

1	Appendix Table of <i>RoboRefer</i>	
2	A Implementation Details and Samples of <i>RefSpatial</i> Dataset	3
3	A.1 2D Web Image	3
4	A.1.1 Multi-Stage Image Filtering	3
5	A.1.2 Pseudo-3D Scene Graphs Construction	5
6	A.1.3 Hierarchical Object Description Generation	6
7	A.1.4 Generating Diverse QA Pairs via Pseudo-3D Scene Graphs	10
8	A.2 3D Embodied Video	13
9	A.2.1 Why Use CA-1M and How to Pre-process It	13
10	A.2.2 Inherent Challenges and Limitations in CA-1M	13
11	A.2.3 Addressing Limitations: Object Annotation and Bounding Box Filtering . .	14
12	A.2.4 3D Object Description Generation and Scene Graph Construction.	15
13	A.2.5 Free Space QA Generation for Object Placement	15
14	A.2.6 Generating Diverse QA via precise 3D annotations	19
15	A.3 Synthetic Data Generation in the simulator	21
16	A.3.1 Indoor Scene Generation	21
17	A.3.2 3D Asset Selection and Preparation	21
18	A.3.3 3D Asset Annotation with LLM	22
19	A.3.4 Scene Population and Data Rendering	24
20	A.3.5 Question-Answer Pair Generation	25
21	B Implementation Details and Samples of <i>RefSpatial-Bench</i>	25
22	C Implementation Details for <i>RoboRefer</i>	26
23	C.1 Architecture	26
24	C.2 Training Data	27
25	C.3 SFT Training Details	27
26	C.4 RFT Training Details	28
27	C.4.1 Sampling Action Groups	28
28	C.4.2 Reward Design and Policy Update	28
29	D Experimental Setting and Details	30
30	D.1 Experiments Compute Resources	30
31	D.2 Spatial Understanding Benchmarks	30
32	D.3 Spatial Referring Benchmarks	30
33	D.4 Simulation Evaluation	31
34	D.5 Real-world Evaluation	31
35	D.5.1 UR5 Manipulation	31
36	D.5.2 G1 Humanoid Mobile Manipulation	31
37	D.6 More Ablation Studies	32

38	E More Demonstrations	32
39	F More Discussion on Limitations and Future Work	33
40	G Broader Impacts	33
41	H Licenses	34

42 A Implementation Details and Samples of *RefSpatial* Dataset

43 In this section, we present a comprehensive exposition of the implementation details and representative
44 data samples underpinning the construction of the *RefSpatial* dataset. As this dataset is intended to
45 equip general VLMs with the ability to adapt to spatial referring tasks, thereby enhancing spatial
46 understanding and reasoning in a bottom-up manner, we meticulously design a multi-data-source
47 generation pipeline. We elaborate on the three core components of this pipeline as follows:

48 • **2D Web Image (Appx. A.1):** We present a 2D data pipeline comprising image filtering, pseudo-3D
49 scene graph construction, hierarchical referential description generation—from coarse categories to
50 fine-grained spatial referents—and diverse QA pair creation.

51 • **3D Embodied Video (Appx. A.2):** This section outlines the 3D data selection process from CA-
52 1M [1], discusses its limitations and mitigation strategies, and presents methods for enriched scene
53 graph construction compared to the 2D data source. We further describe a QA generation framework
54 that leverages detailed 3D annotations (*e.g.*, depth maps, oriented 3D bounding boxes) to capture
55 richer spatial relations. Finally, we detail how to generate QA pairs for the problem of “feasibility
56 assessment for object placement in free space”.

57 • **Synthetic Data from Simulator (Appx. A.3):** We describe how to synthesize 3D scenes, select
58 and annotate digital assets, efficient scene assembly and rendering, and the generation of QA pairs
59 grounded in these simulated scenes.

60 In the following subsections of each section, we detail the employed models, prompt design rationale,
61 data processing steps, filtering criteria, and illustrative examples, providing a clear and thorough
62 overview of the construction pipeline and core technical details of the *RefSpatial* dataset.

63 A.1 2D Web Image

64 A.1.1 Multi-Stage Image Filtering

65 2D Web Images aim to endow the model with basic spatial concepts and comprehensive depth
66 perception across both indoor and outdoor scenes. Here we use OpenImage¹ [2] as 2D data source.

67 **Overall Motivation and Goals for Filtering.** The OpenImages dataset offers a vast collection
68 of 2D internet images (1.7 M images in training split) with extensive visual diversity, but many
69 images, such as text-only graphics, QR codes, medical scans, or abstract art, lack relevance for spatial
70 understanding and referring with reasoning. To curate a dataset tailored for these tasks, we employ
71 a two-stage filtering pipeline: a coarse pre-filtering using SigLIP2 [3], followed by fine-grained
72 selection via Qwen2.5-VL [4], to retain images rich in spatial semantics.

73 **Stage 1: Initial Coarse Filtering.** We employ the `siglip2-giant-opt-patch16-384` model for
74 initial filtering to efficiently discard low-quality or off-theme images (*e.g.*, irrelevant scenes or
75 content lacking multiple everyday objects). This step greatly reduces data volume, streamlining
76 subsequent processing. Specifically, the SigLIP2 model is guided by predefined positive and negative
77 textual labels. Positive labels represent desired image characteristics, while negative labels describe
78 undesired content. For each image, the model computes cosine similarity between its embedding and
79 all label embeddings. The label with the highest similarity is selected; if it belongs to the positive
80 set, the image is retained, otherwise discarded. Label sets are manually refined iteratively to balance
81 recall and precision, ensuring relevance while excluding noise. These labels act as semantic anchors
82 for visual-text alignment. In this stage, only 934k images are qualified to be retained. Positive and
83 negative label lists are provided in Listings 1 and 2. For more details about the compute resources
84 needed in this stage, please see Appx. D.1.

Listing 1: Positive Labels used during SigLIP2 filtering.

```
85 Positive Labels = [  
86     "Mid-distance observation of some objects on a table",  
87     "Some objects on the desktop",  
88     "Distant view of some animals",  
89     "Mid-distance observation of some animals",  
90     "Distant view of one object",
```

¹<https://storage.googleapis.com/openimages/web/index.html>

```

91     "Mid-distance observation of one object",
92     "Distant view of some objects",
93     "Mid-distance observation of some objects",
94     "Distant view of a person",
95     "Mid-distance observation of a person",
96     "Distant view of some people",
97     "Mid-distance observation of some people",
98     "Distant view of indoor scene",
99     "Distant view of outdoor scene",
100    "Distant view of traffic",
101    "Distant view of Urban architecture"
102 ]

```

Listing 2: Negative Labels used during SigLIP2 filtering.

```

103 Negative Labels = [
104     "Macro shot of an animal",
105     "Macro shot of one object",
106     "Macro shot of a person",
107     "Macro shot of flowers",
108     "A piece of text",
109     "A person displayed in front of a white background",
110     "A product displayed in front of a white background",
111     "A screenshot of the graphics user interface",
112     "A dimly lit environment"
113 ]

```

SigLIP2 preserves images rich in object diversity, depth cues, and scene context (indoor/outdoor) through the above labeling process, but its performance declines on certain image types, including:

1. **Paintings/Artworks:** Especially those with visible brushstrokes or canvas textures.
2. **Low-Light Scenes:** Images with minimal illumination and strong shadows.
3. **Grayscale Photographs:** Black-and-white imagery lacking color cues.
4. **Distorted Images:** Those exhibiting warping, mirroring, or other geometric anomalies.
5. **Multi-Scene Collages:** Images containing three or more distinct scenes with hard borders.

SigLIP2 struggles to detect and interpret these categories reliably, highlighting the need for a secondary, fine-grained filtering stage.

Stage 2: Fine-grained Filtering Due to SigLIP2’s limitations in handling certain visual content mentioned above, **we introduce a fine-grained filtering stage using the Qwen2.5-VL-7B model to further improve dataset quality. This step ensured the final images are clear, authentic, and well-suited** for spatial understanding and reasoning required for spatial referring tasks. Qwen2.5-VL processed 934k images filtered by SigLIP2, retaining 846k. Although Qwen2.5-VL offers higher filtering precision, its slower speed necessitated the use of SigLIP2 for initial fast filtering, significantly improving overall efficiency. To ensure accurate and consistent fine-grained filtering, we adopt a structured prompt engineering strategy for the Qwen2.5-VL model. The process begins with a system prompt (See Listing 3) that defines the model’s role as an image analysis expert, specifying key visual attributes to assess and negative categories to detect, and enforcing a strict workflow. For each image, a corresponding user prompt (See Listing 4) instructs the model to determine whether the image belongs to any predefined negative categories. The model’s response follows a structured format: if the segment after the pipe symbol (|) is “Yes”, the image is classified as negative and discarded; otherwise, it is retained. This prompting scheme ensures that the model adheres to a consistent output, enhancing the reliability of filtering outcomes. Please refer to Appx. D.1 for details on the computational requirements of this stage.

Listing 3: System Prompt for Qwen2.5-VL-7B filtering.

```

139 system_prompt = ""
140 You are an image analysis expert. Follow this workflow rigidly:
141
142 1. **Content Analysis**:

```



```

143 - Inspect: Main subjects, artistic style, visual characteristics
144 - Check: Lighting intensity, color channels, geometric integrity,
145       composition structure
146
147 2. **Category Verification** (YES if matches ANY):
148   a) Painting/Artwork - Visible brushstrokes/canvas texture
149   b) Dim Lighting - Very low brightness, heavy shadows
150   c) B&W Photo - Grayscale only (0 color channels)
151   d) Distorted Image - Warping/mirroring anomalies
152   e) Multi-image Collage - >=3 distinct scenes with hard borders
153
154 3. **Structured Response**:
155   Output EXACTLY in this format:
156   "[Analysis sentence]. | Yes/No"
157   - Analysis must contain observable evidence
158   - Final answer MUST use pipe separator
159
160 Examples of VALID responses:
161   "This image is a composite created by stitching together multiple
162     smaller images, with distinct white borders visible between
163     the individual components. | Yes"
164   "This image features vibrant colors, is neither an artistic
165     painting nor a composite of multiple images, and does not
166     conform to any of the specified categories. | No"
167   ""

```

Listing 4: User Prompt for Qwen2.5-VL-7B filtering.

```

168 user_prompt = ""
169 Analyze if this image belongs to ANY of these categories:
170 1. Painting/artwork
171 2. Dim lighting
172 3. Black-and-white
173 4. Geometric distortion
174 5. Multi-image collage
175
176 Respond EXACTLY FORMATTED as:
177 "[Your evidence-based analysis]. | Yes/No"
178 ""

```

Visualizing Overall Filtering Results. Fig. 1 presents a visual overview, showing the effectiveness of our multi-stage filtering pipeline. The first row shows the output of SigLIP2, which effectively removes images lacking spatial semantics, such as macro shots of animals, people, or flowers, textual content, and GUI screenshots. The second row demonstrates how Qwen2.5-VL-7B further eliminates unsuitable categories, including artworks, dimly lit scenes, black-and-white images, geometrically distorted content, and image collages. The third row displays the retained images after both stages, which exhibit rich spatial relationships, confirming their suitability for spatial understanding and reasoning and the overall quality of our dataset.

A.1.2 Pseudo-3D Scene Graphs Construction

Although filtered 2D images contain some spatial cues, QA pairs containing sufficient 3D spatial information (*e.g.*, “near” vs. “far”, distances) derived directly from these 2D images are challenging. Inspired by prior work [5, 6], we construct pseudo-3D scene graphs from 2D images to enhance the generation of QA pairs with rich 3D spatial semantics. In these graphs, nodes represent object attributes, while edges encode inter-object spatial relations. We detail the process of converting 2D images into pseudo-3D scene graphs below.

Object Detection and Annotation. Although the OpenImages dataset provides annotations, its limited vocabulary and coarse labeling are insufficient for open-world scenarios. To address this, we leverage state-of-the-art foundation models for enhanced object detection and annotation. Specifically, our scene graph construction pipeline integrates the Recognize Anything Model (RAM) [7] and GroundingDINO [8] to assign semantic labels and bounding boxes to key objects in filtered raw 2D images. The workflow is broadly as follows:

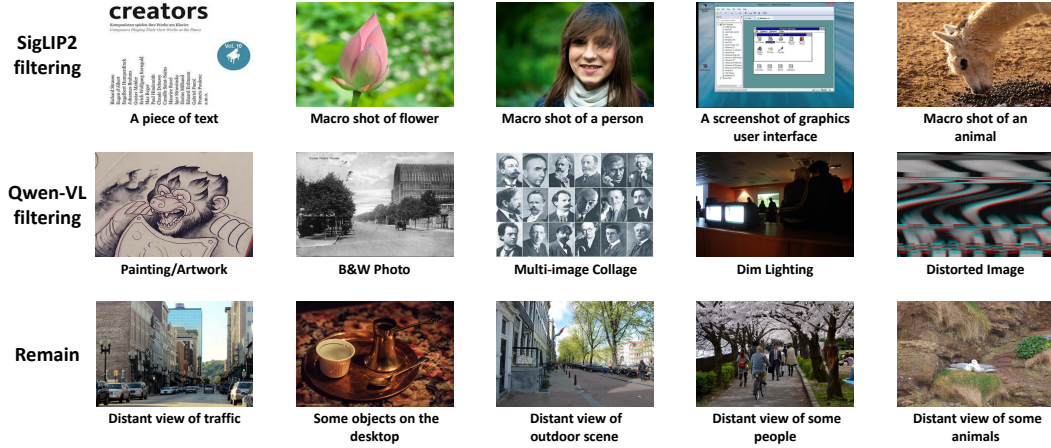


Figure 1: Visual overview of the multi-stage filtering results. Row 1: Images discarded by SigLIP2 due to insufficient spatial context (*e.g.*, close-ups, text). Row 2: Additional filtering by Qwen2.5-VL removes non-natural content (*e.g.*, artwork, collages). Row 3: Remaining high-quality images suitable for spatial understanding and referring.

1. **Semantic Labeling via RAM:** RAM analyzes each image to generate category labels for all recognized objects. Its broad recognition capability ensures comprehensive semantic coverage, guiding subsequent localization.
2. **Bounding Box Localization via GroundingDINO:** The labels from RAM are used as text prompts for GroundingDINO, an open-vocabulary detector that localizes the corresponding objects and outputs precise bounding boxes.

Although recent VLM, *i.e.*, Florence-2-Large supports instruction-based recognition and detection simultaneously, we find that combining RAM with GroundingDINO yields superior results. In Fig. 2, Florence-2 (top) often produces ambiguous or redundant detections (*e.g.*, vague labels, multiple boxes for a single object, or single boxes covering multiple objects), which are unsuitable for precise object referring. In contrast, GroundingDINO+RAM (bottom) generates concise labels and one-to-one object-bounding-box mappings, better satisfying the requirements of referring tasks.

3D-aware information Extraction. To further extract 3D-aware information from 2D images, we adopt UniDepth V2 [9] for metric depth estimation due to its recent state-of-the-art performance, surpassing models such as DepthPro and Metric3Dv2 across multiple benchmarks. For camera intrinsic prediction, we employ WildeCamera [10]. Together, these models enable robust 3D point cloud reconstruction of the scene. Based on previously annotated object bounding boxes, we apply SAM 2.1 [11] to generate instance masks. Each resulting Pseudo-3D scene graph comprises object labels (via RAM), 2D bounding boxes (via GroundingDINO), instance masks (via SAM 2.1), and object-level point clouds (via UniDepth v2 with WildeCamera), resulting in **axis-aligned 3D bounding boxes**. A visualization is provided in Fig. 3.

A.1.3 Hierarchical Object Description Generation

While 3D scene graphs encode basic object categories, real-world scenes often contain multiple instances per category. Prior datasets [12, 13] simplify reference by assuming a single object per category, limiting their utility in spatial referring tasks. To overcome this, **we augment object descriptions with attributes and spatial relations, enabling finer-grained disambiguation among similar instances of the same category.** We present our two-stage generation pipeline below.

Stage 1: Generating Object and Image Dense Descriptions in image space. We first generate detailed descriptions for each detected object and the entire image, employing the Qwen2.5-VL-7B model. This stage involves both object-level captioning and comprehensive image-level captioning. The global captions are essential for providing contextual grounding to downstream large language models (QwQ-32B) during LLM QA generation (detailed in Appx. A.1.4), enhancing the relevance and accuracy of the outputs. Prompt templates are detailed in Listings 5 and 6. In particular, the

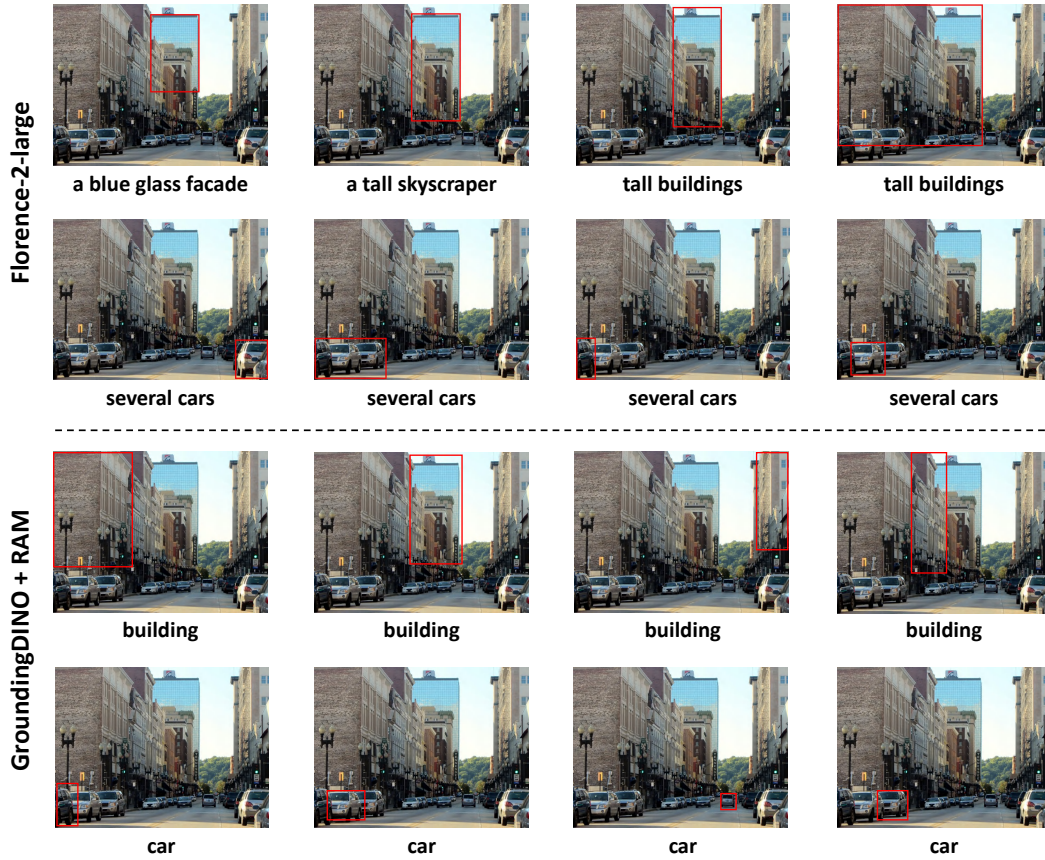


Figure 2: Object detection comparison: Florence-2 (above), GroundingDINO+RAM (below)



Figure 3: Scene graph visualization with image, detected objects, and corresponding point clouds.

233 object_caption_user_text_prompt uses a dynamic placeholder [class_name], which is filled
 234 with the object category predicted by the RAM model (See Appx. A.1.2).

Listing 5: Prompts for Image Caption Generation with Qwen-VL.

```
235 image_caption_system_text_prompt = """
236     You are an expert image analysis assistant. Your task is to
237     generate a detailed and comprehensive description of the image
238     .
```

```

239     Please focus on accurately capturing all visual elements present
240     in the image, including objects, scenery, colors, shapes,
241     textures, and lighting.
242     Your description should be clear, precise, and professional.
243     Additionally, ensure that your description begins with either
244     'this image' or 'the image'.
245     """
246
247     image_caption_user_text_prompt = """
248     Please carefully examine the provided image and generate a
249     detailed description.
250     Include all visible elements such as objects, scenery, colors,
251     shapes, textures, and lighting.
252     Ensure that your description is thorough, accurate, and complete,
253     and that it starts with either 'this image' or 'the image'.
254     """

```

Listing 6: Prompts for Object Caption Generation with Qwen-VL.

```

255     object_caption_system_text_prompt = """
256     You are a visual localization analyzer working with TWO distinct
257     images:
258     1. [POSITION-REFERENCE] (First Image):
259     - Contains ONLY location clues with background
260     - Strictly use ONLY for determining spatial position (left/right/
261       upper/lower/center)
262     - Ignore all visual features except object placement
263
264     2. [DETAIL-SOURCE] (Second Image):
265     - Shows the object's TRUE APPEARANCE without background
266     - Extract EXCLUSIVELY from this: color, texture, material, shape
267     - Never infer details from the first image
268
269     Generate phrase in pattern: [Color][Material][Object] at [Position
270       ]
271     Example: "Matte black laptop on the left" NOT "Red-boxed laptop"
272     """
273
274     object_caption_user_text_prompt = """
275     For the [class_name] marked by red box in FIRST image and fully
276     shown in SECOND image:
277     -> COLOR/MATERIAL: Must come from SECOND image
278     -> POSITION: Only from FIRST image's placement
279     Forbidden actions:
280     x Mention 'red box' or background elements
281     x Use location terms in second image
282     x Combine features across images
283
284     Describe the [class_name] marked by red box in FIRST image and
285     fully shown in SECOND image with this format:
286     [Color][Material/Texture][Object] at [Position]
287     Samples:
288     - "Brushed metal water at bottle left"
289     - "Glossy ceramic mug at upper center"
290     - "Faded denim jacket at lower right"
291     """

```

292 **Stage 2: Generating Object Description with Spatial Cues.** To enhance referential specificity
 293 in object captions, particularly when multiple instances of the same category coexist, we adopt a
 294 heuristic strategy that appends spatially indicated relative information, such as “the third chair from
 295 the front”. This method leverages 3D object positions from the scene graph (See Appx. A.1.2). By
 296 comparing same-category objects along the three principal axes (front–back, left–right, top–bottom),
 297 we identify the axis with the largest spatial variation as the primary arrangement direction to guide
 298 relative spatial reference generation. Once the main sorting axis is identified, we retrieve appropriate

299 templates from a predefined library (see Listing 7) to augment the initial object descriptions. These
 300 templates are designed to capture diverse natural language patterns. For instance, for a row of chairs
 301 arranged left to right, templates may include: “{dense_caption}, which is the {ordinal} {class_name}
 302 from left to right,” or “{dense_caption}, the {ordinal} {class_name} in the left-to-right sequence”.
 303 Here, dense_caption denotes the initial description generated by the Qwen2.5-VL model, ordinal
 304 indicates the object’s position in the sorted sequence, and class_name is the category label predicted
 305 by RAM. This spatially-aware enhancement is applied only when multiple instances of the same
 306 category are detected to avoid redundancy. If an object appears only once, its original dense caption
 307 is used directly. To ensure spatial diversity, we set a variance threshold across the three principal axes;
 308 images with multiple same-category objects but low variance on all axes are discarded, resulting in
 309 a final set of 466k images. By integrating spatial ordering with visual descriptions, this heuristic
 310 enables the generation of precise and discriminative referential expressions, essential for producing
 311 high-quality, unambiguous question-answer pairs.

Listing 7: Templates for Spatial Order Description Enhancement.

```

312 TEMPLATES = {
313     "left_to_right": [
314         "{dense_caption}, which is the {ordinal} {class_name} from
315         left to right",
316         "{dense_caption}, marked as the {ordinal} {class_name} in a
317         left-to-right arrangement",
318     ],
319     "right_to_left": [
320         "{dense_caption}, the {ordinal} {class_name} viewed from the
321         right",
322         "{dense_caption}, the {ordinal} {class_name} from the right",
323     ],
324     "front_to_back": [
325         "{dense_caption}, which appears as the {ordinal} {class_name}
326         when viewed from the front",
327         "{dense_caption}, positioned as the {ordinal} {class_name} in
328         front-to-back order",
329     ],
330     "back_to_front": [
331         "{dense_caption}, which is counted as the {ordinal} {
332         class_name}, starting from the back",
333         "{dense_caption}, the {ordinal} {class_name} in the back-to-
334         front sequence",
335     ],
336     "top_to_bottom": [
337         "{dense_caption}, the {ordinal} {class_name} viewed from the
338         top",
339         "{dense_caption}, placed as the {ordinal} {class_name} when
340         sorted from top to bottom",
341     ],
342     "bottom_to_top": [
343         "{dense_caption}, which ranks as the {ordinal} {class_name} in
344         bottom-to-top order",
345         "{dense_caption}, arranged as the {ordinal} {class_name} when
346         ordered from the bottom",
347     ]
348 }
  
```

349 **Examples of Object and Image Descriptions** This part qualitatively shows the representative
 350 examples with the generated descriptions. As shown in Fig. 4, we present two types of object captions.
 351 The top row shows simple captions produced by Qwen2.5-VL for single-instance object categories,
 352 where spatial ordering is unnecessary. The bottom row includes captions augmented with spatial order
 353 information to distinguish multiple instances of the same category. Additionally, Fig. 5 demonstrates
 354 Qwen2.5-VL’s ability to generate detailed global descriptions of entire images used in the following.



Figure 4: Generated object descriptions. Top: Unique-category captions. Bottom: Spatially-aware captions for the same categories. Red boxes indicate referenced objects.

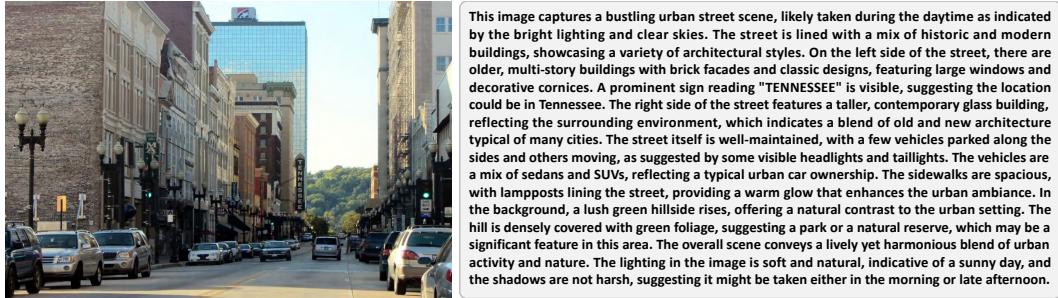


Figure 5: Visualization of generated image detailed descriptions.

355 A.1.4 Generating Diverse QA Pairs via Pseudo-3D Scene Graphs

356 After constructing scene graphs and generating hierarchical object descriptions, we can leverage this
 357 information to generate diverse QA pairs from pseudo-3D scene graphs to support SFT training for
 358 improved single-step spatial understanding.

359 **Template, Choice and Fact QA Generation.** We first adopt a template-based method to generate
 360 structured preliminary QA pairs, multiple-choice questions, and factual statements. The templates
 361 are derived from scene graph information (*e.g.*, object attributes, positions) and refined hierarchical
 362 object descriptions. When designing QA templates using pseudo-3D scene graphs from 2D images,
 363 we explicitly account for the spatial ambiguity inherent in single-view pseudo-3D representations
 364 (*e.g.*, inaccuracies in monocular depth, lack of object orientation, absence of 3D oriented bounding
 365 boxes). Consequently, **our QA templates from 2D data source focus mainly on qualitative spatial**
 366 **understanding and reasoning**, while incorporating quantitative cues only when they can be inferred
 367 reliably from 2D or pseudo-3D signals. The spatial concepts covered in the QA templates fall into
 368 the following categories:

- 369 1. **Relative position relations:** capture spatial layouts (left/right, above/below, front/behind).
- 370 2. **Relative size comparisons:** describe object attributes (*e.g.*, bigger/smaller, taller/shorter,
 371 wider/thinner) often inferred from image-plane projections.
- 372 3. **Quantitative information from 2D or pseudo-3D label:** include spatial reasoning based on
 373 estimated depth maps, 2D object coordinates, and coarse monocular depth approximations.

374 Accordingly, we design diverse QA templates in these types for spatial understanding:

- 375 1. Templates assess spatial and size relations:
 - 376 • Position relation: "Is [A] to the left of [B]?"

- Size comparison: “Which object is larger, [A] or [B]?”
- 2. Templates query 2D point of a uniquely identified object ‘[A]’ (defined in Appx. A.1.3):
 - “Where is [A] located? Provide its 2D coordinates.” (e.g., “Where is the red apple at left, which is the second apple from left to right, located? Provide its 2D coordinates.”)
- 3. Templates query attributes at a specific 2D point ‘[X]’ (formatted as ‘(x, y)’):
 - Depth retrieval: “What is the depth at point [X]?” (e.g., “What is the depth at (0.528, 0.317)?”)
 - Object identification: “Which object is at point [X]?” (e.g., “Which object is at (0.753, 0.839)?”)

We further design fact templates to generate declarative statements, forming a structured basis for prompting Reasoning LLM to produce richer and more natural QA pairs. Example templates include:

1. Approximate depth: “Point [A] and the camera are [X] apart.” (based on depth estimation).
2. Precise 2D object location: “[A] is located at point [X].”, “[A] is to the right of [B].”

Reasoning QA Generation. To generate more natural, complex, and diverse QA pairs beyond templated formats, we leverage a powerful reasoning LLM, QwQ-32B [14]. It takes the factual statements, initial QA pairs, and multiple-choice questions (if available) as input, along with global image captions and precise object descriptions. QwQ-32B then produces more challenging and conversational spatial reasoning QA. The prompt design is shown in Listing 8.

Listing 8: Prompt for QwQ-32B QA Diversification and reasoning QA Generation.

```

"""
You are a helpful assistant tasked with generating spatial reasoning-
based questions and answers from provided descriptions of scenes.

Rules:
1. **We have three types of input information**:
- **[Scene]**: A general description of the entire image, which
  provides context for the objects and their surroundings.
  **Example**:
  [Scene]: The image shows a tranquil lakeside with a small wooden
  dock on the right and calm, reflective water in the center.
  The sky is overcast.
- **[Objects]**: A list containing one or more object labels separated
  by "|".
  **Example**:
  [Objects]: teal glossy water at lower center | green bamboo dock
  at lower right.
- **[Objects Description]**: Provides spatial or comparative details
  between those objects.
  **Example**:
  [Objects Description]: teal glossy water at lower center is taller
  than green bamboo dock at lower right.

2. **When crafting a Question**:
- **Always use the provided [Scene] description as context** to ensure
  the question aligns with the overall image.
- **Mention all object labels from [Objects]** in the question.
- **Do not modify or paraphrase the object labels** they must appear
  exactly as given in '[Objects]'.
- **Do not assume or invent additional scene details** beyond what is
  provided in '[Scene]'.

```

```

432 - **Do not reveal the specific details in [Objects Description]** (
433     like which object is taller, shorter, wider, etc.).
434 - Always generate questions related to the description using the
435     object labels from [Objects].
436 - Each object label in '[Objects]' **must appear exactly once** in the
437     Question.
438 - The question should read from **an observer's perspective**.
439 - The description should always be used to answer and not leak into
440     the question.
441
442 3. **When crafting an Answer**:
443 - **Mention at least one object label from [Objects]** in the answer.
444 - **Use the '[Objects Description]' to provide a correct answer**.
445 - **Ensure the answer is concise, factual, and directly related to the
446     provided '[Scene]' and '[Objects]'**.
447 - **You may restate or summarize the relevant details from '[Objects
448     Description]', but do not introduce new assumptions**.
449
450
451 Here's several examples:
452
453 [Scene]: The image depicts a modern living room with a large window
454     allowing warm sunlight to enter. The room has a wooden floor, a
455     patterned rug in the center, and a coffee table with a few
456     magazines neatly stacked on it. A yellow leather sofa is
457     positioned centrally, facing the television mounted on the
458     opposite wall. To the left of the sofa, a black metal chair with a
459     cushioned seat is placed beside a tall bookshelf filled with an
460     assortment of books and decorative items. The furniture
461     arrangement leaves an open pathway between the sofa and the chair.
462 [Objects]: yellow leather sofa at lower center, black metal chair on
463     the left.
464 [Objects Description]: The path between yellow leather sofa at lower
465     center and black metal chair on the left is 1.5 meters.
466 "Question": You are a cleaning robot that is 1 meter wide. Now you are
467     standing in a living room and see the image; you want to move
468     from here to the door that leads to the backyard. Do you think you
469     can go through the path between the yellow leather sofa at lower
470     center and the black metal chair on the left?
471 "Answer": The path between the yellow leather sofa at lower center and
472     the black metal chair on the left is 1.5 meters, so yes, the
473     robot can go through the path between the yellow leather sofa at
474     lower center and the black metal chair on the left since it is
475     wider than the robot's width.
476
477 [Scene]: The image showcases a modern kitchen with a wooden countertop
478     that extends across the space, separating the cooking area from
479     the dining area. On the left side of the countertop, a fruit bowl
480     holds a variety of fresh produce. A red fresh apple is placed on
481     the left side of the bowl, while a bright fresh orange sits neatly
482     on the right side. Behind the fruit bowl, a glass pitcher filled
483     with orange juice and a stack of white ceramic plates are visible.
484     Natural light streams in from a large window above the sink,
485     reflecting off the stainless steel appliances and giving the space
486     a bright, clean feel.
487 [Objects]: red fresh apple on the left, fresh orange on the right.
488 [Objects Description]: red fresh apple on the left is positioned on
489     the left side of fresh orange on the right.
490 "Question": You see two fruits, a red fresh apple on the left and a
491     fresh orange on the right. Which one is more on the left side?
492 "Answer": The red fresh apple on the left is more on the left.
493
494 Now its your turn!
495 ""

```


496 **Training data visualization.** For specific examples of training data generated from 2D web images
497 and their visualizations, please refer to Appx. E, which contains detailed sample presentations.

498 A.2 3D Embodied Video

499 A.2.1 Why Use CA-1M and How to Pre-process It

500 **Rationale for Selecting CA-1M as the 3D Data Source.** To enable fine-grained spatial reasoning
501 in indoor environments, we adopt Apple’s open-source CA-1M² [1] dataset as our primary 3D data
502 source. CA-1M aligns closely with our objectives due to the following key attributes:

- 503 1. **Dense 2D/3D Annotations:** CA-1M provides per-frame 2D/3D oriented bounding boxes,
504 enabling spatial localization (*e.g.*, 3D spatial occupancy) and accurate interaction modeling.
- 505 2. **Comprehensive Camera and Depth Data:** The inclusion of camera intrinsics, extrinsics, and
506 depth maps supports accurate 3D reconstruction and geometric reasoning.
- 507 3. **Large-Scale Coverage:** Its extensive volume enables training/evaluation of VLMs at scale.

508 **These features offer a strong foundation for constructing 3D scene graphs and generating spatial**
509 **reasoning data involving 3D geometry, object interactions, and egocentric understanding.**

510 **Comparative Analysis with Alternative 3D Datasets.** While several high-quality 3D datasets exist,
511 they fall short of meeting the specific requirements of our project, motivating our selection of CA-1M:

- 512 1. **ARKitScenes** [15]: As a predecessor to CA-1M, it provides only global 3D bounding boxes
513 without per-frame 3D or 2D annotations. Projecting these 3D boxes to 2D yields oversized
514 and contain irrelevant objects. Additionally, its annotations are less comprehensive, and
515 image resolution is also much lower compared to CA-1M.
- 516 2. **ScanNet V2** [16]: Lacks 3D bounding boxes, making object orientation estimation infea-
517 sible. Although EmbodiedScan introduces per-frame 3D bounding boxes, it still lacks 2D
518 annotations, and projected 2D boxes remain imprecise and contain irrelevant objects.
- 519 3. **3RScan** [17]: Suffers from low image quality, hindering spatial information extraction and
520 limiting its usability.

521 In summary, despite some limitations (See Appx. A.2.2), CA-1M provides large-scale egocentric
522 video, per-frame 2D/3D annotations, depth, and camera parameters, **making CA-1M the most**
523 **suitable choice for generating rich and complex 3D spatial data from an embodied perspective.**

524 **Pre-processing.** Video datasets capturing continuous activities, such as CA-1M, exhibit high tem-
525 poral redundancy, as consecutive frames often contain near-identical visual content with minor
526 variations. Processing all frames is computationally intensive and yields redundant samples with
527 limited informational gain for model training. To mitigate this, we adopt a frame sampling strategy,
528 selecting one frame every 20 frames. This reduces redundancy while preserving meaningful scene
529 and viewpoint transitions. The resulting subset maintains scene diversity and supports efficient
530 downstream processing, including 3D scene graph construction and question-answer generation. In
531 Fig. 6, the top row shows four consecutive frames before sampling, revealing minimal variation; the
532 bottom row, sampled at 20-frame intervals, demonstrates significantly greater scene variation.

533 A.2.2 Inherent Challenges and Limitations in CA-1M

534 Although the CA-1M dataset is chosen for its large scale, egocentric perspective, and rich annotations,
535 it presents inherent limitations that undermine scene graph construction if unaddressed.

536 **Ambiguous or Meaningless Object Annotations.** A major issue is the prevalence of ambiguous
537 or semantically insignificant object annotations. Many instances are difficult to interpret, even for
538 humans, due to unclear boundaries or a lack of identifiable semantics. For example, some bounding
539 boxes enclose small, indiscernible regions such as a patch of wall or vague background elements
540 (see Fig. 7). Incorporating such annotations into model training introduces noise, hampers spatial
541 understanding, and may mislead the model’s perception of object relationships.

²<https://github.com/apple/ml-cubifyanything>



Figure 6: Comparison of frame sequences before and after sampling. Top: Four consecutive frames from the original video. Bottom: Corresponding frames sampled every 20 frames, exhibiting increased temporal variation between adjacent frames.

542 **Widespread Absence of Semantic Labels.** Another major limitation of CA-1M is the lack of
 543 semantic labels for most annotated objects. Unlike datasets (*e.g.*, ARKitScenes and ScanNet V2),
 544 which provide object category annotations, CA-1M includes only a few structural categories (*e.g.*,
 545 floors, walls, doors), leaving the majority of object instances unlabeled. Although bounding boxes
 546 indicate the presence of objects (*i.e.*, labeled as “object”), their categories (*e.g.*, “chair”, “table”)
 547 remain unknown. This absence of semantic information hinders spatial understanding, making it
 548 impossible to generate category-dependent queries such as “Is the red chair to the left of the table?”



Figure 7: Example of some meaningless annotations in CA-1M.

549 A.2.3 Addressing Limitations: Object Annotation and Bounding Box Filtering

550 To address the limitations of the CA-1M dataset, particularly the lack of semantic labels and the
 551 prevalence of noisy bounding boxes (See Appx. A.2.2 above), we develop a dedicated annotation and
 552 filtering pipeline. This enhances the dataset’s usability for downstream tasks such as 3D scene graph
 553 construction and spatial reasoning. We detail the two-stage pipeline below.

554 **Stage 1: Initial Annotation and 2D Bounding Box Prediction.** We employ a multi-model combi-
 555 nation strategy to annotate unlabelled objects and refine bounding boxes in CA-1M video frames.
 556 Specifically, we integrated GroundingDINO, RAM, and Florence-2 to perform object semantic
 557 labeling and 2D bounding box prediction. **Our approach yields semantically meaningful and**
 558 **visually coherent object bounding boxes compared to the often difficult-to-discern or ambiguous**
 559 **bounding boxes from the original CA-1M annotations.** To maximize recall, we intentionally lower
 560 the confidence thresholds of GroundingDINO and RAM, ensuring the inclusion of low-confidence
 561 but potentially relevant objects. These candidates are retained for further validation in subsequent
 562 matching and filtering stages.

563 **Stage 2: Bidirectional 2D Bounding Box Matching.** To associate model-predicted 2D bounding
 564 boxes (with semantic labels) with unlabeled boxes in CA-1M, **we propose a bidirectional matching**
 565 **and refinement strategy. This not only enables semantic annotation of meaningful objects but**

also filters out noisy or ambiguous CA-1M bounding boxes, thus enhancing the utility of its 3D annotations (e.g., 3D oriented bounding boxes). The strategy consists of two steps:

1. **Matching CA-1M Bounding Boxes to Model Predictions.** We first match original CA-1M 2D bounding boxes to model-predicted bounding boxes based on the IoU metric. Due to the sparsity, occlusions, and fragmentation of annotation, multiple CA-1M boxes may correspond to a single prediction, resulting in many-to-one matches.
2. **Refining Model Predictions via One-to-One Mapping.** To resolve many-to-one matches, we retain only predicted boxes that match at least one CA-1M bounding box. For each, we assign the CA-1M bounding box with the highest IoU as its unique match. This enforces a one-to-one correspondence, eliminating redundant or weakly aligned CA-1M boxes. The result is a refined set of original bounding boxes with strong semantic alignment.

Visualizing Matching Results. Fig. 8 shows the effectiveness of our bounding box matching procedure. The first row shows successfully matched CA-1M bounding boxes annotated with RAM-predicted object labels, while the second row highlights unmatched bounding boxes, typically corresponding to ambiguously annotated objects. This demonstrates that our method reliably filters out uncertain annotations and assigns semantic labels to clearly identifiable objects. For subsequent scene graph generation, we adopt the model-predicted bounding boxes instead of the original CA-1M ones, as they better align with the visible object extents, facilitating more accurate instance mask extraction via models like SAM2.

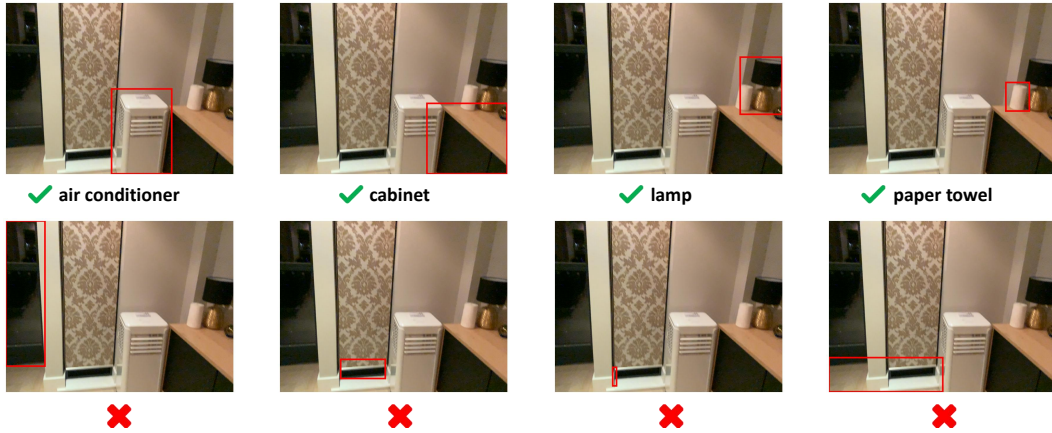


Figure 8: Visual comparison of 2D bounding box matching. The top row shows matched CA-1M bounding boxes with object labels; the bottom row displays unmatched or ambiguous cases.

A.2.4 3D Object Description Generation and Scene Graph Construction.

The process of 3D scene graphs construction largely follows the 2D scene graph pipeline in Appx. A.1.2 and the description generation in Appx. A.1.3. The resulting graphs are structurally similar to those from OpenImages in Fig. 3. The key distinction is CA-1M’s focus on indoor environments and its provision of high-precision geometric data (e.g., ground-truth depth, camera intrinsics/extrinsics, 3D oriented bounding boxes). **The enhanced 3D scene graph outperforms 2D-based counterparts in object localization and spatial relation accuracy, enabling structurally rich and quantitatively grounded 3D QA data across 28 spatial relation types..**

A.2.5 Free Space QA Generation for Object Placement

A key challenge in spatial referring is identifying unoccupied regions suitable for object placement. We propose a multi-step pipeline to address this.

Step 1: Detecting Viable Platforms via Qwen2.5-VL. We first filter out scenes lacking plausible placement surfaces (e.g., only walls or ceilings). To this end, we employ the Qwen2.5-VL-7B, initialized with a system prompt (Listing 9) that specifies its role and procedure. For each image, a user prompt (Listing 10) directs the model to detect candidate images with surfaces such as tables,

600 floors, or shelves. Fig. 9 shows the filtering results. The first row shows scenes without platforms,
601 while the second row shows scenes with platforms. This filtering enables subsequent computationally
602 intensive analysis to focus on semantically relevant scenes for placement queries.

Listing 9: System Prompt for Qwen2.5-VL to identify images with suitable platforms.

```
603 image_have_platform_system_text_prompt = ""
604 You are an expert visual scene understanding assistant.
605
606 Your task is to analyze an image and determine whether it contains **
607 any obvious flat horizontal surfaces** where physical objects can
608 be placed. These include **floors, tabletops, bed surfaces, or
609 other flat and stable areas**.
610
611 IMPORTANT:
612 - If you can see any part of the **floor**, **tabletop**, **bed**, or
613 **similar flat surfaces**, you MUST assume it can support physical
614 objects (\eg, books, boxes, pillows).
615 - Do NOT consider whether the surface is cluttered, partially visible,
616 or obstructed. If the platform exists and is horizontal, assume
617 it can hold objects.
618 - Your answer must be based strictly on visible surfaces.
619
620 You must provide a short reasoning based on visual evidence in the
621 image, followed by a final conclusion.
622
623 Your response MUST strictly follow this format:
624 "[Your analysis]. | Yes/No"
625
626 Examples of valid responses:
627
628 Example 1:
629 The image shows a wooden floor that is flat and unobstructed. And it
630 could potentially support physical objects. | Yes
631
632 Example 2:
633 There is a bed clearly visible in the scene with a flat top surface
634 where items like pillows or books can be placed. | Yes
635
636 Example 3:
637 A rectangular table is visible in the center of the image, providing a
638 flat surface suitable for placing objects. | Yes
639
640 Example 4:
641 The image contains mostly a wall with a window and no visible floor,
642 table, or other flat surfaces. | No
643
644 Do NOT provide any extra commentary or formatting outside this exact
645 format.
646 ""
```

Listing 10: User Prompt for Qwen2.5-VL to identify images with suitable platforms.

```
647 image_have_platform_user_text_prompt = ""
648 Please examine the image and determine whether it contains any **clear
649 horizontal platforms** where physical objects can be placed.
650 These platforms include: **floors, tabletops, bed surfaces, or
651 other flat and stable horizontal areas**.
652
653 Important notes:
654 - As long as **any part** of a **floor**, **tabletop**, **bed surface
655 **, or similar platform is visible in the image, you must assume
656 it is capable of supporting physical objects (such as books, boxes
657 , pillows, etc.).
```

658 - Do NOT consider whether the surface is messy, partially blocked, or
659 whether there's enough space. If the platform exists and is
660 horizontal, you must assume it can hold objects.
661 - Your answer must be based entirely on visible visual evidence in the
662 image.
663
664 You must respond in the following exact format:
665 "[Your analysis]. | Yes/No"
666
667 Refer to the following examples to guide your response:
668
669 Example 1:
670 The image shows a wooden floor that is flat and unobstructed. And it
671 could potentially support physical objects. | Yes
672
673 Example 2:
674 There is a bed clearly visible in the scene with a flat top surface
675 where items like pillows or books can be placed. | Yes
676
677 Example 3:
678 A rectangular table is visible in the center of the image, providing a
679 flat surface suitable for placing objects. | Yes
680
681 Example 4:
682 The image contains mostly a wall with a window and no visible floor,
683 table, or other flat surfaces. | No
684
685 Please strictly follow the format and do not add any extra commentary.
686 ""



Figure 9: Visualization of platform filtering results. Row 1: Examples of images filtered out (lacking platforms). Row 2: Examples of images retained (containing platforms).

687 **Step 2: Gravity Alignment and Top-Down Projection.** To enable top-down (orthographic) scene
688 projection, we apply gravity alignment matrices from the CA-1M to transform point clouds and
689 object 3D bounding box into a consistent frame where gravity uniformly points downward. This
690 normalization allows for projection onto a plane orthogonal to the gravity vector, revealing object
691 layouts and spatial relations more clearly. Fig. 10 shows the point cloud and coordinate axes before
692 and after alignment.

693 **Step 3: Programmatic Platform Association.** After gravity alignment, we associate objects with
694 supporting platforms based on their spatial relationships. For relationships such as “front/behind/left-
695 /right”, we identify the platform supporting a target object by checking the top surfaces of candidate
696 platforms. A platform qualifies if its top surface is within 0.05 meters of the object’s bottom surface
697 and overlaps at least 70% of the object’s bottom area in top-down view. For the “below” relationship,
698 when an object may be suspended in mid-air, we identify the supporting platform as the one located

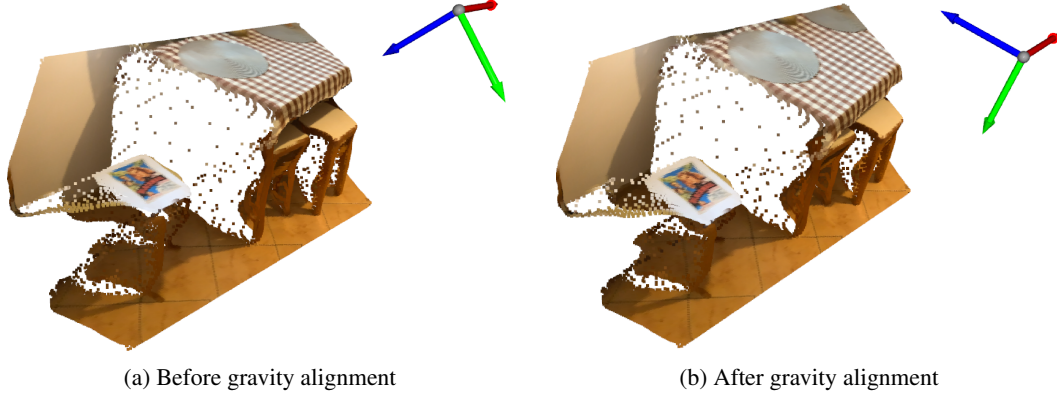


Figure 10: Visualization of gravity alignment. (a) The scene before alignment. (b) The scene after alignment, with the Y-axis oriented along the gravity vector.

beneath it, whose top surface is nearest to the object’s bottom, and whose top-down intersection with the object exceeds 70% of the object’s area. For the “**above**” relationship, the object’s own top surface is treated as the reference platform, eliminating the need for additional platform search. For the “**between**” relationship involving two objects, each object’s supporting platform is determined independently using the same procedure as in the “**front/back/left/right**” cases. The space between them is considered valid only if both share the same supporting platform, ensuring spatial reasoning occurs within a unified physical context.

Step 4: Identifying Other Objects on the Platform. After identifying the target object and its supporting platform, we locate other relevant objects co-occurring on the same platform. For “**front/behind/left/right/between**”, candidate objects are selected based on the following criteria:

1. Their bottom must not be significantly higher than the top of the target object.
2. Their top must remain above the platform surface.
3. Their XZ-plane footprint must intersect with that of the target object.
4. Their volume must not exceed 4.236 times that of the target object.

To mitigate the over-filtering of solid objects during visible point selection via depth matching, we introduce a volumetric constraint (thresholded at approximately 4.236, *i.e.*, $(1/0.618)^3$, the cube of the reciprocal of the golden ratio). This prevents large, potentially hollow objects—such as tables with substantial under-space—from being misclassified as fully occluding based solely on their bounding box volume, which would mislead placement reasoning on the primary platform. Instead, the filter targets small to medium-sized objects, where the bounding box volume more accurately reflects actual occupancy. For such objects, even if hollow, the limited under-space is typically negligible for placement purposes. This behavior is shown in the top-view occupancy maps in Fig. 11a and Fig. 13a, where a large table is excluded due to exceeding the volume threshold. Despite being present in the scene, its hollow geometry allows for usable space beneath, justifying its omission.

For the “**below**” relation, an object on a platform is considered below the target object (when the target might be suspended or on a higher tier) if:

1. Its footprint (XZ-plane projection) intersects with that of the target;
2. Its bottom is no higher than the top of the target;
3. Its top is not below the top surface of its supporting platform.

For the “**above**” relation, an object is considered above the platform of target object if:

1. Its bottom is within 20 cm above the platform’s top surface;
2. Its top is not below the platform’s top surface;
3. Its footprint (projection on the XZ plane) overlaps with that of the platform.

Step 5: Sampling Unoccupied Points in the Top-View Occupancy Map. After identifying the target object, its supporting platform, and adjacent objects, we determine the surrounding free space. This includes regions in front, behind, left, right, above, or below the target. To locate free space in the horizontal directions (“**front/behind/left/right**”), we define a 90° sector centered on the target and oriented in the respective direction. The sector radius is set to the maximum of either the diagonal of the object’s footprint or a fixed 20 cm, ensuring adequate coverage. For vertical directions (“**above/below**”), we project the object’s top or bottom surface onto the supporting platform. To mitigate overestimation from coarse 3D bounding boxes, we shrink the projection to 80% of its original size (centered), reducing overlap with nearby objects and better approximating usable space. To identify free space “**between**” two target objects, we define the search region as the planar area enclosed by the projections of both objects onto their shared supporting surface. This region is evaluated for occupancy by other objects. Across all spatial contexts (“**above/below/between**”), we enforce a minimum free area constraint: the unoccupied region must exceed $0.036m^2$ (half an A4 sheet) in a top-down view. This threshold filters out trivial cases and ensures the queried space can accommodate small objects such as a book or cup. Free space is computed by analyzing object footprints in the gravity-aligned top-down view. In Fig. 11a, 12a, 13a, blue-shaded regions indicate the final sampling areas after accounting for bounding box scaling and occlusions. These visualizations highlight viable unoccupied regions for subsequent placement analysis.

Step 5: Projection and Visibility Filtering. Given candidate points sampled in the top-down view (XZ-plane) as free space or target locations, we project them into the original 2D camera image to assess visibility. Each point is assigned the y coordinate of the platform’s top surface, forming full 3D coordinates. Using camera intrinsics, extrinsics, and gravity alignment, we project these 3D points onto the 2D image plane, as shown in Fig. 11b, 12b, 13b. To determine visibility, we compare the z -coordinate of each 3D point with the corresponding depth value in the aligned depth image. Points are discarded as occluded if this difference exceeds 2.5 cm. For the **front**, **behind**, **left**, and **right** directions, we sample 9,000 points per direction and retain the direction if at least 2,000 points remain visible. For the **above**, **below**, and **between** relations, we sample 10,000 points and require a minimum of 6,000 visible points. We compute the mean position of the remaining visible points as the representative target location. If its depth deviates from the depth image by more than 2.5 cm, we instead choose the nearest point within this threshold. Fig. 11c, 12c, 13c illustrate this process, highlighting visible points and the final selected target (blue circle).



(a) Top-view occupancy map with the target’s right-side search area (b) Sampled points projected onto the 2D image (c) Visible points with the final placement center

Figure 11: Visualization of **right**-side free space identification. (a) Top-view occupancy map with the target’s right-side search area. (b) Projection of sampled candidate points into the image plane. (c) Non-visible points are removed; the final placement center is marked with a blue circle.

A.2.6 Generating Diverse QA via precise 3D annotations

Building on the 2D QA generation pipeline (Appx. A.1.4), we utilize a template-based approach augmented with QwQ-32B to generate diverse QA pairs for 3D scenes. **Unlike 2D images, 3D datasets offer rich and precise spatial annotations—such as depth maps, camera poses, and per-object 3D bounding boxes—which enable the construction of more complex QA with spatial referring with reasoning.** Additionally, QA generation from 3D data accounts for above/below relations in both world and camera reference frames: the former reflects gravity, while the latter corresponds to vertical image orientation. Leveraging these annotations, we design both qualitative and quantitative QA templates grounded in the following spatial concept categories:




Figure 12: Visualization of **bottom**-side free space identification. (a) Top-view occupancy map of the target’s bottom surface on the platform. (b) Projection of candidate points into the image plane. (c) Non-visible points are filtered; the final placement center is indicated.



Figure 13: Visualization of **between**-object free space identification. (a) Top-view occupancy map showing the area between two target objects on the same platform. (b) Projection of sampled points into the image plane. (c) Non-visible points are filtered; the final placement center is indicated.

- 772 1. **Relative Position Relations:** Encompasses spatial relations such as left/right, above/below,
773 front/behind, inside/outside, touching/separated, and near/far. These queries require accurate
774 3D positioning and spatial layout to infer inter-object relationships in physical space.
- 775 2. **Orientation and Rotation Reasoning:** Involves reasoning over face/back direction, ori-
776 entation (horizontal/vertical), and relative angles, using 3D object or camera poses (*e.g.*,
777 orientation vectors, rotation matrices) to infer facing direction or viewpoint shifts.
- 778 3. **Geometric Attribute Comparisons:** Covers attributes like size (big/small), height (tal-
779 l/short), and width (wide/thin). These comparisons rely on true 3D dimensions, mitigating
780 distortions from 2D projections.
- 781 4. **Quantitative Spatial Reasoning:** Involves computing depth, distance, relative angles, and
782 spatial betweenness using precise 3D coordinates and metric reasoning.
- 783 5. **Free Space Reasoning:** Identifies free space above, below, or between objects. As illustrated
784 in Fig. 11, 12, and 13, blue-shaded regions represent unoccupied areas computed from object
785 footprints and platform segmentation. To mitigate overestimation from large bounding boxes,
786 we apply a shrink factor (*e.g.*, 80%) to the projected surfaces for above/below queries.
- 787 6. **Location and Placement Prediction for Spatial Referring:** Involves predicting precise
788 2D coordinates from language descriptions, *e.g.*, “Point to the second chair from the left”
789 — identifying a target object, or “Indicate a free spot to the right of the white box on the
790 second shelf” — selecting a valid placement location. These tasks require accurate 2D-3D
791 projection and fine-grained spatial understanding, forming a vital bridge between visual
792 perception and physical interaction and execution.

793 Building on the structured templates, we design a diverse suite of 3D QA covering spatial reasoning,
794 geometric comparison, viewpoint inference, environmental understanding, and coordinate-level
795 localization. Leveraging QwQ-32B’s powerful capabilities, our pipeline also generates complex
796 reasoning QA pairs that are both structurally diverse and semantically rich.

797 **3D training data visualization.** For specific examples of 3D training data and their visualizations,
798 please refer to Appx. E, which contains detailed sample presentations.

799 A.3 Synthetic Data Generation in the simulator

800 We want to arm our model with multi-step referring capabilities with spatial reasoning. While 2D and
801 3D data enable single-step spatial understanding, they are less scalable for multi-step spatial referring
802 with reasoning. Therefore, we generate synthetic data in the simulator.

803 A.3.1 Indoor Scene Generation

804 **Initial Scene Generation.** We utilize Infinigen [18] to generate a large corpus of indoor
805 scenes. To be specific, we configure the generation process to exclude small objects by setting
806 `compose_indoors.solve_steps_small=0` to avoid pre-existing clutter on target surfaces. This
807 allows us to reserve space for the subsequent placement of our curated 3D assets. This initial step
808 yields over 3k unique indoor scenes.

809 **Scene Filtering.** The generated scenes underwent a rigorous filtering process to ensure their suitability
810 for our downstream tasks. The primary filtering criteria included:

- 811 • **Adequate Tabletop Area:** Selected scenes contain at least one sufficiently large, continuous
812 tabletop surface (*e.g.*, desk, table, counter) suitable for object placement. Scenes with absent
813 or impractically small surfaces are excluded.
- 814 • **Acceptable Lighting Conditions:** Scenes with extreme lighting issues (*e.g.*, darkness,
815 oversaturation, unnatural hues) are discarded to ensure a viable baseline for subsequent
816 lighting adjustments.
- 817 • **Scene Realism and Coherence:** Scenes with severe geometric inconsistencies or implausi-
818 ble layouts are removed to maintain physical plausibility.
- 819 • **Camera Accessibility:** Scenes are required to support feasible camera placement with clear
820 views of the target surfaces. Highly cluttered or confined environments are deprioritized.

821 **Scene Adjustments.** To enhance diversity and control experimental variables, filtered scenes
822 underwent automated modifications:

- 823 • **Lighting Randomization:** Light source intensities (*e.g.*, ceiling lights, lamps) are uniformly
824 scaled within $[0.6I, 1.4I]$, where I denotes the original intensity.
- 825 • **Camera Pose Adjustment:** For each tabletop, camera viewpoints are defined with pitch
826 angles randomly sampled from $[-60^\circ, -30^\circ]$ relative to the tabletop plane, oriented toward
827 the region center.
- 828 • **Camera Height and Distance Variation:** Camera height is uniformly sampled from 0.3–
829 0.8 m above the tabletop. Distance to the target area is adjusted to maintain full visibility,
830 conditioned on surface size and field of view.

831 Some typical cases of scene filtering are shown in Fig. 14.

832 A.3.2 3D Asset Selection and Preparation

833 Our 3D assets are sourced from the Objaverse [19] LVIS dataset and undergo a two-stage filtering
834 process to ensure quality and relevance.

835 **Stage 1: Category Filtering.** We select objects based on LVIS annotations, following categories:

- 836 • Are typically placeable on surfaces.
- 837 • Have a maximum dimension under 1 meter, suitable for tabletop scenarios.

838 **Stage 2: Attribute-based Filtering.** Next, we apply fine-grained filtering using attributes from the
839 Orient300K [20] dataset. Retained assets satisfy the following criteria:

- 840 • **Axis Alignment:** Key features (*e.g.*, edges, handles) align with canonical camera axes.




Figure 14: Cases of Scene Filtering

- **Single Object:** Represents a standalone object, not a scene or object collection.
- **Color Diversity:** Contains colors beyond white or gray.
- **No Ground Plane:** Excludes auxiliary visualization ground planes.
- **High Quality:** Clean, well-constructed geometry without artifacts.
- **Distinguishable Views:** Canonical views (front, back, top, bottom, left, right) exhibit meaningful visual or semantic differences.
- **Reasonable Object:** Represents a common, identifiable object, not an abstract shape or an unidentifiable entity.

This process yields a curated set of over 9k high-quality 3D assets.

Stage 3: Manual Filtering. Due to the suboptimal quality of annotations, we perform manual verification based on the seven defined rules. Since object size estimation relies on LVIS-labeled categories with a $\pm 30\%$ tolerance, we additionally discard instances whose irregular geometry (*e.g.*, entangled cables of a wired mouse) distorts the bounding box and hinders reliable scaling. After filtering, we retain over 3k high-quality 3D assets compliant with our criteria.

Fig.15 shows the screening results for some representative 3D assets, which encompass all the criteria mentioned above.

A.3.3 3D Asset Annotation with LLM

To generate diverse annotations, we leverage OrienText300K’s orientation and caption data, processed via GPT-4o with tailored prompts to extract structured textual attributes for each sample.

Generated Attributes:

- **Orientation Descriptions:** Prepositional phrases indicating the object’s canonical front, salient parts, or intrinsic orientation (*e.g.*, “on the front of ...”, “on the handle side of ...”), suitable for insertion into sentence templates.
- **Color Labels:** A single-word descriptor of the object’s dominant color. If multiple colors are prominent, the attribute is marked as “none” (*e.g.*, “blue”, “none”).
- **Object Labels:** Concise noun phrases specifying the object category (*e.g.*, “coffee mug”, “computer mouse”), usable as subjects or objects in templates.



Figure 15: Cases of 3D assets Filtering

868 • **Category Consistency:** A boolean flag indicating alignment between the object's visual
869 category and its textual description.

870 **LLM Prompt:** The exact prompt used to obtain these annotations is provided in Listing 11.

Listing 11: Prompt of generating brief description, color and orientation preposition.

```
871 Your objective is to generate four distinct labels derived from the
872 provided 3D asset information. Each asset is characterized by the
873 following attributes:
874
875 Category: {category}
876 Detailed description: {description}
877 Direction hint: {direction_hint}
878
879 Based on the information above, you are required to perform
880 the following four tasks:
881
882 Task 1: Consistency Verification
883 Evaluate whether the asset's specified category aligns semantically
884 with its detailed description. Output "True" if consistent,
885 and "False" otherwise.
886
887 Task 2: Object Description Phrase Generation
888 formulate a concise and unambiguous object description phrase. This
889 phrase must encapsulate the primary characteristics of the object,
890 not contain any articles (\eg, "a", "an", "the"), and be
```

```

891 grammatically suitable for use as either the subject or object
892 within a template sentence.
893
894 Task 3: Simplified Directional Phrase Generation
895 Generate a concise and clear simplified directional description phrase
896 based on the provided direction_hint. This phrase must be capable of
897 functioning as a prepositional phrase to describe the relative
898 position of other objects within a template sentence. In the generated
899 phrase, the current object should be represented by "OBJECT".
900
901 Task 4: Color Extraction
902 Extract the dominant color of the asset based on the detailed
903 description, adhering to the following criteria:
904 1. The extracted color must be a single English word in lowercase.
905 2. If the detailed description does not explicitly state a color,
906 return the string "none".
907
908 Please ensure your output strictly adheres to the following JSON
909 format. The output must be a valid JSON object without any
910 supplementary explanations, comments, or introductory/closing
911 remarks.
912 {
913   "consistent": (True/False),
914   "simple_desc": "simplified object description phrase",
915   "simple_dir": "simplified directional description phrase",
916   "color": "the main color"
917 }

```




Figure 16: A Sample of generating object phrases, colors, and orientation descriptions from 3D asset information and orientation hints.

918 A.3.4 Scene Population and Data Rendering

919 The curated 3D assets are then programmatically placed into the filtered and adjusted indoor scenes.
 920 **Asset Placement Strategy:**

- 921 • Assets are placed predominantly on identified tabletop surfaces.
- 922 • To enhance the model’s understanding of object orientation and inter-object relationships,
- 923 placement strategies included:
 - 924 – Increasing the proportion of objects with orientation vectors within the X-Y plane (*e.g.*,
 - 925 laptops, teddy bears, mugs) in many scenes.
 - 926 – Increasing the co-occurrence of objects from the same category that have significantly
 - 927 different features. (*e.g.*, a ceramic cup with a handle and a cola cup).
 - 928 – Ensuring plausible physical arrangements (*e.g.*, no excessive interpenetration, objects
 - 929 placed upright).
- 930 • The number of assets per scene varied from 3 to 9.

931 Rendering Details:

- 932 • **Renderer:** High-fidelity, physically-based images are rendered using the Blender Cycles.
- 933 • **Image Resolution:** Output images are rendered at 960×540 pixels.

- **Render Quality:** The `configure_render_cycles.num_samples` parameter is set to 2048 to achieve high-quality renders while maintaining reasonably controlled noise.

936 A.3.5 Question-Answer Pair Generation

937 Based on the rendered scenes and their corresponding ground truth information (scene graphs, object
938 properties, masks), we generate question-answering (QA) pairs. This process involves two steps:

939 **Step 1: Generating Unique Object Referring Expressions.** For each object in a scene, we formulate
940 unique identifiers or referring expressions based on a combination of its attributes and relationships:

- 941 • **Feature Category:** The semantic class of the object (*e.g.*, “the mug”, “the laptop”).
- 942 • **Color:** The primary color of the object (*e.g.*, “the red mug”).
- 943 • **Left-Right Rank:** Ordinal position from left to right (*e.g.*, “third bottle from left to right”).
- 944 • **Front-Back Rank:** Ordinal position from front to back (*e.g.*, “farthest LEGO minifigure”).
- 945 • **Distance Rank from Anchor:** Ordinal position based on distance from a salient anchor
946 object (*e.g.*, “the closest plate to the blue mug”).
- 947 • **Height Rank:** Ordinal position based on the height (*e.g.*, “the second tallest teddy bear”).

948 These components are combined to create unambiguous references (*e.g.*, “the second red object from
949 the left”). When filling objects into template structures in the subsequent process, we can randomly
950 select from all referring expressions.

951 **Step 2: Applying QA Templates.** Due to the heavy reliance on spatial relationships between objects,
952 we initially developed a template structure to formalize these relationships:

- 953 • **Position:** Location of an object relative to another (*e.g.*, “on the left of the green bottle”).
- 954 • **Orientation:** Questions involving an object’s intrinsic orientation (*e.g.*, “on the handle side
955 of the red mug”).
- 956 • **Distance Queries:** The precise distance (*e.g.*, “0.2 meters to the left of the plate”).
- 957 • **Betweenness:** Identifying an object located between two other objects (*e.g.*, “between the
958 stapler and the telephone”).
- 959 • **Specific Surface Locations:** Locating objects relative to parts of a surface (*e.g.*, “at the far
960 left corner of the table”).

961 QA Template Types:

- 962 • **Locate from Description:** Given a unique referring expression for an object, ask for its
963 location (*e.g.*, “*Give me the position of ...*”).
- 964 • **Identify from Relations:** Provide several spatial relationships an object satisfies and ask to
965 identify the object (*e.g.*, “*Please specify an object on the desktop that satisfies the following*
966 *spatial constraints: ...*”).
- 967 • **Locate Empty Space:** Define a point in an empty area on a surface based on its spatial
968 relationships with surrounding objects, and ask to confirm this empty location (*e.g.*, “*Please*
969 *provide a point in the vacant area on the desktop that simultaneously satisfies the following*
970 *spatial conditions: ...*”).

971 **Generation of Thought Processes (Reasoning Steps):** For each QA pair, a structured thought
972 process or chain of reasoning is also generated. This involved selecting the pertinent pieces of
973 information from the complete scene graph and ground truth data that are necessary to arrive at the
974 answer. This recorded reasoning follows a predefined format.

975 B Implementation Details and Samples of RefSpatial-Bench

976 The *RefSpatial-Bench* benchmark evaluates spatial referring with reasoning in complex 3D indoor
977 scenes. It includes two primary tasks—**Location Prediction** and **Placement Prediction**—as well as
978 an **Unseen** split featuring novel query types absent from existing datasets. Each sample comprises a

manually selected image, a referring caption, and precise mask annotations. The dataset contains 100 samples each for the Location and Placement tasks, and 77 for the Unseen set.

Location Task: Given an indoor scene and a unique referring expression, the model predicts a 2D point indicating the target object. Expressions may reference color, shape, spatial order (e.g., “the second chair from the left”), or spatial anchors.

Placement Task: Given a caption specifying a free space (e.g., “to the right of the white box on the second shelf”), the model predicts a 2D point within that region. Queries often involve complex spatial relations, multiple anchors, hierarchical references, or implied placements.

Unseen Set: This set includes queries with novel spatial reasoning or question types, designed to assess model generalization and compositional reasoning.

Notably, **we introduce a metric termed *reasoning steps (step)* for each text instruction, quantifying the number of anchor objects and their associated spatial relations that effectively constrain the search space.** Specifically, each *step* corresponds to either an explicitly mentioned anchor object or a directional phrase linked to an anchor that greatly reduces ambiguity (e.g., “on the left of”, “above”, “in front of”, “behind”, “between”). We exclude the “viewer” as an anchor and disregard the spatial relation “on”, since it typically refers to an implied surface of an identified anchor, offering minimal disambiguation. Intrinsic attributes of the target (e.g., color, shape, size, or image-relative position such as “the orange box” or “on the right of the image”) also do not count towards *step*.

A higher *step* value indicates increased reasoning complexity, requiring stronger compositional and contextual understanding. Empirically, we find that beyond 5 *steps*, additional qualifiers yield diminishing returns in narrowing the search space. Thus, we cap the *step* value at 5. Instructions with $step \geq 3$ already exhibit substantial spatial complexity. Detailed statistics on *step* distributions and instruction lengths are provided in the Tab. 1. To further show the diversity and reasoning complexity of *RefSpatial-Bench*, we present representative examples from both the Location and Placement tasks. Fig. 33, 34, 35, and 36 show Location queries with varying reasoning step counts, where *RoboRefer* accurately localizes the target object (marked by a blue dot). Similarly, Fig. 37, 38, 39, and 40 show Placement queries involving the identification of free space based on spatial relations. These examples highlight the step-wise complexity of the queries and the effectiveness of *RoboRefer* in addressing challenging spatial referring tasks.

Table 1: Statistics of the *RefSpatial-Bench* benchmark across three splits.

Split	Step	Samples	Avg. Prompt Length
Location	Step 1	30	11.13
	Step 2	38	11.97
	Step 3	32	15.28
	Avg. (All)		12.78
Placement	Step 2	43	15.47
	Step 3	28	16.07
	Step 4	22	22.68
	Step 5	7	22.71
	Avg. (All)		17.68
Unseen	Step 2	29	17.41
	Step 3	26	17.46
	Step 4	17	24.71
	Step 5	5	23.8
	Avg. (All)		19.45

C Implementation Details for *RoboRefer*

C.1 Architecture

We adopt NVILA [21] as base model, including a visual encoder, an LLM, and a multimodal projector.

Table 2: Details about the training datasets used in the SFT and RFT stages. D.A. and S.U.E denote the Depth Alignment and Spatial Understanding Enhancement step in the SFT stage, respectively.

Stage	Categories	Datasets
SFT (D.A)	Spatial	<i>RefSpatial</i> (RGB-D)
SFT (S.U.E)	Spatial General REC	<i>RefSpatial</i> (RGB), <i>RefSpatial</i> (RGB-D), SAT [26], EmbSpatial [27] COCO [28], GQA [29], OCR-VQA [30], TextVQA [31], VG [32], LRV [24] RefCOCO+/g [25]
RFT	Spatial	<i>RefSpatial</i> (RGB-D) w/ Reasoning Processing

Visual Encoder. We use the same image encoder as siglip-so400m-patch14-448 [22] from NVILA [21], supporting 448×448 resolution for richer visual details. Rather than simply resizing the image to a fixed resolution and producing the same number of tokens, this image encoder processes inputs at dynamic resolutions, yielding more visual tokens from higher-resolution images via finer patch division. This enables fine-grained vision-language understanding, crucial for tasks like point prediction that require detailed perception beyond VQA. We further incorporate a dedicated depth encoder, structurally mirroring the image encoder and initialized with its weights. It encodes relative depth maps as special images, providing spatial cues to enhance 3D understanding.

Large Language Model. We adopt the Qwen2 LLM backbone from NVILA [21], which has been fully fine-tuned with extensive data during supervised training. This endows the model with rich visual knowledge, facilitating downstream 3D spatial understanding and reasoning tasks.

Multi-Modal Projector. To align multi-modal representations (*e.g.*, image to language, depth to language), we use linear connectors, the same as NVILA [21], which is better than Q-Former, to allow the LLM to focus on visual understanding and improve generalization. Separate connectors for image and depth embeddings ensure modality-specific processing and prevent cross-modal interference.

C.2 Training Data

Here we highlight the training data used at each stage, including the number of samples per dataset and the overall total. See Tab. 2 for details.

SFT stage. Specifically, in the first step of the SFT stage, *i.e.*, depth alignment, we train a depth projector to align depth and language space using the *RefSpatial* (RGB-D) dataset with 2.5M samples. To increase training efficiency, we slice multi-turn conversations (up to 15 turns per sample), yielding 3.4M samples post-processing to train our model. In the second step, *i.e.*, spatial understanding enhancement via full-parameter fine-tuning—we use both *RefSpatial* (RGB) and *RefSpatial* (RGB-D) datasets, yielding 6.8M samples after slicing. To further improve instruction-following and referring capabilities, we incorporate auxiliary datasets: 965k samples from instruction-tuned data (LLaVA-1.5 [23], LRV [24]), 321k from referring datasets (RefCOCO+/g [25]), 176k from SAT [26] benchmark training sets, and 127k from EmbSpatial [27] benchmark training sets. These additions help bridge distribution gaps between *RefSpatial* and benchmark-style queries. After slicing, the total number of samples used in this stage reaches 8.5M post-slicing.

RFT Stage. In the RFT stage, we train the model using *RefSpatial* data annotated with detailed reasoning processes, including key intermediate steps and final answers. To ensure both training efficiency and effective learning, we use moderately difficult samples (typically involving three reasoning steps), resulting in a 100k-sample dataset.

C.3 SFT Training Details

We formulate the SFT training stage as follows: given a dataset \mathcal{D} consisting of samples in the form of triplets $(\mathcal{O}, \mathcal{Q}, \mathcal{A})$, where \mathcal{O} is a sensor image (either RGB or RGB-D), \mathcal{Q} is a textual question, and \mathcal{A} is the corresponding answer. The answer \mathcal{A} may be a direct response (*e.g.*, a point coordinate) or include intermediate reasoning steps (*e.g.*, perceptual results followed by the final answer). The training objective is to maximize the likelihood of generating the answer given the input pair $(\mathcal{Q}, \mathcal{A})$:

$$\mathcal{L}_{\text{SFT}} = -\mathbb{E}_{(\mathcal{O}, \mathcal{Q}, \mathcal{A}) \sim \mathcal{D}} \sum_{t=1}^T \log \pi_{\theta}(y_t \mid \mathcal{O}, \mathcal{Q}, y_{<t}), \quad (1)$$

where π_{θ} is the model’s token distribution. The output model π_{SFT} serves as the initialization for the next RFT stage, ensuring a robust foundation for reinforcement learning.

To be specific, our SFT consists of two steps. In the first step, depth alignment, only the depth projector is updated by using the *RefSpatial* (RGB-D). We employ a maximum learning rate of 1e-4, a weight decay of 0, and a warm-up ratio of 0.03. The 2B variant is trained with a batch size of 7 per GPU, and the 8B variant with 3, both for one epoch. In the second step of spatial understanding enhancement, we fine-tune all model parameters using the datasets described in Sec. C.2. Training is conducted for one epoch with a maximum learning rate of 5e-5. We use a batch size of 6 per GPU for the 2B model and 2 for the 8B model. Other hyperparameters follow those in the first step. For more details, please refer to NVILA [21] settings during alignment and SFT.

C.4 RFT Training Details

During the RFT stage, we refine π_{SFT} via GRPO [33], a reinforcement learning method designed for efficiency and scalability. Unlike PPO [34], which relies on a costly value network, GRPO estimates relative advantages by comparing intra-group rewards, reducing computation, and simplifying optimization. This makes it well-suited for reasoning-intensive spatial referring tasks. In detail, we modify R1-V [35] to support our 3D-aware architecture. Training is conducted for two epochs with a batch size of 1 per GPU and 8 outputs in GRPO. For details about hyperparameters, see R1-V³.

C.4.1 Sampling Action Groups

Given an input state $s = (\mathcal{O}, \mathcal{Q})$, where \mathcal{O} denotes the visual encoding of the RGB or RGB-D observation and \mathcal{Q} the textual encoding of the question, GRPO samples a set of actions $\{a_1, a_2, \dots, a_N\}$ from the current policy π_{θ} , initialized from π_{SFT} . The sampling process is:

$$a_i \sim \pi_{\theta}(a \mid \mathcal{O}, \mathcal{Q}), \quad \text{for } i = 1, 2, \dots, N \quad (2)$$

This strategy ensures diverse responses, promoting exploration and preventing premature convergence.

C.4.2 Reward Design and Policy Update

Each sampled action a_i is assigned a reward $R(a_i)$ based on verifiable criteria, yielding a reward set r_1, r_2, \dots, r_N . For spatial referring tasks, $R(a_i)$ integrates two outcome-based and our proposed two process-based components. The outcome-based reward functions are defined as follows:

Outcome Format Reward R_{OF} . This component ensures structured and interpretable outputs by requiring the model to a predefined format: reasoning within “<think>...</think>” and the final answer in “<answer>...</answer>”. A reward is assigned 1 for strict compliance, 0 otherwise.

Point L1 Reward R_P . This component evaluates the accuracy of the model’s final point prediction by comparing it with the ground truth from the annotations of *RefSpatial*. Following the criterion inspired by Seg-zero [36], a stricter reward of 1 is assigned if the L1 distance between the predicted and ground-truth points is within 50 pixels; otherwise, the reward is 0.

Notably, most process-based rewards depend on a Process Reward Model (PRM), typically a fine-tuned LLM or VLM tasked with providing feedback. However, applying such an approach in our setting presents two main challenges. (1) LLMs cannot process images, making it impossible to determine whether predicted coordinates match the target object. (2) Although VLMs integrate visual and textual information, prior work [37] has shown they may lack precise visual understanding when dealing with textual coordinates. Since the correct assessment of predicted coordinates is paramount for reward assignment, additional or specialized methods are needed to ensure reliable feedback.

To address this issue, we propose a rule-based process reward for spatial referring that obviates the need for a Process Reward Model. Our approach directly evaluates key intermediate perceptual

³<https://github.com/Deep-Agent/R1-V>

steps using the ground-truth step-wise annotations provided in *RefSpatial*. This contrasts with most existing methods on process-based rewards [38, 39], which emphasize strictly sequential reasoning and rely on a PRM for evaluation. **In contrast, our method employs metric-sensitive rule-based process reward functions to assess intermediate perceptual results in an order-invariant manner.** Our key insight lies in two aspects: (1) **Metric-sensitivity**: Different spatial attributes require distinct metrics due to inherent differences in their representations (*e.g.*, points for positions, vectors for orientations). (2) **Order-invariance**: The reasoning process in spatial referring is not strictly sequential; for instance, identifying the position of the keyboard or the mouse first does not affect the final interpretation of “the free area between the keyboard and the mouse”.

In detail, we have two process-based reward functions:

Process Format Reward R_{PF} . Similar to the Outcome Format Reward strategy, this component enforces a structured and interpretable reasoning process, thereby facilitating accurate reward computation. In particular, the model is required to produce outputs in the following format:

$$[\text{Perception Type}] [\text{Target Object}]: [\text{Value}] \quad (3)$$

where “Perception Type” must be one of three categories: “Position”, “Orientation”, or “Size”. The “Target Object” corresponds to a uniquely identifiable entity (*e.g.*, “the second largest cup” or “the second large cup from large to small”). The “Value” depends on the selected “Perception Type”:

- For “Position”, the value should be a normalized 2D coordinate of the form $[(x, y)]$, where both x and y lie in the interval $[0, 1]$, rounded to three decimal places.
- For “Orientation”, the value is a 3D unit vector (x, y, z) .
- For “Size”, the value represents a scalar measured in meters.

Below are examples to illustrate the expected format:

- [Position] [the second largest cup]: [(0.245, 0.147)]
- [Orientation] [the handle of the second largest cup]: (1.000, 0.000, 0.000)
- [Size] [the second largest cup]: 0.12

Accuracy Reward R_{Acc} . The reward is computed only for steps annotated as key steps in *RefSpatial*. In detail, we use regex matching to determine whether the “Target Object” in the current process format appears in the key-step annotations. If not, the step receives no reward. Since the model has already undergone a cold-start phase in SFT, it can interpret instructions and identify relevant target objects. Thus, a failed match implies that the model cannot accurately refer to the object linguistically, and no reward is assigned. For each perception type, we apply a specific metric to compute the reward: (1) “Position”: If the L1 distance between the predicted point and the ground truth is below 50 pixels, the reward is 1; otherwise, 0. (2) “Orientation”: If the cosine similarity between the predicted and ground-truth vectors exceeds 0.8, the reward is 1; otherwise, 0. (3) “Size”: If the predicted value falls within $\pm 15\%$ of the ground truth, the reward is 1; otherwise, 0.

We prioritize the correctness of the final outcome over intermediate steps. To prevent reward accumulation from multi-step processes, we scale the process reward by 0.25. The final reward function is defined as:

$$r_i = R_{OF}(a_i) + R_P(a_i) + R_{PF}(a_i) + R_{Acc}(a_i) \quad (4)$$

By normalizing the rewards within the sampled group, we obtain the set of relative advantages $\{A_1, A_2, \dots, A_N\}$ defined as

$$A_i = \frac{r_i - \text{mean}(\{r_j\})}{\text{std}(\{r_j\})}, \quad (5)$$

which measures how each reward compares to the mean in units of standard deviation. We then update the policy based on these advantages, reinforcing actions with higher relative advantages while reducing the likelihood of those deemed less effective. To ensure stable reinforcement learning, the update is further constrained by minimizing the KL divergence between the updated policy and its reference counterpart, thereby promoting incremental and controlled policy adjustments.

1137 D Experimental Setting and Details

1138 D.1 Experiments Compute Resources

1139 We conduct experiments on an A100 GPU cluster, with each node equipped with 8 GPUs.

1140 **2D Web Data Coarse Filtering.** We perform the initial coarse filtering of 1.7M OpenImages using
1141 SigLIP2. The process runs on 1 node and takes 8.5 hours and yields 933k high-quality samples.

1142 **2D Web Data Fine-grained Filtering.** We further filter 933K samples using Qwen 2.5-VL 7B to
1143 ensure high visual quality and spatially relevant QA pairs. The process is conducted on 1 node (two
1144 models per GPU) over 2.5 days, yielding approximately 845k high-quality samples.

1145 **Pseudo-3D Scene Graphs Construction.** We construct pseudo-3D scene graphs for 845k samples
1146 using 3 nodes, requiring 10 hours for depth estimation and another 10 hours for camera parameter
1147 extraction, segmentation masks, and point cloud generation. Additionally, we generate object-level
1148 captions for all instances using 4 nodes over 18 hours by using Qwen 2.5-VL.

1149 **Reasoning QA Generation from 2D data source.** To enrich factual statements with contextual
1150 scenarios, we employ QwQ-32B to construct reasoning QA, utilizing 4 nodes over 3.75 days.

1151 **3D Data Filtering and Scene Graphs Construction.** Given the limited amount of 3D data (100k)
1152 and the availability of precise annotations, only 2D bounding box bidirectional matching is required.
1153 This process is completed in 3 hours using 1 node.

1154 **Reasoning QA Generation from 3D data source.** To enrich factual statements with contextual
1155 scenarios, we employ QwQ-32B to construct reasoning QA, utilizing 4 nodes over 1.5 days.

1156 **Synthetic Data Generation in Simulator.** We use 4 RTX 4090 GPUs to generate data for one week.

1157 **Depth Alignment in SFT.** The process is conducted on 10 nodes over 12 hours for 2B variants and 8
1158 nodes over 40 hours for 8B variants. Both variants training use ZeRO-3.

1159 **Spatial Understanding Enhancement in SFT.** The process is conducted on 10 nodes over 2 days
1160 for 2B variants and 10 nodes over nearly 1 week for 8B variants. Both variants training use ZeRO-3.

1161 **Spatial Referring in RFT.** The process is conducted on 1 node over 3 days for 2B variants. However,
1162 our model is over twice as slow as other Qwen 2/2.5-VL-based methods [35, 40], mainly because
1163 they process only a single RGB image during training and can leverage vLLM for group inference
1164 acceleration. In contrast, our method requires RGB-D inputs and modifies the original NVILA
1165 architecture, making it incompatible with vLLM or SGLang acceleration.

1166 D.2 Spatial Understanding Benchmarks

1167 We evaluate several public *single-step spatial understanding* benchmarks, including CV-Bench [41]
1168 (2D Spatial Relation, 3D Depth Order, 3D Distance), the BLINK [42] validation set (Spatial Relation,
1169 Relative Depth), RoboSpatial [13] (configuration), SAT [26], and EmbSpatial [27], following their
1170 official evaluation protocols. We exclude non-spatial tasks from our evaluation, such as 2D Counting
1171 in CV-Bench and Art Style or IQ Test in BLINK. Since all these benchmarks are multiple-choice
1172 tasks, we report accuracy as the evaluation metric.

1173 We compare 3 categories of models: (1) proprietary VLMs, such as Gemini-2.5-Pro [43], which
1174 show strong spatial perception, as shown in the Gemini-Robotics [44] paper; (2) open-source VLMs
1175 trained on general VQA datasets; and (3) spatially specialized models trained on spatially relevant
1176 datasets, offering basic spatial understanding.

1177 D.3 Spatial Referring Benchmarks

1178 We evaluate three recent robotic referring benchmarks—RoboRefIt [45], Where2Place [46], and
1179 RoboSpatial [13], all limited to 2 reasoning steps. Specifically, RoboRefIt concentrates on object lo-
1180 cation referring by leveraging object attributes and spatial relations with anchor objects. Where2Place
1181 further explores how to place objects relative to anchor objects in the camera’s coordinate frame.
1182 RoboSpatial builds on this idea, considering the same form of placement relative to anchor objects,
1183 but many samples in benchmarks consider the object-centric coordinate frame. We also evaluate
1184 more complex multi-step spatial referring on *RefSpatial-Bench*, a challenging benchmark based on

1185 real-world cluttered scenes, as introduced in Appx. B. For all these benchmarks, we use the same
1186 evaluation protocol: we first compute the proportion of predicted points that fall within the ground
1187 truth mask for each sample, and then average the results across all samples to obtain the success rate.

1188 We compare two main categories of models: (1) Proprietary models (*i.e.*, Gemini-2.5-Pro) with
1189 strong spatial referring capabilities, and (2) spatially specialized models trained on spatially relevant
1190 datasets, exhibiting basic spatial referring abilities.

1191 D.4 Simulation Evaluation

1192 We use the same evaluation protocol of Open6DOR V2 introduced in SoFar [20], following the
1193 official repository⁴. Specifically, we only test the position track as this work focuses on location
1194 and placement via spatial referring rather than executing 6DOF manipulation tasks. Notably, we
1195 find that our model achieves nearly 100% success in the perception stage (*i.e.*, determining location
1196 and placement), with failures primarily attributed to motion planning errors such as IK failures or
1197 collision-prone trajectories. We show more demonstrations of simulation evaluation in Appx. E.

1198 D.5 Real-world Evaluation

1199 D.5.1 UR5 Manipulation

1200 We show two demos for UR5⁵ Manipulation: human disturbance and voice interruption. In the human
1201 disturbance case, *RoboRefer* runs at 2.5Hz. Significant shifts in predicted 2D coordinates trigger
1202 motion interruption and re-planning. In the voice interruption case, incoming speech commands are
1203 continuously monitored. Upon detection, the current task is halted. We use the Whisper [47] ASR
1204 model to transcribe speech, which *RoboRefer* processes into new 2D coordinates for task redirection.

1205 For grasping, the 2D coordinates are fed into SAM2 [11] to generate a segmentation mask, which
1206 filters the target object’s point cloud from the scene captured by a third-person Intel RealSense L515⁶
1207 depth camera. The extracted point cloud is input to AnyGrasp [48] to predict a grasp pose in the
1208 camera coordinate frame. Using an eye-to-hand calibration method, the grasp pose is transformed
1209 into the UR5 robot’s base frame for execution.

1210 For placement, *RoboRefer* predicts the 2D placement point, which is converted to 3D coordinates
1211 using the camera’s intrinsic parameters and depth data. The 3D point is then transformed into the
1212 robot’s coordinate system to guide the placement action.

1213 D.5.2 G1 Humanoid Mobile Manipulation

1214 For grasping, we employ a head-mounted Intel RealSense D435⁷ on the Unitree G1 humanoid to
1215 capture RGB-D images, which are processed by *RoboRefer* to extract 2D target coordinates. These
1216 coordinates guide SAM2 [11] to generate a segmentation mask, which filters the third-person D435
1217 point cloud to isolate the target object. The filtered point cloud is then passed to AnyGrasp [48] to
1218 predict a grasp pose in the third-person frame, which is transformed to the robot’s base frame using
1219 known camera-to-robot calibration.

1220 For navigation, the chest-mounted L515 camera continuously captures images used by *RoboRefer* to
1221 detect nearby landmarks (*e.g.*, a table near the robot). The resulting 2D locations, combined with
1222 depth and intrinsics, are projected into 3D world coordinates and integrated into a global map via
1223 FAST_LIO_LOCALIZATION_HUMANOID⁸ for SLAM-based navigation, as shown in Fig 17, 18.

1224 For placement, the head-mounted D435 captures images processed by *RoboRefer* to localize the
1225 target placement region. The corresponding 3D coordinates, computed from depth and intrinsics, are
1226 transformed into the robot’s base frame for accurate placement execution.

⁴https://github.com/Zhangwenyao1/Open6DOR_V2_Execution

⁵<https://www.universal-robots.com/products/ur5e/>

⁶<https://www.intelrealsense.com/lidar-camera-l515/>

⁷<https://www.intelrealsense.com/depth-camera-d435/>

⁸https://github.com/deeplint/FAST_LIO_LOCALIZATION_HUMANOID/tree/main

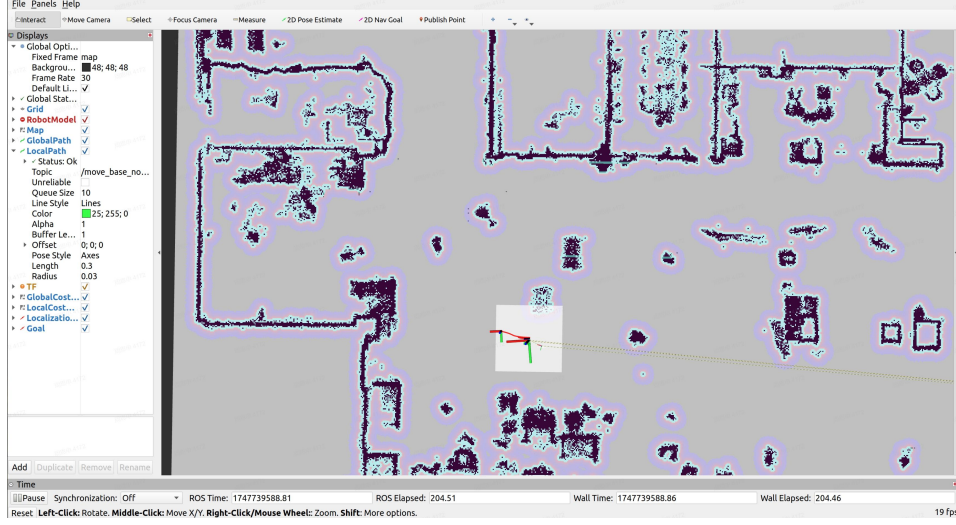


Figure 17: Map Visualization (RViz).

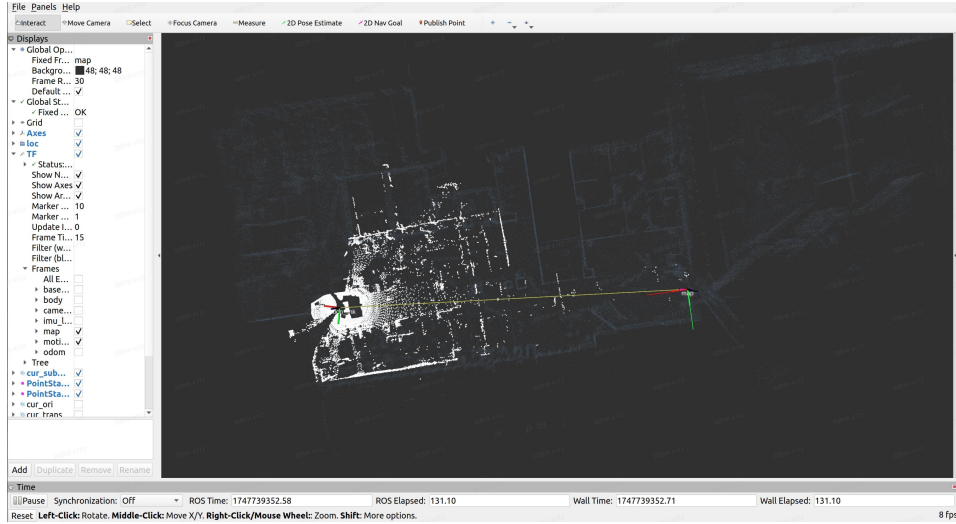


Figure 18: SLAM Navigation (RViz).

1227 D.6 More Ablation Studies

1228 We conduct additional ablation studies to identify which design choices enhance the performance.

1229 **RefSpatial RGB and RGB-D combination for SFT training encourages the image encoder to**
 1230 **learn spatial understanding beyond depth cues.** In Tab. 3, incorporating both RGB and RGB-D
 1231 data from *RefSpatial* in the second stage of SFT training effectively enhances the image encoder’s
 1232 spatial understanding. In contrast, training solely with RGB-D may lead to over-reliance on the depth
 1233 encoder, limiting the image encoder’s ability to learn spatial cues from RGB images alone.

1234 E More Demonstrations

1235 **Visualization of RefSpatial.** We present dataset examples in Fig 20, 21, 22, 23, 24, 25, 26, 27, 28,
 1236 29, 30, 31, which cover 31 distinct types of spatial relationships.

1237 **Visualization of Simulation Evaluation.** We present example rollouts of *RoboRefer* in Fig. 32.

Table 3: Ablation on the combination of RGB/RGB-D from *RefSpatial* for SFT training (*i.e.*, spatial understanding enhancement). Top-1/Top-2 accuracies are represented using **bold text**, and underlines.

Method	Input	CV-Bench [41]			BLINK _{val} [42]	
		2D-Relation	3D-Depth	3D-Distance	2D-Relation	3D-Depth
Only RGB-D from RefSpatial						
RoboRefer-2B-SFT	RGB	87.69	86.83	82.50	79.02	81.45
Combination of RGB and RGB-D from RefSpatial						
RoboRefer-2B-SFT	RGB	<u>96.15</u>	<u>95.83</u>	<u>90.67</u>	<u>83.92</u>	<u>88.71</u>
RoboRefer-2B-SFT	RGB-D	96.31	97.17	90.83	87.41	91.13

1238 **Visualization of *RefSpatial*-Bench.** We present examples of location in Fig. 33, 34, 35, 36 and
1239 placement in Fig. 37, 38, 39, 40 with *RoboRefer* predictions.

1240 **Visualization of Simulator.** We present example rollouts with *RoboRefer* predictions in Fig. 32.

1241 **Visualization of Real-world Evaluation.** We present examples in Fig. 41, 42, 43.

1242 F More Discussion on Limitations and Future Work

1243 Despite achieving promising results, our model still has limitations. In particular, it relies on precise
1244 textual descriptions to pinpoint specific object locations and placement targets, including accurate
1245 references to anchor objects. However, in practical real-world robotics scenarios, human instructions
1246 are often ambiguous. While such utterances may still imply a unique location, resolving them
1247 typically requires sophisticated visual-linguistic reasoning and a process of elimination grounded in
1248 human prior knowledge, capabilities that challenge our current model.

1249 We show two representative but interesting examples that highlight the need for human intent
1250 understanding and visual-linguistic reasoning: **(1) Probabilistic Preference.** As shown in Fig. 19,
1251 humans may refer to a sushi plate as “*pick the one facing the drink*”. In the depicted scene, there
1252 are four drink bottles, yet only the two in the middle align with the second sushi plate from the left
1253 in the farthest row; the leftmost drink aligns with the first plate, and the rightmost with the third.
1254 Despite this ambiguity, people often judge the second plate to be the intended reference due to its
1255 higher likelihood of being aligned with two out of the four drinks, reflecting a probabilistic bias in
1256 interpretation. **(2) Spatial Compatibility.** As shown in Fig. 19, a user might instruct, “*Place another*
1257 *sushi between the plate and the soy sauce dish*”. Although multiple plate–soy sauce pairs exist, only
1258 the pair closest to the observer affords sufficient physical space to place another sushi plate. Thus,
1259 implicit spatial feasibility guides the correct interpretation, even without explicit constraints.

1260 Our model struggles with these cases because *RefSpatial* lacks multi-step referring data that embeds
1261 human priors and intent understanding. As *RefSpatial* is procedurally generated, incorporating
1262 such characteristics at scale remains challenging. Future work may explore procedural synthesis
1263 of intent-aware data or improve model performance via co-training with datasets supporting intent
1264 comprehension, such as PixMo-Points [49].

1265 G Broader Impacts

1266 *RoboRefer* can serve as a versatile visual assistant with advanced spatial understanding and reasoning.
1267 Due to its integration with large language models (LLMs), it inherits both potential benefits and
1268 risks, similar to other VLMs, such as output hallucinations, biases from base models, and heightened
1269 energy consumption associated with model upscaling. Beyond these considerations, *RoboRefer*
1270 can also function as a high-level planner with spatial referring abilities, guiding robots in tasks like
1271 manipulation and navigation. While such capabilities substantially enhance robotic control, they also
1272 pose safety challenges when combined with existing control policies.

1273 Despite these concerns, releasing *RoboRefer* to the broader research community would be highly
1274 advantageous. Open access would foster continued development and refinement of spatial referring



Figure 19: Discussion about the limitation of human prior knowledge and intent understanding.

1275 with reasoning, ultimately benefiting diverse robotics platforms (*e.g.*, robotic arms, humanoids)
 1276 performing various tasks (*e.g.*, manipulation, navigation).

1277 H Licenses

1278 (1) 2D web image data: OpenImages [2] is released under Apache License 2.0.

1279 (2) 3D embodied video data: CA-1M [1] is released under CC-by-NC-ND.

1280 (3) Procedural scene generation: Infinigen [18] is released under BSD 3-Clause License.

1281 (4) 3D digital assets: objaverse-xl [19] is released under Apache License 2.0.



Figure 20

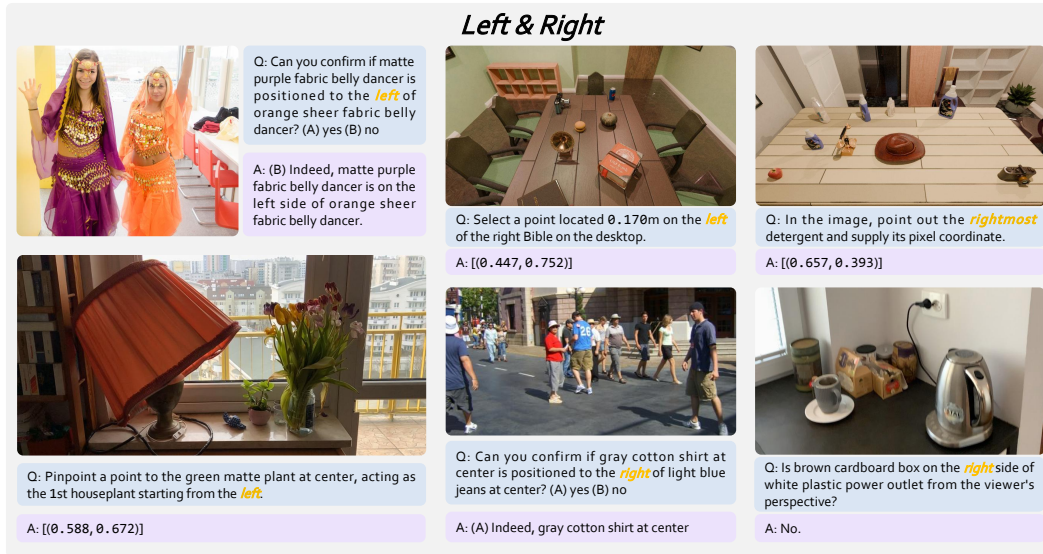


Figure 21

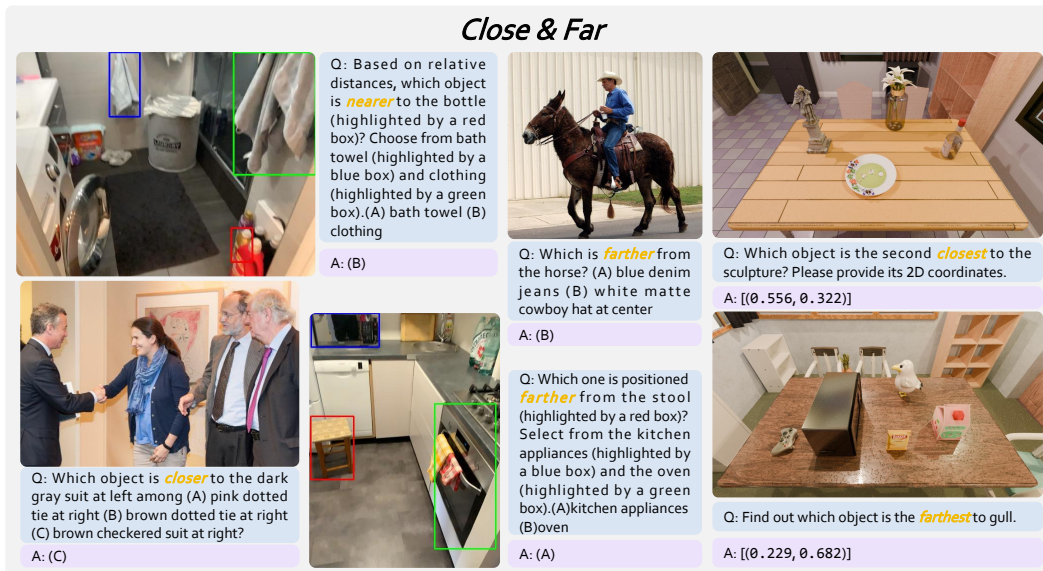


Figure 22



Figure 23

Above & Below (Image)



Q: From the image's perspective, would you say gold textured tray at center is positioned **over** glossy white cabinet at upper center?

A: Yes. Gold textured tray at center appears below white cabinet in the image.



Q: From the image's perspective, is wooden chair at right **below** gold brass saxophone at center? (A) no (B) yes

A: (B) yes.



Q: From the image's perspective, is black iron streetlight at upper left **above** white glossy car at right?

A: Yes. Gold textured tray at center appears below glossy white cabinet at upper center in the image.



Q: From the image's perspective, is wooden chair at right **below** gold brass saxophone at center? (A) no (B) yes

A: (B) Indeed, light wood/wooden chair at right is positioned under gold brass saxophone at center.

Figure 24

Between



Q: A white paper towel can be found **between** translucent plastic bag at center and transparent glass jar at center, which is the second glass jar from left to right, in the front-back direction. Please provide its 2D coordinates.


A: [[0.526, 0.691]]



Q: Please mark a point on the desktop that lies **between** the leftmost object and the beer bottle, and is located on the side with the handle of the nearer cup.


A: [[0.616, 0.805]]

Rotation (Horizon & Vertical)



Q: To directly face white glossy coffee maker at center, should I **rotate upward or downward**?

A: You should rotate down by 18.701 degrees.




Q: If I want to look at the white matte radiator at lower left, should I **rotate to the left or to the right**?

A: It's best to rotate left around 19.945 degrees.


Figure 25

Free Space



Q: In this image, there is a matte table at center. Please find a location within the **vacant space** on the white matte table at center.

A: [[0.539, 0.772]]



Q: Please choose a spot located in the **free space** between the red camera and the right camera.

A: [[0.616, 0.805]]

Corner & Edge



Q: Select a point that is positioned in the free space at the right **edge** on the desktop.

A: [[0.768, 0.655]]



Q: Please denote the thermos bottle positioned at the far left **corner** area on the desktop.

A: [[0.352, 0.285]]

Figure 26

Tall & Short



Q: You are organizing a display of tobacco tins and need to stack them vertically. Given that you have a space that can only accommodate items up to the height of the **taller** object between the red matte tin can at left and the matte plastic peterson university flake at center, which tin should you choose to ensure it fits?

A: The matte plastic peterson university flake at center is taller than the red matte tin can at left, so choosing the matte plastic peterson university flake at center will ensure it fits within the space designated for the taller object.



Q: Where is the **tallest** vodka bottle with a red label located? Please provide its 2D coordinates.

A: $[(0.616, 0.258)]$



Q: Give me position of the **shorter** bottle with a blue cap. Please provide its 2D coordinates.

A: $[(0.383, 0.398)]$



Q: Which object is **taller**, the stone fireplace at center or the wooden photo frame at on the left side?

A: The stone fireplace at center is taller than wooden photo frame.



Q: Is transparent glass jar at center, which is the second glass jar from left to right, **shorter** than transparent glass jar at lower left? (A) no (B) yes

A: (B) Yes.




Q: Which of these two, plastic recycling bin at right or dark blue quilted mattress at lower left, stands **shorter**?

A: With less height is the dark blue quilted mattress at lower left.


Figure 27

Front & Behind




Q: Which object is **closer to the camera**, the white fabric trash bin at center or the wooden photo frame at left?

A: The wooden photo frame at left.




Q: Is the position of metallic chair more **distant** than that of white radiator at upper right? (A) yes (B) no

A: (B) No.




Q: Give me position of the **nearest** blue orange juice. Please provide its 2D coordinates.

A: $[(0.415, 0.523)]$




Q: Among these objects, which one is **nearest** to the camera? (A) blue fabric shirt at center (B) dark blue t-shirt at left (C) light blue jeans at center (D) black cotton shirt at center

A: (B) dark blue t-shirt at left



Q: Point out the lamp, which is found **behind** the light gray fabric couch at lower left.

A: $[(0.127, 0.734)]$



Q: Provide the coordinates of a location in the free space on the desktop that fulfills both of the following spatial criteria: 1. on the left side of the rightmost pepsi can; 2. **behind** the duck.

A: $[(0.476, 0.432)]$

Figure 28



Figure 29

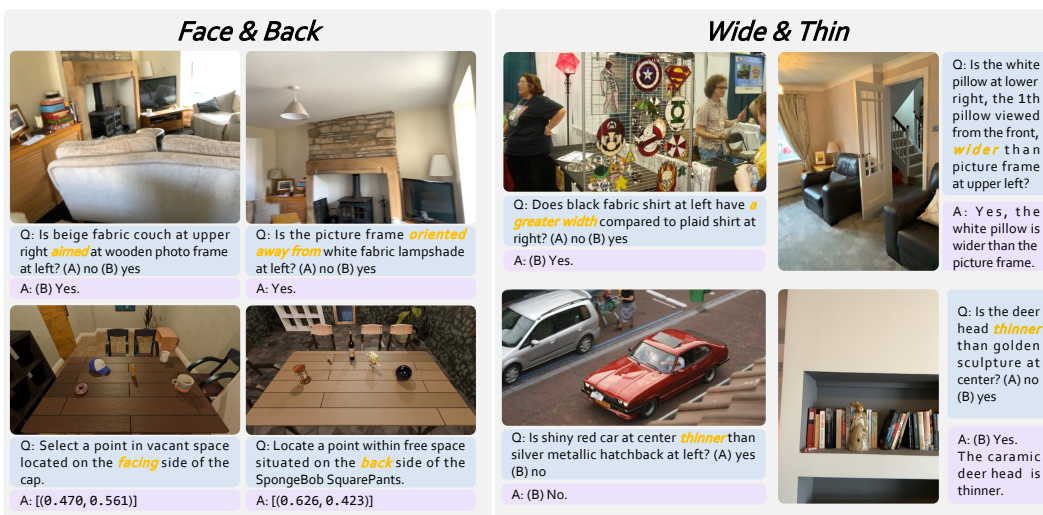



Figure 30



Figure 31



Figure 32: Visualization of RoboRefer’s prediction (blue point) on Open6DOR V2 [20] benchmark.




Reasoning Step = 1

Please point out the **box with the person logo** in the picture.



Reasoning Step = 1

Please point out the **monitor closest to the viewer** in the image.




Reasoning Step = 1

Please point out the **third chair from the left to the right**.




Reasoning Step = 2

Please point to the **top piece of paper** on the **white table**.



Reasoning Step = 1

Please point to the **farthest white cabinet** in the picture.




Reasoning Step = 3

Please point out the **black framed painting on the right of the lamp**.



Reasoning Step = 2

Please point out the **third object from left to right on the closest platform**.



Reasoning Step = 1

Please point to the **wooden plate on the far left of the picture**.



Reasoning Step = 1

Please point out the **sofa on the right side of the picture** that is **closest to the viewer**.



Reasoning Step = 1

Please point to the **rightmost box** in the picture.



Reasoning Step = 1

Please point out the **second silver box from left to right** in the picture.



Reasoning Step = 2


Please point out the **painting hanging on the wall**.

Figure 33: Some Location Examples. The model is asked to identify the object referred to by a prompt. The blue point shows the RoboRefer's prediction (all correct).




Reasoning Step = 2

Please point out the **blue toothbrush** farthest from the faucet.




Reasoning Step = 2

Please point out the **card** closest to the wooden door.



Reasoning Step = 2

Please point to the **third card** from right to left on the cabinet.



Reasoning Step = 2

Please point to the **pillow** closest to the remote controller.



Reasoning Step = 3

Please point out the **object** between the white box and the farthest black pot in the picture.




Reasoning Step = 2

Please point out the **vase** closest to the TV.



Reasoning Step = 3

Please point out the **leftmost black object** on the same platform as the micro-wave oven.



Reasoning Step = 2

Please point out the **orange box** on the white table on the left.



Reasoning Step = 1


Please point out the **yellow five-pointed star** in the picture.

Figure 34: Some Location Examples. The model is asked to identify the object referred to by a prompt. The blue point shows the RoboRefer's prediction (all correct).



Reasoning Step = 2

Please point out the **tallest bottle** on the **black table**.



Reasoning Step = 1

Please point out the **farthest chair** from the viewer.



Reasoning Step = 2

Please point out the **silver bottle** on the **right edge** of the **sink**.




Reasoning Step = 3

Please point out the **object** under the **table**.




Reasoning Step = 2

Please point to the **rightmost blue box** on the **refrigerator**.



Reasoning Step = 2

Please point to the **pillow** closest to the **right nightstand**.




Reasoning Step = 3

Please point out the **cardboard box** under the **bed** which is **closest box to the viewer**.



Reasoning Step = 2

Please point out the **blue object** on the **table** in the picture.



Reasoning Step = 2

Please point to the **white box** closest to the **remote control**.

Figure 35: Some Location Examples. The model is asked to identify the object referred to by a prompt. The blue point shows the RoboRefer's prediction (all correct).



Reasoning Step = 3
Please point out the **second object** from right to left on the platform with the **banana**.




Reasoning Step = 1
Please point out the **farthest chair** from the viewer.



Reasoning Step = 1
Please point out the **orange box** in the picture.




Reasoning Step = 1
Please point out the **remote control** on the right side of the picture.



Reasoning Step = 1
Please point to the **stool** which is the **third stool** from the front.



Reasoning Step = 3
Please point out the **white object** on the **first shelf** of the **cabinet**.



Reasoning Step = 3
Please point out the **plant** on the **far left** of the **second shelf** of the **cabinet**.



Reasoning Step = 3
Please point out the **blue object** on the **third level** of the **wooden shelf**.



Reasoning Step = 2
Please point out the **paper tube** closest to the viewer.

Figure 36: Some Location Examples. The model is asked to identify the object referred to by a prompt. The blue point shows the RoboRefer's prediction (all correct).



Reasoning Step = 4

Please point out the **free space** between the **black water bottle**, the **pot lid**, and the **scissors**.




Reasoning Step = 4

Please point out the **free space** between the **black cloth box** to the **bottom-right** of the **monitor** and the **keyboard**.




Reasoning Step = 3

Please point out the **free spot** halfway between the **blue cup** and the **green cup** to place another identical one.




Reasoning Step = 3

Please point out the **free space** midway between the **first** and **second green cups** from the left.



Reasoning Step = 4

Please point out the **free spot** to the **left** of the **can**, where an object of the same size can be placed at an equal distance between the **can** and the **bag**.



Reasoning Step = 2

Please point out the **free space** in the **direction** the **logo** of the **closest white bottle** to the viewer.



Reasoning Step = 4

Please point out the **free spot**, equidistant from both the **blue bowl** and the **red bowl**, and between them, where a new, similar-sized bowl can be placed.



Reasoning Step = 3


Please point out the **free space** between the **red box** on the left and the **black box**.



Reasoning Step = 1

Please point out the **sofa** on the **right side** of the picture that is closest to the viewer.


Figure 37: Some Placement Examples. The model is asked to identify a valid free space based on spatial reference. The blue point shows the RoboRefer’s prediction (all correct).



Reasoning Step = 2
Please point out the **free area** in the **direction of the handle of the rightmost green cup**.



Reasoning Step = 2
Please point out the **free space** in the **facing direction of the orange box**.



Reasoning Step = 3
Please point out the **free space** in **front of the brown object on the shelf**.



Reasoning Step = 4
Please point out the **free space** on the **table between the pillow and the brown bowl**.




Reasoning Step = 2
Please point out the **free space** **below the white table**.




Reasoning Step = 2
Please point out the **free space** **below the black shelf**.



Reasoning Step = 2
Please point out the **free space** **inside the closest pot**.



Reasoning Step = 2
Please point out the **free space** **below the table**.



Reasoning Step = 2
Please point out the **free space** **between the white chair and the brown shelf**.

Figure 38: Some Placement Examples. The model is asked to identify a valid free space based on spatial reference. The blue point shows the RoboRefer's prediction (all correct).




Reasoning Step = 2

Please point out the **free space** in the **direction of the white side of the can**.



Reasoning Step = 3

Please point out the **free space** **between the scissors and the microwave**.



Reasoning Step = 3

Please point out the **free space** **between the bathtub and the toilet**.



Reasoning Step = 4

Please point out the **free space** in **front of the blue box which is on the top level of the shelf**.




Reasoning Step = 2

Please point out the **free space** in **front of the white vase which is on the top level of the shelf**.




Reasoning Step = 2

Please point out the **free space** **on the right of the farthest pot**.




Reasoning Step = 4

Please point out the **free space** **between the black plate, blue can and closest water glass**.



Reasoning Step = 2


Please point out the **free space** in the **top corner of the table**.




Reasoning Step = 3

Please point out the **free area** on the **table in facing direction of the second chair from the front on the right side**.

Figure 39: Some Placement Examples. The model is asked to identify a valid free space based on spatial reference. The blue point shows the RoboRefer's prediction (all correct).




Reasoning Step = 3
Please point out the **free spot** between the **blue water kettle** and the **orange**.



Reasoning Step = 3
Please point out the **free space** between the **mouse** and the **green cup**.




Reasoning Step = 3
Please point out the **free space** between the **pot** and the **black box**.




Reasoning Step = 2
Please point out the **free area** in the **top-left corner** of the **table**.




Reasoning Step = 3
Please point out the **free space** on the **third level from the top** of the **cat tree**.




Reasoning Step = 2
Please point out the **free space** on the **lowest shelf** of the **cat tree**.



Reasoning Step = 3
Please point out the **free area** on the **table** that the **first chair from the front on the left side** is directly **facing**.



Reasoning Step = 5
Please point out the **free space** on the **table** between the **speaker** to the **right** of the **monitor** and the **mouse**.



Reasoning Step = 4
Please point out the **free space** on the **right part** of the **table** between the **mouse** and the **picture frame**.

Figure 40: Some Placement Examples. The model is asked to identify a valid free space based on spatial reference. The blue point shows the RoboRefer's prediction (all correct).

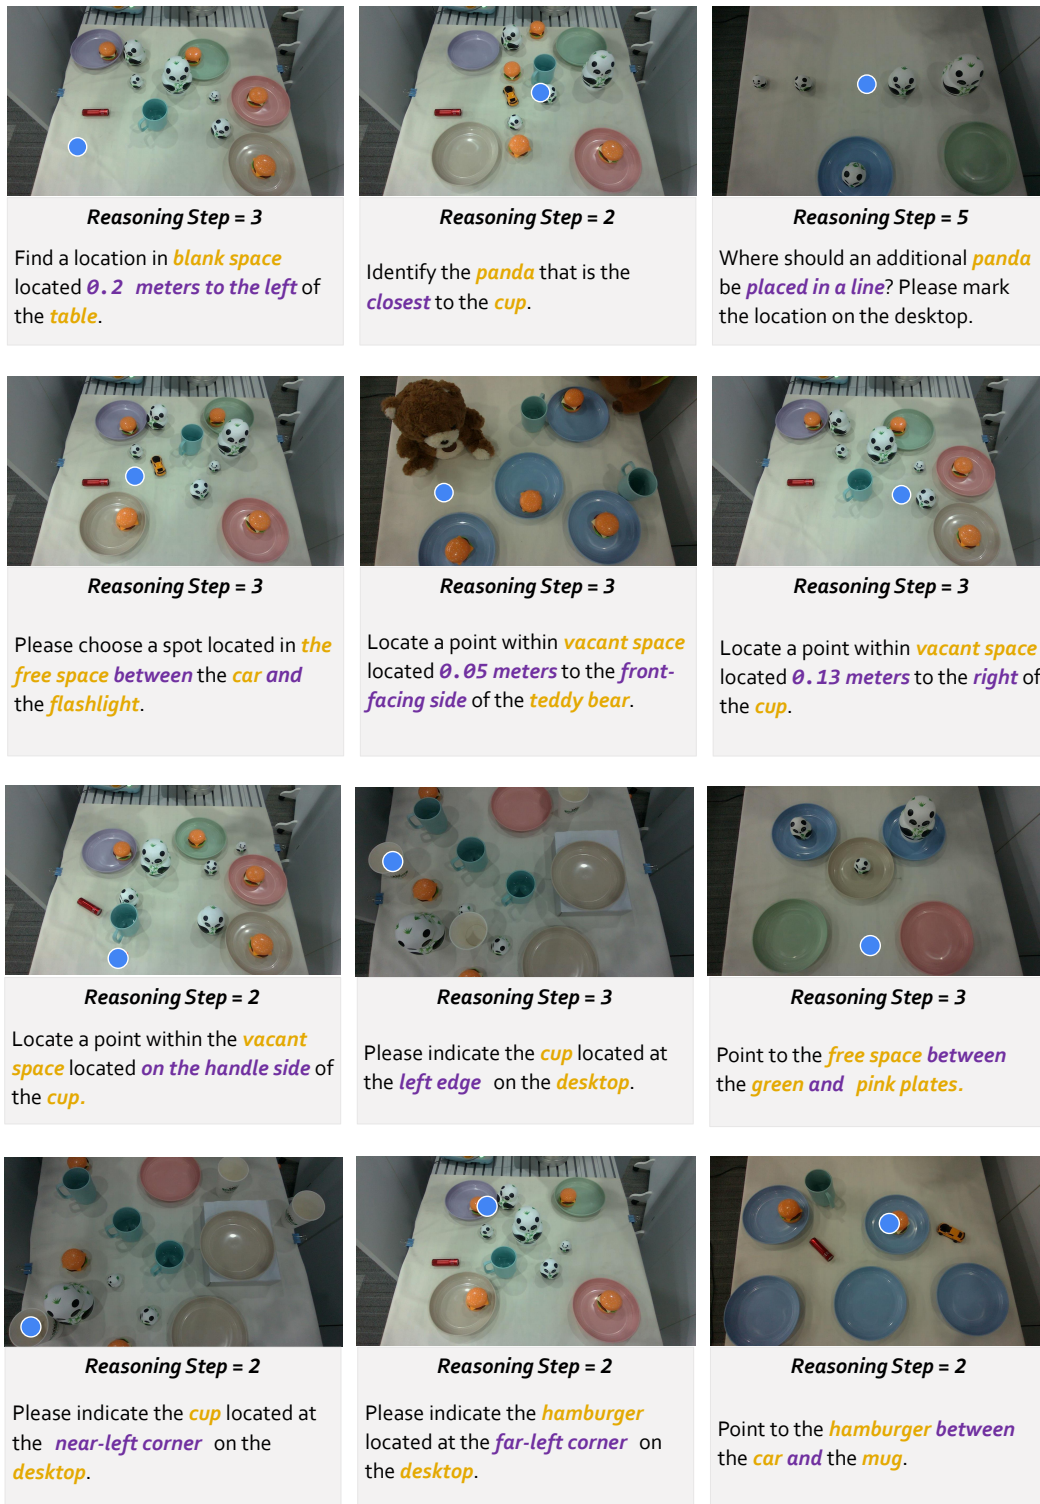


Figure 41: Visualization of RoboRefer’s prediction (blue point) in the real-world evaluation.

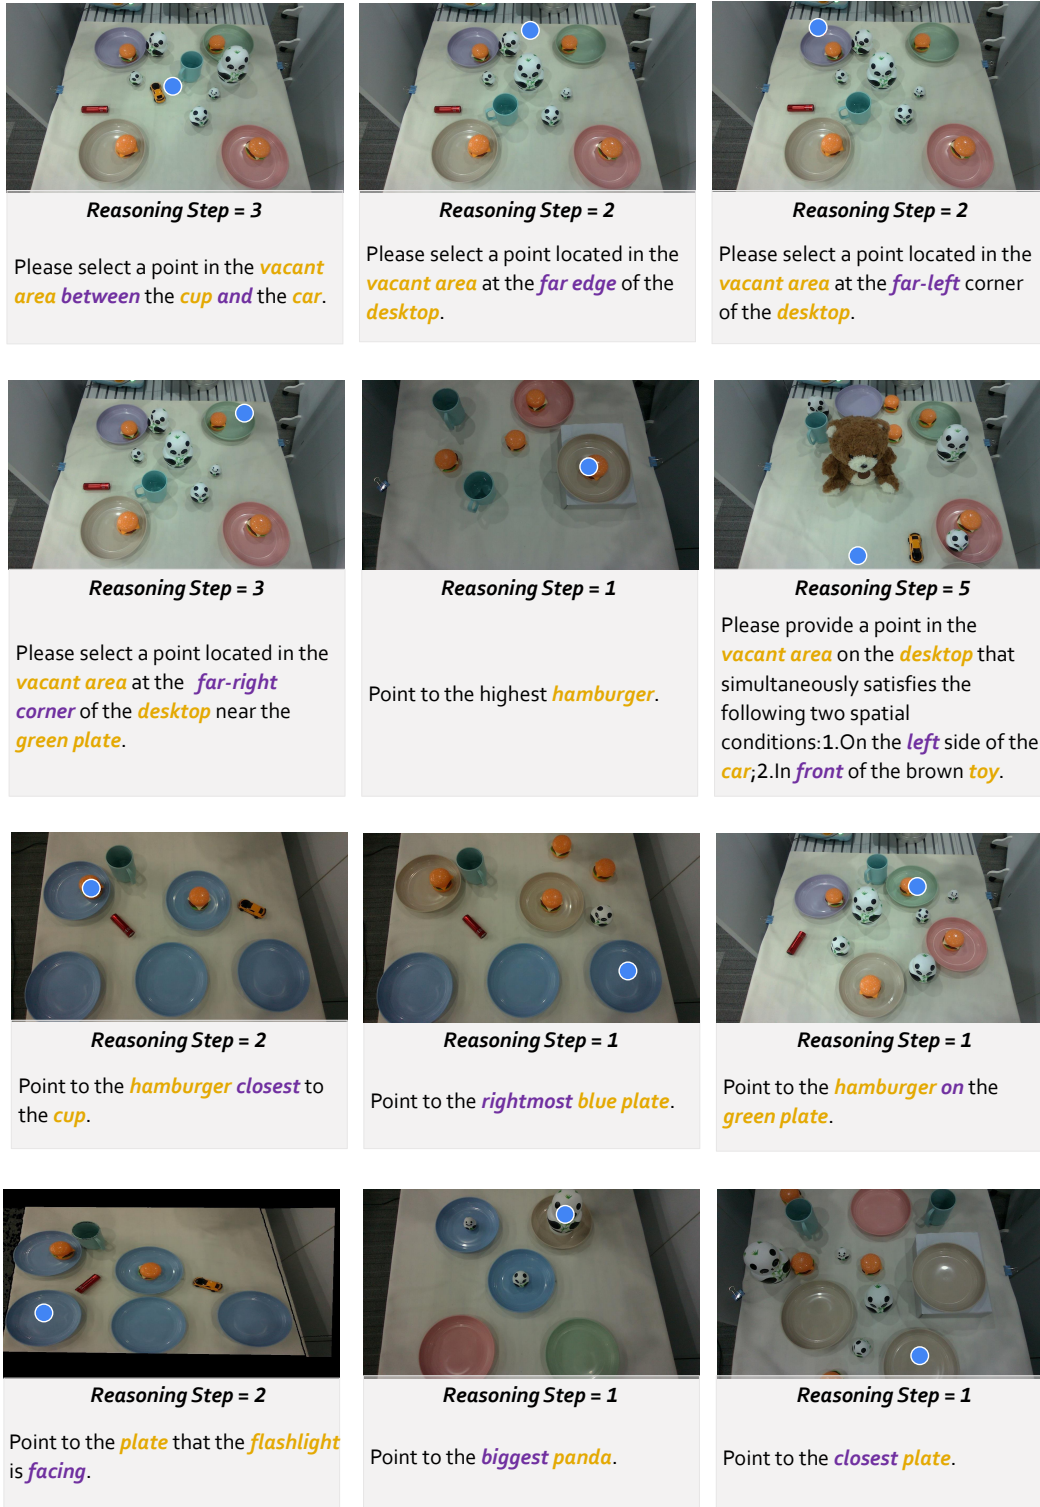



Figure 42: Visualization of RoboRefer’s prediction (blue point) in the real-world evaluation.




Reasoning Step = 1

Point to the **farthest panda**.




Reasoning Step = 1

Point to the **biggest panda**.



Reasoning Step = 1

Point to the **closest plate**.




Reasoning Step = 2

Point to the **hamburger** that is **first close** to the **nearest** mug to the **camera**.




Reasoning Step = 3

Point to the **free space** on the **front-facing** side of the **teddy bear**.




Reasoning Step = 1

Point to the **hamburger** that the **flashlight** is **illuminating**.



Reasoning Step = 1

Point to the **nearest panda**.



Reasoning Step = 2

Point to the **plate illuminated** by the **flashlight**.




Reasoning Step = 1

Point to the **leftmost panda**.




Reasoning Step = 2

Pick the **apple in front of the logo** side of the **leftmost cup**.



Reasoning Step = 3

Place the **apple** on the table, **aligned with the apple row**.



Reasoning Step = 1

Navigate to the **nearest table**.

Figure 43: Visualization of *RoboRefer*'s prediction (blue point) in the real-world evaluation.

References

- [1] Justin Lazarow, David Griffiths, Gefen Kohavi, Francisco Crespo, and Afshin Dehghan. Cubify anything: Scaling indoor 3d object detection. *arXiv preprint arXiv:2412.04458*, 2024.
- [2] Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab Kamali, Stefan Popov, Matteo Mallocci, Alexander Kolesnikov, et al. The open images dataset v4: Unified image classification, object detection, and visual relationship detection at scale. *IJCV*, 2020.
- [3] Michael Tschannen, Alexey Gritsenko, Xiao Wang, Muhammad Ferjad Naeem, Ibrahim Alabdulmohsin, Nikhil Parthasarathy, Talfan Evans, Lucas Beyer, Ye Xia, Basil Mustafa, et al. Siglip 2: Multilingual vision-language encoders with improved semantic understanding, localization, and dense features. *arXiv*, 2025.
- [4] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv*, 2025.
- [5] An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. Spatialrgpt: Grounded spatial reasoning in vision language models. *NIPS*, 2024.
- [6] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *CVPR*, 2024.
- [7] Youcai Zhang, Xinyu Huang, Jinyu Ma, Zhaoyang Li, Zhaochuan Luo, Yanchun Xie, Yuzhuo Qin, Tong Luo, Yaqian Li, Shilong Liu, et al. Recognize anything: A strong image tagging model. In *CVPR*, 2024.
- [8] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *ECCV*, 2024.
- [9] Luigi Piccinelli, Christos Sakaridis, Yung-Hsu Yang, Mattia Segu, Siyuan Li, Wim Abbeloos, and Luc Van Gool. Unidepthv2: Universal monocular metric depth estimation made simpler. *arXiv*, 2025.
- [10] Shengjie Zhu, Abhinav Kumar, Masa Hu, and Xiaoming Liu. Tame a wild camera: in-the-wild monocular camera calibration. *NIPS*, 2023.
- [11] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *ICLR*, 2025.
- [12] Erik Daxberger, Nina Wenzel, David Griffiths, Haiming Gang, Justin Lazarow, Gefen Kohavi, Kai Kang, Marcin Eichner, Yinfei Yang, Afshin Dehghan, et al. Mm-spatial: Exploring 3d spatial understanding in multimodal llms. *arXiv*, 2025.
- [13] Chan Hee Song, Valts Blukis, Jonathan Tremblay, Stephen Tyree, Yu Su, and Stan Birchfield. Robospatial: Teaching spatial understanding to 2d and 3d vision-language models for robotics. *CVPR*, 2025.
- [14] Qwen Team. Qwq-32b: Embracing the power of reinforcement learning, March 2025.
- [15] Gilad Baruch, Zhuoyuan Chen, Afshin Dehghan, Tal Dimry, Yuri Feigin, Peter Fu, Thomas Gebauer, Brandon Joffe, Daniel Kurz, Arik Schwartz, et al. Arkitscenes: A diverse real-world dataset for 3d indoor scene understanding using mobile rgb-d data. *arXiv*, 2021.
- [16] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017.
- [17] Johanna Wald, Torsten Sattler, Stuart Golodetz, Tommaso Cavallari, and Federico Tombari. Beyond controlled environments: 3d camera re-localization in changing indoor scenes. In *ECCV*, 2020.
- [18] Alexander Raistrick, Lingjie Mei, Karhan Kayan, David Yan, Yiming Zuo, Beining Han, Hongyu Wen, Meenal Parakh, Stamatis Alexandropoulos, Lahav Lipson, et al. Infinigen indoors: Photorealistic indoor scenes using procedural generation. In *CVPR*, 2024.

- [19] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *CVPR*, 2023.
- [20] Zekun Qi, Wenyao Zhang, Yufei Ding, Runpei Dong, Xinqiang Yu, Jingwen Li, Lingyun Xu, Baoyu Li, Xialin He, Guofan Fan, et al. Sofar: Language-grounded orientation bridges spatial reasoning and object manipulation. *arXiv*, 2025.
- [21] Zhijian Liu, Ligeng Zhu, Baifeng Shi, Zhuoyang Zhang, Yuming Lou, Shang Yang, Haocheng Xi, Shiyi Cao, Yuxian Gu, Dacheng Li, et al. Nvila: Efficient frontier visual language models. *arXiv*, 2024.
- [22] Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *CVPR*, 2023.
- [23] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *CVPR*, 2024.
- [24] Fuxiao Liu, Kevin Lin, Linjie Li, Jianfeng Wang, Yaser Yacoob, and Lijuan Wang. Mitigating hallucination in large multi-modal models via robust instruction tuning. In *ICLR*, 2024.
- [25] Licheng Yu, Patrick Poirson, Shan Yang, Alexander C Berg, and Tamara L Berg. Modeling context in referring expressions. In *ECCV*, 2016.
- [26] Arijit Ray, Jiafei Duan, Reuben Tan, Dina Bashkirova, Rose Hendrix, Kiana Ehsani, Aniruddha Kembhavi, Bryan A Plummer, Ranjay Krishna, Kuo-Hao Zeng, et al. Sat: Spatial aptitude training for multimodal language models. *arXiv*, 2024.
- [27] Mengfei Du, Binhao Wu, Zejun Li, Xuan-Jing Huang, and Zhongyu Wei. Embspatial-bench: Benchmarking spatial understanding for embodied tasks with large vision-language models. In *ACL (Volume 2: Short Papers)*, 2024.
- [28] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.
- [29] Drew A Hudson and Christopher D Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *CVPR*, 2019.
- [30] Anand Mishra, Shashank Shekhar, Ajeet Kumar Singh, and Anirban Chakraborty. Ocr-vqa: Visual question answering by reading text in images. In *ICDAR*, 2019.
- [31] Amanpreet Singh, Vivek Natarajan, Meet Shah, Yu Jiang, Xinlei Chen, Devi Parikh, and Marcus Rohrbach. Towards vqa models that can read. In *CVPR*, 2019.
- [32] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *IJCV*, 2017.
- [33] Zhihong Shao, Peiyi Wang, Qihao Zhu, Runxin Xu, Junxiao Song, Xiao Bi, Haowei Zhang, Mingchuan Zhang, YK Li, Y Wu, et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv*, 2024.
- [34] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [35] Jingyi Zhang, Jiaxing Huang, Huanjin Yao, Shunyu Liu, Xikun Zhang, Shijian Lu, and Dacheng Tao. R1-vl: Learning to reason with multimodal large language models via step-wise group relative policy optimization. *arXiv*, 2025.
- [36] Yuqi Liu, Bohao Peng, Zhisheng Zhong, Zihao Yue, Fanbin Lu, Bei Yu, and Jiaya Jia. Seg-zero: Reasoning-chain guided segmentation via cognitive reinforcement. *arXiv*, 2025.
- [37] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, et al. Openeqa: Embodied question answering in the era of foundation models. In *CVPR*, 2024.
- [38] Muhammad Khalifa, Rishabh Agarwal, Lajanugen Logeswaran, Jaekyeom Kim, Hao Peng, Moontae Lee, Honglak Lee, and Lu Wang. Process reward models that think. *arXiv*, 2025.
- [39] Runze Liu, Junqi Gao, Jian Zhao, Kaiyan Zhang, Xiu Li, Biqing Qi, Wanli Ouyang, and Bowen Zhou. Can 1b llm surpass 405b llm? rethinking compute-optimal test-time scaling. *arXiv*, 2025.

- 1385 [40] Haozhan Shen, Peng Liu, Jingcheng Li, Chunxin Fang, Yibo Ma, Jiajia Liao, Qiaoli Shen, Zilun
1386 Zhang, Kangjia Zhao, Qianqian Zhang, et al. Vlm-r1: A stable and generalizable r1-style large
1387 vision-language model. *arXiv preprint arXiv:2504.07615*, 2025.
- 1388 [41] Peter Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Adithya Jairam Vedagiri IYER,
1389 Sai Charitha Akula, Shusheng Yang, Jihan Yang, Manoj Middepogu, Ziteng Wang, et al.
1390 Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *NIPS*, 2024.
- 1391 [42] Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A
1392 Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see
1393 but not perceive. In *ECCV*, 2024.
- 1394 [43] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut,
1395 Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly
1396 capable multimodal models. *arXiv*, 2023.
- 1397 [44] Gemini Robotics Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montser-
1398 rat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza,
1399 Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *arXiv*, 2025.
- 1400 [45] Yuhao Lu, Yixuan Fan, Beixing Deng, Fangfu Liu, Yali Li, and Shengjin Wang. V1-grasp: a
1401 6-dof interactive grasp policy for language-oriented objects in cluttered indoor scenes. In *IROS*,
1402 2023.
- 1403 [46] Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan
1404 Murali, Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial
1405 affordance prediction for robotics. *arXiv*, 2024.
- 1406 [47] Alec Radford, Jong Wook Kim, Tao Xu, Greg Brockman, Christine McLeavey, and Ilya
1407 Sutskever. Robust speech recognition via large-scale weak supervision. In *ICML*, 2023.
- 1408 [48] Hao-Shu Fang, Chenxi Wang, Hongjie Fang, Minghao Gou, Jirong Liu, Hengxu Yan, Wenhai
1409 Liu, Yichen Xie, and Cewu Lu. Anygrasp: Robust and efficient grasp perception in spatial and
1410 temporal domains. *IEEE Transactions on Robotics*, 2023.
- 1411 [49] Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park,
1412 Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, et al. Molmo and pixmo:
1413 Open weights and open data for state-of-the-art multimodal models. *arXiv*, 2024.