

# SLAP: Spatial-Language Attention Policies

Anonymous Author(s)

Affiliation

Address

email

1       **Abstract:** Despite great strides in language-guided manipulation, existing work  
2       has been constrained to table-top settings. Table-tops allow for perfect and consis-  
3       tent camera angles, properties are that do not hold in mobile manipulation. Task  
4       plans that involve moving around the environment must be robust to egocentric  
5       views and changes in the plane and angle of grasp. A further challenge is ensur-  
6       ing this is all true while still being able to learn skills efficiently from limited data.  
7       We propose Spatial-Language Attention Policies (SLAP) as a solution. SLAP uses  
8       three-dimensional tokens as the input representation to train a single multi-task,  
9       language-conditioned action prediction policy. Our method shows an 80% success  
10      rate in the real world across eight tasks with a single model, and a 47.5% success  
11      rate when unseen clutter and unseen object configurations are introduced, even  
12      with only a handful of examples per task. This represents an improvement of 30%  
13      over prior work (20% given unseen distractors and configurations). In addition, we  
14      show how SLAPs robustness enables allows us to execute Task Plans from open-  
15      vocabulary instructions using a large language model for multi-step mobile ma-  
16      nipulation. For videos, see the website: <https://robotslap.github.io>

## 17   1 Introduction

18   Transformers have demonstrated impressive results on natural language processing tasks by be-  
19   ing able to contextualize large numbers of tokens over long sequences, and even show substantial  
20   promise for robotics in a variety of manipulation tasks [1, 2, 3]. However, when it comes to using  
21   transformers for *mobile* robots performing long-horizon tasks, we face the challenge of representing  
22   spatial information in a useful way. In other words, we need a fundamental unit of representation -  
23   an equivalent of a “word” or “token” - that can handle spatial awareness in a way that is independent  
24   of the robot’s exact embodiment. We argue this is essential for enabling robots to perform manipula-  
25   tion tasks in diverse human environments, where they need to be able to generalize to new positions,  
26   handle changes in the visual appearance of objects and be robust to irrelevant clutter. In this work,  
27   we propose Spatial-Language Attention Policies (SLAP), that use a point-cloud based tokenization  
28   which can scale to a number of viewpoints, and has a number of advantages over prior work.

29   SLAP tokenizes the world into a varying-length stream of multi-resolution spatial embeddings,  
30   which capture a local context based on PointNet++ [4] features. Unlike ViT-style [1], object-  
31   centric [5, 3], or static 3D grid features [2], our PointNet++-based [4] tokens capture free-form  
32   relations between observed points in space. This means that we can combine multiple camera views  
33   from a moving camera when making decisions and still process arbitrary-length sequences.

34   Our approach leverages a powerful skill representation we refer to as “attention-driven robot poli-  
35   cies” [6, 7, 8, 2, 9] operating on an input-space combining language with spatial information. Unlike  
36   other methods that directly predict robot motor controls [10, 1], these techniques predict goal poses  
37   in Cartesian space and integrate them with a motion planner [6, 8, 2] or conditional low-level poli-  
38   cy [9] to execute goal-driven motion. This approach requires less data, but it still has limitations  
39   such as making assumptions about the input scene’s size and camera position and long training  
40   times [7, 6, 2]. However, these methods fall into a different trap: they make strong assumptions

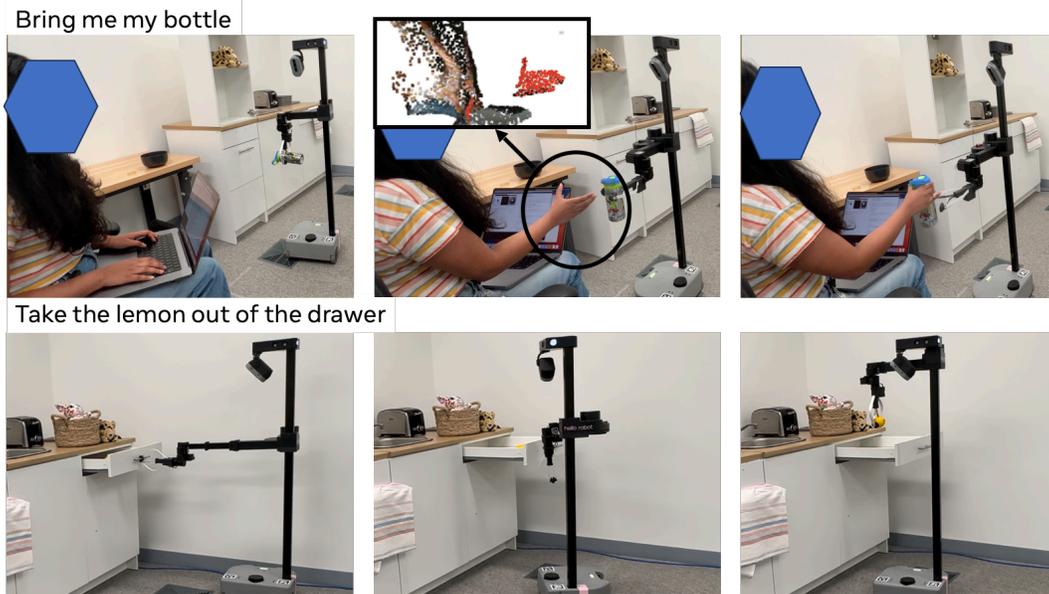


Figure 1: We propose SLAP, which allows us to learn skills for mobile manipulators to accomplish multi-step tasks given natural language goals. Our system works by training a language-conditioned *interaction prediction module*, which will determine which areas of a scene should be interacted with, in addition to an action policy which operates on predicted interaction points. This allows us to scale to more complex scenes, while predicting continuous actions.

41 about how big the input scene is [2], where the camera is [7, 6], and generally take a very long time  
 42 to train [7, 6], meaning they could not be used too quickly teach policies in a new environment.

43 SLAP uses a hybrid policy architecture. The *interaction prediction* module determines which parts  
 44 of the tokenized environment the robot focuses on, and a *relative action* module predicts parameters  
 45 of continuous motion with respect to the interaction features in the world. SLAP generalizes better  
 46 to unseen positions and orientations, as well as distractors, while being unrestricted by workspace  
 47 size, and camera placement assumption, using fewer demonstrations and training in roughly a day.

48 We evaluate SLAP on two robot platforms. First, on a Franka Panda we can perform a direct skills  
 49 comparison to the current state-of-the-art, PerAct, [2], where we demonstrate better performance  
 50 with 80% success rate on 8 static real-world tasks on held-out scene configurations and a 47.5%  
 51 success rate tested with out-of-distribution objects. Second, unlike prior work, we move beyond  
 52 the stationary camera views and robot arms of a table-top setting, and demonstrate SLAP on the  
 53 Hello Robot Stretch RE-2 mobile manipulator with an ego-centric camera and 6-DoF end-effector  
 54 configuration. In this setting, we also include task planning to successfully execute natural language  
 55 task instructions with 10 demonstrations over 5 learnt skills and 3 heuristic skills (Fig. 1).

## 56 2 Related Work

57 **Attention-Based Policies.** Attention-based policies have been widely studied in prior research and  
 58 have been found to have superior data efficiency, generalization, and the ability to solve previously  
 59 unsolvable problems [11, 9, 6, 12, 2, 13]. However, these approaches often rely on strong assump-  
 60 tions about the robot’s workspace, such as modeling the entire workspace as a 2D image [12, 6, 7, 8]  
 61 or a 3D voxel cube with predetermined scene bounds [2, 9]. This restricts their applicability and  
 62 may lead to issues related to camera positioning, workspace location, and discretization size. Ad-  
 63 ditionally, these works can be seen, at least partly, as applications of object detection systems like  
 64 Detic [14] or 3DETR [15], but they lack the manipulation component.

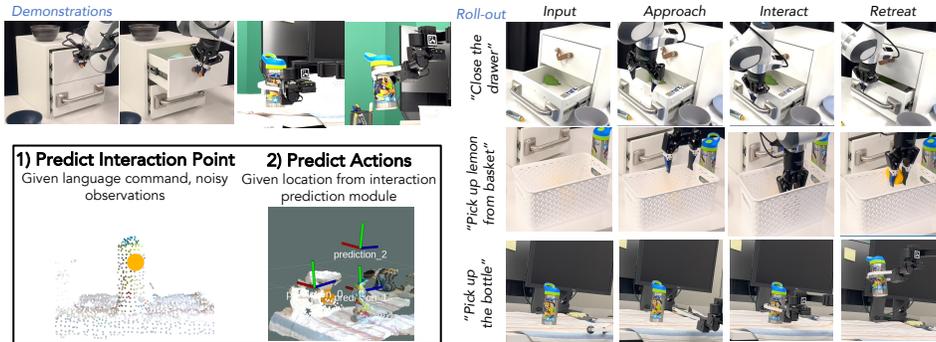


Figure 2: Spatial Language Attention Policies (SLAP) learn language-conditioned skills from few demonstrations in a wide variety of cluttered scenes. SLAP has two components: an “interaction prediction” module which localizes relevant features in a scene, and an “action prediction” module which uses local context to predict an executable action.

65 Compared to previous works, some recent studies focus on unstructured point clouds [11, 16]. These  
 66 approaches demonstrate improved data efficiency and performance compared to traditional behavior cloning.  
 67 For instance, Where2Act [11] and VAT-Mart [16] predict interaction trajectories, while  
 68 UMPNet [17] supports closed-loop 6DoF trajectories. They share a common framework: a generalizable  
 69 method to predict the interaction location and then predict local motion for the robot.

70 **Training Quickly with Attention-Based Policies.** CLIPort [7] and PerAct [2] are attention-based  
 71 policies similar to Transporter Nets [6]. While fitting our definition of attention-based policies,  
 72 they confine their workspace, use a rigid grid-like structure and treat action prediction as a discrete  
 73 classification. While still a limited workspace, SPOT [12], demonstrated the usefulness of 2D  
 74 attention-based policies for fast RL training, including sim-to-real transfer, and Zeng et al. [6] have  
 75 shown these policies are valuable for certain real-world tabletop tasks like kitting.

76 **Manipulation of Unknown Objects.** Manipulation of unknown objects includes segmentation [18,  
 77 19], grasping [20, 21], placement [22], and multi-object rearrangement either from a goal image [23,  
 78 24] or from language instructions [25, 26]. These approaches rely, generally, on first segmenting  
 79 relevant objects out, and then predicting how to grasp them and where to move them using separate  
 80 purpose-built models, including for complex task and motion planning [27].

81 **Language and Robotics.** Language is a natural and powerful way to specify goals for multi-task  
 82 robot systems. Several recent works [10, 1, 28] use a large-language model for task planning to combine  
 83 sequential low-level skills and assume to learn the low-level skills with IL or RL. To realistically  
 84 handle language task diversity, we need to learn these skills quickly. SLAP is more sample-efficient  
 85 than prior IL or RL approaches. In PaLM-E [3], textual and multi-modal tokens are interleaved as  
 86 inputs to the Transformer for handling language and multimodal data to generate high-level plans  
 87 for robotics tasks. Our approach is a spatial extension of this strategy.

88 **Language for Low-Level Skills.** A number of works have shown how to learn low-level language-  
 89 conditioned skills, e.g. [7, 2, 1, 29]. Like our work, Mees et al. [29] predicts 6DoF end effector goal  
 90 positions end-to-end and sequences them with large language models. They predict a 2D affordance  
 91 heatmap and depth from RGB; We do not predict depth, but specifically look at robustness and  
 92 generalization, where theirs is trained from play data in mostly-fixed scenes. Shridhar et al. [2]  
 93 predict a 3D voxelized grid and show strong real-world performance with relatively few examples,  
 94 but don’t look at out-of-domain generalization and are limited to a coarse voxelization of the world.

### 95 3 Approach

96 Most manipulation tasks necessarily involve interacting with environment objects [11]. We define an  
 97 ‘atomic skill’ as a task that can be specified by an interaction point, and a sequence of relative offsets  
 98 from this interaction point. For example, `pick('mug')` is an atomic skill as it can be defined in

99 terms of an interaction point on the ‘mug’ and subsequent relative waypoints for approach, grasp,  
 100 and lift actions. Similarly, `open('drawer')` is an atomic skill for which the interaction point is  
 101 on the drawer handle, and relative waypoints from it can be defined for approach, grasp, and pull.

102 We train a two-phase language-conditioned policy  $\pi(x, l)$ , which takes visual observation  $x$  and  
 103 a language command  $l$  as inputs and predicts an *interaction point*  $p_I$ , as well as a set of *relative*  
 104 *motions*, which are offsets from this point, instead of absolute coordinates. However, any realistic  
 105 task given to a home robot by a user typically involves more than one atomic skill. Our system breaks  
 106 down a high-level natural language task description ( $\mathcal{T}$ ) into a sequence of atomic skill descriptions  
 107  $\{l_j\}$  and uses them to condition the atomic skill motion policies. Our full paradigm is as follows:

$$\begin{aligned} \mathcal{T} &\rightarrow \{l_0, \dots, l_n\} \rightarrow \{\pi_j(x_j, l_j)\}_n \\ \forall j \in n, \quad \pi_j &:= (\pi_I, \pi_R), \text{ where:} \\ \pi_I(x_j, l_j) &\rightarrow p_I \quad (3D \text{ interaction point}) \\ \pi_R(x_j, l_j, p_I) &\rightarrow \{\mathbf{a}\}_m \quad (\text{sequence of actions}) \end{aligned}$$

108 The interaction point  $p_I$  is predicted by an **Interaction Prediction Module**  $\pi_I$ , and the continuous  
 109 component of the action by a **Relative Action Module**  $\pi_R$ . The Interaction Prediction Module  $\pi_I$   
 110 predicts *where the robot should attend to*; it is a specific location in the world, where the robot will  
 111 be interacting with the object as a part of its skill, as shown in Fig. 2.  $\pi_R$  predicts a relative action  
 112 sequence with respect to this contact point in the Cartesian space. These actions are then provided  
 113 as input to a low-level controller to execute the trajectory. These models are trained using labeled  
 114 expert demonstrations; a complete overview of the training process is shown in Fig. 4. Overall, the  
 115 system outputs a sequence of end-effector actions  $\mathbf{a}$ .

### 116 3.1 Scene Representation

117 The input observation  $x$  is a structured point-cloud (PCD) in the robot’s base-frame, constructed  
 118 by combining the inputs from a sequence of pre-defined scanning actions. This point cloud is then  
 119 preprocessed by voxelizing at a 1mm resolution to remove duplicate points from overlapping camera  
 120 views. The pointcloud is then used as input into both  $\pi_I$  and  $\pi_R$ .

121 For  $\pi_I$ , we perform a second voxelization, this time at 5mm resolution. This creates the down-  
 122 sampled set of points  $P$ , such that the interaction point  $\hat{p}_I \in P$ . This means  $\pi_I$  has a consistently  
 123 high-dimensional input and action space - for a robot looking at its environment with a set of  $N$   
 124 aggregated observations, this can be 5000-8000 input “tokens” representing the scene.

125 While SLAP discretizes the world similar to prior work [30, 31, 2], we can do so selectively, at a  
 126 higher resolution, and capture fine local details even in large scenes. We couple this with a set-based  
 127 learning formulation which allows us to attend to fine details in a data-efficient manner.

### 128 3.2 Interaction Prediction Module

129 We use our insight about tasks being shaped around an interaction point to make learning more  
 130 robust and more efficient: instead of predicting the agent’s motion directly, we formulate our  $\pi_I$   
 131 to solely focus on predicting a specific point  $p_I \in P$ , representing a single 5mm voxel that is  
 132 referred to as the “interaction point”. This formulation is akin to learning object affordance and  
 133 can be thought of as similar to prior work like Transporter Nets in 2D [6]. We hypothesize that  
 134 predicting attention directly on visual features, even for manipulation actions, will make SLAP  
 135 more general. We use a PerceiverIO [32] backbone to process the data, based on prior work on  
 136 language-conditioned real-world policies [2].

137 We first pass our input point cloud through two *modified set abstraction* layers [4] which result in a  
 138 sub-sampled point-cloud with each point’s feature capturing the local spatial structure around it at  
 139 two different resolutions. This encourages the classifier to pay attention to *local structures* rather  
 140 than a specific point that may not be visible in real-world settings. We concatenate the CLIP [33]  
 141 tokenized natural language command with the encoded point cloud to create an input sequence.

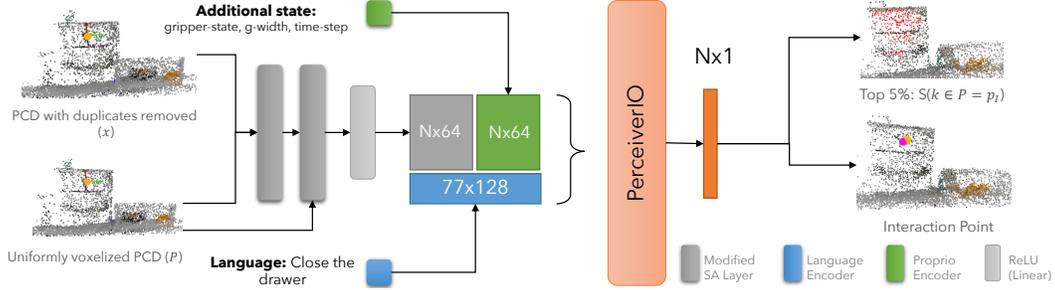


Figure 3: An overview of the architecture of the interaction prediction module. The point cloud is downsampled to remove duplicates and encoded using two modified set-abstraction layers. The SA layers generate a local spatial embedding which is concatenated with proprioceptive features - in our case, the current gripper state. Both spatial and language features are concatenated and input into a PerceiverIO transformer backbone. We then predict an interaction score per spatial feature and the  $\text{argmax}$  is chosen as the interaction site for command  $l$ .

142 Each point  $i \in P$  in the point cloud is assigned a score with respect to task  $\tau_j$  which results in the  
 143 interaction point for that task,  $p_I^j := \underset{x,y,z}{\text{argmax}}(S(i = p_I^j | l, x, P, \mathcal{D}^j))$ , where  $\mathcal{D}^j$  is the set of expert  
 144 demonstrations provided for task  $\tau_j$ . The IPM architecture overview is provided in Fig. 3. Note we  
 145 also use semantic features from Detic in the Stretch experiment for training SLAP as an additional  
 146 feature channel apart from the color-channels.

147 **Modified Set Abstraction Layer.** The default SA layer as introduced by Qi et al. [4] uses farthest  
 148 point sampling (FPS) to determine which locations feature vectors are created. FPS ensures that  
 149 subsampled point-cloud is a good representation of a given scene, without any guarantees about the  
 150 granularity. However, it’s very sensitive to the number of points selected - in most PointNet++-based  
 151 policies, a fixed number of points are chosen using FPS [4]. However, SLAP must adapt to scenes  
 152 of varying sizes, possibly with multiple views, and still not miss small details.

153 We propose an alternative PointNet++ set abstraction layer, which computes embeddings based on  
 154 the original and an *evenly* downsampled version of the point-cloud,  $P$ . This results in a denser  
 155 spatial embedding by considering a subset of all points within a certain radius of each-point in the  
 156 downsampled point-cloud. This downsampled set of points guarantees we can attend to even small  
 157 features, and allows us to predict an interaction point  $p_I$  from the PointNet++ aggregated features.

### 158 3.3 Relative Action Module

159 The relative action module relies on the interaction point predicted by the classifier and operates on  
 160 a cropped point cloud,  $x_R$ , around this point to predict the actions associated with this sequence.  
 161 As in the interaction prediction module, the model uses a cascade of modified *set abstraction* layers  
 162 as the backbone to compute a multi-resolution encoding feature over the cropped point cloud. We  
 163 train three multi-head regressors (described further below) over these features to predict the actions  
 164 for the overall task. Specifically,  $\pi_R$  has three heads, one for each component of the relative action  
 165 space: gripper activation  $g$ , position offset  $\delta p$ , and orientation  $q$ . Our LSTM-based architecture  
 166 (details in B.1) can predict skills with variable number of actions (3,4 in our experiments).

167 Positions and orientations associated with the interaction action generally tend to be much closer to  
 168 the crop-center thus we train one model per action to encourage each action to be learned according  
 169 to its own distribution. Also note that the cropped input point-cloud is not perfectly centered at the  
 170 ground truth interaction point  $\hat{p}_I$ , but rather with some noise added:  $\hat{p}_I' = \hat{p}_I + \mathcal{N}(0, \sigma)$ . This is  
 171 done to force the action predictor to be robust to sub-optimal interaction point predictions by the  
 172 interaction predictor module during real-world roll-outs. Thus, for each part of the action sequence,  
 173 the keyframe position is calculated as:  $p = p_I + \delta p$ . When acting, the robot will move to  $(p, q)$  via  
 174 a motion planner, and then will send a command to the gripper to set its state to  $g$ .

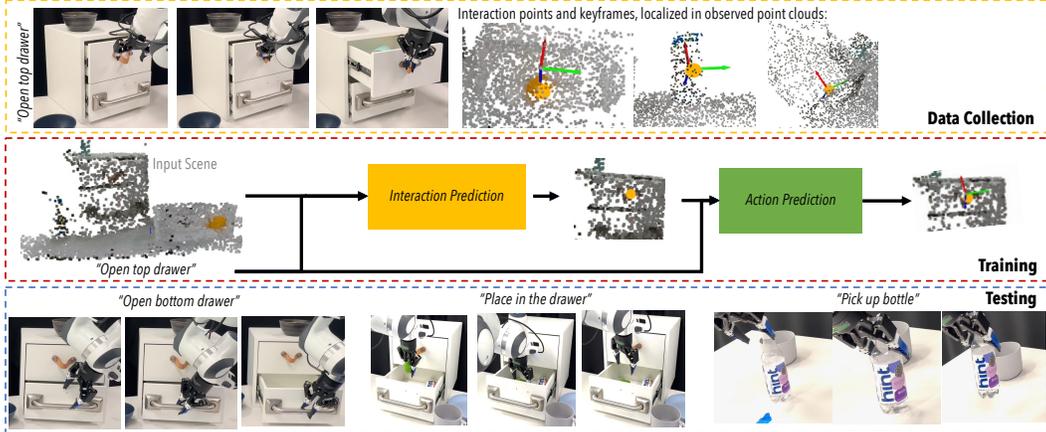


Figure 4: Illustration of the complete process for training SLAP. Demonstrations are collected and used to train the Interaction Prediction module and the Action Prediction Module separately.

### 175 3.4 Training SLAP

176 To collect data, an expert operator guides the robot through a trajectory, pressing a button to record  
 177 *keyframes* representing crucial parts of a task. At each keyframe, we record the associated expert  
 178 action  $\hat{a} = (\delta p, \hat{q}, \hat{g})$ . We assume that low-level controllers exist - in our case, we use Polymetis [34]  
 179 for the Franka arm and Hello Robot’s controllers<sup>1</sup> for Stretch. Example tasks are shown in Fig. 6.

180 **Interaction Prediction Module.** We train  $\pi_I$  with a cross-entropy loss, predicting the interaction  
 181 point  $p_I$  from the downsampled set of coarse voxels  $P$ . We additionally apply what we call a  
 182 *locality loss* ( $L_{loc}$ ), as per prior work [35]. Conceptually, we want to penalize points the further they  
 183 are from the contact point, both to encourage learning relevant features as well as to aid in ignoring  
 184 distractors. To achieve this, we define the locality loss as:  $L_{loc} = \sum_{k \in P} \text{softmax}(f_k) \|\hat{p}_I, k\|^2$ ,  
 185 where  $f_k$  is the output of the transformer for point  $k \in P$ . The *softmax* turns  $f_k$  into attention  
 186 over the points, meaning that  $L_{loc}$  can be interpreted as a weighted average of the square distances.  
 187 Points further from  $\hat{p}_I$  are therefore encouraged to have lower classification scores. Combining our  
 188 two losses, we obtain  $L_I = CE(P, \hat{p}_I) + \frac{w}{|P|} L_{loc}$ , where  $w$  is a scaling constant that implicitly  
 189 defines how much spread to allow in our points.

190 **Relative Action Module.** To train  $\pi_R$ , we use the weighted sum of three different losses. We train  
 191  $a = (p, q, g) = \pi_R(x_R)$  with an L2 loss over the  $\delta p$ , a quaternion distance metric for the loss on  $q$   
 192 based on prior work [36] and binary cross-entropy loss for gripper action classification (Sec. A.3).

### 193 3.5 Task Planning

194 Consider a natural language instruction from a user such as ‘put away the bottle on a table’. We  
 195 decompose it to a sequence of atomic skills as: `goto(‘bottle’)`, `pick_up(‘bottle’)`,  
 196 `goto(‘table’)`, and `place_on(‘table’)`. We programmatically create natural language  
 197 and code templates for 16 task families and generated a dataset of 500k samples. We use  
 198 LLaMA [37] models for in-context learning [38, 39] and adapter fine-tuning [40] to learn the map-  
 199 ping between natural language task instructions to the corresponding sequence of atomic skills.

## 200 4 Experiments

201 We report the success rate of our model for 8 real-world manipulation tasks in Table 2, and compare  
 202 it against prior baselines trained using the same labeling scheme. Overall, we see an improvement of  
 203 1.6x over our best comparative baseline, PerAct [2]. We test each model under two different condi-  
 204 tions: *Seen setting assumption*; i.e. those with seen distractor objects and objects placed roughly in

<sup>1</sup>[https://github.com/hello-robot/stretch\\_ros](https://github.com/hello-robot/stretch_ros)

Skill Name	Seen		Unseen	
	PerAct	SLAP	PerAct	SLAP
Open bottom drawer	00%	<b>80%</b>	00%	<b>60%</b>
Open top drawer	60%	<b>80%</b>	<b>40%</b>	<b>40%</b>
Close drawer	<b>100%</b>	<b>100%</b>	<b>40%</b>	<b>40%</b>
Pick lemon from basket	60%	<b>80%</b>	10%	<b>40%</b>
Pick bottle	<b>60%</b>	<b>60%</b>	<b>60%</b>	40%
Place into the drawer	60%	<b>80%</b>	40%	<b>60%</b>
Place into the basket	40%	<b>100%</b>	10%	<b>60%</b>
Place into the bowl	40%	<b>60%</b>	00%	<b>40%</b>
Average Success Rate	50%	80%	27.5%	47.5%
Improvement		1.6x		1.7x

Table 2: SLAP and PerAct [2] performance on real world Franka manipulation tasks. We evaluate both seen scenes (seen object positions and distractors), but in different arrangements, and unseen scenes with previously-unseen object positions and distractors. SLAP is notably better overall in both conditions.

205 the same range of positions and orientations as in the training data in any relative arrangement (in-  
 206 cluding unseen). Second, we test under *unseen setting assumptions*; i.e. those with unseen distractor  
 207 objects and the implicated object placed significantly out of the range of positions and orientations  
 208 already seen. We run 5 tests per scene setting per skill per model and report the percentage suc-  
 209 cess numbers in Table 2. We compare our model against Perceiver-Actor (PerAct) [2]. We train  
 210 each model for the same number of training steps and choose the SLAP model based on the best  
 211 validation loss. For PerAct, we use the last checkpoint, per their testing practices [2].

212 We also run a per-skill evaluation of SLAP on Stretch under the *unseen setting assumption* (see  
 213 Fig. 3) accomplished by adding unseen distractor objects to the scene and moving the robot base  
 214 position within reachable distance of the object. Note demonstrations were taken on a different  
 215 robot than the one policies were deployed on.

#### 216 4.1 Longitudinal Task Execution on Stretch

217 We trained a multi-task model for the Stretch robot for five skills using 10  
 218 demonstrations each. This model was deployed in an end-to-end system  
 219 which operates over code-list generated by a task-planner (as in Sec. 3.5).  
 220 We ran 5 prompts end-to-end with 4 to 8 skills each, using ground truth plans  
 221 - we verify the viability of generating these task plans in §4.2. These exper-  
 222 iments are done under the *unseen setting* with the robot spawning anywhere  
 223 with respect to the objects. For fair evaluation in low-data regime, we add  
 224 some structure by specifying orthogonal viewing direction for objects. Once  
 225 the agent finds the object of interest it fires an initial prediction using SLAP  
 226 to find most promising interaction point. This prediction happens under any  
 227 dynamic viewing angle of the object (we assume the robot can see the object). This dynamic pre-  
 228 diction and pre-programmed viewing angle is used to *approach* the object at an orthogonal viewing  
 229 angle where the model fires an actionable prediction for the full skill execution. We observe adding  
 230 semantic features from Detic significantly improves IPM performance with unseen distractors (80%  
 231 against 47.5% in Table 2) however we see failures when relative position is significantly perturbed.

#### 232 4.2 Task planning with in-context learning and fine-tuning LLaMa

233 Previous work has shown the strength of language models as zero-shot planners [41], a result  
 234 strengthened by improved techniques for “in-context learning” or prompting [42]. To verify that  
 235 models can produce task plans with the skills we defined, we experiment with both in-context learn-  
 236 ing (IC) [43] of LLaMA [37] and adapter fine-tuning (FT) [40].

Skill Name	SLAP (5 tries)
Open Drawer	60%
Close Drawer	100%
Take bottle	80%
Pour into bowl	80%

Table 3: SLAP on a mobile manipulator using a multi-task model across 4 skills. With semantic predictions added to our feature space, we see the model perform better on unseen scenes with new distractors and unseen relative position of the robot with respect to the scene

	Success
Total	68.5
Heuristic	66.0
Learned	80.0

Table 1: End-to-end performance. Learned skills outperform heuristics except due to Detic failures.



Figure 6: Examples of tasks executed on a Franka arm through our trained model in a clean setting. We trained numerous tasks (left) and tested on both seen and unseen objects (right).

					Task	Lat.
	LLaMa	Verb	Noun	Acc.	Corr.	(sec.)
237	IC 7B	83	81	76	61	16.4
238	IC 30B	81	81	76	62	27.6
239						
240	FT 7B	100	98	99	91	19.5
241						
242						

Table 5 presents the models’ verb, noun, and combined accuracies. Task Correctness is a binary score if the entire prediction was correct, and finally, latency is measured in seconds on a single A6000 with 16 GB RAM.<sup>2</sup> High Task Correctness from a small model verifies the compatibility of our skills with LLM task planning.

Figure 5: Fine-tuning (FT) outperforms in-context learning (IC) for same latency.

### 243 4.3 Ablations

244 **Hybrid vs Monolithic Architecture (Table-top).** For the same number of epochs, SLAP does  
 245 better than PerAct on 6 of 8 tasks when tested in in-distribution setting and 5 of 8 tasks when tested  
 246 in out-of-distribution settings. PerAct performs equally well as our model for 2 of 8 tasks on our  
 247 in-distribution scenes. Similarly, for our “hard” generalization scenes, PerAct performed equally  
 248 well in two cases, and actually outperformed SLAP when picking up a bottle. In failure cases,  $\pi_R$   
 249 predicted the correct trajectory, but not with respect to the right part of the object.

250 **Unseen Scene Generalization.** We see a drop in the success rate for both PerAct and SLAP when  
 251 tested on out-of-distribution settings. PerAct would often predict the correct approach actions, but  
 252 then it would fail to grasp accurately. With SLAP, however, we saw that  $p_I$  was predicted fairly  
 253 accurately, but the regressor would fail for out-of-distribution object placements specifically because  
 254 of bad orientation prediction. When  $\pi_I$  failed, it was because the position and orientation of the  
 255 target object was dramatically different, *and* unseen distractors confused it. We see better results for  
 256 SLAP under Stretch setting due to the addition of semantic features from Detic.

## 257 5 Conclusion

258 We proposed a method for learning visual-language policies for decision making in complex envi-  
 259 ronments. SLAP is a novel architecture which combines the *structure* of a point-cloud based input  
 260 with *semantics* from language and accompanying demonstrations to predict continuous end-effector  
 261 actions for manipulation tasks. We demonstrate SLAP on two hardware platforms, including an  
 262 end-to-end evaluation on a mobile manipulator, something not present in prior work.

### 263 5.1 Limitations

264 SLAP has high variance in out-of-distribution situations, resulting in complete failure if  $\pi_I$  fails to  
 265 correctly identify the context. For  $\pi_R$ , multimodal or noisy data still poses issues; replacing  $\pi_R$   
 266 with a policy which can better handle this data, e.g. Diffusion Policies [45]. Overall system has  
 267 multiple points-of-failure due to heuristic policies, unaligned language and vision models; end-to-  
 268 end trainable architectures and cross-modal alignment could help.

<sup>2</sup>Adaptor fine-tuning increases the model size by  $\sim 6\%$ , which accounts for the additional latency compared to IC. We use standard inference libraries so results are comparable, but not optimized for runtime [44].

## 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. *arXiv preprint arXiv:2212.06817*, 2022.
- [2] M. Shridhar, L. Manuelli, and D. Fox. Perceiver-actor: A multi-task transformer for robotic manipulation. *arXiv preprint arXiv:2209.05451*, 2022.
- [3] 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. *arXiv preprint arXiv:2303.03378*, 2023.
- [4] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [5] W. Yuan, C. Paxton, K. Desingh, and D. Fox. Sornet: Spatial object-centric representations for sequential manipulation. In *Conference on Robot Learning*, pages 148–157. PMLR, 2022.
- [6] A. Zeng, P. Florence, J. Tompson, S. Welker, J. Chien, M. Attarian, T. Armstrong, I. Krasin, D. Duong, V. Sindhwani, et al. Transporter networks: Rearranging the visual world for robotic manipulation. In *Conference on Robot Learning*, pages 726–747. PMLRG, 2021.
- [7] M. Shridhar, L. Manuelli, and D. Fox. Cliport: What and where pathways for robotic manipulation. In *Conference on Robot Learning*, pages 894–906. PMLR, 2022.
- [8] C. Huang, O. Mees, A. Zeng, and W. Burgard. Visual language maps for robot navigation. *arXiv preprint arXiv:2210.05714*, 2022.
- [9] S. James and A. J. Davison. Q-attention: Enabling efficient learning for vision-based robotic manipulation. *IEEE Robotics and Automation Letters*, 7(2):1612–1619, 2022.
- [10] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, K. Gopalakrishnan, K. Hausman, A. Herzog, et al. Do as i can, not as i say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [11] K. Mo, L. J. Guibas, M. Mukadam, A. Gupta, and S. Tulsiani. Where2act: From pixels to actions for articulated 3d objects. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6813–6823, 2021.
- [12] A. Hundt, B. Killeen, N. Greene, H. Wu, H. Kwon, C. Paxton, and G. D. Hager. ““Good robot!”: Efficient reinforcement learning for multi-step visual tasks with sim to real transfer. *IEEE Robotics and Automation Letters*, 5(4):6724–6731, 2020.
- [13] P.-L. Guhur, S. Chen, R. Garcia, M. Tapaswi, I. Laptev, and C. Schmid. Instruction-driven history-aware policies for robotic manipulations. *arXiv preprint arXiv:2209.04899*, 2022.
- [14] X. Zhou, R. Girdhar, A. Joulin, P. Krähenbühl, and I. Misra. Detecting twenty-thousand classes using image-level supervision. In *ECCV*, 2022.
- [15] I. Misra, R. Girdhar, and A. Joulin. An end-to-end transformer model for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2906–2917, 2021.
- [16] R. Wu, Y. Zhao, K. Mo, Z. Guo, Y. Wang, T. Wu, Q. Fan, X. Chen, L. Guibas, and H. Dong. Vat-mart: Learning visual action trajectory proposals for manipulating 3d articulated objects. *arXiv preprint arXiv:2106.14440*, 2021.
- [17] Z. Xu, Z. He, and S. Song. Umpnet: Universal manipulation policy network for articulated objects. *arXiv preprint arXiv:2109.05668*, 2021.

- 312 [18] C. Xie, Y. Xiang, A. Mousavian, and D. Fox. Unseen object instance segmentation for robotic  
313 environments. *IEEE Transactions on Robotics*, 37(5):1343–1359, 2021.
- 314 [19] Y. Xiang, C. Xie, A. Mousavian, and D. Fox. Learning rgb-d feature embeddings for unseen  
315 object instance segmentation. In *Conference on Robot Learning*, pages 461–470. PMLR, 2021.
- 316 [20] A. Murali, A. Mousavian, C. Eppner, C. Paxton, and D. Fox. 6-dof grasping for target-driven  
317 object manipulation in clutter. In *2020 IEEE International Conference on Robotics and Au-*  
318 *tomation (ICRA)*, pages 6232–6238. IEEE, 2020.
- 319 [21] M. Sundermeyer, A. Mousavian, R. Triebel, and D. Fox. Contact-graspnet: Efficient 6-dof  
320 grasp generation in cluttered scenes. In *2021 IEEE International Conference on Robotics and*  
321 *Automation (ICRA)*, pages 13438–13444. IEEE, 2021.
- 322 [22] C. Paxton, C. Xie, T. Hermans, and D. Fox. Predicting stable configurations for semantic  
323 placement of novel objects. In *Conference on Robot Learning*, pages 806–815. PMLR, 2022.
- 324 [23] A. H. Qureshi, A. Mousavian, C. Paxton, M. C. Yip, and D. Fox. Nerp: Neural rearrangement  
325 planning for unknown objects. *arXiv preprint arXiv:2106.01352*, 2021.
- 326 [24] A. Goyal, A. Mousavian, C. Paxton, Y.-W. Chao, B. Okorn, J. Deng, and D. Fox. Ifor: Iter-  
327 ative flow minimization for robotic object rearrangement. In *Proceedings of the IEEE/CVF*  
328 *Conference on Computer Vision and Pattern Recognition*, pages 14787–14797, 2022.
- 329 [25] W. Liu, C. Paxton, T. Hermans, and D. Fox. Structformer: Learning spatial structure for  
330 language-guided semantic rearrangement of novel objects. In *2022 International Conference*  
331 *on Robotics and Automation (ICRA)*, pages 6322–6329. IEEE, 2022.
- 332 [26] W. Liu, T. Hermans, S. Chernova, and C. Paxton. Structdiffusion: Object-centric diffusion for  
333 semantic rearrangement of novel objects. *arXiv preprint arXiv:2211.04604*, 2022.
- 334 [27] A. Curtis, X. Fang, L. P. Kaelbling, T. Lozano-Pérez, and C. R. Garrett. Long-horizon manip-  
335 ulation of unknown objects via task and motion planning with estimated affordances. In *2022*  
336 *International Conference on Robotics and Automation (ICRA)*, pages 1940–1946. IEEE, 2022.
- 337 [28] K. Lin, C. Agia, T. Migimatsu, M. Pavone, and J. Bohg. Text2motion: From natural language  
338 instructions to feasible plans. *arXiv preprint arXiv:2303.12153*, 2023.
- 339 [29] O. Mees, J. Borja-Diaz, and W. Burgard. Grounding language with visual affordances over  
340 unstructured data. *arXiv preprint arXiv:2210.01911*, 2022.
- 341 [30] V. Blukis, C. Paxton, D. Fox, A. Garg, and Y. Artzi. A persistent spatial semantic representation  
342 for high-level natural language instruction execution. In *Conference on Robot Learning*, pages  
343 706–717. PMLR, 2022.
- 344 [31] S. Y. Min, D. S. Chaplot, P. Ravikumar, Y. Bisk, and R. Salakhutdinov. Film: Following  
345 instructions in language with modular methods. *arXiv preprint arXiv:2110.07342*, 2021.
- 346 [32] A. Jaegle, S. Borgeaud, J.-B. Alayrac, C. Doersch, C. Ionescu, D. Ding, S. Koppula, D. Zoran,  
347 A. Brock, E. Shelhamer, et al. Perceiver io: A general architecture for structured inputs &  
348 outputs. *arXiv preprint arXiv:2107.14795*, 2021.
- 349 [33] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell,  
350 P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervi-  
351 sion. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021.
- 352 [34] Y. Lin, A. S. Wang, G. Sutanto, A. Rai, and F. Meier. Polymetis. [https://](https://facebookresearch.github.io/fairo/polymetis/)  
353 [facebookresearch.github.io/fairo/polymetis/](https://facebookresearch.github.io/fairo/polymetis/), 2021.
- 354 [35] S. Powers, A. Gupta, and C. Paxton. Evaluating continual learning on a home robot, 2023.

- 355 [36] C. Paxton, Y. Bisk, J. Thomason, A. Byravan, and D. Foxl. Prospection: Interpretable plans  
356 from language by predicting the future. In *2019 International Conference on Robotics and  
357 Automation (ICRA)*, pages 6942–6948. IEEE, 2019.
- 358 [37] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière,  
359 N. Goyal, E. Hambro, F. Azhar, A. Rodriguez, A. Joulin, E. Grave, and G. Lample. Llama:  
360 Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- 361 [38] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal, A. Neelakantan,  
362 P. Shyam, G. Sastry, A. Askell, S. Agarwal, A. Herbert-Voss, G. Krueger, T. Henighan,  
363 R. Child, A. Ramesh, D. M. Ziegler, J. Wu, C. Winter, C. Hesse, M. Chen, E. Sigler,  
364 M. Litwin, S. Gray, B. Chess, J. Clark, C. Berner, S. McCandlish, A. Radford, I. Sutskever,  
365 and D. Amodei. Language models are few-shot learners. *ArXiv*, 2020.
- 366 [39] E. Akyürek, D. Schuurmans, J. Andreas, T. Ma, and D. Zhou. What learning algorithm is  
367 in-context learning? investigations with linear models, 2023.
- 368 [40] E. Hu, Y. Shen, P. Wallis, Z. Allen-Zhu, Y. Li, L. Wang, and W. Chen. Lora: Low-rank  
369 adaptation of large language models, 2021.
- 370 [41] W. Huang, P. Abbeel, D. Pathak, and I. Mordatch. Language models as zero-shot planners: Ex-  
371 tracting actionable knowledge for embodied agents. *arXiv preprint arXiv:2201.07207*, 2022.
- 372 [42] O. Ram, Y. Levine, I. Dalmedigos, D. Muhlgay, A. Shashua, K. Leyton-Brown, and Y. Shoham.  
373 In-Context Retrieval-Augmented Language Models, Jan. 2023. URL [http://arxiv.org/  
374 abs/2302.00083](http://arxiv.org/abs/2302.00083). arXiv:2302.00083 [cs].
- 375 [43] J. Wei, J. Wei, Y. Tay, D. Tran, A. Webson, Y. Lu, X. Chen, H. Liu, D. Huang, D. Zhou, and  
376 T. Ma. Larger language models do in-context learning differently, Mar. 2023. URL [http:  
377 //arxiv.org/abs/2303.03846](http://arxiv.org/abs/2303.03846). arXiv:2303.03846 [cs].
- 378 [44] J. Fernandez, J. Kahn, C. Na, Y. Bisk, and E. Strubell. The Framework Tax: Disparities  
379 Between Inference Efficiency in Research and Deployment. *ArXiv*, 2023. URL [https:  
380 //arxiv.org/abs/2302.06117](https://arxiv.org/abs/2302.06117).
- 381 [45] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song. Diffusion policy:  
382 Visuomotor policy learning via action diffusion. *arXiv preprint arXiv:2303.04137*, 2023.
- 383 [46] C. C. Kemp, A. Edsinger, H. M. Clever, and B. Matulevich. The design of stretch: A com-  
384 pact, lightweight mobile manipulator for indoor human environments. In *2022 International  
385 Conference on Robotics and Automation (ICRA)*, pages 3150–3157. IEEE, 2022.
- 386 [47] B. Akgun, M. Cakmak, J. W. Yoo, and A. L. Thomaz. Trajectories and keyframes for kines-  
387 thetic teaching: A human-robot interaction perspective. In *Proceedings of the seventh annual  
388 ACM/IEEE international conference on Human-Robot Interaction*, pages 391–398, 2012.
- 389 [48] A. Handa, A. Allshire, V. Makoviychuk, A. Petrenko, R. Singh, J. Liu, D. Makoviichuk,  
390 K. Van Wyk, A. Zhurkevich, B. Sundaralingam, et al. Dextreme: Transfer of agile in-hand  
391 manipulation from simulation to reality. *arXiv preprint arXiv:2210.13702*, 2022.
- 392 [49] J. Mahler, J. Liang, S. Niyaz, M. Laskey, R. Doan, X. Liu, J. A. Ojea, and K. Goldberg. Dex-  
393 net 2.0: Deep learning to plan robust grasps with synthetic point clouds and analytic grasp  
394 metrics. *arXiv preprint arXiv:1703.09312*, 2017.
- 395 [50] N. M. M. Shafiullah, C. Paxton, L. Pinto, S. Chintala, and A. Szlam. Clip-fields: Weakly  
396 supervised semantic fields for robotic memory. *arXiv preprint arXiv:2210.05663*, 2022.
- 397 [51] B. Bolte, A. Wang, J. Yang, M. Mukadam, M. Kalakrishnan, and C. Paxton. Usa-net: Unified  
398 semantic and affordance representations for robot memory. *arXiv preprint arXiv:2304.12164*,  
399 2023.

# 401 Appendix

## Table of Contents

---

404	<b>A Training</b>	<b>12</b>
405	A.1 Data Collection and Annotation . . . . .	12
406	A.2 Data Processing . . . . .	12
407	A.3 Action Prediction Losses . . . . .	13
408	A.4 Skill Weighting . . . . .	13
409	<b>B Relative Action Module</b>	<b>14</b>
410	B.1 Relative Action Model: MLP Implementation . . . . .	14
411	<b>C Tasks</b>	<b>15</b>
412	<b>D Skills</b>	<b>15</b>
413	D.1 In vs. Out Of Distribution . . . . .	15
414	D.2 Language Annotations . . . . .	17
415	D.3 Out of distribution Results from SLAP . . . . .	17
416	D.4 Motion Planning Failures . . . . .	17
417	<b>E Additional Analysis</b>	<b>19</b>
418	E.1 Visualizing the Learned Attention . . . . .	19
419	E.2 Language Generalization . . . . .	19
420	<b>F Additional Related Work</b>	<b>19</b>

---

## 424 A Training

425 Below is expanded information on our training, algorithm, and data processing to improve repro-  
426 ducibility.

### 427 A.1 Data Collection and Annotation

428 When collecting an episode with the Franka arm, we first scan the scene with a pre-defined list of  
429 scanning positions to collect an aggregated  $x$ . In our case, we make no assumption as to what or how  
430 many these are, or how large the resulting input point cloud  $x$  is. With the Hello Robot Stretch [46],  
431 we collect data based on exactly where the robot is looking.

432 Then, we collect demonstration data using kinesthetic teaching for the Franka arm (demonstrator  
433 physically moves the robot) and via controller teleoperation for the Stretch robot. The demonstrator  
434 moves the arm through the trajectory associated with each task, explicitly recording the *keyframes*  
435 [47] associated with action execution. These represent the salient moments within a task – the  
436 bottlenecks in the tasks’ state space, which can be connected by our low-level controller.

### 437 A.2 Data Processing

438 We execute each individual skill open-loop based on an initial observation. We use data augmenta-  
439 tion to make sure even with relatively few examples, we still see good generalization performance.

440 **Data Augmentation.** Prior work in RGB-D perception for robotic manipulation (e.g. [18, 48]) has  
441 extensively used a variety of data augmentation tricks to improve real-world performance. In this

work, we use three different data augmentation techniques to randomize the input scene  $x$  used to train  $p_I = \pi_I(x, l)$ :

- *Elliptical dropout*: Random ellipses are dropped out from the depth channel to emulate occlusions and random noise, as per prior work [49, 18]. Number of ellipses are sampled from a Poisson distribution with mean of 10.
- *Multiplicative Noise*: Again as per prior work [49, 18, 22], we add multiplicative noise from a gamma process to the depth channel.
- *Additive Noise*: Gaussian process noise is added to the points in the point-cloud. Parameters for the Gaussian distribution are sampled uniformly from given ranges. This is to emulate the natural frame-to-frame point-cloud noise that occurs in the real-world.
- *Rotational Randomization*: Similar to prior work [2, 22, 25], we rotate our entire scene around the z-axis within a range of  $\pm 45$  degrees to help force the model to learn rotational invariance.
- *Random cropping*: with  $p = 0.75$ , we randomly crop to a radius around  $\hat{p}_I + \delta$ , where  $\delta$  is a random translation sampled from a Gaussian distribution. The radius to crop is randomly sampled in (1, 2) meters.

**Data Augmentation for  $\pi_R$ .** We crop the relational input  $x_R \subset x$  around the ground-truth  $p_I$ , using a fixed radius  $r = 0.1m$ . We implement an additional augmentation for learning our action model. Since  $p_I$  is chosen from the discretized set of downsampled points  $P$ , we might in principle be limited to this granularity of response. Instead, we randomly shift both  $p_I$  and the positional action  $\delta p$  by some uniformly-sampled offset  $\delta r \in \mathbb{R}^3$ , with up to  $0.025m$  of noise. This lets  $\pi_R$  adapt to interaction prediction errors of up to several centimeters.

### A.3 Action Prediction Losses

Following [36] for the orientation, we can compute the angle between two quaternions  $\theta$  as:

$$\theta = \cos^{-1}(2\langle \hat{q}_1, \hat{q}_2 \rangle^2). \quad (1)$$

We can remove the cosine component and use it as a squared distance metric between 0 and 1. We then compute the position and orientation loss as:

$$L_R = \lambda_p \|\delta p - \hat{\delta p}\|_2^2 + \lambda_q (1 - \langle \hat{q}, q \rangle) \quad (2)$$

where  $\lambda_p$  and  $\lambda_q$  are weights on the positional and orientation components of the loss, set to 1 and  $1e - 2$  respectively.

Predicting gripper action is a classification problem trained with a cross-entropy loss. For input we use the task’s language description embedding and proprioceptive information about the robot as input, i.e.  $s = (l, g_{act}, g_w, ts)$  where  $g_{act}$  is 1 if gripper is closed and 0 otherwise,  $g_w$  is the distance between fingers of the gripper and  $ts$  is the time-step. The gripper action loss is then:

$$L_g = \lambda_g CE(g, \hat{g}) \quad (3)$$

where  $\lambda_g$  is the weight on cross-entropy loss set to 0.0001. The batch-size is set at 1 for this implementation.

We train  $\pi_I$  and  $\pi_R$  separately for  $n = 85$  epochs. At each epoch, we compare validation performance to the current best - if validation did not improve, we reset performance to the last best model.

### A.4 Skill Weighting

In Stretch experiments, we used a wide range of skills with different error tolerances and corresponding variances. As a result,

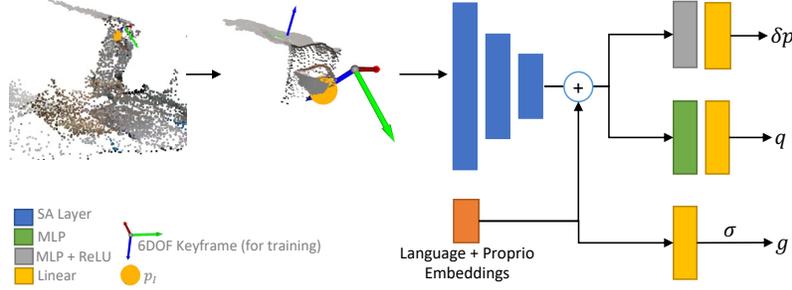


Figure 7: Regression model architecture with separate heads for each output. The point-cloud is cropped around the interaction point with some perturbation and passed to a cascade of set abstraction layers. Encoded spatial features are then concatenated with language and proprioception embeddings to predict position offset of action from interaction point, absolute orientation and gripper action as a boolean.

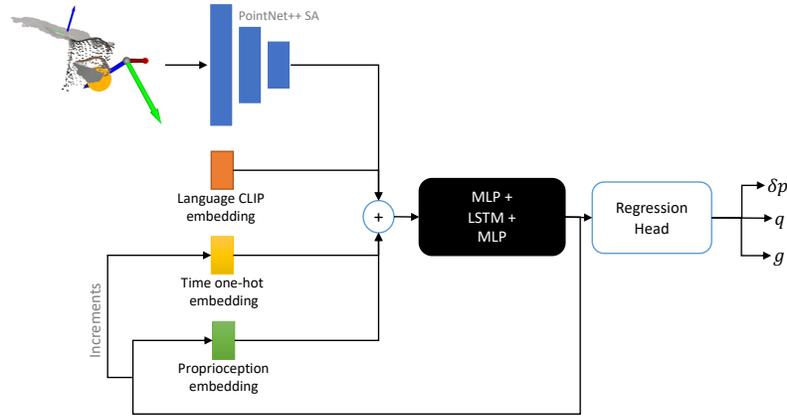


Figure 8: LSTM-based regression model architecture based on the regression head and PointNet++ embeddings introduced in Fig. 7. LSTM-based architecture shows higher stability in learning action distribution with wider distribution due to the conditioning effect.

## 482 B Relative Action Module

483 In our work, the Relative Action Module  $\pi_R$  is assumed to be some *local* policy which predicts  
 484 end-effector poses. In our case, we implement two different versions of this policy, one which was  
 485 used on the static Franka manipulator and one which was implemented on the Stretch. In both cases:

- 486 • The policy predicts an *end effector pose* relative to the predicted interaction point from  $\pi_I$
- 487 • The policy is conditioned on a *local crop* around this interaction point.

### 488 B.1 Relative Action Model: MLP Implementation

489 Fig. 7 gives an overview of the MLP version of the regression model. The model takes in the  
 490 cropped point cloud (augmented during training as discussed in Sec. A.2. We saw that injecting  
 491 random noise to the interaction point during training allowed the policy to, at test time, recover from  
 492 failures (because it predicted an interaction point *near* the correct area, instead of at the correct  
 493 position). However, we observed this architecture suffered when positional distribution of actions  
 494 varied widely with respect to the interaction point position across tasks. Thus we implemented a  
 495 LSTM version based on the MLP version which exhibited better performance in learning wider  
 496 action distribution, based on our initial experiments.



Figure 9: Within distribution objects used at training time and out-of-distribution objects introduced during testing in our experiments.

## 497 C Tasks

498 We generate a dataset with more than 500k samples using natural language descriptions and corre-  
 499 sponding atomic skill templates. We consider the following task families:

- 500 ‘bring x from y articulated’
- 501 ‘bring x from y articulated to pour in z’
- 502 ‘bring x from y articulated to wipe z’
- 503 ‘bring x from y surface’
- 504 ‘bring x from y surface to pour in z’
- 505 ‘bring x from y surface to pour in z then place on w surface’
- 506 ‘bring x from y surface to wipe z’
- 507 ‘move x from y articulated to z articulated’
- 508 ‘move x from y articulated to z surface’
- 509 ‘move x from y surface to z articulated’
- 510 ‘move x from y surface to z surface’
- 511 ‘take x from human to pour in z and place on y surface’
- 512 ‘take x from human to pour in z’
- 513 ‘take x from human to wipe z’
- 514 ‘take x from human to z articulated’
- 515 ‘take x from human to z surface’

516 For each task family, we define a corresponding template containing the sequence of atomic skills.  
 517 To populate these templates and generate the data, we create a list of more than 150 movable objects  
 518 kitchen objects, surfaces like `table`, `kitchen counter` and articulated objects like `drawer`,  
 519 `cabinets`. For `pour` skill, we create a list of “spillable” items such as `cup of coffee`, or  
 520 `bowl of jelly beans`. Similarly, for `wipe` skill, we have a list of items to wipe with such as  
 521 `sponge`, or `brush`.

## 522 D Skills

523 Here we refer to atomic skills learned by SLAP as simple tasks or “tasks”. This allows us to discuss  
 524 corresponding “actions” that are defined in terms of the relative offset from the interaction point.

### 525 D.1 In vs. Out Of Distribution

526 We used a number of objects for our Franka manipulation experiments, which included both in- and  
 527 out-of-distribution objects. One goal of SLAP is to show that our methods generalize much better  
 528 than others to different types of scenes and different levels of clutter.



Figure 10: Seen objects and unseen distractors used in longitudinal experiments with Stretch.

529 Every real-world task scene had a sub-sample of all within-distribution objects.

- 530 1. Open the top drawer
- 531 • *Task:* Grab the small loop and pull the drawer open. Drawer configuration within
  - 532 training data is face-first with slight orientation changes
  - 533 • *Action labeling:* Approach the loop, grab the loop, pull the drawer out
  - 534 • *Success metric:* When the drawer is open by 50% or more
- 535 2. Open the bottom drawer
- 536 • *Task:* Grab the cylindrical handle and pull the drawer open. Drawer configuration
  - 537 within training data is face-first with slight orientation changes. Note significantly
  - 538 different grasp is required than for top drawer
  - 539 • *Action labeling:* Approach the handle, grab it, pull the drawer out
  - 540 • *Success metric:* When the drawer is open by 50% or more
- 541 3. Close the drawer
- 542 • *Task:* This task is unqualified, i.e. the instructor does not say whether to close the top
  - 543 or bottom drawer instead the agent must determine which drawer needs closing from
  - 544 its state and close it. Align the gripper with the front of whichever drawer is open
  - 545 and push it closed. The training set always has only one of the drawers open, in a
  - 546 front-facing configuration with small orientation changes
  - 547 • *Action labeling:* Approach drawer from the front, make contact, push until closed
  - 548 • *Success metric:* When the drawer is closed to within 10% of its limit or when arm is
  - 549 maximally stretched out to its limit (when the drawer is kept far back)
- 550 4. Place inside the drawer
- 551 • *Task:* Approach an empty spot inside the drawer and place whatever is in hand inside
  - 552 it
  - 553 • *Action labeling:* Top-down approach pose on top of the drawer, move to make contact
  - 554 with the surface and release the object, move up for retreat
  - 555 • *Success metric:* Object should be inside the drawer
- 556 5. Pick lemon from the basket
- 557 • *Task:* Reach into the basket where lemon is placed and pick up the lemon
  - 558 • *Action labeling:*
  - 559 • *Success metric:* Lemon should be in robot's gripper
  - 560 • *Considerations:* Since the roll-out is open-loop and a lemon is spherical in nature, a
  - 561 trial was assigned success if the lemon rolled out of hand upon contact after the 2nd
  - 562 action. This was done consistently for both PerAct and SLAP.

- 563 6. Place in the bowl
- 564 • *Task*: Place whatever is in robot’s hand into the bowl receptacle
  - 565 • *Action labeling*: Approach action on top of the bowl, interaction action inside the
  - 566 bowl with gripper open, retreat action on top of the bowl
  - 567 • *Success metric*: The object in hand should be inside the bowl now
- 568 7. Place in the basket
- 569 • *Task*: Place the object in robot’s hand into the basket
  - 570 • *Action labeling*: Approach action on top of the free space in basket, interaction action
  - 571 inside the basket with gripper open, retreat action on top of the basket
  - 572 • *Success metric*: The object is inside the basket
- 573 8. Pick up the bottle
- 574 • *Task*: Pick up the bottle from the table
  - 575 • *Action labeling*: Approach pose in front of the robot with open gripper, grasp pose
  - 576 with gripper enclosing the bottle and gripper closed, retreat action at some height
  - 577 from previous action with grippers closed
  - 578 • *Success metric*: The bottle should be in robot’s gripper off the table

579 Notably, success for opening drawers is if the drawer is 50% open after execution; this is because  
 580 sometimes the drawer is too close to the robot’s base for it to open fully with a fixed-base Franka  
 581 arm.

## 582 D.2 Language Annotations

583 Below we include the list of language annotations used in our experiments. Table 4 shows the  
 584 language that was used to train the model; we’re able to show some robustness to different language  
 585 expressions. We performed a set of experiments on held-out, out-of-distribution language despite  
 586 this not being the focus of our work; this test language is shown in Table 5.

## 587 D.3 Out of distribution Results from SLAP

588 We show more results for attention point predicted by  $\pi_I$  in Fig. 11. For the placing task, the agent  
 589 has never seen a heavily cluttered drawer inside before, but it is able to find flat space which indicates  
 590 placing affordance. For the bottle picking task, this sample has a lemon right next to the bottle which  
 591 changes the shape of the point-cloud around the bottle. We see that  $\pi_I$  is able to find an interaction  
 592 point albeit with placement different from expert and lower down on the bottle. Similarly the open  
 593 top drawer sample has more heavy clutter on and around the drawer to test robustness.

594 Fig. 12 shows the prediction and generated trajectory for picking up a previously unseen bottle Note  
 595 that while the models are able to detect the out of distribution bottle, the trajectory actually fails due  
 596 to bottle being much wider and requiring more accuracy in grasping.

## 597 D.4 Motion Planning Failures

598 Our evaluation system has a simple motion planner which is not collision aware as a result we saw  
 599 a number of task failures for both the models. However, we note that the frequency of task failures  
 600 due to motion planning problems was higher for PerAct. We think it is because PerAct predicts  
 601 each action of the same task as an entirely separate prediction trial, while SLAP forces continuity on  
 602 the relative motions for the same task by centering them around the interaction point (see Fig. 12).  
 603 That said, we also note with an advanced motion planner PerAct will not run into such issues as  
 604 seen during our evaluations. Authors note in their own paper their heavy reliance on good motion  
 605 planning solutions [2].

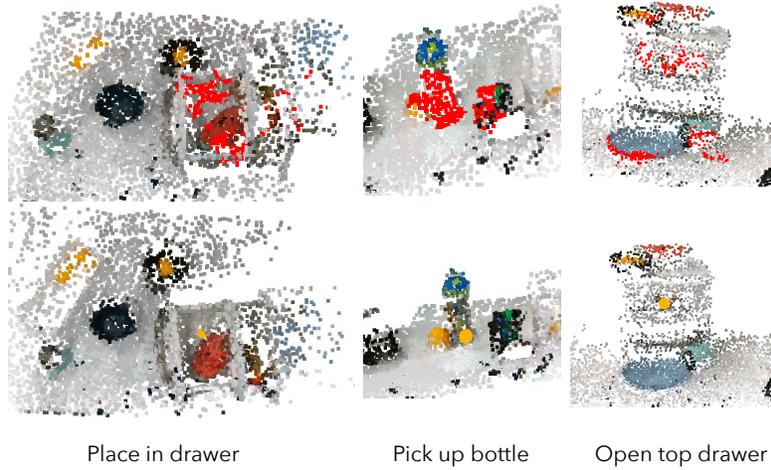


Figure 11: Examples of out of distribution predictions made by  $\pi_I$ . We show that it is able to handle heavy clutter around the implicated object to predict interaction points. Note that the prediction for bottle picking is sub-optimal in this example.

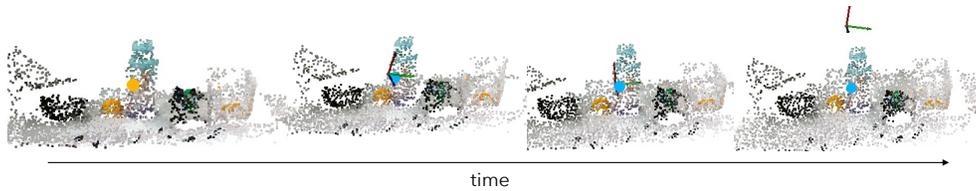


Figure 12: A generalization example of success for our model. The new bottle has same shape as the within distribution bottle but is much taller, different in color and wider in girth. The model is able to predict the interaction site and a feasible trajectory around it. We note though the execution of this trajectory was a failure; due to wider girth of the bottle the predicted grasp was not accurate enough to enclose the object.

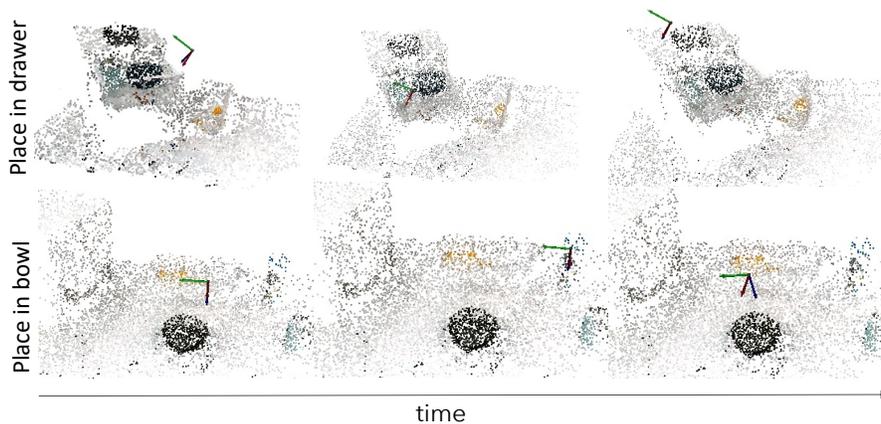


Figure 13: Examples of failure cases for our baseline, PerAct, for the “place in drawer” and “place in bowl” tasks. In the top example, the gripper is moved from drawer’s side towards inside, instead of from the top as demonstrated by expert. The gripper ends up pushing off the drawer to the side as our motion-planner is not collision-aware. Note that SLAP does not exhibit such behaviors as  $\pi_R$  implicitly learns the collision constraints present in demonstrated data. In the bottom example, each action prediction is disjointed from previous and semantically wrong.

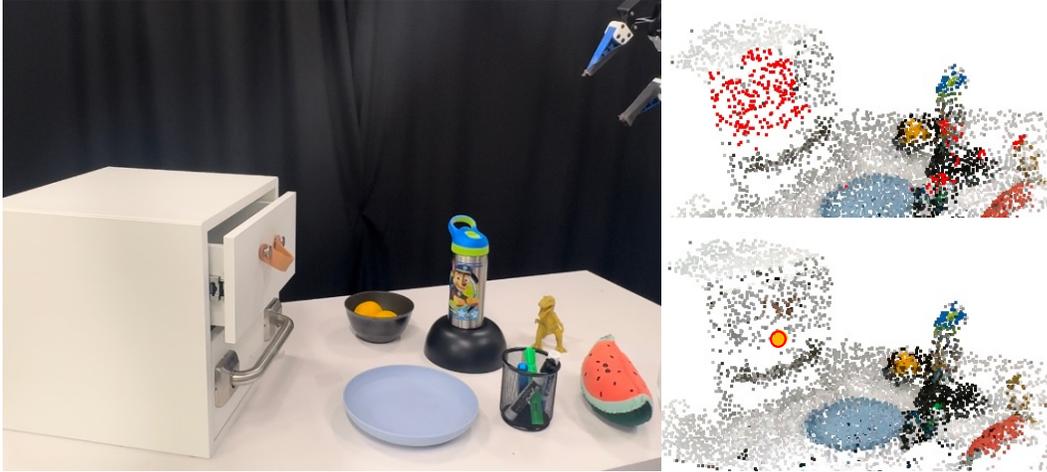


Figure 14: An out-of-distribution SLAP failure example where an extreme sideways configuration of the drawer is paired with unseen distractors for the “open top drawer” skill. We can see the attention mask ranking other distractors in its top 5% and failing to choose an optimal interaction point.

## 606 E Additional Analysis

### 607 E.1 Visualizing the Learned Attention

608 Since we use scores to choose the final interaction point, our classifier model is naturally inter-  
 609 pretable, being able to highlight points of interest in a scene. We visualize this attention by selecting  
 610 the points with the highest 5% of interaction score given a language command  $l$ .

### 611 E.2 Language Generalization

612 By using pretrained CLIP language embeddings to learn our spatial attention module  $\pi_I$ , our model  
 613 can generalize to unseen language to some extent. We tested this by running an experiment where we  
 614 evaluate performance on in-distribution scene settings, prompted by a held-out list of language ex-  
 615 pressions. We choose three representative tasks for this experiment and run 10 tests with 2 different  
 616 language phrasings.

## 617 F Additional Related Work

618 We note some other related work that’s relevant to SLAP, but not as directly relevant.

619 **Vision-Language Navigation.** Similar representations are often used to predict subgoals for explo-  
 620 ration in vision-language navigation [30, 31, 8, 50, 51]. HLSM builds a voxel map [30], whereas  
 621 FiLM builds a 2D representation and learns to predict where to go next [31]. VLMaps proposes an  
 622 object-centric solution, creating a set of candidate objects to move to [8], while CLIP-Fields learns  
 623 an implicit representation which can be used to make predictions about point attentions in responds  
 624 to language queries [50], but does not look at manipulation. Similarly, USA-Net [51] generates a 3d  
 625 representation with a lot of semantic features.

626

Task Name	Training Annotations
pick up the bottle	pick up a bottle from the table pick up a bottle
pick up a lemon	grab my water bottle pick the lemon from inside the white basket grab a lemon from the basket on the table hand me a lemon from that white basket
place lemon in bowl	place the lemon from your gripper into the bowl add the lemon to a bowl on the table put the lemon in the bowl
place in the basket	place the object in your hand into the basket put the object into the white basket place the thing into the basket on the table
open bottom drawer	open the bottom drawer of the shelf on the table pull the second drawer out open the lowest drawer
close the drawer	close the drawers push in the drawer close the drawer with your gripper
open top drawer	open the top drawer of the shelf on the table pull the first drawer out open the highest drawer
place in the drawer	put it into the drawer place the object into the open drawer add the object to the drawer

Table 4: Examples of language used to train the model.

Task Name	Held-Out Test Annotations
Pick up the bottle	Grab the bottle from the table Pick up the water bottle
Open the top drawer	Pull top drawer out Open the first drawer
Place into the drawer	Add to the drawer Put inside the drawer

Table 5: Examples of out-of-distribution language annotations used for evaluation