

ROBOSPATIAL: Teaching Spatial Understanding to 2D and 3D Vision-Language Models for Robotics

Anonymous CVPR submission

Paper ID 0000

Abstract

Spatial understanding is a crucial capability that enables robots to perceive their surroundings, reason about their environment, and interact with it meaningfully. In modern robotics, these capabilities are increasingly provided by vision-language models. However, these models face significant challenges in spatial reasoning tasks, as their training data are based on general-purpose image datasets that often lack sophisticated spatial understanding. For example, datasets frequently do not capture reference frame comprehension, yet effective spatial reasoning requires understanding whether to reason from ego-, world-, or object-centric perspectives. To address this issue, we introduce ROBOSPATIAL, a large-scale dataset for spatial understanding in robotics. It consists of real indoor and tabletop scenes, captured as 3D scans and egocentric images, and annotated with rich spatial information relevant to robotics. The dataset includes 1M images, 5k 3D scans, and 3M annotated spatial relationships, and the pairing of 2D egocentric images with 3D scans makes it both 2D- and 3D- ready. Our experiments show that models trained with ROBOSPATIAL outperform baselines on downstream tasks such as spatial affordance prediction, spatial relationship prediction, and robot manipulation.

1. Introduction

The rise of vision-language models (VLMs) has opened new opportunities for agents to interpret and act upon the visual world using natural language. VLMs have been adopted across a range of embodied settings, notably in robotics and augmented reality (AR). In robotics, they have enabled grounded scene understanding [13, 66], manipulation [7], and policy code generation [29, 48], while in AR, they support tasks like object labeling [50], action recognition [16, 17], and temporal grounding [22].

VLMs can recognize objects, classify scenes, and even provide general descriptions that capture high-level at-




Figure 1. ROBOSPATIAL dataset facilitates 3D spatial reasoning for robot manipulation. This illustration demonstrates how a model trained on ROBOSPATIAL enables human-aligned spatial reasoning within the correct reference frame, supporting task grounding, planning, and detection for manipulation tasks.

tributes. However, despite significant recent advancements, VLMs [30, 35, 42] still fall short in spatial understanding [26, 38, 47, 61]. They struggle with tasks that require interpreting nuanced spatial relationships between objects, such as describing where one object is in relation to another or determining the best location to place an item within a specific condition. For example, while a model might accurately describe a “bowl on the table,” it lacks the ability to reason about where on the table the bowl is, where it should go to ensure accessibility or stability, or how it might fit among other objects. Furthermore, a critical limitation of existing VLM training datasets is their inability to capture *reference frame* understanding (*ref. frame*) — the way we interpret spatial relationships changes drastically depending on whether we’re viewing from a first-person perspective, focusing on specific objects, or observing the entire scene, all of which are essential for real-world interactions.

These limitations highlight an ongoing challenge: bridging the gap between surface-level scene description and the deeper spatial comprehension necessary for intuitive interaction. Several recent efforts aim to address this by explicitly training VLMs on spatial reasoning tasks, yet many fall

| Dataset | 3D scans | Embodied | Ref. frames | Compatibility | Domain | #Scans | #Images | #Spatial QAs |
|-----------------------|----------|----------|-------------|---------------|------------------|--------|---------|--------------|
| EmbSpatial-Bench [10] | ✓ | ✓ | ✗ | ✗ | Indoor | 277 | 2k | 4k |
| Visual Spatial [33] | ✗ | ✗ | ✓ | ✗ | MSCOCO | 0 | 10k | 10k |
| SpatialRGPT-Bench [6] | ✗ | ✗ | ✓ | ✓ | Indoor, AV | 0 | 1.4k | 1.4k |
| BLINK-Spatial [15] | ✗ | ✗ | ✓ | ✗ | Generic | 0 | 286 | 286 |
| What's up [26] | ✗ | ✗ | ✗ | ✗ | Generic | 0 | 5k | 10k |
| Spatial-MM [47] | ✗ | ✗ | ✓ | ✗ | Generic | 0 | 2.3k | 2.3k |
| ROBOSPATIAL | ✓ | ✓ | ✓ | ✓ | Indoor, tabletop | 5k | 1M | 3M |

Table 1. Comparison with other spatial reasoning datasets that include object-centric spatial relationships.

short of the demands posed by embodied or robotics settings. For example, SpatialVLM [5] and SpatialRGPT [6] train VLMs to answer questions about distances and spatial relationships between objects, advancing spatial understanding at a conceptual level. However, these models are trained on datasets comprised of images from the internet, with annotations generated by perception models. As a result, they struggle to generalize to *embodied* images—those captured by robot cameras within real-world environments, which often lack identifiable cues for an absolute scale. Pointing models, such as RoboPoint [62] and more recently Molmo [9], take a different approach by training VLMs to produce grounded 2D coordinates that pinpoint object locations or free space within a scene. However, these models lack understanding of real-world constraints, such as inferring object-centric reference frames for perspective-invariant reasoning, or accounting for the space required to place various objects. This results, for example, in failing to predict whether the gray bowl can fit in front of the car in Figure 1.

This paper hypothesizes that a primary bottleneck limiting the effectiveness of VLMs in robotics is the scarcity of suitable training data, as highlighted by Table 1. To address this, we introduce **ROBOSPATIAL**, a dataset designed specifically to facilitate spatial understanding in VLMs for robotic applications. The proposed approach leverages annotated indoor scene and tabletop RGBD datasets, transforming them into targeted question-answer pairs designed to probe spatial reasoning skills critical for robotics.

We categorize the questions into three types, each serving a distinct purpose. **Spatial context** focuses on identifying empty space or support surfaces in the environment that can accommodate other objects. These questions are formulated as point-prediction tasks, challenging the model to determine appropriate locations within free space where an object can be placed—for example, “Where on the table can I put the plate?” **Spatial compatibility** builds on the identified empty space to assess whether a given area can feasibly support the placement of a specific object, ensuring sufficient size and fit. These questions are posed in a binary format, such as “Can the chair be placed in front of the table?” **Spatial configuration** examines the relative spatial relationships between two objects. These questions

use a binary format to determine whether a spatial relation holds, such as “Is the mug to the left of the laptop?”

To enhance the model’s ability to interpret spatial instructions from different perspectives, each question-answer pair in ROBOSPATIAL is posed from three distinct reference perspectives/frames: (a) **Ego-centric** from the observer’s perspective at the camera pose, (b) **World-centric** grounded in a global world frame, and (c) **Object-centric** based on a reference frame attached to the focal object. This multi-frame approach enables models to handle complex spatial instructions more flexibly, preparing them to better generalize to dynamic robotic contexts. Applying our methodology to existing indoor scene and tabletop datasets, we generate both a comprehensive training dataset and a benchmark for spatial question answering in robotics. ROBOSPATIAL contains around **1M** images, **5k** 3D scans, and **3M** annotated spatial relationships, with paired 2D egocentric images and 3D scans to make it both 2D- and 3D- ready.

To validate the effectiveness of ROBOSPATIAL, comprehensive experiments were conducted using multiple state-of-the-art (SOTA) 2D and 3D VLMs. Results demonstrate that models trained on ROBOSPATIAL exhibit significantly improved spatial reasoning capabilities, consistently outperforming baseline methods on the evaluation benchmark **ROBOSPATIAL-Val**, a held-out validation subset derived from the heuristically generated ROBOSPATIAL dataset. To further assess the generalization and robustness of these trained VLMs, additional evaluations were conducted using three complementary benchmarks: **ROBOSPATIAL-Home**, a manually collected dataset consisting of paired RGB and depth images, and two external benchmarks, BLINK-Spatial [15] and SpatialBot [2]. These benchmarks rigorously test spatial reasoning skills in practical robotic tasks, including object rearrangement and contextual question answering in indoor environments, while also examining the models’ capacity to generalize to novel spatial reasoning scenarios beyond the original training data. Across all benchmarks, models trained on ROBOSPATIAL consistently outperformed baseline methods, demonstrating the broad utility of the dataset. Leveraging the 3D-ready design of ROBOSPATIAL, direct comparisons between the spatial reasoning performance of 2D and 3D VLMs were also performed. Although initial results indicate potential advan-

tages for 3D models, differences in pretraining data and base LLM architectures among models render the comparison inconclusive. ROBOSPATIAL is specifically designed to support both 2D and 3D research directions, enabling future studies to address these differences more conclusively.

Our contributions are threefold:

- A new training dataset, ROBOSPATIAL, comprising images and 3D scans paired with spatial questions and answers, accompanied by an evaluation benchmark, ROBOSPATIAL-Val, a held-out validation set. Additionally, we introduce ROBOSPATIAL-Home, a manually collected and annotated dataset designed specifically for assessing real-world spatial reasoning in indoor environments. These datasets uniquely incorporate multiple reference frames, object-object spatial relationships, object-space relationships, and object compatibility. We make the data and code for generating the dataset from 3D annotated scenes publicly available.
- VLMs trained on ROBOSPATIAL demonstrate superior spatial reasoning, outperforming SOTA baselines on language-guided robot manipulation and indoor scene question answering.
- Comprehensive experiments assessing spatial reasoning capabilities in both 2D and 3D VLMs, comparing the difference between SOTA VLMs in real-world spatial tasks.

2. Related Work

VLMs for Robotics. Vision-language models (VLMs) have emerged as pivotal tools in robotics, enabling systems to interpret and act upon complex visual and textual information. By integrating visual perception with language understanding, VLMs facilitate more intuitive human-robot interactions and enhance autonomous decision-making capabilities. Recent advancements have demonstrated the potential of VLMs in various robotic applications. For instance, vision-language-action models (VLAs) [27, 41, 65] enable robots to interpret and execute complex instructions and output executable robot actions. Additionally, VLMs like GPT-4v [42] have been utilized for high-level task planning [55], allowing robots to generate detailed action sequences from natural language instructions. Furthermore, VLMs have been used for keypoint/mask prediction [21, 40, 59], error analysis [11, 49], grasp pose prediction [19]. Despite these advancements, integrating VLMs [2, 6, 62] into robotic systems presents challenges. One significant hurdle is the need for precise spatial reasoning to navigate and manipulate objects effectively. While VLMs excel in understanding and generating language, their ability to comprehend and reason about spatial relationships in dynamic environments remains limited [57, 60, 61]. Therefore, ROBOSPATIAL aims to tackle this gap by presenting a large scale pretraining and evaluation setup for teaching spatial understanding to VLM for robotics.

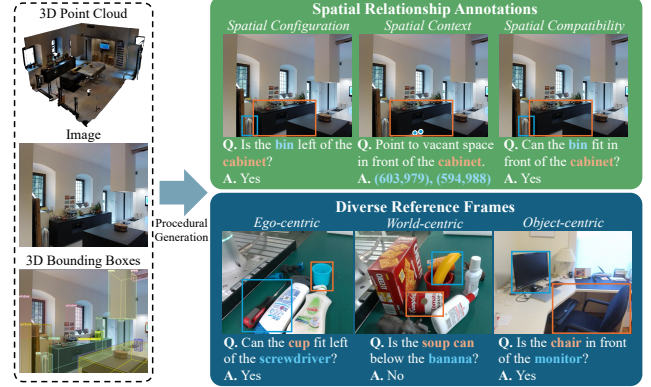


Figure 2. Overview of the ROBOSPATIAL dataset. We automatically generate spatial relationship annotations from existing datasets with 3D point clouds, egocentric images, and 3D bounding box annotations. We create question/answer pairs covering three classes of spatial relationships, three spatial reference frames, and both binary (yes/no) and numeric (e.g. 2D image points) answers. From 1M images and 5k scans, we generate over 3M spatial question/answer pairs.

Spatial Understanding with VLMs. Spatial understanding has been implicitly and explicitly part of various vision and question answering tasks [1, 15, 23–25, 28, 46, 51]. While many benchmarks and methods have been proposed, they often come with limitations: some focus exclusively on simulations [53] or generic images [5, 6, 15, 26, 33, 43, 44, 47], others are difficult to evaluate due to their reliance on free-form text outputs [10, 32, 53], some rely on complete 3D scans [32, 37, 39, 64], and others do not account for reference frames [5, 6, 15, 32, 37, 39, 44, 64]. Furthermore, many fail to address actionable, robotics-relevant spatial relationships such as spatial compatibility and context [10, 15, 26, 32, 44, 47, 58].

Inspired by prior works on spatial reasoning [26, 33]—where the impact of reference frames and spatial configurations was explored in generic images [23, 31]—we extend spatial understanding to a robotics-specific context with actionable spatial relationships such as spatial compatibility and spatial context. Our aim is to enable direct application to robotic workflows, such as task planning and verification.

To achieve this, we have developed a large-scale 2D/3D ready training dataset using our automated data generation pipeline. We further show how ROBOSPATIAL can be used to teach spatial reasoning to a suite of vision-language models (VLMs) in in-domain and out-of-domain spatial reasoning datasets. We hope these resources lower the barrier to entry for exploring spatial understanding tailored to robotics.

3. Approach

We begin by explaining the selection of three spatial relationships: spatial context, spatial compatibility, and spatial configuration. Next, we describe the data generation pipeline used to construct the ROBOSPATIAL. Figure 2 provides an overview of the dataset.

3.1. Spatial Relationships

The dataset is organized around three core spatial relationships that we believe address the essential aspects of spatial reasoning for robotic tasks: *spatial context*, *spatial compatibility*, and *spatial configuration*. *Context* allows robots to assess the relationship between objects and their surrounding space, facilitating the identification of empty or occupied areas, which is relevant for downstream applications such as path planning and obstacle avoidance. *Compatibility* focuses on whether objects can coexist or interact without conflict in a given space, which is vital for object placement, assembly, and operational safety. *Configuration* enables robots to understand and interpret the relative positioning of objects, which is crucial for directing navigation, manipulation, and interaction within complex environments. Together, these spatial relationships provide a more nuanced and practical framework for robotic applications than metrics like distance—which is hard to normalize across different scales, environments, and tasks—thereby enabling robots to perform complex tasks with greater reliability.

3.2. Dataset Generation

The goal of the data construction pipeline is to generate a large-scale, high-accuracy spatial relationship dataset with minimal human intervention, using automatic heuristics grounded in 3D geometry and 2D image views.

The pipeline takes as input a scene dataset \mathcal{D}_s that contains RGB images, camera poses (both extrinsic and intrinsic parameters), and oriented 3D bounding box annotations with semantic object labels. The output is a spatial reasoning dataset \mathcal{D} , where each entry $d_i = \langle I_i, q_i, a_i, l_i \rangle$ consists of an image I_i , a question q_i , an answer a_i , and a reference frame label $l_i \in \{\text{ego}, \text{world}, \text{object}\}$. Each question is derived from one of three spatial reasoning categories: spatial configuration, spatial context, or spatial compatibility. To support reliable object reference resolution, we also generate an auxiliary object grounding dataset that links object descriptions to 2D bounding boxes.

To improve clarity and reproducibility, we describe the data generation pipeline in two main stages. We also separate the reasoning logic used for extracting 3D relationships from that used for generating 2D image-space targets.

3.2.1. Stage 1: 3D Spatial Relation Extraction

The first stage involves extracting spatial relationships between objects or between objects and free space, based on 3D geometry. Each spatial relation is defined as $s_i = \langle I_i, a_i, t_i, r_i, l_i \rangle$, where I_i is the source image, a_i is the anchor object, t_i is the target object or a sampled point in free space, $r_i \in \{\text{left}, \text{right}, \text{above}, \text{below}, \text{front}, \text{behind}\}$ is the relation preposition, and $l_i \in \{\text{ego}, \text{world}, \text{object}\}$ denotes the reference frame.

We use oriented 3D bounding boxes, provided by the source dataset, to compute spatial relationships. Each bounding box includes both the 3D location and heading of the object. The object’s orientation is defined by the heading vector of the bounding box, aligned with the object’s front-facing direction. Using this orientation, we determine the appropriate directional region (e.g., front, left) relative to the reference frame. For instance, a relation such as “in front of (anchor object) (object frame)” refers to the positive direction along the anchor object’s heading vector. These relationships are calculated independently for each of the three reference frames: the world frame is aligned with the dataset-level coordinate system; the ego frame is defined by the camera pose (i.e., camera-centered); and the object frame is defined by the local orientation of the anchor object.

The camera extrinsics are used to transform coordinates between reference frames. Although the method does not require point clouds or meshes, it relies on camera intrinsics and extrinsics to project between 2D and 3D and to ensure consistent reference frame reasoning. For each spatial configuration task, we evaluate all visible object pairs that appear uniquely in the image, avoiding duplicate instances to minimize ambiguity. The resulting relationships are binary (True/False) and specify whether the spatial condition holds for the given object pair.

3.2.2. Stage 2: 2D Spatial Point and Region Sampling

In the second stage, we generate 2D image-space annotations for spatial context and spatial compatibility tasks. These rely on the 3D bounding box layout and calibrated camera parameters to map spatial relationships into image coordinates.

For spatial context, we construct a top-down occupancy map of the scene by marking regions occupied by 3D bounding boxes. We then randomly sample 3D points in empty space that lie in a specified directional relation to the anchor object, following the same frame-dependent heuristics as in the configuration task. These points are projected into the image plane using the camera intrinsics. To ensure the points are valid, we filter out samples that are obstructed or occluded based on line-of-sight from the camera. Specifically, we perform raycasting from the camera center to each sampled 3D point, and discard points whose rays

intersect any occupied bounding box volumes before reaching the target location. The final answer is a list of 2D (x, y) image coordinates that satisfy the spatial context constraint.

Spatial compatibility extends this idea by checking whether a target object can fit within the sampled region. We simulate placing a virtual bounding box, matching the size of the target object, at the candidate location on the ground plane. A region is considered compatible if the simulated placement does not intersect with any existing bounding boxes in the scene and provides at least a 10 cm margin along each axis. The simulation allows for translation and in-plane rotation of the object. The answer to this task is binary (True/False), indicating whether the region can accommodate the object.

3.2.3. Question-Answer Generation

Once the spatial relations $\{s_i\}$ have been extracted, we generate corresponding question-answer pairs $\{d_i\}$ using structured templates. Each question follows the format: {TARGET} {RELATION} {ANCHOR} {REF. FRAME} where the relation and frame are defined in Section 3.2.1. To ensure that models learn from visual grounding rather than linguistic priors, we use deterministic templates that avoid ambiguity and minimize reliance on commonsense.

Each spatial relation type—context, compatibility, and configuration—has a corresponding question format. Configuration and compatibility tasks result in binary (True/False) answers. Context questions produce a list of valid 2D coordinates in image space.

Correctly resolving which object is being referred to in a spatial question is essential for reliable spatial understanding. To reduce errors arising from incorrect object identification, we additionally generate an auxiliary object grounding dataset that links object descriptions to 2D bounding boxes in the image. These grounding annotations are derived by projecting 3D bounding boxes into image space using camera intrinsics and extrinsics. This supervision helps models more accurately resolve references during spatial reasoning and is included during training. See Appendix B.3 for details.

Using this pipeline, we generate around 3 million spatial relationships and their associated question-answer pairs. This scale is an order of magnitude larger than prior spatial reasoning datasets (see Table 1).

4. Experiments

4.1. Setup

We apply the data generation pipeline to three scene datasets—ScanNet [8], Matterport3D [4], and 3RScan [56]—and two tabletop datasets—HOPE [54] and GraspNet-1B [12]. We retrieve 3D bounding box annotations and embodied images from EmbodiedScan [58], and generate a large-scale spatial reasoning dataset covering

| Dataset | Type | Splits | Images | QA pairs |
|----------|------------|------------|-------------|----------|
| Indoor | Train | 4916 scans | 883k images | 2.7M |
| | Validation | 40 scans | 1k images | 3k |
| Tabletop | Train | 190 scenes | 76k images | 220k |
| | Validation | 77 scenes | 355 images | 3k |

Table 2. Dataset splits for indoor and tabletop dataset. Detailed data statistics are in the Appendix.

diverse indoor environments: larger scenes for navigation and smaller object-centric setups for manipulation.

In total, ROBOSPATIAL includes approximately 3M spatial QA pairs across 5k 3D scans and 1M images. (Table 2 provides a breakdown; values are rounded to the nearest thousand for clarity.)

4.1.1. Trained 2D/3D VLMs

2D VLMs. We evaluate several vision-language models (VLMs) using RGB-only image inputs. Our selected base VLMs are VILA-1.5-8B [30] and LLaVA-NeXT-8B [35]. We also include three specialized models: SpaceLLaVA-13B (a community version of SpatialVLM [5]), RoboPoint-13B [62] (trained to predict points in empty space given an object reference), and Molmo-7B [9] (designed for pointing and counting). We also include GPT-4o [42] as a closed-source baseline. We omit models such as SpatialRGPT [6] that depend on external mask inputs, as they bypass the object grounding challenge.

3D VLMs. Models operating over 3D data must handle richer, more complex spatial representations. We include two models that process 3D inputs: 3D-LLM [18], which reconstructs colored 3D point clouds from multi-view RGB images, and LEO [20], which operates on segmented colored point clouds of individual objects. These models allow us to explore spatial reasoning when models consume RGBD or point cloud representations directly.

Fine-tuning. We evaluate models in both zero-shot and fine-tuned settings, using ROBOSPATIAL to fine-tune open-source models. To mitigate failure cases arising from poor object grounding, we also include an auxiliary grounding dataset during training, which provides additional supervision for object reference resolution. This auxiliary dataset does not contribute to spatial reasoning performance. (See Appendix for ablation experiments.)

4.1.2. Spatial Understanding Evaluation

We evaluate spatial reasoning capabilities using ROBOSPATIAL-Val, a held-out validation subset of ROBOSPATIAL sampled from scans that are entirely unseen during training. This benchmark comprises 6,000 heuristically generated questions from our data generation pipeline, with 2,000 questions per spatial relation type. Questions fall into two categories: binary yes/no questions and coordinate prediction tasks. For yes/no questions, we report accuracy. For co-

| Model | Indoor | | | Tabletop | | | Average | | |
|--------------------------------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|---------------|
| | Configuration | Context | Compatibility | Configuration | Context | Compatibility | Indoor | Tabletop | Total |
| <i>Open-source VLMs</i> | | | | | | | | | |
| 2D VLMs | | | | | | | | | |
| VILA [30] | 54.7 | 18.3 | 56.3 | 45.1 | 13.2 | 53.8 | 43.1 | 37.4 | 40.2 |
| +ROBOSPATIAL | 71.4 ↑ | 45.9 ↑ | 77.2 ↑ | 71.8 ↑ | 43.7 ↑ | 73.3 ↑ | 64.8 ↑ | 62.9 ↑ | 63.9 ↑ |
| LLaVA-NeXT [35] | 48.9 | 12.5 | 32.7 | 48.3 | 8.4 | 30.9 | 31.4 | 29.2 | 30.3 |
| +ROBOSPATIAL | 69.3 ↑ | 41.3 ↑ | 70.5 ↑ | 70.7 ↑ | 44.8 ↑ | 66.1 ↑ | 60.4 ↑ | 60.5 ↑ | 60.5 ↑ |
| SpaceLLaVA [5] | 52.6 | 15.3 | 49.0 | 66.5 | 12.2 | 60.1 | 38.9 | 46.2 | 43.6 |
| +ROBOSPATIAL | 76.0 ↑ | 50.7 ↑ | 76.6 ↑ | 74.9 ↑ | 46.4 ↑ | 70.5 ↑ | 67.8 ↑ | 63.6 ↑ | 65.7 ↑ |
| RoboPoint [62] | 39.0 | 41.4 | 38.3 | 37.9 | 31.6 | 45.2 | 39.6 | 38.2 | 38.9 |
| +ROBOSPATIAL | 72.2 ↑ | 68.9 ↑ | 72.1 ↑ | 70.3 ↑ | 61.7 ↑ | 78.4 ↑ | 71.0 ↑ | 70.1 ↑ | 70.6 ↑ |
| 3D VLMs | | | | | | | | | |
| 3D-LLM [18] | 54.5 | 8.1 | 53.6 | 59.2 | 10.6 | 57.4 | 37.6 | 42.4 | 40.0 |
| +ROBOSPATIAL | 76.3 ↑ | 35.4 ↑ | 77.5 ↑ | 76.2 ↑ | 46.8 ↑ | 75.0 ↑ | 63.1 ↑ | 66.0 ↑ | 64.6 ↑ |
| LEO [20] | 56.1 | 11.3 | 58.3 | 60.8 | 11.1 | 59.3 | 41.9 | 43.7 | 42.8 |
| +ROBOSPATIAL | 80.2 ↑ | 56.7 ↑ | 82.5 ↑ | 78.1 ↑ | 55.2 ↑ | 78.9 ↑ | 73.1 ↑ | 70.7 ↑ | 71.9 ↑ |
| <i>Not available for fine-tuning</i> | | | | | | | | | |
| 2D VLMs | | | | | | | | | |
| Molmo [9] | 40.6 | 48.2 | 60.0 | 61.5 | 35.8 | 54.6 | 49.6 | 50.6 | 50.1 |
| GPT-4o [42] | 63.5 | 25.1 | 59.4 | 62.3 | 27.9 | 66.8 | 49.3 | 52.3 | 50.8 |

Table 3. Results of existing 2D/3D VLMs on a held-out validation split (ROBOSPATIAL-Val) of images and scans. All methods, for all tasks, perform better (↑) when fine-tuned on ROBOSPATIAL. The best result for each column is bolded.

| Model | ROBOSPATIAL-Home | | | BLINK | SpatialBench |
|--------------------------------------|------------------|---------------|---------------|---------------|---------------|
| | Configuration | Context | Compatibility | Accuracy | Accuracy |
| 2D VLMs | | | | | |
| VILA [30] | 57.8 | 0.0 | 69.0 | 72.7 | 53.0 |
| +ROBOSPATIAL | 65.9 ↑ | 15.6 ↑ | 78.0 ↑ | 79.7 ↑ | 73.6 ↑ |
| LLaVA-NeXT [35] | 68.3 | 0.0 | 70.5 | 71.3 | 55.9 |
| +ROBOSPATIAL | 78.9 ↑ | 19.7 ↑ | 80.1 ↑ | 79.0 ↑ | 70.6 ↑ |
| SpaceLLaVA [5] | 61.0 | 2.5 | 61.0 | 76.2 | 47.1 |
| +ROBOSPATIAL | 71.6 ↑ | 13.1 ↑ | 72.4 ↑ | 81.8 ↑ | 67.7 ↑ |
| RoboPoint [62] | 69.9 | 19.7 | 70.5 | 63.6 | 44.1 |
| +ROBOSPATIAL | 78.0 ↑ | 31.1 ↑ | 81.0 ↑ | 70.6 ↑ | 64.7 ↑ |
| 3D VLMs | | | | | |
| 3D-LLM [18] | 39.8 | 0.0 | 35.2 | N/A | N/A |
| +ROBOSPATIAL | 55.2 ↑ | 8.2 ↑ | 52.3 ↑ | N/A | N/A |
| LEO [20] | 51.2 | 0.0 | 38.1 | N/A | N/A |
| +ROBOSPATIAL | 64.2 ↑ | 10.0 ↑ | 57.1 ↑ | N/A | N/A |
| <i>Not available for fine-tuning</i> | | | | | |
| Molmo [9] | 58.6 | 0.1 | 18.1 | 67.1 | 55.9 |
| GPT-4o [42] | 77.2 | 5.7 | 58.1 | 76.2 | 70.6 |

Table 4. Results on an out-of-domain test split comparing prior art VLMs. The results show improved (↑) spatial understanding capabilities on similar domains. Bolded number is the best result for the column.

ordinate predictions, we evaluate whether the model’s predicted 3D location lies within the convex hull of a reference point set derived from scene geometry.

While this convex hull criterion provides a well-defined geometric target, it is arguably overly strict—*e.g.*, predictions near but just outside the boundary are marked incorrect. As a result, reported scores represent a conservative estimate of each model’s spatial understanding. ROBOSPATIAL-Val serves as the primary benchmark for comparing

2D and 3D VLMs trained on ROBOSPATIAL, enabling controlled evaluation within the same data distribution. Results are presented in Table 3.

4.1.3. Cross-Dataset Generalization Evaluation

To assess generalization across environment types, we partition the training data into indoor-scene and tabletop subsets. Models are trained on one type and evaluated on held-out datasets from the other. Despite differing object distributions and scene layouts, we observe a positive synergy between indoor and tabletop environments: training on one environment type improves spatial reasoning on the other, as shown in Table 5.

4.1.4. Out-of-Domain Evaluation

To test the out-of-domain transferability of ROBOSPATIAL-trained models, we evaluate on three benchmarks: ROBOSPATIAL-Home, BLINK [15], and SpatialBench [2]. ROBOSPATIAL-Home contains 350 manually written spatial questions over diverse real-world RGBD scenes captured with an iPhone equipped with a depth sensor. We curated this benchmark to evaluate generalization to novel indoor settings with previously unseen objects. BLINK is a visual reasoning benchmark consisting of binary spatial questions involving relationships such as “next to,” “touching,” and “on top.” BLINK allows us to assess the ability of models to assess generalization of spatial reasoning to unseen language configurations. We evaluate only the *spatial* portion of BLINK, as that aligns with the core focus of ROBOSPATIAL. SpatialBench, introduced in the Spa-

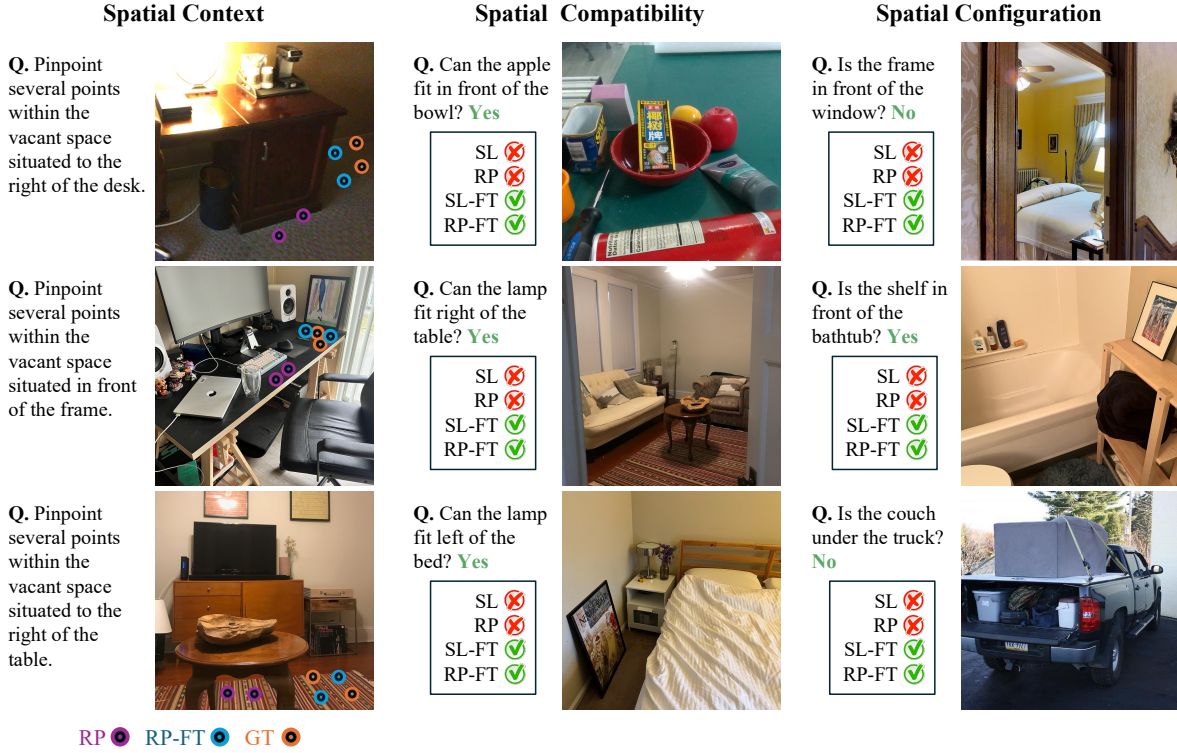


Figure 3. In-domain (ROBOSPATIAL-Val, top) and out-of-domain (ROBOSPATIAL-Home, BLINK [15], middle and bottom) results for ROBOSPATIAL-trained models. Two models shown: SL (SpaceLLaVA [5]) and RP (RoboPoint [62]); the -FT suffix indicates fine-tuning on ROBOSPATIAL. Correct answers in green. All images except bottom-right in the out-of-domain rows are from ROBOSPATIAL-Home.

tialBot [2] paper, uses RGB inputs and tests spatial understanding across several categories. We focus on the *position* category, which most directly aligns with our emphasis on spatial localization and placement. Together, these benchmarks complement ROBOSPATIAL-Val by covering a broader range of visual and linguistic variation.

4.2. Results

We evaluate the effectiveness of ROBOSPATIAL in improving spatial reasoning capabilities in VLMs across held-out and out-of-domain benchmarks. In this section, we focus on analyzing the model’s generalization and understanding of spatial relationships. We address the following questions:

How well does ROBOSPATIAL training generalize to unseen spatial relationships? Although ROBOSPATIAL consists of template-generated QA pairs with a fixed set of spatial prepositions, we observe in Tab. 4 that models trained on it can generalize to spatial relationships not explicitly included in the training set. This is particularly evident in evaluations on the BLINK dataset [15], which contains diverse prepositions such as “under,” “next to,” and “far away.” We attribute this generalization to the fact that ROBOSPATIAL encompasses all six principal directions in 3D space (along the x, y, and z axes). Generalizing to new

prepositions often requires mapping linguistic expressions (e.g., “on top of,” “under”) to these spatial primitives—a task at which LLMs are naturally proficient. For example, “on top of” often refer to “above” in a world-centric frame, while “under” maps to “below.” Moreover, prepositions such as “next to” or “beside” imply proximity between objects. Because ROBOSPATIAL includes questions that require generating points near a reference object, it implicitly teaches the concept of closeness. This enables trained models to understand these proximity-based relationships, even if they are not explicitly represented during training.

Do ROBOSPATIAL-trained models understand nuanced perspectives? Spatial references in natural language often imply specific reference frames. For instance, “in front of the car” typically refers to the direction of the car’s front hood. In ROBOSPATIAL-Home, we omit explicit frame specifications in the questions to evaluate whether models can align with the implicit reference frame intended by the questioner. We find that models trained with ROBOSPATIAL can often infer the correct frame of reference, suggesting that they have learned to associate object geometries and orientations with spatial language. Figure 3 shows examples such as “Is the frame in front of the window?”, where the model accurately identifies the intended spatial relation.

| | Indoor ↓ Tabletop | Tabletop ↓ Indoor |
|----------------|-------------------------|-------------------------|
| RoboPoint [62] | 38.7 | 38.2 |
| +ROBOSPATIAL | 48.9 ↑ | 51.3 ↑ |
| LEO [20] | 41.9 | 43.7 |
| +ROBOSPATIAL | 47.2 ↑ | 54.5 ↑ |

Table 5. Cross-dataset generalization results between indoor and tabletop environments.

| Model | Success (%) |
|--------------------------------------|-------------|
| <i>Open-source</i> | |
| LLaVA-NeXT [35] | 23.7 |
| + ROBOSPATIAL | 52.6 ↑ |
| RoboPoint [62] | 44.7 |
| + ROBOSPATIAL | 46.2 ↑ |
| <i>Not available for fine-tuning</i> | |
| Molmo [9] | 43.8 |
| GPT-4o [42] | 46.9 |

Table 6. Robot experiment results.




Figure 4. Robotics experiments: the red dot shows the model output (if not present, the model failed to provide a valid point in the image); green dots are used to show when a model outputs multiple points. The robot motion generator, cuRobo [52], is used to grasp the item referenced by the generated point. The *spatial*-prefix indicates model trained with ROBOSPATIAL.

Are 3D VLMs better at learning spatial relationships than 2D VLMs? The findings in Tab. 3 suggest that 3D VLMs tend to outperform 2D counterparts in spatial reasoning tasks, likely due to their ability to directly utilize depth information. However, this comparison is not entirely fair: models like 3D-LLM [18] and LEO [20] are pretrained on RGB-D indoor scan datasets, some of which overlap with the environments used in the source datasets (e.g., Matterport3D, ScanNet). This gives them prior exposure to scene geometry and object layouts, which may bias their performance. To support more controlled and fair comparisons in the future, we designed ROBOSPATIAL to be compatible with both 2D and 3D modalities, allowing researchers to investigate the impact of modality, architecture, and pre-training data under unified evaluation protocols.

4.3. Real Robot Experiments

We design a suite of tabletop manipulation tasks requiring spatial reasoning. The setup includes a Kinova Jaco robot [3], paired with a ZED2 camera for RGB-D perception. The robot system implements actions to *pick* and *place* objects on the table using cuRobo [52] for motion planning. Tasks include spatial questions that require a yes/no answer, and pick-and-place instructions that require successfully controlling the robot to complete the task. We adopt a modular design, where the VLM is queried for spatial understanding, and the resulting predictions (e.g., target points) are passed to a separate motion planning system for execution. We use a range of simple, unambiguous objects—colored cubes, cylinders, food items, and toys—to

ensure the challenge lies in spatial understanding rather than object recognition (Figure 4). In total, we conducted over 200 model queries. Details of the questions and scene configurations are provided in the Appendix D.5. We evaluate the following VLMs: LLaVA-NeXT [35] and RoboPoint [62], both with and without ROBOSPATIAL training; and two strong baselines, Molmo [9] and GPT-4o [42]. Table 6 and Figure 4 present the results.

Experiments show that LLaVA-NeXT fine-tuned on ROBOSPATIAL achieves the highest success rate across all models. Training with ROBOSPATIAL enhances spatial understanding in 2D VLMs, enabling the model to correctly interpret instructions such as “place in front of the pony,” where placement is aligned with the pony’s head direction. It also demonstrates sensitivity to object scale, as in the task “place in front of the orange juice box,” where the model places the object at a reasonable distance. In contrast, baseline models such as RoboPoint frequently place objects too far from the target, likely due to limited understanding of spatial proximity. We also observe that spatial failures in 2D VLMs often stem from errors in projecting 2D predictions into 3D. Even a small 2-pixel shift in image space can translate to a 5–10 cm error in the physical world, which is significant in manipulation tasks. Nonetheless, models trained on ROBOSPATIAL produce more accurate predictions, reducing these failure cases and showing the benefit of dataset-driven improvements. Interestingly, GPT-4o performs comparably to ROBOSPATIAL-trained RoboPoint. We attribute this to GPT-4o’s broader language understanding and instruction-following ability, which partially compensates for its lack of task-specific spatial training. Looking forward, promising directions include investigating how viewpoint affects 2D spatial predictions, and developing 3D VLMs that can reason over partial point clouds—removing the need for complete 3D scans and making deployment in real-world systems more feasible.

5. Conclusion

We introduce ROBOSPATIAL, ROBOSPATIAL-Val, and ROBOSPATIAL-Home, a large-scale 2D/3D spatial understanding training and evaluation dataset tailored for robotics. Experimental results show that models trained with ROBOSPATIAL are able to understand spatial relationships, generalize to unseen relationships, and infer nuanced reference frames, making them applicable in a wide range of tasks that require spatial understanding. We further demonstrate the real-world applicability of ROBOSPATIAL with robot experiments. In addition, our automatic data generation pipeline can be used to extend the dataset to new data sources and spatial relations. We show that ROBOSPATIAL has the potential to serve as a foundation for broader applications in robotics which require spatial understanding.

References

- [1] Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. Scanqa: 3d question answering for spatial scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 3
- [2] Wenxiao Cai, Yaroslav Ponomarenko, Jianhao Yuan, Xiaoli Li, Wankou Yang, Hao Dong, and Bo Zhao. Spatialbot: Precise spatial understanding with vision language models. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2025. 2, 3, 6, 7
- [3] Alexandre Campeau-Lecours, Hugo Lamontagne, Simon Latour, Philippe Fauteux, Véronique Maheu, François Boucher, Charles Deguire, and Louis-Joseph Caron L'Ecuyer. Kinova modular robot arms for service robotics applications. *Int. J. Robot. Appl. Technol.*, 5(2):49–71, 2017. 8
- [4] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision (3DV)*, 2017. 5, 16
- [5] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brain Ichter, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. Spatialvlm: Endowing vision-language models with spatial reasoning capabilities. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 14455–14465, 2024. 2, 3, 5, 6, 7, 15, 16
- [6] An-Chieh Cheng, Hongxu Yin, Yang Fu, Qiushan Guo, Ruihan Yang, Jan Kautz, Xiaolong Wang, and Sifei Liu. Spatialrgpt: Grounded spatial reasoning in vision-language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. 2, 3, 5
- [7] Open X-Embodiment Collaboration, Abby O'Neill, Abdul Rehman, Abhinav Gupta, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, Albert Tong, Alex Bewley, Alex Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Ankit Gupta, Andrew Wang, Andrey Kolobov, Anikait Singh, Animesh Garg, Aniruddha Kembhavi, Annie Xie, Anthony Brohan, Antonin Raffin, Archit Sharma, Arefeh Yavary, Arhan Jain, Ashwin Balakrishna, Ayzaan Wahid, Ben Burgess-Limerick, Beomjoon Kim, Bernhard Schölkopf, Blake Wulfe, Brian Ichter, Cewu Lu, Charles Xu, Charlotte Le, Chelsea Finn, Chen Wang, Chenfeng Xu, Cheng Chi, Chenguang Huang, Christine Chan, Christopher Agia, Chuer Pan, Chuyuan Fu, Coline Devin, Danfei Xu, Daniel Morton, Danny Driess, Daphne Chen, Deepak Pathak, Dhruv Shah, Dieter Buechler, Dinesh Jayaraman, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Ethan Foster, Fangchen Liu, Federico Ceola, Fei Xia, Feiyu Zhao, Felipe Vieira Frujeri, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Gilbert Feng, Giulio Schiavi, Glen Berseth, Gregory Kahn, Guangwen Yang, Guanzhi Wang, Hao Su, Hao-Shu Fang, Haochen Shi, Henghui Bao, Heni Ben Amor, Henrik I Christensen, Hiroki Furuta, Homanga Bharadhwaj, Homer Walke, Hongjie Fang, Huy Ha, Igor Mordatch, Ilija Radosavovic, Isabel Leal, Jacky Liang, Jad Abou-Chakra, Jaehyung Kim, Jaimyn Drake, Jan Peters, Jan Schneider, Jasmine Hsu, Jay Vakil, Jeannette Bohg, Jeffrey Bingham, Jeffrey Wu, Jensen Gao, Jiaheng Hu, Jiajun Wu, Jialin Wu, Jiankai Sun, Jianlan Luo, Jiayuan Gu, Jie Tan, Jihoon Oh, Jimmy Wu, Jingpei Lu, Jingyun Yang, Jitendra Malik, João Silvério, Joey Hejna, Jonathan Bozher, Jonathan Tompson, Jonathan Yang, Jordi Salvador, Joseph J. Lim, Junhyek Han, Kaiyuan Wang, Kanishk Rao, Karl Pertsch, Karol Hausman, Keegan Go, Keerthana Gopalakrishnan, Ken Goldberg, Kendra Byrne, Kenneth Oslund, Kento Kawaharazuka, Kevin Black, Kevin Lin, Kevin Zhang, Kiana Ehsani, Kiran Lekkala, Kirsty Ellis, Krishan Rana, Krishnan Srinivasan, Kuan Fang, Kunal Pratap Singh, Kuo-Hao Zeng, Kyle Hatch, Kyle Hsu, Laurent Itti, Lawrence Yunliang Chen, Lerrel Pinto, Li Fei-Fei, Liam Tan, Linxi "Jim" Fan, Lionel Ott, Lisa Lee, Luca Weihs, Magnum Chen, Marion Lepert, Marius Memmel, Masayoshi Tomizuka, Masha Itkina, Mateo Guaman Castro, Max Spero, Maximilian Du, Michael Ahn, Michael C. Yip, Mingtong Zhang, Mingyu Ding, Minh Heo, Mohan Kumar Srirama, Mohit Sharma, Moo Jin Kim, Naoaki Kanazawa, Nicklas Hansen, Nicolas Heess, Nikhil J Joshi, Niko Sunderhauf, Ning Liu, Norman Di Palo, Nur Muhammad Mahi Shafullah, Oier Mees, Oliver Kroemer, Osbert Bastani, Panag R Sanketi, Patrick "Tree" Miller, Patrick Yin, Paul Wohlhart, Peng Xu, Peter David Fagan, Peter Mitrano, Pierre Sermanet, Pieter Abbeel, Priya Sundaresan, Qiuyu Chen, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Mart'in-Mart'in, Rohan Bajjal, Rosario Scalise, Rose Hendrix, Roy Lin, Runjia Qian, Ruohan Zhang, Russell Mendonca, Rutav Shah, Ryan Hoque, Ryan Julian, Samuel Bustamante, Sean Kirmani, Sergey Levine, Shan Lin, Sherry Moore, Shikhar Bahl, Shivin Dass, Shubham Sonawani, Shubham Tulsiani, Shuran Song, Sichun Xu, Siddhant Halder, Siddharth Karamcheti, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Subramanian Ramamoorthy, Sudeep Dasari, Suneel Belkale, Sungjae Park, Suraj Nair, Suvir Mirchandani, Takayuki Osa, Tanmay Gupta, Tatsuya Harada, Tatsuya Matsushima, Ted Xiao, Thomas Kollar, Tianhe Yu, Tianli Ding, Todor Davchev, Tony Z. Zhao, Travis Armstrong, Trevor Darrell, Trinity Chung, Vidhi Jain, Vikash Kumar, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiangyu Chen, Xiaolong Wang, Xinghao Zhu, Xinyang Geng, Xiyuan Liu, Xu Liangwei, Xuanlin Li, Yansong Pang, Yao Lu, Yecheng Jason Ma, Yejin Kim, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Yilin Wu, Ying Xu, Yixuan Wang, Yonatan Bisk, Yongqiang Dou, Yoonyoung Cho, Youngwoon Lee, Yuchen Cui, Yue Cao, Yueh-Hua Wu, Yujin Tang, Yuke Zhu, Yunchu Zhang, Yunfan Jiang, Yunshuang Li, Yunzhu Li, Yusuke Iwasawa, Yutaka Matsuo, Zehan Ma, Zhuo Xu, Zichen Jeff Cui, Zichen Zhang, Zipeng Fu, and Zipeng Lin. Open X-Embodiment: Robotic learning datasets and RT-X models. <https://arxiv.org/abs/2310.08864>, 2023. 1
- [8] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: 640 641 642 643 644 645 646 647 648 649 650 651 652 653 654 655 656 657 658 659 660 661 662 663 664 665 666 667 668 669 670 671 672 673 674 675 676 677 678 679 680 681 682 683 684 685 686 687 688 689 690 691 692 693 694 695

- Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017. 5, 16
- [9] Matt Deitke, Christopher Clark, Sangho Lee, Rohun Tripathi, Yue Yang, Jae Sung Park, Mohammadreza Salehi, Niklas Muennighoff, Kyle Lo, Luca Soldaini, Jiasen Lu, Taira Anderson, Erin Bransom, Kiana Ehsani, Huong Ngo, YenSung Chen, Ajay Patel, Mark Yatskar, Chris Callison-Burch, Andrew Head, Rose Hendrix, Favyen Bastani, Eli VanderBilt, Nathan Lambert, Yvonne Chou, Arnavi Chheda, Jenna Sparks, Sam Skjonsberg, Michael Schmitz, Aaron Samat, Byron Bischoff, Pete Walsh, Chris Newell, Piper Wolters, Tanmay Gupta, Kuo-Hao Zeng, Jon Borchardt, Dirk Groeneveld, Jen Dumas, Crystal Nam, Sophie Lebrecht, Caitlin Wittlif, Carissa Schoenick, Oscar Michel, Ranjay Krishna, Luca Weihs, Noah A. Smith, Hannaneh Hajishirzi, Ross Girshick, Ali Farhadi, and Aniruddha Kembhavi. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *arXiv preprint arXiv:2409.17146*, 2024. 2, 5, 6, 8, 16, 19, 20
- [10] Mengfei Du, Binhao Wu, Zejun Li, Xuanjing Huang, and Zhongyu Wei. EmbSpatial-bench: Benchmarking spatial understanding for embodied tasks with large vision-language models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 346–355, Bangkok, Thailand, 2024. Association for Computational Linguistics. 2, 3
- [11] Jiafei Duan, Wilbert Pumacay, Nishanth Kumar, Yi Ru Wang, Shulin Tian, Wentao Yuan, Ranjay Krishna, Dieter Fox, Ajay Mandlekar, and Yijie Guo. AHA: A vision-language-model for detecting and reasoning over failures in robotic manipulation. In *The Thirteenth International Conference on Learning Representations*, 2025. 3
- [12] Hao-Shu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11444–11453, 2020. 5, 14, 16
- [13] Kuan Fang, Fangchen Liu, Pieter Abbeel, and Sergey Levine. Moka: Open-world robotic manipulation through mark-based visual prompting. *Robotics: Science and Systems (RSS)*, 2024. 1
- [14] Chaoyou Fu, Peixian Chen, Yunhang Shen, Yulei Qin, Mengdan Zhang, Xu Lin, Jinrui Yang, Xiawu Zheng, Ke Li, Xing Sun, et al. Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394*, 2023. 17
- [15] Xingyu Fu, Yushi Hu, Bangzheng Li, Yu Feng, Haoyu Wang, Xudong Lin, Dan Roth, Noah A. Smith, Wei-Chiu Ma, and Ranjay Krishna. Blink: Multimodal large language models can see but not perceive. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 148–166. Springer, 2024. 2, 3, 6, 7, 18, 20
- [16] Xinyu Gong, Sreyas Mohan, Naina Dhingra, Jean-Charles Bazin, Yilei Li, Zhangyang Wang, and Rakesh Ranjan. Mmg-ego4d: Multi-modal generalization in egocentric action recognition. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6481–6491, 2023. 1
- [17] Kristen Grauman, Andrew Westbury, Eugene Byrne, Zachary Chavis, Antonino Furnari, Rohit Girdhar, Jackson Hamburger, Hao Jiang, Miao Liu, Xingyu Liu, Miguel Martin, Tushar Nagarajan, Ilija Radosavovic, Santhosh Kumar Ramakrishnan, Fiona Ryan, Jayant Sharma, Michael Wray, Mengmeng Xu, Eric Zhongcong Xu, Chen Zhao, Siddhant Bansal, Dhruv Batra, Vincent Cartillier, Sean Crane, Tien Do, Morrie Doulaty, Akshay Erapalli, Christoph Feichtenhofer, Adriano Fragomeni, Qichen Fu, Abraham Gebreselasie, Cristina González, James Hillis, Xuhua Huang, Yifei Huang, Wenqi Jia, Weslie Khoo, Jáchym Kolář, Satwik Kotur, Anurag Kumar, Federico Landini, Chao Li, Yanghao Li, Zhenqiang Li, Karttikeya Mangalam, Raghava Modhugu, Jonathan Munro, Tullie Murrell, Takumi Nishiyasu, Will Price, Paola Ruiz Puentes, Merrey Ramazanov, Leda Sari, Kiran Somasundaram, Audrey Southerland, Yusuke Sugano, Ruijie Tao, Minh Vo, Yuchen Wang, Xindi Wu, Takuma Yagi, Ziwei Zhao, Yunyi Zhu, Pablo Arbeláez, David Crandall, Dima Damen, Giovanni Maria Farinella, Christian Fuegen, Bernard Ghanem, Vamsi Krishna Ithapu, C. V. Jawahar, Hanbyul Joo, Kris Kitani, Haizhou Li, Richard Newcombe, Aude Oliva, Hyun Soo Park, James M. Rehg, Yoichi Sato, Jianbo Shi, Mike Zheng Shou, Antonio Torralba, Lorenzo Torresani, Mingfei Yan, and Jitendra Malik. Ego4d: Around the world in 3,000 hours of egocentric video. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18973–18990, 2022. 1
- [18] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. In *Advances in Neural Information Processing Systems*, 2023. NeurIPS. 5, 6, 8, 15, 16
- [19] Haoxu Huang, Fanqi Lin, Yingdong Hu, Shengjie Wang, and Yang Gao. Copa: General robotic manipulation through spatial constraints of parts with foundation models. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 9488–9495, 2024. 3
- [20] Jiangyong Huang, Silong Yong, Xiaojian Ma, Xiongkun Linghu, Puhao Li, Yan Wang, Qing Li, Song-Chun Zhu, Baoxiong Jia, and Siyuan Huang. An embodied generalist agent in 3d world. In *Proceedings of the International Conference on Machine Learning (ICML)*, 2024. 5, 6, 8, 15, 16
- [21] Wenlong Huang, Chen Wang, Yunzhu Li, Ruohan Zhang, and Li Fei-Fei. Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. In *8th Annual Conference on Robot Learning*, 2024. 3
- [22] Yifei Huang, Jilan Xu, Baoqi Pei, Yuping He, Guo Chen, Mingfang Zhang, Lijin Yang, Zheng Nie, Jinyao Liu, Guoshun Fan, Dechen Lin, Fang Fang, Kunpeng Li, Chang Yuan, Xinyuan Chen, Yaohui Wang, Yali Wang, Yu Qiao, and Limin Wang. An egocentric vision-language model based portable real-time smart assistant, 2025. 1
- [23] Drew A. Hudson and Christopher D. Manning. Gqa: A new dataset for real-world visual reasoning and compositional question answering. In *Proceedings of the IEEE/CVF*

- Conference on Computer Vision and Pattern Recognition (CVPR), pages 6700–6709, 2019. 3
- [24] Baoxiong Jia, Yixin Chen, Huangyue Yu, Yan Wang, Xuesong Niu, Tengyu Liu, Qing Li, and Siyuan Huang. Sceneverse: Scaling 3d vision-language learning for grounded scene understanding. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2024.
- [25] Justin Johnson, Bharath Hariharan, Laurens van der Maaten, Li Fei-Fei, C. Lawrence Zitnick, and Ross Girshick. Clevr: A diagnostic dataset for compositional language and elementary visual reasoning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [26] Amita Kamath, Jack Hessel, and Kai-Wei Chang. What’s “up” with vision-language models? investigating their struggle with spatial reasoning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2023. 1, 2, 3
- [27] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan P Foster, Pannag R Sanketi, Quan Vuong, Thomas Kollar, Benjamin Burchfiel, Russ Tedrake, Dorsa Sadigh, Sergey Levine, Percy Liang, and Chelsea Finn. Openvla: An open-source vision-language-action model. In *Proceedings of the 8th Conference on Robot Learning (CoRL)*, pages 2679–2713. PMLR, 2025. 3, 14
- [28] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A Shamma, et al. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *International journal of computer vision*, 123:32–73, 2017. 3
- [29] Jacky Liang, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, and Andy Zeng. Code as policies: Language model programs for embodied control. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 9493–9500, 2023. 1
- [30] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. VILA: On Pre-training for Visual Language Models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26679–26689, Los Alamitos, CA, USA, 2024. IEEE Computer Society. 1, 5, 6, 15, 16
- [31] Tsung-Yi Lin, Michael Maire, Serge Belongie, Lubomir Bourdev, Ross Girshick, James Hays, Pietro Perona, Deva Ramanan, C. Lawrence Zitnick, and Piotr Dollár. Microsoft COCO: Common objects in context. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2014. 3
- [32] Xiongkun Linghu, Jiangyong Huang, Xuesong Niu, Xiaojian Ma, Baoxiong Jia, and Siyuan Huang. Multi-modal situated reasoning in 3d scenes. In *Advances in Neural Information Processing Systems*, 2024. NeurIPS. 3
- [33] Fangyu Liu, Guy Emerson, and Nigel Collier. Visual spatial reasoning. *Transactions of the Association for Computational Linguistics*, 11:635–651, 2023. 2, 3
- [34] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Advances in Neural Information Processing Systems*, 2023. 15
- [35] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 26296–26306, 2024. 1, 5, 6, 8, 15, 16, 19, 20
- [36] Yuan Liu, Haodong Duan, Yuanhan Zhang, Bo Li, Songyang Zhang, Wangbo Zhao, Yike Yuan, Jiaqi Wang, Conghui He, Ziwei Liu, Kai Chen, and Dahua Lin. Mmbench: Is your multi-modal model an all-around player? In *European Conference on Computer Vision (ECCV)*, 2024. 17
- [37] Xiaojian Ma, Silong Yong, Zilong Zheng, Qing Li, Yitao Liang, Song-Chun Zhu, and Siyuan Huang. Sqa3d: Situated question answering in 3d scenes. In *International Conference on Learning Representations*, 2023. 3
- [38] Arjun Majumdar, Anurag Ajay, Xiaohan Zhang, Pranav Putta, Sriram Yenamandra, Mikael Henaff, Sneha Silwal, Paul Mcvay, Oleksandr Maksymets, Sergio Arnaud, Karmesh Yadav, Qiyang Li, Ben Newman, Mohit Sharma, Vincent Berges, Shiqi Zhang, Pulkit Agrawal, Yonatan Bisk, Dhruv Batra, Mrinal Kalakrishnan, Franziska Meier, Chris Paxton, Sasha Sax, and Aravind Rajeswaran. Openeqa: Embodied question answering in the era of foundation models. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 1
- [39] Yunze Man, Liang-Yan Gui, and Yu-Xiong Wang. Situational awareness matters in 3d vision language reasoning. In *CVPR*, 2024. 3
- [40] Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, Jacky Liang, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, Quan Vuong, Tingnan Zhang, Tsang-Wei Edward Lee, Kuang-Huei Lee, Peng Xu, Sean Kirmani, Yuke Zhu, Andy Zeng, Karol Hausman, Nicolas Heess, Chelsea Finn, Sergey Levine, and Brian Ichter. Pivot: iterative visual prompting elicits actionable knowledge for vlms. In *Proceedings of the International Conference on Machine Learning (ICML)*. JMLR.org, 2024. 3
- [41] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Charles Xu, Jianlan Luo, Tobias Kreiman, You Liang Tan, Lawrence Yunliang Chen, Pannag Sanketi, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, 2024. 3
- [42] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, Irwan Bello, Jake Berdine, Gabriel Bernadett-Shapiro, Christopher Berner, Lenny Bogdonoff, Oleg Boiko, Madelaine Boyd, Anna-Luisa Brakman, Greg Brockman, Tim Brooks, Miles Brundage, Kevin Button, Trevor Cai, Rosie Campbell, Andrew Cann, Brittany Carey, Chelsea Carlson, Rory Carmichael, Brooke Chan, Che Chang, Fotis Chantzis, Derek Chen, Sully Chen, Ruby Chen, Jason Chen, Mark Chen, Ben Chess, Chester Cho, Casey Chu, Hyung Won Chung, Dave Cummings, Jeremiah Cur-

| | | | |
|-----|---|--|------|
| 928 | rier, Yunxing Dai, Cory Decareaux, Thomas Degry, Noah | Tao Xu, Sarah Yoo, Kevin Yu, Qiming Yuan, Wojciech | 986 |
| 929 | Deutsch, Damien Deville, Arka Dhar, David Dohan, Steve | Zaremba, Rowan Zellers, Chong Zhang, Marvin Zhang, | 987 |
| 930 | Dowling, Sheila Dunning, Adrien Ecoffet, Atty Eleti, | Shengjia Zhao, Tianhao Zheng, Juntang Zhuang, William | 988 |
| 931 | Tyna Eloundou, David Farhi, Liam Fedus, Niko Felix, | Zhuk, and Barret Zoph. Gpt-4 technical report, 2024. 1, | 989 |
| 932 | Simón Posada Fishman, Juston Forte, Isabella Fulford, Leo | 3, 5, 6, 8, 16, 18, 19, 20 | 990 |
| 933 | Gao, Elie Georges, Christian Gibson, Vik Goel, Tarun Gogi- | [43] Navid Rajabi and Jana Kosecka. Towards grounded visual | 991 |
| 934 | neni, Gabriel Goh, Rapha Gontijo-Lopes, Jonathan Gordon, | spatial reasoning in multi-modal vision language models, | 992 |
| 935 | Morgan Grafstein, Scott Gray, Ryan Greene, Joshua Gross, | 2024. 3 | 993 |
| 936 | Shixiang Shane Gu, Yufei Guo, Chris Hallacy, Jesse Han, | [44] Kanchana Ranasinghe, Satya Narayan Shukla, Omid Pour- | 994 |
| 937 | Jeff Harris, Yuchen He, Mike Heaton, Johannes Heidecke, | saeed, Michael S. Ryoo, and Tsung-Yu Lin. Learning to lo- | 995 |
| 938 | Chris Hesse, Alan Hickey, Wade Hickey, Peter Hoeschele, | calize objects improves spatial reasoning in visual-llms. In | 996 |
| 939 | Brandon Houghton, Kenny Hsu, Shengli Hu, Xin Hu, Joost | <i>Proceedings of IEEE/CVF Conference on Computer Vision</i> | 997 |
| 940 | Huizinga, Shantanu Jain, Shawn Jain, Joanne Jang, Angela | <i>and Pattern Recognition (CVPR)</i> , pages 12977–12987, 2024. | 998 |
| 941 | Jiang, Roger Jiang, Haozhun Jin, Denny Jin, Shino Jomoto, | 3 | 999 |
| 942 | Billie Jonn, Heewoo Jun, Tomer Kaftan, Łukasz Kaiser, | [45] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang | 1000 |
| 943 | Ali Kamali, Ingmar Kanitscheider, Nitish Shirish Keskar, | Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman | 1001 |
| 944 | Tabarak Khan, Logan Kilpatrick, Jong Wook Kim, Christina | Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junt- | 1002 |
| 945 | Kim, Yongjik Kim, Jan Hendrik Kirchner, Jamie Kiros, | ing Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao- | 1003 |
| 946 | Matt Knight, Daniel Kokotajlo, Łukasz Kondraciuk, An- | Yuan Wu, Ross Girshick, Piotr Dollar, and Christoph Feicht- | 1004 |
| 947 | drew Kondrich, Aris Konstantinidis, Kyle Kopic, Gretchen | enhofer. SAM 2: Segment anything in images and videos. In | 1005 |
| 948 | Krueger, Vishal Kuo, Michael Lampe, Ikai Lan, Teddy Lee, | <i>The Thirteenth International Conference on Learning Rep-</i> | 1006 |
| 949 | Jan Leike, Jade Leung, Daniel Levy, Chak Ming Li, Rachel | <i>resentations</i> , 2025. 17 | 1007 |
| 950 | Lim, Molly Lin, Stephanie Lin, Mateusz Litwin, Theresa | [46] Leonard Salewski, A. Sophia Koepke, Hendrik P. A. Lensch, | 1008 |
| 951 | Lopez, Ryan Lowe, Patricia Lue, Anna Makanju, Kim Mal- | and Zeynep Akata. Clevr-x: A visual reasoning dataset for | 1009 |
| 952 | facini, Sam Manning, Todor Markov, Yaniv Markovski, | natural language explanations. In <i>xxAI - Beyond explainable</i> | 1010 |
| 953 | Bianca Martin, Katie Mayer, Andrew Mayne, Bob Mc- | <i>Artificial Intelligence</i> , pages 85–104. Springer, 2022. 3 | 1011 |
| 954 | Grew, Scott Mayer McKinney, Christine McLeavey, Paul | [47] Fatemeh Shiri, Xiao-Yu Guo, Mona Golestan Far, Xin Yu, | 1012 |
| 955 | McMillan, Jake McNeil, David Medina, Aalok Mehta, Jacob | Reza Haf, and Yuan-Fang Li. An empirical analysis on spa- | 1013 |
| 956 | Menick, Luke Metz, Andrey Mishchenko, Pamela Mishkin, | tial reasoning capabilities of large multimodal models. In | 1014 |
| 957 | Vinnie Monaco, Evan Morikawa, Daniel Mossing, Tong Mu, | <i>Proceedings of the 2024 Conference on Empirical Meth-</i> | 1015 |
| 958 | Mira Murati, Oleg Murk, David Mély, Ashvin Nair, Rei- | <i>ods in Natural Language Processing</i> , pages 21440–21455, | 1016 |
| 959 | ichiro Nakano, Rajeev Nayak, Arvind Neelakantan, Richard | Miami, Florida, USA, 2024. Association for Computational | 1017 |
| 960 | Ngo, Hyeonwoo Noh, Long Ouyang, Cullen O’Keefe, Jakob | Linguistics. 1, 2, 3 | 1018 |
| 961 | Pachocki, Alex Paino, Joe Palermo, Ashley Pantuliano, Gi- | [48] Ishika Singh, Valts Blukis, Arsalan Mousavian, Ankit Goyal, | 1019 |
| 962 | ambattista Parascandolo, Joel Parish, Emy Parparita, Alex | Danfei Xu, Jonathan Tremblay, Dieter Fox, Jesse Thomason, | 1020 |
| 963 | Passos, Mikhail Pavlov, Andrew Peng, Adam Perelman, Fil- | and Animesh Garg. Progprompt: Generating situated robot | 1021 |
| 964 | ipe de Avila Belbute Peres, Michael Petrov, Henrique Ponde | task plans using large language models. In <i>2023 IEEE In-</i> | 1022 |
| 965 | de Oliveira Pinto, Michael, Pokorný, Michelle Pokrass, | <i>ternational Conference on Robotics and Automation (ICRA)</i> , | 1023 |
| 966 | Vitchyr H. Pong, Tolly Powell, Alethea Power, Boris Power, | pages 11523–11530, 2023. 1 | 1024 |
| 967 | Elizabeth Proehl, Raul Puri, Alec Radford, Jack Rae, Aditya | [49] Chan Hee Song, Jihyung Kil, Tai-Yu Pan, Brian M. Sadler, | 1025 |
| 968 | Ramesh, Cameron Raymond, Francis Real, Kendra Rim- | Wei-Lun Chao, and Yu Su. One step at a time: Long-horizon | 1026 |
| 969 | bach, Carl Ross, Bob Rotsted, Henri Roussez, Nick Ryder, | vision-and-language navigation with milestones. In <i>Proceed-</i> | 1027 |
| 970 | Mario Saltarelli, Ted Sanders, Shibani Santurkar, Girish Sas- | <i>ings of the IEEE/CVF Conference on Computer Vision and</i> | 1028 |
| 971 | try, Heather Schmidt, David Schnurr, John Schulman, Daniel | <i>Pattern Recognition (CVPR)</i> , pages 15482–15491, 2022. 3 | 1029 |
| 972 | Selsam, Kyla Sheppard, Toki Sherbakov, Jessica Shieh, | [50] Alessandro Suglia, Claudio Greco, Katie Baker, Jose L. Part, | 1030 |
| 973 | Sarah Shoker, Pranav Shyam, Szymon Sidor, Eric Sigler, | Ioannis Papaioannou, Arash Eshghi, Ioannis Konstas, and | 1031 |
| 974 | Maddie Simens, Jordan Sitkin, Katarina Slama, Ian Sohl, | Oliver Lemon. AlanaVLM: A multimodal embodied AI | 1032 |
| 975 | Benjamin Sokolowsky, Yang Song, Natalie Staudacher, Fe- | foundation model for egocentric video understanding. In | 1033 |
| 976 | lipe Petroski Such, Natalie Summers, Ilya Sutskever, Jie | <i>Findings of the Association for Computational Linguistics:</i> | 1034 |
| 977 | Tang, Nikolas Tezak, Madeleine B. Thompson, Phil Tillet, | <i>EMNLP 2024</i> , pages 11101–11122, Miami, Florida, USA, | 1035 |
| 978 | Amin Tootoonchian, Elizabeth Tseng, Preston Tuggle, Nick | 2024. Association for Computational Linguistics. 1 | 1036 |
| 979 | Turley, Jerry Tworek, Juan Felipe Cerón Uribe, Andrea Val- | [51] Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Hua- | 1037 |
| 980 | lone, Arun Vijayvergiya, Chelsea Voss, Carroll Wainwright, | jun Bai, and Yoav Artzi. A corpus for reasoning about natu- | 1038 |
| 981 | Justin Jay Wang, Alvin Wang, Ben Wang, Jonathan Ward, | ral language grounded in photographs. In <i>Proceedings of the</i> | 1039 |
| 982 | Jason Wei, CJ Weinmann, Akila Welihinda, Peter Welin- | <i>57th Annual Meeting of the Association for Computational</i> | 1040 |
| 983 | der, Jiayi Weng, Lilian Weng, Matt Wiethoff, Dave Willner, | <i>Linguistics</i> , pages 6418–6428, Florence, Italy, 2019. Associ- | 1041 |
| 984 | Clemens Winter, Samuel Wolrich, Hannah Wong, Lauren | ation for Computational Linguistics. 3 | 1042 |
| 985 | Workman, Sherwin Wu, Jeff Wu, Michael Wu, Kai Xiao, | | |

- [52] Balakumar Sundaralingam, Siva Kumar Sastry Hari, Adam Fishman, Caelan Garrett, Karl Van Wyk, Valts Blukis, Alexander Millane, Helen Oleynikova, Ankur Handa, Fabio Ramos, et al. cuRobo: Parallelized collision-free minimum-jerk robot motion generation. *arXiv preprint arXiv:2310.17274*, 2023. 8, 17
- [53] Emilia Szymanska, Mihai Dusmanu, Jan-Willem Buurlage, Mahdi Rad, and Marc Pollefeys. Space3D-Bench: Spatial 3D Question Answering Benchmark. In *European Conference on Computer Vision (ECCV) Workshops*, 2024. 3
- [54] Stephen Tyree, Jonathan Tremblay, Thang To, Jia Cheng, Terry Mosier, Jeffrey Smith, and Stan Birchfield. 6-dof pose estimation of household objects for robotic manipulation: An accessible dataset and benchmark. In *International Conference on Intelligent Robots and Systems (IROS)*, 2022. 5, 14, 16
- [55] Naoki Wake, Atsushi Kanehira, Kazuhiro Sasabuchi, Jun Takamatsu, and Katsushi Ikeuchi. Gpt-4v(ision) for robotics: Multimodal task planning from human demonstration. *IEEE Robotics and Automation Letters*, 9(11):10567–10574, 2024. 3
- [56] Johanna Wald, Armen Avetisyan, Nassir Navab, Federico Tombari, and Matthias Niessner. Rio: 3d object instance re-localization in changing indoor environments. In *Proceedings IEEE International Conference on Computer Vision (ICCV)*, 2019. 5, 16
- [57] Jiayu Wang, Yifei Ming, Zhenmei Shi, Vibhav Vineet, Xin Wang, Yixuan Li, and Neel Joshi. Is a picture worth a thousand words? delving into spatial reasoning for vision language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 3
- [58] Tai Wang, Xiaohan Mao, Chenming Zhu, Runsen Xu, Ruiyuan Lyu, Peisen Li, Xiao Chen, Wenwei Zhang, Kai Chen, Tianfan Xue, Xihui Liu, Cewu Lu, Dahua Lin, and Jiangmiao Pang. Embodiedscan: A holistic multi-modal 3d perception suite towards embodied ai. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 3, 5
- [59] Youngsun Wi, Mark Van der Merwe, Pete Florence, Andy Zeng, and Nima Fazeli. Calamari: Contact-aware and language conditioned spatial action mapping for contact-rich manipulation. In *7th Annual Conference on Robot Learning*, 2023. 3
- [60] Liuchang Xu, Shuo Zhao, Qingming Lin, Luyao Chen, Qianqian Luo, Sensen Wu, Xinyue Ye, Hailin Feng, and Zhenhong Du. Evaluating large language models on spatial tasks: A multi-task benchmarking study, 2024. 3
- [61] Yutaro Yamada, Yihan Bao, Andrew K. Lampinen, Jungo Kasai, and Ilker Yildirim. Evaluating spatial understanding of large language models, 2024. 1, 3
- [62] Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial affordance prediction in robotics. In *8th Annual Conference on Robot Learning*, 2024. 2, 3, 5, 6, 7, 8, 15, 16, 17, 19, 20
- [63] Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoyi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, Cong Wei, Botao Yu, Ruibin Yuan, Renliang Sun, Ming Yin, Boyuan Zheng, Zhenzhu Yang, Yibo Liu, Wenhao Huang, Huan Sun, Yu Su, and Wenhua Chen. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024. 17
- [64] Yue Zhang, Zhiyang Xu, Ying Shen, Parisa Kordjamshidi, and Lifu Huang. SPARTUN3d: Situated spatial understanding of 3d world in large language model. In *The Thirteenth International Conference on Learning Representations*, 2025. 3
- [65] Brianna Zitkovich, Tianhe Yu, Sichun Xu, Peng Xu, Ted Xiao, Fei Xia, Jialin Wu, Paul Wohlhart, Stefan Welker, Ayzaan Wahid, Quan Vuong, Vincent Vanhoucke, Huong Tran, Radu Soricut, Anikait Singh, Jaspiar Singh, Pierre Sermanet, Pannag R. Sanketi, Grecia Salazar, Michael S. Ryoo, Krista Reymann, Kanishka Rao, Karl Pertsch, Igor Mordatch, Henryk Michalewski, Yao Lu, Sergey Levine, Lisa Lee, Tsang-Wei Edward Lee, Isabel Leal, Yuheng Kuang, Dmitry Kalashnikov, Ryan Julian, Nikhil J. Joshi, Alex Irpan, Brian Ichter, Jasmine Hsu, Alexander Herzog, Karol Hausman, Keerthana Gopalakrishnan, Chuyuan Fu, Pete Florence, Chelsea Finn, Kumar Avinava Dubey, Danny Driess, Tianli Ding, Krzysztof Marcin Choromanski, Xi Chen, Yevgen Chebotar, Justice Carbajal, Noah Brown, Anthony Brohan, Montserrat Gonzalez Arenas, and Kehang Han. Rt-2: Vision-language-action models transfer web knowledge to robotic control. In *Proceedings of The 7th Conference on Robot Learning*, pages 2165–2183. PMLR, 2023. 3
- [66] Alex Zook, Fan-Yun Sun, Josef Spjut, Valts Blukis, Stan Birchfield, and Jonathan Tremblay. Grs: Generating robotic simulation tasks from real-world images, 2024. 1

Appendices

In this supplementary material, we present additional details and clarifications that are omitted in the main text due to space constraints.

- [Appendix A](#) Limitations.
- [Appendix B](#) Dataset Details.
- [Appendix C](#) Implementation Details.
- [Appendix D](#) More Results.

A. Limitations

While ROBOSPATIAL significantly improves spatial reasoning capabilities in VLMs, certain design choices naturally introduce trade-offs and areas for future exploration.

First, the dataset relies on a top-down occupancy map to identify and annotate empty regions for spatial context and compatibility tasks. This approach simplifies reasoning about object placement on horizontal surfaces and enables efficient data generation, but it currently does not support spatial questions involving containment—such as whether an object can fit inside or under another object—which would require more detailed volumetric modeling.

Second, although the models are deployed on a real robot using a modular approach, we do not yet explore tighter forms of integration such as training it jointly with robot trajectories [27]. Investigating these alternatives could enhance downstream policy learning and enable more seamless end-to-end systems.

Finally, ROBOSPATIAL focuses on indoor and tabletop scenes containing objects commonly encountered in household environments, and does not include humans or animals. This reflects the nature of source datasets and our emphasis on robot object manipulation. While this limits coverage of social or dynamic interaction scenarios, trained models still generalize well to out-of-distribution benchmarks like BLINK, which include humans and animals—suggesting that the learned spatial representations are broadly transferable.

B. Dataset Details

B.1. Dataset Statistics

We provide the full dataset statistics in Tab. 7. For all training, we use only 900,000 spatial relationships, sampled equally across all datasets, due to computational constraints. We further experiment on the effect of data scaling on Tab. 9 and explain the results. Notably, HOPE [54] and GraspNet-1B [12] contain similar tabletop images captured from different perspectives, resulting in lower dataset diversity for the tabletop environment. We plan to enhance the diversity of ROBOSPATIAL by incorporating additional tabletop datasets.

B.2. Choice of Spatial Relationships

In designing the dataset, we focused on spatial relationships that directly impact robotic perception, planning, and interaction: context, compatibility, and configuration. These were selected to reflect the core spatial reasoning challenges that robots encounter when operating in complex, real-world environments.

We intentionally excluded tasks such as object counting, as we consider them to fall outside the scope of spatial understanding. While counting is an important visual reasoning skill, it does not require reasoning about spatial relations between objects or between objects and their environment. For example, determining that “three cups are on the table” is a perceptual task rather than a spatial reasoning one. As such, counting may complement but does not substitute for the types of relational reasoning we target. We leave the integration of counting tasks into spatial benchmarks as future work.

Similarly, we exclude tasks that rely solely on distance measurements. Although distance is a fundamental spatial quantity, it is difficult to define consistently across different environments, object scales, and robot embodiments. Absolute distances can vary significantly between indoor and outdoor scenes, small and large objects, or different robot perspectives, making them hard to normalize or interpret in a general way. Moreover, distance alone often lacks the relational semantics required for higher-level reasoning—for example, understanding that an object is behind, above, or in front of others. ROBOSPATIAL instead focuses on spatial relationships that are more invariant, interpretable, and transferable across diverse robotic scenarios.

That said, the data generation pipeline is general and could readily support auxiliary tasks involving object counting or distance estimation if desired. These metrics may serve as useful complements in future extensions of the benchmark or as auxiliary supervision signals in model training.

B.3. Object Grounding Dataset

To support accurate spatial understanding, we generate an auxiliary dataset for object grounding. Many spatial reasoning tasks assume that the model can correctly identify which object is being referred to in the scene. However, in practice, this can be a major source of error—especially in cluttered environments or when multiple instances of the similar object type are present.

The grounding dataset provides direct supervision to help models learn to associate text descriptions with specific objects in the image. For each image, we include a set of object descriptions (*e.g.*, “the keyboard” or “the chair”) paired with the corresponding 2D bounding box of the object in the image. These 2D boxes are projected from the annotated 3D bounding boxes using camera intrinsics and




Figure 5. An example of generated top-down map of the image from 3D bounding boxes.

extrinsics.

A total of 100k grounding QA pairs are generated and used during training to reduce reference ambiguity and improve object identification accuracy in spatial tasks. While not part of the main spatial reasoning taxonomy, grounding accuracy is a prerequisite for answering spatial questions correctly, and we find that including this data helps reduce errors caused by incorrect object identification.

B.4. Dataset Generation Details

The dataset generation pipeline is detailed in the main text (subsection 3.2), which introduces a two-stage process for computing 3D spatial relationships and projecting them into 2D image space. Here, we expand on implementation details not covered in the main paper and provide clarification on the reasoning logic used in spatial annotation.

Reference Frame Annotation. For each spatial configuration question, we label relationships from three perspectives: ego-centric (camera view), object-centric (based on object heading), and world-centric (aligned with the dataset’s global frame). To compute object-centric directions, we use the heading vector of each oriented 3D bounding box to define the “front” of the object. Left, right, behind, and front relations are then assigned accordingly. World-centric annotations modify vertical relationships (above/below) using global z -coordinates to reflect elevation.

Surface Detection and Free Space Sampling. To identify support surfaces such as tables, counters, or floors, we use GPT-4o to select candidate objects that are likely to support placement. A top-down occupancy map is constructed from bounding boxes in the scene Fig. 5. We sample 3D points in unoccupied regions and project them into the image plane for spatial context tasks. Points are filtered via occlusion checks using raycasting, ensuring sampled points are visible and unobstructed.

Compatibility Check and Object Placement. For spatial compatibility, we simulate placing a virtual object bounding box at candidate locations. The placement must fit without intersecting other objects and must allow a clearance of

at least 10 cm in all axes. We allow in-plane rotation and translation to test flexible placement. This provides a binary label (True/False) indicating whether the object can be compatibly placed in the region.

Output Format. Though ROBOSPATIAL uses point prediction for ease of integration with robot setups, the pipeline also supports mask-based outputs and can be extended in future work.

C. Implementation Details

C.1. Model Training

We further explain the training details for all 2D and 3D VLMs trained on ROBOSPATIAL. For all models, we perform instruction tuning using the model weights from public repositories. All training is done using 8 Nvidia H100 GPUs, with the training time between 20 and 40 hours.

C.2. Model Setup

VILA [30] We initialize the model from Efficient-Large-Model/Llama-3-VILA1.5-8B on Hugging Face. We use the fine-tuning script from the VILA GitHub repository to train the model using the default hyperparameters.

LLaVA-NeXT [35] We initialize the model from Imms-lab/llama3-llava-next-8b on Hugging Face. We use the LLaVA-Next fine-tuning script from the LLaVA-Next repository using the default hyperparameters.

SpaceLLaVA [5] As official code and weights for SpatialVLM [5] is not released, we use a community implementation which is endorsed by SpatialVLM [5] authors. We initialize the model from remyxai/SpaceLLaVA from Hugging Face. We use LLaVA-1.5 finetuning script from LLaVa [34] repository using the default hyperparameters.

RoboPoint [62] We initialize the model from wentao-yuan/robopoint-v1-vicuna-v1.5-13b on Hugging Face. We use the fine-tuning script provided in the RoboPoint [62] GitHub repository to train the model using the default hyperparameters.

3D-LLM [18] We initialize the model using the pretrain_blip2_sam_flant5x1_v2.pth checkpoint downloaded from the official GitHub repository. Since the model requires preprocessing of multiview images, we follow the author’s pipeline to process multiview images from the environments. Because the model does not accept image input, we append the following text in front of the question to ensure the model understands the perspective from which the question is being asked: “I am facing ANCHOR OBJECT.” We use the default hyperparameters and train the model for 20 epochs per the author’s guidelines. We choose the best model based on validation accuracy.

LEO [20] We initialize the model from the sft_noact.pth checkpoint downloaded from the official GitHub repository. Since LEO supports dual image and 3D point cloud in-

| Category | Dataset | Split | Scans | Images | Configuration Q | Context Q | Compatibility Q |
|----------|------------------|------------|------------|--------|-----------------|-----------|-----------------|
| Indoor | Matterport3D [4] | Train | 1859 scans | 236243 | 298439 | 298439 | 298439 |
| | | Validation | 10 scans | 200 | 200 | 200 | 200 |
| | ScanNet [8] | Train | 1514 scans | 280402 | 299039 | 299039 | 299039 |
| | | Validation | 12 scans | 400 | 400 | 400 | 400 |
| | 3RScan [56] | Train | 1543 scans | 366755 | 298839 | 298839 | 298839 |
| | | Validation | 18 scans | 400 | 400 | 400 | 400 |
| Tabletop | HOPE [54] | Train | 60 scenes | 50050 | 36817 | 36817 | 36817 |
| | | Validation | 47 scenes | 235 | 500 | 500 | 500 |
| | GraspNet-1B [12] | Train | 130 scenes | 25620 | 36817 | 36817 | 36817 |
| | | Validation | 30 scenes | 120 | 500 | 500 | 500 |

Table 7. Full dataset statistics for indoor and tabletop datasets.

| Model | Indoor | | | Tabletop | | | Average | | |
|--------------------------------------|--------------|----------------|---------------|--------------|----------------|---------------|--------------|--------------|--------------|
| | Ego-centric | Object-centric | World-centric | Ego-centric | Object-centric | World-centric | Indoor | Tabletop | Total |
| <i>Open-source VLMs</i> | | | | | | | | | |
| 2D VLMs | | | | | | | | | |
| VILA [30] | 55.9 | 40.5 | 32.9 | 43.6 | 39.7 | 28.9 | 43.1 | 37.4 | 40.2 |
| +ROBOSPATIAL | 74.3↑ | 57.8↑ | 62.3↑ | 70.3↑ | 58.1↑ | 60.3↑ | 64.8↑ | 62.9↑ | 63.9↑ |
| LLaVA-Next [35] | 35.2 | 24.3 | 34.7 | 36.4 | 28.5 | 22.7 | 31.4 | 29.2 | 30.3 |
| +ROBOSPATIAL | 75.4↑ | 54.1↑ | 68.8↑ | 67.9↑ | 54.7↑ | 58.9↑ | 60.4↑ | 60.5↑ | 60.5↑ |
| SpaceLLaVA [5] | 40.6 | 36.0 | 30.1 | 52.3 | 32.8 | 53.5 | 38.9 | 46.2 | 43.6 |
| +ROBOSPATIAL | 78.5↑ | 60.6↑ | 64.3↑ | 73.0↑ | 49.5↑ | 68.3↑ | 67.8↑ | 63.6↑ | 65.7↑ |
| RoboPoint [62] | 41.9 | 36.2 | 40.7 | 46.2 | 30.5 | 37.9 | 39.6 | 38.2 | 38.9 |
| +ROBOSPATIAL | 76.4↑ | 58.3↑ | 78.3↑ | 76.7↑ | 62.6↑ | 71.0↑ | 71.0↑ | 70.1↑ | 70.6↑ |
| 3D VLMs | | | | | | | | | |
| 3D-LLM [18] | 28.9 | 38.3 | 45.6 | 38.9 | 35.7 | 52.6 | 37.6 | 42.4 | 40.0 |
| +ROBOSPATIAL | 60.7↑ | 52.1↑ | 76.5↑ | 57.9↑ | 62.8↑ | 77.3↑ | 63.1↑ | 66.0↑ | 64.6↑ |
| LEO [20] | 46.9 | 30.6 | 48.2 | 41.4 | 34.3 | 55.4 | 41.9 | 43.7 | 42.8 |
| +ROBOSPATIAL | 68.1↑ | 71.6↑ | 79.6↑ | 71.4↑ | 60.2↑ | 80.5↑ | 73.1↑ | 70.7↑ | 71.9↑ |
| <i>Not available for fine-tuning</i> | | | | | | | | | |
| 2D VLMs | | | | | | | | | |
| Molmo [9] | 50.4 | 50.8 | 47.6 | 64.4 | 33.6 | 53.8 | 49.6 | 50.6 | 50.1 |
| GPT-4o [42] | 52.9 | 38.7 | 56.3 | 62.5 | 30.7 | 63.7 | 49.3 | 52.3 | 50.8 |

Table 8. Results of per frame accuracy of existing 2D/3D VLMs on a ROBOSPATIAL-Val. All methods, for all tasks, perform better (↑) when fine-tuned on ROBOSPATIAL. The best result for each column is bolded.

| | 100K | 300K | 900k (Default) | 1.8M | 3M (Full) |
|-----------------|------|------|----------------|------|-----------|
| LLaVa-Next [35] | 38.1 | 46.7 | 60.5 | 65.8 | 72.4 |

Table 9. Results of scaling experiment on LLaVa-Next [35] with varied number of spatial relationship annotations. Average accuracy on ROBOSPATIAL-Val is reported.

| | MMMU _{val} | MME _p | MME _c | MMBench _{dev} |
|--------------|---------------------|------------------|------------------|------------------------|
| LLaVA-NeXT | 39.4 | 1561.8 | 305.4 | 71.6 |
| +ROBOSPATIAL | 39.8 | 1604.5 | 293.2 | 71.6 |

Table 10. Evaluation on general-purpose multimodal benchmarks (MMMU, MME, MMBench) to assess whether training on ROBOSPATIAL affects commonsense and factual reasoning.

put, we input both of them and modify the question as in 3D-LLM. We use the default hyperparameters and train the model for 10 epochs per the author’s guidelines, and choose the best model based on validation accuracy.

We could not fine-tune Molmo [9] from allenai/Molmo-

7B-D-0924 or GPT-4o [42] from the gpt-4o-2024-08-06 API due to the unavailability of the fine-tuning script at the time of this work, thus we use them as a zero-shot baselines.

| | Base | Auxiliary | ROBOSPATIAL | Both |
|------------|------|-----------|-------------|------|
| LLaVA-NeXT | 30.3 | 32.4 | 51.8 | 60.5 |

Table 11. Ablation study evaluating the impact of the auxiliary grounding dataset on ROBOSPATIAL-Val.

D. More Results

D.1. Accuracy Per Reference Frame

We show the results per frame in Tab. 8 for ROBOSPATIAL-Val. From the results, we can see a distinct difference between 2D and 3D VLMs in understanding the world-centric frame before training with ROBOSPATIAL. Baseline 2D VLMs have trouble understanding the world-centric frame, which involves understanding elevation, while 3D VLMs comparatively excel at it. Furthermore, we can see that since baseline 3D VLMs are trained on point clouds without information of perspective, their accuracy in ego-centric and object-centric frames is lower. However, with ROBOSPATIAL training, we were able to teach the 3D VLMs to think in a certain frame, thus considerably improving their performance on ego-centric and object-centric frames. However, we hypothesize that, due to their design—specifically, the lack of a means to visually inject perspective information since they require complete 3D point clouds—3D VLMs still lag behind 2D VLMs on ego-centric and object-centric frames.

D.2. Data Scaling

In Tab. 9, we experiment with scaling the number of annotations while keeping images fixed. We found that even though the number of images stays consistent, increasing the number of annotations can improve performance. For future work, we plan to apply the data generation pipeline to a diverse set of indoor and tabletop environments to further improve the performance of the models.

D.3. Commonsense Knowledge Retention

To ensure that training on ROBOSPATIAL does not degrade a model’s general reasoning or commonsense capabilities, we evaluate the RoboSpatial-trained model on a suite of standard multimodal benchmarks: MMMU [63], MME [14], and MMBench [36]. As shown in Table 10, the ROBOSPATIAL-trained model maintains or slightly improves performance across all benchmarks, suggesting that spatial fine-tuning preserves broader knowledge capabilities.

D.4. Ablation of the Auxiliary Grounding Dataset

As shown in Table 11, training on the auxiliary dataset alone yields a small improvement over the base model (+2.1), but it falls far short of the gains achieved with ROBOSPATIAL, which is explicitly designed to teach spatial reason-

ing. This confirms that grounding supervision alone is insufficient for spatial understanding. However, combining both datasets leads to the best performance, suggesting that improving object localization can complement spatial supervision when jointly trained.

D.5. Robot Experiments Details

D.5.1. Robot Setup

For picking, we find which object the point maps to using SAM 2 [45] and execute the picking behavior on that object. For placing, we simply compute the 3D coordinate based on the depth value at that pixel and place the object at that coordinate. There were no failures due to cuRobo [52] failing. The experiments were purposely designed to consist of behaviors that our robot system can handle in order to avoid introducing irrelevant factors. The picking behavior consists of computing a top-down grasp pose and reaching it with cuRobo [52]. To compute the grasp pose:

1. We estimate the major axis of the object’s point cloud in top-down view using PCA.
2. The grasp orientation is orthogonal to the major axis.
3. The grasp height is based on the highest point in the object’s point cloud minus an offset of 3cm. This heuristic ensures the system can grip long objects.

The placing behavior is the same as picking, except that an area within 5cm of the placement coordinate is used as the point cloud for estimating orientation and height, and a vertical height offset is added to account for the height at which the object was picked.

D.5.2. Additional Results

We present additional results from the robot experiments in Fig. 6. We observe that models trained with ROBOSPATIAL consistently outperform baseline models in most cases, even though the prompt is not optimized for ROBOSPATIAL-trained models. This demonstrates that the power of VLMs enables templated language to generalize to language unseen during training while maintaining spatial understanding capabilities. However, even with ROBOSPATIAL training, the models struggle with understanding stacked items, indicating a need for further data augmentation with diverse layouts. In a few cases, ROBOSPATIAL training adversely affects performance, especially with RoboPoint [62]. We hypothesize that mixing the dataset with RoboPoint training data and ROBOSPATIAL training data may lead to unforeseen side effects, particularly in grounding objects. Nevertheless, we demonstrate that ROBOSPATIAL training enhances VLM’s spatial understanding in real-life robotics experiments, even with freeform language.

1423 D.6. More Qualitative Examples

1424 Fig. 7 present additional qualitative comparisons between
1425 models trained on ROBOSPATIAL. The findings demon-
1426 strate that models trained on ROBOSPATIAL consistently
1427 exhibit spatial understanding in the challenging ROBOSPA-
1428 TIAL-Home dataset, even outperforming closed models like
1429 GPT-4o [42]. However, we observed that object ground-
1430 ing is a crucial prerequisite for spatial understanding; the
1431 improvement is often hindered by the model’s inability to
1432 ground objects in cluttered scenes, where GPT-4o performs
1433 more effectively. Additionally, we show that the ROBOSPA-
1434 TIAL-trained model successfully generalizes to unseen spa-
1435 tial relationships in BLINK-Spatial [15], including those in-
1436 volving distance, such as “touching.”















| | | | | | | | | | | | | | | | | | | | | | | | | | | | |
|---|--|-----------------|---|--------------------|---|----------------|---|-------------------|---|-----------|---|-------------|---|--|---|-----------------|---|--------------------|---|----------------|---|-------------------|---|-----------|---|-------------|---|
|  | <p>Question: pick lone object</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✗</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✓</td></tr><tr><td>GPT-4o [42]</td><td>✗</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✗ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✓ | GPT-4o [42] | ✗ |  | <p>Question: Is there room to slot the pancake mix in the middle of the row of boxes</p> <table><tr><td>LLaVa-Next [35]</td><td>✓</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✗</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✓</td></tr><tr><td>GPT-4o [42]</td><td>✓</td></tr></table> | LLaVa-Next [35] | ✓ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✗ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✓ | GPT-4o [42] | ✓ |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
|  | <p>Question: Is there space in the white container for the orange juice box</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✗</td></tr><tr><td>RoboPoint-FT [62]</td><td>✗</td></tr><tr><td>Molmo [9]</td><td>✗</td></tr><tr><td>GPT-4o [42]</td><td>✓</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✗ | RoboPoint-FT [62] | ✗ | Molmo [9] | ✗ | GPT-4o [42] | ✓ |  | <p>Question: alphabet soup fit in the purple box</p> <table><tr><td>LLaVa-Next [35]</td><td>✓</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✗</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✗</td></tr><tr><td>GPT-4o [42]</td><td>✓</td></tr></table> | LLaVa-Next [35] | ✓ | LLaVa-Next-FT [35] | ✗ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✗ | GPT-4o [42] | ✓ |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
|  | <p>Question: pick object behind the middle container</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✗</td></tr><tr><td>Molmo [9]</td><td>✗</td></tr><tr><td>GPT-4o [42]</td><td>✗</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✗ | Molmo [9] | ✗ | GPT-4o [42] | ✗ |  | <p>Question: pick shortest object</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✓</td></tr><tr><td>GPT-4o [42]</td><td>✓</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✓ | GPT-4o [42] | ✓ |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
|  | <p>Question: place object in container behind pop-corn</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✗</td></tr><tr><td>GPT-4o [42]</td><td>✗</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✗ | GPT-4o [42] | ✗ |  | <p>Question: place the object inside the smallest box</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✓</td></tr><tr><td>GPT-4o [42]</td><td>✗</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✓ | GPT-4o [42] | ✗ |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
|  | <p>Question: can the robot directly pick the red orange peaches can without disturbing other objects?</p> <table><tr><td>LLaVa-Next [35]</td><td>✓</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✗</td></tr><tr><td>RoboPoint-FT [62]</td><td>✗</td></tr><tr><td>Molmo [9]</td><td>✓</td></tr><tr><td>GPT-4o [42]</td><td>✓</td></tr></table> | LLaVa-Next [35] | ✓ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✗ | RoboPoint-FT [62] | ✗ | Molmo [9] | ✓ | GPT-4o [42] | ✓ |  | <p>Question: is there an object that is not in a stack?</p> <table><tr><td>LLaVa-Next [35]</td><td>✓</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✓</td></tr><tr><td>GPT-4o [42]</td><td>✓</td></tr></table> | LLaVa-Next [35] | ✓ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✓ | GPT-4o [42] | ✓ |
| LLaVa-Next [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
|  | <p>Question: can the macaroni and cheese be placed on top of cheez-it without touching other objects?</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✗</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✗</td></tr><tr><td>GPT-4o [42]</td><td>✓</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✗ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✗ | GPT-4o [42] | ✓ |  | <p>Question: is there space to place one of the cans on the cheez-it box?</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✗</td></tr><tr><td>RoboPoint [62]</td><td>✗</td></tr><tr><td>RoboPoint-FT [62]</td><td>✗</td></tr><tr><td>Molmo [9]</td><td>✗</td></tr><tr><td>GPT-4o [42]</td><td>✗</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✗ | RoboPoint [62] | ✗ | RoboPoint-FT [62] | ✗ | Molmo [9] | ✗ | GPT-4o [42] | ✗ |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
|  | <p>Question: place on the object to the left of macaroni and cheese</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✓</td></tr><tr><td>RoboPoint [62]</td><td>✓</td></tr><tr><td>RoboPoint-FT [62]</td><td>✓</td></tr><tr><td>Molmo [9]</td><td>✓</td></tr><tr><td>GPT-4o [42]</td><td>✗</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✓ | RoboPoint [62] | ✓ | RoboPoint-FT [62] | ✓ | Molmo [9] | ✓ | GPT-4o [42] | ✗ |  | <p>Question: pick the highest object on the stack of two objects</p> <table><tr><td>LLaVa-Next [35]</td><td>✗</td></tr><tr><td>LLaVa-Next-FT [35]</td><td>✗</td></tr><tr><td>RoboPoint [62]</td><td>✗</td></tr><tr><td>RoboPoint-FT [62]</td><td>✗</td></tr><tr><td>Molmo [9]</td><td>✗</td></tr><tr><td>GPT-4o [42]</td><td>✗</td></tr></table> | LLaVa-Next [35] | ✗ | LLaVa-Next-FT [35] | ✗ | RoboPoint [62] | ✗ | RoboPoint-FT [62] | ✗ | Molmo [9] | ✗ | GPT-4o [42] | ✗ |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✓ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| LLaVa-Next-FT [35] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| RoboPoint-FT [62] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| Molmo [9] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |
| GPT-4o [42] | ✗ | | | | | | | | | | | | | | | | | | | | | | | | | | |

Figure 6. Additional robot experiments. A green check mark indicates that the model answered correctly. The -FT suffix denotes a model trained with ROBOSPATIAL. The questions are purposely not cleaned to reflect realistic language inputs.


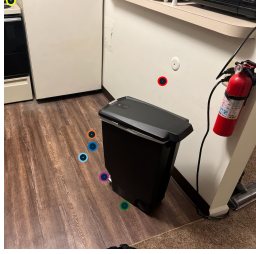



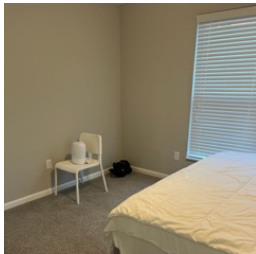
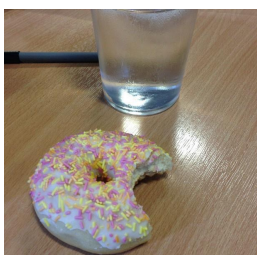
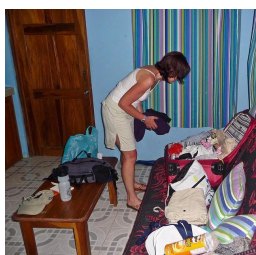
| | | | |
|---|---|--|---|
|  | <p>Question: Pinpoint several points within the vacant space situated to the left of the pot.</p> <p>Answer</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ●</p> <p>RoboPoint [62] ●</p> <p>RoboPoint-FT [62] ●</p> <p>Molmo [9] ●</p> <p>GPT-4o [42] ●</p> |  | <p>Question: Pinpoint several points within the vacant space situated behind the trash bin.</p> <p>Answer</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ●</p> <p>RoboPoint [62] ●</p> <p>RoboPoint-FT [62] ●</p> <p>Molmo [9] ●</p> <p>GPT-4o [42] ●</p> |
|  | <p>Question: Can the lamp fit in front of the shelf?</p> <p>Answer Yes</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ✓</p> <p>RoboPoint [62] ✗</p> <p>RoboPoint-FT [62] ✓</p> <p>Molmo [9] ✗</p> <p>GPT-4o [42] ✗</p> |  | <p>Question: Can the pot fit above the fridge?</p> <p>Answer Yes</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ✓</p> <p>RoboPoint [62] ✗</p> <p>RoboPoint-FT [62] ✓</p> <p>Molmo [9] ✗</p> <p>GPT-4o [42] ✗</p> |
|  | <p>Question: Is the lamp above the shelf?</p> <p>Answer Yes</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ✓</p> <p>RoboPoint [62] ✗</p> <p>RoboPoint-FT [62] ✓</p> <p>Molmo [9] ✗</p> <p>GPT-4o [42] ✓</p> |  | <p>Question: Is the chair behind the bed?</p> <p>Answer Yes</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ✓</p> <p>RoboPoint [62] ✗</p> <p>RoboPoint-FT [62] ✓</p> <p>Molmo [9] ✗</p> <p>GPT-4o [42] ✗</p> |
|  | <p>Question: Is the dining table touching the donut?</p> <p>Answer Yes</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ✓</p> <p>RoboPoint [62] ✗</p> <p>RoboPoint-FT [62] ✓</p> <p>Molmo [9] ✗</p> <p>GPT-4o [42] ✗</p> |  | <p>Question: Is the couch under the suitcase?</p> <p>Answer Yes</p> <p>LLaVa-Next [35] ✗</p> <p>LLaVa-Next-FT [35] ✓</p> <p>RoboPoint [62] ✗</p> <p>RoboPoint-FT [62] ✓</p> <p>Molmo [9] ✗</p> <p>GPT-4o [42] ✗</p> |

Figure 7. Qualitative results on spatial reasoning benchmarks. The -FT suffix denotes a model trained with ROBOSPATIAL. The first three rows show examples from ROBOSPATIAL-Home, covering spatial context, spatial compatibility, and spatial configuration. For spatial context questions, only the first predicted point from each model is shown. The fourth row shows generalization to unseen spatial relationships on the Blink-Spatial [15] dataset, demonstrating that the ROBOSPATIAL-trained model can transfer to unseen relationships.