2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)
October 14-18, 2024. Abu Dhabi, UAE

Open6DOR: Benchmarking Open-instruction 6-DoF Object Rearrangement and A VLM-based Approach

Yufei Ding[1,2*], Haoran Geng[1,3*], Chaoyi Xu[2], Xiaomeng Fang[4], Jiazhao Zhang[1,4], Songlin Wei[1,2], Qiyu Dai[1], Zhizheng Zhang[2], He Wang[1,2,4†]

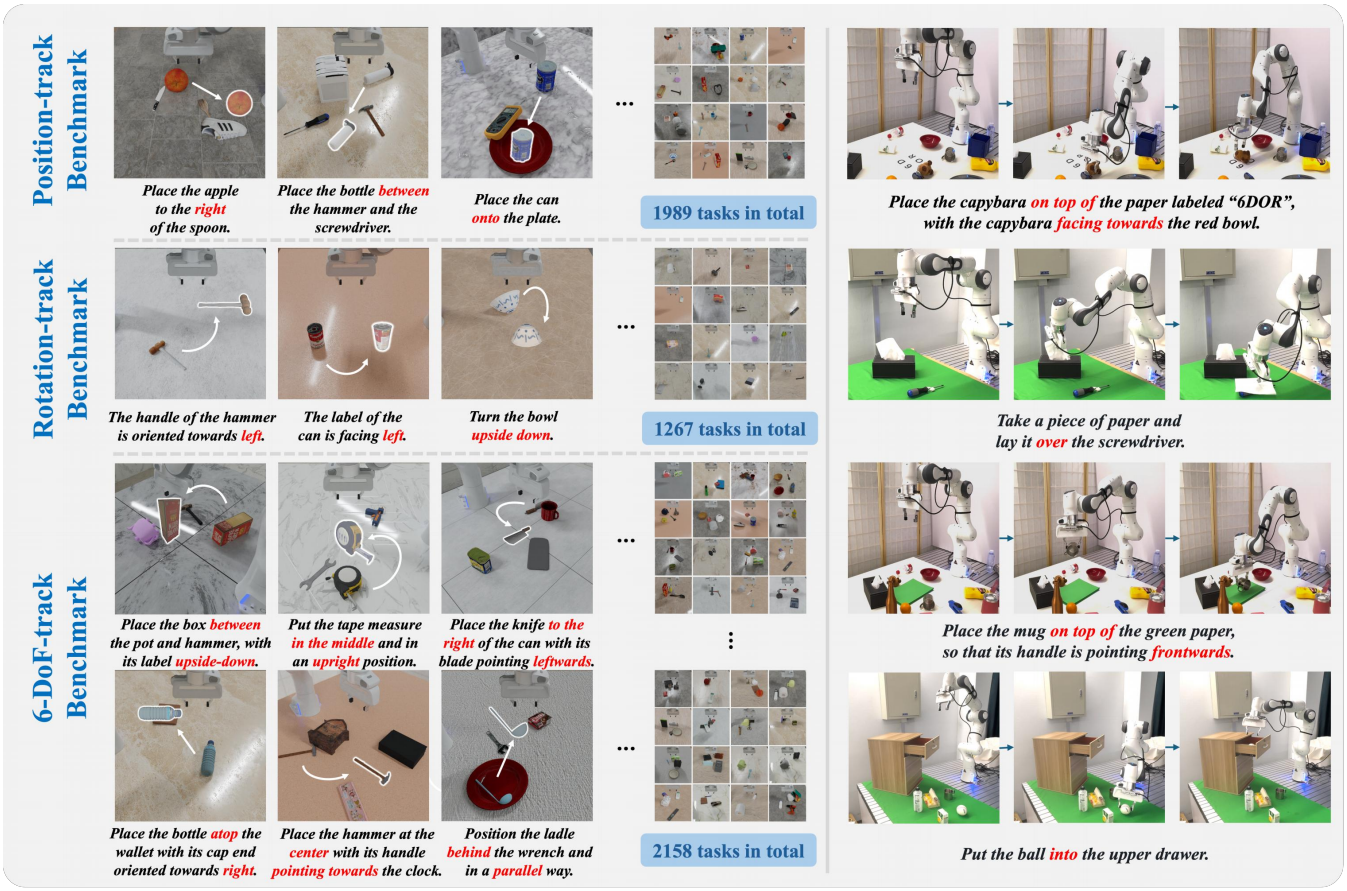

Fig. 1: **Open6DOR Benchmark and Real-world Experiments.** We introduce a challenging and comprehensive benchmark for Open-instruction 6-DoF object rearrangement tasks, termed Open6DOR. Following this, we propose a zero-shot and robust method, Open6DOR-GPT, which proves effective in demanding simulation environments and real-world scenarios.

***Abstract*— The integration of large-scale Vision-Language Models (VLMs) with embodied AI can greatly enhance the generalizability and the capacity to follow open instructions for robots. However, existing studies on object manipulation are not up to full consideration of the 6-DoF requirements, let alone establishing a comprehensive benchmark. In this paper, we propel the pioneer construction of the benchmark and approach for Open-instruction 6-DoF Object Rearrangement (Open6DOR). Specifically, we collect a synthetic dataset of 200+ objects and carefully design 5400+ Open6DOR tasks. These tasks are divided into the Position-track, Rotation-track, and 6-DoF-track for evaluating different embodied agents in predicting the positions and rotations of target objects.**

Besides, we also propose a VLM-based approach for Open6DOR, named Open6DOR-GPT, which empowers GPT-4V with 3D-awareness and simulation-assistance while exploiting its strengths in generalizability and instruction-following. We compare the existing embodied agents with our Open6DOR-GPT on the proposed Open6DOR benchmark and find that Open6DOR-GPT achieves the state-of-the-art performance. We further show the impressive performance of Open6DOR-GPT in diverse real-world experiments.

* Equal contribution.
[1] Peking University.
[2] Galbot
[3] University of California, Berkeley
[4] Beijing Academy of Artificial Intelligence.
Corresponding author: hewang@pku.edu.cn

I. Introduction

Large-scale multimodal models [8, 26] pre-trained on web-scale data have revolutionized numerous fields beyond what was previously imaginable, enabling open-vocabulary text understanding and 2D visual perception. The pursuit to bring general intelligence into the robotic realm stands at an exciting yet nascent stage, calling for stronger capabilities in 3D-aware perception, semantic understanding and complex reasoning.

The advent of large-scale embodied models, exemplified by the RT series [2, 3, 15] and VoxPoser [17], has demonstrated considerable progress in mobile or fixed-station pick-and-place operations. While these models are capable of rearranging the object positions following human instructions, they fall short of satisfying full 6-DoF object placement instructions that involve specified 3D rotations. This limitation renders them incompetent at many practical robotic applications, where both object position and orientation are essential. For instance, in our daily life we often need a water bottle to be placed upright, while on the shelves in retail stores, goods should face the same direction. Moreover, previous works [3, 15, 17] are often evaluated on their own robots in their own scenes with self-reported performance and nonstandard evaluation metrics. The absence of a standard evaluation protocal condone cherry-picking, obstruct comparative assessment, and thus, hinder the iterative enhancement of effective approaches.

In this paper, we target the task of Open-instruction 6-DoF Object Rearrangement, referred to as Open6DOR, which requires embodied agents to move the target objects according to open instructions that specify its 6-DoF pose. Open6DOR represents a fundamental skill for robotic manipulation tasks, presenting significant challenges in integrating instruction comprehension, 3D visual perception, and motion planning capabilities. Specifically, we promote the envelope of Open6DOR from two perspectives:

1) **Benchmark construction:** We construct a standardized benchmark, namely Open6DOR Benchmark, which comprises 5414 tasks designed with more than 200 objects across diverse categories in simulation environments. For comprehensive evaluation, we divide the Open6DOR benchmark into the position-track, rotation-track, and 6-DoF-track, each providing manually configured tasks along with comprehensive and quantitative 3D annotations. These tracks enable independent or combined assessments of translational, rotational, and overall performance.

2) **VLM-based approach:** We propose a VLM-based approach for Open6DOR tasks. Due to the aforementioned challenges of Open6DOR, all prior works, such as VoxPoser [17] and Dream2Real [20], fail to fulfill Open6DOR's 6-DoF requirements adequately. Among these efforts, Dream2Real [20] attempts to consider position and rotation dimensions simultaneously by imagining randomly rearranged scenes and leveraging VLM as an evaluator. This leads to almost intolerable time costs resulted from numerous renderings and VLM inferences, as well as unsatisfactory results due to the VLM's limited 3D perception, which renders it an incompetent critic. In contrast, we propose Open6DOR-GPT, which explicitly integrates 3D information from the initial scene into GPT-4V with equipped auxiliary modules and decomposes the translational and rotational determinations. In this way, we augment GPT-4V with 3D understanding capabilities and improve efficiency by reducing the determination space with decoupled modeling and simulation-assistance. Open6DOR-GPT achieves state-of-the-art performance in both benchmark evaluation and real-world experiments.

II. Related Work

A. Object Rearrangement Benchmarks

Benchmarking object rearrangement is extremely challenging and requires extensive annotation of ground truth placement. Existing object pick-and-place benchmarks [9, 33, 36] leverage pre-built environments such as Repilca-CAD [28] or hand-crafted scenes and reconstructed objects [4]. These benchmarks contain annotations about the initial position and target position of each object, which are used to evaluate placement accuracy. Aside from simulator benchmarks, real-world benchmarks [23] directly evaluate baselines in real robots. Despite their ability to assess real-world performance, these benchmarks fail to accurately replicate the testing environments for all baselines and are not permanently available due to hardware limitations. Different from existing works, we propose a 6-DoF object rearrangement benchmark that comprehensively evaluates 6-DoF placement, providing both position and rotation annotations.

B. Object Rearrangement Methods

Object rearrangement [1] requires an embodied agent to reposition objects based on specific instructions. Early methods [11, 12] use task-and-motion-planning (TAMP), which relies on predefined action primitives and near-perfect scene knowledge for trajectory sampling, making it computationally inefficient and unscalable for complex scenarios. To enhance generalizability and efficiency, recent research has shifted towards learning-based approaches [9, 14, 25, 29, 35], trained in simulators for high-level planning [9, 25, 29] or low-level actions [14, 35] prediction. Despite strong performance in simulators, they suffer from a severe sim-to-real gap [25]. Recent advanced methods leverage large language models (LLM) [6] or vision language models (VLM) [22, 27] for real-world deployment. A part of these approaches [16, 20] construct carefully designed prompts to describe the environments, and query off-the-shelf LLMs or VLMs for placement guidance before execution. Other methods [32] train LLMs from self-collected data and directly output low-level actions. Departing from existing works that predominantly focus on location, our method also emphasizes rotation, leading to a 6-DoF rearrangement method.

C. VLM for Open-instruction 6-DoF Tasks

Large models trained on internet-scale data have demonstrated great potential in high-level planning [13, 17, 21]. The recent advent of VLMs further bridges the gap between visual perception and textual interpretation, empowering embodied agents with semantic understanding of scenes and instructions [19] to perform 6-DoF tasks. Some of the prior works [17] leverage VLMs to compose 3D value maps, planning robot trajectories based on instructions; while other methods [20] employ VLMs as evaluators, generating goal states in the form of images for VLM to assess. However, both approaches fail to consider complex tasks that strictly specify the rotation of an object. Moreover, Dream2Real suffers from excessive time-consumption and VLM's inaccurate judgment. In contrast, our method addresses the rotation

and position aspect of the 6-DoF problem in a decoupled way, enhancing VLM's decision-making capabilities while expediting the inference process.

III. Open6DOR Benchmark

A. Open6DOR Task Formulation

We aim to identify a shared, fundamental component within complex embodied problems and, based on that, concisely formulate a elementary task, which we name Open6DOR. Open-instruction object rearrangement refers to the process wherein an embodied agent repositions objects within a scene from an initial state, following specific instructions. Specifically, a 6-DoF (Degrees of Freedom) object rearrangement task focuses on repositioning objects in a 6-DoF space, which includes both orientational and translational movement. We define each of these pick-and-place processes as an Open6DOR task, where a single target object is moved from its initial pose to a goal pose, guided by an open-vocabulary instruction. The input includes a single-view RGB-D image of the initial scene, denoted as I_{rgbd}, alongside a task instruction $\tilde{I}$ that specifies the desired pose of a target object in the scene. Based on these, the model is required to output the goal position P_{goal} and goal rotation R_{goal} of the target object. The Open6DOR task lies at the core of various long-horizon or complex-scene problems, simultaneously evaluating a model's capabilities in instruction following, 3D perception, and semantic understanding.

B. Open6DOR Benchmark Overview

The Open6DOR Benchmark is specifically designed for Open6DOR tasks grounded in simulation environment. To ensure comprehensive evaluation, we provide three specialized tracks of benchmark: Rotation-track $\mathcal{B}_r$, Position-track $\mathcal{B}_p$, and 6-DoF-track $\mathcal{B}_{6DOR}$. $\mathcal{B}_r$ encompasses tasks achievable through a singular rotational movement—for example, '*place the cup upside down*'. $\mathcal{B}_p$ focuses on object repositioning, like '*put the cup between A and B*', without specification of the object's orientation. Meanwhile, $\mathcal{B}_{6DOR}$ integrates both rotation and position requirements, such as '*place the mug in front of A with its handle pointing towards the left*'. Overall, the Open6DOR Benchmark consists of 5k+ tasks, featuring intricate configurations, realistic scenes, comprehensive annotations, and interactive environment. The construction process involved substantial manual effort, and we will elaborate on its individual components below.

Asset collection. The synthetic object dataset $\mathcal{O}_s$ comprises 200+ items spanning 70+ distinct categories. Originally derived from YCB [5] and Objaverse-XL [7], the objects are carefully filtered with a 10% retention rate to ensure they are physically intact and semantically suitable for table-top placement. We also adjust the scale of all objects and use a standardized format of mesh representation. The objects are then classified into semantic categories and tagged with instruction labels for the convenience of future analysis.

Task configuration. Each task in the Open6DOR Benchmark is set within a table-top environment, where multiple objects are placed randomly. Initial object poses are carefully configured to avoid issues such as model clipping, range exceeding or unstable placement. Moreover, we manually design diverse instructions based on the target object, including positional and rotational requirements. The tasks are further reviewed to prevent occlusion or infeasible settings, ending with a total of 5k+ tasks, with statistics shown in Tab. I.

Annotation and evaluation. For rotational assessment, we manually annotate each goal pose of a specific object as quaternions. The axis of symmetry, if present, is also specified to represent rotational equivalence. These annotations enable calculation of deviations between predicted and ground-truth rotations. For positional evaluation, we design heuristic functions to judge whether the spatial arrangement conforms to the instruction. In a task like '*place A to the left of B*', we compare the predicted position of A with the position of the reference object B to verify its correctness.

Simulation setting. The Open6DOR Benchmark is based on Isaac Gym [24], offering an interactive environment and executable platform. All tasks can be directly loaded into the simulator, in which users may control a robotic arm to complete the tasks. We also provide motion-planning APIs that generate actions based on a goal pose, implemented with cuRobo [30] and IsaacGym Motion Planing Library [24].

Rendering augmentation. To enhance observation realism, we propose a rendering API based on Blender [10]. Using such high-quality rendering, we generate a single-vew RGB-D image dataset, which serves as observation input for models. Additionally, the API enables customization of camera positions, lighting conditions, and background textures to accomodate personalized observation settings.

C. Position-track Benchmark

Task setting. The Position-track Benchmark includes 1989 tasks, each set in a table-top scene that contains 2-6 objects. We render a realistic RGB-D image for each of the scenes, resulting in a single-view RGB-D image dataset $\mathcal{V}_p$. For position instructions $\mathcal{I}_p$, we design three levels that evaluate the understanding of: basic directions (Level 0) such as *Left, Right, Top, Behind, Front*; spatial relations (Level 1) such as *Between, Center*; and customized commands (Level 2) like '*put A into B*'. The task instructions adhere to a uniform format, such as '*place A in front of B*', where A and B are specified based on the context of individual scenes.

Evaluation metrics. We assess the predicted goal position P_{goal} according to the annotated position range. A position that falls into that range is considered correct; otherwise, it is wrong. For instance, in the *Left* task category, we verify whether the predicted position is to the left of the reference object (indicated by a smaller y-axis coordinate).

D. Rotation-track Benchmark

Task setting. The Rotation-track Benchmark consists of 1267 diverse tasks, each set in a scene containing a randomly placed object from $\mathcal{O}_s$. We provide a single-view RGB-D image dataset $\mathcal{V}_r$ featuring the scenes. For the instruction input $\tilde{I}$, we label each object with 1-5 rotation-specified

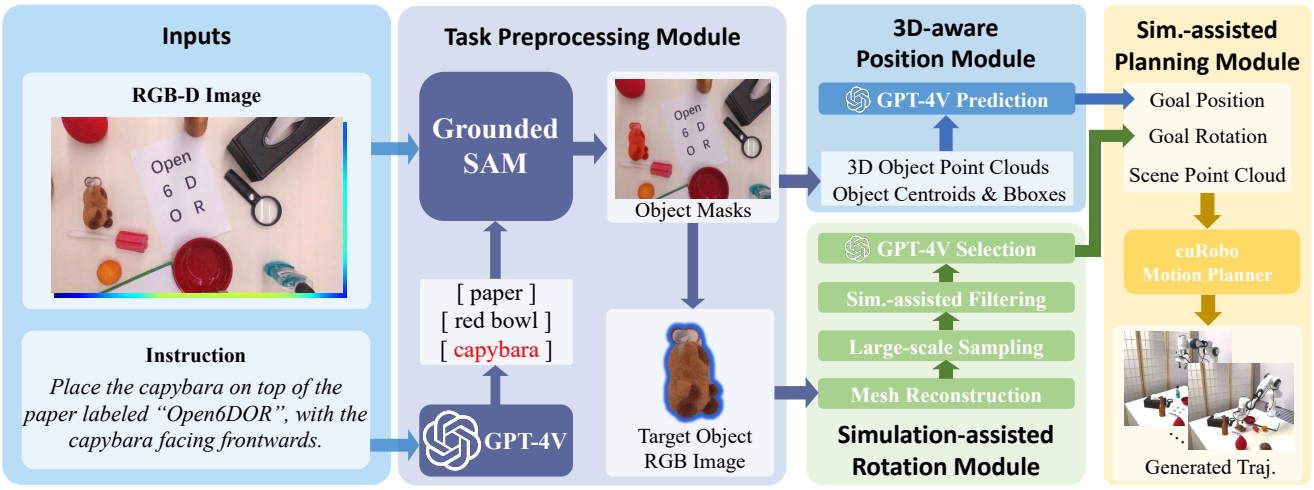

Fig. 2: **Method Overview.** Open6DOR-GPT takes the RGB-D image and instruction as input and outputs the corresponding robot motion trajectory. Firstly, the preprocessing module extracts the object names and masks. Then, two modules simultaneously predict the position and rotation of the target object in a decoupled way. Finally, the planning module generates a trajectory for execution.

Track	Position-track								Rotation-track			6DoF-track
Level	Level 0					Level 1		Level 2	Level 0	Level 1	Level 2	-
Task Catog.	Left	Right	Top	Behind	Front	Between	Center	Customized	Geometric	Directional	Semantic	-
Task Stat.	341	318	238	351	323	217	191	10	526	511	230	2158
Benchmark Stat.	1989								1267			2158

TABLE I: **Statistics of Open6DOR 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.

instructions $\mathcal{I}_r$ based on its features. The instructions are categorized into 3 levels that progressively increase in difficulty. Level 0 includes basic instructions that are related to the geometric shape of the object, such as *Upright* and *Upside down*. Level 1 generally requires a higher understanding of direction and orientation, such as '*handle to the left*'. Level 2 contains harder instructions concerning semantic and textual information of the object, such as '*label forth*' and '*characters right side up*'.

Evaluation metrics. We assess the predicted goal rotation R_{goal} by comparing it with the annotated ground-truth rotation. Specifically, we define a rotation range for each task based on the annotations and object geometry. Rotation results that fall within this range are considered correct; otherwise, they are wrong.

E. 6-DoF-track Benchmark

Task setting. The 6-DoF-track Benchmark comprises 2158 tasks, providing a comprehensive evaluation that jointly assesses the rotation and position performance of an Open6DOR task. The formulation of the RGB-D scene image I_{rgbd} aligns with that of the Position-track. For instruction $\tilde{I}$, we combine instructions from $\mathcal{I}_p$ and $\mathcal{I}_r$, forming instructions that specify both the position and rotation of an object. Each instruction is paired with an RGB-D scene image as the task input, and we exclude the incompatible pairs to ensure that the tasks are well-defined and performable.

Evaluation metrics. We evaluate the quality of a 6-DoF pose from two perspectives: rotation and position. Specifically, we manually annotate the desired rotation and position of the target object based on the instruction. We consider a task successful only when it satisfies both criteria.

IV. Open6DOR-GPT

A. Method Overview

As shown in Fig. 2, we enhance GPT-4V [26]'s capabilities to address the challenges of the Open6DOR task in a decomposed way. Initially, the Task Preprocessing Module deciphers $\tilde{I}$ based on the I_{rgbd} and feeds the resulting images to the Position Module and Rotation Module respectively. Within the two modules, we empower GPT-4V with 3D awareness and simulation assistance, thereby effectively outputting the predicted goal position P_{goal} and rotation R_{goal}. Finally, the Simulation-assisted Planning Module identifies a suitable grasping pose and plans out an optimal action trajectory to accomplish the task. We will first introduce each module of our proposed system in subsection *B-E* to explain how an Open6DOR task is accomplished. We then elaborate on how the system tackles long-horizon tasks with multiple rounds of operations.

B. Task Preprocessing Module

With the single-view RGB-D Image I_{rgbd} and the task instruction $\tilde{I}$ as input, this module leverages GPT-4V to

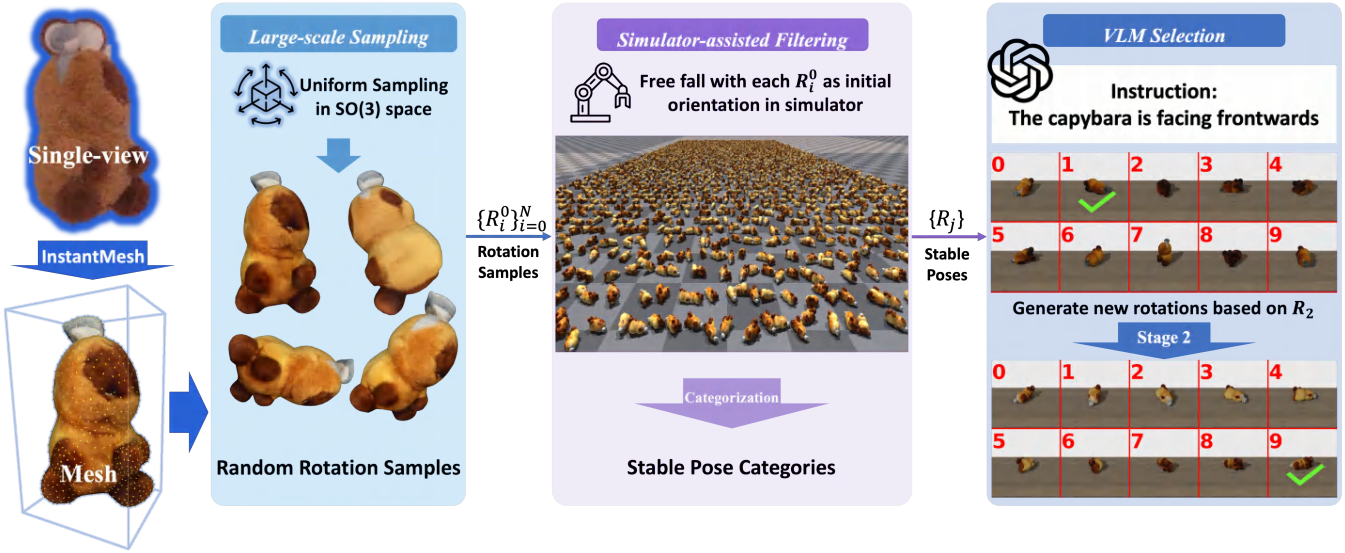

Fig. 3: **Simulation-assisted Rotation Module.** Firstly, a textured mesh is reconstructed from the single-view image of the target object. Then, we employ large-scale sampling to obtain multiple rotation samples. This sample set is then narrowed down through a simulation-assisted filtering process to derive several stable pose categories. Finally, we generate rendered images of the pose candidates, from which GPT-4V selects the optimal goal rotation.

interpret the instruction and identifies object names$\{O_i^{name}\}$, which in turn triggers GroundedSAM [18] to generate a set of labeled masks. Based on the masked Image I_{mask}, the RGB image of the target object I_{object} is extracted. These images are used in subsequent modules.

C. 3D-aware Position Module

Taking the masked RGB-D image I_{mask} and task instruction $\tilde{I}$ as input, the 3D-aware Position Module $\mathcal{M}_p$ determines and outputs the goal position.

To incorporate three-dimensional (3D) data into GPT-4V's understanding, our approach utilizes back-projection based on I_{mask} to generate a 3D masked point cloud, symbolized as PC_i^{3d}. This computation includes determining the centroid Center_i^{3d} and bounding box Bbox_i^{3d} of the point cloud associated with the queried object.

$$PC_i^{3d} = \text{BackProj}(I_{rgbd}(\text{Mask}_i^{2d})) \tag{1}$$

$$\text{Center}_i^{3d}, \text{Bbox}_i^{3d} = \text{Mean}(PC_i^{3d}), \text{Max}(PC_i^{3d}) - \text{Min}(PC_i^{3d})$$

These spatial attributes are then integrated back into the prompt for GPT-4V, facilitating the model to accurately ascertain the goal position for the target object P_{goal}.

D. Simulation-assisted Rotation Module

As illustrated in Fig. 3, with the single-view RGB image of the target object I_{object} and the task instruction $\tilde{I}$ as input, the rotation module would output the goal rotation R_{goal} for the object. We first reconstruct the target object from I_{object} using InstantMesh [34] , resulting in a textured mesh denoted as M. The reconstruction process is followed by four phases: (1) large-scale sampling (2) simulation-assisted filtering (3) rotation categorization (4) GPT-4V selection.

Large-scale sampling. In Phase 1, we randomly sample N rotations $\{R_i^0\}_{i=0}^N$ as initial inputs for subsequent phases. We set $N = 3600$ and use Uniform Sampling in SO(3) space (Special Orthogonal Group in 3D space) to ensure diversity.

Simulation-assisted filtering. Now that we have a large pool of rotation candidates $\{R_i^0\}$, the goal of Phase 2 is to filter out the unreasonable candidates and narrow down the sample pool efficiently. To accomplish this task, we first examine the stability of the rotation candidates by incorporating a physics simulator through which all the unstable poses are excluded. To be specific, each R_i^0 is applied to a replica of M, denoted by M_i, as its initial rotation in the simulator, amounting to N actors $\{M_i^0\}_{i=0}^N$ in total. Then, all the actors are dropped from a low height, landing on the ground with diverse ending poses. We record the relative rotation from the original mesh M to the ending pose of M_i as R_i^t. By now, we have narrowed down the originally random and irregular distribution of $\{R_i^0\}_{i=0}^N$ to a more condensed space $\{R_i^t\}_{i=0}^N$ of geometrically stable rotations. By simulating the dropping process within 10 seconds, we avoid the time-consuming inferences of vision language models while accurately extracting stable poses of M.

Rotation categorization. From Phase 2, we've obtained a set of rotations $\{R_i^t\}_{i=0}^N$ that guarantee the stability of the object. However, these rotations are unevenly cluttered around several distinct centers, each representing a stable pose. In order to categorize these rotations and extract a representative for each stable pose, we propose a criterion by which we judge whether two rotations R_a and R_b could be classified into the same category. In brief, we regard R_a and R_b as identical if they meet any one of the following two criteria: (1) the relative rotation from R_a to R_b is small in magnitude (2) R_a and R_b represents symmetrical poses that are transferrable via a rotation along the z-axis (perpendicular to the table surface). We define a threshold for each of these criteria and calculate whether R_a and R_b could be considered as identical. Through this method, we classify the rotations into several clusters and represent each category with a single rotation, which largely reduces the total number of rotation candidates for the task. Therefore, we successfully narrow

down $\{R_i^t\}_{i=0}^N$ to a small set: $\{R_j\}_{i=0}^n$, in which each rotation represents a distinct stable pose.

GPT-4V selection. During the last phase of our rotation engine, we feed the filtered set of rotations $\{R_j\}_{i=0}^n$ in the form of the 2D image along with the original instruction to GPT-4V and let it select a candidate as the goal rotation R_{goal}. To transform $\{R_j\}$ into a modal that VLM could easily understand, we apply each R_j to M and render the image of the object loaded on a table accordingly. The images are then arranged together into a collage, with an index mark on the upper left of each grid. Empirically, we found that this strategic approach of numbering and segmenting the images boosts the performance of GPT-4V in selecting the right answer. To further enhance our method, we employ a two-stage strategy that resamples a set of rotation candidates based on the rotation GPT-4V has selected in Stage 1. After this second round of adjustment, the goal rotation R_{goal} is determined and outputted.

E. Simulation-assisted Planning Module

Utilizing the predicted goal position P_{goal} and goal rotation R_{goal}, the planning module formulates an effective execution strategy with simulation assistance. Firstly, the Grasp Detection Model, GSNet [31], takes the refined point cloud PC_{refined} as input and generates a series of scored grasping pose candidates $\{(G_j, s_j)\}$. From $\{G_j\}$, GPT-4V selects valid grasping poses that rest on the target object by leveraging the object bounding box Bbox^{3d} derived from the 3D-aware Position, resulting in a ranked set of $\{\tilde{G}_j\}$.

$$\{(G_j, s_j)\} = \text{GSNet}(PC_{\text{refined}}) \tag{2}$$

$$\{\tilde{G}_j\} = \text{Sorted}_s(\{(G_j)\}_{\text{Bbox}^{3d}}) \tag{3}$$

Next, we use cuRobo [30] as the motion planner, which enumerates $\{\tilde{G}_j\}$ within the simulator based on their score rankings $\{s_j\}$, and identifies a trajectory that optimizes both grasping and placement, denoted as $\mathcal{T}$. Finally, the robot employs its control system to accomplish execution according to $\mathcal{T}$.

F. Long-horizon Open6DOR Execution

With the framework outlined in previous sections, our system is capable of managing individual Open6DOR tasks. For long-horizon rearrangement tasks, we first employ GPT-4V to evaluate whether multiple Open6DOR steps are required. If the task necessitates multiple steps, GPT-4V leverages high-level planning to divide the instruction $\tilde{I}$ into several discrete execution steps $\{\tilde{I}_i\}$. Subsequently, for each individual step I_i, we apply the previously described methodology to carry out the task. Upon the completion of each Open6DOR sequence, the overall task is considered complete. This approach ensures a systematic and efficient handling of complex rearrangement tasks, breaking them down into manageable steps that are executed with precision.

Success Rate (%)	Level 0	Level 1	Level 2	Overall
GPT-4V [26]	46.8	39.1	50.0	45.2
Dream2Real* [20]	17.2	11.0	-	15.9
VoxPoser* [17]	35.6	21.7	0.0	32.6
VoxPoser(VLM)* [17]	37.2	19.9	0.0	33.5
Open6DOR-GPT	**78.6**	**60.3**	**80.0**	**74.9**

TABLE II: **Results on Position-track Benchmark.** We compared our approach against several benchmarks for positioning proposals. This includes: (1) GPT-4V [26], utilizing pixel input to predict object placement and employing depth for 3D location. (2) A tailored Dream2Real [20] baseline for our task. (3,4) VoxPoser [17] original and adapted versions, aligning with our goals. Our tests include GPT-4V's Large Language Model (LLM) and Vision-Language Model (VLM) setups, with an asterisk denoting ground-truth data usage as reference baselines.

V. Experiments

A. Results on Position-track Benchmark

We evaluate the performance of our position module and several baselines on the Position-track Benchmark. As shown in Table II, both Dream2Real and GPT-4V demonstrate incompetence at precise position determination. VoxPoser [17], another baseline, yields unsatisfactory performance due to reliance on Large Language Models (LLM) without visual inputs. But even when adapted to a VLM-assisted version and incorporating image data, VoxPoser(VLM)* fails to gain significant improvements. Comparatively, our approach markedly surpasses all these baselines by over 30 percent, demonstrating superior and consistent performance on the Position-track Benchmark.

B. Results on Rotation-track Benchmark

Our Rotation Module comprises four phases aimed at enhancing GPT-4V [26] through a simulation-assisted sample-and-filter mechanism. To evaluate the effectiveness of each phase, we conduct ablation studies using the Rotation-track of Open6DOR Benchmark, with results detailed in Table III. We compare our approach with Dream2Real [20], replacing their CLIP Model with GPT-4V to ensure fairness. As shown in the first row of the table, directly querying GPT-4V yields an unsatisfactory success rate. Substituting our module with Dream2Real's method also leads to a noticeable performance decline. However, upon incorporating the Simulation-assisted Filtering Phase, we observe a noticeable performance increase as GPT-4V is able to choose from a confined set of rotation candidates. Further enhancements are achieved by integrating a 2-stage VLM Selection Phase. Notably, the combined module outperforms Dream2Real* by 10 percent.

C. Results on 6-DoF Benchmark

We evaluate our entire pipeline using the 6DoF-track of Open6DOR Benchmark. The evaluation of rotational, positional, and joint performance are presented in Table IV. For baseline methods, we found a limited number of works addressing the 6DoF problem and chose to compare with Dream2Real [20]. However, the original Dream2Real method

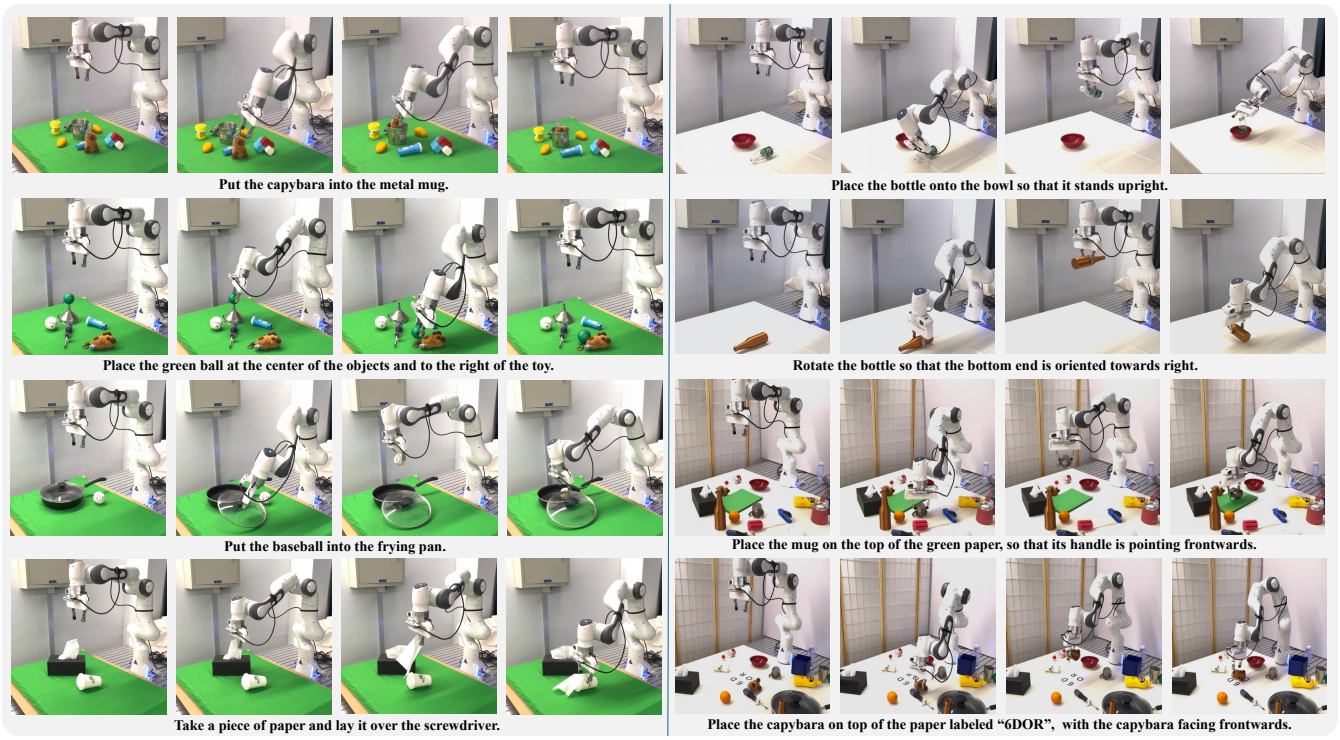

Fig. 4: **Real-world Experiments.** We ground Open6DOR-GPT in real-world settings and conduct various tasks including long-horizon ones, demonstrating its zero-shot generalization potential across challenging tasks.

Success Rate(%)	Level 0	Level 1	Level 2	Overall
GPT-4V [26]	9.1	6.9	11.7	9.2
Dream2Real* [20]	37.3	27.6	26.2	31.3
S-F + GPT-4V	41.1	30.7	30.4	38.4
Open6DOR-GPT (S-F + 2-Stage 4V)	**45.7**	**32.5**	**49.8**	**41.1**
Open6DOR-GPT*	**48.2**	**32.6**	**60.0**	**44.1**

TABLE III: **Results on Rotation-track Benchmark.** Quantitative comparison with a refined version of Dream2Real [20] method (replacing CLIP Model with GPT-4V), and ablation studies of different phases in the Rotation Module. 'S-F' stands for 'Sampling-Filtering'. '*' means using ground-truth mesh instead of reconstructed ones. The first three rows ablate Phase1-4, Phase3-4, and Stage2 in Phase 4, respectively.

Success Rate (%)	Rotation	Position	Overall	Time Cost(s)
Dream2Real [20]	-	-	-	>700
Dream2Real* [20]	18.7	26.2	13.5	358.3
Open6DOR-GPT	**40.0**	**84.8**	**35.6**	**126.3**

TABLE IV: **Results on 6-DoF-track Benchmark.** We compare our method with an optimized version of Dream2Real [20] on the 6DoF Benchmark. The three columns depict the quality of the goal pose in terms of rotation, position, and overall performance.

is excessively time-consuming, requiring over 10 minutes per task for completion. To fully assess Dream2Real on the 6DoF-track Benchmark (containing 2k+ tasks), we skip the scene-scanning and reconstruction section of their method, using ground-truth mesh and image instead. Despite the disadvantage of the mesh quality on our side (our method uses reconstructed mesh from the original image), Open6DOR-GPT significantly outperforms Dream2Real* by over 20 percent. Our approach also demonstrates better efficiency compared to baseline approaches.

D. Real-world Experiments

In our real-world experiments, we leverage a Franka Panda arm with a parallel gripper and mount a Realsense D415 camera to its end for image capturing. To comprehensively demonstrate the performance of our approach, we design tasks of varying difficulty levels: (1) place objects to the target position (2) place objects to the target rotation (3) place objects to the target position and rotation. We employ diverse objects with different geometries, textures, and materials, including transparent and specular ones. As shown in Fig. 4, our zero-shot method is able to tackle challenging Open6DOR scenarios and demonstrates strong potential in long-horizon tasks.

VI. Conclusion

In this paper, we pioneer the establishment of the Open6DOR benchmark and VLM-based approach, addressing the need for a comprehensive evaluation and a foregoing method exploration in open-instruction 6-DoF object rearrangement. Our synthetic benchmark, comprising over 200 objects and 5400 tasks, offers a standardized framework for evaluating the capabilities of embodied agents in simulation environments. Additionally, our Open6DOR-GPT approach achieves state-of-the-art performance, augmenting GPT-4V with 3D awareness and simulation assistance. As for the current limitations, while Open6DOR-GPT significantly improves position and rotation handling, it does not achieve

real-time performance, and rotation understanding remains suboptimal. We look forward to future improvements to our benchmarks, especially for real-world extensions.

References

[1] Batra, D., Chang, A.X., Chernova, S., Davison, A.J., Deng, J., Koltun, V., Levine, S., Malik, J., Mordatch, I., Mottaghi, R., et al.: Rearrangement: A challenge for embodied ai. arXiv preprint arXiv:2011.01975 (2020) 2

[2] Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Chen, X., Choromanski, K., Ding, T., Driess, D., Dubey, A., Finn, C., Florence, P., Fu, C., Arenas, M.G., Gopalakrishnan, K., Han, K., Hausman, K., Herzog, A., Hsu, J., Ichter, B., Irpan, A., Joshi, N., Julian, R., Kalashnikov, D., Kuang, Y., Leal, I., Lee, L., Lee, T.W.E., Levine, S., Lu, Y., Michalewski, H., Mordatch, I., Pertsch, K., Rao, K., Reymann, K., Ryoo, M., Salazar, G., Sanketi, P., Sermanet, P., Singh, J., Singh, A., Soricut, R., Tran, H., Vanhoucke, V., Vuong, Q., Wahid, A., Welker, S., Wohlhart, P., Wu, J., Xia, F., Xiao, T., Xu, P., Xu, S., Yu, T., Zitkovich, B.: Rt-2: Vision-language-action models transfer web knowledge to robotic control (2023) 2

[3] Brohan, A., Brown, N., Carbajal, J., Chebotar, Y., Dabis, J., Finn, C., Gopalakrishnan, K., Hausman, K., Herzog, A., Hsu, J., et al.: Rt-1: Robotics transformer for real-world control at scale. arXiv preprint arXiv:2212.06817 (2022) 2

[4] Calli, B., Singh, A., Bruce, J., Walsman, A., Konolige, K., Srinivasa, S., Abbeel, P., Dollar, A.M.: Yale-cmu-berkeley dataset for robotic manipulation research. The International Journal of Robotics Research **36**(3), 261–268 (2017) 2

[5] Calli, B., Walsman, A., Singh, A., Srinivasa, S., Abbeel, P., Dollar, A.M.: Benchmarking in manipulation research: Using the yale-cmu-berkeley object and model set. IEEE Robotics amp; Automation Magazine **22**(3), 36–52 (Sep 2015) 3

[6] Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y., et al.: A survey on evaluation of large language models. ACM Transactions on Intelligent Systems and Technology (2023) 2

[7] Deitke, M., Liu, R., Wallingford, M., Ngo, H., Michel, O., Kusupati, A., Fan, A., Laforte, C., Voleti, V., Gadre, S.Y., VanderBilt, E., Kembhavi, A., Vondrick, C., Gkioxari, G., Ehsani, K., Schmidt, L., Farhadi, A.: Objaverse-xl: A universe of 10m+ 3d objects (2023) 3

[8] Driess, D., Xia, F., Sajjadi, M.S.M., Lynch, C., Chowdhery, A., Ichter, B., Wahid, A., Tompson, J., Vuong, Q., Yu, T., Huang, W., Chebotar, Y., Sermanet, P., Duckworth, D., Levine, S., Vanhoucke, V., Hausman, K., Toussaint, M., Greff, K., Zeng, A., Mordatch, I., Florence, P.: Palm-e: An embodied multimodal language model. In: arXiv preprint arXiv:2303.03378 (2023) 1

[9] Ehsani, K., Han, W., Herrasti, A., VanderBilt, E., Weihs, L., Kolve, E., Kembhavi, A., Mottaghi, R.: Manipulathor: A framework for visual object manipulation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4497–4506 (2021) 2

[10] Foundation, B.: Blender. https://www.blender.org/ (2024), accessed: 2024-09-11 3

[11] Garrett, C.R., Chitnis, R., Holladay, R., Kim, B., Silver, T., Kaelbling, L.P., Lozano-Pérez, T.: Integrated task and motion planning. Annual review of control, robotics, and autonomous systems **4**, 265–293 (2021) 2

[12] Garrett, C.R., Lozano-Pérez, T., Kaelbling, L.P.: Pddlstream: Integrating symbolic planners and blackbox samplers via optimistic adaptive planning. In: Proceedings of the International Conference on Automated Planning and Scheduling. vol. 30, pp. 440–448 (2020) 2

[13] Geng, H., Wei, S., Deng, C., Shen, B., Wang, H., Guibas, L.: Sage: Bridging semantic and actionable parts for generalizable articulated-object manipulation under language instructions. arXiv preprint arXiv:2312.01307 (2023) 2

[14] Gu, J., Chaplot, D.S., Su, H., Malik, J.: Multi-skill mobile manipulation for object rearrangement. arXiv preprint arXiv:2209.02778 (2022) 2

[15] Gu, J., Kirmani, S., Wohlhart, P., Lu, Y., Arenas, M.G., Rao, K., Yu, W., Fu, C., Gopalakrishnan, K., Xu, Z., et al.: Rt-trajectory: Robotic task generalization via hindsight trajectory sketches. arXiv preprint arXiv:2311.01977 (2023) 2

[16] Hu, C., Hudson, S., Ethier, M., Al-Sharman, M., Rayside, D., Melek, W.: Sim-to-real domain adaptation for lane detection and classification in autonomous driving. In: 2022 IEEE Intelligent Vehicles Symposium (IV). pp. 457–463. IEEE (2022) 2

[17] Huang, W., Wang, C., Zhang, R., Li, Y., Wu, J., Fei-Fei, L.: Voxposer: Composable 3d value maps for robotic manipulation with language models. arXiv preprint arXiv:2307.05973 (2023) 2, 6

[18] Huang, W., Xia, F., Shah, D., Driess, D., Zeng, A., Lu, Y., Florence, P., Mordatch, I., Levine, S., Hausman, K., Ichter, B.: Grounded decoding: Guiding text generation with grounded models for robot control (2023) 5

[19] Jiang, Y., Gupta, A., Zhang, Z., Wang, G., Dou, Y., Chen, Y., Fei-Fei, L., Anandkumar, A., Zhu, Y., Fan, L.: Vima: General robot manipulation with multimodal prompts. arXiv preprint arXiv:2210.03094 (2022) 2

[20] Kapelyukh, I., Ren, Y., Alzugaray, I., Johns, E.: Dream2real: Zero-shot 3d object rearrangement with vision-language models. arXiv preprint arXiv:2312.04533 (2023) 2, 6, 7

[21] Li, X., Zhang, M., Geng, Y., Geng, H., Long, Y., Shen, Y., Zhang, R., Liu, J., Dong, H.: Manipllm: Embodied multimodal large language model for object-centric robotic manipulation (2023) 2

[22] Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. Advances in neural information processing systems **36** (2024) 2

[23] Liu, Z., Liu, W., Qin, Y., Xiang, F., Gou, M., Xin, S., Roa, M.A., Calli, B., Su, H., Sun, Y., et al.: Ocrtoc: A cloud-based competition and benchmark for robotic grasping and manipulation. IEEE Robotics and Automation Letters **7**(1), 486–493 (2021) 2

[24] Makoviychuk, V., Wawrzyniak, L., Guo, Y., Lu, M., Storey, K., Macklin, M., Hoeller, D., Rudin, N., Allshire, A., Handa, A., State, G.: Isaac gym: High performance gpu-based physics simulation for robot learning (2021), https://arxiv.org/abs/2108.10470 3

[25] Ni, T., Ehsani, K., Weihs, L., Salvador, J.: Towards disturbance-free visual mobile manipulation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5219–5231 (2023) 2

[26] OpenAI: Gpt-4 technical report (2023) 1, 4, 6, 7

[27] Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021) 2

[28] Straub, J., Whelan, T., Ma, L., Chen, Y., Wijmans, E., Green, S., Engel, J.J., Mur-Artal, R., Ren, C., Verma, S., et al.: The replica dataset: A digital replica of indoor spaces. arXiv preprint arXiv:1906.05797 (2019) 2

[29] Sun, C., Orbik, J., Devin, C.M., Yang, B.H., Gupta, A., Berseth, G., Levine, S.: Fully autonomous real-world reinforcement learning with applications to mobile manipulation. In: Conference on Robot Learning. pp. 308–319. PMLR (2022) 2

[30] Sundaralingam, B., Hari, S.K.S., Fishman, A., Garrett, C., Wyk, K.V., Blukis, V., Millane, A., Oleynikova, H., Handa, A., Ramos, F., Ratliff, N., Fox, D.: curobo: Parallelized collision-free minimum-jerk robot motion generation (2023) 3, 6

[31] Wang, C., Fang, H.S., Gou, M., Fang, H., Gao, J., Lu, C.: Graspness discovery in clutters for fast and accurate grasp detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15964–15973 (2021) 6

[32] Wu, H., Jing, Y., Cheang, C., Chen, G., Xu, J., Li, X., Liu, M., Li, H., Kong, T.: Unleashing large-scale video generative pre-training for visual robot manipulation. arXiv preprint arXiv:2312.13139 (2023) 2

[33] Xia, F., Shen, W.B., Li, C., Kasimbeg, P., Tchapmi, M.E., Toshev, A., Martín-Martín, R., Savarese, S.: Interactive gibson benchmark: A benchmark for interactive navigation in cluttered environments. IEEE Robotics and Automation Letters **5**(2), 713–720 (2020) 2

[34] Xu, J., Cheng, W., Gao, Y., Wang, X., Gao, S., Shan, Y.: Instantmesh: Efficient 3d mesh generation from a single image with sparse-view large reconstruction models. arXiv preprint arXiv:2404.07191 (2024) 5

[35] Zhang, J., Gireesh, N., Wang, J., Fang, X., Xu, C., Chen, W., Dai, L., Wang, H.: Gamma: Graspability-aware mobile manipulation policy learning based on online grasping pose fusion. arXiv preprint arXiv:2309.15459 (2023) 2

[36] Zheng, K., Chen, X., Jenkins, O.C., Wang, X.E.: Vlmbench: A compositional benchmark for vision-and-language manipulation (2022) 2