

# Rapid Motor Adaptation for Robotic Manipulator Arms

## Supplementary Material

### A. Environment Randomization Details

The sample range for the environment parameters during training and evaluation is shown in Tab. 2.

Rand. type	Params.	Training	Testing
Env. Var.	Scale	[0.70, 1.20]	[0.56, 1.44]
	Density	[0.50, 5.00]	[0.40, 6.00]
	Friction	[0.50, 1.10]	[0.40, 1.32]
	Shape	All	All*
Ex. Dist.	Force	[0.00, 2.00]	[0.00, 2.40]
Obs. Nois.	Obj. Pos.	[-0.005, 0.005]	[-0.006, 0.006]
	Obj. Rot.	[-10, 10]	[-12, 12]
	Joint Pos.	[-0.005, 0.005]	[-0.006, 0.006]

Table 2. Summary of environment randomization for training and testing phases. The table specifies the range of uniform distribution from which the parameters are sampled. Environment Variations (Env. Var.) encompass object properties such as object scale multiplier (which is multiplied on top of the default object scale), object density multiplier, object coefficient of friction, and object shape. \*The pick and place agent is trained on all the YCB objects and evaluated on both YCB and EGAD objects. The faucet turning agent is trained and tested on the same set of faucets from the PartNet-Mobility dataset. External Disturbance (Ex. Dist.) simulates environmental perturbations through force applied onto the object. Observation Noise (Obs. Nois.) accounts for inaccuracies in the object position, rotation estimation and joint position reading.

### B. Example Objects from the Tasks

Fig. 4 shows example objects from the YCB object dataset<sup>2</sup>[7] and Faucets from the PartNet-Mobility dataset used in the Faucet Turning task[37]. For EGAD objects, shape complexity is measured using the entropy of a histogram based on the angular deficit at each vertex of a 3D mesh. Grasp difficulty is estimated using the 75th percentile of the Ferrari-Canny quality metric calculated from sampled antipodal grasps on each object

### C. Rewards


We use the default dense reward functions from Maniskill2 [12]. We give an overview of the reward function used in Pick and Place. And we refer to the reward functions for Faucet Turning and Peg Insertion to ManiSkill2’s repository<sup>3</sup>.

<sup>2</sup>Figure reproduced from [www.ycbbenchmarks.com](http://www.ycbbenchmarks.com) with permission.

<sup>3</sup>[github.com/haosulab/ManiSkill2](https://github.com/haosulab/ManiSkill2)



(a) Objects in the YCB dataset



(b) Sample Faucets

Figure 4. Figure (a) illustrates the YCB object we use in the Pick and Place task. Figure (b), (c) shows sample faucets for the Faucet Turning task.

For Pick and Place, the objective is to move the objects to be within 2.5 cm of the goal position, and the robot is static. Once the goal is achieved, the agent receives a reward of 10. Otherwise, the reward at time step is the sum of the following three components:

$$r_t = \begin{cases} 10 & \text{if succ}_t, \\ 1 - \tanh(3\Delta_t) + 3 \cdot \mathbb{I}_{\text{is\_grasped}, t} & \\ + 3 \cdot (1 - \tanh(3\|\mathbf{x}_t^{\text{obj}} - \mathbf{x}_t^{\text{goal}}\|)) & \text{o/w.} \end{cases} \quad (4)$$

where  $\text{succ}_t$  is a binary indicator of success at time step  $t$ , which equals 1 if the task is successfully completed, otherwise 0;  $\Delta_t = \max(\|\mathbf{x}_t^{\text{tcp}} - \mathbf{x}_t^{\text{obj}}\| - \|\mathbf{d}_{\text{bbox}}\|, 0)$  represents the distance between the tool center point (TCP) of the robot’s end effector and the object at time step  $t$ , after accounting for the bounding box dimensions of the object;  $\|\mathbf{x}_t^{\text{tcp}} - \mathbf{x}_t^{\text{obj}}\|$  is the Euclidean norm representing the distance between the TCP and the object at time step  $t$ ;  $\|\mathbf{d}_{\text{bbox}}\|$  is the Euclidean norm representing the dimensions of the object’s bounding box;  $\mathbb{I}_{\text{is\_grasped}, t}$  is an indicator function that equals 1 if the object is grasped at time step  $t$ , otherwise 0;  $\|\mathbf{x}_t^{\text{obj}} - \mathbf{x}_t^{\text{goal}}\|$  is the Euclidean norm measuring the distance between the object

and the goal location at time step  $t$ .

## D. Sample Evaluation Videos

We showcase an example trajectory of our method, RMA<sup>2</sup>, vs. the domain randomization baseline from the evaluation of each of the four tasks and a lemon-picked representative failure episode from RMA<sup>2</sup> at <https://youtu.be/Vii6kx3-sTQ>. We found the most common failure mode is a failure to insert the peg precisely into the hole under disturbances, or a failed initial grasp (in pick and place and faucet turning), likely due to uncertainty about the object geometry under RL training.