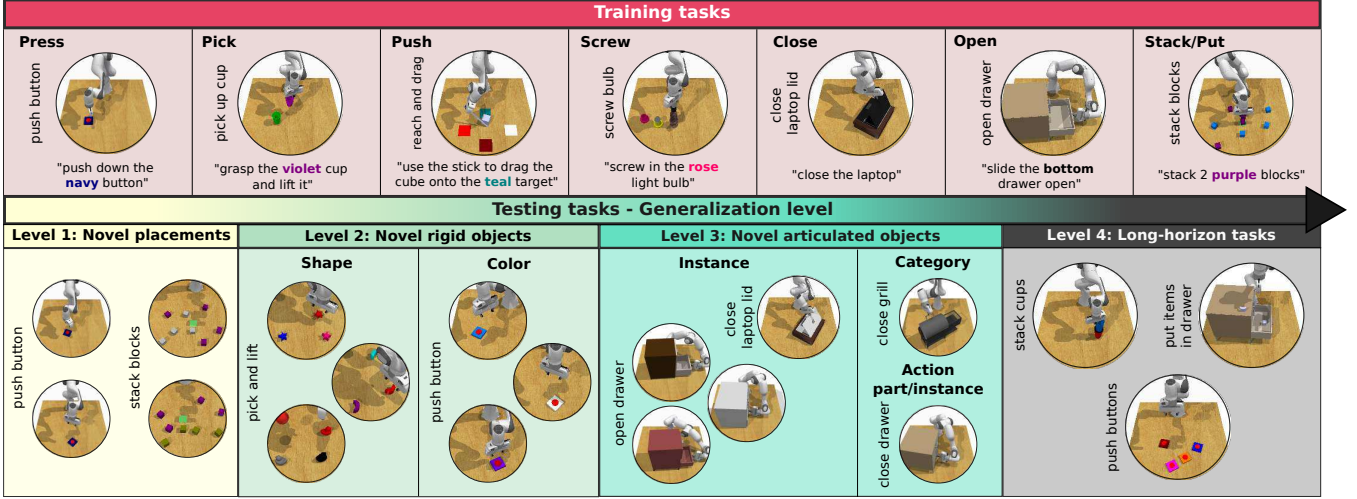


# Towards Generalizable Vision-Language Robotic Manipulation: A Benchmark and LLM-guided 3D Policy

Ricardo Garcia<sup>\*†</sup>, Shizhe Chen<sup>\*†</sup>, Cordelia Schmid<sup>†</sup>



**Fig. 1: GemBench benchmark for vision-language robotic manipulation.** Top: GemBench comprises 16 training tasks with 31 variations, covering seven action primitives. Bottom: The testing set includes 44 tasks with 92 variations, which are organized into four progressively more challenging levels to systematically evaluate generalization capabilities.

**Abstract**—Generalizing language-conditioned robotic policies to new tasks remains a significant challenge, hampered by the lack of suitable simulation benchmarks. In this paper, we address this gap by introducing GemBench, a novel benchmark to assess generalization capabilities of vision-language robotic manipulation policies. GemBench incorporates seven general action primitives and four levels of generalization, spanning novel placements, rigid and articulated objects, and complex long-horizon tasks. We evaluate state-of-the-art approaches on GemBench and also introduce a new method. Our approach 3D-LOTUS leverages rich 3D information for action prediction conditioned on language. While 3D-LOTUS excels in both efficiency and performance on seen tasks, it struggles with novel tasks. To address this, we present 3D-LOTUS++, a framework that integrates 3D-LOTUS’s motion planning capabilities with the task planning capabilities of LLMs and the object grounding accuracy of VLMs. 3D-LOTUS++ achieves state-of-the-art performance on novel tasks of GemBench, setting a new standard for generalization in robotic manipulation. The benchmark, codes and trained models are available at <https://www.di.ens.fr/willow/research/gembench/>.

## I. INTRODUCTION

Vision-language robotic manipulation aims to train policies performing complex tasks based on visual inputs and language instructions [1], [2], [3], [4]. Given the vast diversity of real-world tasks, collecting data for every possible

task is prohibitively expensive. Hence, it is critical to develop models that can effectively generalize to novel tasks.

A number of approaches have been proposed to improve the generalization ability of robotic policies. One prominent direction focuses on pretraining visual and language representations on large-scale web data for robotics [5], [6], [7], [8]. While these representations are more powerful, they do not directly translate into generalizable policies for new tasks. Another approach [1], [9], [10], [11], [12] pretrains entire robotic policies by combining robot data with web data, but these methods remain constrained by the limited size and diversity of available robot datasets [13]. More recently, foundation models such as large language models (LLMs) [14] and vision-language models (VLMs) [15] have been employed to enhance generalization of robot policies. For example, CaP [16] uses LLMs to generate codes that executes primitive actions, and VoxPoser [3] composes value maps with foundation models for action execution with classic motion planning. However, these approaches focus on simpler tasks like pick-and-place and often struggle with more complex manipulation like rotating objects.

Furthermore, there is a notable lack of systematic benchmarks to evaluate the generalization capabilities of models. Though several simulation-based benchmarks have been developed for vision-language robotic manipulation [26], [21], [19], they primarily evaluate models on tasks seen during training, neglecting the crucial aspect of generalization. A

<sup>\*</sup>The authors contributed equally to this work.

<sup>†</sup>Inria, École normale supérieure, CNRS, PSL Research University  
firstname.lastname@inria.fr

**TABLE I: Comparison of benchmarks for vision-and-language robotic manipulation.** *Multi-skill*: covering multiple action primitives beyond pick and place. *atc-obj*: tasks involve interactions with articulated objects. *Test generalization level to attr-obj*: unseen size, color or texture of the object, *act-obj*: unseen action object combination, *inst* or *cate*: same action for unseen object instance or category respectively and *long-horizon*: unseen combination of multiple seen actions and objects.

Benchmark	Simulator	Physics	# Train task(var)	# Test task(var)	Multi-skills	atc-obj	attr-obj	Test generalization level			
								act-obj	inst	cate	long-horizon
RLBench-74Task [17]	RLBench	✓	74 (74)	74 (74)	✓	✓	✗	✗	✗	✗	✗
RLBench-18Task [18]	RLBench	✓	18 (249)	18 (249)	✓	✓	✗	✗	✗	✗	✗
VLMBench [19]	RLBench	✓	8 (233)	8 (374)	✓	✓	✓	✗	✓	✗	✗
ALFRED [20]	AI2-THOR	✗	7 (21,023)	7 (1,529)	✓	✓	✓	✗	✗	✗	✓
Calvin [21]	PyBullet	✓	34	34 (1000)	✓	✓	✓	✗	✗	✗	✓
Ravens [22]	PyBullet	✓	10	10	✗	✗	✓	✗	✓	✓	✗
Arnold [23]	Isaac Sim	✓	8 (3571)	8 (800)	✓	✓	✓	✓	✓	✗	✗
VIMA-Bench [24]	Ravens	✓	13	17	✗	✗	✓	✗	✓	✓	✓
Colosseum [25]	RLBench	✓	20 (280)	20 (20,371)	✓	✓	✓	✗	✗	✗	✗
GemBench (Ours)	RLBench	✓	16 (31)	44 (92)	✓	✓	✓	✓	✓	✓	✓

few exceptions like VIMA-Bench [24] and Colosseum [25] attempt to evaluate generalization abilities. However, VIMA-Bench is limited to relatively simple action skills, while Colosseum focuses more on environment perturbations like lighting changes rather than generalizing to new tasks.

In this work, we introduce **GEMBench** - a new simulation benchmark to evaluate **GENERALIZATION** capabilities for vision-language robotic **MANIPULATION**. As shown in Figure 1, GemBench features two key improvements compared to prior work. First, it incorporates a wider range of complex tasks, involving seven core action skills: *press*, *pick*, *push*, *screw*, *close*, *open* and *put*. Secondly, it introduces a comprehensive suite of four generalization levels with increasing difficulty, focusing on generalization to novel placements, rigid objects, articulated objects, and long-horizon tasks respectively.

To tackle the problem, we first propose a new **3D** robotic manipulation policy - **3D-LOTUS** with **Language-conditioned point cloud transformer**. By leveraging a strong 3D backbone and an effective action representation, 3D-LOTUS achieves state-of-the-art performance on existing vision-language manipulation benchmark [18] and Level 1 of GemBench, while significantly improving training efficiency. Nevertheless, 3D-LOTUS struggles to generalize to new tasks in GemBench, mainly due to limitations in planning for new tasks and grounding new objects. Therefore, we propose an enhanced version **3D-LOTUS++** which integrates foundation models to boost generalization capabilities. Specifically, LLMs are employed for task planning, decomposing tasks into step-by-step actionable plans, and VLMs are used for object grounding which can localize new objects mentioned in the plan. With the grounded object and the primitive action in the plan, 3D-LOTUS serves as the motion controller to generate action trajectories. Experimental results demonstrate that 3D-LOTUS++ significantly improves generalization, outperforming 3D-LOTUS on Levels 2 to 4 of GemBench. To summarize, our contributions are three fold:

- We introduce a new benchmark GemBench to systematically evaluate generalization in vision-language robotic manipulation across four complexity levels.
- We propose an effective manipulation policy 3D-LOTUS

and enhance its generalization ability with foundation models for task planning and object grounding (3D-LOTUS++).

- Our models establish state of the arts on existing benchmark and GemBench, and also work reliably on a real robot.

## II. RELATED WORK

**Robotic manipulation benchmark.** Significant progress has been made in the development of robot simulators such as RLBench [26], AI2-THOR [27] and Isaac Sim [28]. Leveraging these simulators, various benchmarks have emerged for robotic manipulation. Early benchmarks [17], [18] train and test policies on the same task set, overlooking the critical aspect of generalization to unseen scenarios. To address this, more recent benchmarks [19], [20], [21], [22], [23], [25], [24] have introduced generalization evaluations on new compositions of objects and colors, new object shapes, or even long-horizon tasks. Among them, VIMA-Bench [24] and Colosseum [25] are most similar to our work, aiming to systematically evaluate different generalization abilities. However, VIMA-Bench [24] is limited to pick-and-place tasks using a suction gripper, while Colosseum [25] emphasizes generalization to environment perturbations such as changes in lighting and camera angles. In contrast, GemBench covers more complex action skills and evaluate generalization to entirely new tasks rather than perturbations of the seen tasks. Table I provides a comprehensive comparison of these benchmarks. We can observe that GemBench is most general among all of them.

**Vision-and-language robotic manipulation.** Learning robotic manipulation conditioned on vision and language has received increasing attention [29], [30], [31]. The high dimensionality of manipulation action space makes it challenging to directly use reinforcement learning (RL) in training [32]. Therefore, most works rely on imitation learning (IL) [33], [1], [17], [18], [34], [2], [8], [4], [35] using scripted trajectories [26] or tele-operation data [13]. Visual representation plays a crucial role in policy learning. Existing works [33], [1], [36], [17], [34], [37] mainly use 2D images to predict actions, though recent works have begun exploring 3D representations [38], [18], [2], [4], [35], [8]. In this work,

we take advantage of rich spatial information of 3D point cloud for motion planning, while leaving object grounding in 2D to benefit from the generalization strength of pretrained 2D models [15], [39], [40].

**Foundation models for robotics.** Learning-based robotic policies often struggle to generalize to new scenarios [41]. Inspired by the remarkable generalization capabilities of foundation models [39], [15], [42], recent work explores how to leverage these models for planning, perception and control in robotics. Huang *et al.* [43] use LLMs to decompose high-level tasks into sub-steps. To ground plans in the visual world, SayCan [44] combine LLMs with value functions of pretrained skills. ViLa [45] replaces LLMs with a multi-modal LLM GPT-4V [46]. CaP [16] instructs LLMs to write code which call tools for perception and control. However, these approaches rely on predefined motion skills, limiting applicability to broader tasks. To address this, VoxPoser [3] use LLMs to construct 3D voxel maps of affordance, constraint, rotation and velocity, which are fed into traditional motion planning algorithms to plan a trajectory. Nevertheless, VoxPoser only provides a coarse-grained understanding of the scene and struggles with precise robot control. In this work, we propose to combine the generalization ability of foundation models with strong motion control capabilities of 3D policies for robotic manipulation.

### III. GEMBENCH: GENERALIZABLE VISION-LANGUAGE ROBOTIC MANIPULATION BENCHMARK

This paper introduces the GemBench benchmark to systematically evaluate generalization capabilities of vision-and-language robotic manipulation policies. It is built upon the RL Bench simulator [26], which provides a wide range of visually and physically realistic tasks together with a framework to generate scripted demonstrations. In the following, we first describe training tasks in GemBench in Sec III-A and then present four levels of generalization for evaluation in Sec III-B. The details of the proposed GemBench are presented in Sec VI-A in the appendix.

#### A. Training tasks

We select 16 tasks (31 variations<sup>1</sup>) from existing RL Bench benchmarks [26], [18] to capture a diverse range of action primitives beyond simple pick-and-place. These tasks, shown in Figure 1 (top), include seven action primitives: press, pick, push, screw, close, open, and stack/put. Examples of training tasks are *push button*, *pick up cup*, *reach and drag cube*, *screw light bulb in*, *close laptop lid*, *open drawer*, *stack blocks*. Task variations cover 20 objects (e.g., cube, cup, fridge), 20 colors (e.g., red, blue, violet), and 3 object parts (e.g., top, middle, bottom). The training set is sufficiently diverse and should enable a robot to generalize to new tasks, such as novel attribute-object compositions, new action-object pairings, or even entirely new shapes.

<sup>1</sup>The RL Bench tasks typically involve manipulating fixed-shape objects, such as pressing a button, with variations including different object colors (e.g., red or blue button), parts (e.g., top or middle drawer), and manipulation sequences (e.g., press a red button followed by a green one).

#### B. Testing tasks with four levels of generalization

As shown in Figure 1 (bottom), the test set includes four levels of generalization that progressively increase the difficulty of vision-language robotic manipulation. The test set consists of a total of 44 tasks (92 variations), with 23 tasks selected from the original RL Bench 100 tasks [26] and 21 newly scripted tasks. The test levels differ in object shape and color, object articulation and horizon of the tasks.

**Level 1 - Novel placements:** This level consists of the same 16 tasks (31 variations) as the training set, but with new object placements randomly sampled within the robot’s workspace. Additionally, some tasks feature new distractor objects with different colors. The objective is to evaluate whether a policy can perform well on seen tasks with minor configuration changes.

**Level 2 - Novel rigid objects:** This level comprises 15 unseen tasks (28 variations) where the robot interacts with novel rigid objects using actions such as press, pick, and put. There are two categories for generalization to rigid objects: 1) Novel object-color compositions. For instance, training tasks only manipulate yellow button and rose bulb, while the test task requires to operate a rose button. This level includes 20 new object-color compositions. 2) Novel object shapes. For example, picking a cube is learned in training, but the testing task is lifting a toy or a star-shaped item. There are 8 new object shapes in the evaluation.

**Level 3 - Novel articulated objects:** This set includes 18 new tasks (21 variations) where the robot interacts with articulated objects. Three categories are proposed: 1) Novel action-part compositions. For example, if trained to open bottom drawer and put item in middle shelf, it now needs to open middle drawer. There are 8 novel compositions. 2) Novel instances. For example, after being trained on a three-drawer unit, it must generalize to one with four drawers. This set includes 11 new object instances. 3) Novel categories. For example, it is trained to close a laptop lid but must generalize to closing a grill lid. This includes 2 new object categories.

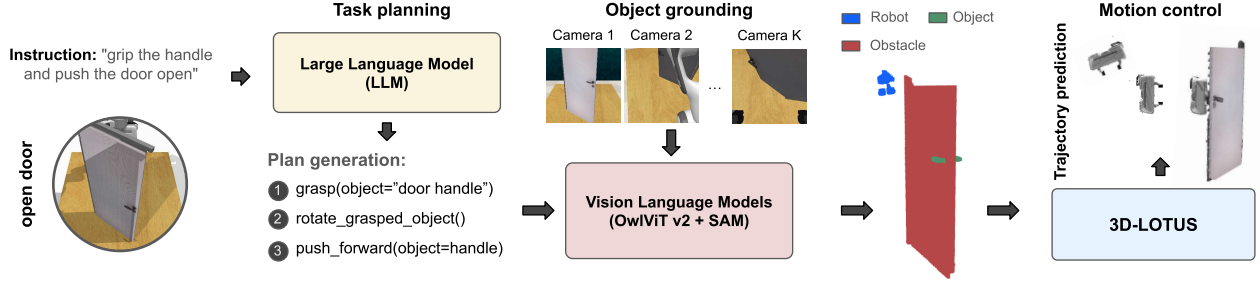
**Level 4 - Novel Long-horizon tasks:** This level presents the greatest challenge as it requires the robot to combine multiple actions learned during training. It includes 6 long-horizon tasks (12 variations). For example, the task “put items in drawer” involves a sequence of actions such as opening the drawer, picking up a sequence of items (a cube, a cylinder and a moon), and placing them inside according to the order specified by the variation instruction.

## IV. METHOD

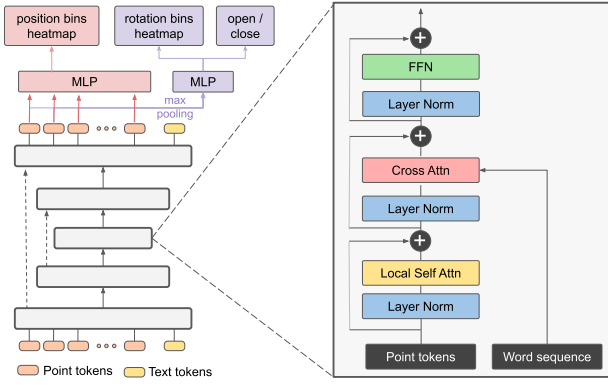
#### A. Problem formulation

The goal is to learn a policy  $\pi(a_t|O_t, L)$  for robotic manipulation, where  $L$  is a language instruction,  $O_t \in \mathbb{O}$ ,  $a_t \in \mathbb{A}$  are the observation and action at step  $t$  respectively, with  $\mathbb{O}$  and  $\mathbb{A}$  denoting the observation and action space.

The observation space  $\mathbb{O}$  includes aligned RGB-D images from  $K$  cameras, along with the robot’s proprioceptive state consisting of joint and gripper poses. We assume the intrinsic and extrinsic camera parameters are known.



**Fig. 2: Overview of 3D-LOTUS++ framework.** It leverages generalization capabilities of foundation models for planning and perception, and strong action execution ability of 3D-LOTUS to perform complex tasks.



**Fig. 3: 3D-LOTUS architecture.** It takes point cloud and text as input to predict the next action.

The action space  $\mathbb{A}$  comprises the gripper’s position  $a_t^p \in \mathbb{R}^3$ , rotation  $a_t^r \in \mathbb{R}^3$ , and open state  $a_t^o \in \{0, 1\}$  indicating if the gripper is open or closed. We utilize waypoint representation [38], [47], [17], [18] for action sequences. Inverse kinematics is used to move the robot from  $a_{t-1}$  to  $a_t$ .

### B. 3D-LOTUS policy

As 3D point clouds provide rich spatial information to perceive object shapes and positions, we propose to effectively exploit the 3D information for robotic manipulation. The 3D-LOTUS policy is a **3D** robotic manipulation policy with language-conditioned point cloud transformer. Figure 3 illustrates the architecture of the policy.

**Point cloud preprocessing.** We follow PolarNet [2] to project multi-view RGB-D images into a unified point cloud in world coordinates, then downsample this point cloud to one point per  $1\text{cm}^3$  voxel [48]. We exclude points outside the robot’s workspace, and points on the robotic arm using a CAD model and joint poses of the arm, as these contribute little for manipulation. The resulting point cloud  $V$  of  $n$  points only covers the objects and robot gripper, significantly reducing the number of points and improving speed without compromising performance. Each point  $v_i \in V$  consists of XYZ coordinates  $v_i^p$  and additional feature  $v_i^o$  such as RGB color and relative height to the table.

**Language-conditioned point cloud transformer.** We employ point cloud transformer v3 (PTV3) [49] as backbone to

encode point cloud  $V$ . PTV3 adopts a U-Net [50] architecture with downsampling and upsampling blocks to efficiently compute point embeddings, where each block consists of transformer layers. For more details, please refer to the PTV3 paper [49]. We explore two variants to incorporate language information into the PTV3 model. Assuming the language instruction  $L$  is encoded by a CLIP text encoder [39] and represented as a sequence of word embeddings  $(w_1, \dots, w_L)$ . The first variant employs the *adaptive normalization approach* [51]. We compute a global language embedding  $\bar{w}$  by weighted averaging  $(w_1, \dots, w_L)$ , and use  $\bar{w}$  to directly regress dimension-wise scale and shift parameters for each normalization layer in PTV3. The second variant utilizes a more conventional *cross-attention mechanism* [52]. A cross-attention layer is added after each self-attention layer in PTV3, enabling each point to attend to the entire sequence of word embeddings  $(w_1, \dots, w_L)$ .

**Action prediction.** Let  $v_i^e$  denote the final point embedding for each point  $v_i$  after the language-conditioned PTV3 model. We propose a new classification-based approach for action prediction, in contrast to the regression-based approach [17], [2], [8] or inefficient position classification over the whole 3D workspace [18], [4]. For position prediction, we predict the gripper’s location along the  $X, Y, Z$  axes separately. We define sequential bins  $v_{i,k,j}$  centered at each point’s location  $v_i^p$  for each axis  $k \in \{X, Y, Z\}$ , with  $j \in [-m, m]$  representing the bin index. Each bin has a size of  $b$ , so the position of bin  $v_{i,k,j}$  along  $k$ -axis is given by  $v_{i,k}^p + b \times j$ . Using  $v_i^e$ , we predict a heatmap for these bins at each point, and concatenate the bins across all points to form the final heatmap for each axis. During inference, we select the bin with the highest probability to determine the position for each axis. For rotation prediction, we also discretize the Euler angle for each axis into bins and use classification to predict the angles. The open state prediction remains a binary classification task. We apply max pooling over all points  $v_i^e$  to predict rotation and open state. The cross entropy loss is employed to train position, rotation and open state classification. Further details are available in Sec VI-B in the appendix.

### C. 3D-LOTUS++ policy

To accomplish a task specified by an instruction  $L$  like ‘open the door’, the end-to-end policy 3D-LOTUS integrates

multiple components into a single action prediction step. This includes task planning (e.g., first grasp the door handle), object grounding (e.g., localize the door handle in 3D space), and motion control (move to the localized door handle). This integration not only complicates error diagnosis, but also poses challenges in generalization to unseen scenarios, such as ‘open a new door’ or ‘close the door’.

To alleviate the above limitations, we propose enhancing 3D-LOTUS with foundation models. Existing LLMs [14] and VLMs [15], [39], [40], are able to generalize to unseen scenarios due to the training on massive data. Therefore, we introduce a modular framework that disentangles task planning, object grounding and motion control, leveraging the generalization capabilities of foundation models alongside the action execution abilities of 3D-LOTUS to achieve more generalizable robotic manipulation. Figure 2 illustrates the overall framework, comprising three modules: task planning with LLMs, object grounding with VLMs, and motion control with a modified version of 3D-LOTUS.

**Task planning with LLM.** Task planning aims to decompose the instruction  $L$  into a sequence of steps  $l_1, \dots, l_T$ . Each step corresponds to an action primitive that interacts with an object. In this work, we define six action primitives for object manipulation, covering a broad range of tasks, namely *grasp(object)*, *move-grasped-object(target)*, *push-down(object)*, *push-forward(object, target)*, *release()* and *rotate-grasped-object()*. We utilize the LLM LLaMa3-8B [14] for task planning due to its strong commonsense knowledge and language reasoning capabilities. By providing prompts for the task requirement and several in-context examples, we guide the LLM to generate an plan for instruction  $L$ . Figure 2 presents an example of generated plans for the task of opening a door. The detailed prompts for LLMs are presented in Sec VI-C in the appendix.

**Object grounding with VLMs.** This module aims to localize an object given its text description in the generated plan. To achieve this, we leverage state-of-the-art VLMs to ensure robust generalization to new objects. First, we employ the open-vocabulary object detector OWLv2 [40] to detect bounding boxes with high objectiveness scores for each RGB image. OWLv2 also generates a semantic embedding for each bounding box, which is aligned with text embeddings from the CLIP text encoder [39]. Next, we use the Segment Anything Model (SAM) [15] to segment the object within each bounding box. This segmentation mask, combined with the corresponding RGB-D image, yields a 3D point cloud for each bounding box. To merge observations of the same object from different cameras, we compare semantic embeddings and point cloud distances. Pairs of objects are merged if their semantic and point cloud distances are below certain thresholds. In this way, we obtain object-centric representations for all objects in the scene, each object containing a merged point cloud and an averaged semantic embedding. Given the text description of an object, we compute its text embedding via CLIP and measure cosine similarities between this text embedding and all object semantic embeddings. The object with the highest cosine similarity is selected as the match.

**Motion control.** Given the action primitive name and input point cloud, the motion control module predicts a trajectory of actions to execute this actionable step. 3D-LOTUS can be easily modified for this purpose.

First, we change the input point feature  $v_i^o$  by leveraging the output of the object grounding module, which segments the manipulated object and/or target location. This allows us to categorize points into four types: goal object, goal target, robot and obstacle. We treat points that do not belong to goal object, target and robot as obstacles. We then learn a look-up table to encode each point label and use this as point feature  $v_i^o$  instead of RGB colors in addition to the XYZ coordinates  $v_i^p$ . The new point feature help the model focus on geometry rather than textures during motion planning, thereby enhancing generalization to objects with novel textures.

Second, instead of predicting a single action, we should generate a sequence of actions to complete the planned step. To achieve this, we introduce a look-up table to encode the timestep index of actions in the trajectory, denoted as  $\{x_t\}_{t=1}^s$ . We concatenate the time embedding  $x_t$  with the final point embedding  $v_i^o$  to predict action for each timestep. The action prediction head is shared across all timesteps. As the number of actions varies for different plans, we also predict a stop probability to indicate whether the trajectory should terminate at the current time step.

## V. EXPERIMENTS

### A. Experimental setup

**Evaluation setup.** We follow prior work [18], [2] and use  $K = 4$  cameras positioned at the front, left shoulder, right shoulder and wrist, with an image resolution of  $256 \times 256$ . For training, we generate 100 demonstrations for each task variation, leading to a dataset of 3,100 demonstrations. During testing, we use different random seeds from training to ensure initial scene configurations are distinct from the training data. We evaluate 20 episodes per task variation per seed, and run the evaluation with 5 seeds, which results in  $20 \times 5 \times 92$  evaluation episodes in total. The maximum number of steps per episode is set to 25. We measure the task performance by success rate (SR) of the evaluation episodes, which is 1 for success and 0 for failure of an episode. The average SR and standard derivation across seeds are reported.

**Implementation details.** For the 3D-LOTUS model, we use 5 downsampling-upsampling blocks, each containing 1 transformer layer. The hidden sizes for these blocks are 64, 128, 256, 512, 768, respectively. For action prediction, the number of bins for position is 30 ( $m = 15$ ) and bin size  $b = 1\text{cm}$ . The number of bins for rotation is 72 with bin size of  $5^\circ$ . In the modified 3D-LOTUS for trajectory prediction, the maximum trajectory length  $s$  is set as 5. We train 3D-LOTUS with batch size of 8 and initial learning rate of  $1e-4$  for 150k iterations with linear learning rate decay. Training takes around 11 hours on a single Nvidia A100 GPU. A validation set of 20 episodes per task variation on Level 1 (different from testing episodes) is used to select the best checkpoint, evaluated every 10k iterations.

**TABLE II:** Performance on RL-Bench-18Task. The Avg. Rank denotes the averaged rank of the model across tasks. Training time is the number of V100 GPU days for training.

	Avg. SR $\uparrow$	Avg. Rank $\downarrow$	Train time $\downarrow$
C2F-ARM-BC [38]	20.1	8.6	-
Hiveformer [17]	45.3	6.9	-
PolarNet [2]	46.4	6.4	8.9
PerAct [18]	49.4	6.2	128.0
RVT [34]	62.9	4.4	8.0
Act3D [4]	65.0	4.3	40.0
RVT2 [37]	81.4	2.4	6.6
3D diffuser actor [35]	81.3	2.3	67.6
3D-LOTUS	<b>83.1<math>\pm</math>0.8</b>	<b>2.2</b>	<b>2.2<sup>3</sup></b>

**TABLE III:** Performance on four levels of GemBench.

Method	L1	L2	L3	L4
Hiveformer [17]	60.3 $\pm$ 1.5	26.1 $\pm$ 1.4	35.1 $\pm$ 1.7	0.0 $\pm$ 0.0
PolarNet [2]	77.7 $\pm$ 0.9	37.1 $\pm$ 1.4	38.5 $\pm$ 1.7	0.1 $\pm$ 0.2
3D diffuser actor [35]	91.9 $\pm$ 0.8	43.4 $\pm$ 2.8	37.0 $\pm$ 2.2	0.0 $\pm$ 0.0
RVT-2 [37]	89.1 $\pm$ 0.8	51.0 $\pm$ 2.3	36.0 $\pm$ 2.2	0.0 $\pm$ 0.0
3D-LOTUS	<b>94.3<math>\pm</math>1.4</b>	49.9 $\pm$ 2.2	38.1 $\pm$ 1.1	0.3 $\pm$ 0.3
3D-LOTUS++	68.7 $\pm$ 0.6	<b>64.5<math>\pm</math>0.9</b>	<b>41.5<math>\pm</math>1.8</b>	<b>17.4<math>\pm</math>0.4</b>

**Baselines.** We run four state-of-the-art methods on GemBench, including two 2D image based models<sup>2</sup> (Hiveformer [17] and RVT-2 [37]), and two 3D-based models (PolarNet [2] and 3D diffuser actor [35]). All the baselines use CLIP text encoder [39], while only 3D diffuser actor employs visual representations pretrained on large-scale datasets. We use official codes provided by the authors to validate the training pipeline on RL-Bench-18Task benchmark. After reproducing the results on RL-Bench-18Task, we apply the same configuration to train on our GemBench benchmark.

### B. Comparison with state of the arts

**RL-Bench-18Task.** We first evaluate on the widely used RL-Bench-18Task benchmark [18], which contains the same 18 tasks (249 variations) for training and testing. The results are summarized in Table II, with a breakdown of performance on individual tasks provided in Table IX in the appendix. The 3D-LOTUS policy achieves state-of-the-art performance using significantly less training time, demonstrating strong action execution capability.

**GemBench.** Table III presents the results of different models across the four generalization levels in GemBench. Detailed results on individual tasks are shown in Sec VI-D in the appendix. As expected, Level 1 which only involves novel object placements, is easiest. The performance trend of different models are similar to those in RL-Bench-18Task. In Levels 2 to 4, we observe a significant drop in performance for the state-of-the-art methods, highlighting the limitations of existing methods in unseen generalization. Generalizing skills for articulated objects (Level 3) proves to be more

<sup>2</sup>To be noted, the two methods utilize RGB images as input to the model, but depth images are still used in post-processing to predict actions.

<sup>3</sup>We use 1 V100 GPU for training for fair comparison of training time.

**TABLE IV:** Ablation of 3D-LOTUS components.

Action	Condition	L1	L2	L3	L4
Regression	AdaptiveNorm	83.3 $\pm$ 0.7	29.3 $\pm$ 1.9	34.5 $\pm$ 1.0	0.0 $\pm$ 0.0
Classification	AdaptiveNorm	90.8 $\pm$ 0.7	47.8 $\pm$ 0.6	37.9 $\pm$ 1.5	0.0 $\pm$ 0.0
Classification	CrossAttn	<b>94.3<math>\pm</math>1.4</b>	<b>49.9<math>\pm</math>2.2</b>	<b>38.1<math>\pm</math>1.1</b>	<b>0.3<math>\pm</math>0.3</b>

**TABLE V:** Ablations on 3D-LOTUS++ modules.

Task	Object Grounding	L1	L2	L3	L4
GT	GT	<b>92.6<math>\pm</math>0.7</b>	<b>80.1<math>\pm</math>0.5</b>	<b>47.8<math>\pm</math>1.4</b>	<b>31.5<math>\pm</math>1.1</b>
GT	VLM	71.0 $\pm$ 1.7	66.3 $\pm$ 0.9	46.0 $\pm$ 1.5	19.4 $\pm$ 1.5
LLM	VLM	68.7 $\pm$ 0.6	64.5 $\pm$ 0.9	41.5 $\pm$ 1.8	17.4 $\pm$ 0.4

challenging than for rigid objects (Level 2). Level 4, which features long-horizon tasks, is most difficult; the performance of all state-of-the-art methods drop to close to a 0% success rate. 3D-LOTUS++ significantly outperforms previous methods on more challenging generalization levels. Note that its performance is lower than the SOTA methods on Level 1. This can be explained by the zero-shot grounding models, which struggles to distinguish some objects in seen tasks such as ‘tuna can’ and ‘soup can’. The performance on Level 4 is suboptimal. Refer to the ablation study for an analysis.

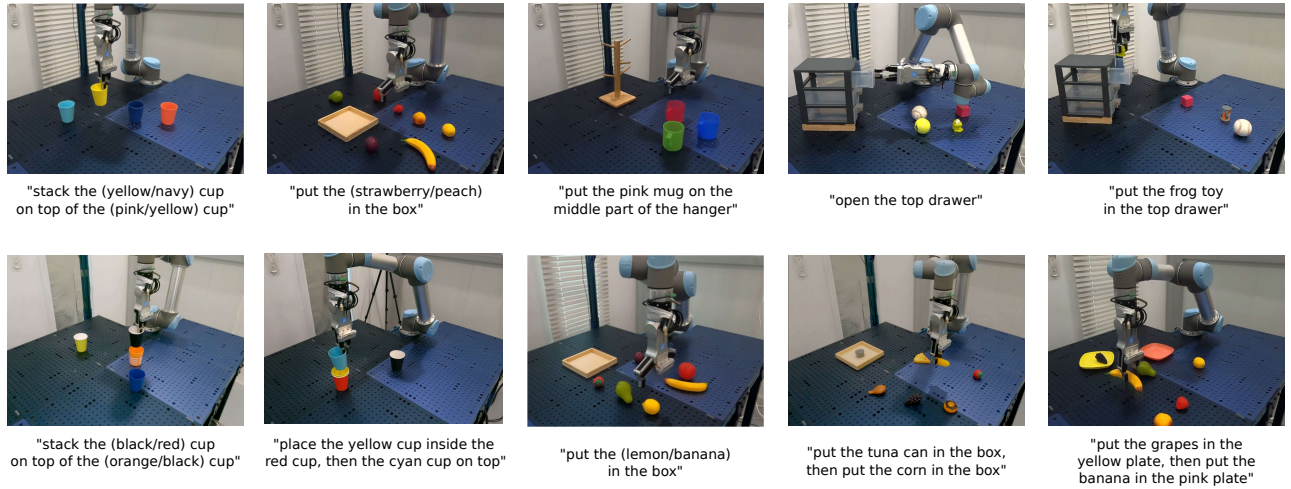
### C. Ablations

**3D-LOTUS components.** In Table IV, we ablate different components of 3D-LOTUS. Row 1 uses regression for action prediction. Its performance is worse than classification in Row 2 for all levels. Furthermore, it requires more iterations to converge. Row 2 and 3 compare two variants for language conditioning. The cross attention outperforms adaptive normalization method, though at a higher computation cost.

**3D-LOTUS++ modules.** The proposed 3D-LOTUS++ allows for detailed error analysis by isolating each module — task planning, object grounding, and motion control. To facilitate this, we manually annotate groundtruth task plans and object grounding labels for each task, and evaluate the model’s performance with and without the ground truth information. The results are shown in Table V. We can see that the primary bottleneck in Levels 1 and 2 is the object grounding module, where groundtruth object labels improve the performance by a large margin. In Levels 3 and 4, however, even when provided with all groundtruth information, the performance remains suboptimal. The primary issue lies within the 3D-LOTUS motion control policy, specifically its struggle to generalize to long-horizon tasks where initial robot configurations deviate substantially from the training data. Task planning in Level 4 also suffers from reduced accuracy as the LLM operates without visual input. This lack of visual awareness can lead to incorrect assumptions about the environment. For example, for the task ‘take shoes out of the box,’ the LLM cannot determine if the box is open or closed, potentially leading to an inefficient or failed plan.”

### D. Real world experiments

We further perform real world evaluations of our models.



**Fig. 4: Real robot tasks variations.** The top row illustrates task variations used for model training. The bottom row presents new task variations to assess model’s generalization capabilities on the real robot.

**TABLE VI: Performance of seen tasks with real robot.**

Task	PolarNet	3D-LOTUS
Stack yellow cup in pink cup	<b>10/10</b>	9/10
Stack navy cup in yellow cup	9/10	<b>10/10</b>
Put strawberry in box	7/10	<b>10/10</b>
Put peach in box	<b>8/10</b>	<b>8/10</b>
Open drawer	6/10	<b>9/10</b>
Put item in drawer	1/10	<b>3/10</b>
Hang mug	6/10	<b>8/10</b>
Avg.	6.7/10	<b>8.1/10</b>

**TABLE VII: Performance of unseen tasks with real robot.**

Task	3D-LOTUS	3D-LOTUS++
Stack red cup in yellow cup	0/10	<b>8/10</b>
Stack black cup in orange cup	0/10	<b>7/10</b>
Place the yellow cup inside the red cup, then the cyan cup on top	0/10	<b>7/10</b>
Put lemon in box	0/10	<b>9/10</b>
Put banana in box	0/10	<b>7/10</b>
Put tuna can in box, then corn in box	0/10	<b>8/10</b>
Put grapes in yellow plate, then banana in pink plate	0/10	<b>9/10</b>
Avg.	0/10	<b>7.9/10</b>

**Experimental setup.** Our real robot setup includes a 6-DoF UR5 robotic arm equipped with three RealSense d435 cameras. We consider 7 variations across 5 tasks during training: stack cup (yellow in pink or navy in yellow), put fruit (strawberry or peach) in box, open drawer, put item in drawer and hang mug. For each task variation, we collect 20 human demonstrations via tele-operation. Then, we evaluate on the same 7 seen task variations with different objects placements and evaluate generalization capabilities on 7 new unseen task variations: put fruit (lemon and banana) in box, put food (tuna can then corn) in box and put fruits in plates (grapes in the yellow plate and banana in the pink plate). These tasks are illustrated in Figure 4. For each task variation we run models 10 times and report the success rate.

Table VI shows that our new method, 3D-LOTUS, outperforms PolarNet, achieving an average success rate of 8.1/10 compared to PolarNet’s average performance of 6.7/10. However, when applying the same 3D-LOTUS model to new task variations, we observe a complete failure to generalize to the new objects and instructions as shown in Table VII. In contrast, our improved model, 3D-LOTUS++, successfully addresses these new task variations, achieving an average success rate of 7.9/10.

## VI. CONCLUSION

In this work, we introduce a new benchmark and method for generalizable vision-language robotic manipulation. The proposed benchmark GemBench systematically evaluates four generalization levels: new placements, new rigid objects, new articulated objects, and long-horizon tasks. To improve generalization ability, we introduce 3D-LOTUS++, a modular framework that leverages foundation models for task planning and object grounding alongside a strong 3D-based motion plan policy 3D-LOTUS. Extensive experiments demonstrate the effectiveness of 3D-LOTUS++ on novel tasks. Our ablation studies highlight object grounding as a critical bottleneck and reveal the limitations of the motion control policy in complex scenarios. Future work will focus on addressing these two issues.

**Acknowledgements.** This work was partially supported by the HPC resources from GENCI-IDRIS (Grant 20XX-AD011012122 and AD011014846). It was funded in part by the French government under management of Agence Nationale de la Recherche as part of the “France 2030” program, reference ANR-23-IACL-0008 (PR[AI]RIE-PSAI projet), and the ANR project VideoPredict (ANR-21-FAI1-0002-01). Cordelia Schmid would like to acknowledge the support by the Körber European Science Prize.

## REFERENCES

- [1] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, J. Dabis, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, J. Hsu *et al.*, “RT-1: Robotics transformer for real-world control at scale,” in *RSS*, 2023.
- [2] S. Chen, R. Garcia, C. Schmid, and I. Laptev, “PolarNet: 3D point clouds for language-guided robotic manipulation,” in *CoRL*, 2023.
- [3] W. Huang, C. Wang, R. Zhang, Y. Li, J. Wu, and L. Fei-Fei, “VoxPoser: Composable 3D value maps for robotic manipulation with language models,” in *CoRL*, 2023.
- [4] T. Gervet, Z. Xian, N. Gkanatsios, and K. Fragkiadaki, “Act3D: 3D feature field transformers for multi-task robotic manipulation,” in *CoRL*, 2023.
- [5] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta, “R3M: A universal visual representation for robot manipulation,” in *CoRL*, 2022.
- [6] I. Radosavovic, T. Xiao, S. James, P. Abbeel, J. Malik, and T. Darrell, “Real-world robot learning with masked visual pre-training,” in *CoRL*, 2023.
- [7] S. Karamcheti, S. Nair, A. S. Chen, T. Kollar, C. Finn, D. Sadigh, and P. Liang, “Language-driven representation learning for robotics,” 2019.
- [8] S. Chen, R. Garcia, I. Laptev, and C. Schmid, “SUGAR: Pre-training 3D visual representations for robotics,” in *CVPR*, 2024.
- [9] A. Brohan, N. Brown, J. Carbajal, Y. Chebotar, X. Chen, K. Choro-manski, T. Ding, D. Driess, A. Dubey, C. Finn *et al.*, “RT-2: Vision-language-action models transfer web knowledge to robotic control,” in *CoRL*, 2023.
- [10] D. Driess, F. Xia, M. S. Sajjadi, C. Lynch, A. Chowdhery, B. Ichter, A. Wahid, J. Tompson, Q. Vuong, T. Yu *et al.*, “PALM-E: An embodied multimodal language model,” in *ICML*, 2023.
- [11] O. M. Team, D. Ghosh, H. Walke, K. Pertsch, K. Black, O. Mees, S. Dasari, J. Hejna, T. Kreiman, C. Xu *et al.*, “Octo: An open-source generalist robot policy,” 2024.
- [12] M. J. Kim, K. Pertsch, S. Karamcheti, T. Xiao, A. Balakrishna, S. Nair, R. Rafailov, E. Foster, G. Lam, P. Sanketi *et al.*, “OpenVLA: An open-source vision-language-action model,” *CoRL*, 2024.
- [13] Q. Vuong, S. Levine, H. R. Walke, K. Pertsch, A. Singh, R. Doshi, C. Xu, J. Luo, L. Tan, D. Shah *et al.*, “Open X-Embodiment: Robotic learning datasets and RT-X models,” in *CoRL*, 2023.
- [14] Meta, “Llama 3 model card,” 2024.
- [15] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, P. Dollár, and R. Girshick, “Segment anything,” in *ICCV*, 2023.
- [16] J. Liang, W. Huang, F. Xia, P. Xu, K. Hausman, B. Ichter, P. Florence, and A. Zeng, “Code as policies: Language model programs for embodied control,” in *ICRA*, 2023.
- [17] P.-L. Guhur, S. Chen, R. Garcia Pinel, M. Tapaswi, I. Laptev, and C. Schmid, “Instruction-driven history-aware policies for robotic manipulations,” in *CoRL*, 2023.
- [18] M. Shridhar, L. Manuelli, and D. Fox, “Perceiver-actor: A multi-task transformer for robotic manipulation,” in *CoRL*, 2023.
- [19] K. Zheng, X. Chen, O. C. Jenkins, and X. Wang, “VLMbench: A compositional benchmark for vision-and-language manipulation,” in *NeurIPS*, 2022.
- [20] M. Shridhar, J. Thomason, D. Gordon, Y. Bisk, W. Han, R. Mottaghi, L. Zettlemoyer, and D. Fox, “ALFRED: A benchmark for interpreting grounded instructions for everyday tasks,” in *CVPR*, 2020.
- [21] O. Mees, L. Hermann, E. Rosete-Beas, and W. Burgard, “CALVIN: A benchmark for language-conditioned policy learning for long-horizon robot manipulation tasks,” *IEEE RA-L*, 2022.
- [22] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani, and J. Lee, “Transporter networks: Rearranging the visual world for robotic manipulation,” in *CoRL*, 2020.
- [23] R. Gong, J. Huang, Y. Zhao, H. Geng, X. Gao, Q. Wu, W. Ai, Z. Zhou, D. Terzopoulos, S.-C. Zhu *et al.*, “ARNOLD: A benchmark for language-grounded task learning with continuous states in realistic 3d scenes,” in *CVPR*, 2023.
- [24] Y. Jiang, A. Gupta, Z. Zhang, G. Wang, Y. Dou, Y. Chen, L. Fei-Fei, A. Anandkumar, Y. Zhu, and L. Fan, “VIMA: robot manipulation with multimodal prompts,” in *ICML*, 2023.
- [25] W. Pumacay, I. Singh, J. Duan, R. Krishna, J. Thomason, and D. Fox, “The COLOSSEUM: A benchmark for evaluating generalization for robotic manipulation,” *arXiv preprint arXiv:2402.08191*, 2024.
- [26] S. James, Z. Ma, D. R. Arrojo, and A. J. Davison, “RLBench: The robot learning benchmark & learning environment,” *IEEE RA-L*, 2020.
- [27] E. Kolve, R. Mottaghi, W. Han, E. VanderBilt, L. Weihs, A. Herrasti, M. Deitke, K. Ehsani, D. Gordon, Y. Zhu *et al.*, “AI2-Thor: An interactive 3d environment for visual ai,” *arXiv:1712.05474*, 2017.
- [28] V. Makovychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa, and G. State, “Isaac Gym: High performance gpu based physics simulation for robot learning,” in *NeurIPS Datasets and Benchmarks*, 2021.
- [29] L. Shao, T. Migimatsu, Q. Zhang, K. Yang, and J. Bohg, “Concept2Robot: Learning manipulation concepts from instructions and human demonstrations,” *IJRR*, 2021.
- [30] C. Lynch, A. Wahid, J. Tompson, T. Ding, J. Betker, R. Baruch, T. Armstrong, and P. Florence, “Interactive language: Talking to robots in real time,” *IEEE RA-L*, 2023.
- [31] S. Stepputtis, J. Campbell, M. Phielipp, S. Lee, C. Baral, and H. Ben Amor, “Language-conditioned imitation learning for robot manipulation tasks,” in *NeurIPS*, 2020.
- [32] D. Kalashnikov, A. Irpan, P. Pastor, J. Ibarz, A. Herzog, E. Jang, D. Quillen, E. Holly, M. Kalakrishnan, V. Vanhoucke *et al.*, “Scalable deep reinforcement learning for vision-based robotic manipulation,” in *CoRL*, 2018.
- [33] E. Jang, A. Irpan, M. Khansari, D. Kappler, F. Ebert, C. Lynch, S. Levine, and C. Finn, “BC-Z: Zero-shot task generalization with robotic imitation learning,” in *CoRL*, 2022.
- [34] A. Goyal, J. Xu, Y. Guo, V. Blukis, Y.-W. Chao, and D. Fox, “RVT: Robotic view transformer for 3D object manipulation,” in *CoRL*, 2023.
- [35] T.-W. Ke, N. Gkanatsios, and K. Fragkiadaki, “3D Diffuser Actor: Policy diffusion with 3D scene representations,” in *CoRL*, 2024.
- [36] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, “Diffusion policy: Visuomotor policy learning via action diffusion,” *RSS*, 2023.
- [37] A. Goyal, V. Blukis, J. Xu, Y. Guo, Y.-W. Chao, and D. Fox, “RVT2: Learning precise manipulation from few demonstrations,” in *RSS*, 2024.
- [38] S. James, K. Wada, T. Laidlow, and A. J. Davison, “Coarse-to-fine q-attention: Efficient learning for visual robotic manipulation via discretisation,” in *CVPR*, 2022.
- [39] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark *et al.*, “Learning transferable visual models from natural language supervision,” in *ICML*, 2021.
- [40] M. Minderer, A. A. Gritsenko, and N. Houlsby, “Scaling open-vocabulary object detection,” in *NeurIPS*, A. Oh, T. Naumann, A. Globerson, K. Saenko, M. Hardt, and S. Levine, Eds., 2023.
- [41] T. Yu, T. Xiao, A. Stone, J. Tompson, A. Brohan, S. Wang, J. Singh, C. Tan, D. M. J. Peralta, B. Ichter, K. Hausman, and F. Xia, “Scaling robot learning with semantically imagined experience,” *arXiv:2302.11550*, 2023.
- [42] OpenAI, “GPT-4 technical report,” *arXiv:2302.11550*, 2023.
- [43] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch, “Language models as zero-shot planners: Extracting actionable knowledge for embodied agents,” in *ICML*, 2022.
- [44] A. Brohan, Y. Chebotar, C. Finn, K. Hausman, A. Herzog, D. Ho, J. Ibarz, A. Irpan, E. Jang, R. Julian *et al.*, “Do as I can, not as I say: Grounding language in robotic affordances,” in *CoRL*, 2023.
- [45] Y. Hu, F. Lin, T. Zhang, L. Yi, and Y. Gao, “Look before you leap: Unveiling the power of GPT-4V in robotic vision-language planning,” *arXiv:2311.17842*, 2023.
- [46] G. OpenAI, “GPT-4V(ision) system card,” *preprint*, 2023.
- [47] S. Liu, S. James, A. J. Davison, and E. Johns, “Auto-Lambda: Disentangling dynamic task relationships,” *TMLR*, 2022.
- [48] Q.-Y. Zhou, J. Park, and V. Koltun, “Open3D: A modern library for 3D data processing,” *arXiv:1801.09847*, 2018.
- [49] X. Wu, L. Jiang, P.-S. Wang, Z. Liu, X. Liu, Y. Qiao, W. Ouyang, T. He, and H. Zhao, “Point Transformer v3: Simpler, faster, stronger,” in *CVPR*, 2024.
- [50] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional networks for biomedical image segmentation,” in *MICCAI*, 2015.
- [51] S. Dubey, S. K. Singh, and B. Chaudhuri, “AdaNorm: Adaptive gradient norm correction based optimizer for CNNs,” in *WACV*, 2023.
- [52] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, “Attention is all you need,” in *NeurIPS*, 2017.

## APPENDIX

### A. The proposed GemBench benchmark

Table VIII presents all the tasks and variations used in training and the four generalization levels in testing in the proposed GemBench.

### B. Details of 3D-LOTUS policy

1) *Point cloud preprocessing*: We automatically filter out irrelevant points from the point cloud during preprocessing. To remove background and table points, we define the robot’s workspace and the table height, excluding all points outside these boundaries. For the robotic arm, we assign 3D bounding boxes to each of its links. Using the robot’s proprioceptive state, we transform these bounding boxes based on the known poses and remove any points within them. As a result, the remaining point cloud contains only objects and the robotic gripper. Figure 5 illustrates the point cloud before and after the point removal step.

2) *Action prediction*: We discretize the groundtruth position  $a_t^p$  and rotation  $a_t^r$  to train the model. For position prediction, let  $b_{t,k}^p \in \mathbb{R}^{n \times 2m}$  represent the position of the concatenated bins for all points along the  $k$  axis, where  $n$  is the number of points and  $2m$  is the number of bins per point. We calculate the Euclidean distance between each bin and  $a_{t,k}^p$ , defining the score for bin  $i$  as:

$$\hat{p}_{t,k,i} = \begin{cases} 0, & \text{if } \|b_{t,k,i}^p - a_{t,k}^p\|_2^2 > 0.01 \text{ or } b_{t,k,i}^p \in \mathbb{B} \\ \frac{1}{\|b_{t,k,i}^p - a_{t,k}^p\|_2^2}, & \text{otherwise.} \end{cases} \quad (1)$$

where  $\mathbb{B}$  denotes the set of points that belong to the robot arm and gripper, and their scores are set to zero to ensure that the gripper’s position is predicted based only on the objects in the scene. The groundtruth probability of position along the  $k$  axis is then obtained by normalizing the scores  $\hat{p}_{t,k,i}$  via L1 norm. For rotation prediction, we simply use one-hot label along each axis.

### C. Details of 3D-LOTUS++ policy

1) *Task planning*: Figure 6 illustrates the prompts used in LLMs for task planning. For each task variation in the training set, we craft a corresponding example as shown in Figure 7. During inference, for each new instruction, we use SentenceBert [?] to compute the sentence embedding of the instruction and compare it to the existing instructions. We then select the top 20 examples with the highest similarities to the query instruction as in-context examples to the LLM.

2) *Object grounding*: We use VLMs to detect the location of queried objects except for tasks requiring grounding different parts of articulated object like bottom drawer and top shelf. This limitation arises because our VLMs can only detect the whole object like drawer, but cannot ground the target level of the drawer. Therefore, we further leverage the LLM to predict the height range of the target object. We first obtain the overall height of the target object based on VLM’s prediction, then we use prompts presented in Figure 8 to guide the LLM in predicting the height range.

### D. Detailed evaluation results

**RLBench-18Task**. Table IX presents the detailed results of different models on each task in RLBench-18Task benchmark. Our 3D-LOTUS policy achieves better performances especially on tasks requiring high precision such as insert peg, place cups and stack cups.

**GemBench**. Table X to XIII show the results of different models on the four generalization levels in GemBench respectively.

### E. Detail task specification

We describe each task in detail along with its variations below, and highlight the newly created tasks in GemBench.

#### Push Button

**Filename**: push\_button.py

**Task**: Push button with the specified color.

**New/Modified**: No.

**Variations per level**: Level 1: 0 (maroon), 3 (navy) and 4 (yellow). Level 2: 13 (azure), 15 (rose) and 17 (white).

**Objects**: 1 button.

**Success Metric**: The button is completely pressed.

#### Close Fridge

**Filename**: close\_fridge.py

**Task**: Close the fridge door.

**New/Modified**: No.

**Objects**: 1 fridge.

**Variations per level**: Level 1: 0.

**Success Metric**: The revolute joint of the fridge door has been rotated so that the door is closed and in contact with the fridge cabinet.

#### Close Laptop Lid

**Filename**: close\_laptop\_lid.py

**Task**: Grasp the laptop lid and rotate to close the laptop.

**New/Modified**: No.

**Objects**: 1 box, 1 laptop.

**Variations per level**: Level 1: 0.

**Success Metric**: The revolute joint of the laptop lid has been rotated so that the laptop is closed.

#### Close Microwave

**Filename**: close\_microwave.py

**Task**: Close the microwave door.

**New/Modified**: No.

**Objects**: 1 microwave.

**Variations per level**: Level 1: 0.

**Success Metric**: The revolute joint of the microwave has been rotated so that the microwave door is closed.

#### Open Door

**Filename**: open\_door.py

**Task**: Pick up the door handle and open the door by pushing.

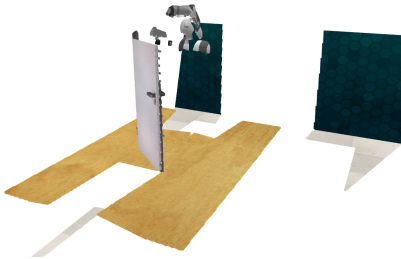
**New/Modified**: Yes, modified to add new instructions.

**Objects**: 1 door with a handle.

**Variations per level**: Level 1: 0.

**TABLE VIII: Training and testing tasks & variations in GemBench.** The testing tasks contain four levels of generalization, where Level 1 evaluates the generalization to novel placements, Level 2 novel rigid objects, Level 3 novel articulated objects, and Level 4 novel long-horizon tasks.

Train / Level 1		Level 2		Level 3		Level 4	
Task	Variation	Color	Shape	Instance	Category	Action-Part	Long-horizon
Press	Push button	maroon button navy button yellow button	azure button rose button white button	Lamp on			2 buttons 3 buttons 4 buttons
Pick	Pick and lift	red block lime block cyan block	teal block violet block black block	red cylinder red star red moon			
	Pick up cup	magenta cup silver cup orange cup	gray cup olive cup purple cup	red toy			
Push	Slide block	green target blue target	pink target yellow target				
	Reach and drag	teal target black target	cyan target navy target				
Screw	Close jar	azure jar violet jar	blue jar green jar				
	Screw bulb	rose bulb white bulb	lime bulb maroon bulb				
Close	Close fridge	fridge		fridge2	grill	door	
	Close laptop lid	laptop lid		laptop lid2		box	
	Close microwave	microwave		microwave2		drawer	
Open	Open door	door		door2	toilet	fridge	Take shoes
	Open box	box		box2		laptop lid	out of box
	Open drawer	bottom drawer top drawer		drawer2, drawer3 long drawer w/ 4 levels		microwave middle drawer	Put 3 items in drawer
Put/ Stack	Stack blocks	2 gray blocks 2 olive blocks 2 purple blocks	2 orange blocks 2 silver blocks 2 magenta blocks				Stack 3-4 blocks Stack 2 cups
	Put groceries	crackers box soup can		mustard bottle sugar box			Put all groceries
	Put money	bottom shelf middle shelf		Put cube in bottom shelf		top shelf	



(a) The original point cloud.



(b) Removing background and table.



(c) Removing robot arm.

**Fig. 5: Automatic point removal.** We use geometry information to automatically filter out irrelevant points in the scene.

**Success Metric:** The handle has been rotated 25° to unlock the door and the revolute joint of the door has been rotated 25°.

#### Open Box

**Filename:** open\_box.py

**Task:** Open the box lid.

**New/Modified:** Yes, modified to add new instructions.

**Objects:** 1 box.

**Variations per level:** Level 1: 0.

**Success Metric:** The revolute joint of the box lid has rotated 90° so that it is in the open configuration.

#### Open Drawer

**Filename:** open\_drawer.py

**Task:** Open one of the three drawers: top, middle, or bottom.

**New/Modified:** No.

**Objects:** 1 drawer.

**Variations per level:** Level 1: 0 (bottom) and 2 (top) and

I would like you to help me write Python code to control a robot arm operating in a tabletop environment. Please complete the code every time when I give you new query and a list of objects visible at the initial step. Pay attention to appeared patterns in the given context code. Be thorough and thoughtful in your code. Do not include any import statement. Do not repeat my question. Do not provide any text explanation (comment in code is okay).

You are only allowed to use the following action primitives that a robotic arm can perform:

1. `'grasp(object)'`: Grasp the specified object. Ensure that the robot gripper is open and not holding any other object before grasping. The robot gripper can only grasp one object at a time. After grasping, the robot gripper will close and securely hold the object. Return the grasped object.
2. `'move_grasped_object(target)'`: Move the grasped object to the specified target. Ensure that the robot gripper is closed and holding an object before moving. After moving, the robot gripper will still hold the object. The target can be a text description of a specified place, the location of previous objects, or a direction such as up, down, forward and out for small movements in those directions. Return the grasped object.
3. `'rotate_grasped_object()'`: Rotate the gripper while holding the object. Ensure that the robot gripper is holding an object before performing the rotation. After rotating, the gripper will still hold the object. Return the grasped object.
4. `'push_down(object)'`: Push down the specified object vertically, such as a button. The robot gripper does not hold the specified object but may hold other objects. Return the grasped object.
5. `'push_forward(object, target)'`: Push forward the specified object towards a target place. If no target is specified, the object will be pushed forward by a small distance. The robot gripper does not hold the specified object but may hold other objects. Return the grasped object.
6. `'release()'`: Open the gripper to release an object. Ensure the object is held in the gripper before releasing. After releasing, the gripper is open and not holding any object.

It's essential to stick to the format of these basic skills. When creating a plan, replace object or target inside the function with text descriptions or previously returned objects. Do not use objects not visible in the scene, but the robot can discover more objects through for example opening box or drawer. Generate step-by-step plans. Do not use for loop.

I will first give you the context of the code below:

**Fig. 6: Prompts used in LLM for task planning.**

**TABLE IX: Multi-Task Performance on RL Bench.** We report the success rate on all tasks in RL Bench-18Task [18] benchmark. Our 3D-LOTUS outperforms all methods while having higher training speed.

Models	Avg. Success $\uparrow$	Avg. Rank $\downarrow$	Train time (in days) $\downarrow$	Inf. Speed (in fps) $\uparrow$	Close Jar	Drag Stick	Insert Peg	Meat off Grill	Open Drawer	Place Cups	Place Wine
C2F-ARM-BC [38], [18]	20.1	8.6	-	-	24	24	4	20	20	0	8
HiveFormer [17]	45.3	6.9	-	-	52.0	76.0	0.0	<b>100.0</b>	52.0	0.0	80
PolarNet [2]	46.4	6.4	8.9	-	36.0	92.0	4.0	<b>100.0</b>	84.0	0.0	40
PerAct [18]	49.4	6.2	128.0	4.9	55.2 $\pm$ 4.7	89.6 $\pm$ 4.1	5.6 $\pm$ 4.1	70.4 $\pm$ 2.0	88.0 $\pm$ 5.7	2.4 $\pm$ 3.2	44.8 $\pm$ 7.8
RVT [34]	62.9	4.4	8.0	11.6	52.0 $\pm$ 2.5	99.2 $\pm$ 1.6	11.2 $\pm$ 3.0	88.0 $\pm$ 2.5	71.2 $\pm$ 6.9	4.0 $\pm$ 2.5	91.0 $\pm$ 5.2
Act3D [4]	65.0	4.3	40.0	-	92.0	92.0	27.0	94.0	<b>93.0</b>	3.0	80
RVT-2 [37]	81.4	2.4	6.6	20.6	<b>100.0</b> $\pm$ 0.0	99.0 $\pm$ 1.7	40.0 $\pm$ 0.0	99.0 $\pm$ 1.7	74.0 $\pm$ 11.8	38.0 $\pm$ 4.5	<b>95.0</b> $\pm$ 3.3
3D diffuser actor [35]	81.3	2.3	67.6	-	96.0 $\pm$ 2.5	<b>100.0</b> $\pm$ 0.0	65.6 $\pm$ 4.1	96.8 $\pm$ 1.6	89.6 $\pm$ 4.1	24.0 $\pm$ 7.6	93.6 $\pm$ 4.8
3D-LOTUS (ours)	<b>83.1</b>	<b>2.2</b>	<b>2.2</b>	9.5	96.0 $\pm$ 0.0	<b>100.0</b> $\pm$ 0.0	<b>69.6</b> $\pm$ 3.6	98.4 $\pm$ 2.2	85.6 $\pm$ 7.3	<b>40.8</b> $\pm$ 12.1	91.2 $\pm$ 6.6
Models	Push Buttons	Put in Cupboard	Put in Drawer	Put in Safe	Screw Bulb	Slide Block	Sort Shape	Stack Blocks	Stack Cups	Sweep to Dustpan	Turn Tap
C2F-ARM-BC [38], [18]	72	0	4	12	8	16	8	0	0	0	68
HiveFormer [17]	84	32.0	68.0	76.0	8.0	64.0	8.0	8.0	0.0	28.0	80
PolarNet [2]	96	12.0	32.0	84.0	44.0	56.0	12.0	4.0	8.0	52.0	80
PerAct [18]	92.8 $\pm$ 3.0	28.0 $\pm$ 4.4	51.2 $\pm$ 4.7	84.0 $\pm$ 3.6	17.6 $\pm$ 2.0	74.0 $\pm$ 13.0	16.8 $\pm$ 4.7	26.4 $\pm$ 3.2	2.4 $\pm$ 2.0	52.0 $\pm$ 0.0	88.0 $\pm$ 4.4
Act3D [4]	99	51.0	90.0	95.0	47.0	93.0	8.0	12.0	9.0	92.0	94
RVT [34]	<b>100.0</b> $\pm$ 0.0	49.6 $\pm$ 3.2	88.0 $\pm$ 5.7	91.2 $\pm$ 3.0	48.0 $\pm$ 5.7	81.6 $\pm$ 5.4	36.0 $\pm$ 2.5	28.8 $\pm$ 3.9	26.4 $\pm$ 8.2	72.0 $\pm$ 0.0	93.6 $\pm$ 4.1
RVT-2 [37]	<b>100.0</b> $\pm$ 0.0	66.0 $\pm$ 4.5	96.0 $\pm$ 0.0	96.0 $\pm$ 2.8	88.0 $\pm$ 4.9	92.0 $\pm$ 2.8	35.0 $\pm$ 7.1	<b>80.0</b> $\pm$ 2.8	69.0 $\pm$ 5.9	<b>100.0</b> $\pm$ 0.0	99.0 $\pm$ 1.7
3D diffuser actor [35]	98.4 $\pm$ 2.0	<b>85.6</b> $\pm$ 4.1	96.0 $\pm$ 3.6	<b>97.6</b> $\pm$ 2.0	82.4 $\pm$ 2.0	97.6 $\pm$ 3.2	<b>44.0</b> $\pm$ 4.4	68.3 $\pm$ 3.3	47.2 $\pm$ 8.5	84.0 $\pm$ 4.4	<b>99.2</b> $\pm$ 1.6
3D-LOTUS (ours)	<b>100.0</b> $\pm$ 0.0	78.4 $\pm$ 4.6	<b>97.6</b> $\pm$ 3.6	95.2 $\pm$ 3.4	<b>88.8</b> $\pm$ 3.4	<b>99.2</b> $\pm$ 1.8	34.4 $\pm$ 4.6	58.4 $\pm$ 8.3	<b>75.2</b> $\pm$ 7.7	96.0 $\pm$ 2.8	90.4 $\pm$ 4.6

Level 3: 1(middle).

drawer is fully extended.

**Success Metric:** The prismatic joint of the specified

```

1 # query: push the maroon button.
2 button = push_down(object="maroon button")
3
4 # query: close fridge.
5 fridge_door = push_forward(object="fridge door")
6
7 # query: close laptop lid.
8 laptop_lid = grasp(object="laptop lid")
9 laptop_lid = move_grasped_object(target="down")
10 release()
11
12 # query: close microwave.
13 microwave_door = push_forward(object="microwave door")
14
15 # query: open the door.
16 door_handle = grasp(object="door handle")
17 door_handle = rotate_grasped_object()
18 door_handle = push_forward(object=door_handle)
19
20 # query: open box.
21 box_lid = grasp(object="box lid")
22 box_lid = move_grasped_object(target="up")
23 release()
24
25 # query: open bottom drawer.
26 bottom_handle = grasp(object="bottom drawer handle")
27 bottom_handle = move_grasped_object(target="out")
28 release()
29
30 # query: lift the cyan block up to the target.
31 cyan_cube = grasp(object="cyan cube")
32 cyan_cube = move_grasped_object(target="red ball")
33
34 # query: lift the orange cup.
35 orange_cup = grasp(object="orange cup")
36 orange_cup = move_grasped_object(target="up")
37
38 # query: pick up and set down 2 purple blocks on top of each other.
39 purple_cube_1 = grasp(object="purple cube")
40 purple_cube_1 = move_grasped_object(target="green square")
41 release()
42 purple_cube_2 = grasp(object="purple cube", not=[purple_cube_1])
43 purple_cube_2 = move_grasped_object(target=purple_cube_1)
44 release()
45
46 # query: put the crackers box in the cupboard.
47 crackers_box = grasp(object="crakers box")
48 crackers_box = move_grasped_object(target="cupboard")
49 release()
50
51 # query: leave the money on the middle shelf on the safe.
52 money = grasp(object="money")
53 money = move_grasped_object(target="middle shelf")
54 release()
55
56 # query: push the block until it is sitting on top of the green target.
57 cube = push_forward(object="red cube", target="green square")
58
59 # query: use the stick to drag the cube onto the teal target.
60 stick = grasp(object="stick")
61 cube = push_forward(object="gray cube", target="teal square")
62
63 # query: screw on the violet jar lid.
64 lid = grasp(object="gray lid")
65 lid = move_grasped_object(target="violet jar")
66 lid = rotate_grasped_object()
67 release()
68
69 # query: screw in the white light bulb.
70 bulb = grasp(object="white light bulb")
71 bulb = move_grasped_object(target="brown lamp")
72 bulb = rotate_grasped_object()
73 release()

```

**Fig. 7: In-context examples of each training task for task planning.**

Please help me define the height range for different levels of an articulated object. I will provide you the target level and the total height of the object. Your task is to output two numbers representing the height range for the target level. Pay attention to the appeared patterns in the given examples. Do not repeat my question. Do not provide any text explanation.

The examples are as follows:

target: bottom drawer handle  
height: 0.4  
target height range: [0.1, 0.2]

target: top drawer handle  
height: 0.4  
target height range: [0.3, 0.4]

target: bottom shelf  
height: 0.5  
target height range: [0, 0.1]

target: middle shelf  
height: 0.5  
target height range: [0.15, 0.25]

**Fig. 8: Prompts used in LLMs to predict the height range of an object.**

**TABLE X: Performance on GemBench Level 1.**

Method	Avg.	Close Fridge+0	Close Jar+15	Close Jar+16	CloseLaptop Lid+0	Close Microwave+0	LightBulb In+17	LightBulb In+19	Open Box+0	Open Door+0	Open Drawer+0
Hiveformer [17]	60.3 $\pm$ 1.5	96 $\pm$ 4.2	64 $\pm$ 13.9	92 $\pm$ 2.7	90 $\pm$ 3.5	88 $\pm$ 7.6	12 $\pm$ 4.5	13 $\pm$ 6.7	4 $\pm$ 4.2	53 $\pm$ 15.2	15 $\pm$ 12.2
PolarNet [2]	77.6 $\pm$ 0.9	99 $\pm$ 2.2	99 $\pm$ 2.2	99 $\pm$ 2.2	95 $\pm$ 3.5	98 $\pm$ 2.7	72 $\pm$ 12.5	71 $\pm$ 6.5	32 $\pm$ 11.5	69 $\pm$ 8.9	61 $\pm$ 12.4
3D diffuser actor [35]	91.9 $\pm$ 0.8	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	<b>99</b> $\pm$ 2.2	<b>100</b> $\pm$ 0.0	85 $\pm$ 5.0	88 $\pm$ 2.7	11 $\pm$ 2.2	96 $\pm$ 4.2	82 $\pm$ 9.1
RVT-2 [37]	89.0 $\pm$ 0.8	77 $\pm$ 11.0	97 $\pm$ 4.5	98 $\pm$ 2.7	77 $\pm$ 13.0	<b>100</b> $\pm$ 0.0	<b>93</b> $\pm$ 5.7	<b>91</b> $\pm$ 8.2	7 $\pm$ 4.5	<b>98</b> $\pm$ 4.5	<b>93</b> $\pm$ 5.7
3D-LOTUS (ours)	<b>94.3</b> $\pm$ 3.5	96 $\pm$ 3.7	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	98 $\pm$ 2.5	98 $\pm$ 4.0	84 $\pm$ 7.4	85 $\pm$ 9.5	<b>99</b> $\pm$ 2.0	77 $\pm$ 2.5	83 $\pm$ 8.7
3D-LOTUS++ (ours)	68.7 $\pm$ 0.6	95 $\pm$ 0.0	<b>100</b> $\pm$ 0.0	99 $\pm$ 2.0	28 $\pm$ 2.5	87 $\pm$ 5.1	55 $\pm$ 10.5	45 $\pm$ 8.9	55 $\pm$ 8.9	79 $\pm$ 9.7	68 $\pm$ 12.5
Method	Open Drawer+2	Pick& Lift+0	Pick& Lift+2	Pick& Lift+7	PickUp Cup+8	PickUp Cup+9	PickUp Cup+11	Push Button+0	Push Button+3	Push Button+4	PutIn Cupboard+0
Hiveformer [17]	59 $\pm$ 7.4	86 $\pm$ 4.2	92 $\pm$ 6.7	93 $\pm$ 2.7	83 $\pm$ 7.6	69 $\pm$ 12.9	61 $\pm$ 19.8	84 $\pm$ 11.9	68 $\pm$ 6.7	87 $\pm$ 7.6	34 $\pm$ 8.2
PolarNet [2]	90 $\pm$ 7.1	92 $\pm$ 9.1	84 $\pm$ 7.4	88 $\pm$ 5.7	82 $\pm$ 7.6	79 $\pm$ 4.2	72 $\pm$ 10.4	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	99 $\pm$ 2.2	52 $\pm$ 7.6
3D diffuser actor [35]	<b>97</b> $\pm$ 4.5	<b>99</b> $\pm$ 2.2	99 $\pm$ 2.2	99 $\pm$ 2.2	96 $\pm$ 2.2	97 $\pm$ 4.5	98 $\pm$ 2.7	98 $\pm$ 2.7	96 $\pm$ 4.2	98 $\pm$ 2.7	85 $\pm$ 5.0
RVT-2 [37]	94 $\pm$ 4.2	<b>99</b> $\pm$ 2.2	98 $\pm$ 2.7	<b>100</b> $\pm$ 0.0	<b>99</b> $\pm$ 2.2	<b>99</b> $\pm$ 2.2	<b>99</b> $\pm$ 2.2	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	88 $\pm$ 8.4
3D-LOTUS (ours)	93 $\pm$ 6.0	<b>99</b> $\pm$ 2.0	<b>100</b> $\pm$ 0.0	99 $\pm$ 2.0	97 $\pm$ 4.0	96 $\pm$ 3.7	94 $\pm$ 4.9	99 $\pm$ 2.0	99 $\pm$ 2.0	<b>100</b> $\pm$ 0.0	<b>89</b> $\pm$ 5.8
3D-LOTUS++ (ours)	75 $\pm$ 4.5	97 $\pm$ 6.0	94 $\pm$ 3.7	93 $\pm$ 5.1	86 $\pm$ 8.0	88 $\pm$ 6.8	91 $\pm$ 4.9	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	1 $\pm$ 2.0
Method	PutIn Cupboard+3	PutMoney InSafe+0	PutMoney InSafe+1	Reach& Drag+14	Reach& Drag+18	Slide Block+0	Slide Block+1	Stack Blocks+30	Stack Blocks+36	Stack Blocks+39	
Hiveformer [17]	74 $\pm$ 6.5	85 $\pm$ 3.5	88 $\pm$ 2.7	37 $\pm$ 5.7	32 $\pm$ 7.6	99 $\pm$ 2.2	91 $\pm$ 12.4	6 $\pm$ 5.5	7 $\pm$ 4.5	6 $\pm$ 4.2	
PolarNet [2]	<b>88</b> $\pm$ 4.5	93 $\pm$ 4.5	95 $\pm$ 5.0	99 $\pm$ 2.2	99 $\pm$ 2.2	<b>100</b> $\pm$ 0.0	0 $\pm$ 0.0	34 $\pm$ 10.8	30 $\pm$ 9.4	36 $\pm$ 12.9	
3D diffuser actor [35]	82 $\pm$ 11.5	<b>95</b> $\pm$ 5.0	98 $\pm$ 2.7	<b>100</b> $\pm$ 0.0	99 $\pm$ 2.2	<b>100</b> $\pm$ 0.0	89 $\pm$ 4.2	88 $\pm$ 7.6	85 $\pm$ 6.1	89 $\pm$ 5.5	
RVT-2 [37]	80 $\pm$ 6.1	93 $\pm$ 8.4	96 $\pm$ 8.5	85 $\pm$ 10.0	94 $\pm$ 2.2	<b>100</b> $\pm$ 0.0	37 $\pm$ 6.7	88 $\pm$ 5.7	<b>93</b> $\pm$ 2.7	88 $\pm$ 11.5	
3D-LOTUS (ours)	72 $\pm$ 11.2	94 $\pm$ 3.7	<b>99</b> $\pm$ 2.0	99 $\pm$ 2.0	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	<b>94</b> $\pm$ 5.8	91 $\pm$ 6.6	<b>90</b> $\pm$ 4.5	
3D-LOTUS++ (ours)	2 $\pm$ 2.5	22 $\pm$ 6.8	16 $\pm$ 4.9	94 $\pm$ 3.7	62 $\pm$ 8.7	<b>100</b> $\pm$ 0.0	65 $\pm$ 5.5	86 $\pm$ 5.8	20 $\pm$ 4.5	28 $\pm$ 13.6	

### Pick and Lift

**Filename:** pick\_and\_lift.py

**Task:** Pick a colored cube and lift it to a red sphere target.

**New/Modified:** No.

**Objects:** 3 colored cubes, one with the specified color and the other two with different colors as distractors.

**Variations per level:** Level 1: 0 (red), 2 (lime) and 7 (cyan) and Level 2: 14 (teal), 16 (violet) and 18 (black).

**Success Metric:** The cube of the specified color is grasped and next to the target red sphere.

**Task:** Pick the cup with the specified color and lift it from the table.

**New/Modified:** No.

**Objects:** 3 tall colored cups.

**Variations per level:** Level 1: 8 (magenta), 9 (silver) and 11 (orange) and Level 2: 10 (gray), 12 (olive) and 13 (purple).

**Success Metric:** The cup of the specified color is grasped and lifted from the table.

### Stack Blocks

**Filename:** stack\_blocks.py

**Task:** Stack  $N$  blocks of the specified color on the green

**TABLE XI: Performance on GemBench Level 2.**

Method	Avg.	Push Button+13	Push Button+15	Push Button+17	Pick& Lift+14	Pick& Lift+16	Pick& Lift+18	PickUp Cup+10	PickUp Cup+12	PickUp Cup+13
Hiveformer	26.1 $\pm$ 1.4	97 $\pm$ 2.7	85 $\pm$ 10.0	88 $\pm$ 2.7	21 $\pm$ 6.5	9 $\pm$ 4.2	8 $\pm$ 6.7	30 $\pm$ 7.1	22 $\pm$ 13.5	26 $\pm$ 10.6
PolarNet	37.1 $\pm$ 1.4	100 $\pm$ 0.0	<b>100</b> $\pm$ 0.0	85 $\pm$ 7.9	3 $\pm$ 4.5	1 $\pm$ 2.2	0 $\pm$ 0.0	48 $\pm$ 11.0	46 $\pm$ 8.9	16 $\pm$ 6.5
3D diffuser actor	43.4 $\pm$ 2.8	87 $\pm$ 13.0	81 $\pm$ 6.5	60 $\pm$ 9.4	9 $\pm$ 4.2	18 $\pm$ 9.1	0 $\pm$ 0.0	84 $\pm$ 5.5	60 $\pm$ 11.7	62 $\pm$ 13.0
RVT-2	51.0 $\pm$ 2.3	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	47 $\pm$ 7.6	29 $\pm$ 9.6	8 $\pm$ 4.5	81 $\pm$ 8.2	59 $\pm$ 9.6	72 $\pm$ 9.7
3D-LOTUS (ours)	49.9 $\pm$ 2.2	99 $\pm$ 2.0	<b>100</b> $\pm$ 0.0	<b>100</b> $\pm$ 0.0	3 $\pm$ 2.5	18 $\pm$ 8.7	33 $\pm$ 9.3	<b>89</b> $\pm$ 3.7	78 $\pm$ 8.7	57 $\pm$ 7.5
3D-LOTUS++ (ours)	<b>64.5</b> $\pm$ 0.9	99 $\pm$ 2.0	<b>100</b> $\pm$ 0.0	99 $\pm$ 2.0	<b>94</b> $\pm$ 3.7	<b>96</b> $\pm$ 3.7	<b>95</b> $\pm$ 3.2	79 $\pm$ 4.9	<b>89</b> $\pm$ 9.7	<b>84</b> $\pm$ 10.2
Method	Stack Blocks+24	Stack Blocks+27	Stack Blocks+33	Slide Block+2	Slide Block+3	Close Jar+3	Close Jar+4	LightBulb In+1	LightBulb In+2	Lamp On+0
Hiveformer	0 $\pm$ 0.0	4 $\pm$ 4.2	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	4 $\pm$ 4.2	0 $\pm$ 0.0	7 $\pm$ 4.5
PolarNet	1 $\pm$ 2.2	2 $\pm$ 2.7	6 $\pm$ 8.2	0 $\pm$ 0.0	0 $\pm$ 0.0	20 $\pm$ 10.6	82 $\pm$ 5.7	22 $\pm$ 11.5	17 $\pm$ 8.4	<b>14</b> $\pm$ 10.8
3D diffuser actor	<b>66</b> $\pm$ 13.9	82 $\pm$ 2.7	50 $\pm$ 14.6	0 $\pm$ 0.0	0 $\pm$ 0.0	23 $\pm$ 16.8	82 $\pm$ 5.7	51 $\pm$ 17.8	<b>60</b> $\pm$ 10.0	7 $\pm$ 7.6
RVT-2	18 $\pm$ 4.5	<b>56</b> $\pm$ 16.7	45 $\pm$ 13.7	0 $\pm$ 0.0	1 $\pm$ 2.2	7 $\pm$ 7.6	77 $\pm$ 5.7	<b>68</b> $\pm$ 14.4	6 $\pm$ 6.5	0 $\pm$ 0.0
3D-LOTUS (ours)	13 $\pm$ 8.1	40 $\pm$ 9.5	<b>69</b> $\pm$ 5.8	0 $\pm$ 0.0	0 $\pm$ 0.0	71 $\pm$ 5.8	90 $\pm$ 4.5	24 $\pm$ 4.9	41 $\pm$ 8.6	0 $\pm$ 0.0
3D-LOTUS++ (ours)	22 $\pm$ 9.3	<b>83</b> $\pm$ 7.5	59 $\pm$ 3.7	<b>27</b> $\pm$ 9.8	<b>5</b> $\pm$ 3.2	<b>98</b> $\pm$ 2.5	<b>96</b> $\pm$ 3.7	56 $\pm$ 9.7	43 $\pm$ 7.5	2 $\pm$ 2.0
Method	Reach& Drag+5	Reach& Drag+7	PutCube InSafe+0	Pick&Lift Cylinder+0	Pick&Lift Star+0	Pick&Lift Moon+0	Pick&Lift Toy+0	PutIn Cupboard+7	PutIn Cupboard+8	
Hiveformer	1 $\pm$ 2.2	0 $\pm$ 0.0	4 $\pm$ 2.2	78 $\pm$ 5.7	73 $\pm$ 7.6	88 $\pm$ 2.7	87 $\pm$ 4.5	0 $\pm$ 0.0	0 $\pm$ 0.0	
PolarNet	61 $\pm$ 8.2	10 $\pm$ 6.1	<b>40</b> $\pm$ 14.1	93 $\pm$ 6.7	88 $\pm$ 8.4	93 $\pm$ 6.7	90 $\pm$ 3.5	0 $\pm$ 0.0	0 $\pm$ 0.0	
3D diffuser actor	0 $\pm$ 0.0	64 $\pm$ 6.5	3 $\pm$ 2.7	<b>99</b> $\pm$ 2.2	43 $\pm$ 17.9	91 $\pm$ 9.6	30 $\pm$ 9.4	0 $\pm$ 0.0	<b>3</b> $\pm$ 4.5	
RVT-2	91 $\pm$ 2.2	<b>89</b> $\pm$ 6.5	6 $\pm$ 5.5	98 $\pm$ 2.7	<b>98</b> $\pm$ 4.5	<b>94</b> $\pm$ 4.2	78 $\pm$ 8.4	0 $\pm$ 0.0	0 $\pm$ 0.0	
3D-LOTUS (ours)	<b>95</b> $\pm$ 4.5	18 $\pm$ 10.8	25 $\pm$ 5.5	88 $\pm$ 8.7	69 $\pm$ 6.6	80 $\pm$ 8.4	<b>96</b> $\pm$ 3.7	0 $\pm$ 0.0	0 $\pm$ 0.0	
3D-LOTUS++ (ours)	94 $\pm$ 2.0	64 $\pm$ 12.4	37 $\pm$ 5.1	91 $\pm$ 2.0	94 $\pm$ 3.7	29 $\pm$ 6.6	71 $\pm$ 2.0	1 $\pm$ 2.0	0 $\pm$ 0.0	

**TABLE XII: Performance on GemBench Level 3.**

Method	Avg.	Close Door+0	Close Box+0	Close Fridge2+0	CloseLaptop Lid2+0	Close Microwave2+0	Open Door2+0	Open Box2+0
Hiveformer	35.1 $\pm$ 1.7	0 $\pm$ 0.0	1 $\pm$ 2.2	34 $\pm$ 9.6	52 $\pm$ 9.1	15 $\pm$ 7.1	32 $\pm$ 11.5	5 $\pm$ 3.5
PolarNet	38.5 $\pm$ 1.7	0 $\pm$ 0.0	0 $\pm$ 0.0	78 $\pm$ 5.7	26 $\pm$ 8.2	74 $\pm$ 6.5	33 $\pm$ 6.7	23 $\pm$ 8.4
3D diffuser actor	37.0 $\pm$ 2.2	0 $\pm$ 0.0	0 $\pm$ 0.0	<b>97</b> $\pm$ 2.7	23 $\pm$ 6.7	88 $\pm$ 7.6	<b>86</b> $\pm$ 7.4	<b>67</b> $\pm$ 9.8
RVT-2	36.0 $\pm$ 2.2	1 $\pm$ 2.2	2 $\pm$ 2.7	72 $\pm$ 6.7	42 $\pm$ 14.0	71 $\pm$ 8.9	79 $\pm$ 6.5	5 $\pm$ 6.1
3D-LOTUS (ours)	38.1 $\pm$ 1.1	0 $\pm$ 0.0	<b>58</b> $\pm$ 8.1	36 $\pm$ 9.7	<b>54</b> $\pm$ 10.7	85 $\pm$ 7.1	42 $\pm$ 6.8	11 $\pm$ 6.6
3D-LOTUS++ (ours)	<b>41.5</b> $\pm$ 1.8	<b>1</b> $\pm$ 2.0	29 $\pm$ 8.6	93 $\pm$ 2.5	50 $\pm$ 9.5	<b>99</b> $\pm$ 2.0	52 $\pm$ 10.3	16 $\pm$ 8.0
Method	Open Drawer2+0	Open Drawer3+0	OpenDrawer Long+0	OpenDrawer Long+1	OpenDrawer Long+2	OpenDrawer Long+3	Toilet SeatUp+0	Open Fridge+0
Hiveformer	59 $\pm$ 11.9	39 $\pm$ 11.9	78 $\pm$ 8.4	82 $\pm$ 4.5	49 $\pm$ 4.2	57 $\pm$ 11.5	6 $\pm$ 4.2	0 $\pm$ 0.0
PolarNet	<b>91</b> $\pm$ 4.2	29 $\pm$ 8.2	84 $\pm$ 11.9	<b>88</b> $\pm$ 5.7	<b>63</b> $\pm$ 8.4	37 $\pm$ 7.6	2 $\pm$ 2.7	4 $\pm$ 2.2
3D diffuser actor	19 $\pm$ 8.2	1 $\pm$ 2.2	15 $\pm$ 5.0	35 $\pm$ 13.7	26 $\pm$ 9.6	<b>79</b> $\pm$ 12.9	0 $\pm$ 0.0	<b>7</b> $\pm$ 5.7
RVT-2	81 $\pm$ 11.9	0 $\pm$ 0.0	<b>84</b> $\pm$ 8.2	39 $\pm$ 10.8	11 $\pm$ 8.9	75 $\pm$ 6.1	7 $\pm$ 5.7	0 $\pm$ 0.0
3D-LOTUS (ours)	90 $\pm$ 3.2	22 $\pm$ 8.1	56 $\pm$ 13.9	33 $\pm$ 11.2	17 $\pm$ 8.1	75 $\pm$ 6.3	0 $\pm$ 0.0	4 $\pm$ 5.8
3D-LOTUS++ (ours)	70 $\pm$ 5.5	<b>41</b> $\pm$ 4.9	72 $\pm$ 4.0	52 $\pm$ 10.8	23 $\pm$ 8.1	78 $\pm$ 5.1	<b>8</b> $\pm$ 5.1	0 $\pm$ 0.0
Method	OpenLaptop Lid+0	Open Microwave+0	PutMoney InSafe+2	Open Drawer+1	Close Drawer+0	Close Grill+0		
Hiveformer	<b>100</b> $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	83 $\pm$ 5.7	44 $\pm$ 10.8		
PolarNet	<b>100</b> $\pm$ 0.0	0 $\pm$ 0.0	1 $\pm$ 2.2	4 $\pm$ 4.2	29 $\pm$ 11.9	42 $\pm$ 11.5		
3D diffuser actor	<b>100</b> $\pm$ 0.0	0 $\pm$ 0.0	2 $\pm$ 4.5	0 $\pm$ 0.0	66 $\pm$ 7.4	<b>65</b> $\pm$ 13.7		
RVT-2	93 $\pm$ 5.7	0 $\pm$ 0.0	0 $\pm$ 0.0	<b>6</b> $\pm$ 2.2	78 $\pm$ 8.4	9 $\pm$ 4.2		
3D-LOTUS (ours)	<b>100</b> $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	<b>87</b> $\pm$ 8.1	29 $\pm$ 6.6		
3D-LOTUS++ (ours)	86 $\pm$ 6.6	0 $\pm$ 0.0	<b>13</b> $\pm$ 8.1	0 $\pm$ 0.0	69 $\pm$ 5.8	19 $\pm$ 13.9		

platform. There are always 4 blocks of the specified color, and 4 distractor blocks of another color. The block colors are sampled from the full set of 20 color instances.

**New/Modified:** No.

**Objects:** 8 color blocks (4 are distractors), and 1 green platform

**Variations per level:** Level 1: 30 (2 gray blocks), 36 (2 olive blocks), 39 (2 purple blocks) and Level 2: 24 (2 orange blocks), 27 (2 silver blocks) and 33 (2 magenta blocks).

**Success Metric:**  $N$  blocks are inside the area of the green platform.

### Put Groceries in Cupboard

**Filename:** put\_groceries\_in\_cupboard.py

**Task:** Grab the specified object and put it in the cupboard above. The scene always contains 9 YCB objects that are randomly placed on the tabletop.

**New/Modified:** Yes, modified the object names to include object category and therefore the instructions.

**Objects:** 9 YCB objects, and 1 cupboard (that hovers in the air like magic).

**Variations per level:** Level 1: 0 (crackers box) and 3 (soup

**TABLE XIII:** Performance on GemBench Level 4.

Method	Avg.	Push Buttons4+1	Push Buttons4+2	Push Buttons4+3	TakeShoes OutOfBox+0	PutItems InDrawer+0	PutItems InDrawer+2
Hiveformer	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
PolarNet	0.1 $\pm$ 0.2	1 $\pm$ 2.2	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
3D diffuser actor	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
RVT-2	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
3D-LOTUS (ours)	0.3 $\pm$ 0.3	3 $\pm$ 4.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
3D-LOTUS++ (ours)	<b>17.4</b> $\pm$ 0.4	<b>76</b> $\pm$ 7.4	<b>49</b> $\pm$ 8.6	<b>37</b> $\pm$ 8.1	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0

Method	PutItems InDrawer+4	Tower4+1	Tower4+3	Stack Cups+0	Stack Cups+3	PutAllGroceries InCupboard+0
Hiveformer	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
PolarNet	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
3D diffuser actor	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
RVT-2	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
3D-LOTUS (ours)	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0
3D-LOTUS++ (ours)	0 $\pm$ 0.0	<b>17</b> $\pm$ 10.8	<b>30</b> $\pm$ 13.4	0 $\pm$ 0.0	0 $\pm$ 0.0	0 $\pm$ 0.0

can) and Level 2: 7 (mustard bottle) and 8 (sugar box).

**Success Metric:** The specified object is inside the cupboard.

#### Put Money in Safe

**Filename:** put\_money\_in\_safe.py

**Task:** Pick up the stack of money and put it inside the safe on the specified shelf. The shelf has three placement locations: top, middle, bottom.

**New/Modified:** No.

**Objects:** 1 stack of money, and 1 safe.

**Variations per level:** Level 1: 0 (bottom) and 1 (middle) and Level 2: 2 (top).

**Success Metric:** The stack of money is on the specified shelf inside the safe.

#### Slide Block to Color Target

**Filename:** slide\_block\_to\_color\_target\_peract.py

**Task:** Slide the block on to one of the colored square targets. The target colors are limited to red, blue, pink, and yellow.

**New/Modified:** Yes, modified as in RLBench-18Task [18]. The original slide block to target.py task contained only one target. Three other targets were added to make a total of 4 variations.

**Objects:** 1 block and 4 colored target squares.

**Variations per level:** Level 1: 0 (green) and 1 (blue) and Level 2: 2 (pink) and 3 (yellow).

**Success Metric:** Some part of the block is inside the specified target area.

#### Reach and Drag

**Filename:** reach\_and\_drag\_peract.py

**Task:** Grab the stick and use it to drag the cube on to the specified colored target square. The target colors are sampled from the full set of 20 color instances.

**New/Modified:** Yes, modified as in RLBench-18Task [18]. The original reach\_and\_drag.py task contained only one target. Three other targets were added with randomized colors.

**Objects:** 1 block, 1 stick, and 4 colored target squares.

**Variations per level:** Level 1: 14 (teal) and 18 (black) and Level 2: 5 (cyan) and 7 (navy).

**Success Metric:** Some part of the block is inside the specified target area.

#### Close Jar

**Filename:** close\_jar.py

**Task:** Put the lid on the jar with the specified color and screw the lid in. The jar colors are sampled from the full set of 20 color instances.

**New/Modified:** No.

**Objects:** 1 jar lid and 2 colored jars.

**Variations per level:** Level 1: 15 (azure) and 16 (violet) and Level 2: 3 (blue) and 4 (green).

**Success Metric:** The lid is on top of the specified jar and the Franka gripper is not grasping anything.

#### Screw Bulb

**Filename:** light\_bulb\_in.py

**Task:** Pick up the light bulb from the specified holder, and screw it into the lamp stand. The colors of holder are sampled from the full set of 20 color instances. There are always two holders in the scene – one specified and one distractor holder.

**Modified:** No.

**Objects:** 2 light bulbs, 2 holders, and 1 lamp stand.

**Variations per level:** Level 1: 17 (rose) and 19 (white) and Level 2: 1 (lime) and 2 (maroon).

**Success Metric:** The bulb from the specified holder is inside the lamp stand dock.

#### Close Door

**Filename:** close\_door.py

**Task:** Grab the handle of the door and pull to close the door.

**Modified:** Yes, modified to add new instructions.

**Objects:** 1 door with a handle.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the door is in the close configuration.

### Close Box

**Filename:** close\_box.py

**Task:** Grab the box lid and rotate to close the box.

**New/Modified:** Yes, modified to add new instructions.

**Objects:** 1 box.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the box lid has rotated so that it is in the closed configuration.

### Close Drawer

**Filename:** close\_drawer.py

**Task:** Close one of the three drawers: top, middle, or bottom.

**New/Modified:** Yes, modified to add new instructions.

**Objects:** 1 drawer unit with 3 drawers.

**Variations per level:** Level 3: 0.

**Success Metric:** The prismatic joint of the specified drawer is fully contracted.

### Open Fridge

**Filename:** open\_fridge.py

**Task:** Grab the handle of the fridge door and pull to open the fridge.

**New/Modified:** Yes, modified to add new instructions.

**Objects:** 1 fridge.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the fridge door has rotated 70 degrees so that the fridge is open.

### Open Laptop Lid

**Filename:** open\_laptop\_lid.py

**Task:** Grab the laptop lid and open it.

**New/Modified:** Yes, we modified the close laptop lid task to add a new task where the laptop is already close and the goal is to open the lid instead of closing it.

**Objects:** 1 laptop and 1 box.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the laptop lid reaches its minimum value so that the laptop is open.

### Open Microwave

**Filename:** open\_microwave.py

**Task:** Grab the handle to pull and open the microwave door.

**New/Modified:** Yes, modified to add new instructions.

**Objects:** 1 microwave.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the microwave has rotated at least 80°.

### Put Cube In Safe

**Filename:** put\_cube\_in\_safe.py

**Task:** Pick up the cube and put it inside the safe on the specified shelf. The shelf has three placement locations: top, middle, bottom.

**New/Modified:** Yes, new task based on put money in safe where the bank note is changed by a cube.

**Objects:** 1 cube and 1 safe.

**Variations per level:** Level 1: 0.

**Success Metric:** The cube is on the specified shelf inside the safe.

### Close Fridge2

**Filename:** close\_fridge2.py

**Task:** Close the fridge door by pushing it.

**New/Modified:** Yes, new task based on close fridge where the fridge mesh and texture are completely changed.

**Objects:** 1 microwave.

**Variations per level:** Level 1: 0.

**Success Metric:** The revolute joint of the fridge door is in closed configuration and the door is in contact with the fridge cabinet.

### Close Laptop Lid2

**Filename:** close\_laptop\_lid2.py

**Task:** Grasp the laptop lid and rotate to close the laptop.

**New/Modified:** Yes, a new task based on close laptop lid task where the laptop is already close and the goal is to open the lid instead of closing it.

**Objects:** 1 box, 1 laptop.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the laptop lid is in the close configuration.

### Close Microwave2

**Filename:** close\_microwave2.py

**Task:** Close the microwave door.

**New/Modified:** Yes, new task based on close microwave task where the microwave and handle meshes and textures are completely changed.

**Objects:** 1 microwave.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the microwave is in the close configuration.

### Open Door2

**Filename:** open\_door2.py

**Task:** Pick up the door handle and open the door by pushing.

**New/Modified:** Yes, a new task based on open door task where the door, door frame and handle meshes and textures are completely changed.

**Objects:** 1 door with a handle.

**Variations per level:** Level 3: 0.

**Success Metric:** The handle has been rotated 25° to unlock the door and the revolute joint of the door has been rotated 20°.

### Open Box2

**Filename:** open\_box2.py

**Task:** Open the box lid.

**New/Modified:** Yes, a new task based on open box task where the box mesh and texture are changed.

**Objects:** 1 box with a lid.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the box lid has rotated 90° so that it is in the open configuration.

### Open Drawer2

**Filename:** open\_drawer2.py

**Task:** Open one of the three drawers: top, middle, or bottom.

**New/Modified:** Yes, a new task based on open drawer task where the drawer handles and cabinet meshes, colors and textures are changed.

**Objects:** 1 drawer.

**Variations per level:** Level 3: 0 (bottom).

**Success Metric:** The prismatic joint of the specified drawer is fully extended.

### Open Drawer3

**Filename:** open\_drawer3.py

**Task:** Open one of the three drawers: top, middle, or bottom.

**New/Modified:** Yes, a new task based on open drawer task where the drawer handles and cabinet meshes, colors and textures are changed and different to Open Drawer2.

**Objects:** 1 drawer.

**Variations per level:** Level 3: 0 (bottom).

**Success Metric:** The prismatic joint of the specified drawer is fully extended.

### Open Drawer Long

**Filename:** open\_drawer\_long.py

**Task:** Open one of the four drawers: top, top middle, bottom middle or bottom.

**New/Modified:** Yes, a new task based on open drawer task where the drawer unit has 4 drawers and the handles and cabinet meshes, colors and textures are changed.

**Objects:** 1 drawer.

**Variations per level:** Level 3: 0 (bottom), 1 (bottom middle), 2 (top middle) and 3 (top).

**Success Metric:** The prismatic joint of the specified drawer is fully extended.

### Lamp On

**Filename:** lamp\_on.py

**Task:** Press the button to light on the lamp.

**New/Modified:** No.

**Objects:** 1 lamp, 1 button.

**Variations per level:** Level 2: 0.

**Success Metric:** The button is pressed and the lamp light is on.

### Close Grill

**Filename:** close\_grill.py

**Task:** Close the grill.

**New/Modified:** No.

**Objects:** 1 grill.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the grill lid is in the close configuration.

### Toilet Seat Up

**Filename:** toilet\_seat\_up.py

**Task:** Grasp the toilet seat and rotate it to move it up the toilet.

**New/Modified:** No.

**Objects:** 1 toilet.

**Variations per level:** Level 3: 0.

**Success Metric:** The revolute joint of the toilet seat is in the open configuration.

### Pick and Lift Cylinder

**Filename:** pick\_and\_lift\_cylinder.py

**Task:** Pick and lift the cylinder with the specified color.

**New/Modified:** Yes, a new task based on pick and lift task where cubes are changed with cylinders.

**Objects:** 3 colored cylinders, one with the specified color and the other two with different colors as distractors.

**Variations per level:** Level 2: 0 (red).

**Success Metric:** The cylinder of the specified color is grasped and placed in the target red sphere.

### Pick and Lift Star

**Filename:** pick\_and\_lift\_star.py

**Task:** Pick and lift the star with the specified color.

**New/Modified:** Yes, a new task based on pick and lift task where cubes are changed with stars.

**Objects:** 3 colored stars, one with the specified color and the other two with different colors as distractors.

**Variations per level:** Level 2: 0 (red).

**Success Metric:** The star of the specified color is grasped and placed in the target red sphere.

### Pick and Lift Moon

**Filename:** pick\_and\_lift\_moon.py

**Task:** Pick and lift the moon with the specified color.

**New/Modified:** Yes, a new task based on pick and lift task where cubes are changed with moons.

**Objects:** 3 colored moons, one with the specified color and the other two with different colors as distractors.

**Variations per level:** Level 2: 0 (red).

**Success Metric:** The moon of the specified color is grasped and placed in the target red sphere.

### Pick and Lift Toy

**Filename:** pick\_and\_lift\_moon.py

**Task:** Pick and lift the duck rubber toy with the specified color

**New/Modified:** Yes, a new task based on pick and lift task where cubes are changed with duck rubber toys.

**Objects:** 3 colored duck rubber toys, one with the specified color and the other two with different colors as distractors.

**Variations per level:** Level 2: 0 (red).

**Success Metric:** The duck rubber toy of the specified color is grasped and placed in the target red sphere.

#### Push Buttons4

**Filename:** push\_buttons4.py

**Task:** Push 2/3/4 buttons in the given color order.

**New/Modified:** Yes, a new task based on push button where there are 4 buttons instead of one, and the goal is change to follow a sequence of 2/3/4 buttons colors.

**Objects:** 4 colored buttons.

**Variations per level:** Level 4: 1 (navy then teal), 2 (green then yellow then rose), 3 (maroon then blue then orange then magenta).

**Success Metric:** The buttons are pressed consecutively in the same order as specified by the instructions.

#### Take Shoes Out Of Box

**Filename:** take\_shoes\_out\_of\_box.py

**Task:** Open the box and then take both shoes and place them in the table.

**New/Modified:** No.

**Objects:** 1 box and 2 shoes.

**Variations per level:** Level 4: 0.

**Success Metric:** Both shoes are outside of the box and on top of the table.

#### Put Items In Drawer

**Filename:** put\_items\_in\_drawer.py

**Task:** Put a cube, a cylinder and a moon in the specified order and inside one of the three drawers: top, middle and bottom.

**New/Modified:** Yes. The task is a new long horizon task that combines open drawer task and put in drawer.

**Objects:** 1 cube, 1 cylinder, 1 moon and a drawer.

**Variations per level:** Level 4: 0 (cube then cylinder then moon in bottom drawer), 2 (cube then cylinder then moon in top drawer) and 4 (cube then moon then cylinder in middle drawer).

**Success Metric:** The cube, the cylinder and the moon have been placed in order inside the specified drawer.

#### Tower4

**Filename:** tower4.py

**Task:** Stack 2, 3 or 4 cubes to create a tower on top of the green platform.

**New/Modified:** Yes, a new task based on stack blocks reducing the number of cubes to 4 blocks of different colors without repetitions.

**Objects:** 4 cubes of different colors and a green platform.

**Variations per level:** Level 4: 1 (stack white then teal then blue) and 3 (stack orange then gray then lime then rose).

**Success Metric:** The cubes are stacked in the correct order and on top of the green platform.

#### Stack Cups

**Filename:** stack\_cups.py

**Task:** Stack all cups on top of the specified color cup. The cup colors are sampled from the full set of 20 color instances. The scene always contains three cups.

**New/Modified:** No.

**Objects:** 3 tall cups.

**Variations per level:** Level 4: 0 (red) and 3 (green).

**Success Metric:** All other cups are inside the specified cup.

#### Put All Groceries In Cupboard

**Filename:** put\_all\_groceries\_in\_cupboard.py

**Task:** Grab each one of the groceries and put them inside the cupboard until all of them are inside the cupboard.

**New/Modified:** Yes, modified to add new instructions and improve the groceries names.

**Objects:** 6 YCB objects (1 crackers box, 1 chocolate jello box, 1 strawberry jello box, 1 soup can, 1 spam can, 1 mustard bottle and 1 sugar box), and 1 cupboard (that hovers in the air like magic).

**Variations per level:** Level 4: 0.

**Success Metric:** All the groceries are set inside the cupboard.