelligence</i> , volume 36,		
778	pages 11321–11329.		
779	Anna Shusterman and Peggy Li. 2016. Frames of ref-		
780	erence in spatial language acquisition . <i>Cognitive</i>		
781	<i>Psychology</i> , 88:115–161.		
782	Thora Tenbrink. 2011. Reference frames of space and		
783	time in language . <i>Journal of Pragmatics</i> , 43(3):704–		
784	722. <i>The Language of Space and Time</i> .		
785	Nikola Vukovic and John N. Williams. 2015. Individ-		
786	ual differences in spatial cognition influence mental		
787	simulation of language . <i>Cognition</i> , 142:110–122.		

Case	Context	QA-CP	QA-RP	T2I
Cow Case	792	3168	3168	3168
Car Case	128	512	512	512
Box Case	120	120	120	120
Pen Case	488	488	488	488
Total	1528	4288	4288	4288

Table 4: Dataset Statistic of FoREST A-split dataset. QA-CP and QA-RP are question-answering with camera perspective and question-answering with relatum perspective, respectively. T2I refers to prompt used in Text-to-Image experiments.

FoR Class	Context	QA-CP	QA-RP	T2I
External Relative	1528	4288	3680	4288
External Intrinsic	920	3680	3680	3680
Internal Intrinsic	128	512	512	0
Internal Relative	248	632	512	0
Total	2824	9112	8384	7968

Table 5: Dataset Statistic of FoREST C-split dataset. QA-CP and QA-RP are question-answering with camera perspective and question-answering with relatum perspective, respectively. T2I refers to prompt used in Text-to-Image experiments.

A Additional details of FoREST Dataset

A.1 Dataset Statistics

The FoREST dataset statistic is provided in the Table 4 for A-split and Table 5 for C-split.

A.2 Details Creation of FoREST dataset

We define the nine categories of objects selected in our dataset as indicated below in Table 6. We select sets of locatum and relatum based on the properties of each class to cover four cases of frame of reference defined in Section 3.2. Notice that we also consider the appropriateness of the container; for example, the car should not contain the bus.

Based on the selected locatum and relatum. To create an A-split spatial expression, we substitute the actual locatum and relatum objects in the Spatial Relation template. After obtaining the A-split contexts, we create their counterparts using the perspective/topology clauses to make the counterparts in the C-split. Then, we obtain the I-A and I-C split by applying the directional template to the first occurrence of relatum when it has intrinsic directions. The directional templates are "that is facing towards," "that is facing backward," "that is facing to the left," and "that is facing to the right." All the templates are in the Table 7. We then construct the scene configuration from each modified spatial expression and send it to the simulator developed using Unity3D. Eventually, the simulator

produces four visualization images for each scene configuration.

A.2.1 Simulation Details

The simulation starts with randomly placing the relatum into the scene with the orientation based on the given scene configuration. We randomly select the orientation by given scene configuration, $[-40, 40]$ for front, $[40, 140]$ for left, $[140, 220]$ for back, and $[220, 320]$ for right. Then, we create the locatum from the relatum position and move it in the spatial relation provided. If the frame of reference is relative, we move the locatum based on the camera’s orientation. Otherwise, we move it from the relatum’s orientation. Then, we check the camera’s visibility of both objects. If one of them is not visible, we repeat the process of generating the relatum until the correct placement is achieved. After getting the proper placement, we randomly choose the background from 6 backgrounds. Eventually, we repeat the procedures four times for one configuration.

A.2.2 Object Models and Background

For the object models and background, we find it from the unity asset store². All of them are free and available for download. All of the 3D models used are shown in Figure 4.

A.2.3 Textual templates

All manually created templates used to create FoREST are given in Table 7.

A.2.4 GPT template

We prompt GPT with an initial question using a specific template (the first template in Table 8) to generate four additional variations. Each generated template was manually reviewed and validated before being added to the set of questions in our corpus. All templates are listed in Table 8. Question templates for asking the spatial relation from relatum view are edited by replacing references from the camera to the relatum object.

B Comparison between Qwen2 and Qwen2-VL in QA setting

According to the A-split results in Table 1, Qwen2VL demonstrates a more balanced preference pattern and reasonable performance compared to Qwen2. This suggests that, unlike Qwen2, which often extracts spatial relations without accounting

²<https://assetstore.unity.com>

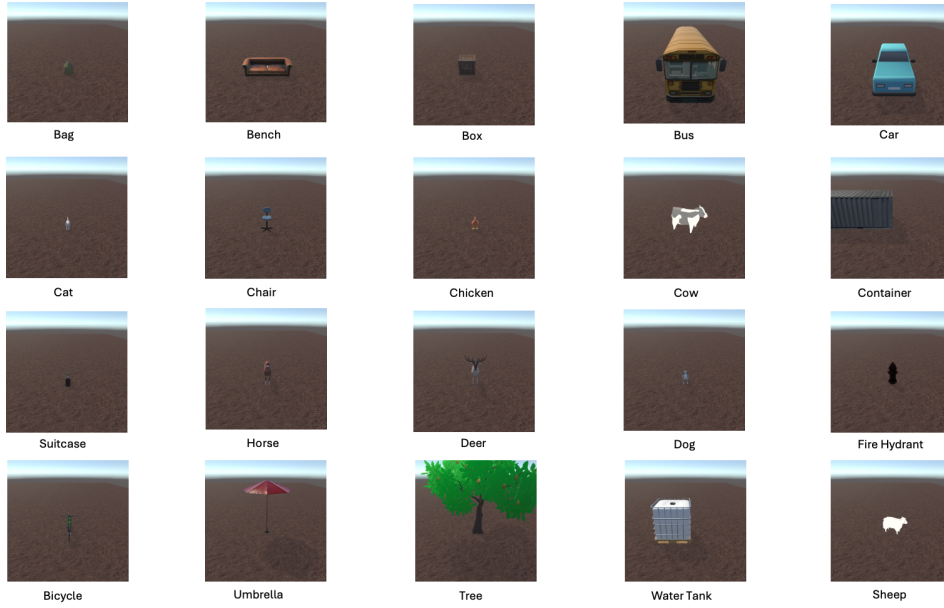


Figure 4: All 3d models used to generate visualizations for FoREST.

Category	Object(s)	Intrinsic Direction	Possible Container
Small object without intrinsic directions	umbrella, bag, suitcase, fire hydrant	✗	✗
Big object with intrinsic directions	bench, chair	✓	✗
Big object without intrinsic direction	water tank	✗	✗
Container	box, container	✗	✓
Small animal	chicken, dog, cat	✓	✗
Big animal	deer, horse, cow, sheep	✓	✗
Small vehicle	bicycle	✓	✗
Big vehicle	bus, car	✓	✓
Tree	tree	✗	✗

Table 6: All selected objects with two properties: intrinsic direction, affordance of being container

Qwen2					Qwen2VL								
True Label	Left	0.21	0.37	0.16	0.26	0-shot		Left	0.3	0.37	0.14	0.19	
	Right	0.38	0.22	0.24	0.16			Left	0.36	0.28	0.18	0.18	
	Front	0.15	0.18	0.51	0.17			Left	0.072	0.066	0.64	0.22	
	Back	0.21	0.19	0.14	0.47			Left	0.067	0.086	0.27	0.58	
True Label	Left	0.28	0.46	0.061	0.19	SG + CoT		Left	0.27	0.44	0.13	0.16	
	Right	0.42	0.28	0.19	0.1			Left	0.44	0.27	0.16	0.13	
	Front	0.078	0.16	0.63	0.13			Left	0.14	0.16	0.49	0.21	
	Back	0.12	0.065	0.2	0.62			Left	0.13	0.14	0.22	0.51	
		Left	Right	Front	Back			Left	Right	Front	Back		
		Predicted Label						Predicted Label					

Figure 5: Confusion matrices of spatial relation answers when Qwen2 and Qwen2-VL must adapt FoR in the 0-shot and (SG+CoT) settings.

for perspective, Qwen2VL appears capable of reasoning about perspective shifts as well. This observation is further supported by the C-split results in Table 2, where Qwen2VL significantly outperforms Qwen2 in cases requiring adapting the spatial relations mentioned between different perspectives in most settings. These results support our hypothesis that visual training enhances the model’s ability to interpret perspective changes.

However, despite these improvements, Qwen2VL still struggles to provide reasoning from the relatum’s perspective on questions in CoT and SG+CoT settings. This suggests that while the model leverages visual training to reason about perspective and perform perspective shifts when necessary, its reasoning is more accurate when questions are framed from the camera’s viewpoint. However, it struggles with questions asked from the relatum’s perspective. This limitation is expected, as training data—particularly image captioning datasets—are typically annotated from

Spatial Relation Templates	{locatum} is in front of {relatum} {locatum} is on the left of {relatum} {locatum} is to the left of {relatum} {locatum} is behind of {relatum} {locatum} is back of {relatum} {locatum} is on the right of {relatum} {locatum} is to the right of {relatum}
Topology Templates	within {relatum} and inside {relatum} and outside of {relatum}
Perspective Templates	from {relatum}'s view relative to {relatum} from {relatum}'s perspective from my perspective from my point of view relative to observer
Orientation Templates	{relatum} facing toward that camera {relatum} is facing away from the camera. {relatum} facing left relative to the camera {relatum} facing right relative to the camera

Table 7: All templates used to create FoREST dataset.

Question Templates	In the camera view, how is {locatum} positioned in relation to {relatum}? Based on the camera perspective, where is the {locatum} from the {relatum}'s position? From the camera perspective, what is the relation of the {locatum} to the {relatum}? Looking through the camera perspective, how does {locatum} appear to be oriented relative to {relatum}'s position? Based on the camera angle, where is {locatum} located with respect to {relatum}'s location?
--------------------	--

Table 8: All templates used to create FoREST dataset.

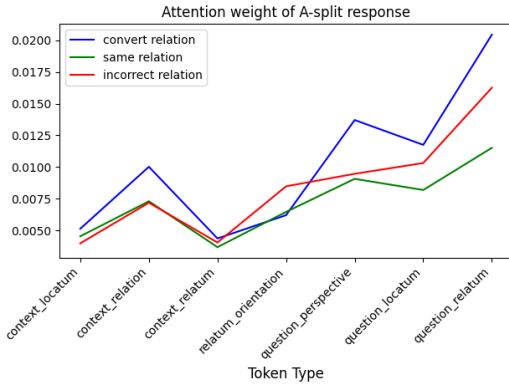


Figure 6: Attention weights of input context for Question Answering in A-split. Considering the following semantic concepts: locatum mention in context, relation in context, relatum in context, relatum orientation, perspective in question, locatum in question, and relatum in question. Three patterns are considered: (1) incorrect answers, (2) correct answers assuming a shared perspective between context and question, and (3) correct answers assuming different perspectives.

a human, camera-based perspective, restricting the model’s ability to generalize across different viewpoints.

C Attention Analysis

C.1 Experiment Setting

We conducted an attention-based interpretability analysis using Qwen2 in a 0-shot setting, where only the context and question were provided. We focused on analyzing attention weights across key

semantic elements in both A-split and C-split. We analyzed the attention weights for various groups of tokens that convey the following concepts: locatum mention occurring in the context, relation in context, relatum in context, perspective in context, relatum orientation, perspective in question, locatum in question, and relatum occurring in the question.

We categorized response patterns into three scenarios: (1) the model answers incorrectly, (2) the model answers correctly with the assumption that the same perspective holds for context and question, and (3) the model answers correctly with the assumption that context and question have different perspectives. Note that we separate the analysis for both splits because the conceptual tokens are different, that is, the perspective mentioned in the context is only for the C-split.

C.2 Results

A-split. According to Figure 6, we found that, for the correct responses, the model paid the least attention to the question’s perspective in the *same perspective*. Meanwhile, for the *different perspective* case, attention significantly increased to tokens related to the relation, the question perspective, and the relatum in the question. This result illustrates that these tokens played an important role in adapting spatial relations for perspective shifts. However, when the model increased attention weight on relatum orientation compared to other tokens, more

errors occurred in answering the questions.

C-split. According to Figure 7, in cases that require a shift in perspective, we observed that the model pays significantly more attention to the *spatial relation in context*, the *question perspective*, and the *relatum in the question* concepts necessary for answering correctly. However, the model largely overlooks the orientation tokens, which also play a crucial role in resolving these questions. This lack of attention to orientation may majorly contribute to the model’s failures in perspective-shifting scenarios. In cases where the relation mentioned remains the same in both the context and the answer, we observed that the model, in incorrect predictions, tends to assign higher attention to the relatum in the question. In contrast, correct predictions are associated with more balanced attention across tokens. Maintaining moderate attention to all tokens appears to help the model consistently respond with the correct relation, repeating it from the context.

D Quantitative Examples

D.1 Example of QA Bias in the A-split

To illustrate example biases in the A-split, we present three types of model behavior. The first, shown in Figure 8, occurs when the model assumes a shared perspective between context and question, allowing it to extract spatial relations directly without FoR reasoning. This scenario is the most common behavior of Qwen2. The second, illustrated in Figure 9, arises when the model assumes the context always reflects the relatum’s perspective, favoring intrinsic over relative interpretation. Lastly, Figure 10 shows the model assuming the context always adopts the camera perspective.

D.2 Incorrect QA reasoning in the C-split

Our results demonstrate that models struggle with converting between different perspectives when required to answer the questions, especially when the correct answer is either left or right. This difficulty contributes to their performance, as discussed in research question 2 in the paper (RQ2). To further illustrate LLMs’ struggle, we provide quantitative examples of failure cases, focusing specifically on the models’ most challenging instances in Figure 11.

E Template Variation in Question-Answering

We report average accuracy on the C-split across all

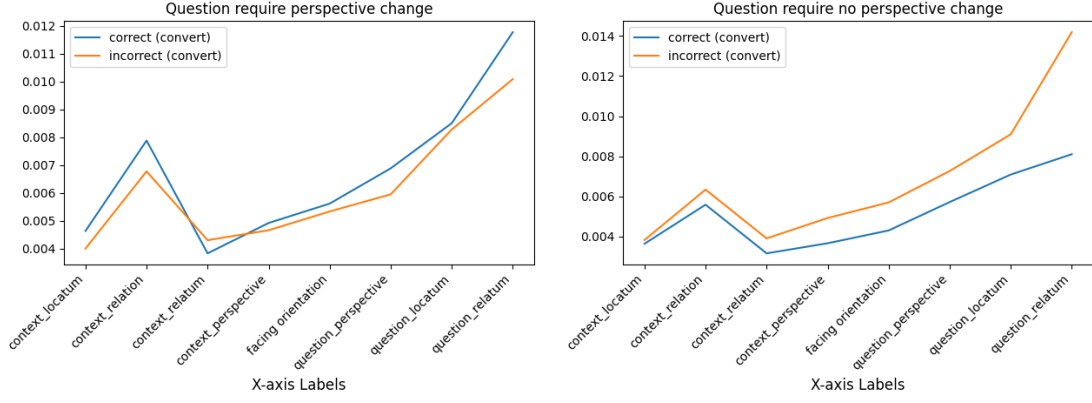
question templates for Qwen2 and GPT-4o in Table 9. While results show notable variation across templates, no single template consistently outperforms others across all tasks. Template effectiveness depends on both the prompting technique and the underlying model. Simpler templates generally yield higher accuracy, though more complex expressions may reduce overall performance. Nevertheless, the comparative trend remains consistent, as SG+CoT often outperforms CoT in overall improvement across both models. Notably, SG prompting yields the smallest performance variation across templates.

F VISOR-uncond Score

VISOR_{uncond} provides the overall spatial relation score, including images with object generation errors. Since it is less focused on evaluating spatial interpretation than VISOR_{cond}, which assesses explicitly the text-to-image model’s spatial reasoning, we report VISOR_{uncond} results here in the Table 10 rather than in the main paper. The results are similar to the pattern observed in VISOR_{uncond} that the based models(SD-1.5 and SD-2.1) perform better in the relative frame of reference, while the layout-to image models, i.e. GLIGEN, are better in the intrinsic frame of reference.

G Analyze the improvements in SG-prompting for T2I.

To further explain improvements of SG-prompting in T2I task, we assess the generated bounding boxes in the I-C split for left and right relations relative to the camera since these can be evaluated using only bounding boxes without depth information. As seen in Table 11, our SG prompting improved Llama3-70B’s by 3.48%, while Llama3-8B saw a slight decrease of 0.22%. This evaluation was conducted on all generated layouts from the I-C split, which differs from the evaluation subset of images used for VISOR_{cond} in Table 3. We report the layout_{cond} score for a consistent evaluation in the same table. Layout_{cond} shows that Llama3-8B improves within the same evaluation subset with VISOR_{cond}. Overall, by incorporating FoR information through SG layout diffusion, Llama3 generates better spatial configurations, enhancing image generation performance.



(a) Questions that require changing perspective.

(b) Questions that do not require changing perspective.

Figure 7: Attention weights of input context for Question Answering in C-split. Considering the following semantic concepts: locatum mention in context, relation in context, relatum in context, perspective in context, relatum orientation, perspective in question, locatum in question, and relatum in question. Two patterns are considered: (1) incorrect answers, (2) correct answers. We separate the questions that require and not-require perspective changing.

Model	T0	T1	T2	T3	T4	Avg.	Highest Δ
Qwen2 (0-shot)	70.23	61.36	69.35	65.55	67.97	66.91	8.87
Qwen2 (4-shot)	65.15	64.01	68.80	63.24	65.46	65.34	5.57
Qwen2 (CoT)	71.26	66.11	72.78	63.40	71.87	69.10	9.37
Qwen2 (SG + CoT)	71.15	70.92	72.89	70.11	71.69	71.36	2.78
GPT-4o (0-shot)	64.02	61.14	73.16	65.77	56.66	64.19	16.50
GPT-4o (4-shot)	65.85	52.23	70.55	61.64	58.33	61.78	18.31
GPT-4o (CoT)	80.98	69.65	84.81	83.24	74.98	78.78	15.16
GPT-4o (SG + CoT)	85.74	85.02	84.10	86.53	87.57	85.79	3.47

Table 9: Average accuracy across different templates. T_i refers to the generated QA template i in Table 7. Δ measures the difference in accuracy between two different templates.

Model	VISOR(%)					
	uncond (I)	uncond (R)	uncond (avg)	uncond (I)	uncond (R)	uncond (avg)
	I-A-Split			I-C-Split		
SD-1.5	45.43	33.22	43.51	35.06	35.68	35.40
SD-2.1	62.87	43.90	59.89	45.98	46.59	46.31
Llama3-8B + GLIGEN	46.74	38.16	45.39	33.98	39.36	36.89
Llama3-70B + GLIGEN	54.33	46.89	53.17	38.04	46.04	42.37
Llama3-8B + SG + GLIGEN (Our)	51.83	43.24	50.48	36.28	44.43	40.70
Llama3-70B + SG + GLIGEN (Our)	58.92	47.44	57.12	38.23	48.62	43.86

Table 10: $VISOR_{uncond}$ score on the I-A-Split and I-C-Split where I refer to the Cow Case and Car Case where relatum has intrinsic directions, and R refer to the Box Case and Pen case where relatum lacks intrinsic directions, avg is mirco-average of I and R . cond and uncond are explained in Section 5.1.

Context: A suitcase is on the left of a dog. The dog is facing toward the camera.	
Question: Based on the dog angle , where is suitcase located with respect to dog's location?	Question: Based on the camera perspective , where is the suitcase from the dog's position?
Answer: Left	Answer: Left

Figure 8: The first example in A-split of FoREST, where the model correctly answers both perspectives. In this example, the model responds with the spatial relation in the context, assuming all questions and context have the same perspective.

Context: A chicken is in front of a cat. The cat is facing away from the camera	
Question: Based on the cat view , where is chicken located with respect to cat's location?	Question: Based on the camera perspective , where is the chicken from the cat's position?
Answer: Front	Answer: Back

Figure 9: The second example in A-split of FoREST, where the model correctly answers both perspectives. In this example, the model assumes the context has a relatum perspective.

H Frame of Reference Identification

We evaluate the LLMs' performance in recognizing the FoR classes from given spatial expressions. The LLMs receive spatial expression, denoted as T , and output one FoR class, FoR , from the valid set of

FoR classes, $FoR \in \{\text{external relative, external intrinsic, internal intrinsic, internal relative}\}$. All in-context learning examples are in the Appendix K.

H.1 Experimental Setting

Zero-shot model. We follow the regular setting of *zero-shot* prompting. We only provide instruc-

Context: A bench is to the right of a sheep. The sheep is facing toward the camera.	
Question: Looking through the sheep's perspective , how does bench appear to be oriented relative to sheep's position?	Question: Based on the camera perspective , where is the bench from the sheep's position?
Answer: Left	Answer: Right

Figure 10: The third example in A-split of FoREST, where the model correctly answers both perspectives. In this example, the model assumes the context has a camera perspective.

Context: A bag is in front of a sheep from my point of view. The sheep is facing left relative to me.	
Topology: The bag is not contained by the sheep.	
Distance: There is some distance between the bag and the sheep.	
Direction: From the observer's perspective, the bag is located in front of the sheep.	
Frame of Reference: external relative.	
Question: From the sheep's perspective, what is the relation of the bag to the sheep?	
Explanation: Based on the frame of reference (external relative), the front relation is created from the observer's perspective. However, the question asks for the relation from the sheep's perspective. Since the sheep is facing left relative to the observer, the bag being in front of the sheep from the observer's perspective means it is to the right of the sheep from the sheep's perspective. Answer: right	
Answer: Right ❌	GT-Answer: Left

Figure 11: An example of an incorrect SG+CoT case, where the model generates incorrect reasoning and mistakenly converts left to right.

tion to LLM with spatial context. The instruction prompt briefly explains each class of the FoR and candidate answers for the LLM. We called the LLM with the instruction prompt and T to find F .

Few-shot model. We manually craft four spatial expressions for each FoR class. To avoid creating bias, each spatial expression is ensured to fit in only one FoR class. These expressions serve as examples of our *few-shot* setting. We provide these examples in addition to the instruction as a part of the prompt, followed by T and query F from the LLM.

Chain-of-Thought (CoT) model. To create CoT (Wei et al., 2023) examples, we modify the prompt to require reasoning before answering. Then, we manually crafted reasoning explanations with the necessary information for each example used in few-shot. Finally, we call the LLMs, adding modified instructions to updated examples, followed by T and query F .

Spatial-Guided Prompting (SG) model. We hypothesize that the general spatial relation types defined in Section 2 can provide meaningful information for recognizing FoR classes. For instance, a topological relation, such as "inside," is intuitively

Model	Layout	Layout _{cond}
Llama3-8B	85.26	88.84
Llama3-8B + SG	85.04	88.86
Llama3-70B	88.47	93.16
Llama3-70B + SG	91.95	95.45

Table 11: Layout accuracy where spatial relations are left or right relative to the camera. Layout is evaluated for all generated layouts in I-C split while Layout_{cond} uses the same testing examples as VISOR_{cond}.

Model	inherently clear		require template	
	CoT	SG	CoT	SG
Llama3-70B	19.84	44.64 (↑ 24.80)	76.72	87.39 (↑ 10.67)
Qwen2-72B	58.20	84.22 (↑ 26.02)	88.36	93.86 (↑ 10.67)
GPT-4o	12.50	29.17 (↑ 16.67)	87.73	90.74 (↑ 3.01)

Table 12: The comparison between CoT and SG prompting in C-split separated by inherently clear / required template to be clear.

associated with an internal FoR. Therefore, we propose Spatial-Guided Prompting to direct the model in identifying the type of relations before querying F . We revise the prompting instruction to guide the model in considering these three aspects. Then, we manually explain these three aspects. We specify the relation's origin from the context for direction relations, such as "the left direction is relative to the observer." We hypothesize that this information helps the model distinguish between intrinsic and relative FoR. Next, we specify whether the locatum is inside or outside the relatum for topological relations. This information should help distinguish between internal and external FoR classes. Lastly, we provide the potential quantitative distance, e.g., far. This quantitative distance further encourages identifying the correct topological and directional relations. Eventually, we insert these new explanations in examples and call the model with the updated instructions followed by T to query F .

H.2 Evaluation Metrics

We report the accuracy of the model on the multi-class classification task. Note that the expressions in A-split can have multiple correct answers. Therefore, we consider the prediction correct when it is in one of the valid FoR classes for the given spatial expression.

H.3 Results

H.3.1 FoR Inherently Bias in LLMs

C-split. The *zero-shot* setting reflects the LLMs' inherent bias in identifying FoR. Table 13 presents the accuracy for each FoR class in C-split, where sentences explicitly include information about

Model	A-split	C-Split				
		ER-C-Split	EC-Split	IC-Split	IR-C-Split	Avg.
Gemma2-9B (0-shot)	94.17	94.24	35.98	53.91	57.66	60.45
Gemma2-9B (4-shot)	59.58	55.89(↓ 38.34)	72.61(↑ 36.63)	74.22(↑ 20.31)	54.44(↓ 3.23)	64.29(↑ 3.84)
Gemma2-9B (CoT)	60.49	60.49(↓ 33.74)	60.54(↑ 24.57)	87.50(↑ 33.59)	54.03(↓ 3.63)	65.64(↑ 5.20)
Gemma2-9B (SG)(Our)	72.67	65.87(↓ 28.37)	65.54(↑ 29.57)	53.12(↓ 0.78)	95.97(↑ 38.31)	70.13(↑ 9.68)
Llama3-8B (0-shot)	60.21	32.20	90.11	75.78	0.00	49.52
Llama3-8B (4-shot)	60.14	47.77(↑ 15.58)	54.35(↓ 35.76)	100.00(↑ 24.22)	41.13(↑ 41.13)	60.81(↑ 11.29)
Llama3-8B (CoT)	61.32	61.06(↑ 28.86)	97.28(↑ 7.17)	100.00(↑ 24.22)	36.29(↑ 36.29)	73.66(↑ 24.14)
Llama3-8B (SG) (Our)	62.95	63.29(↑ 31.09)	94.57(↑ 4.46)	100.00(↑ 24.22)	43.55(↑ 43.55)	75.35(↑ 25.83)
Llama3-70B (0-shot)	84.23	74.08	9.57	92.19	68.55	61.10
Llama3-70B (4-shot)	78.47	81.81(↑ 7.72)	64.89(↑ 55.33)	100.00(↑ 7.81)	75.81(↑ 7.26)	80.63(↑ 19.53)
Llama3-70B (CoT)	69.11	72.05(↓ 2.03)	97.07(↑ 87.50)	100.00(↑ 7.81)	79.44(↑ 10.89)	87.14(↑ 26.04)
Llama3-70B (SG) (Our)	76.50	78.21(↑ 4.12)	97.61(↑ 88.04)	100.00(↑ 7.81)	72.18(↑ 3.63)	87.00(↑ 25.90)
Qwen2-7B (0-shot)	83.64	79.97	59.24	77.34	40.73	64.32
Qwen2-7B (4-shot)	61.12	50.52(↓ 29.45)	65.76(↑ 6.52)	93.75(↑ 16.41)	56.05(↑ 15.32)	66.52(↑ 2.20)
Qwen2-7B (CoT)	72.12	70.81(↓ 9.16)	63.80(↓ 4.57)	99.22(↑ 21.88)	51.61(↑ 10.89)	71.36(↑ 7.04)
Qwen2-7B (SG)	70.61	68.00(↓ 11.98)	71.20(↑ 11.96)	88.28(↑ 10.94)	57.26(↑ 16.53)	71.18(↑ 6.86)
Qwen2-72B (0-shot)	64.46	62.70	100.00	100.00	39.11	75.45
Qwen2-72B (4-shot)	79.12	78.73(↑ 16.03)	99.35(↓ 0.65)	87.50(↓ 12.50)	87.10(↑ 47.98)	88.17(↑ 12.72)
Qwen2-72B (CoT)	88.54	88.87(↑ 26.18)	89.57(↓ 10.43)	93.75(↓ 6.25)	83.47(↑ 44.35)	88.91(↑ 13.46)
Qwen2-72B (SG)	90.51	90.18(↑ 27.49)	93.26(↓ 6.74)	98.44(↓ 1.56)	85.08(↑ 45.97)	91.74(↑ 16.29)
GPT3.5 (0-shot)	83.11	88.15	17.50	70.31	41.13	54.27
GPT3.5 (4-shot)	61.25	48.95(↓ 39.20)	62.72(↑ 45.22)	100.00(↑ 29.69)	28.63(↓ 12.50)	60.07(↑ 5.80)
GPT3.5 (CoT)	66.55	66.62(↓ 21.53)	96.85(↑ 79.35)	100.00(↑ 29.69)	50.81(↑ 9.68)	78.57(↑ 24.30)
GPT3.5 (SG) (Our)	70.61	73.30(↓ 14.86)	92.93(↑ 75.43)	99.22(↑ 28.91)	49.19(↑ 8.06)	78.66(↑ 24.39)
GPT4o (0-shot)	73.82	71.27	98.80	100.00	70.56	85.16
GPT4o (4-shot)	66.23	67.87(↓ 3.40)	98.70(↓ 0.11)	100.00(↑ 0.00)	78.63(↑ 8.06)	86.30(↑ 1.14)
GPT4o (CoT)	72.44	72.77(↑ 1.51)	100.00(↑ 1.20)	100.00(↑ 0.00)	73.79(↑ 3.23)	86.64(↑ 1.48)
GPT4o (SG) (Our)	76.44	74.67(↑ 3.40)	97.72(↓ 1.09)	100.00(↑ 0.00)	68.55(↓ 2.02)	85.23(↑ 0.08)

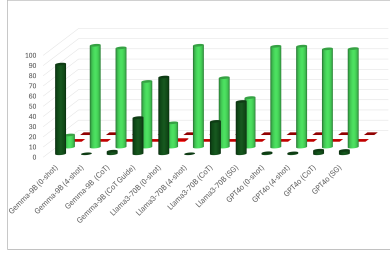
Table 13: Accuracy results report from FoR Identification with LLMs. The correct prediction is one of the valid FoR classes for the given spatial expression. All FoR classes are external relative (ER), external intrinsic (EI), internal intrinsic (II), and internal relative (IR).

topology and perspectives. We found that some models strongly prefer specific FoR classes. Notably, Gemma2-9B achieves a near-perfect accuracy on external relative FoR but performs poorly on other classes, especially external intrinsic, indicating a notable bias towards external relative. In contrast, GPT4o and Qwen2-72B perform exceptionally in both intrinsic FoR classes. However, they perform poorly in the relative FoRs. We also

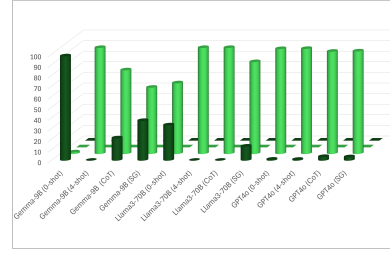
A-split. We examine the FoR bias in the A-split. Based on the results in Table 13, we plotted the top-3 models’ results (Gemma2-9B, Llama3-70B, and GPT4o) for a more precise analysis in Figures 12. The plots show the frequencies of each FoR category. According to the plot, Gemma and GPT have strong biases toward external relative and external intrinsic, respectively. This bias helps Gemma2 perform well in the A-split since all spatial expressions can be interpreted as external relative. However, GPT4o’s bias leads to errors when intrinsic FoRs aren’t valid, as in the Box and Pen cases (see plots (c) and (d)). Llama3 exhibits different behavior, showing a bias based on the relatum’s properties, specifically the relatum’s affordance as a container. In cases where relatum cannot serve as containers, i.e., Cow and Pen cases, Llama3 favors external relative. Conversely, Llama3 tends to favor external intrinsic when the relatum has the potential to be a container.

H.3.2 Behavior with ICL variations

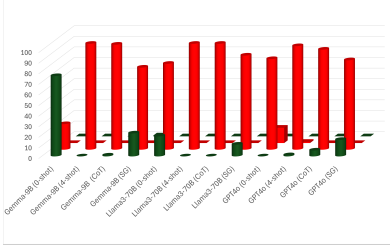
C-split. We evaluate the models’ behavior under various in-context learning (ICL) methods. As observed in Table 13, the *few-shot* method improves the performance of the *zero-shot* method across multiple LLMs by reducing their original bias toward specific classes. Reducing the bias, however, lowers the performance in some cases, such as the performance of Gemma 2 in ER class. One noteworthy observation is that while the *CoT* prompting generally improves performance in larger LLMs, it is counterproductive in smaller models for some FoR classes. This suggests that the smaller models have difficulty inferring FoR from the longer context. This negative effect also appears in SG prompting, which uses longer explanations. Despite performance degradation in particular classes of small models, SG prompting performs exceptionally well across various models and achieves outstanding performance with Qwen2-72B. We further investigate the performance of CoT and SG prompting. As shown in Table 12, CoT exhibits a substantial difference in performance between contexts with inherently clear FoR and contexts requiring the template to clarify FoR ambiguity. This implies that CoT heavily relies on the specific template to identify FoR classes. In contrast, SG prompting demonstrates a smaller gap between these two scenarios and significantly enhances per-



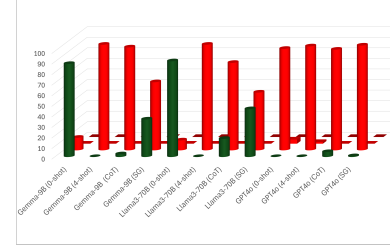
(a) Results of Cow Case in A-Split.



(b) Results of Car Case in A-Split.



(c) Results of Box Case in A-Split.



(d) Results of Pen Case in A-Split.

Figure 12: Red shows the wrong FoR identifications, and green shows the correct ones. The dark color is for relative FoRs, while the light color is for intrinsic FoRs. The round shape is for the external FoRs, while the square is for internal FoRs. The depth of the plots shows the four FoRs, i.e., *external relative*, *external intrinsic*, *internal intrinsic*, and *internal relative*, **from front to back**.

formance over CoT in inherently clear FoR contexts. Therefore, guiding the model to provide characteristics regarding topological, distance, and directional types of relations improves FoR comprehension.

A-split. We use the same Figure 12 to observe the behavior when applying ICL. The A-split shows minimal improvement with ICL variations, though some notable changes are observed. With *few-shot*, all models show a strong bias toward external intrinsic FoR, even when the relatum lacks intrinsic directions, i.e., Box and Pen cases. This bias appears even in Gemma2-9B, which usually behaves differently. This suggests that the models pick up biases from the examples despite efforts to avoid such patterns. However, *CoT* reduces some bias, leading LLMs to revisit relative, which is generally valid across scenarios. In Gemma2, the model predicts relative FoR where the relatum has intrinsic directions, i.e., Cow and Car cases. Llama3 behaves similarly in cases where the relatum cannot act as a container, i.e., Cow and Pen cases. GPT4o, however, does not depend on the relatum’s properties and shows slight improvements across all cases. Unlike *CoT*, our SG prompting is effective in all scenarios. It significantly reduces biases while following a similar pattern to *CoT*. Specifically, SG prompting increases external relative predictions for Car and Cow in Gemma2-9B, and for Cow and Pen in Llama3-70B. Nevertheless, GPT4o shows

only a slight bias reduction. However, Our proposed method improves the overall performance of most models, as shown in Table 13. The Llama3-70B behaviors are also seen in LLama3-8B and GPT3.5. The plots for LLama3-8B and GPT3.5 are in Figure 14.

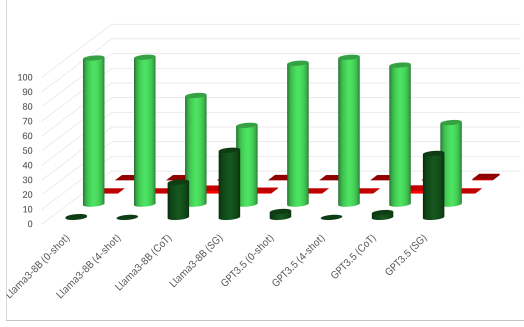
H.3.3 Experiment with different temperatures

We conducted additional experiments to further investigate the impact of temperature on the biased interpretation of the model in the A-split of our dataset. As presented in Table 14, comparing distinct temperatures (0 and 1) revealed a shift in the distribution. The frequencies of the classes experienced a change of up to 10%. However, the magnitude of this change is relatively minor, and the relative preferences for most categories remained unchanged. Specifically, the model exhibited the highest frequency responses for the cow, car, and pen cases, even with higher frequencies in certain settings. Consequently, a high temperature does not substantially alter the diversity of LLMs’ responses to this task, which is an intriguing finding.

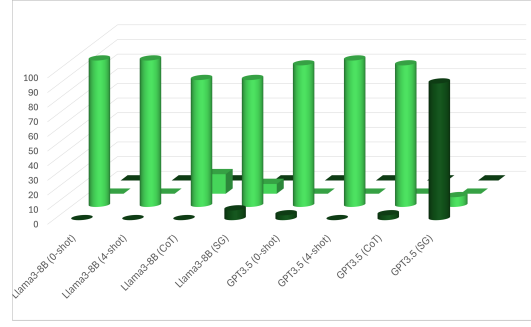
I Study of Human bias in A-split

I.1 Experiment Setting

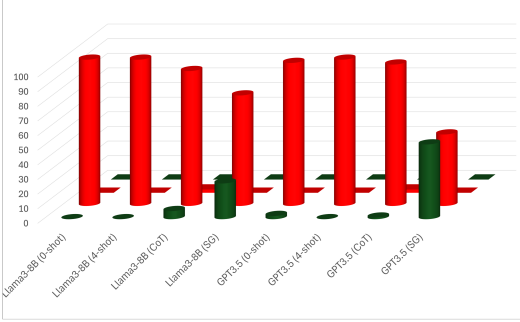
We provide a Google Form with instructions in Listing 1, followed by 150 randomly ordered context-question pairs. These are sampled based on rela-



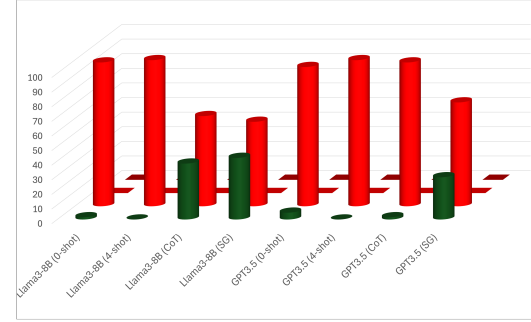
(a) Results of Cow Case in A-Split.



(b) Results of Car Case in A-Split.



(c) Results of Box Case in A-Split.



(d) Results of Pen Case in A-Split.

Figure 13: Red shows the wrong FoR identifications, and green shows the correct ones. The dark color is for relative FoRs, while the light color is for intrinsic FoRs. The round shape is for the external FoRs, while the square is for internal FoRs. The depth of the plots shows the four FoRs, i.e., external relative, external intrinsic, internal intrinsic, and internal relative, from front to back. This plot is the result of the rest of LLMs.

tum category: 25 questions per case from the A-split, covering four cases (Cow, Car, Box, Pen) for camera-perspective questions, and two cases (Cow, Car) for relatum-perspective questions, as Box and Pen lack intrinsic direction. No time constraints were imposed, and all participants were compensated at the standard research assistant rate. All participants were informed about the study details at the end and asked for consent to use their responses in the analysis. After acquiring the results, we utilize the same metric defined in Section 5.1. Then, we average the accuracy of each participant to derive the human baseline based on this study.

I.2 Results

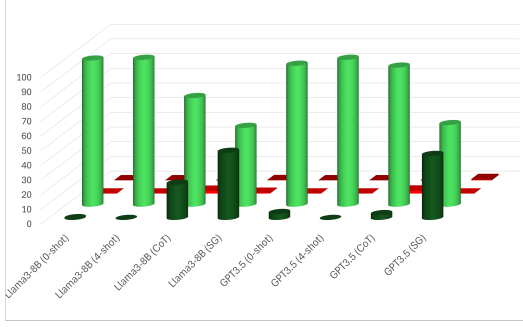
Individual human results are shown in Table 15. These results indicate that humans rely heavily on their assumptions when interpreting ambiguous frame-of-reference (FoR) contexts. Most participants favored the intrinsic FoR in A-split scenarios, though this preference varied across individuals, suggesting that background or prior visual experience may influence FoR interpretation. This aligns with findings from cognitive studies on how humans describe spatial relations that recently seen

examples may influence (Coventry et al., 2018). However, our results contrast with other studies suggesting that humans tend to adopt an egocentric (relative) FoR in scenes with fewer landmarks and an allocentric (intrinsic) FoR in those with more landmarks (Forster et al., 2023). Since our study does not provide visual input, we cannot infer how participants mentally reconstruct scenes to answer the questions with only two objects in the scene. The next notable point is that while Box and Pen cases appear straightforward for models (Table 1), humans may still struggle with questions in these categories. Humans do not exhibit the same limitations as LLMs in perspective-shifting tasks. When the context and question differ in perspective, humans still achieve over 90% accuracy, whereas LLMs often struggle with this type of reasoning.

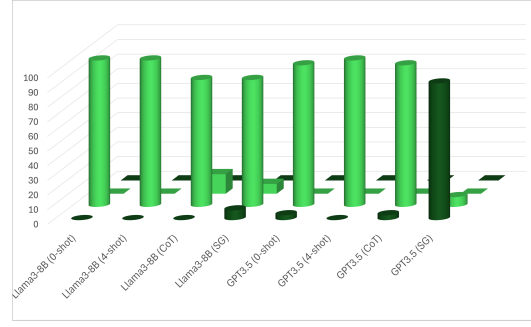
```

"""
Instruction:
You will be provided with a scene
description that describes the
spatial relationship between two
objects.
The scene description may include object
orientation, such as A is facing
toward the camera, to indicate the
direction A is facing.

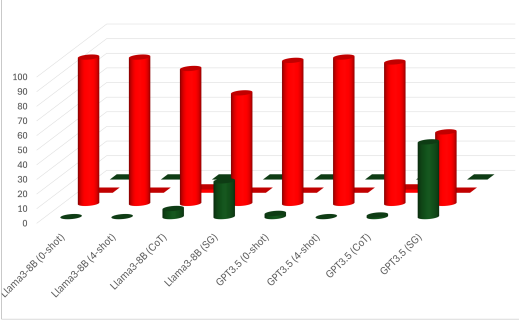
```



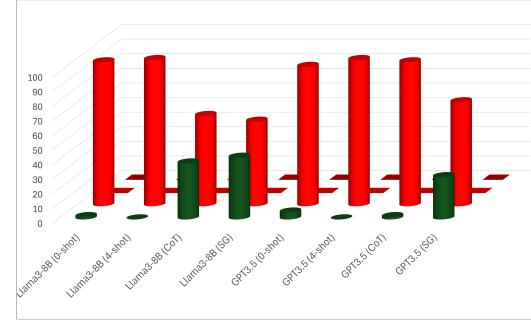
(a) Results of Cow Case in A-Split.



(b) Results of Car Case in A-Split.



(c) Results of Box Case in A-Split.



(d) Results of Pen Case in A-Split.

Figure 14: Red shows the wrong FoR identifications, and green shows the correct ones. The dark color is for relative FoRs, while the light color is for intrinsic FoRs. The round shape is for the external FoRs, while the square is for internal FoRs. The depth of the plots shows the four FoRs, i.e., external relative, external intrinsic, internal intrinsic, and internal relative, from front to back. This plot is the result of the rest of LLMs.

```

You will then receive a question asking
about the spatial relationship
between A and B from either the
camera's perspective or an object's
perspective.
Your task is to answer the question
based on your understanding of the
given spatial relationship.
" " "

```

Listing 1: Instruction for collecting human results on QA using A-split of FoREST dataset.

J Additional Results with GPT-o4-mini-high

We include results from GPT-4o-mini-high, a model configured for enhanced reasoning, for comparison with our previous results in Table 1 and Table 2. In this experiment, the reasoning-effort parameter is set to *high*. As shown in Table 16, the model exhibits a balanced preference between relative and intrinsic interpretations. Similar to Qwen2VL, it demonstrates strong reasoning ability, achieving high accuracy even when FoR adaptation is required due to its visual reasoning capability. Notably, GPT-o4-mini-high performs well on questions from the relatum perspective, even outper-

forms the GPT-4o (SG+CoT) in Table 17. However, its performance declines on camera-perspective questions. We believe this is possibly due to training emphasizing alternative perspectives over the camera view, which may reduce their performance from the camera perspective. Therefore, even with a strong reasoning model, the perspective shift remains challenging in the textual domain.

K LLM Prompt

K.1 FoR Identification

We provide the prompting for each in-context learning. The prompting for *zero-shot* and *few-shot* is provided in Listing 2. The instruction answer for these two in-context learning is “Answer only the category without any explanation. The answer should be in the form of {Answer: Category.}”

For the Chain of Thought (CoT), we only modified the instruction answer to “Answer only the category with an explanation. The answer should be in the form of {Explanation: Explanation Answer: Category.}” Similarly to CoT, we only modified the instruction answer to “Answer only the category with an explanation regarding topological, distance,

Model	ER		EI		II		IR	
	temp-0	temp-1	temp-0	temp-1	temp-0	temp-1	temp-0	temp-1
Cow Case								
0-shot	75.38	87.12	23.86	12.50	0.76	0.13	0.00	0.25
4-shot	0.00	15.66	100.00	84.34	0.00	0.00	0.00	0.00
CoT	31.82	49.87	68.18	49.87	0.00	0.13	0.00	0.13
SG	51.39	70.45	48.61	29.42	0.00	0.00	0.00	0.13
Box Case								
0-shot	22.50	41.67	77.50	58.33	0.00	0.13	0.00	0.25
4-shot	0.00	0.00	100.00	100.00	0.00	0.00	0.00	0.00
CoT	0.00	5.83	100.00	94.17	0.00	0.00	0.00	0.00
SG	11.67	33.33	88.33	66.67	0.00	0.00	0.00	0.00
Car Case								
0-shot	55.20	68.24	49.01	31.15	0.79	0.61	0.00	0.00
4-shot	0.60	5.94	99.40	94.06	0.00	0.00	0.00	0.00
CoT	19.64	38.52	80.16	61.27	0.20	0.20	0.00	0.00
SG	44.25	56.97	55.75	43.03	0.00	0.00	0.00	0.00
Pen Case								
0-shot	90.62	96.88	9.38	3.12	0.00	0.61	0.00	0.00
4-shot	0.00	7.03	100.00	92.97	0.00	0.00	0.00	0.00
CoT	17.19	28.91	82.81	71.09	0.20	0.20	0.00	0.00
SG	48.31	57.81	54.69	42.19	0.00	0.00	0.00	0.00

Table 14: The results between two different temperatures of Llam3-70B on the A-split of FoREST. The number shows the percentage frequency of responses from the model.

Model	Camera perspective									Relatum perspective						
	Cow			Car			Box	Pen	Avg.	Cow			Car			Avg.
	R%	I%	Acc.	R%	I%	Acc.	Acc.	Acc.	Acc.	R%	I%	Acc.	R%	I%	Acc.	Acc.
P #1	22.7	77.3	88.0	8.3	91.7	96.0	92.0	96.0	93.0	20.8	79.2	96.0	4.2	95.8	96.0	96.0
P #2	8.7	91.3	92.0	0.0	100.0	96.0	36.0	56.0	70.0	20.0	80.0	100.0	16.7	83.3	96.0	98.0
P #3	78.3	21.7	92.0	75.0	25.0	96.0	88.0	96.0	93.0	83.3	16.7	96.0	87.5	12.5	96.0	96.0
Human-baseline	36.6	63.4	90.7	27.8	72.2	96.0	72.0	82.7	85.3	41.4	58.6	97.3	36.1	63.9	96.0	96.7

Table 15: QA accuracy in the A-Split of human study. R% and I% represent the percentage the model assumes relative or intrinsic FoR for ambiguous expression, explained in Section 5.1. Acc is the accuracy, and Avg is the micro-average of accuracy. (1): 0-shot, (2): 4-shot, (3): CoT, and (4): SG+CoT.

and direction aspects. The answer should be in the form of {Explanation: Explanation Answer: Category.}', respectively. The example responses are provided in Listing 5 for Spatial Guided prompting.

```
# Instruction to find frame of reference
class of given context
"""
Instruction:
You specialize in language and spatial
relations, specifically in the frame
of context (multiple perspectives
in the spatial relation). Identify
the frame of reference category
given the following context. There
are four classes of the frame of
reference (external intrinsic,
internal intrinsic, external
relative, internal relative). Note
that the intrinsic direction refers
to whether the model has the front/
back by itself. (Example: a bird,
human. Counter Example: a ball, a
box). "

External intrinsic. The spatial
description of an entity A relative
to another entity B, where (1) A is
not contained by B, (2) the spatial
relation is based on B's facing
orientation (intrinsic direction) if
B has one.

Internal intrinsic. The spatial
```

description of an entity A relative
to another entity B, where (1) A is
contained by B, (2) the spatial
relation is based on B's facing
orientation (intrinsic direction) if
B has one.

External relative. The spatial
description of an entity A relative
to another entity B, where (1) A is
not contained by B, (2) The spatial
relation is relative to entity C,
where C is not involved in the
spatial relation and only provides
the perspective. An example of C is
the camera.

Internal relative. The spatial
description of an entity A relative
to another entity B, where (1) A is
contained by B, (2) The spatial
relation is relative to entity C,
where C is not involved in the
spatial relation and only provides
the perspective. An example of C is
the camera.

```
{Instruction answer}
# Normal Instruction answer: Find the
frame of reference category without
any explanations. The answer should
be in the form of Answer: Category.
# CoT Instruction answer: Find the frame
of reference category with an
explanation. The answer should be in
```

Model	Camera perspective									Relatum perspective						
	Cow			Car			Box	Pen	Avg.	Cow			Car			Avg.
	R%	I%	Acc.	R%	I%	Acc.	Acc.	Acc.	Acc.	R%	I%	Acc.	R%	I%	Acc.	Acc.
GPT-4o (1)	84.3	15.3	94.5	88.5	11.0	97.3	99.2	99.8	95.6	21.6	78.0	91.6	16.1	83.5	90.5	91.4
GPT-4o (2)	69.0	30.6	76.6	80.3	19.2	89.5	100.0	100.0	81.5	29.0	70.5	74.7	30.9	68.7	77.5	75.1
GPT-4o (3)	41.5	58.3	92.3	38.2	61.6	91.0	100.0	99.8	93.2	33.9	65.8	93.9	32.0	67.6	93.9	93.9
GPT-4o (4)	26.0	73.9	79.2	27.7	72.1	79.4	96.7	94.3	81.4	16.2	83.4	95.5	19.2	80.4	94.8	95.4
o4-mini-high	68.0	31.6	92.0	69.9	29.7	94.7	100.0	99.4	93.4	58.5	41.3	86.9	50.2	49.5	93.3	87.8
Human-baseline	36.6	63.4	90.7	27.8	72.2	96.0	72.0	82.7	85.3	41.4	58.6	97.3	36.1	63.9	96.0	96.7

Table 16: Additional results of QA accuracy in the A-Split with GPT-o4-mini-high. R% and I% represent the percentage the model assumes relative or intrinsic FoR for ambiguous expression, explained in Section 5.1. Acc is the accuracy, and Avg is the micro-average of accuracy. (1): 0-shot, (2): 4-shot, (3): CoT, and (4): SG+CoT.

Model	Camera perspective						Relatum perspective					
	ER (CP)	EI (RP)	II (RP)	IR (CP)	Avg.		ER (CP)	EI (RP)	II (RP)	IR (CP)	Avg.	
GPT-4o (0-shot)	79.7	45.1	39.5	90.2	64.2		46.9	88.5	98.2	34.8	67.5	
GPT-4o (4-shot)	68.0	52.6	60.7	74.1	61.8		44.9	98.2	100.0	37.5	71.2	
GPT-4o (CoT)	81.7	76.1	82.4	71.5	78.8		53.0	91.1	90.6	50.8	71.9	
GPT-4o (SG+CoT)	97.9	72.2	72.7	93.4	85.8		48.9	96.3	95.9	36.1	71.8	
o4-mini-high	68.7	73.1	72.1	79.0	71.4		78.9	91.9	93.9	62.7	84.5	

Table 17: Additional results of QA accuracy in the C-Split with GPT-o4-mini-high. ER, EI, II, and IR denote external relative, external intrinsic, internal intrinsic, and internal relative FoRs. Avg represents the micro-average accuracy. CP refers to context with camera perspective, while RP denotes context with relatum perspective.

```

the form of Explanation:
Explanation. Answer: Category.
# SG Instruction answer: Find the frame
of reference category with an
explanation. Additionally, include
three more explanations relating to
topology, distance, and direction
from the given context. The answer
should be in the form of \"
Explanation: Explanation. Topology:
Explanation. Distance: Explanation.
Direction: Explanation. Answer:
Category.

Context: {spatial exprssion}
"""

```

Listing 2: Prompt for finding the frame of reference class of given context.

```

## These are four examples used in few-
shot, CoT, and SG
# Example 1: The bird is outside the car
and in front of the car relative to
the car.
# Few-shot response: Answer: external
intrinsic
# CoT response: Explanation: The car
does not contain the bird, and the
spatial relation is relative to the
car's facing orientation. The
intrinsic direction of the car is
relevant in this context, as the
description \"in front of\" implies
a direction relative to the sheep's
front or back. Answer: external
intrinsic.
# SG response: Explanation: The car does
not contain the bird, and the
spatial relation is relative to the
car's facing orientation. The
intrinsic direction of the car is
relevant in this context, as the
description \"in front of\" implies

```

```

a direction relative to the sheep's
front or back. Topology: The bird is
outside the car. Distance: There is
little distance between the bird
and the car. Direction: The bird is
in the front direction of the car (
entity B). Answer: external
intrinsic.

# Example 2: The bird is inside of the
car and left of the car from the car
's perspective.
# Few-shot response: Answer: internal
intrinsic
# CoT response: Explanation: In this
scenario, the bird is contained
within the car, meaning it is an
internal relationship. Additionally,
the description specifies the bird's
position relative to the car,
indicating that the spatial relation
is relative to the car itself,
considering its intrinsic direction.
Answer: internal intrinsic.
# SG response: Explanation: In this
scenario, the bird is contained
within the car, meaning it is an
internal relationship. Additionally,
the description specifies the bird's
position relative to the car,
indicating that the spatial relation
is relative to the car itself,
considering its intrinsic direction.
Topological: The bird is inside the
car. Distance: The bird is in the
same location as the car. Direction:
The bird is in the left direction
of the car (entity B). Answer:
internal intrinsic.

# Example 3: The box is inside and at
the back of the room.
# Few-shot response: Answer: internal
relative

```

1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489

1490 # CoT response: Explanation: In this
1491 context, the box (entity A) is
1492 described as inside the room (entity
1493 B), fulfilling the first condition
1494 for an internal frame of reference.
1495 However, the spatial relation given
1496 (" at the back") cannot originate
1497 from the room since the room does
1498 not have intrinsic direction,
1499 implying it is relative to an
1500 external observer or reference point
1501 (C), such as a camera or a person
1502 viewing the scene. Answer: internal
1503 relative
1504 # SG response: Explanation: In this
1505 context, the box (entity A) is
1506 described as inside the room (entity
1507 B), fulfilling the first condition
1508 for an internal frame of reference.
1509 However, the spatial relation given
1510 (" at the back") cannot originate
1511 from the room since the room does
1512 not have intrinsic direction,
1513 implying it is relative to an
1514 external observer or reference point
1515 (C), such as a camera or a person
1516 viewing the scene. Topological: The
1517 box is inside the room. Distance:
1518 The box is in the same location as
1519 the room. Direction: From the
1520 observer's perspective, the box is
1521 located at the back of the room.
1522 Answer: internal relative
1523
1524 # Example 4: A phone is on the left of a
1525 tablet from my perspective.
1526 # Few-shot response: Answer: external
1527 relative
1528 # CoT response: Explanation: In this
1529 context, the spatial relation is
1530 described from the perspective of an
1531 observer (C) who is not involved in
1532 the spatial relation. The phone (A)
1533 is not contained by the tablet (B),
1534 and the spatial relation "left of
1535 \" is relative to the observer's
1536 perspective, not the tablet's facing
1537 orientation. Answer: external
1538 relative.
1539 # SG response: Explanation: In this
1540 context, the spatial relation is
1541 described from the perspective of an
1542 observer (C) who is not involved in
1543 the spatial relation. The phone (A)
1544 is not contained by the tablet (B),
1545 and the spatial relation "left of
1546 \" is relative to the observer's
1547 perspective, not the tablet's facing
1548 orientation. Topological: The phone
1549 is not contained by the tablet.
1550 Distance: There is some distance
1551 between the phone and the tablet.
1552 Direction: From the observer's
1553 perspective, the phone is located to
1554 the left of the tablet. Answer:
1555 external relative.

Listing 3: Spatial expression examples using for few-shots, Chain-of-Thought (CoT), and Spatial Guide (SG) prompting

K.2 Question Answering Example

```
'''
#Context: The bird is outside the car
and in front of the car relative to
the car. The car is facing toward
the camera.
#Question: Based on the camera's
perspective, where is the bird from
the car position in the scene?
#SG information (used in SG + CoT):
Frame of reference Explanation: The
car does not contain the bird, and
the spatial relation is relative to
the car's facing orientation. The
intrinsic direction of the car is
relevant in this context, as the
description "in front of" implies
a direction relative to the sheep's
front or back. Topology: The bird is
outside the car. Distance: There is
little distance between the bird
and the car. Direction: The bird is
in the front direction of the car (
entity B). Frame of reference:
external intrinsic.
#Normal response: Answer: front
#CoT / (SG + COT) response:Explanation:
Based on the context, the bird's
position is in the front direction
of the car. The car is facing toward
the camera. Then, the car's front
direction is the camera's front
direction. Therefore, the bird's
position is in front of the car's
position from the camera's
perspective. Answer: front

#Context: The bird is inside the car and
left of the car from the car's
perspective. The car is facing to
the right relative to the camera. #
Question: Based on the camera's
perspective, where is the bird from
the car's position?
#SG information (used in SG + CoT):
Frame of reference Explanation: In
this scenario, the bird is contained
within the car, meaning it is an
internal relationship. Additionally,
the description specifies the bird's
position relative to the car,
indicating that the spatial relation
is relative to the car itself,
considering its intrinsic direction.
Topological: The bird is inside the
car. Distance: The bird is in the
same location as the car. Direction:
The bird is in the left direction
of the car (entity B). Frame of
reference: internal intrinsic.
#Normal response: Answer: back
#CoT / (SG + COT) response: Explanation:
Based on the context, the bird's
position is in the left direction of
the car. The car is facing to the
right. Then, the car's left
direction is the camera's back
direction. Therefore, the bird's
position is to the back of the car's
```

```

1627         position from the camera's
1628         perspective. Answer: back
1629
1630 Context: The box is inside and to the
1631 right of the room from the observer'
1632 s perspective.
1633 Question: From the observer's
1634 perspective, what is the spatial
1635 relation of the box to the room?
1636 #SG information (used in SG + CoT):
1637 Frame of reference Explanation: In
1638 this context, the box (entity A) is
1639 described as inside the room (entity
1640 B), fulfilling the first condition
1641 for an internal frame of reference.
1642 However, the spatial relation given
1643 (\ " at the back\ ") cannot originate
1644 from the room since the room does
1645 not have intrinsic direction,
1646 implying it is relative to an
1647 external observer or reference point
1648 (C), such as a camera or a person
1649 viewing the scene. Topological: The
1650 box is inside the room. Distance:
1651 The box is in the same location as
1652 the room. Direction: From the
1653 observer's perspective, the box is
1654 located at the back of the room.
1655 Frame of reference: internal
1656 relative.
1657 #Normal response: Answer: right
1658 #CoT / (SG + COT) response: Explanation:
1659 Based on the context, the box is to
1660 the right of the room from the
1661 camera's direction. Therefore, the
1662 box's position is to the right of
1663 the room's position from the
1664 observer's perspective. Answer:
1665 right
1666
1667 Context: A phone is to the left of a
1668 tablet from my perspective. The
1669 tablet is facing to the right.
1670 Question: From my perspective, what
1671 is the spatial relation of the phone
1672 to the tablet?
1673 #SG information (used in SG + CoT):
1674 Frame of Reference Explanation: In
1675 this context, the spatial relation
1676 is described from the perspective of
1677 an observer (C) who is not involved
1678 in the spatial relation. The phone
1679 (A) is not contained by the tablet (
1680 B), and the spatial relation \ "left
1681 of\ " is relative to the observer's
1682 perspective, not the tablet's facing
1683 orientation. Topological: The phone
1684 is not contained by the tablet.
1685 Distance: There is some distance
1686 between the phone and the tablet.
1687 Direction: From the observer's
1688 perspective, the phone is located to
1689 the left of the tablet. Frame of
1690 Reference: external relative.
1691 #Normal response: Answer: left
1692 #CoT / (SG + COT) response: Explanation:
1693 Based on the context, the phone is
1694 to the left of the tablet from my
1695 perspective. The direction of the
1696 tablet is not relevant in this

```

```

context since the left relation is
from my perspective. Therefore, from
my perspective, the phone is to the
left of the tablet. Answer: left
'''

```

Listing 4: Spatial expression examples using for few-shots, Chain-of-Thought (CoT), and Spatial Guide (SG) prompting for question-answering.

K.3 Text to Layout

```

# Instruction for generating
    bounding box
"""
Your task is to generate the bounding
boxes of objects mentioned in the
caption.
The image is size 512x512. The bounding
box should be in the format of (x, y
, width, height). Please considering
the frame of reference of caption
and direction of reference object if
possible. If needed, you can make
the reasonable guess.
"""

```

Listing 5: Prompt for generating bounding coordinates to use as the layout for layout-to-image models.