

# SoFAR: Language-Grounded Orientation Bridges Spatial Reasoning and Object Manipulation

Zekun Qi<sup>13¶</sup> Wenyao Zhang<sup>237¶</sup> Yufei Ding<sup>34¶</sup> Runpei Dong<sup>5</sup> Xinqiang Yu<sup>3</sup>  
 Jingwen Li<sup>4</sup> Lingyun Xu<sup>4</sup> Baoyu Li<sup>5</sup> Xialin He<sup>5</sup> Guofan Fan<sup>1</sup> Jiazhao Zhang<sup>3</sup> Jiawei He<sup>3</sup>  
 Jiayuan Gu<sup>6</sup> Xin Jin<sup>7</sup> Kaisheng Ma<sup>1</sup> Zhizheng Zhang<sup>3†</sup> He Wang<sup>34†</sup> Li Yi<sup>189†</sup>

<sup>¶</sup>Equal Contribution <sup>†</sup>Corresponding authors

<sup>1</sup>Tsinghua University <sup>2</sup>Shanghai Jiao Tong University <sup>3</sup>Galbot <sup>4</sup>Peking University <sup>5</sup>UIUC  
<sup>6</sup>ShanghaiTech University <sup>7</sup>Eastern Institute of Technology <sup>8</sup>Shanghai Qi Zhi Institute <sup>9</sup>Shanghai AI Laboratory  
<https://qizekun.github.io/sofar>

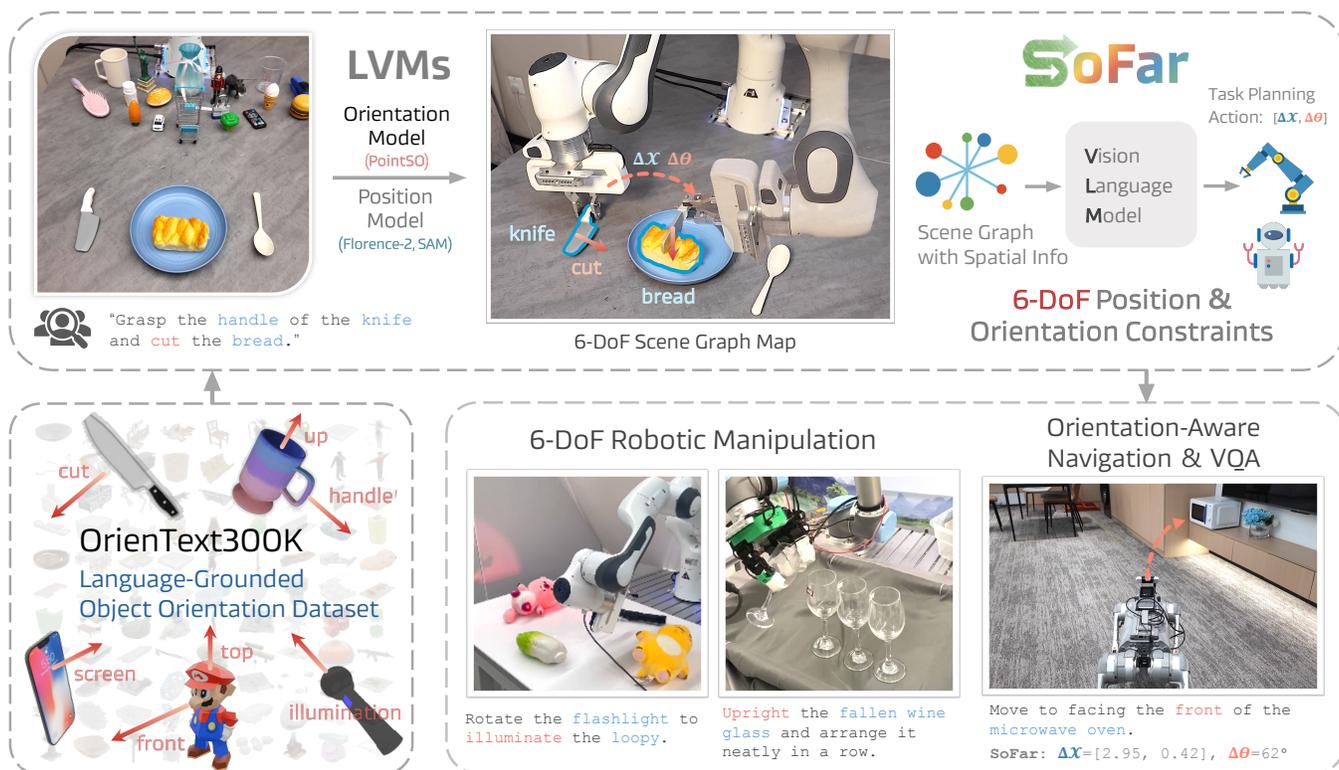


Fig. 1: **Overview.** We present SoFAR, the first system to incorporate object orientation into spatial reasoning and robotic manipulation. We introduce the concept of *Semantic Orientation*, which refers to natural language-grounded object orientations, such as the “cutting” direction of a knife or the “handle” direction of a cup. To support this, we construct OrientText300K, a large-scale dataset of object-text-orientation pairs. SoFAR enables accurate 6-DoF object understanding and excels in tasks such as visual question answering, object manipulation, and navigation.

**Abstract**—Spatial intelligence is a critical component of embodied AI, promoting robots to understand and interact with their environments. While recent advances have enhanced the ability of VLMs to perceive object locations and positional relationships, they still lack the capability to precisely understand object orientations—a key requirement for tasks involving fine-grained manipulations. Addressing this limitation not only requires geometric reasoning but also an expressive and intuitive way to represent orientation. In this context, we propose that natural language offers a more flexible representation space than canonical frames, making it particularly suitable for instruction-following robotic systems. In this paper, we introduce the concept of

semantic orientation, which defines object orientations using natural language in a reference-frame-free manner (e.g., the “plug-in” direction of a USB or the “handle” direction of a knife). To support this, we construct OrientText300K, a large-scale dataset of 3D models annotated with semantic orientations that link geometric understanding to functional semantics. By integrating semantic orientation into a VLM system, we enable robots to generate manipulation actions with both positional and orientational constraints. Extensive experiments in simulation and real world demonstrate that our approach significantly enhances robotic manipulation capabilities, e.g., 48.7% accuracy on Open6DOR and 74.9% accuracy on SIMPLER.

## I. INTRODUCTION

Open-world spatial intelligence is crucial for embodied AI, as a robot must understand not only “what” an object is but also its precise “where” for effective interaction. To this end, vision-language-models (VLMs) [1, 30, 68, 79] with spatial understanding [10, 12, 15] that comprehend spatial concepts [70, 142] and relationships [41, 66, 135] have been built. These models incorporate spatial knowledge into their architecture design or training data, enabling them to perform tasks such as distinguishing left from right [12, 15], counting objects [10, 33], and even planning for position-only manipulations [10, 160]. Despite the remarkable achievements, we ask: What is the missing cornerstone of such spatial understanding? Given the original intent of “seeing is for doing” [40], *how can we push spatial understanding further?*

We observe that current VLMs struggle with understanding object **orientation**, making them insufficient for generic robot manipulation planning. Consider some everyday scenarios: inserting a pen into a pen holder, righting a tilted wine glass, or plugging a cord into a power strip. Previous approaches [10, 12, 15] primarily focused on understanding “where is the pen” or “where is the wine glass” while ignoring their orientations, making them insufficient for accomplishing these seemingly simple object manipulation tasks.

More importantly, different orientations of an object hold varying semantic significance. The capability of connecting specific orientations to their semantic meanings is essential for language-guided robot manipulations. For example, inserting a pen into a pen holder requires aligning the pen tip with the direction of the pen holder’s opening; righting a wine glass necessitates aligning the glass’s top with the z-axis in the world coordinate frame; and plugging into a power strip involves understanding the “insertion” direction, which is perpendicular to the power strip’s surface. However, translating a specific language description into a desired object orientation is challenging for existing VLMs.

To move forward, we introduce **language-grounded orientation that bridges spatial reasoning and object manipulation**, characterized by the following:

- **From Position Awareness to Orientation Awareness.** While prior works [10, 12, 15] emphasize position relationship, orientation understanding is equally critical for defining the full six degrees of freedom (6-DoF) of object or end-effector poses [18, 34, 71, 141, 146, 158]. Orientation awareness involves understanding object orientations and their relationships in the open world, enabling robots to complete tasks requiring precise alignment and rearrangement together with position awareness.
- **From Orientation to Semantic Orientation.** Traditional orientation, defined relative to a base frame or template model [18, 67, 122, 141, 146], is insufficient for open-world manipulation guided by language instructions [57, 119, 127]. We introduce semantic orientation, linking orientational vectors of an object to open-vocabulary prompts (*e.g.*, the “handle” direction of a knife or “plug-in”

direction of a USB). This bridges geometric reasoning with functional semantics, enabling robots to interpret task-specific orientation changes.

Achieving such spatial awareness requires addressing two key challenges: acquiring semantic orientation knowledge in the open world and integrating it with VLMs. To tackle the first, we propose PointSO, a generalizable cross-modal 3D Transformer [29, 107, 109, 134], which serves as a robust and versatile framework for open-world spatial orientation understanding. To train PointSO effectively, we construct OrientText300K, a large-scale orientation-text paired dataset curated from internet sources. This dataset, devoid of expensive robot data, is automatically labeled by prompting GPT-4o [100] with extensive and diverse language-grounded semantic orientation queries. These queries encompass *intra-object* spatial understanding and *inter-object* interaction-related semantics, such as manipulation orientations. OrientText300K comprises over 350K 3D models of diverse everyday objects. Powered by OrientText300K, PointSO can reliably infer semantic orientation for an arbitrary object without being restricted to a known category or instance.

We further develop an integrated reasoning system, SOFAR, to coordinate our proposed PointSO and existing position foundation models for achieving more comprehensive spatial understanding, where Florence-2 [147] and SAM [65] handle the object positions, while our PointSO focuses on understanding and outputting orientations complimentary. Specifically, we parse an input RGB-D observation as an orientation-aware 3D scene graph using SAM-segmented object point clouds and our PointSO. The RGB-D observation, together with the scene graph, is then input to the VLM, which outputs a chain-of-thought [140] spatial reasoning for both position and orientation commands. These commands can then serve as visual planning outcomes to support robotic manipulation tasks.

To assess our system, we introduce Open6DOR V2, a large-scale robot manipulation benchmark designed for 6-DoF object rearrangement in simulation. This benchmark demands robust positional and orientational reasoning in open-world settings and supports both open-loop and closed-loop robotic control. Our experiments demonstrate that our system considerably outperforms state-of-the-art vision-language models and popular vision-language-action (VLA) models—even those trained with extensive and costly robot trajectories. These performance gains are also observed in real-world experiments. Additionally, we establish a new spatial visual-question-answering benchmark, that confirms the system’s exceptional open-world spatial reasoning capabilities.

In summary, our contributions are fourfold. First, we introduce PointSO, an orientation base model that infers the semantic directions of novel objects in an open-world context. Second, we curate OrientText300K, a large-scale 3D model dataset annotated with semantic directions to support the training of orientation models. Third, we develop an integrated system that enhances powerful VLMs with advanced spatial understanding, facilitating robot manipulations that require both positional and orientational spatial knowledge.



Fig. 2: Data Construction of OrientText300K.

Fourth, we establish 6-DoF SpatialBench & Open6DOR V2, an orientation-aware spatial visual-question-answering benchmark and a comprehensive robot manipulation benchmark designed to evaluate both open-loop and closed-loop open-world 6D rearrangement strategies. Extensive experiments demonstrate the superior performance of our method together with a series of perception and robot manipulation benchmarks.

## II. SEMANTIC ORIENTATION: CONNECTING LANGUAGE AND OBJECT ORIENTATION

### A. Definition of Semantic Orientation

Traditionally, the orientation of an object is defined within a reference frame, using quaternions or Euler angles to represent relative rotations. Intuitively, object orientations commonly correspond to some specific semantics in most interactive behaviors. This aligns with the fact that humans typically understand an object’s orientation in a more semantic, reference-free way. For instance, when plugging a plug into a charger, we accomplish the action of “plugging in” by matching the metal prongs’ direction with the outward direction of the charger’s socket. Drawing on this observation, we define an object’s *Semantic Orientation* as follows. Given an object  $X$  and a description  $\ell$ , the corresponding semantic orientation  $\mathbf{s}_\ell^X \in S(2)$  is an object-centric direction represented as a unit vector semantically matching the description  $\ell$ .

$$\mathbf{s}_\ell^X = \mathcal{F}(X, \ell). \quad (1)$$

$\ell$  is an open-vocabulary language description that should have a clear semantic correspondence to a general orientation (e.g., front, top), an object part (e.g., handle, cap), or a specific manipulation goal (e.g., pour out, plug-in).

For an object  $X$ , it may have multiple semantic orientations corresponding to different functions or attributes via changing the language description  $\ell$ , forming a semantic orientation set  $S_X = \{\mathbf{s}_{\ell_1}^X, \mathbf{s}_{\ell_2}^X, \dots, \mathbf{s}_{\ell_n}^X\}$ . Based on this set, the rotation of  $X$  can be characterized by transforming its semantic orientations.

### B. Robotic Manipulation via Semantic Orientation

Semantic orientations are powerful representations that help characterize various orientation-related knowledge. What relates

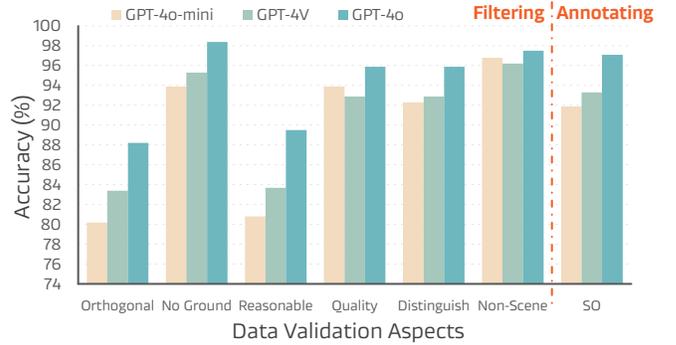


Fig. 3: OrientText300K data filtering and annotating accuracy on human-labeled validation sets. SO denotes the annotation quality for semantic orientation. All VLMs achieve high accuracy and GPT-4o consistently yields the best result.

most to robot manipulation is the knowledge of object reorientation. Given an initial observation of an object  $X$  and a task command  $c$  that specifies the desired reorientation, semantic orientations can be used to determine the necessary object rotation. First, we identify task-related semantic orientation descriptions  $\{\ell^c\}$  from the task command  $c$  whose desired directions are clearly outlined in the command. For instance, a command such as “turn a bottle over and put it on the ground” necessitates identifying the “up direction of the bottle” as a semantic orientation, with the desired direction being  $(0, 0, -1)$  in a world coordinate system where the z-axis is perpendicular to the ground. Then, by extracting the semantic orientations from the initial observation  $X$  and calculating the rotation needed to align these with the desired directions, we can effectively determine how the object should be reoriented.

Beyond object reorientation, semantic orientations can be linked to traditional instance-level and category-level orientations, and even facilitate cross-category orientation. Specifically, by orthogonalizing a set of semantic orientations, we can establish reference frames that define instance-level orientations for individual objects. For objects within the same category, using a consistent set of linguistic descriptions aligns their semantic orientation sets, resulting in category-level consistent references from which category-level orientations can be derived. Moreover, applying the same linguistic descriptors across different categories creates cross-category consistent semantic orientation sets, enabling the derivation of cross-category reference frames. Thus, our approach builds orientation understanding on semantic orientations, with an emphasis on learning to estimate these directions for open-world objects based on open-ended descriptions.

### C. OrientText300K: Orientation-Text Paired Data at Scale

Our goal is to learn an *orientation model* on a large-scale 3D model dataset so that it can identify semantic orientations in open-world scenarios. To achieve it, we first introduce OrientText300K, a newly curated 3D model dataset with diverse and extensive semantic orientation labels for training our language-conditional orientation model.

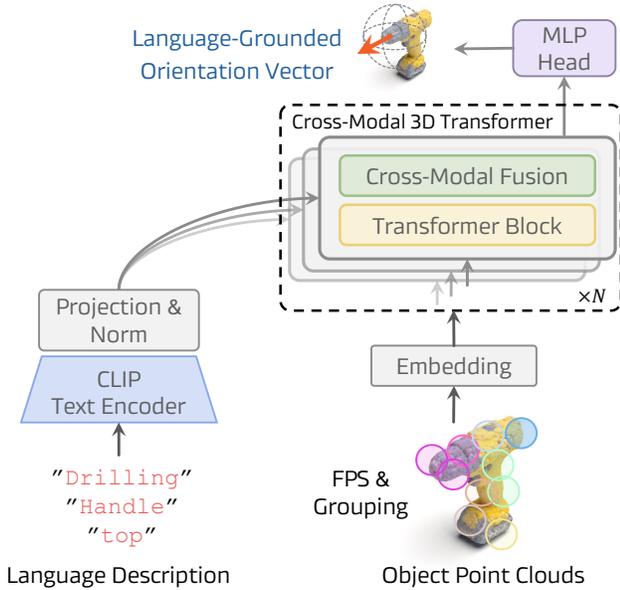


Fig. 4: **PointSO architecture** for SO prediction. PointSO takes 3D and language embeddings as inputs of a  $B$ -layer plain Transformer model, and uses a simple token-level sum as cross-modal fusion in all layers. An MLP head predicts the corresponding orientation vector normalized in a unit sphere.

1) *Scalability Analysis*: Before discussing the data curation process, it is important to note that semantic orientations are typically defined within the six *standard orthogonal views* of an object from a canonical setup. For example, the orientation for placing a pen into a pencil holder is opposite to the holder’s upright direction. Object canonicalization often involves specifying the up and front directions of an object. Fortunately, most 3D models available in web datasets are already canonicalized, except for potential axis flipping. Based on this observation, we propose to scale up the semantic orientation annotations by leveraging readily available web 3D datasets with automatic labeling using GPT. Specifically, we can use the six orthogonal directions as candidates and generate descriptions from rendered multi-view images. This associates various language descriptions with salient directions on objects, providing the supervision needed for learning an orientation model.

2) *Data Source*: To scale up, we build the OrienText300K dataset from Objaverse [22], which originally contains  $\sim 800K$  Internet 3D models across substantial categories. However, such Internet data contains a lot of noisy annotations or low-quality samples that cannot be used. Based on Blender, we carefully set up the light and rendered more than 8M high-fidelity rendering images.

3) *Data Filtering*: To clean the data that can better be used for generating semantic orientation annotations, we first clean the data by using a dedicated filtering strategy that filters data to preserve samples that satisfy the following 6 requirements. ❶ Standard orthogonal view only. Samples in random views will be filtered. ❷ Clean objects without the ground for auxiliary

visualization. ❸ Reasonable objects that have sufficient spatial reasoning potentials. ❹ High-quality objects. Low-quality objects such as blurry and wrong samples are filtered. ❺ Distinguishable objects. Abstract objects such as meaningless solids are filtered. ❻ Non-scene objects. Samples that describe a 3D scene are filtered for object-centric understanding purposes.

However, it is non-trivial to conduct filtering on such big data using manual labor. Inspired by recent works showing large VLMs are human-aligned 2D or 3D image-based judges [103, 143, 170], we employ GPT-4o [100] by prompting requirements above. To be specific, the multi-view images of 3D objects are concatenated together with our designed prompts into GPT-4o, and GPT-4o will decide whether one sample should be filtered. The filtered dataset yields 350K+ clean samples, significantly reducing data noise.

4) *Data Annotation*: Like data filtering, we annotate OrienText300K with semantic orientations using GPT-4o. GPT-4o will take rendered object views as input and generate the language description and one best-grounded direction from the 6 standard orthogonal direction candidates. We render the six *standard* orthogonal view images in Blender, which are concatenated as input of GPT-4o. Besides, four 45-degree *oblique* orthogonal views are rendered and concatenated as an additional input, ensuring more robust annotation.

5) *Quality Validation*: To validate the accuracy of our prompted GPT agents as the data filter and annotator, we construct a validation set containing 208 samples with manually labeled filtering criteria and semantic orientation labels, respectively. From Fig. 3, we observe that GPT-4o achieves an average accuracy of 88.3% and 97.1% accuracy on filtering and annotating, respectively. This provides a quality guarantee of our OrienText300K.

#### D. PointSO: A Cross-Modal 3D Transformer for Semantic Orientation Prediction

We introduce PointSO, a plain Transformer-based architecture [134] with cross-modal 3D-language fusion as our orientation model. As illustrated in Fig. 4, PointSO takes the object’s 3D point clouds and a language description as inputs, and predicts the corresponding semantic orientation.

1) *3D and Language Embeddings*: Given an object’s point cloud  $X = \{\mathbf{x}_i \in \mathbb{R}^3 | i = 1, 2, \dots, N\}$  with  $N$  3D points defined in (x, y, z) Cartesian space, and an arbitrary language description  $\ell$ , we first embed both into discrete token embeddings. For the 3D point clouds, we follow [29, 107, 159] to first sample  $N_s$  seed points using farthest point sampling (FPS) and then group inputs with KNN for point feature embedding with a local geometric extraction network such as lightweight PointNet [104, 105]. An MLP head is used which maps a special [CLS] token [31] to a predicted direction. As for the language inputs, we adopt OpenAI’s CLIP [115] and use the global token as cross-modal fusion inputs. Like Vision Transformer [31], we explore three model configurations, *i.e.*, small, base, and large versions. Each configuration is of a different number of Transformer layers, dimensions, and CLIP models, detailed in the Appendix B.

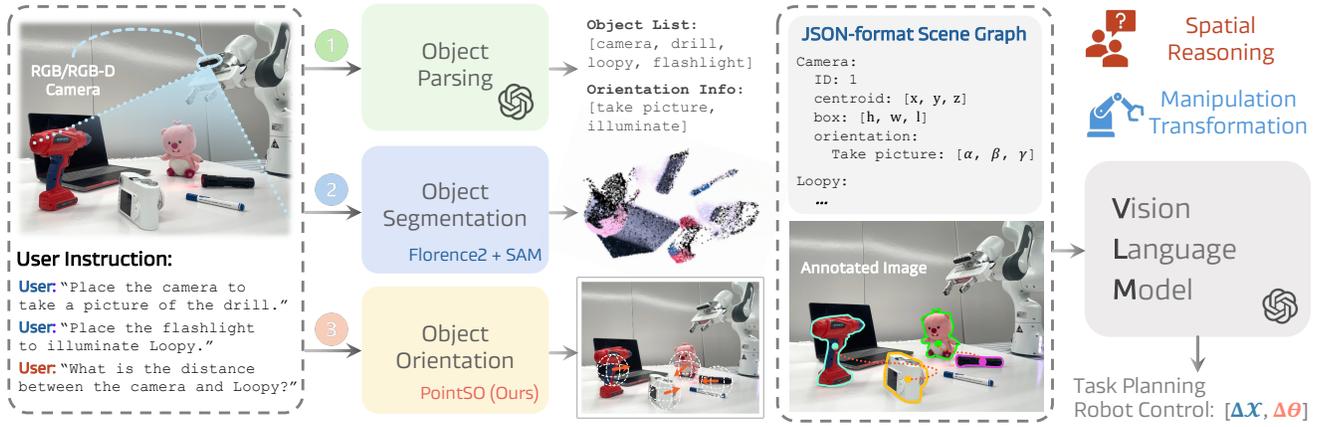


Fig. 5: **Overview of SOFAR system.** SOFAR takes RGB-D images as inputs, where the depth images can be obtained from depth sensor or metric 3D prediction [157]. Given the language instruction, SOFAR prompts the VLM to obtain task-oriented object phases and semantic orientation descriptions. Then, SOFAR leverage foundation models Florence-2 [147] and SAM [65] to segment depth point clouds and our PointSO (Section II-D) to obtain semantic orientations. Summarizing 3D object-centric information, an orientation-aware scene graph is constructed and encoded into languages (Section III). The VLM takes the RGB image and the scene graph and outputs the queried spatial understanding VQA or translation for manipulation.

2) *Cross-Modal Fusion:* We adopt a dense layer-wise feature fusion strategy by injecting the global text features into the Transformer layers of the 3D Transformer, where a cross-modal fusion is conducted. Generally, this fusion operation can be implemented in various fashions such as cross-attention, adapter, or concatenating features along spatial/channel dimensions. Empirically, however, we find that the sum of the text token to every point token is the most simple but effective solution (see Appendix H2). This may be due to the relatively short length of the languages, where a per-token sum enhances attention to languages.

3) *Optimization:* Let  $\mathcal{F}_{SO}$  represent the PointSO model parameterized by  $\theta_{SO}$  (the CLIP is kept frozen and thus its parameters are not included). Given every object point cloud  $X_i \in \mathcal{D}_{\text{OrientText300K}}$  in the OrientText300K dataset, where each object is labeled with a language set  $L_i = \{\ell_j^i, j = 1, 2, \dots, Q\}$  and the corresponding ground truth semantic orientation set,  $S_i = \{s_j^i, j = 1, 2, \dots, Q\}$ . The optimization is to minimize the negative cosine similarity  $\mathcal{L}_{\cos}(\mathbf{v}, \mathbf{k}) = 1 - \frac{\mathbf{v} \cdot \mathbf{k}}{\|\mathbf{v}\| \cdot \|\mathbf{k}\|}$  between predicted and the ground truth semantic orientations:

$$\min_{\theta_{SO}} \sum_{X_i \in \mathcal{D}_{\text{OrientText300K}}} \sum_{\ell_j^i \in L_i} \mathcal{L}_{\cos}(\mathcal{F}_{SO}(X_i, \ell_j^i), s_j^i). \quad (2)$$

### III. SOFAR: SEMANTIC ORIENTATION BRIDGES SPATIAL REASONING AND OBJECT MANIPULATION

Our proposed PointSO model now paves the off-the-shelf for object-centric spatial orientation understanding. However, it is still unclear how to leverage such object-centric spatial understanding for scene-level spatial reasoning both in the digital world (e.g., orientation-aware visual question answering, VQA) and in the physical world (e.g., robot manipulations). To enable such applications, we build an integrated reasoning

system where a powerful VLM acts as an agent and reasons about the scene while communicating with off-the-shelf models including PointSO and SAM [65]. Fig. 5 illustrates an overview of our proposed framework, aiming at **Semantic Orientation For Autonomous Robotic manipulation (SOFAR)**. SOFAR consumes an RGB-D image and a language query as input and first leverages off-the-shelf models including SAM and PointSO to convert the image into an orientation-aware 3D scene graph. Then SOFAR leverages a VLM agent to produce planning outcomes based upon the scene graph and the input language query, which can be later used for robot manipulation. We will introduce the construction of the orientation-aware 3D scene graph in Section III-A and how to perform spatial-aware task reasoning and plan for robot manipulation in Section III-B.

#### A. Orientation-Aware Scene Graph from RGB-D

To convert the input RGB-D image into an orientation-aware 3D scene graph, we first segment the RGB image to obtain object-level 3D point clouds using SAM [65] and then construct a scene graph with object-attribute nodes.

1) *Task-Oriented Object Segmentation:* Given a language query  $\mathcal{Q}$ , we first prompt a VLM model  $\mathcal{F}_{VLM}$  to abstract the task-oriented *object phrase set*. Thus, a set  $\mathcal{P} = \{p_i | i = 1, 2, \dots, M\}$  with  $M$  object phrases in language will be generated from  $\mathcal{Q}$ . With set  $\mathcal{P}$ , we use language-conditioned object segmentation with SAM to obtain an object set  $\mathcal{X} = \{X_i | i = 1, 2, \dots, M\}$ , where  $X_i$  is the 3D point cloud of the  $i$ -th object. Besides, we assign individual IDs to objects which are used for Set-of-Mark (SoM) prompting [151] on VLM's image input. Next, we prompt the VLM to generate every object's corresponding task-oriented language description set  $L_i$ . We predict the semantic orientation using pretrained PointSO for each description in the description set  $L_i$ , forming the semantic orientation set  $S_i$  for the  $i$ -th object.



Place the water bottle on the right into the box and arrange it in a 3x3 pattern.



Grasp the handle of the knife and cut the bread.



Aim the camera at the toy truck.



Put the chili into the basket.



Pick the highest box and place it on the right.



Rotate the flashlight to illuminate the loopy.



Pull out a tissue.



Take out the test tube with the green solution.



Open the microwave.



Open the cabinet.



Pick the box and place it to the right of the doll.



Pick up the cabbage and place it in the basket.



Pick up the Lego block and place it between the two toy trucks.



Rotate Loopy to face the yellow dragon doll.



Right the bottle.



Right the fallen wine glass and arrange it neatly in a row with the other wine glasses.

Fig. 6: **Qualitative results** of real world language-grounded manipulation. SOFAR doesn't require any robotic data training or human annotation of task-specific prompts, and it can generalize across various **embodiments**, tasks and environments.

TABLE I: **6-DoF object rearrangement evaluation** on Our Proposed Open6DOR V2 Benchmark.

Method	Position Track			Rotation Track				6-DoF Track			Time Cost (s)
	Level 0	Level 1	Overall	Level 0	Level 1	Level 2	Overall	Position	Rotation	Overall	
<i>Perception Tasks on Issac Sim [95]</i>											
GPT-4V [98]	46.8	39.1	45.2	9.1	6.9	11.7	9.2	-	-	-	-
Dream2Real [62]	17.2	11.0	15.9	37.3	27.6	26.2	31.3	26.2	18.7	13.5	358.3s
VoxPoser [54]	35.6	21.7	32.6	-	-	-	-	-	-	-	-
Open6DOR-GPT [28]	78.6	60.3	74.9	45.7	32.5	49.8	41.1	84.8	40.0	35.6	126.3 s
<b>SO FAR-LLaVA</b>	86.3	57.9	78.7	62.5	30.2	67.1	48.6	83.0	48.2	40.3	9.6s
<b>SO FAR</b>	<b>96.0</b>	<b>81.5</b>	<b>93.0</b>	<b>68.6</b>	<b>42.2</b>	<b>70.1</b>	<b>57.0</b>	<b>92.7</b>	<b>52.7</b>	<b>48.7</b>	<b>8.5s</b>
<i>Execution Tasks on Libero [76]</i>											
Octo [126]	51.2	32.1	47.2	10.7	18.3	29.9	17.2	45.6	8.0	8.0	-
OpenVLA [64]	51.6	32.4	47.6	11.0	18.5	30.6	17.6	46.2	8.2	8.2	-
<b>SO FAR</b>	<b>72.1</b>	<b>47.6</b>	<b>67.0</b>	<b>28.3</b>	<b>18.3</b>	<b>34.7</b>	<b>25.7</b>	<b>63.7</b>	<b>25.6</b>	<b>18.4</b>	<b>40s</b>

2) *Orientation-Aware 3D Scene Graph*: From segmented object set  $\mathcal{X}$ , we construct an orientation-aware scene graph  $\mathcal{G} = (\mathbf{V}, \mathbf{E})$  with  $M$  object nodes. Each node  $\mathbf{o}_i \in \mathbf{V}$  consists of following *semantic* and *spatial* attributes: ❶ object name  $p_i$  with an individual ID  $i$  that distinguishes instances; ❷ object position as 3D centroid coordinate  $\mathbf{c}_i = (x, y, z) \in \mathbb{R}^3$  from segmented depth; ❸ object’s 3D bounding box size  $\mathbf{b}_i = (h, w, l) \in \mathbb{R}^3$ ; ❹ semantic orientation set  $S_i$  and the corresponding task-oriented language-description set  $L_i$ . Each edge  $\mathbf{e}_{ij} \in \mathbf{E}$  stores the relative translation and bounding box size ratio between the connected two objects  $\mathbf{o}_i$  and  $\mathbf{o}_j$ .

### B. Spatial-Aware Task Reasoning

We convert the orientation-aware scene graph  $\mathcal{G}$  into descriptive language and input this, along with the RGB image  $I$  and query  $\mathcal{Q}$ , into a VLM. With spatial knowledge accurately encoded in the scene graph, the VLM can leverage its robust image and language understanding capabilities to produce high-quality spatial reasoning results.

1) *Chain-of-Thought Spatial Reasoning*: Since most of the robot manipulation problems involving rigid bodies can be abstract as changing the position and orientation of objects, we especially focus on instructing the VLM to output a goal transformation as the plan for robot manipulation. To realize this, we use chain-of-thought (CoT) reasoning [140] that regularizes the VLM to derive a transformation given a language-described manipulation goal in 3 steps: i) scene analysis of the language query  $\mathcal{Q}$  and object nodes  $\mathbf{V}$ ; ii) manipulation analysis to provide an analytical calculation process of the target object’s desired position and orientation; iii) output task-desired object position  $\tilde{\mathbf{c}}_i$  and semantic orientation set  $\tilde{S}_i$  for each object. Afterward, given each object’s initial position  $\mathbf{c}_i$  and semantic orientation set  $S_i$ , the 6-DoF transformation matrix  $\mathbf{P}_i$  can be derived. Specifically, the translation transformation  $\mathbf{t}_i = \tilde{\mathbf{c}}_i - \mathbf{c}_i$ , and we solve the rotation transformation  $\mathbf{R}_i$  from  $S_i$  and  $\tilde{S}_i$  using Kabsch-Umeyama algorithm [60, 61, 132].

2) *Low-Level Motion Execution*: Similar to CoPa [51], our model includes task-oriented grasping and task-aware motion planning. We first segment the manipulated objects or parts

using Florence-2 [147] and SAM [65] to obtain the object point cloud, and then we use GSNet [38] to generate grasp pose candidates. We select the most effective grasp pose by balancing the grasp score and the angle between the robot’s approaching direction and the world frame’s z-axis. Conditioned on the text instruction, SO FAR predicts the target object’s translation and rotation matrix, which define the transformation from the grasp pose to the placement pose. An open-source motion planning module [121] is then used to generate a collision-free trajectory. In addition, we set the initial joint position as the midpoint to achieve smooth motion while reducing collisions between the manipulated object and the environment.

## IV. EXPERIMENTS

### A. Benchmarks

We propose two benchmarks to demonstrate the effectiveness of our SO FAR in spatial reasoning and robotic manipulation.

1) *Open6DOR V2*: For simulation experiments, we choose the Open6DOR[28] Benchmark to comprehensively evaluate our spatial understanding abilities. Beyond the perception tasks that it originally proposes, we further construct an execution track to enable comparison with close-loop policies. We name the new combined benchmark Open6DOR V2.

- **Perception tasks.** In line with Open6DOR’s definition, the model takes an RGB-D scene image along with a language instruction as input and directly outputs the translation and orientation of the target object.
- **Execution tasks.** We replicate Open6DOR scenes and ground them into a robosuite simulation environment for execution, excluding single-object scenes to better assess the understanding of spatial relationships. The model takes the RGB-D image with a language instruction as input and completes the entire execution process. We build it based on robosuite [173] and adopt the format established by LIBERO [76]. Evaluation is conducted according to the final position and orientation of the target object.

2) *6-DoF SpatialBench*: To further evaluate spatial understanding that requires 6-DoF awareness, we propose a VQA

TABLE II: **SimplerEnv [74] simulation valuation results for the Google Robot setup.** We present success rates for the “Variant Aggregation” and “Visual Matching” approaches. Top-1 & Top-2 accuracies are represented using different colors, bold text, and underlines. OXE: Open X-Embodiment dataset [96].

Google Robot Evaluation Setup	Policy	Training Data	Pick Coke Can				Move Near	Open / Close Drawer			Average
			Horizontal Laying	Vertical Laying	Standing	Average	Average	Open	Close	Average	
Variant Aggregation	RT-1-X [96]	OXE	0.569	0.204	0.698	0.490	0.323	0.069	<b>0.519</b>	0.294	0.397
	RT-2-X [174]	OXE	<u>0.822</u>	<u>0.754</u>	<u>0.893</u>	<u>0.823</u>	<b>0.792</b>	<b>0.333</b>	0.372	<b>0.353</b>	<u>0.661</u>
	Octo-Base [126]	OXE	0.005	0.000	0.013	0.006	0.031	0.000	0.021	0.011	0.012
	OpenVLA [64]	OXE	0.711	0.271	0.653	0.545	0.477	0.158	0.195	0.177	0.411
	<b>SoFAR</b>	<b>Zero-Shot</b>	<b>0.861</b>	<b>0.960</b>	<b>0.901</b>	<b>0.907</b>	<u>0.740</u>	<u>0.200</u>	<u>0.394</u>	<u>0.297</u>	<b>0.676</b>
Visual Matching	RT-1-X [96]	OXE	<b>0.820</b>	0.330	0.550	0.567	0.317	<b>0.296</b>	<b>0.891</b>	<b>0.597</b>	0.534
	RT-2-X [174]	OXE	0.740	<u>0.740</u>	<u>0.880</u>	<u>0.787</u>	<u>0.779</u>	0.157	0.343	0.250	<u>0.606</u>
	Octo-Base [126]	OXE	0.210	0.210	0.090	0.170	0.042	0.009	0.444	0.227	0.168
	OpenVLA [64]	OXE	0.270	0.030	0.190	0.163	0.462	0.194	0.518	0.356	0.277
	<b>SoFAR</b>	<b>Zero-Shot</b>	<u>0.770</u>	<b>1.000</b>	<b>1.000</b>	<b>0.923</b>	<b>0.917</b>	<u>0.227</u>	<u>0.578</u>	<u>0.403</u>	<b>0.749</b>

TABLE III: **SimplerEnv [74] simulation evaluation results for the WidowX + Bridge setup.** We report both the final success rate (“Success”) along with partial success (e.g., “Grasp Spoon”). Top-1 & Top-2 accuracies are represented using different colors, bold text, and underlines. OXE: Open X-Embodiment dataset [96]. Bridge: BridgeData V2 dataset [136].

Policy	Training Data	Put Spoon on Towel		Put Carrot on Plate		Stack Green Block on Yellow Block		Put Eggplant in Yellow Basket		Average
		Grasp Spoon	Success	Grasp Carrot	Success	Grasp Green Block	Success	Grasp Eggplant	Success	
RT-1-X [8]	OXE	0.167	0.000	0.208	0.042	0.083	0.000	0.000	0.000	0.011
Octo-Base [126]	OXE	0.347	0.125	<u>0.528</u>	0.083	0.319	0.000	0.667	0.431	0.160
Octo-Small [126]	OXE	<b>0.778</b>	<u>0.472</u>	<u>0.278</u>	0.097	0.403	0.042	<u>0.875</u>	0.569	0.300
OpenVLA [64]	OXE	0.041	0.000	0.333	0.000	0.125	0.000	0.083	0.041	0.010
RoboVLM [72]	OXE	0.375	0.208	0.333	0.250	0.083	0.083	0.000	0.000	0.135
RoboVLM [72]	Bridge	0.542	0.292	<u>0.250</u>	0.250	0.458	0.125	0.583	<u>0.583</u>	0.313
SpatialVLA [112]	OXE	0.250	0.208	0.417	0.208	0.583	0.250	0.792	0.708	0.344
SpatialVLA [112]	Bridge	0.208	0.167	0.292	0.250	<u>0.625</u>	<u>0.292</u>	<b>1.000</b>	<b>1.000</b>	<u>0.427</u>
<b>SoFAR</b>	<b>Zero-Shot</b>	<u>0.625</u>	<b>0.583</b>	<b>0.750</b>	<b>0.667</b>	<b>0.917</b>	<b>0.708</b>	0.667	0.375	<b>0.583</b>

benchmark for spatial understanding evaluation, named 6-DoF SpatialBench. Previous benchmarks [10, 15, 33, 124] for assessing VLMs primarily focus on understanding spatial positions, with little attention given to orientation. Additionally, most assessments emphasize relative and imprecise spatial relationships (e.g., “to the left,” “nearest”) while lacking quantitative metrics. In contrast, 6-DoF SpatialBench focuses on both positional and orientational understanding, encompassing 223 manual annotated samples, divided into the position track and orientation track depending on the question. Each task includes an RGB image of a scene and multiple-choice questions with four options. Specifically, the tasks cover numerical queries (e.g., counting), positional relationships (e.g., left / right), and orientation (the facing direction of an object). The model needs to analyze the image and select the correct answer from four options. All questions and ground-truth answers are carefully designed through **human annotation**.

### B. 6-DoF object rearrangement evaluation in Simulation

We conduct experiments on the proposed Open6DOR V2 benchmark, with results presented in Table I. In perception tasks, we compare our results against the same baselines used in the original Open6DOR [28] experiments. SoFAR outperforms all baselines, demonstrating effective spatial understanding and zero-shot generalizability. In the execution tasks, we record the initial and final poses of the objects to evaluate execution success. We use the original pretrained Octo [126] and the LIBERO-finetuned OpenVLA [64] as baselines, conducting all experiments in the same LIBERO environment where OpenVLA was fine-tuned to minimize the domain gap. Despite this, both Octo and OpenVLA show lower success rates, indicating their poor generalizability. In contrast, SoFAR achieves about 40% success, even with a vanilla execution implementation. It is worth noting that some of the objects are inherently difficult to grasp, which significantly hampers execution success.

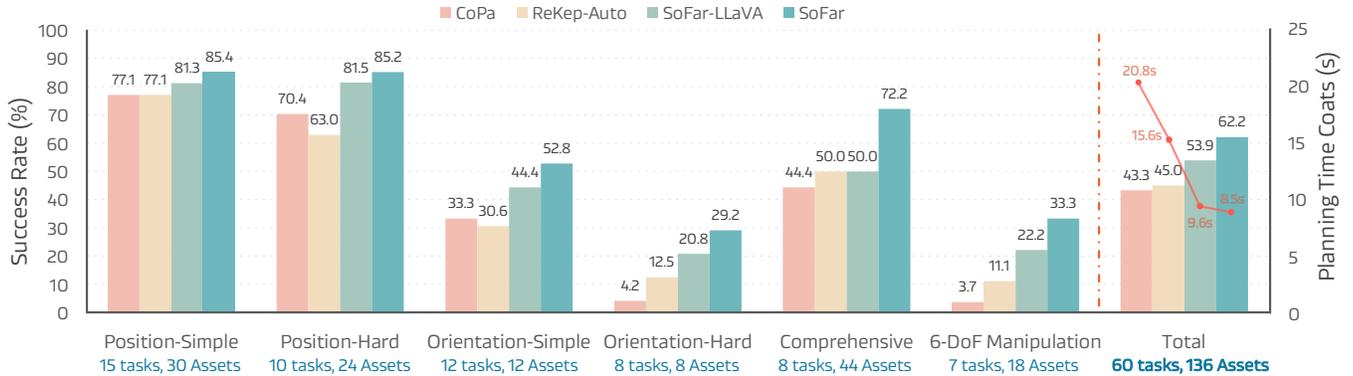


Fig. 7: **Quantitative evaluation** of zero-shot real-world language-grounded rearrangement. We design **60** diverse real-world tasks involving over **100** diverse objects (detailed in Table XIII). The Success (%) is obtained with 3 trials per task and method.

TABLE IV: **Spatial comprehension evaluation** on our proposed 6-DoF SpatialBench. Depth-Esti: Use monocular depth estimation methods like Metric3D [157] or Moge [138]. rel.: Relative metric evaluation, abs.: Absolute metric evaluation.

Method	Depth-Esti	Position		Orientation		Total
		rel.	abs.	rel.	abs.	
<i>Blind Evaluation with Large Language Models</i>						
GPT-3.5-Turbo [9]	✗	24.5	24.9	26.7	27.5	25.7
GPT-4-Turbo [99]	✗	27.2	27.3	29.2	27.9	27.8
<i>General Vision Language Models</i>						
LLaVA-1.5 [80]	✗	30.9	24.5	28.3	25.8	27.2
GPT-4o-mini [100]	✗	33.3	26.9	32.5	23.8	31.0
GPT-4V [98]	✗	37.7	32.7	36.7	27.5	33.9
GPT-4o [100]	✗	49.4	28.4	44.2	25.8	36.2
<i>Vision Language Models with Spatial Awareness</i>						
SpaceLLaVA [12]	✗	32.4	30.5	30.9	24.9	28.2
SpaceMantis [12]	✗	33.6	29.2	27.2	25.0	28.9
SpatialBot [10]	✓	50.9	21.6	39.6	22.9	32.7
RoboPoint [160]	✗	43.8	30.8	33.8	25.8	33.5
<b>SO FAR</b>	✓	<b>59.6</b>	<b>33.8</b>	<b>54.6</b>	<b>31.3</b>	<b>43.9</b>

We call for more robust execution policies and manipulation strategies, such as prehensile grasping and adaptive techniques, to demonstrate better performance on Open6DOR V2.

### C. Zero-shot Object Manipulation Evaluation in Simulation

SIMPLER [74] is a suite of open-source simulated evaluation environments designed for real-world robot manipulation setups. SIMPLER provides a standardized platform for benchmarking manipulation tasks, emphasizing reproducibility and alignment with real-world scenarios. We conduct quantitative evaluations of SO FAR’s zero-shot execution performance on Google Robot tasks & Widow-X tasks and compare it to baselines including Octo [126], OpenVLA [64] and more concurrent works [72, 112]. The robot follows the planned trajectory generated by the planning module, as described in Sec. III-B2, to execute the task. Furthermore, leveraging the error detection

TABLE V: **Semantic Orientation evaluation** on our proposed OrientText300K dataset test split.

Method	45°	30°	15°	5°	Average
PointSO-S	77.34	74.22	67.97	60.94	70.12
<b>PointSO-B</b>	<b>79.69</b>	<b>77.34</b>	<b>70.31</b>	<b>62.50</b>	<b>72.46</b>
<b>PointSO-L</b>	<b>81.25</b>	<b>78.13</b>	<b>72.66</b>	<b>65.63</b>	<b>74.42</b>

TABLE VI: **Zero-shot Semantic Direction evaluation** of robustness on OrientText300K test split. Single-View: randomly select a camera viewpoint within the unit sphere and generate a **single viewpoint** within the FoV on polar coordinates. Jitter: Gaussian jittering with noise  $\epsilon \sim \mathcal{N}(0, \sigma^2)$  and  $\sigma = 0.01$ . Rotate: random SO(3) rotation sampling over X-Y-Z Euler angle  $(\alpha, \beta, \gamma) \sim \mathcal{U}(-\theta, \theta)$  and  $\theta = \pi$ . All: All the corruptions.

Method	OrientText300K-C Variants			
	Single-View	Jitter	Rotate	All
PointSO-S	72.66	76.56	73.43	67.19
<b>PointSO-B</b>	<b>75.00</b>	<b>78.90</b>	<b>75.78</b>	<b>71.09</b>
<b>PointSO-L</b>	<b>76.56</b>	<b>81.25</b>	<b>77.34</b>	<b>74.22</b>

and re-planning capabilities of VLMs [100, 125], we can make multiple attempts following a single-step execution failure to approximate achieve a closed-loop effect. For fairness, we limit the maximum number of attempts to three. Detailed visualizations and analyses are provided in the Appendix C. As shown in Tables II and III, despite the training data for Octo and OpenVLA including Google Robot tasks, SO FAR demonstrates superior zero-shot performance compared to most baselines.

### D. Zero-shot Real-world Manipulation

1) *Hardware Setup*: We set up a real-world tabletop environment utilizing a Franka Panda robotic arm equipped with a parallel gripper. For perception, we integrate a single RGB-D camera (Intel RealSense D415) mounted on the end-effector. The details of the environmental visualization and additional real-world robot setup are provided in the Appendix A.



Fig. 8: **Realworld Orientation-Aware Navigation.** We present both the third-person view and the egocentric view, annotating the predicted orientation of the interacted objects.

2) *Tasks and Evaluations:* To comprehensively evaluate the generalization of SOFAR, we design **60 diverse real-world experimental tasks involving over 100 diverse objects**. Following Open6DOR [28]. These tasks are categorized into three tracks—position, orientation, and comprehensive & 6-DoF, each with simple and hard levels. The position-simple track focuses on basic spatial relationships (front/back/left/right), while the position-hard track involves more complex spatial concepts (between/center/customized). The orientation-simple track targets object part orientations, whereas the orientation-hard track requires precise angle judgments for upright or flipped objects. Comprehensive tasks demand intricate instruction understanding and long-horizon operations, and 6-DoF tasks require simultaneous position and orientation control. Each task is executed three times to ensure statistical reliability. Detailed visualization and analysis are provided in the Appendix A.

3) *Results:* The quantitative results are shown in Fig. 7, SOFAR outperform the baseline across all tracks, particularly in the challenging orientation and comprehensive & 6-DoF tasks. Meanwhile, SOFAR also utilizes the minimal planning time overhead. In addition, SOFAR is not limited to a single embodiment. In qualitative experiments, we also test different embodiments, such as a dexterous hand and a suction cup, as shown in Fig. 6. Additional robot setups and generalization experiments are included in the Appendix A.

#### E. Visual Question Answer on 6-DoF SpatialBench

We evaluate SOFAR on our proposed 6-DoF SpatialBench, as shown in Table IV. We choose several VLMs and comparable methods as baselines. SOFAR achieves superior performance across both position-track and orientation-track, outperforming baselines by over 18%.

#### F. Semantic Orientation Prediction

Using free-text descriptions to extract semantic orientations from object point clouds is challenging. In Objaverse [22], we manually annotate 128 diverse objects and form an OrientText300K test split to evaluate the directional accuracy of PointSO. We trained different model variants on OrientText300K, and the results in Table V, report varying accuracies

between  $45^\circ \sim 5^\circ$ . In the real world, acquiring complete point clouds is difficult even impractical. To assess the robustness, we introduce random rotations, single-sided point clouds, and Gaussian noise. The accuracy at  $45^\circ$ , as shown in Table VI, demonstrates the model’s performance under these conditions.

#### G. Additional Applications

1) *Cross Embodiment Generalization:* Our approach determines grasp poses by generating masks and plans the target pose and transformation using our PointSO and large language model. It does not rely on trajectory data specific to any robotic arm, making SOFAR embodiment-agnostic. Fig. 6 illustrates the diverse embodiments employed in our real-world experiments. Leveraging the GSNet [137] algorithm based on LeapHand [116], we perform 6-DoF object manipulation experiments on dexterous hands. We conduct three position-related and three rotation-related experiments. Leveraging the PointSO and large language models, SOFAR is capable of performing complex 6-DoF manipulation tasks, such as “Upright the fallen wine glass and arrange it neatly in a row with the other wine glasses.”

2) *Long Horizon Planning & Close-Loop Execution:* Similar to ReKep [56], SOFAR leverages VLMs [100, 125] to perform long-horizon decomposition of complex tasks and employs dual-system VLMs [100, 125] to determine the success of execution between tasks and subtasks, enabling closed-loop execution. When a discrepancy between the results and expectations is detected, SOFAR re-perceives and re-executes the current subtask. Detailed visualization and analysis can be found in Appendices C and D.

3) *Orientation-Aware Robotic Navigation:* *Semantic orientation* can not only be applied to manipulation tasks but also to robotic navigation tasks. As shown in Fig. 8, we conduct experiments using a robotic dog for orientation-aware robotic navigation tasks. Unlike traditional navigation approaches, the dog is required to face a specific direction during navigation. This orientation-aware constraint enhances the navigation process by ensuring precise alignment with the desired orientation, thereby improving task performance in scenarios where directionality is critical.

## V. RELATED WORKS

### A. Vision-Language Models for Spatial Understanding

Vision-Language Models (VLMs) are rapidly being developed in research community, driven by the storming lead in extending GPT-style [9, 113, 114] Large Language Models (LLMs) like LLaMA [130, 131] to VLMs [1, 4, 23, 30, 32, 79, 80, 123, 129, 165, 167]. SpatialVLM [12] pioneers this direction by constructing VQA data in spatial understanding from RGB-D, which is used for training an RGB-only VLM. Following SpatialVLM, SpatialRGPT [15] extends RGB-based spatial understanding to RGB-D by constructing spatial understanding data using 3D scene graphs. SpatialBot [10] explores RGB-D spatial reasoning through hierarchical depth-based reasoning. Some other works propose visual prompting for improving GPT-4V’s spatial understanding [77, 88, 151]. Meanwhile, another line of works explores VLMs using 3D representations such as point clouds for 3D scene [42, 50] and object-centric [109, 110, 149] understanding. Despite the remarkable progress, these works are limited in 3-DoF understanding which is not actionable. In contrast, we explore spatial understanding in 6-DoFs from RGB-D via VLMs. Unlike vanilla 3D scene graphs used by SpatialRGPT for data construction, we propose orientation-aware 3D scene graphs realized by our proposed PointSO. In addition, we formulate spatial understanding as graph learning, where the scene graph nodes are directly input during inference.

### B. Language-Grounded Robot Manipulation

Language-grounded robot Manipulation adopts the human language as a general instruction interface. Existing works can be categorized into two groups: **i)** *End-to-end* models like RT-series [5, 8, 174] built upon unified cross-modal Transformers with tokenized actions [7, 45, 73, 78, 92, 117, 168], large vision-language-action (VLA) models built from VLMs [64], or 3D representations [13, 160, 169]. Training on robot data such as Open X-Embodiment [96] and DROID [63], a remarkable process has been made. However, the data *scale* is still limited compared to in-the-wild data for training VLMs. **ii)** *Decoupled* high-level reasoning and low-level actions in large VLMs and small off-the-shelf policy models, primitives [39, 51, 53–55, 57, 59, 75, 90, 94, 153, 161], or articulated priors [52, 71]. Our SOFAR lies in this group, where an open-world generalization property emerges from VLMs and our proposed PointSO empowered by orientation-aware spatial understanding.

## VI. LIMITATIONS

One notable limitation for decoupled systems like SOFAR is that the execution may fail due to a sub-module error, *i.e.*, robots may place target objects with an error transformation because of unstable grasping or inaccurate visual perception. For example, the pen will be placed in an unexpected pose due to the rotation during execution. Future works include integrating scalable data and more advanced models and exploring the potential of combining end-to-end and such decoupled methods, and expanding SOFAR to more applications.

## VII. CONCLUSIONS

In this paper, we propose semantic orientation, a language-grounded representation that defines object orientations via intuitive descriptors (*e.g.*, “plug-in direction”), bridging geometric reasoning and functional semantics. To enable this, we introduce OrientText300K, a large-scale dataset of 3D models annotated with semantic orientation. Through PointSO and the integrated SOFAR system, we significantly enhance robotic manipulation capabilities, as demonstrated by strong performance in both simulated and real-world experiments.

## REFERENCES

- [1] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katie Millican, Malcolm Reynolds, Roman Ring, Eliza Rutherford, Serkan Cabi, Tengda Han, Zhitao Gong, Sina Samangooei, Marianne Monteiro, Jacob Menick, Sebastian Borgeaud, Andrew Brock, Aida Nematzadeh, Sahand Sharifzadeh, Mikolaj Binkowski, Ricardo Barreira, Oriol Vinyals, Andrew Zisserman, and Karen Simonyan. Flamingo: a visual language model for few-shot learning. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2022. 2, 11
- [2] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2016. 27
- [3] Jinze Bai, Shuai Bai, Shusheng Yang, Shijie Wang, Sinan Tan, Peng Wang, Junyang Lin, Chang Zhou, and Jingren Zhou. Qwen-vl: A frontier large vision-language model with versatile abilities. *arXiv preprint arXiv:2308.12966*, 2023. 24, 27
- [4] Yutong Bai, Xinyang Geng, Karttikeya Mangalam, Amir Bar, Alan L. Yuille, Trevor Darrell, Jitendra Malik, and Alexei A. Efros. Sequential modeling enables scalable learning for large vision models. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2024. 11
- [5] Suneel Belkhal, Tianli Ding, Ted Xiao, Pierre Sermanet, Quon Vuong, Jonathan Tompson, Yevgen Chebotar, Debidatta Dwibedi, and Dorsa Sadigh. RT-H: action hierarchies using language. *CoRR*, abs/2403.01823, 2024. 11
- [6] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri S. Chatterji, Annie S. Chen, Kathleen Creel, Jared Quincy Davis, Dorottya Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah D. Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain,

- Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark S. Krass, Ranjay Krishna, Rohith Kuditipudi, and et al. On the opportunities and risks of foundation models. *CoRR*, abs/2108.07258, 2021. [26](#)
- [7] Konstantinos Bousmalis, Giulia Vezzani, Dushyant Rao, Coline Manon Devin, Alex X. Lee, Maria Bauza Villalonga, Todor Davchev, Yuxiang Zhou, Agrim Gupta, Akhil Raju, Antoine Laurens, Claudio Fantacci, Valentin Dalibard, Martina Zambelli, Murilo Fernandes Martins, Rugile Pevceviute, Michiel Blokzijl, Misha Denil, Nathan Batchelor, Thomas Lampe, Emilio Parisotto, Konrad Zolna, Scott Reed, Sergio Gómez Colmenarejo, Jonathan Scholz, Abbas Abdolmaleki, Oliver Groth, Jean-Baptiste Regli, Oleg Sushkov, Thomas Rothörl, Jose Enrique Chen, Yusuf Aytar, David Barker, Joy Ortiz, Martin Riedmiller, Jost Tobias Springenberg, Raia Hadsell, Francesco Nori, and Nicolas Heess. Robocat: A self-improving generalist agent for robotic manipulation. *Transactions on Machine Learning Research*, 2024. ISSN 2835-8856. [11](#)
- [8] Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Jasmine Hsu, Julian Ibarz, Brian Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J. Joshi, Ryan Julian, Dmitry Kalashnikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deeksha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl Pertsch, Jornell Quiambao, Kanishka Rao, Michael S. Ryoo, Grecia Salazar, Pannag R. Sanketi, Kevin Sayed, Jaspier Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong T. Tran, Vincent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu, and Brianna Zitkovich. RT-1: robotics transformer for real-world control at scale. In *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*, 2023. [8](#), [11](#)
- [9] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Nee-lakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2020. [9](#), [11](#), [31](#)
- [10] Wenxiao Cai, Yaroslav Ponomarenko, Jianhao Yuan, Xiaoqi Li, Wankou Yang, Hao Dong, and Bo Zhao. Spatialbot: Precise spatial understanding with vision language models. *CoRR*, abs/2406.13642, 2024. [2](#), [8](#), [9](#), [11](#), [22](#), [23](#), [24](#)
- [11] Matthew Chang, Theophile Gervet, Mukul Khanna, Sri-ran Yenamandra, Dhruv Shah, So Yeon Min, Kavitha Shah, Chris Paxton, Saurabh Gupta, Dhruv Batra, Roozbeh Mottaghi, Jitendra Malik, and Devendra Singh Chaplot. GOAT: GO to Any Thing. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. [30](#)
- [12] Boyuan Chen, Zhuo Xu, Sean Kirmani, Brian Ichter, Danny Driess, Pete Florence, Dorsa Sadigh, Leonidas Guibas, and Fei Xia. SpatialVLM: Endowing vision-language models with spatial reasoning capabilities. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2024. [2](#), [9](#), [11](#)
- [13] Shizhe Chen, Ricardo Garcia Pinel, Cordelia Schmid, and Ivan Laptev. Polarnet: 3d point clouds for language-guided robotic manipulation. In *Conference on Robot Learning, CoRL 2023, 6-9 November 2023, Atlanta, GA, USA*, volume 229 of *Proceedings of Machine Learning Research*, pages 1761–1781. PMLR, 2023. [11](#)
- [14] Zixuan Chen, Xialin He, Yen-Jen Wang, Qiayuan Liao, Yanjie Ze, Zhongyu Li, S Shankar Sastry, Jiajun Wu, Koushil Sreenath, Saurabh Gupta, et al. Learning smooth humanoid locomotion through lipschitz-constrained policies. *arXiv preprint arXiv:2410.11825*, 2024. [30](#)
- [15] 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 model. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2024. [2](#), [8](#), [11](#)
- [16] Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xingang Wang, and Ying Shan. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16901–16911, 2024. [26](#), [27](#)
- [17] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive Whole-Body Control for Humanoid Robots. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. [30](#)
- [18] Jaime Corsetti, Davide Boscaini, Changjae Oh, Andrea Cavallaro, and Fabio Poiesi. Open-vocabulary object 6d pose estimation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 18071–18080. IEEE, 2024. [2](#)
- [19] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017. [27](#)
- [20] Wenliang Dai, Junnan Li, Dongxu Li, Anthony Meng Huat Tiong, Junqi Zhao, Weisheng Wang, Boyang Li, Pascale Fung, and Steven C. H. Hoi. Instructblip: Towards general-purpose vision-language models with instruction tuning. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2023. [24](#)
- [21] Matt Deitke, Ruoshi Liu, Matthew Wallingford, Huong

- Ngo, Oscar Michel, Aditya Kusupati, Alan Fan, Christian Laforte, Vikram Voleti, Samir Yitzhak Gadre, Eli VanderBilt, Aniruddha Kembhavi, Carl Vondrick, Georgia Gkioxari, Kiana Ehsani, Ludwig Schmidt, and Ali Farhadi. Objaverse-xl: A universe of 10m+ 3d objects. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2023. 26, 30
- [22] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2023. 4, 10, 27
- [23] 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, Yen-Sung 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 Sarnat, 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 B. Girshick, Ali Farhadi, and Aniruddha Kembhavi. Molmo and pixmo: Open weights and open data for state-of-the-art multimodal models. *CoRR*, abs/2409.17146, 2024. 11
- [24] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *AAAI Conf. Artif. Intell. (AAAI)*, 2021. 27
- [25] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics, 2019. 29
- [26] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. PLA: language-driven open-vocabulary 3d scene understanding. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2023. 29
- [27] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. Lowis3d: Language-driven open-world instance-level 3d scene understanding. *IEEE Trans. Pattern Anal. Mach. Intell. (TPAMI)*, pages 1–16, 2024. 29
- [28] Yufei Ding, Haoran Geng, Chaoyi Xu, Xiaomeng Fang, Jiazhao Zhang, Songlin Wei, Qiyu Dai, Zhizheng Zhang, and He Wang. Open6DOR: Benchmarking open-instruction 6-dof object rearrangement and a VLM-based approach. In *IEEE/RSJ Int. Conf. Intell. Robot. and Syst. (IROS)*, 2024. 7, 8, 10, 26, 27, 29
- [29] Runpei Dong, Zekun Qi, Linfeng Zhang, Junbo Zhang, Jianjian Sun, Zheng Ge, Li Yi, and Kaisheng Ma. Autoencoders as cross-modal teachers: Can pretrained 2d image transformers help 3d representation learning? In *Int. Conf. Learn. Represent. (ICLR)*, 2023. 2, 4, 27, 29, 30
- [30] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, Xiangwen Kong, Xiangyu Zhang, Kaisheng Ma, and Li Yi. DreamLLM: Synergistic multimodal comprehension and creation. In *Int. Conf. Learn. Represent. (ICLR)*, 2024. 2, 11
- [31] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *Int. Conf. Learn. Represent. (ICLR)*, 2021. 4, 29
- [32] Danny Driess, Fei Xia, Mehdi S. M. Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, Wenlong Huang, Yevgen Chebotar, Pierre Sermanet, Daniel Duckworth, Sergey Levine, Vincent Vanhoucke, Karol Hausman, Marc Toussaint, Klaus Greff, Andy Zeng, Igor Mordatch, and Pete Florence. Palm-e: An embodied multimodal language model. In *Int. Conf. Mach. Learn. (ICML)*, 2023. 11
- [33] Mengfei Du, Binhao Wu, Zejun Li, Xuanjing Huang, and Zhongyu Wei. Embspatial-bench: Benchmarking spatial understanding for embodied tasks with large vision-language models. *arXiv preprint arXiv:2406.05756*, 2024. 2, 8, 22, 23, 24
- [34] Jiafei Duan, Wentao Yuan, Wilbert Pumacay, Yi Ru Wang, Kiana Ehsani, Dieter Fox, and Ranjay Krishna. Manipulate-anything: Automating real-world robots using vision-language models. In *8th Annual Conference on Robot Learning*, 2024. 2
- [35] Ben Eisner, Harry Zhang, and David Held. Flowbot3d: Learning 3d articulation flow to manipulate articulated objects. *arXiv preprint arXiv:2205.04382*, 2022. 24
- [36] Nico Engel, Vasileios Belagiannis, and Klaus Dietmayer. Point transformer. *IEEE Access*, 9:134826–134840, 2021. 27, 29
- [37] Guofan Fan, Zekun Qi, Wenkai Shi, and Kaisheng Ma. Point-gcc: Universal self-supervised 3d scene pre-training via geometry-color contrast. *CoRR*, abs/2305.19623, 2023. 29
- [38] Haoshu Fang, Chenxi Wang, Minghao Gou, and Cewu Lu. Graspnet-1billion: A large-scale benchmark for general object grasping. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 11441–11450. Computer Vision Foundation / IEEE, 2020. 7

- [39] Kuan Fang, Fangchen Liu, Pieter Abbeel, and Sergey Levine. MOKA: Open-World Robotic Manipulation through Mark-Based Visual Prompting. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. 11
- [40] Li Fei-Fei. *The Worlds I See: Curiosity, Exploration, and Discovery at the Dawn of AI*. Flatiron books: a moment of lift book, 2023. 2
- [41] Li Fei-Fei, Asha Iyer, Christof Koch, and Pietro Perona. What do we perceive in a glance of a real-world scene? *Journal of vision*, 7(1):10–10, 2007. 2
- [42] Rao Fu, Jingyu Liu, Xilun Chen, Yixin Nie, and Wenhan Xiong. Scene-llm: Extending language model for 3d visual understanding and reasoning. *CoRR*, abs/2403.11401, 2024. 11
- [43] JJ Gibson. The theory of affordances. *Perceiving, acting and knowing: Towards an ecological psychology/Erlbaum*, 1977. 29
- [44] Mohit Goyal, Sahil Modi, Rishabh Goyal, and Saurabh Gupta. Human hands as probes for interactive object understanding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 3283–3293. IEEE, 2022. 29
- [45] Markus Grotz, Mohit Shridhar, Yu-Wei Chao, Tamim Asfour, and Dieter Fox. Peract2: Benchmarking and learning for robotic bimanual manipulation tasks. In *CoRL 2024 Workshop on Whole-body Control and Bimanual Manipulation: Applications in Humanoids and Beyond*, 2024. 11
- [46] Abdullah Hamdi, Silvio Giancola, and Bernard Ghanem. MVTN: multi-view transformation network for 3d shape recognition. In *Int. Conf. Comput. Vis. (ICCV)*, pages 1–11. IEEE, 2021. 27
- [47] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022. 29, 30
- [48] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. In *Annu. Conf. Robot. Learn. (CoRL)*, 2024. 30
- [49] Xialin He, Runpei Dong, Zixuan Chen, and Saurabh Gupta. Learning getting-up policies for real-world humanoid robots. *arXiv preprint arXiv:2502.12152*, 2025. 30
- [50] 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 *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2023. 11
- [51] Haoxu Huang, Fanqi Lin, Yingdong Hu, Shengjie Wang, and Yang Gao. Copa: General robotic manipulation through spatial constraints of parts with foundation models. *CoRR*, abs/2403.08248, 2024. 7, 11, 23, 28
- [52] Siyuan Huang, Haonan Chang, Yuhan Liu, Yimeng Zhu, Hao Dong, Peng Gao, Abdeslam Boularias, and Hongsheng Li. A3vlm: Actionable articulation-aware vision language model. *arXiv preprint arXiv:2406.07549*, 2024. 11
- [53] Wenlong Huang, Fei Xia, Ted Xiao, Harris Chan, Jacky Liang, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Tomas Jackson, Noah Brown, Linda Luu, Sergey Levine, Karol Hausman, and Brian Ichter. Inner monologue: Embodied reasoning through planning with language models. In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pages 1769–1782. PMLR, 2022. 11
- [54] Wenlong Huang, Chen Wang, Ruohan Zhang, Yunzhu Li, Jiajun Wu, and Li Fei-Fei. Voxposer: Composable 3d value maps for robotic manipulation with language models. In *Annu. Conf. Robot. Learn. (CoRL)*, 2023. 7, 31
- [55] Wenlong Huang, Fei Xia, Dhruv Shah, Danny Driess, Andy Zeng, Yao Lu, Pete Florence, Igor Mordatch, Sergey Levine, Karol Hausman, and Brian Ichter. Grounded decoding: Guiding text generation with grounded models for embodied agents. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. 11
- [56] Wenlong Huang, Chen Wang, Yunzhu Li, Ruohan Zhang, and Li Fei-Fei. Rekep: Spatio-temporal reasoning of relational keypoint constraints for robotic manipulation. In *Annu. Conf. Robot. Learn. (CoRL)*, 2024. 10, 22, 23, 28, 29, 30, 31
- [57] Brian Ichter, Anthony Brohan, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Alexander Herzog, Daniel Ho, Julian Ibarz, Alex Irpan, Eric Jang, Ryan Julian, Dmitry Kalashnikov, Sergey Levine, Yao Lu, Carolina Parada, Kanishka Rao, Pierre Sermanet, Alexander Toshev, Vincent Vanhoucke, Fei Xia, Ted Xiao, Peng Xu, Mengyuan Yan, Noah Brown, Michael Ahn, Omar Cortes, Nicolas Sievers, Clayton Tan, Sichun Xu, Diego Reyes, Jarek Rettinghouse, Jornell Quiambao, Peter Pastor, Linda Luu, Kuang-Huei Lee, Yuheng Kuang, Sally Jesmonth, Nikhil J. Joshi, Kyle Jeffrey, Rosario Jauregui Ruano, Jasmine Hsu, Keerthana Gopalakrishnan, Byron David, Andy Zeng, and Chuyuan Kelly Fu. Do as I can, not as I say: Grounding language in robotic affordances. In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pages 287–318. PMLR, 2022. 2, 11, 29
- [58] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carneyn, et al. Openai o1 system card. *arXiv preprint arXiv:2412.16720*, 2024.

- [59] Hanxiao Jiang, Binghao Huang, Ruihai Wu, Zhuoran Li, Shubham Garg, Hooshang Nayyeri, Shenlong Wang, and Yunzhu Li. Roboexp: Action-conditioned scene graph via interactive exploration for robotic manipulation. In *Annu. Conf. Robot. Learn. (CoRL)*, 2024. 11
- [60] Wolfgang Kabsch. A solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 32(5):922–923, 1976. 7
- [61] Wolfgang Kabsch. A discussion of the solution for the best rotation to relate two sets of vectors. *Acta Crystallographica Section A: Crystal Physics, Diffraction, Theoretical and General Crystallography*, 34(5): 827–828, 1978. 7
- [62] Ivan Kapelyukh, Yifei Ren, Ignacio Alzugaray, and Edward Johns. Dream2real: Zero-shot 3d object rearrangement with vision-language models. In *IEEE International Conference on Robotics and Automation, ICRA 2024, Yokohama, Japan, May 13-17, 2024*, pages 4796–4803. IEEE, 2024. 7
- [63] Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis, Peter David Fagan, Joey Hejna, Masha Itkina, Marion Lepert, Yecheng Jason Ma, Patrick Tree Miller, Jimmy Wu, Suneel Belkhale, Shivin Dass, Huy Ha, Arhan Jain, Abraham Lee, Youngwoon Lee, Marius Memmel, Sungjae Park, Ilija Radosavovic, Kaiyuan Wang, Albert Zhan, Kevin Black, Cheng Chi, Kyle Beltran Hatch, Shan Lin, Jingpei Lu, Jean Mercat, Abdul Rehman, Pannag R Sanketi, Archit Sharma, Cody Simpson, Quan Vuong, Homer Rich Walke, Blake Wulfe, Ted Xiao, Jonathan Heewon Yang, Arefeh Yavary, Tony Z. Zhao, Christopher Agia, Rohan Baijal, Mateo Guaman Castro, Daphne Chen, Qiuyu Chen, Trinity Chung, Jaimyn Drake, Ethan Paul Foster, Jensen Gao, David Antonio Herrera, Minh Heo, Kyle Hsu, Jiaheng Hu, Donovan Jackson, Charlotte Le, Yunshuang Li, Roy Lin, Zehan Ma, Abhiram Maddukuri, Suvir Mirchandani, Daniel Morton, Tony Nguyen, Abigail O’Neill, Rosario Scalise, Derick Seale, Victor Son, Stephen Tian, Emi Tran, Andrew E. Wang, Yilin Wu, Annie Xie, Jingyun Yang, Patrick Yin, Yunchu Zhang, Osbert Bastani, Glen Berseth, Jeannette Bohg, Ken Goldberg, Abhinav Gupta, Abhishek Gupta, Dinesh Jayaraman, Joseph J Lim, Jitendra Malik, Roberto Martín-Martín, Subramanian Ramamoorthy, Dorsa Sadigh, Shuran Song, Jiajun Wu, Michael C. Yip, Yuke Zhu, Thomas Kollar, Sergey Levine, and Chelsea Finn. DROID: A Large-Scale In-The-Wild Robot Manipulation Dataset. In *Proceedings of Robotics: Science and Systems*, Delft, Netherlands, July 2024. 11
- [64] Moo Jin Kim, Karl Pertsch, Siddharth Karamcheti, Ted Xiao, Ashwin Balakrishna, Suraj Nair, Rafael Rafailov, Ethan Paul Foster, Grace Lam, Pannag 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. *CoRR*, abs/2406.09246, 2024. 7, 8, 9, 11
- [65] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 2, 5, 7, 24, 26
- [66] Ranjay Krishna, Yuke Zhu, Oliver Groth, Justin Johnson, Kenji Hata, Joshua Kravitz, Stephanie Chen, Yannis Kalantidis, Li-Jia Li, David A. Shamma, Michael S. Bernstein, and Li Fei-Fei. Visual genome: Connecting language and vision using crowdsourced dense image annotations. *Int. J. Comput. Vis.*, 123(1):32–73, 2017. 2
- [67] Yann Labbé, Lucas Manuelli, Arsalan Mousavian, Stephen Tyree, Stan Birchfield, Jonathan Tremblay, Justin Carpentier, Mathieu Aubry, Dieter Fox, and Josef Sivic. Megapose: 6d pose estimation of novel objects via render & compare. In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pages 715–725. PMLR, 2022. 2
- [68] Junnan Li, Dongxu Li, Caiming Xiong, and Steven C. H. Hoi. BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In *Int. Conf. Mach. Learn. (ICML)*, 2022. 2
- [69] Junnan Li, Dongxu Li, Silvio Savarese, and Steven C. H. Hoi. BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models. *CoRR*, abs/2301.12597, 2023. 24
- [70] Li-Jia Li and Li Fei-Fei. What, where and who? classifying events by scene and object recognition. In *IEEE 11th International Conference on Computer Vision, ICCV 2007, Rio de Janeiro, Brazil, October 14-20, 2007*, pages 1–8. IEEE Computer Society, 2007. 2
- [71] Xiaoqi Li, Mingxu Zhang, Yiran Geng, Haoran Geng, Yuxing Long, Yan Shen, Renrui Zhang, Jiaming Liu, and Hao Dong. Manipllm: Embodied multimodal large language model for object-centric robotic manipulation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 18061–18070. IEEE, 2024. 2, 11, 23, 24
- [72] Xinghang Li, Peiyan Li, Minghuan Liu, Dong Wang, Jirong Liu, Bingyi Kang, Xiao Ma, Tao Kong, Hanbo Zhang, and Huaping Liu. Towards generalist robot policies: What matters in building vision-language-action models. *arXiv preprint arXiv:2412.14058*, 2024. 8, 9
- [73] Xinghang Li, Minghuan Liu, Hanbo Zhang, Cunjun Yu, Jie Xu, Hongtao Wu, Chilam Cheang, Ya Jing, Weinan Zhang, Huaping Liu, Hang Li, and Tao Kong. Vision-language foundation models as effective robot imitators.

- In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. 11
- [74] Xuanlin Li, Kyle Hsu, Jiayuan Gu, Karl Pertsch, Oier Mees, Homer Rich Walke, Chuyuan Fu, Ishikaa Lunawat, Isabel Sieh, Sean Kirmani, et al. Evaluating real-world robot manipulation policies in simulation. *arXiv preprint arXiv:2405.05941*, 2024. 8, 9, 23, 30
- [75] 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 2023, London, UK, May 29 - June 2, 2023*, pages 9493–9500. IEEE, 2023. 11
- [76] Bo Liu, Yifeng Zhu, Chongkai Gao, Yihao Feng, qiang liu, Yuke Zhu, and Peter Stone. LIBERO: Benchmarking knowledge transfer for lifelong robot learning. In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023. 7, 23, 30
- [77] Dingning Liu, Xiaomeng Dong, Renrui Zhang, Xu Luo, Peng Gao, Xiaoshui Huang, Yongshun Gong, and Zhihui Wang. 3daxiesprompts: Unleashing the 3d spatial task capabilities of GPT-4V. *CoRR*, abs/2312.09738, 2023. 11
- [78] Hao Liu, Lisa Lee, Kimin Lee, and Pieter Abbeel. Instruction-following agents with jointly pre-trained vision-language models. *CoRR*, abs/2210.13431, 2022. 11
- [79] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2023. 2, 11, 24, 27
- [80] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2024. 9, 11
- [81] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In *European Conference on Computer Vision*, pages 38–55. Springer, 2025. 25, 26, 27
- [82] Yunze Liu, Yun Liu, Che Jiang, Kangbo Lyu, Weikang Wan, Hao Shen, Boqiang Liang, Zhoujie Fu, He Wang, and Li Yi. HOI4D: A 4d egocentric dataset for category-level human-object interaction. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022. 30
- [83] Yunze Liu, Changxi Chen, and Li Yi. Interactive humanoid: Online full-body motion reaction synthesis with social affordance canonicalization and forecasting. *ArXiv*, 2023. 30
- [84] Yunze Liu, Junyu Chen, Zekai Zhang, Jingwei Huang, and Li Yi. Leaf: Learning frames for 4d point cloud sequence understanding. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 604–613, 2023. 30
- [85] Yunze Liu, Changxi Chen, Zifan Wang, and Li Yi. Crossvideo: Self-supervised cross-modal contrastive learning for point cloud video understanding. *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 12436–12442, 2024. URL <https://api.semanticscholar.org/CorpusID:267027937>. 30
- [86] Ze Liu, Zheng Zhang, Yue Cao, Han Hu, and Xin Tong. Group-free 3d object detection via transformers. In *Int. Conf. Comput. Vis. (ICCV)*, 2021. 29
- [87] Shiyang Lu, Haonan Chang, Eric Pu Jing, Abdeslam Boularias, and Kostas E. Bekris. OVIR-3D: open-vocabulary 3d instance retrieval without training on 3d data. In *Annu. Conf. Robot. Learn. (CoRL)*, 2023. 29
- [88] Chenyang Ma, Kai Lu, Ta-Ying Cheng, Niki Trigoni, and Andrew Markham. SpatialPIN: Enhancing spatial reasoning capabilities of vision-language models through prompting and interacting 3d priors. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 11
- [89] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel transformer for 3d object detection. In *Int. Conf. Comput. Vis. (ICCV)*, 2021. 29
- [90] Weixin Mao, Weiheng Zhong, Zhou Jiang, Dong Fang, Zhongyue Zhang, Zihan Lan, Fan Jia, Tiancai Wang, Haoqiang Fan, and Osamu Yoshie. Robomatrix: A skill-centric hierarchical framework for scalable robot task planning and execution in open-world. *arXiv preprint arXiv:2412.00171*, 2024. 11
- [91] Daniel Maturana and Sebastian A. Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *IEEE/RSJ Int. Conf. Intell. Robot. and Syst. (IROS)*, pages 922–928. IEEE, 2015. 27
- [92] Atharva Mete, Haotian Xue, Albert Wilcox, Yongxin Chen, and Animesh Garg. Quest: Self-supervised skill abstractions for learning continuous control. *CoRR*, abs/2407.15840, 2024. 11
- [93] Kaichun Mo, Leonidas J Guibas, Mustafa Mukadam, Abhinav Gupta, and Shubham Tulsiani. Where2act: From pixels to actions for articulated 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6813–6823, 2021. 24
- [94] 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 *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. 11
- [95] NVIDIA. Nvidia isaac sim. <https://developer.nvidia.com/isaac-sim>, 2021. 7
- [96] Abby O’Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee,

- Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, Albert Tung, Alex Bewley, Alexander Herzog, Alex Irpan, Alexander Khazatsky, Anant Rai, Anchit Gupta, Andrew Wang, 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 Büchler, Dinesh Jayaraman, Dmitry Kalashnikov, Dorsa Sadigh, Edward Johns, Ethan Paul Foster, Fangchen Liu, Federico Ceola, Fei Xia, Feiyu Zhao, Freek Stulp, Gaoyue Zhou, Gaurav S. Sukhatme, Gautam Salhotra, Ge Yan, Gilbert Feng, Giulio Schiavi, Glen Berseth, Gregory Kahn, Guanzhi Wang, Hao Su, Haoshu Fang, Haochen Shi, Henghui Bao, Heni Ben Amor, Henrik I. Christensen, Hiroki Furuta, 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, 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 Booyer, Jonathan Tompson, Jonathan Yang, Jordi Salvador, Joseph J. Lim, Junhyek Han, Kaiyuan Wang, Kanishka 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 Sünderhauf, Ning Liu, Norman Di Palo, Nur Muhammad (Mahi) Shafiullah, Oier Mees, Oliver Kroemer, Osbert Bastani, Pannag R. Sanketi, Patrick Tree Miller, Patrick Yin, Paul Wohlhart, Peng Xu, Peter David Fagan, Peter Mitrano, Pierre Sermanet, Pieter Abbeel, Priya Sundaesan, Qiuyu Chen, Quan Vuong, Rafael Rafailov, Ran Tian, Ria Doshi, Roberto Martín-Martín, Rohan Baijal, 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 D. Sonawani, Shuran Song, Sichun Xu, Siddhant Haldar, Siddharth Karamcheti, Simeon Adebola, Simon Guist, Soroush Nasiriany, Stefan Schaal, Stefan Welker, Stephen Tian, Subramanian Ramamoorthy, Sudeep Dasari, Suneel Belkhale, 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, Vincent Vanhoucke, Wei Zhan, Wenxuan Zhou, Wolfram Burgard, Xi Chen, Xiaolong Wang, Xinghao Zhu, Xinyang Geng, Xiyuan Liu, Liangwei Xu, Xuanlin Li, Yao Lu, Yecheng Jason Ma, Yejin Kim, Yevgen Chebotar, Yifan Zhou, Yifeng Zhu, Yilin Wu, Ying Xu, Yixuan Wang, Yonatan Bisk, 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, and Zipeng Lin. Open x-embodiment: Robotic learning datasets and RT-X models : Open x-embodiment collaboration. In *IEEE International Conference on Robotics and Automation, ICRA 2024, Yokohama, Japan, May 13-17, 2024*, pages 6892–6903. IEEE, 2024. 8, 11
- [97] OpenAI. Introducing chatgpt. 2022. URL <https://openai.com/blog/chatgpt>. 27
- [98] OpenAI. Gpt-4v(ision) system card, 2023. URL <https://openai.com/research/gpt-4v-system-card>. 7, 9, 24
- [99] OpenAI. GPT-4 technical report. *CoRR*, abs/2303.08774, 2023. URL <https://openai.com/research/gpt-4>. 9
- [100] OpenAI. Introducing gpt-4o and more tools to chatgpt free users. 2024. URL <https://openai.com/index/gpt-4o-and-more-tools-to-chatgpt-free/>. 2, 4, 9, 10, 24, 26
- [101] Yatian Pang, Wenxiao Wang, Francis E. H. Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Eur. Conf. Comput. Vis. (ECCV)*, 2022. 29
- [102] Songyou Peng, Kyle Genova, Chiyu "Max" Jiang, Andrea Tagliasacchi, Marc Pollefeys, and Thomas A. Funkhouser. Openscene: 3d scene understanding with open vocabularies. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2023. 29
- [103] Yuang Peng, Yuxin Cui, Haomiao Tang, Zekun Qi, Runpei Dong, Jing Bai, Chunrui Han, Zheng Ge, Xiangyu Zhang, and Shu-Tao Xia. Dreambench++: A human-aligned benchmark for personalized image generation. *CoRR*, abs/2406.16855, 2024. 4
- [104] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 77–85, 2017. 4, 26, 27
- [105] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning

- on point sets in a metric space. In *Adv. Neural Inform. Process. Syst. (NIPS)*, pages 5099–5108, 2017. 4, 27
- [106] William Qi, Ravi Teja Mullanpudi, Saurabh Gupta, and Deva Ramanan. Learning to move with affordance maps. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. 29
- [107] Zekun Qi, Runpei Dong, Guofan Fan, Zheng Ge, Xiangyu Zhang, Kaisheng Ma, and Li Yi. Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. In *Int. Conf. Mach. Learn. (ICML)*, 2023. 2, 4, 27, 29, 30
- [108] Zekun Qi, Muzhou Yu, Runpei Dong, and Kaisheng Ma. VPP: efficient conditional 3d generation via voxel-point progressive representation. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2023. 27
- [109] Zekun Qi, Runpei Dong, Shaochen Zhang, Haoran Geng, Chunrui Han, Zheng Ge, Li Yi, and Kaisheng Ma. Shapellm: Universal 3d object understanding for embodied interaction. In *Computer Vision - ECCV 2024 - 18th European Conference, Milan, Italy, September 29-October 4, 2024, Proceedings, Part XLIII*, volume 15101 of *Lecture Notes in Computer Science*, pages 214–238. Springer, 2024. 2, 11, 27
- [110] Zhangyang Qi, Ye Fang, Zeyi Sun, Xiaoyang Wu, Tong Wu, Jiaqi Wang, Dahua Lin, and Hengshuang Zhao. Gpt4point: A unified framework for point-language understanding and generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26417–26427, 2024. 11
- [111] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Abed Al Kader Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2022. 27
- [112] Delin Qu, Haoming Song, Qizhi Chen, Yuanqi Yao, Xinyi Ye, Yan Ding, Zhigang Wang, JiaYuan Gu, Bin Zhao, Dong Wang, et al. Spatialvla: Exploring spatial representations for visual-language-action model. *arXiv preprint arXiv:2501.15830*, 2025. 8, 9
- [113] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. 2018. 11
- [114] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9, 2019. 11
- [115] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *Int. Conf. Mach. Learn. (ICML)*, volume 139 of *Proceedings of Machine Learning Research*, pages 8748–8763. PMLR, 2021. 4, 29, 30
- [116] Kenneth Shaw, Ananye Agarwal, and Deepak Pathak. Leap hand: Low-cost, efficient, and anthropomorphic hand for robot learning. *arXiv preprint arXiv:2309.06440*, 2023. 10
- [117] Mohit Shridhar, Lucas Manuelli, and Dieter Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. In *Conference on Robot Learning, CoRL 2022, 14-18 December 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pages 785–799. PMLR, 2022. 11
- [118] Tianmin Shu, Michael S. Ryoo, and Song-Chun Zhu. Learning social affordance for human-robot interaction. In *Int. Joint Conf. Artif. Intell. (IJCAI)*, 2016. 29
- [119] Austin Stone, Ted Xiao, Yao Lu, Keerthana Gopalakrishnan, Kuang-Huei Lee, Quan Vuong, Paul Wohlhart, Sean Kirmani, Brianna Zitkovich, Fei Xia, Chelsea Finn, and Karol Hausman. Open-world object manipulation using pre-trained vision-language models. In *Conference on Robot Learning, CoRL 2023, 6-9 November 2023, Atlanta, GA, USA*, volume 229 of *Proceedings of Machine Learning Research*, pages 3397–3417. PMLR, 2023. 2
- [120] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik G. Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Int. Conf. Comput. Vis. (ICCV)*, 2015. 27
- [121] Ioan A. Şucan, Mark Moll, and Lydia E. Kavraki. The Open Motion Planning Library. *IEEE Robotics & Automation Magazine*, 19(4):72–82, December 2012. doi: 10.1109/MRA.2012.2205651. <https://ompl.kavrakilab.org>. 7
- [122] Jiaming Sun, Zihao Wang, Siyu Zhang, Xingyi He, Hongcheng Zhao, Guofeng Zhang, and Xiaowei Zhou. Onepose: One-shot object pose estimation without CAD models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pages 6815–6824. IEEE, 2022. 2
- [123] Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yueze Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Emu: Generative pretraining in multimodality. In *Int. Conf. Learn. Represent. (ICLR)*, 2024. 11
- [124] Emilia Szymanska, Mihai Dusmanu, Jan-Willem Burchard, Mahdi Rad, and Marc Pollefeys. Space3d-bench: Spatial 3d question answering benchmark. *arXiv preprint arXiv:2408.16662*, 2024. 8
- [125] Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 9, 10, 24, 26, 27
- [126] Octo Model Team, Dibya Ghosh, Homer Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Lawrence Yunliang Chen, Pannag San-

- keti, Quan Vuong, Ted Xiao, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. *CoRR*, abs/2405.12213, 2024. 7, 8, 9
- [127] Stefanie Tellex, Thomas Kollar, Steven Dickerson, Matthew R. Walter, Ashis Gopal Banerjee, Seth J. Teller, and Nicholas Roy. Understanding natural language commands for robotic navigation and mobile manipulation. In *AAAI Conf. Artif. Intell. (AAAI)*, 2011. 2
- [128] Stefanie Tellex, Nakul Gopalan, Hadas Kress-Gazit, and Cynthia Matuszek. Robots that use language. *Annu. Rev. Control. Robotics Auton. Syst.*, 3:25–55, 2020. 29
- [129] Shengbang Tong, Ellis Brown, Penghao Wu, Sanghyun Woo, Manoj Middepogu, Sai Charitha Akula, Jihan Yang, Shusheng Yang, Adithya Iyer, Xichen Pan, Austin Wang, Rob Fergus, Yann LeCun, and Saining Xie. Cambrian-1: A fully open, vision-centric exploration of multimodal llms. *CoRR*, abs/2406.16860, 2024. 11
- [130] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. Llama: Open and efficient foundation language models. *CoRR*, abs/2302.13971, 2023. 11, 27
- [131] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023. 11, 27
- [132] Shinji Umeyama. Least-squares estimation of transformation parameters between two point patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, 13(04):376–380, 1991. 7
- [133] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, pages 1588–1597, 2019. 27
- [134] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Adv. Neural Inform. Process. Syst. (NIPS)*, pages 5998–6008, 2017. 2, 4, 29
- [135] Johanna Wald, Helisa Dharmo, Nassir Navab, and Federico Tombari. Learning 3d semantic scene graphs from 3d indoor reconstructions. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 3960–3969. Computer Vision Foundation / IEEE, 2020. 2
- [136] Homer Rich Walke, Kevin Black, Tony Z Zhao, Quan Vuong, Chongyi Zheng, Philippe Hansen-Estruch, Andre Wang He, Vivek Myers, Moo Jin Kim, Max Du, et al. Bridgedata v2: A dataset for robot learning at scale. In *Conference on Robot Learning*, pages 1723–1736. PMLR, 2023. 8
- [137] Chenxi Wang, Haoshu Fang, Minghao Gou, Hongjie Fang, Jin Gao, and Cewu Lu. Graspness discovery in clutters for fast and accurate grasp detection. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 15944–15953. IEEE, 2021. 10
- [138] Ruicheng Wang, Sicheng Xu, Cassie Dai, Jianfeng Xiang, Yu Deng, Xin Tong, and Jiaolong Yang. Moge: Unlocking accurate monocular geometry estimation for open-domain images with optimal training supervision. *arXiv preprint arXiv:2410.19115*, 2024. 9
- [139] Zehan Wang, Ziang Zhang, Tianyu Pang, Chao Du, Hengshuang Zhao, and Zhou Zhao. Orient anything: Learning robust object orientation estimation from rendering 3d models. *arXiv preprint arXiv:2412.18605*, 2024. 22, 30
- [140] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In *Adv. Neural Inform. Process. Syst. (NeurIPS)*, 2022. 2, 7, 31
- [141] Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. Foundationpose: Unified 6d pose estimation and tracking of novel objects. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2024, Seattle, WA, USA, June 16-22, 2024*, pages 17868–17879. IEEE, 2024. 2
- [142] Jiajun Wu. Physical scene understanding. *AI Mag.*, 45 (1):156–164, 2024. 2
- [143] Tong Wu, Guandao Yang, Zhibing Li, Kai Zhang, Ziwei Liu, Leonidas J. Guibas, Dahua Lin, and Gordon Wetzstein. Gpt-4v(ision) is a human-aligned evaluator for text-to-3d generation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2024. 4
- [144] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*,

- pages 1912–1920, 2015. [27](#)
- [145] Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN: A simulated part-based interactive environment. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2020. [23](#), [24](#)
- [146] Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. In *Robotics: Science and Systems XIV, Carnegie Mellon University, Pittsburgh, Pennsylvania, USA, June 26-30, 2018*, 2018. [2](#)
- [147] Bin Xiao, Haiping Wu, Weijian Xu, Xiyang Dai, Houdong Hu, Yumao Lu, Michael Zeng, Ce Liu, and Lu Yuan. Florence-2: Advancing a unified representation for a variety of vision tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4818–4829, 2024. [2](#), [5](#), [7](#), [24](#), [25](#), [26](#), [27](#)
- [148] Saining Xie, Jiatao Gu, Demi Guo, Charles R. Qi, Leonidas J. Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *Eur. Conf. Comput. Vis. (ECCV)*, volume 12348 of *Lecture Notes in Computer Science*, pages 574–591. Springer, 2020. [29](#)
- [149] Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. *CoRR*, abs/2308.16911, 2023. [11](#)
- [150] Zhenjia Xu, Zhanpeng He, and Shuran Song. Universal manipulation policy network for articulated objects. *IEEE robotics and automation letters*, 7(2):2447–2454, 2022. [24](#)
- [151] Jianwei Yang, Hao Zhang, Feng Li, Xueyan Zou, Chunyuan Li, and Jianfeng Gao. Set-of-mark prompting unleashes extraordinary visual grounding in GPT-4V. *CoRR*, abs/2310.11441, 2023. [5](#), [11](#), [27](#)
- [152] Jihan Yang, Shusheng Yang, Anjali W Gupta, Rilyn Han, Li Fei-Fei, and Saining Xie. Thinking in space: How multimodal large language models see, remember, and recall spaces. *arXiv preprint arXiv:2412.14171*, 2024. [30](#)
- [153] Sherry Yang, Yilun Du, Seyed Kamyar Seyed Ghasemipour, Jonathan Tompson, Leslie Pack Kaelbling, Dale Schuurmans, and Pieter Abbeel. Learning interactive real-world simulators. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*. OpenReview.net, 2024. [11](#)
- [154] Sriram Yenamandra, Arun Ramachandran, Karmesh Yadav, Austin Wang, Mukul Khanna, Theophile Gervet, Tsung-Yen Yang, Vidhi Jain, Alexander William Clegg, John Turner, et al. Homerobot: Open-vocabulary mobile manipulation. *arXiv preprint arXiv:2306.11565*, 2023. [30](#)
- [155] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM Trans. Graph.*, 35(6):1–12, 2016. [27](#)
- [156] Li Yi, Hao Su, Xingwen Guo, and Leonidas J. Guibas. Syncspecnn: Synchronized spectral CNN for 3d shape segmentation. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2017. [27](#)
- [157] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from A single image. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, pages 9009–9019. IEEE, 2023. [5](#), [9](#)
- [158] Tsuneo Yoshikawa. Manipulability of robotic mechanisms. *The international journal of Robotics Research*, 4(2):3–9, 1985. [2](#)
- [159] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *IEEE/CVF Conf. Comput. Vis. Pattern Recog. (CVPR)*, 2022. [4](#), [29](#)
- [160] Wentao Yuan, Jiafei Duan, Valts Blukis, Wilbert Pumacay, Ranjay Krishna, Adithyavairavan Murali, Arsalan Mousavian, and Dieter Fox. Robopoint: A vision-language model for spatial affordance prediction for robotics. *CoRR*, abs/2406.10721, 2024. [2](#), [9](#), [11](#)
- [161] Zhecheng Yuan, Tianming Wei, Shuiqi Cheng, Gu Zhang, Yuanpei Chen, and Huazhe Xu. Learning to manipulate anywhere: A visual generalizable framework for reinforcement learning. In *8th Annual Conference on Robot Learning*, 2024. [11](#)
- [162] Jiazhao Zhang, Kunyu Wang, Shaoan Wang, Minghan Li, Haoran Liu, Songlin Wei, Zhongyuan Wang, Zhizheng Zhang, and He Wang. Uni-navid: A video-based vision-language-action model for unifying embodied navigation tasks. *arXiv preprint arXiv:2412.06224*, 2024. [23](#)
- [163] Junbo Zhang, Runpei Dong, and Kaisheng Ma. CLIP-FO3D: learning free open-world 3d scene representations from 2d dense CLIP. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023 - Workshops, Paris, France, October 2-6, 2023*, pages 2040–2051. IEEE, 2023. [29](#)
- [164] Linfeng Zhang, Xin Chen, Runpei Dong, and Kaisheng Ma. Region-aware knowledge distillation for efficient image-to-image translation. In *Brit. Mach. Vis. Conf. (BMVC)*, 2023. [29](#)
- [165] Renrui Zhang, Jiaming Han, Aojun Zhou, Xiangfei Hu, Shilin Yan, Pan Lu, Hongsheng Li, Peng Gao, and Yu Qiao. Llama-adapter: Efficient fine-tuning of language models with zero-init attention. In *Int. Conf. Learn. Represent. (ICLR)*, 2024. [11](#)
- [166] Shaochen Zhang, Zekun Qi, Runpei Dong, Xiuxiu Bai, and Xing Wei. Positional prompt tuning for efficient 3d representation learning. *arXiv preprint*

- arXiv:2408.11567*, 2024. [29](#)
- [167] Liang Zhao, En Yu, Zheng Ge, Jinrong Yang, Haoran Wei, Hongyu Zhou, Jianjian Sun, Yuang Peng, Runpei Dong, Chunrui Han, and Xiangyu Zhang. Chatspot: Bootstrapping multimodal llms via precise referring instruction tuning. In *Int. Joint Conf. Artif. Intell. (IJCAI)*, 2024. [11](#)
- [168] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In *Robotics: Science and Systems XIX, Daegu, Republic of Korea, July 10-14, 2023*, 2023. [11](#)
- [169] Haoyu Zhen, Xiaowen Qiu, Peihao Chen, Jincheng Yang, Xin Yan, Yilun Du, Yining Hong, and Chuang Gan. 3d-vla: A 3d vision-language-action generative world model. In *Forty-first International Conference on Machine Learning, ICML 2024, Vienna, Austria, July 21-27, 2024*. OpenReview.net, 2024. [11](#)
- [170] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Advances in Neural Information Processing Systems 36: Annual Conference on Neural Information Processing Systems 2023, NeurIPS 2023, New Orleans, LA, USA, December 10 - 16, 2023*, 2023. [4](#)
- [171] Chengliang Zhong, Yuhang Zheng, Yupeng Zheng, Hao Zhao, Li Yi, Xiaodong Mu, Ling Wang, Pengfei Li, Guyue Zhou, Chao Yang, et al. 3d implicit transporter for temporally consistent keypoint discovery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3869–3880, 2023. [24](#)
- [172] Deyao Zhu, Jun Chen, Xiaoqian Shen, Xiang Li, and Mohamed Elhoseiny. Minigt-4: Enhancing vision-language understanding with advanced large language models. In *Int. Conf. Learn. Represent. (ICLR)*, 2024. [24](#)
- [173] Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martín-Martín, Abhishek Joshi, Kevin Lin, Abhiram Maddukuri, Soroush Nasiriany, and Yifeng Zhu. robo-suite: A modular simulation framework and benchmark for robot learning. In *arXiv preprint arXiv:2009.12293*, 2020. [7](#)
- [174] 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 T. 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 *Conference on Robot Learning, CoRL 2023, 6-9 November 2023, Atlanta, GA, USA*, volume 229 of *Proceedings of Machine Learning Research*, pages 2165–2183. PMLR, 2023. [8](#), [11](#)

# Contents

<b>I Introduction</b>	<b>2</b>	<b>VII Conclusions</b>	<b>11</b>
<b>II Semantic Orientation: Connecting Language and Object Orientation</b>	<b>3</b>	<b>Appendix</b>	<b>23</b>
II-A Definition of Semantic Orientation . . . . .	3	<b>Appendix A: Robot Setups</b>	<b>23</b>
II-B Robotic Manipulation via Semantic Orientation . . . . .	3	A Simulation Robot Setups . . . . .	23
II-C OrientText300K: Orientation-Text Paired Data at Scale . . . . .	3	B Real World Robot Setups . . . . .	23
II-C1 Scalability Analysis . . . . .	4	<b>Appendix B: Additional Experiments</b>	<b>23</b>
II-C2 Data Source . . . . .	4	A Articulated Objects Manipulation Evaluation . . . . .	23
II-C3 Data Filtering . . . . .	4	B Spatial VQA on EmbSpatial-Bench [33] & SpatialBot-Bench [10] . . . . .	23
II-C4 Data Annotation . . . . .	4	C Close-Loop Execution Experiment . . . . .	23
II-C5 Quality Validation . . . . .	4	D Long Horizon Object Manipulation Experiment . . . . .	24
II-D PointSO: A Cross-Modal 3D Transformer for Semantic Orientation Prediction . . . . .	4	E In the Wild Evaluation of Semantic Orientation . . . . .	24
II-D1 3D and Language Embeddings . . . . .	4	F Cross-View Generalization . . . . .	25
II-D2 Cross-Modal Fusion . . . . .	5	G Failure Case Distribution Analysis . . . . .	25
II-D3 Optimization . . . . .	5	H Ablation Study . . . . .	26
<b>III SoFAR: Semantic Orientation Bridges Spatial Reasoning and Object Manipulation</b>	<b>5</b>	H1 Scaling Law . . . . .	26
III-A Orientation-Aware Scene Graph from RGB-D . . . . .	5	H2 Cross-Modal Fusion Choices . . . . .	26
III-A1 Task-Oriented Object Segmentation . . . . .	5	H3 Open Vocabulary Object Detection Module . . . . .	26
III-A2 Orientation-Aware 3D Scene Graph . . . . .	7	<b>Appendix C: Additional Implementation Details</b>	<b>26</b>
III-B Spatial-Aware Task Reasoning . . . . .	7	A Detail Real World Experiment Results . . . . .	26
III-B1 Chain-of-Thought Spatial Reasoning . . . . .	7	B PointSO Model Details . . . . .	26
III-B2 Low-Level Motion Execution . . . . .	7	C SoFar-LLaVA Model Details . . . . .	27
<b>IV Experiments</b>	<b>7</b>	D ChatGPT API Costs . . . . .	27
IV-A Benchmarks . . . . .	7	<b>Appendix D: Additional Benchmark Statistic Analysis</b>	<b>27</b>
IV-A1 Open6DOR V2 . . . . .	7	A 6-DoF SpatialBench Analysis . . . . .	27
IV-A2 6-DoF SpatialBench . . . . .	7	B Open6DOR V2 Analysis . . . . .	27
IV-B 6-DoF object rearrangement evaluation in Simulation . . . . .	8	<b>Appendix E: Additional Related Works</b>	<b>27</b>
IV-C Zero-shot Object Manipulation Evaluation in Simulation . . . . .	9	A 3D Representation Learning . . . . .	27
IV-D Zero-shot Real-world Manipulation . . . . .	9	<b>Appendix F: Additional Discussions</b>	<b>29</b>
IV-D1 Hardware Setup . . . . .	9	A Relation to Affordance & 6-DoF Pose Estimation . . . . .	29
IV-D2 Tasks and Evaluations . . . . .	10	B Comparison to Concurrent Works . . . . .	29
IV-D3 Results . . . . .	10	B1 Comparison with ReKep [56] . . . . .	29
IV-E Visual Question Answer on 6-DoF SpatialBench . . . . .	10	B2 Comparison with Orient Anything [139] . . . . .	30
IV-F Semantic Orientation Prediction . . . . .	10	C Future Works . . . . .	30
IV-G Additional Applications . . . . .	10	<b>Appendix G: Additional Visualizations</b>	<b>30</b>
IV-G1 Cross Embodiment Generalization . . . . .	10	A Robotic Manipulation . . . . .	30
IV-G2 Long Horizon Planning & Close-Loop Execution . . . . .	10	B 6-DoF SpatialBench . . . . .	30
IV-G3 Orientation-Aware Robotic Navigation . . . . .	10	C System Prompts . . . . .	31
<b>V Related Works</b>	<b>11</b>		
V-A Vision-Language Models for Spatial Understanding . . . . .	11		
V-B Language-Grounded Robot Manipulation . . . . .	11		
<b>VI Limitations</b>	<b>11</b>		

APPENDIX  
APPENDIX A  
ROBOT SETUPS

A. Simulation Robot Setups

To ensure fairness, we utilize the same Franka Panda arm for evaluations in both the LIBERO [76] and our Open6DOR V2 benchmarks. For SIMPLER [74], we use the Google Robot exclusively to conduct the baseline experiments, adhering to all configurations outlined in SIMPLER, as presented in Table II.

B. Real World Robot Setups

As for manipulation tasks, in Fig. 9, we perform 6-DoF rearrangement tasks using the Franka Panda equipped with a gripper and the UR robot arm with a LeapHand, while articulated object manipulation is conducted using the Flexiv arm equipped with a suction tool. All the robot arms mount a Realsense D415 camera to its end for image capturing.

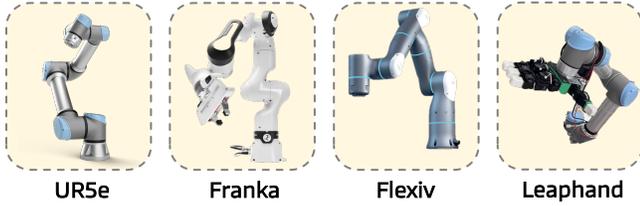


Fig. 9: The robots used in our real-world experiments.

In Fig. 10, we present the workspace and robotic arm for real-world 6-DoF rearrangement. Unlike Rekep [56], CoPa [51] et al., we utilize only a single RealSense D415 camera. This setup significantly reduces the additional overhead associated with environmental setup and multi-camera calibration, and it is more readily reproducible.

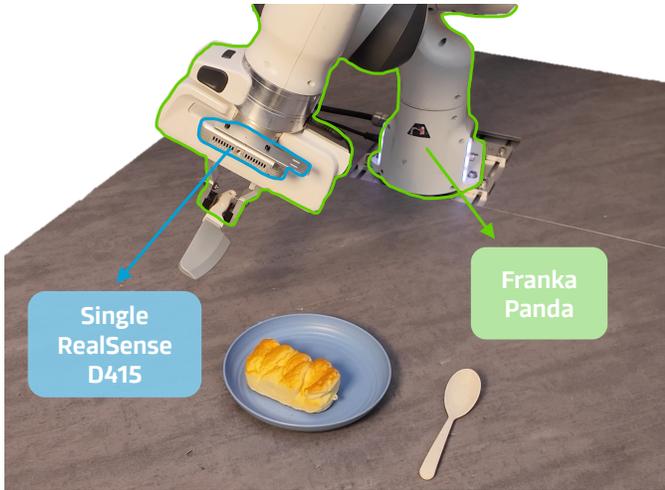


Fig. 10: 6-DoF rearrangement robot setup.

As for navigation tasks, we provide a visualization of our robotic dog in Fig. 11. Following Uni-Navid [162], our robotic dog is Unitree GO2 and we mount a RealSense D455 camera on the head of the robotic dog. Here, we only use the RGB

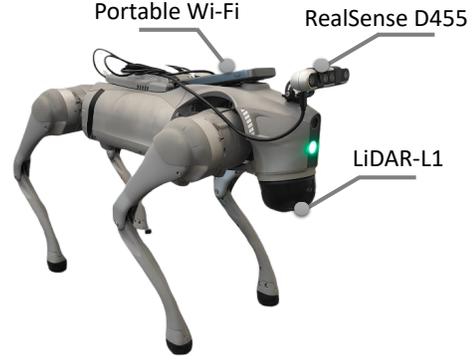


Fig. 11: **Navigation robot setup.** We use Unitree GO2 as our embodiment, and we mount RealSense D455, a portable Wi-Fi and a LiDAR-L1. Note that, our model only takes RGB frames as input. The portable Wi-Fi is used for communication with the remote server and the Lidar is used for the local controller API of Unitree Dog.

frames with a resolution of  $640 \times 480$  in the setting of  $90^\circ$  HFOV. We also mount a portable Wi-Fi at the back of the robot dog, which is used to communicate with the remote server (send captured images and receive commands). Unitree GO2 is integrated with a LiDAR-L1, which is only used for local motion planning.

APPENDIX B  
ADDITIONAL EXPERIMENTS

A. Articulated Objects Manipulation Evaluation

We further integrate SOFAR with articulated object manipulation, as illustrated in Table VII, and evaluate its practicality in robotic manipulation tasks using the PartNet-Mobility Dataset within the SAPIEN [145] simulator. Our experimental setup follows ManipLLM [71], employing the same evaluation metrics. Specifically, we directly utilize the segmentation centers provided by SAM as contact points, leverage PointSO to generate contact directions, and use VLM to determine subsequent motion directions. The results demonstrate significant improvements over the baseline. Notably, our model achieves this performance without dividing the data into training and testing sets, operating instead in a fully zero-shot across most tasks. This underscores the robustness and generalization of our approach.

B. Spatial VQA on EmbSpatial-Bench [33] & SpatialBot-Bench [10]

To further demonstrate SOFAR’s spatial reasoning capabilities, we conducted Spatial VQA tests within the EmbSpatial-Bench [33] and SpatialBot-Bench [10]. As shown in Tables VIII and IX, SOFAR significantly outperformed all baselines, achieving more than a 20% performance improvement in EmbSpatial-Bench.

C. Close-Loop Execution Experiment

We demonstrate the closed-loop replan capabilities of SOFAR within Simpler-Env [74] in Fig. 12. The instruction

TABLE VII: **Zeroshot articulate object manipulation evaluation** within the SAPIEN [145] simulator using PartNet-Mobility Dataset. Notably, while the baseline methods use distinct training and testing splits, our model achieves robust performance without fine-tuning on the SAPIEN samples.

Method																
Where2Act [93]	0.26	0.36	0.19	0.27	0.23	0.11	0.15	0.47	0.14	0.24	0.13	0.12	0.56	0.68	0.07	0.40
UMPNet [150]	0.46	0.43	0.15	0.28	0.54	0.32	0.28	0.56	0.44	0.40	0.10	0.23	0.18	0.54	0.20	0.42
FlowBot3D [35]	0.67	0.55	0.20	0.32	0.27	0.31	0.61	<b>0.68</b>	0.15	0.28	0.36	0.18	0.21	0.70	0.18	0.26
Implicit3D [171]	0.53	0.58	0.35	0.55	0.28	<b>0.66</b>	0.58	0.51	0.52	0.57	0.45	0.34	0.41	0.54	0.39	0.43
ManipLLM [71]	0.68	0.64	0.36	0.77	0.43	0.62	0.65	0.61	0.65	0.52	0.53	<b>0.40</b>	0.64	0.71	<b>0.60</b>	<b>0.64</b>
<b>SoFAR</b>	<b>0.75</b>	<b>0.88</b>	<b>0.43</b>	<b>0.85</b>	<b>0.60</b>	0.54	<b>0.75</b>	0.49	<b>0.58</b>	<b>0.72</b>	<b>0.69</b>	0.42	<b>0.70</b>	<b>0.81</b>	0.58	0.63

Method					AVG											AVG
Where2Act [93]	0.13	0.18	0.13	0.40	0.26	0.18	0.35	0.38	0.28	0.05	0.21	0.17	0.20	0.15	0.15	0.21
UMPNet [150]	0.22	0.33	0.26	0.64	0.35	0.42	0.20	0.35	0.42	0.29	0.20	0.26	0.28	0.25	0.15	0.28
FlowBot3D [35]	0.17	0.53	0.29	0.42	0.37	0.23	0.10	0.60	0.39	0.27	0.42	0.28	0.51	0.13	0.23	0.32
Implicit3D [171]	0.27	0.65	0.20	0.33	0.46	0.45	0.17	0.80	0.53	0.15	0.69	0.41	0.31	0.30	0.31	0.41
ManipLLM [71]	<b>0.41</b>	<b>0.75</b>	0.44	0.67	0.59	0.38	0.22	0.81	<b>0.86</b>	0.38	<b>0.85</b>	0.42	<b>0.83</b>	0.26	0.38	0.54
<b>SoFAR</b>	0.35	0.68	<b>0.62</b>	<b>0.73</b>	<b>0.64</b>	<b>0.68</b>	<b>0.45</b>	<b>0.90</b>	0.77	<b>0.55</b>	0.79	<b>0.48</b>	0.80	<b>0.56</b>	<b>0.44</b>	<b>0.64</b>

TABLE VIII: Zero-shot performance of LVLMs in EmbSpatialBench [33]. **Bold** indicates the best results.

Model	Generation	Likelihood
BLIP-2 [69]	37.99	35.71
InstructBLIP [20]	38.85	33.41
MiniGPT4 [172]	23.54	31.70
LLaVA-1.6 [79]	35.19	38.84
GPT-4V [98]	36.07	-
Qwen-VL-Max [3]	49.11	-
<b>SoFAR</b>	<b>70.88</b>	-

TABLE IX: **Zero-shot performance of LVLMs in SpatialBotBench [10]**. SpatialBot-3B: SpatialBot-Phi2-3B-RGB, SpatialBot-8B: SpatialBot-Llama3-8B-RGB.

Model	Pos	Exist	Count	Reach	Size	Avg
ChatGPT-4o [100]	70.6	85.0	84.5	51.7	<b>43.3</b>	67.0
SpatialBot-3B [10]	64.7	80.0	88.0	<b>61.7</b>	28.3	64.5
SpatialBot-8B [10]	55.9	80.0	<b>91.2</b>	40.0	20.0	57.4
<b>SoFAR</b>	<b>76.5</b>	<b>87.5</b>	80.0	57.5	40.0	<b>68.3</b>

for both tasks is “pick the coke can” In Fig. 12 (a), the model initially misidentified the coke can as a Fanta can. After correction by the VLM, the model re-identified and located the correct object. In Fig. 12 (b), the model accidentally knocks over the Coke can during motion due to erroneous motion planning. Subsequently, the model re-plans and successfully achieves the grasp.

#### D. Long Horizon Object Manipulation Experiment

Fig. 13 illustrates the execution performance of our model on long-horizon tasks. Through the VLM [100, 125], complex instructions such as “making breakfast” and “cleaning up the desktop” can be decomposed into sub-tasks. In the second example, we deliberately chose complex and uncommon objects as assets, such as “Aladdin’s lamp” and “puppets”, but SOFAR is able to successfully complete all tasks.

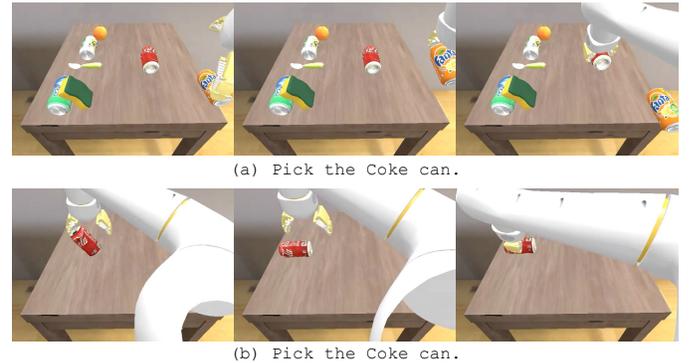


Fig. 12: **Close-loop execution of our SOFAR.**

#### E. In the Wild Evaluation of Semantic Orientation

We provide a qualitative demonstration of the accuracy of PointSO under in-the-wild conditions, as shown in Fig. 14, where the predicted Semantic Orientation is marked in the images. We obtained single-sided point clouds by segmenting objects using Florence-2 [147] and SAM [65] and fed them into PointSO. It can be observed that our model achieves good performance across different views, objects, and instructions, which proves the effectiveness and generalization of PointSO.

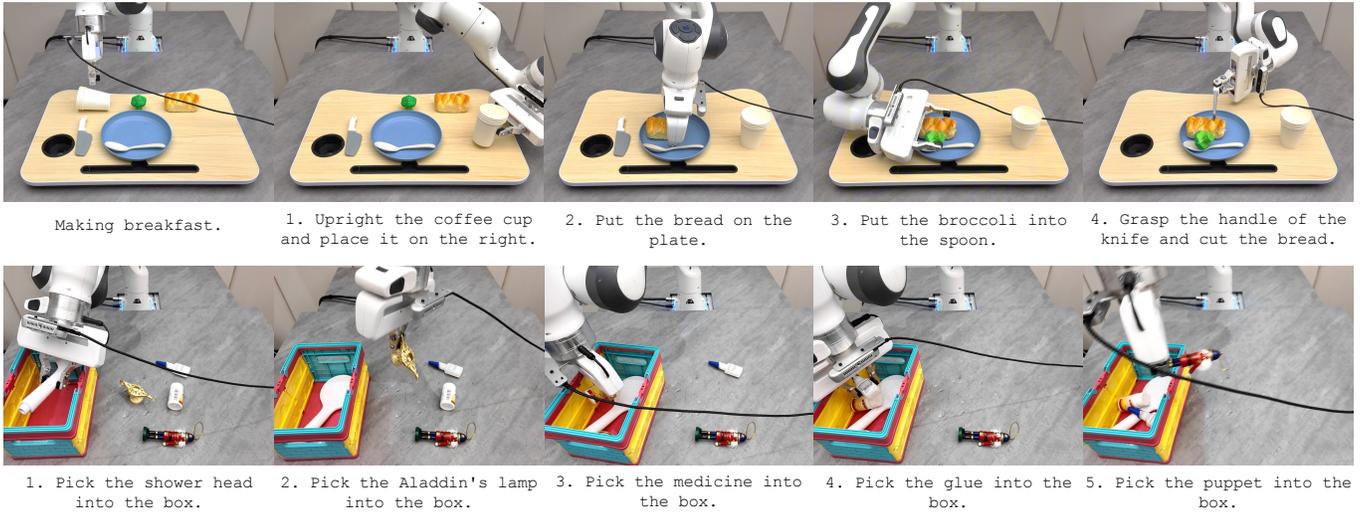


Fig. 13: Long-horizon object manipulation experiment of our SOFAR.

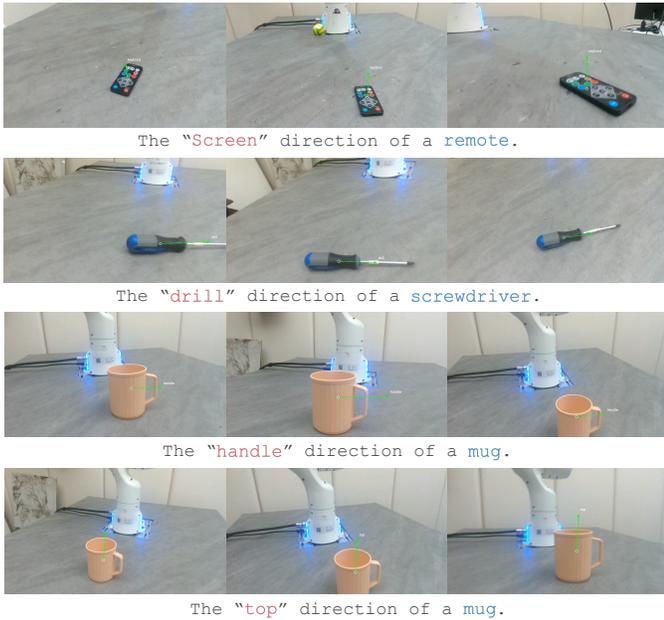


Fig. 14: In the wild evaluation of PointSO.

#### F. Cross-View Generalization

SOFAR gets point clouds in the world coordinate system using an RGB-D camera to obtain grasping poses, and it is not limited to a fixed camera perspective. In addition, PointSO generates partial point clouds from different perspectives through random camera views to serve as data augmentation for training data, which also generalizes to camera perspectives in the real world. Fig. 15 illustrates SOFAR’s generalization capability for 6-DoF object manipulation across different camera poses. It can be observed that whether it’s a front view, side view, or ego view, SOFAR can successfully execute the “upright the bottle” instruction.

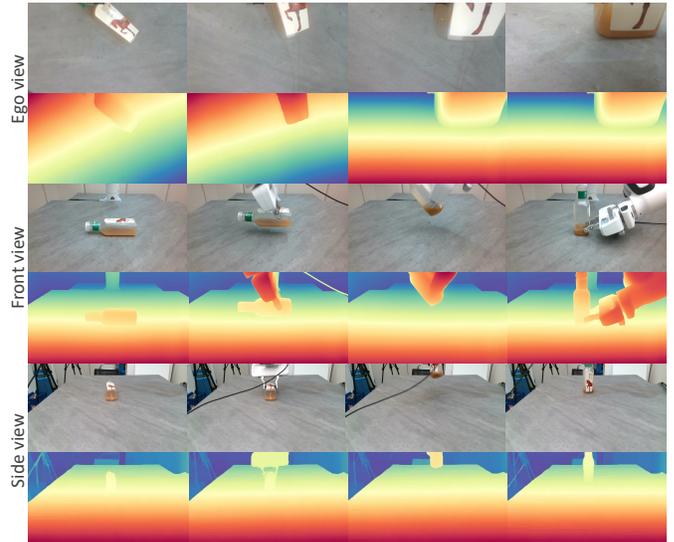


Fig. 15: Cross view generalization of our SOFAR.

#### G. Failure Case Distribution Analysis

Based on the failure cases from real-world experiments, we conducted a quantitative analysis of the failure case distribution for SOFAR, with the results shown in Fig. 16. It can be observed that 31% of the failures originated from grasping issues, including objects being too small, inability to generate reasonable grasping poses, and instability after grasping leading to sliding or dropping. Next, 23% were due to incorrect or inaccurate Semantic Orientation prediction. For tasks such as upright or upside - down, highly precise angle estimation ( $\pm 5^\circ$ ) is required for smooth execution. Object analysis and detection accounted for approximately 20% of the errors. The instability of open-vocabulary detection modules like Florence2 [147] and Grounding DINO [81] often led to incorrect detection of out-of-distribution objects or object parts. In addition, since our Motion

TABLE X: **Data scaling property** of semantic orientation with different training data scales evaluated on OrientText300K test split. All experiments are conducted with PointSO-Base.

Data Scale	45°	30°	15°	5°	Average
5%	57.03	46.09	39.84	27.34	42.58
10%	61.72	53.13	43.75	30.47	47.27
50%	76.56	72.66	66.41	56.25	67.97
100%	<b>79.69</b>	<b>77.34</b>	<b>70.31</b>	<b>62.50</b>	<b>72.46</b>

Planning did not take into account the working space range of the robotic arm and potential collisions of the manipulated object, occasional deadlocks and collisions occurred during motion. Finally, there were issues with the Task Planning of the VLM [100, 125]. For some complex Orientations, the VLM occasionally failed to infer the required angles and directions to complete the task. Employing a more powerful, thought-enabled VLM [58] might alleviate such errors.

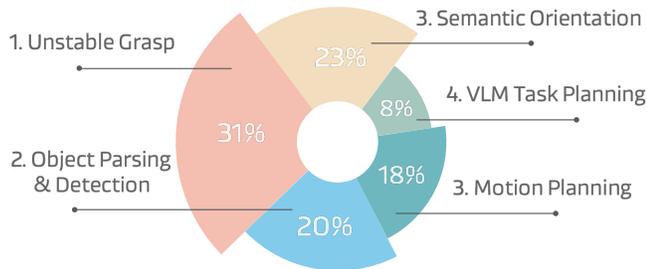


Fig. 16: **Failure case distribution analysis of our SOFAR.**

#### H. Ablation Study

1) *Scaling Law*: The scaling capability of models and data is one of the most critical attributes today and a core feature of foundation models [6]. We investigate the performance of PointSO across different data scales, as illustrated in Table X. We obtain the subset for OrientText300K from Objaverse-LVIS, which consists of approximately 46,000 3D objects with category annotations. The selection was based on the seven criteria mentioned in the main text. Objects meeting all seven criteria formed the strict subset, comprising around 15k objects. When including objects without textures and those of lower quality, the total increases to approximately 26k objects. It can be seen that the increase in data volume is the most significant factor driving the performance improvement of PointSO. It can be anticipated that with further data expansion, such as Objaverse-XL [21], PointSO will achieve better performance.

2) *Cross-Modal Fusion Choices*: We further conduct an ablation study on the multi-modal fusion methods in PointSO, testing commonly used feature fusion techniques such as cross-attention, multiplication, addition, and concatenation, as shown in Table XI. The results indicate that simple addition achieves the best performance. This may be attributed to the fact that instructions in the semantic domain are typically composed of short phrases or sentences, and the text CLS token already encodes sufficiently high-level semantic information.

TABLE XI: **Ablation of multi-modal fusion** in PointSO. All the experiments are under the PointSO-Base variant.

Data Scale	45°	30°	15°	5°	Average
Cross-attention	74.22	70.31	63.28	57.03	66.21
Multiplication	74.22	69.53	60.16	56.25	65.04
Addition	<b>79.69</b>	<b>77.34</b>	<b>70.31</b>	<b>62.50</b>	<b>72.46</b>
Concat	66.41	60.94	52.34	43.75	55.86

3) *Open Vocabulary Object Detection Module*: SOFAR utilize a detection foundation model to localize the interacted objects or parts, then generate masks with SAM [65]. Although not the SOTA performance on the COCO benchmark, Florence-2 [147] exhibits remarkable generalization in in-the-wild detection tasks, even in simulator scenarios. Table XII illustrates the performance of various detection modules in Open6DOR [28] Perception, where Florence-2 achieves the best results and outperforms Grounding DINO [81] and YOLO-World [16].

## APPENDIX C

### ADDITIONAL IMPLEMENTATION DETAILS

#### A. Detail Real World Experiment Results

To fully demonstrate the generalization of SOFAR rather than cherry-picking, we carefully design 60 different real-world experimental tasks, covering more than 100 different and diverse objects. Similar to the Open6DOR [28] benchmark in the simulator, we divide these 60 tasks into three parts: position-track, orientation-track, and the most challenging comprehensive & 6-DoF-track. Each track is further divided into simple and hard levels. The position-simple track includes tasks related to front & back & left & right spatial relationships, while the position-hard track includes tasks related to between, center, and customized. The orientation-simple track includes tasks related to the orientation of object parts, and the orientation-hard track includes tasks related to whether the object is upright or flipped (with very strict requirements for angles in both upright and flipped cases). Comprehensive tasks involve complex instruction understanding and long-horizon tasks; 6-DoF tasks simultaneously include requirements for both object position and orientation instructions. In Table XIII, we present the complete task instructions, as well as the performance metrics of SOFAR and the baseline. Due to the large number of tasks, we performed each task three times. It can be seen that SOFAR achieves the best performance in all tracks, especially in the orientation-track and comprehensive & 6-DoF-track. We also show all the objects used in the real-world experiments in Fig. 17, covering a wide range of commonly and uncommonly used objects in daily life.

#### B. PointSO Model Details

For PointSO, we utilize FPS + KNN to perform patchify and employ a small PointNet [104] as the patch encoder. Subsequently, a standard Transformer encoder is adopted as the backbone, followed by a single linear layer to map the output to a three-dimensional vector space. All parameter

TABLE XII: Ablation study of open vocabulary detection modules on Open6DOR [28] perception tasks.

Method	Position Track			Rotation Track				6-DoF Track			Time Cost (s)
	Level 0	Level 1	Overall	Level 0	Level 1	Level 2	Overall	Position	Rotation	Overall	
YOLO-World [16]	59.0	37.7	53.3	48.3	36.1	62.0	44.9	53.4	44.6	27.8	<b>7.4s</b>
Grounding DINO [81]	92.2	71.5	86.7	64.7	41.1	69.8	55.5	87.2	51.6	44.6	9.2s
Florence-2 [147]	<b>96.0</b>	<b>81.5</b>	<b>93.0</b>	<b>68.6</b>	<b>42.2</b>	<b>70.1</b>	<b>57.0</b>	<b>92.7</b>	<b>52.7</b>	<b>48.7</b>	<b>8.5s</b>



Fig. 17: The real-world assets used in our real-world experiments. More than 100 diverse objects are used in our 6-DoF rearrangement experiments.

configurations follow prior work on point cloud representation learning [29, 107, 109]. Detailed hyperparameter and model configurations are provided in Tables XIV and XV.

### C. SoFar-LLaVA Model Details

In addition to leveraging the extensive knowledge and strong generalization capabilities of closed-source/open-source pretrained VLMs [3, 97, 125] for zero-shot or in-context learning, SOFAR can also enhance the planning performance of open-source models through visual instruction tuning for rapid fine-tuning. The pipeline of the model is illustrated in Fig. 18. A JSON-formatted 6-DoF scene graph, processed through a text tokenizer, along with the image refined by SoM [151], is fed into an LLM (e.g., LLaMA [130, 131]) for supervised fine-tuning [79]. In the Open6DOR [28] task, we supplement the training dataset with additional samples retrieved and manually annotated from Objaverse [22], ensuring alignment with the object categories in the original benchmark. This dataset includes approximately 3,000 6-DoF object manipulation instructions. Using this data, we construct dialogue-style training data based on ChatGPT and train the SOFAR-LLaVA model. The training hyperparameters are detailed in Table XV. Similarly, we finetune PointSO on this training dataset and achieve superior performance on the Open6DOR task.

### D. ChatGPT API Costs

The knowledge of OrientText300K is derived from the annotations of 3D modelers on Sketchfab, combined with Chat-

GPT’s filtering and comprehension capabilities. To generate semantic direction annotations, we filter the 800K dataset of Objaverse [22] and apply ChatGPT to approximately 350K of the filtered data to generate semantic text-view index pairs. The OpenAI official API was used for these calls, with the GPT-4o version set to 2024-08-06 and the output format configured as JSON. The total cost for debugging and execution amounted to approximately \$10K.

## APPENDIX D

### ADDITIONAL BENCHMARK STATISTIC ANALYSIS

#### A. 6-DoF SpatialBench Analysis

We conduct a statistical analysis of the manually constructed 6-DoF SpatialBench, with category comparisons and word cloud visualizations shown in Fig. 19. We collect diverse image data from the internet, encompassing scenes such as indoor, outdoor, and natural landscapes. The questions may involve one or multiple objects, with varying levels of uncertainty in image resolution. Most importantly, we are the first to propose a VQA benchmark for orientation understanding, focusing on both quantitative and qualitative evaluation of orientation.

#### B. Open6DOR V2 Analysis

Open6DOR V2 builds upon Open6DOR V1 by removing some incorrectly labeled data and integrating assets and metrics into Libero, enabling closed-loop policy evaluation. The detailed number of tasks is presented in Table XVI, comprising over 4,500 tasks in total. Notably, we remove level 2 of the position track in Open6DOR V1 [28] because it requires manual inspection, which is not conducive to open-source use and replication by the community. Besides, due to the randomness of object drops in the scene, approximately 8% of the samples already satisfy the evaluation metrics in their initial state.

## APPENDIX E

### ADDITIONAL RELATED WORKS

#### A. 3D Representation Learning

Research on 3D Representation Learning encompasses various methods, including point-based [104, 105], voxel-based [91], and multiview-based approaches [46, 120]. Point-based methods [36, 111] have gained prominence in object classification [133, 144] due to their sparsity yet geometry-informative representation. On the other hand, voxel-based methods [24, 108, 156] offer dense representation and translation invariance, leading to a remarkable performance in object detection [19] and segmentation [2, 155]. The evolution of

TABLE XIII: Detailed zero-shot real-world 6-DoF rearrangement results.

Task	CoPa [51]	ReKep-Auto [56]	SoFAR-LLaVA (Ours)	SoFAR (Ours)
<i>Positional Object Manipulation</i>				
Move the soccer ball to the right of the bread.	2/3	3/3	3/3	<b>3/3</b>
Place the doll to the right of the lemon.	3/3	3/3	3/3	<b>3/3</b>
Put the pliers on the right side of the soccer ball.	1/3	1/3	3/3	<b>2/3</b>
Move the pen to the right of the doll.	3/3	2/3	3/3	<b>3/3</b>
Place the carrot on the left of the croissant.	2/3	3/3	2/3	<b>2/3</b>
Move the avocado to the left of the baseball.	3/3	2/3	2/3	<b>3/3</b>
Pick the pepper and place it to the left of the charger.	1/3	2/3	2/3	<b>2/3</b>
Place the baseball on the left side of the mug.	3/3	2/3	2/3	<b>3/3</b>
Arrange the flower in front of the potato.	2/3	3/3	2/3	<b>3/3</b>
Put the volleyball in front of the knife.	3/3	3/3	3/3	<b>3/3</b>
Place the ice cream cone in front of the potato.	2/3	3/3	2/3	<b>3/3</b>
Move the bitter melon to the front of the forklift.	2/3	1/3	2/3	<b>2/3</b>
Place the orange at the back of the stapler.	3/3	2/3	3/3	<b>3/3</b>
Move the panda toy to the back of the shampoo bottle.	2/3	3/3	3/3	<b>2/3</b>
pick the pumpkin and place it behind the pomegranate.	<b>3/3</b>	2/3	1/3	2/3
Place the basketball at the back of the board wipe.	2/3	2/3	3/3	<b>2/3</b>
Put the apple inside the box.	3/3	2/3	3/3	<b>3/3</b>
Place the waffles on the center of the plate.	3/3	2/3	3/3	<b>3/3</b>
Move the hamburger into the bowl.	2/3	2/3	2/3	<b>3/3</b>
Pick the puppet and put it into the basket.	1/3	2/3	2/3	<b>2/3</b>
Drop the grape into the box.	2/3	3/3	3/3	<b>2/3</b>
Put the doll between the lemon and the USB.	2/3	2/3	2/3	<b>3/3</b>
Set the duck toy in the center of the cart, bowl, and camera.	2/3	1/3	2/3	<b>2/3</b>
Place the strawberry between the Coke bottle and the glue.	2/3	2/3	3/3	<b>3/3</b>
Put the pen behind the basketball and in front of the vase.	2/3	1/3	2/3	<b>2/3</b>
Total success rate	74.7%	72.0%	81.3%	<b>85.3%</b>
<i>Oriental Object Manipulation</i>				
Turn the yellow head of the toy car to the right.	2/3	2/3	1/3	<b>2/3</b>
Adjust the knife handle so it points to the right.	2/3	1/3	2/3	<b>2/3</b>
Rotate the cap of the bottle towards the right.	2/3	2/3	2/3	<b>2/3</b>
Rotate the tip of the screwdriver to face the right.	0/3	0/3	1/3	<b>1/3</b>
Rotate the stem of the apple to the right.	0/3	1/3	1/3	<b>2/3</b>
Turn the front of the toy car to the left.	0/3	0/3	2/3	<b>2/3</b>
Rotate the cap of the bottle towards the left.	2/3	1/3	1/3	<b>2/3</b>
Adjust the pear's stem to the right.	1/3	1/3	1/3	<b>1/3</b>
Turn the mug handle to the right.	1/3	1/3	2/3	<b>2/3</b>
Rotate the handle of the mug to towards right.	2/3	1/3	<b>2/3</b>	1/3
Rotate the box so the text side faces forward.	0/3	1/3	0/3	<b>1/3</b>
Adjust the USB port to point forward.	0/3	0/3	1/3	<b>1/3</b>
Set the bottle upright.	0/3	1/3	0/3	<b>1/3</b>
Place the coffee cup in an upright position.	1/3	1/3	2/3	<b>2/3</b>
Upright the statue of liberty	0/3	0/3	<b>1/3</b>	0/3
Stand the doll upright.	0/3	1/3	0/3	<b>1/3</b>
Right the Coke can.	0/3	0/3	1/3	<b>1/3</b>
Flip the bottle upside down.	0/3	0/3	0/3	<b>1/3</b>
Turn the coffee cup upside down.	0/3	0/3	1/3	<b>1/3</b>
Invert the shampoo bottle upside down.	0/3	0/3	0/3	<b>0/3</b>
Total success rate	21.7%	23.3%	35.0%	<b>43.3%</b>
<i>Comprehensive 6-DoF Object Manipulation</i>				
Pull out a tissue.	3/3	3/3	2/3	<b>3/3</b>
Place the right bottle into the box and arrange it in a 3x3 pattern.	0/3	0/3	0/3	<b>1/3</b>
Take the tallest box and position it on the right side.	1/3	1/3	3/3	<b>3/3</b>
Grasp the error bottle and put it on the right side.	1/3	2/3	1/3	<b>2/3</b>
Take out the green test tube and place it between the two bottles.	2/3	2/3	3/3	<b>3/3</b>
Pack the objects on the table into the box one by one.	1/3	1/3	0/3	<b>1/3</b>
Rotate the loopy doll to face the yellow dragon doll	0/3	1/3	1/3	<b>1/3</b>
Right the fallen wine glass and arrange it neatly in a row.	0/3	0/3	0/3	<b>0/3</b>
Grasp the handle of the knife and cut the bread.	0/3	0/3	0/3	<b>1/3</b>
Pick the baseball into the cart and turn the cart to facing right.	0/3	0/3	1/3	<b>2/3</b>
Place the mug on the left of the ball and the handle turn right.	0/3	0/3	1/3	<b>1/3</b>
Aim the camera at the toy truck.	1/3	0/3	1/3	<b>1/3</b>
Rotate the flashlight to illuminate the loopy.	0/3	0/3	1/3	<b>1/3</b>
Put the pen into the pen container.	0/3	1/3	0/3	<b>1/3</b>
Pour out chips from the chips cylinder to the plate.	0/3	1/3	1/3	<b>1/3</b>
Total success rate	20.0%	26.7%	33.3%	<b>48.9%</b>

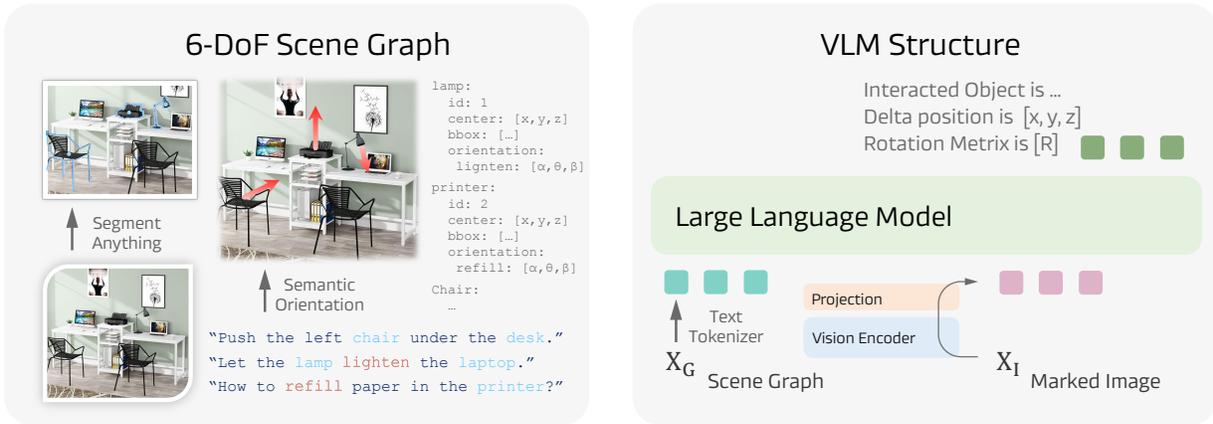


Fig. 18: Pipeline of SOFAR-LLaVA, a fine-tuned vision language model based on visual instruction tuning.

TABLE XIV: Details of PointSO model variants. This table format follows Dosovitskiy et al. [31].

Model	CLIP	Layers	Hidden size	MLP size	Heads	#Params
Small	ViT-B/32	12	256	1024	4	11.4M
Base	ViT-B/32	12	384	1536	6	19.0M
Large	ViT-B/32	12	512	2048	8	43.6M

attention mechanisms [134, 164] has also contributed to the development of effective representations for downstream tasks, as exemplified by the emergence of 3D Transformers [36, 86, 89]. Notably, 3D self-supervised representation learning has garnered significant attention in recent studies. PointContrast [148] utilizes contrastive learning across different views to acquire discriminative 3D scene representations. Innovations such as Point-BERT [159] and Point-MAE [101] introduce masked modeling [25, 47] pretraining into the 3D domain. ACT [29] pioneers cross-modal geometry understanding through 2D or language foundation models such as CLIP [115] or BERT [25]. Following ACT, RECON [107] further proposes a learning paradigm that unifies generative and contrastive learning. PPT [166] highlights the significance of positional encoding in 3D representation learning. Additionally, leveraging foundation vision-language models like CLIP [29, 115] has spurred the exploration of a new direction in open-world 3D representation learning. This line of work seeks to extend the applicability and adaptability of 3D representations in diverse and open-world/vocabulary scenarios [26, 27, 37, 87, 102, 163].

## APPENDIX F ADDITIONAL DISCUSSIONS

### A. Relation to Affordance & 6-DoF Pose Estimation

Conceptually, this semantic orientation is a counterpart of *affordance* [43, 44, 106, 118] but beyond, as SO and affordance all present potential actions and interactions with objects. However, SO also contains the spatial understanding of intra-object part-level attributes more than affordance learning. Compared to vanilla 6-DoF pose estimation, our proposed SO

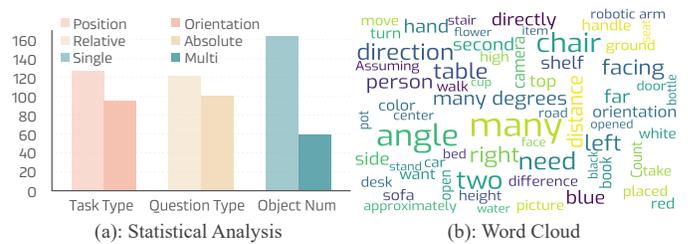


Fig. 19: 6-DoF SpatialBench statistics. (a) Statistical analysis of the task type, question type, and object relation. (b) Word cloud visualization.

combined with the 3-DoF translation understanding has the same DoF completeness. The difference is, our proposed SO is grounded by languages, making it useful for open-world manipulation requiring complicated spatial reasoning [28, 57, 128]. In addition, our Semantic Orientation can be auto-labeled from Internet 3D data that achieves higher scalability, introduced in the next section.

### B. Comparison to Concurrent Works

1) *Comparison with ReKep [56]*: Recently, ReKep has succeeded in executing complex robotic tasks, such as long-horizon manipulation, based on the relationships and constraints between spatial key points. Its structural design offers many insights that SOFAR can draw upon, yet it also presents several issues: (1) Overly customized prompt engineering. ReKep requires manually designed complex system prompts for each task during inference. While this approach may be described as “no training”, it cannot be considered a true zero-shot transfer. In contrast, SOFAR achieves genuine zero-shot transfer by eliminating the need for any human involvement during inference; (2) Using constraints based solely on key points fails to capture the full 6-DoF pose integrity of objects. For example, in the “pouring water” task, merely bringing the spout of the kettle close to the cup may lead to incorrect solutions, such as the kettle overturning; (3) ReKep requires all key points to be present in the first frame, and each step of the process—from

TABLE XV: Training recipes for PointSO and SoFAR-LLaVA.

Config	PointSO			SoFAR-LLaVA	
	Small	Base	Large	Finetune	SFT
optimizer	AdamW	AdamW	AdamW	AdamW	AdamW
learning rate	5e-5	5e-5	2e-5	5e-5	2e-5
weight decay	5e-2	5e-2	5e-2	5e-2	0
learning rate scheduler	cosine	cosine	cosine	cosine	cosine
training epochs	300	300	300	50	2
warmup epochs	10	10	10	5	0.03
batch size	256	256	256	256	128
drop path rate	0.2	0.2	0.2	0.2	-
number of points	10000	10000	10000	10000	-
number of point patches	512	512	512	512	-
point patch size	32	32	32	32	-
augmentation	Rot&Part&Noise	Rot&Part&Noise	Rot&Part&Noise	Rotation	-
GPU device	8×H800	8×H800	8×H800	8×H800	8×H800

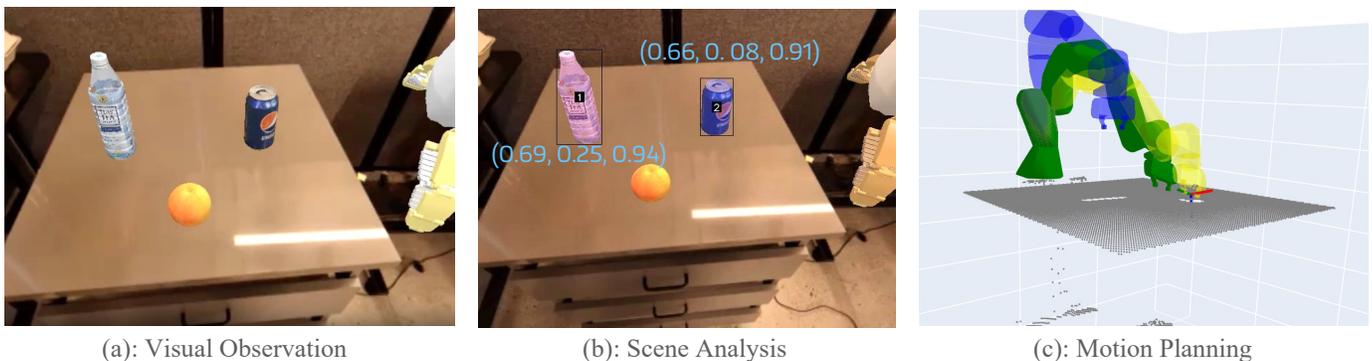


Fig. 20: An example of SoFAR how to finish “move near” task in SIMPLER [74].

mask extraction to feature dimensionality reduction, clustering, and filtering—introduces additional hyperparameters.

2) *Comparison with Orient Anything [139]*: Recently, Orient Anything also highlighted the importance of orientation in spatial perception and adopted a training data construction approach similar to Our PointSO. Our primary distinction lies in semantic orientation, which is language-conditioned orientation. In contrast, Orient Anything is limited to learning basic directions such as “front” and “top”. By aligning with textual information, semantic orientation better enhances spatial perception, understanding, and robotic manipulation.

### C. Future Works

Future work includes further expanding the OrientText300K with larger datasets like Objaverse-XL [21], enhancing the performance of semantic orientation through self-supervised learning and pretraining methods [29, 47, 107, 115], and demonstrating its effectiveness in a broader range of robotic scenarios, such as navigation [11], mobile manipulation [154], lifelong learning [76], spatio-temporal reasoning [56, 84, 85, 152], humanoid [14, 17, 48, 49], and human-robot interaction [82, 83].

## APPENDIX G ADDITIONAL VISUALIZATIONS

### A. Robotic Manipulation

As shown in Fig. 20, we present a visualization of executing a task named “move near”. According to the input image and task instruction - “move blue plastic bottle near pepsi can”, SoFAR can predict the center coordinate of the target object (bottle) and relative target (pepsi can), and it would infer the place coordinate and produce a series of grasp pose.

### B. 6-DoF SpatialBench

To further evaluate 6-DoF spatial understanding, we construct a 6-DoF SpatialBench. We present examples of question-answer pairs from the 6-DoF SpatialBench, with quantitative and qualitative questions shown in Figs. 21 and 22, respectively. The benchmark we constructed is both challenging and practical, potentially involving calculations based on the laws of motion, such as “Assuming a moving speed of 0.5 m/s, how many seconds would it take to walk from here to the white flower?” Moreover, it covers a wide range of spatially relevant scenarios across both indoor and outdoor environments.

TABLE XVI: **Statistics of Open6DOR V2 Benchmark.** The entire benchmark comprises three independent tracks, each featuring diverse tasks with careful annotations. The tasks are divided into different levels based on instruction categories, with statistics demonstrated above.

Track	Position-track						Rotation-track			6-DoF-track	Total	
Level	Level 0			Level 1			Level 0	Level 1	Level 2	-	-	
Task Catog.	Left	Right	Top	Behind	Front	Between	Center	Geometric	Directional	Semantic	-	-
Task Stat.	296	266	209	297	278	193	159	318	367	134	1810	4535
Benchmark Stat.	1698						1027			1810	4535	

### C. System Prompts

Prompt engineering significantly enhances ChatGPT’s capabilities. The model’s understanding and reasoning abilities can be greatly improved by leveraging techniques such as Chain-of-Thought [140] and In-Context Learning [9]. Figs. 23 and 24 illustrate the system prompt we used in constructing OrientText300K. Fig. 25, Fig. 26, and Fig. 27 illustrate the system prompt we used when evaluating SOFAR on Open6DOR (simulation), object manipulation (both simulation and real worlds), and VQA, respectively. Note that different from previous methods [54, 56], SOFAR does not require complicated in-context examples.

[Task Type: Position Question Type: Absolute]

[Question]: Count from right to left and start at 1, which two of the red flower pots are the group of people in the middle of?

[A]: "4 and 5"

[B]: "2 and 3"

[C]: "1 and 2"

[D]: "3 and 4"

[Answer]: C



[Task Type: Orientation Question Type: Absolute]

[Question]: If you want to align the orientations of the two chairs, what is the minimum angle you need to rotate the chair on the right?

[A]: "75°"

[B]: "55°"

[C]: "35°"

[D]: "15°"

[Answer]: C



[Task Type: Position Question Type: Absolute]

[Question]: Assuming a moving speed of 0.5 m/s, how many seconds would it take to walk from here to the white flower?

[A]: "3s"

[B]: "5s"

[C]: "7s"

[D]: "10s"

[Answer]: B



[Task Type: Orientation Question Type: Absolute]

[Question]: How many white chairs are facing the window?

[A]: "2"

[B]: "4"

[C]: "1"

[D]: "3"

[Answer]: A



Fig. 21: Visualization example of 6-DoF SpatialBench data samples.

[Task Type: Position Question Type: Relative]

[Question]: Which side of the steps is narrower?

[A]: "the left"

[B]: "the right"

[C]: "the middle"

[D]: "the same"

[Answer]: B



[Task Type: Orientation Question Type: Relative]

[Question]: Which direction does the handle of the cup in the upper right corner point to?

[A]: "left"

[B]: "right"

[C]: "front"

[D]: "back"

[Answer]: A



[Task Type: Position Question Type: Relative]

[Question]: How many compartments are there in the heart-shaped grid for storing books?

[A]: "5"

[B]: "3"

[C]: "6"

[D]: "4"

[Answer]: A



[Task Type: Orientation Question Type: Relative]

[Question]: If you are a driver driving a car on the road from near to far, which direction will you turn to?

[A]: "first turn left and then left"

[B]: "first turn right and then left"

[C]: "first turn left and then right"

[D]: "first turn right and then right"

[Answer]: C



Fig. 22: Visualization example of 6-DoF SpatialBench data samples.

**[System Prompt]**

You are an expert AI assistant for 3D object understanding.

The user imported a potentially uncalibrated 3D model into Blender and placed cameras in front, back, left, right, top, and bottom to render images, labeled from 1 to 6.

You are required to infer the entire 3D object based on these images and determine its attributes.

Your task is to assess the following attributes for each 3D model and respond with "true" or "false" for each question:

**Axis Alignment:** Determine whether the object is horizontally and vertically aligned across all views. Key features (e.g., edges, handles, or other distinct elements) of the object must be perpendicular or parallel to the cameras. Respond "true" if all views are aligned with the axis, "false" if not.

**Scene or Collection:** Determine whether the 3D model represents a 3D scene or a collection of independent objects (e.g. a room, outdoor scene, or multiple independent objects). Respond with "true" if it does, and "false" if it only contains a single object.

**White:** Determine whether the 3D model only has single white or gray colors, and lacks any other colors. Respond with "true" if it is white or gray, and "false" if it has any other colors (e.g., black or yellow).

**Ground:** Determine whether the 3D model includes a ground plane for auxiliary visualization. Respond with "true" if it does, and "false" if it only has the object.

**High Quality:** Determine whether the 3D model is a refined, well-constructed mesh without defects, such as point noise or streaking artifacts commonly found in low-quality RGBD scans. Respond with "true" if the mesh is clean and smooth, and "false" if it contains noise, roughness, or visual artifacts.

**Distinguishable Views:** Determine whether the 3D model has distinguishable views, or has clear semantic information in certain views (e.g., some 3D object has clear front, top directions). Respond with "true" if the 6 views show noticeable differences or have clear semantic information in certain views, and "false" if the views appear identical and there is no obvious semantic information on all views.

**Reasonable Object:** Determine whether the 3D model represents a common, recognizable, meaningful object. Respond with "true" if it is, and "false" if it is abstract, confused, or unrecognizable.

You need to first analyze the 3D object detail, and then output its correct attributes.

**[User]**

Standard Views:

Oblique Views: (Only for reference)

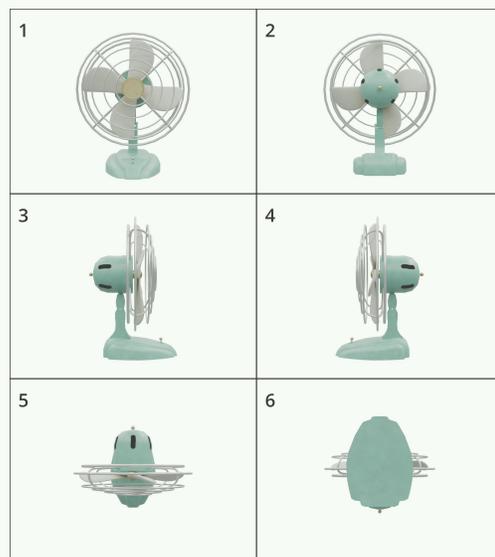


Fig. 23: The system prompt of ChatGPT-4o used for filtering Objaverse data.

**[System Prompt]**

You are a visual assistant specializing in interpreting 3D objects from multiple perspectives. You will receive 6 images of a 3D object from standard views (front, back, left, right, top, bottom), presented in random order. Typically, image 5 corresponds to the top view. Your task is to generate an instruction-index pair that identifies a meaningful semantic direction for the object, based on its function or commonly understood orientation. The instruction can be a verb, noun, adjective, or phrase, and must clearly relate to the object's function or orientation in everyday use. Ensure the direction is clear, objective, and uniquely meaningful.

**Examples:**

For a pen, the instruction might be "pen cap", and the index is the image with the pen cap facing the camera. For a cup, the instruction might be "handle", and the index is the image with the cup handle facing the camera. For a phone, the instruction might be "screen", and the index is the image with the phone screen facing the camera. For a table, the instruction might be "on", and the index is the image with the tabletop facing the camera. For a power outlet, the instruction might be "plug-in". Based on common knowledge, its semantic orientation is perpendicular to the power outlet's plane, along the direction of the power outlet's slots, and therefore, the index is the image with the power outlet pinholes' plane facing the camera. For a desk, the instruction might be "open the drawer". Based on common sense, the robot would need to pull the drawer open. The semantic direction corresponds to the direction of the drawer's extension, hence the index is the image with the drawer handle facing the camera. For a microphone, the instruction might be "speak", the semantic orientation is along the direction of the microphone's head, and therefore, the index is the image with the microphone head facing the camera.

You need to first analyze the category, attributes, characteristics, state, and usage of this 3D object in detail, and then output a pair of instructions and index.

When it is challenging to generate complex instructions, or when multiple views of the object are too similar to produce a unique instruction, you can use simpler instructions, such as "top" or "front".

The output format is as follows:

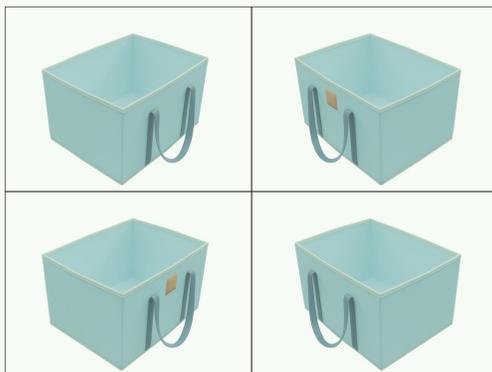
Analysis: "..."

Instruction: "..."

Index: 1-6

**[User]**

Oblique Views: (Only for reference)



Standard Views:



Fig. 24: The system prompt of ChatGPT-4o used for generating Semantic Direction-Index pairs.

### [Parsing System Prompt]

You are an assistant specialized in interpreting tabletop pick-and-place instructions for robotic manipulation. Your main goals are to identify relevant objects and analyze necessary orientations.

#### Key Objectives

1. Object Identification: Identify and list the objects mentioned in the instruction. Exclude the table itself.
2. Orientation Analysis: For the object needs to pick&place, determine any required orientation crucial to the task's success. If orientation isn't specified, leave the orientation list empty.
3. Direction Terms: Limit directional terms to these two categories:
  - Object Parts: e.g., "handle", "pen cap", "top"
  - Interaction Actions: e.g., "pour out", "open"Terms must be single words, not phrases or sentences.  
You must analysis both the instruction and the image to determine the object's direction attributes.
4. Disambiguation of Identification: If instructions reference vague objects (e.g., "else object", "all objects"), use visual information to clarify.
5. Disambiguation of Orientation: If the instructions describe complex rotation like "upright", you can interpret them as ensuring an object's relevant part is aligned with the z-axis (e.g., "bottle cap", "top").  
This disambiguation utilizes world knowledge, as we define the far-to-near direction as the x-axis, the left-to-right direction as the y-axis, and the bottom-to-top direction as the z-axis.  
Similarly, place an object to point forward means that the "top" of the object is oriented along the x-axis.

### [Reasoning System Prompt]

You are an assistant for spatial intelligence and robotic operations, specializing in pick-and-place tasks. Your role is to process robotic commands to pick a object and place it in a specific location.

#### Input Context:

1. Pick & Place Command: A directive specifying which object to pick and where to place it, including any specific pose requirements.
2. picked\_object\_info: A dictionary with the picked object's position in the world coordinate system.
  - Coordinates: Object center and bounding box in 3D (x, y, z), where:
    - x: Extends from far to near. Objects closer to the observer have larger x-values
    - y: Extends from left to right. Objects further to the right have larger y-values
    - z: Extends upward. Objects positioned higher have larger z-values
3. other\_objects\_info: A list of dictionaries with the position of other objects in the scene.
  - Coordinates: Object center and bounding box in 3D (x, y, z), same in the world coordinate system.

#### Objective:

1. Generate target placement position: Based on the spatial location descriptions provided in the instructions (e.g., 'behind,' 'between,' 'left,' etc.), as well as each object's center and bounding box (bbox), analyze and calculate the appropriate placement for the picked object.
  - front indicates positioning the object at an x-coordinate slightly larger than the reference object's x maximum.
  - right indicates positioning the object at a y-coordinate slightly larger than the reference object's y maximum.
  - between indicates positioning the object at the midpoint between two reference objects.

### [User]

Place the knife behind the clipboard on the table.  
And rotate the handle of the knife to left.



Fig. 25: The system prompt of Open6DOR tasks.

### [Parsing System Prompt]

You are a spatially intelligent AI specializing in interpreting objects, spatial directions, and interaction semantics for tasks involving spatial understanding or robotic manipulation.

The user will input an image and an instruction. Analyze user instruction and provide:

Objects: List involved objects using concise nouns or phrases, without any adjectives (e.g., the "top drawer" should be listed as "drawer").

Semantic Directions: Identify essential spatial or action-related terms, categorized as:

- Object Parts: e.g., "handle", "lid", "top".
- Action Terms: e.g., "pour out", "open".

Guidelines:

Focus on key spatial or action contexts for task completion.

Use implicit spatial conventions (Certain user instructions need to satisfy implicit constraints related to position and orientation.) if practical.

Avoid numeric values or absolute positions.

Only specify object-centric pose relationships, not inter-object positions (such as left, right, front, behind).

### [Reasoning System Prompt]

You are a robotic spatial intelligence and manipulation assistant, specialized in interpreting commands and scene structures for robotic object manipulation.

Your task is to analyze the user's directive and scene graph to guide the robot in identifying objects, computing spatial transformations, and producing step-by-step guidance for manipulation tasks.

Input Context:

1. Manipulation Command: A directive specifying which object to pick and where to place it, including any specific pose requirements.
2. Scene Graph: A dictionary with the scene objects' position and orientation in the world coordinate system.
  - Coordinates: Object center and bounding box in 3D  $(x, y, z)$ , where:
    - $x$ : Extends from near to far. Objects further to the observer have larger  $x$ -values
    - $y$ : Extends from right to left. Objects further to the left have larger  $y$ -values
    - $z$ : Extends upward. Objects positioned higher have larger  $z$ -values
  - Orientations of the object's parts (e.g., 'screen', 'handle') in 3D space.
    - $(1, 0, 0)$ : Points forward along the  $x$ -axis
    - $(0, 1, 0)$ : Points left along the  $y$ -axis
    - $(0, 0, 1)$ : Points upward along the  $z$ -axis

Objective: To process each command, follow these steps:

Target Identification: Identify the object to be picked up or manipulated.

Final Position: Specify the intended final position of the object after manipulation, in terms of  $x, y, z$  coordinates.

Orientation Mapping: For each semantic direction provided, compute the final orientation of the manipulated object in the world coordinate system.

### [User]

Open top drawer.



Fig. 26: The system prompt of general manipulation tasks.

### [Parsing System Prompt]

You are a spatially intelligent, embodied AI brain specialized in spatial and interactive understanding, tasked with interpreting objects, spatial directions, and relevant interaction semantics in response to the user's queries. The user provides commands or questions related to spatial intelligence or robotic manipulation, often with an image input. Your job is to analyze the given instruction and provide a list of objects involved in the task, alongside semantic orientations needed to complete the instruction effectively. You should focus on the key interaction directions required for successful completion without specifying numeric values or absolute positions, as these will be calculated by an expert model later.

Guidelines:

1. Focus on Semantic Directions: Define directions concisely using single terms that fall into one of these two categories:
  - Object Parts (e.g., "handle", "screen", "top")
  - Action-Oriented Terms (e.g., "pour out", "plug-in", "open")
2. Optimize for Simplicity: Choose terms that provide essential spatial or action context while remaining simple and intuitive for the model. Use only the most relevant directions or parts needed to complete the user's task.
3. Analysis: When necessary, use implicit spatial conventions where appropriate to ensure a practical output for the model.
4. Only object-centric pose related: Distinguish which object relationships are determined by position (such as left, right, front, behind) and which are determined by object pose, and we only focus on the direction of object centric pose.

### [Reasoning System Prompt]

You are a spatial intelligence assistant specialized in understanding 3D visual scenes and answering spatial reasoning questions.

The user will provide:

Image: An image of the scene.

Question: User question about the spatial relationships between objects in the scene.

Scene Graph: Additional information about the objects, including:

- id: object ID
- object name: object category
- center: 3D coordinates of the object's center
- bounding box: 3D bounding box coordinates
- orientation: object directions in 3D space

All the coordinates are in the camera coordinate system, where:

- x-axis: Extends from left to right in the image, objects located right have larger x-values
- y-axis: Extends from bottom to top in the image, objects located at top of the image have larger y-values
- z-axis: Extends from near to far in the image, objects located further away have larger z-values

You need to focus mainly on the image, the scene graph information is just for reference.

Avoid providing answers such as "cannot determine." Instead, provide the most likely answer based on the information available.

### [User]

How far between the left bottle and the right bottle?



Fig. 27: The system prompt of visual-question-answering tasks.