

MANISKILL-HAB: A BENCHMARK FOR LOW-LEVEL MANIPULATION IN HOME REARRANGEMENT TASKS

Anonymous authors

Paper under double-blind review

ABSTRACT

High-quality benchmarks are the foundation for embodied AI research, enabling significant advancements in long-horizon navigation, manipulation and rearrangement tasks. However, as frontier tasks in robotics get more advanced, they require faster simulation speed, more intricate test environments, and larger demonstration datasets. To this end, we present MS-HAB, a holistic benchmark for low-level manipulation and in-home object rearrangement. First, we provide a GPU-accelerated implementation of the Home Assistant Benchmark (HAB). We support realistic low-level control and achieve over 3x the speed of previous magical grasp implementations at similar GPU memory usage. Second, we train extensive reinforcement learning (RL) and imitation learning (IL) baselines for future work to compare against. Finally, we develop a rule-based trajectory filtering system to sample specific demonstrations from our RL policies which match predefined criteria for robot behavior and safety. Combining demonstration filtering with our fast environments enables efficient, controlled data generation at scale.

1 INTRODUCTION

An important goal of embodied AI is to create robots that can solve manipulation tasks in home-scale environments. Recently, faster and more realistic simulation, home-scale rearrangement benchmarks, and large robot datasets have provided important platforms to accelerate research towards this goal. However, there remains a need for all of these features in one unified benchmark.

We present **MS-HAB**¹², a holistic, open-sourced, home-scale manipulation benchmark with four key features: (1) fast simulation with realistic physics and manipulation, including low-level control, for efficient training, evaluation, and dataset generation, (2) home-scale manipulation tasks through the Home Assisitant Benchmark (HAB) (Szot et al., 2021), (3) extensive baselines for future work to compare against, and (4) scalable, controlled data generation using an automated, rule-based trajectory filtering system.

Fast Manipulation Simulation with Realistic Physics and Rendering: Using ManiSkill3 (Tao et al., 2024), we implement a GPU-accelerated version of the HAB (Szot et al., 2021), an apartment-scale rearrangement benchmark containing three long-horizon tasks using the Fetch mobile manipulator (ZebraTechnologies, 2024). While the original HAB uses magical grasp (teleport closest object within 15cm to the gripper), we require realistic grasping.

The MS-HAB environments support low-level control for realistic grasping, manipulation, and interaction, while the original Habitat 2.0 implementation does not support such kind of low-level control. Furthermore, by scaling parallel environments, MS-HAB environments achieve over 4000 samples per second (SPS) while the robot actively collides with multiple dynamic objects and the environment renders 2 128x128 RGB-D images — 3x faster than Habitat 2.0 at similar GPU memory usage. This significant speedup allows us to scale up training, evaluation, and data generation.

Reinforcement Learning (RL) Baselines: Online RL provides a promising framework to learn from online interaction without needing preexisting demonstration data. As in Gu et al. (2023a), we train individual mobile manipulation skills and chain them to solve long-horizon tasks. We

¹Code: <https://github.com/anonsubmit0/maniskill-hab>

²Website: <https://sites.google.com/view/maniskill-hab>

hand-engineer dense rewards for the Fetch embodiment, designed for low-level control with mobile manipulation. Furthermore, we train manipulation policies overfit to one specific object’s geometry, outperforming all-object policies when grasping many objects or in conditions with tight tolerances depending on object geometry. Leveraging our fast environments, we run extensive RL baselines, training 150 policies across 3 seeds (50 policies/seed) with 1.83 billion environment samples.

Automated Event Labeling and Trajectory Categorization: We use privileged information from the simulator to distill trajectories into chronologically ordered lists of events (e.g. Pick events include ‘Contact (object)’, ‘Grasped’, ‘Dropped’, ‘Success’, and ‘Excessive (robot) Collisions’). Using these events lists, we define mutually exclusive, collectively exhaustive success and failure modes. For example, Pick success mode “reach success \wedge cumulative robot collisions $< 5000 \text{ N}$ \wedge object not dropped” requires events list (Contact, Grasp, Success) and forbids events ‘Dropped’ and ‘Excessive Collisions’. We filter our dataset by selecting demonstrations labeled with success modes that guarantee particular behaviors (e.g. pick without dropping) and safety constraints (e.g. cumulative robot collisions $< 5000 \text{ N}$). Furthermore, we provide trajectory categorization statistics for all baselines in Appendix A.5 so future work can gear its methodology to solve frequent failure modes discovered by our policies.

Dataset Generation and Imitation Learning (IL) Baselines: When generating our dataset, we use trajectory categorization to filter demonstrations without needing manual labor, and we provide Imitation Learning (IL) baselines using our dataset. Our results show that selecting demonstrations with particular behavior biases IL policies towards that behavior. Paired with our fast simulator, users can generate massive datasets and control demonstration type in fast wall-clock time.

Summary of Contributions: The contributions of MS-HAB are summarized as follows: 1) GPU-accelerated HAB implementation which supports realistic low-level control and achieves over 4000 SPS while interacting and rendering, 2) extensive RL and IL baselines, 3) automated event labeling and trajectory categorization, providing success and failure mode statistics for all baseline policies, and 4) efficient, controlled vision-based robot dataset generation at scale.

2 RELATED WORK

Simulators and Scene-Level Embodied AI Platforms: Earlier scene-level simulators focus on navigation and simple interaction with realistic visuals (Savva et al., 2019). Other simulators add kinematic object state transitions (Kolve et al., 2017; Li et al., 2021), significant scene randomization (Nasiriany et al., 2024; Deitke et al., 2022), soft-body physics and audio (Gan et al., 2022), flexible and deformable materials, object composition rules, and so on (Li et al., 2022). However such complicated features often slow down simulation speed.

Habitat 2.0 forgoes additional features, supporting rigid-body dynamics, articulations, and magical grasping, to achieve best-in-class single-process scene-level simulation speed (Szot et al., 2021). However, it is constrained by the limited parallelization of CPU simulation.

Other simulators focus on low-level, contact-rich control in simpler settings (James et al., 2020; Zhu et al., 2020; Xiang et al., 2020). ManiSkill3 in particular achieves state-of-the-art GPU simulation speed (Tao et al., 2024), however its suite of tasks are simpler than the Home Assistant Benchmark (HAB) (Szot et al., 2021), which we implement for MS-HAB.

Scalable Demonstration Datasets: Real-world robot datasets are promising for direct deployment to the real world (Brohan et al., 2023). However, these initiatives are limited in scaling and use cases due to small-scale toy setups (Ebert et al., 2022), vision-only data (Dasari et al., 2019), or requiring massive coordinated (et al., 2024; Khazatsky et al., 2024) or distributed (Mandlekar et al., 2018) human effort over many months or even years. Furthermore, real robot datasets cannot efficiently generate new data, and do not support online sampling.

Generative interactive world models allow some interactivity on similarly realistic data by generating new frames based on provided actions (Yang et al., 2024). However, these models suffer from artifacts and long-term memory issues which rule out home-scale rearrangement, and low frame rates make training high-frequency low-level control policies intractable. Furthermore, neither real-robot datasets nor generative world models currently support querying privileged data from a simulator, which is necessary for MS-HAB’s automated event labeling and trajectory categorization.

Meanwhile, classical physical simulation supports data generation from a variety of sources (expert teleoperated, suboptimal human, etc), and machine-generated data is largely scalable; however, datasets like Fu et al. (2020); Mandlekar et al. (2021); Gu et al. (2023b) only support smaller-scale continuous control tasks.

RoboCasa combines different aspects of above approaches (Nasiriany et al., 2024): a physical simulator, diverse AI-generated textures and models, 1250 human-teleoperated demonstrations, and MimicGen to scale data (Mandlekar et al., 2023). However, RoboCasa achieves only 31.9 SPS *without rendering*, does not support filtering trajectories by behavior, and its demonstrations alternate between manipulation and navigation. Meanwhile, we achieve 125x faster simulation *while rendering* 2 128x128 RGB-D images *and interacting* with multiple dynamic objects. We also support automated filtering under customizable constraints, and our demonstrations use whole-body control.

Skill Chaining: Chen et al. (2023) and Lee et al. (2021) use finetuning methods to bias the initial and terminal state distributions to increase handoff success while skill chaining. However, these methods are applied to tasks with unchanging order (e.g. furniture assembly, block orient/grasp/insert). Meanwhile, Gu et al. (2023a) formulate composable and reusable skills with mobility to create greater overlap in initial and terminal state distributions, achieving better results than stationary manipulation. However, Gu et al. (2023a) uses magical grasp, while we include additional considerations for low-level grasping, such as sampling grasp poses from Pick policies, and overfitting object manipulation policies to specific object geometries.

3 PRELIMINARIES

3.1 TASKS, SUBTASKS, AND POLICIES

The Home Assistant Benchmark (HAB) (Szot et al., 2021) includes three long-horizon tasks which involve rearranging objects from the YCB dataset (Çalli et al., 2015):

- **TidyHouse:** Move 5 target objects to different open receptacles (e.g. table, counter, etc).
- **PrepareGroceries:** Move 2 objects from the opened fridge to goal positions on the counter, then 1 object from the counter to the fridge.
- **SetTable:** Move 1 bowl from the closed drawer to the dining table and 1 apple from the closed fridge to the same dining table.

To solve these tasks, Szot et al. (2021) define parameterized skills: Pick, Place, Open Fridge/Drawer, Close Fridge/Drawer, and Navigate. For each skill, we define corresponding subtasks. Successful low-level grasping is heavily dependent on an object’s pose. So, depending on the subtask, the simulator provides ground-truth pose $x_{pose} = [x_{rot}|x_{pos}]$ for target object x , ground-truth handle position a_{pos} for target articulation a , or 3D goal position g_{pos} , updated each timestep during manipulation. Each subtask also fails if the robot cumulative force reaches beyond a set threshold. For more details, see Appendix A.1. We provide brief descriptions of the subtasks below:

- **Pick**[a , optional](x_{pose}): pick object x (from articulation a , if provided).
- **Place**[a , optional](x_{pose} , g_{pos}): place object x in goal g (in articulation a , if provided)
- **Open**[a](a_{pos}): open articulation a with handle at a_{pos}
- **Close**[a](a_{pos}): close articulation a with handle at a_{pos}
- **Nav**(* pos): navigate to *

From a reinforcement learning perspective, we formulate each long-horizon task as a standard Markov Decision Process (MDP) which can be described as a tuple $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{T}, \rho, \gamma)$ with continuous state space \mathcal{S} , action space \mathcal{A} , scalar reward function $\mathcal{R} : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$, environment dynamics function $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$, initial state distribution ρ , and discount factor $\gamma \in [0, 1]$. Then, as in Gu et al. (2023a), define a subtask ω as a smaller MDP $(\mathcal{S}, \mathcal{A}_\omega, \mathcal{R}_\omega, \mathcal{T}, \rho_\omega, \gamma)$ derived from M . For each task M with subtask ω , we train low-level control policy $\pi_\omega : \mathcal{S} \rightarrow \mathcal{A}_\omega$ with RL or IL.

In this work, we study a partially observable variant of each task, where the policy must use 2 128x128 depth images to infer collisions and obstructions. We train different policies for each

task/subtask combination (i.e., TidyHouse Pick, PrepareGroceries Pick, etc). Additionally, we train Pick and Place RL policies to overfit to specific objects, i.e., one policy for each task/subtask/object combination. Since this work focuses on low-level control, we use a teleport for the Navigation subtask. Additional details on training policies and teleport navigation are provided in Sec. 5.1.

3.2 SKILL CHAINING

Similar to Szot et al. (2021), we split each task into a sequence of subtasks using a perfect task planner. The sequences are defined below:

- **TidyHouse:** For $(x_i, g_i) \in \{(x_0, g_0), \dots, (x_4, g_4)\}$, complete:
 $\text{Nav}(x_i, pos) \rightarrow \text{Pick}(x_i, pose) \rightarrow \text{Nav}(g_i, pos) \rightarrow \text{Place}(x_i, pose, g_i, pos)$
- **PrepareGroceries:** For $(x_i, g_i) \in \{(x_0, g_0), (x_1, g_1), (x_2, g_2)\}$, complete:
 $\text{Nav}(x_i, pos) \rightarrow \text{Pick}_{\text{Fr}[i \leq 1]}(x_i, pose) \rightarrow \text{Nav}(g_i, pos) \rightarrow \text{Place}_{\text{Fr}[i=2]}(x_i, pose, g_i, pos)$
- **SetTable:** For $(x_i, g_i, a_i) \in \{(x_0, g_0, \text{Dr}), (x_0, g_0, \text{Fr})\}$, complete:
 $\text{Nav}(a_i, pos) \rightarrow \text{Open}_{a_i}(a_i, pos) \rightarrow \text{Nav}(x_i, pos) \rightarrow \text{Pick}(x_i, pose) \rightarrow \text{Nav}(g_i, pos) \rightarrow \text{Place}(x_i, pose, g_i, pos) \rightarrow \text{Nav}(a_i, pos) \rightarrow \text{Close}_{a_i}(a_i, pos)$

3.3 TRAIN AND VALIDATION SPLITS

The ReplicaCAD dataset (Szot et al., 2021) serves as the source for our apartment scenes. It comprises 105 scenes divided into 5 macro-variations, each containing 21 micro-variations. Macro-variations alter the layout of large furniture items such as refrigerators and kitchen counters, while micro-variations modify the placement of smaller furnishings like chairs and TV stands. The dataset is split into three parts: 3 macro-variations for training, 1 for validation, and 1 for testing. However, as the test split is not publicly accessible, our study utilizes only the train and validation splits.

Furthermore, for each long-horizon task, HAB provides 10,000 training episode configurations and 1,000 validation configurations. These configurations specify initial poses for YCB objects and define target objects, articulations, and goals. Importantly, these configurations exclusively utilize ReplicaCAD scenes from their respective splits.

4 ENVIRONMENT DESIGN AND BENCHMARKS

By scaling parallel environments with GPU simulation, MS-HAB achieves 4000 SPS on a benchmark involving representative interaction with dynamic objects — 3x Habitat 2.0’s implementation. Our environments support realistic low-level control for successful grasping, manipulation, and interaction, while the Habitat 2.0 environments do not support such kind of low-level control. This section outlines environment design choices which leverage GPU acceleration and benchmarks MS-HAB against Habitat’s implementation.

4.1 ENVIRONMENT DESIGN

Evaluation and Training Environments: First, we provide the base evaluation environment, SequentialTask, which supports executing different subtasks simultaneously on GPU. We perform physics simulation and rendering for all environments in parallel, then slice data by subtask to compute success and fail conditions. It does not support dense reward or spawn selection/rejection.

Second, we provide training environments for each subtask, $\{\text{SubtaskName}\}\text{SubtaskTrain}$, which extend the main evaluation environment. Each training environment provides dense rewards hand-engineered for the Fetch embodiment, supports spawning with randomization and rejection, and incorporates any additional features needed for training specific subtask skills.

Observation Space: We include target object pose, goal position, and TCP pose relative to the base, an indicator of whether the target object is grasped, 128x128 head and arm RGB-D images, and robot proprioception. For our experiments, we use only depth images. As is standard for the ManiSkill suite of tasks, the simulator computes ground-truth poses. We keep a consistent observation space across all subtasks via masking to support different subtasks running in parallel.

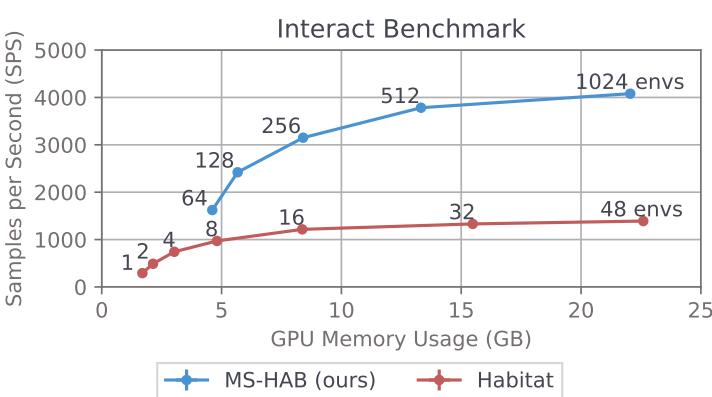


Figure 1: Interact benchmark comparing MS-HAB (ours) with Habitat. Each data point is annotated with the number of parallel environments used. SPS and GPU memory usage for each data point are averaged over 10 seeds; error bars representing 95% CIs are plotted, but are too small to see. Thanks to GPU acceleration, MS-HAB scales parallel environments to achieve approximately 3x the performance of Habitat while using similar GPU memory.

Action Space: We fully actuate the Fetch embodiment’s arm, torso, and head pan/tilt joints. We support joint-based controllers and end-effector-based controllers. For our experiments, we use a PD joint delta position controller for the arm, torso, and head joints. The agent provides linear and angular velocity to control the base. The action space is normalized to $[-1, 1]$.

Additional Details: Our environments load the ReplicaCAD dataset provided by Habitat 2.0. However, since Habitat 2.0 uses magical grasp, the original ReplicaCAD dataset’s collision meshes do not include handles for the kitchen drawers and fridge. So, we alter these collision meshes to include handles based on the provided visual meshes. We additionally provide navigable position meshes for the Fetch embodiment with Trimesh, as ManiSkill3 does not currently support navmeshes.

4.2 BENCHMARKING

We adapt Habitat 2.0’s Interact benchmark, which originally had the Fetch robot execute a precomputed trajectory to collide with two dynamic objects (Szot et al., 2021). While we retain the same precomputed trajectory, assets, and scene configuration, we modify the robot’s initial pose and disable magical grasp, allowing it to interact with five objects instead. Our setup includes two mounted 128x128 RGB-D cameras, with a simulation frequency of 100Hz and a control frequency of 20Hz (the standard for low-level control in ManiSkill3). We collect observation data from vectorized environments at each `step()` call. Our benchmarking is conducted on a machine equipped with a 16-core/32-thread Intel i9-12900KS processor and an Nvidia RTX 4090 GPU with 24 GB VRAM.

It is important to note that running the *exact* same episode in different simulators is exceedingly difficult since different simulation backends will result in interactions and collisions behaving slightly differently. Still, the full rollout is similar in both simulators, and the measured performance increase of MS-HAB in an interactive setting is significant.

Habitat’s Additional Optimizations: While the Habitat simulator already has best-in-class single process simulation speed, it provides optional additional optimizations: concurrent rendering and auto sleep. However, their experiments suggest that concurrent rendering can negatively impact train performance (Szot et al., 2021), so we enable auto-sleep and disable concurrent rendering.

Benchmark Analysis: Per Fig. 1, while Habitat achieves significantly stronger performance per parallel environment, its peak performance is limited to 1397.65 ± 11.02 SPS due to the limited parallelization of CPU simulation. Meanwhile, by scaling up to 1024 environments, MS-HAB is able to achieve 4109.40 ± 26.36 SPS — 2.94x the simulation speed of Habitat. Furthermore, our environments support realistic low-level control for successful grasping, manipulation, and interaction, while the Habitat 2.0 environments do not support such kind of low-level control.

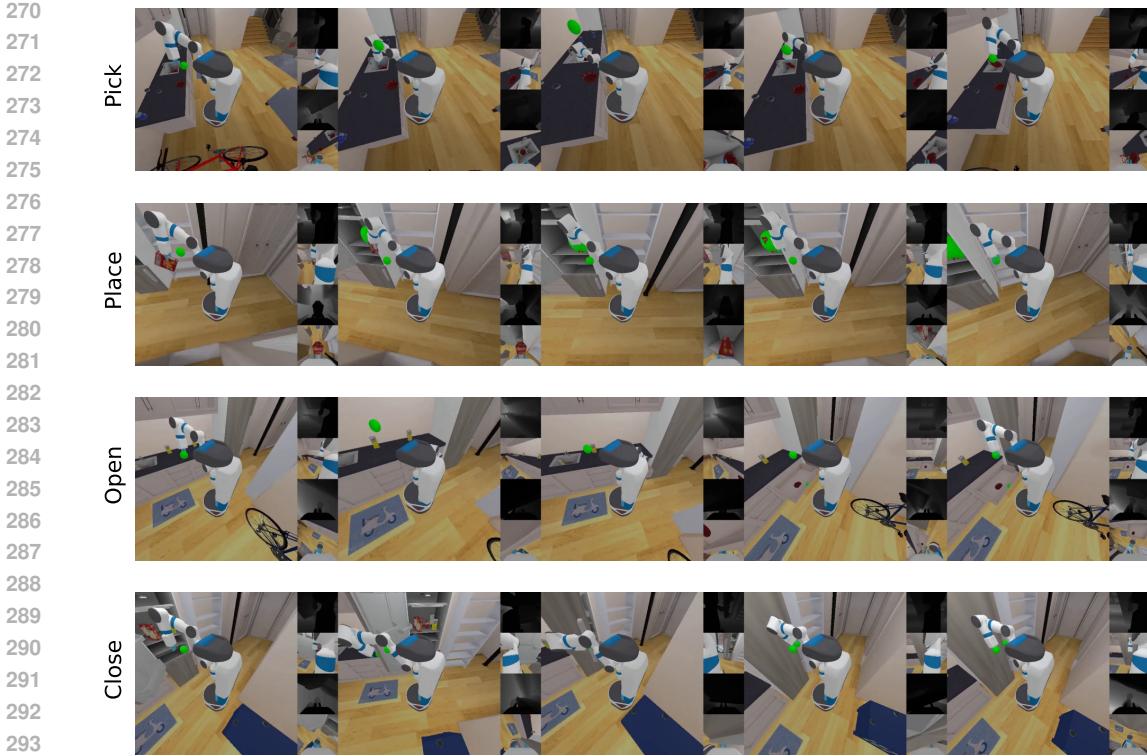


Figure 2: Renders of low-level, whole-body control policies solving Pick, Place, Open, and Close subtasks. We render 1 512x512 image and 4 128x128 sensor images. Note the base’s moving position relative to surroundings. Goal spheres are invisible to sensors. Full videos in supplementary.

5 METHODOLOGY

5.1 TRAINING REINFORCEMENT LEARNING POLICIES

We choose Reinforcement Learning (RL) to learn our subtask policies as RL does not require prior demonstration data, and it can take advantage of our highly parallelized environments to solve tasks in fast wall-clock time. We use a similar subtask formulation as M3, which trains mobile manipulation skills to solve each subtask from a region of spawn points.

Pick: Without magical grasp, our Pick policies must learn grasp poses which are valid, stable, and reachable within the kinematic constraints of the mobile Fetch robot. Furthermore, the policy must learn action sequences which can reach these grasp poses and retrieve the target object within the specified horizon while keeping the robot under the cumulative collision force limit.

As a result, learning successful grasping for multiple objects with different geometries — in addition to whole body control with collision constraints — is difficult. So, we opt to train individual Pick policies for each object, thereby overfitting to the geometry of that object. Our experiments show these per-object Pick policies achieve improved subtask success rates compared to all-object policies when handling many objects with varied geometries. In other words, we train a unique per-object Pick policy for every task/subtask/object combination.

Place: We train per-object Place policies as well. Our experiments show that, in settings where object geometry is more important (e.g. placing in a fridge with tighter tolerances), per-object Place policies reach higher success rates than all-object policies.

Additionally, without magical grasp, there is not a ground-truth means of spawning the robot while grasping an object. So, we train our Pick policies before Place, then we sample grasp poses from our Pick policies to initialize the robot in the Place subtask.

324 **Open and Close:** Following (Gu et al., 2023a), we train different Open and Close policies for the
 325 kitchen drawer and the fridge. Furthermore, we find that opening the kitchen drawer is particularly
 326 difficult due to its small handle. So, we perturb the initial state distribution of our Open kitchen
 327 drawer subtask during training to accelerate learning: 10% of the time, we initialize the kitchen
 328 drawer opened 20% of the way. During evaluation, we do not alter the initial states.

329 **Algorithms and Hyperparameters:** We stack 3 consecutive frames for image observations to han-
 330 dle partial observability.
 331

332 We train Pick and Place using SAC (Haarnoja et al., 2018; Xing, 2022) with a 1m replay buffer
 333 size. Visual observations are encoded by D4PG’s 4-layer CNN (Barth-Maron et al., 2018) and
 334 concatenated with state observations. Actor and critic networks are 3-layer MLPs and the critic has
 335 LayerNorm to avoid value divergence (Ball et al., 2023). We train Pick with 50M timesteps and
 336 Place with 25M timesteps.

337 We train Open and Close using PPO (Schulman et al., 2017; Huang et al., 2022). Visual observations
 338 are encoded by a NatureCNN (Mnih et al., 2015) and concatenated with state observations. The actor
 339 and critic networks are 2-layer MLPs. We train Open Fridge with 15M timesteps, Open Drawer with
 340 50M timesteps, Close Fridge with 25M timesteps, and Close Drawer with 15M timesteps.

341 We train 3 seeds for each task/subtask/object combination, evaluating on 189 episodes every 100,000
 342 train samples. We select the checkpoint with highest evaluation success once rate as our final policy.

343 **Metrics:** We run 1000 episodes for every evaluation run (task/subtask evaluation, ablations, etc).

344
 345 We evaluate subtask policies (Pick, Place, Open, Close) by success once rate (%), which is the
 346 percentage of trajectories that achieve success at least once in an episode with 200 timesteps. We
 347 evaluate long-horizon task success (TidyHouse, PrepareGroceries, SetTable) by Progressive Com-
 348 pletion Rate (%). Here, the success of each subtask requires the success of every previous subtask.
 349 Hence, the completion rate of the final subtask is the completion rate of the entire long-horizon task.

350 Furthermore, since we are primarily interested in low-level control and manipulation, we replace
 351 navigation with a simple teleport. The robot is teleported to the target location with the same arm,
 352 base pose, and spawn location randomizations as in subtask training, described in Appendix A.1.
 353 When completing a long-horizon task, we move to the next subtask as soon as the currently running
 354 subtask achieves success.

356 5.2 AUTOMATED TRAJECTORY CATEGORIZATION AND DATASET GENERATION

357
 358 Thanks to fast simulation environments, we can quickly generate 10s to 100s of thousands of demon-
 359 strations. However, our experiments show that our Imitation Learning (IL) policies are sensitive to
 360 demonstration behavior. To filter out “suboptimal” demonstrations, we use privileged information
 361 from our simulator to group demonstrations into mutually exclusive, collectively exhaustive success
 362 and failure modes without significant manual labor. Furthermore, we use this trajectory labeling
 363 system to identify types and causes of failure in our baseline policies in Sec. 6.2 and Appendix A.5.

364 **Example of Pick Subtask:** We provide a high-level overview of trajectory labeling on the Pick
 365 subtask. For detailed definitions of events and labels, see Appendix A.5. First, we define “events”
 366 which occur at any timestep t : 1) Contact: nonzero robot/target pairwise force, 2) Grasped: object
 367 not grasped at step $t-1$ and grasped at step t , 3) Dropped: object grasped at step $t-1$ and not grasped
 368 at step t , and 4) Excessive Collisions: robot cumulative force exceeds 5000 N. For Pick trajectory
 369 $\tau_{\text{pick}} = (s_0, a_0, \dots, s_n, a_n)$, we create chronologically ordered event list $E_{\text{pick}} = (e_1, \dots, e_k)$.

370 Next, we define success and failure modes. For example, one success mode is “straightforward
 371 success” with $E_{\text{pick}} = (\text{Contact}, \text{Grasp}, \text{Success})$, requiring success without dropping the object
 372 or colliding too much. One failure mode is “dropped failure,” defined as $(\text{Excessive Collisions} \notin E_{\text{pick}}) \wedge (Dropped \in E_{\text{pick}}) \wedge (i < j \text{ for maximal } i, j \text{ such that } e_i = \text{Grasped}, e_j = \text{Dropped})$.
 373 “Dropped failure” trajectories fail because the robot irrecoverably drops the target object.

374
 375 In generating our Pick subtask dataset, we apply filters to include only “straightforward success”
 376 trajectories. These trajectories are characterized by the absence of dropping and minimal collisions.
 377 As the dataset generation code is publicly available, users have the flexibility to create their own
 datasets with custom constraints tailored to their specific requirements.

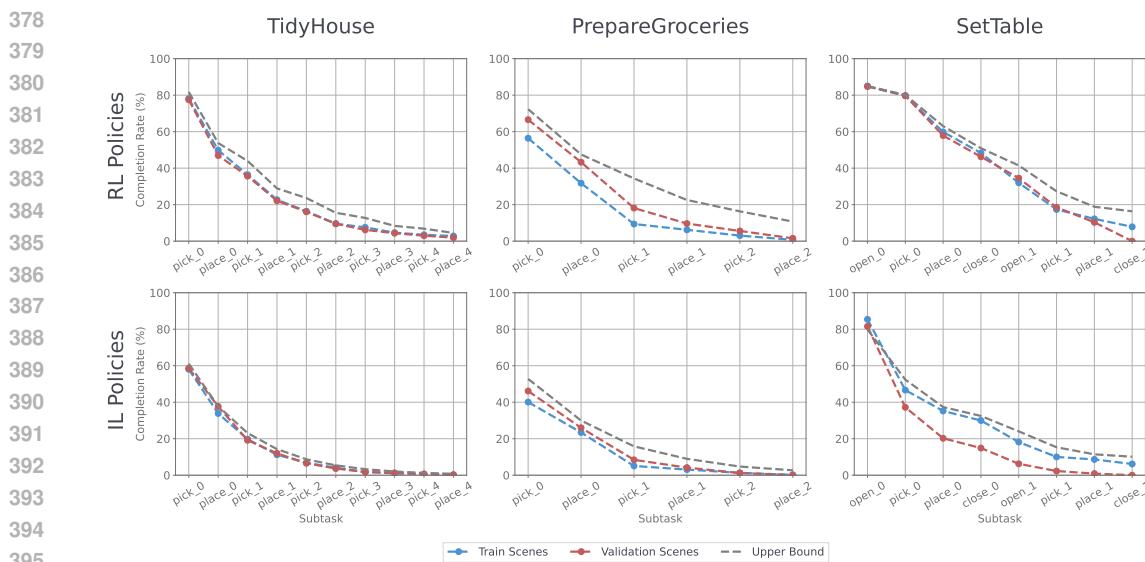


Figure 3: Long-horizon task progressive completion rates (%) on train and validation splits averaged over 1000 episodes. Furthermore, we provide an ‘upper bound’ on performance based on the success rates of each subtask policy. Best viewed zoomed.

Imitation Learning Baselines: We train IL baselines on our dataset using Behavior Cloning (BC) (Bain & Sammut, 1995; Ross et al., 2011; Daftry et al., 2016) (implementation based on Huang et al. (2022)). Visual observations are encoded by the 5-layer CNN from Bojarski et al. (2016), then concatenated with state observations. The actor network is a 3-layer MLP.

Dataset Size: With fast environments and automated trajectory filtering, users can generate as many 10s or 100s of thousands of demonstrations as needed in a controlled manner. For our baselines, we generate 1000 demonstrations per task/subtask/object combination using our per-object RL policies on the train split. The filters used are listed in Appendix A.5.1, with definitions in Appendix A.5.2.

6 RESULTS

6.1 BASELINES

Fig. 3 shows the RL and IL policies’ progressive completion rate. We provide an optimistic upper bound on progressive completion rate by (incorrectly) assuming that the completion of each subtask is independent of every other subtask, thus directly multiplying subtask success once rates. Table 1 shows success once rate for individual subtasks. We find 4 major avenues for improvement.

First, our optimistic upper bound shows low expected success rate on the long-horizon tasks. Even with per-object RL policies, our low-level mobile manipulation subtasks are difficult to train on dense reward, and improving subtask success rate is the most direct way to improve overall task completion rate. Second, TidyHouse and SetTable RL baselines have some gap between upper bound and real completion rate, indicating potential handoff issues or disturbance to prior target objects in success states. Meanwhile, the PrepareGroceries RL baseline has a large drop in completion rate during the second Pick_{Fr} subtask, indicating that the first Pick_{Fr} causes too much disturbance to objects in the fridge. So, improving policy performance in cluttered spaces is important. Third, our IL policies perform notably worse in Pick and Place, indicating a need for methods and architectures which can handle multimodalities in the data. Finally, while most RL policies generalize well to the validation split, the Close Fridge policy completely fails on validation scenes because the fridge door opens into a wall, preventing the arm from reaching the handle. This is not an issue with magical grasping (Gu et al., 2023a), indicating that low-level control may need more scene diversity.

432
 433 Table 1: Subtask success once rates for RL and IL baselines. The RL-Per vs All column shows
 434 the difference in per-object RL policy performance and its all-object counterpart. We do not train
 435 all-object policies for Open or Close subtasks.

TASK	SUBTASK	SPLIT	RL-PER	RL-ALL	IL	RL-PER vs ALL
TidyHouse	Pick	Train	81.75	71.63	61.11	+10.12
		Val	77.48	68.15	59.03	+9.33
Prepare Groceries	Place	Train	65.77	63.69	61.81	+2.08
		Val	65.97	66.07	63.79	-0.10
SetTable	Pick	Train	66.57	51.88	44.64	+14.69
		Val	72.32	62.10	52.78	+10.22
SetTable	Place	Train	60.22	53.37	50.00	+6.85
		Val	65.67	58.63	56.75	+7.04
SetTable	Pick	Train	80.85	75.69	60.71	+5.16
		Val	88.49	79.86	72.62	+4.63
SetTable	Place	Train	73.31	72.82	71.23	+0.49
		Val	67.06	68.25	62.20	-1.19
SetTable	Open _{Fr}	Train	83.43	-	74.01	-
		Val	88.10	-	53.67	-
SetTable	Open _{Dr}	Train	84.92	-	79.86	-
		Val	84.52	-	78.57	-
SetTable	Close _{Fr}	Train	86.81	-	86.90	-
		Val	0.00	-	0.00	-
SetTable	Close _{Dr}	Train	88.79	-	88.39	-
		Val	89.29	-	87.60	-

460 6.2 ABLATIONS

461 6.2.1 RL POLICIES: ALL-OBJECT VS PER-OBJECT POLICIES

463 The goal of training per-object RL policies for Pick and Place is to improve subtask success rate
 464 since policies with higher success rates allow us to generate successful demonstrations under more
 465 initialization conditions. To verify this, we run two ablations.

466 **Does training per-object Pick and Place policies improve subtask success rate compared to all-
 467 object policies?** Per Table 1, per-object policies perform notably better in TidyHouse and Prepare
 468 Groceries Pick, which involve 9 objects, with more modest improvement in SetTable Pick, which
 469 has only 2 objects. Per-object policies perform significantly better in PrepareGroceries Place, which
 470 involves placing with tight tolerances on a cluttered fridge shelf, while performance differences are
 471 negligible in TidyHouse and SetTable Place, which only involve open receptacles. So, per-object
 472 Pick and Place policies learn improved manipulation when grasping a greater variety of objects, or
 473 when manipulating objects in areas with tighter constraints.

474 **Are per-object policies necessary to learn grasping for certain objects in the Pick subtask?** In
 475 Table 2, we run our automated trajectory labeling system on Pick YCB object #003, the Cracker
 476 Box (Çalli et al., 2015). The Fetch robot’s parallel gripper can only grasp the Cracker Box along its
 477 shortest dimension, so the set of valid grasp poses are highly dependent on the object’s pose relative
 478 to the robot. The all-object policy is 1.88-2.42x more likely to fail to excessive collisions and 1.87-
 479 12.37x more likely to fail to grasp the object, indicating that overfitting to a specific geometry is
 480 important for our RL policies to learn grasping on difficult geometries. For more detailed trajectory
 481 labeling definitions and statistics, please see Appendix A.5.

483 6.2.2 IL POLICIES: LABELING AND FILTERING DATASET TRAJECTORIES

485 **IL Policies: Can we control the behavior of our IL policies by filtering for specific demonstra-
 486 tions?** Our PrepareGroceries Place RL policies have two similarly frequent success modes: place

486
 487 Table 2: Trajectory labeling on Pick Cracker Box with all and per-object RL policies. We group
 488 the trajectories into four categories: success once (**S-Once**), excessive collision failure (**F-Col**),
 489 cannot grasp failure (**F-Grasp**), and other failure modes (**F-Other**). We provide the percentage of
 490 trajectories which fall into each category, and each row sums to 100% (barring any rounding errors).

	TASK	SPLIT	TYPE	S-ONCE	F-COL	F-GRASP	F-OTHER
TidyHouse	Train	RL-All	29.46	34.52	28.17	7.85	
		RL-Per	71.63	17.26	2.48	8.63	
	Val	RL-All	33.73	33.13	24.50	8.64	
		RL-Per	73.41	16.67	1.98	7.94	
Prepare Groceries	Train	RL-All	11.51	60.62	16.17	11.70	
		RL-Per	51.98	25.10	8.63	14.29	
	Val	RL-All	14.19	57.24	26.88	1.69	
		RL-Per	56.15	30.46	9.72	3.67	

501
 502 in goal (release the object within 15cm of g_{pos}) and drop to goal (release beyond 15cm). Although
 503 MS-HAB does not simulate state transitions like breaking, placing objects without dropping is a
 504 desirable, safe robot behavior to avoid excessive damage.

505 We generate 3 datasets with 500 demonstrations per object: 1) place in goal only, 2) drop in goal
 506 only, and 3) 50/50 split (“place”, “drop”, and “split”). We fit IL policies to each dataset and run
 507 trajectory labeling to determine policy behavior, shown in Table 3. The place and drop policies
 508 show bias towards executing place and drop trajectories respectively, but still perform the opposite
 509 behavior somewhat frequently. The split policy is somewhat biased towards dropping, likely because
 510 the 1-dim gripper action to drop is easier to learn under MSE loss than a 7-dim arm action to place.

511 Thus, data filtering can generally influence IL policy behavior, but additional methods are needed to
 512 fully control behavior (e.g. online finetuning, reward relabeling, more advanced IL methods, etc).

513
 514 Table 3: Success once rate (**S-Once**, %) and ratio of “place in goal” to “drop to goal” trajectories
 515 (**Place : Drop**). Note that some success trajectories are not labeled place in goal or drop to goal, as
 516 there are other possible success modes described in Appendix A.5.

	FILTERS	SPLIT	S-ONCE	PLACE : DROP
Place in goal	Train	45.73	3.17 : 1	
	Val	54.46	2.55 : 1	
Drop to goal	Train	49.21	1 : 2.22	
	Val	51.19	1 : 2.86	
50/50 Split	Train	50.30	1 : 1.71	
	Val	55.56	1 : 1.41	

531 7 CONCLUSION AND LIMITATIONS

532
 533 We present MS-HAB a holistic home-scale rearrangement benchmark including a GPU-accelerated
 534 implementation of the HAB which supports realistic low-level control, extensive RL and IL base-
 535 lines, systematic evaluation using our trajectory labeling system, and demonstration filtering for
 536 efficient, controlled data generation at scale. However, there is significant room for improvement
 537 on our baselines, and we do not claim transfer to real robots; both of these we leave to future work.
 538 Whole-body low-level control under constraints in cluttered environments, long-horizon skill chain-
 539 ing, and scene-level rearrangement are challenging for current robot learning methods; we hope our
 540 benchmark and dataset aid the community in advancing these research areas.

540 8 REPRODUCIBILITY STATEMENT
 541

542 We open source all code for environments, training, evaluation, and data generation through an
 543 anonymized repository: <https://github.com/anonsubmit0/maniskill-hab>.

544 At the time of writing, the data used in this paper (490GB) is uploading to HuggingFace. While the
 545 data is uploading, users can generate data using the provided scripts in the source code. Depending
 546 on internet speed, generating may be faster than downloading.

548 549 REFERENCES
 550

551 Michael Bain and Claude Sammut. A framework for behavioural cloning. In Koichi Furukawa,
 552 Donald Michie, and Stephen H. Muggleton (eds.), *Machine Intelligence 15, Intelligent Agents*
 553 [St. Catherine’s College, Oxford, UK, July 1995], pp. 103–129. Oxford University Press, 1995.

554 Philip J. Ball, Laura M. Smith, Ilya Kostrikov, and Sergey Levine. Efficient online reinforce-
 555 ment learning with offline data. In Andreas Krause, Emma Brunskill, Kyunghyun Cho, Barbara
 556 Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *International Conference on Machine*
 557 *Learning, ICML 2023, 23-29 July 2023, Honolulu, Hawaii, USA*, volume 202 of *Proceedings of*
 558 *Machine Learning Research*, pp. 1577–1594. PMLR, 2023. URL <https://proceedings.mlr.press/v202/ball23a.html>.

559
 560 Gabriel Barth-Maron, Matthew W. Hoffman, David Budden, Will Dabney, Dan Horgan, Dhruva TB,
 561 Alistair Muldal, Nicolas Heess, and Timothy P. Lillicrap. Distributed distributional deterministic
 562 policy gradients. In *6th International Conference on Learning Representations, ICLR 2018, Van-*
 563 *couver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net,
 564 2018. URL <https://openreview.net/forum?id=SyZipzbCb>.

565
 566 Mariusz Bojarski, Davide Del Testa, Daniel Dworakowski, Bernhard Firner, Beat Flepp, Prasoon
 567 Goyal, Lawrence D. Jackel, Mathew Monfort, Urs Muller, Jiakai Zhang, Xin Zhang, Jake Zhao,
 568 and Karol Zieba. End to end learning for self-driving cars. *CoRR*, abs/1604.07316, 2016. URL
 569 <http://arxiv.org/abs/1604.07316>.

570
 571 Anthony Brohan, Noah Brown, Justice Carbajal, Yevgen Chebotar, Joseph Dabis, Chelsea Finn,
 572 Keerthana Gopalakrishnan, Karol Hausman, Alexander Herzog, Jasmine Hsu, Julian Ibarz, Brian
 573 Ichter, Alex Irpan, Tomas Jackson, Sally Jesmonth, Nikhil J. Joshi, Ryan Julian, Dmitry Kalash-
 574 nikov, Yuheng Kuang, Isabel Leal, Kuang-Huei Lee, Sergey Levine, Yao Lu, Utsav Malla, Deek-
 575 sha Manjunath, Igor Mordatch, Ofir Nachum, Carolina Parada, Jodilyn Peralta, Emily Perez, Karl
 576 Pertsch, Jornell Quiambao, Kanishka Rao, Michael S. Ryoo, Grecia Salazar, Pannag R. Sanketi,
 577 Kevin Sayed, Jaspiar Singh, Sumedh Sontakke, Austin Stone, Clayton Tan, Huong T. Tran, Vin-
 578 cent Vanhoucke, Steve Vega, Quan Vuong, Fei Xia, Ted Xiao, Peng Xu, Sichun Xu, Tianhe Yu,
 579 and Brianna Zitkovich. RT-1: robotics transformer for real-world control at scale. In Kostas E.
 580 Bekris, Kris Hauser, Sylvia L. Herbert, and Jingjin Yu (eds.), *Robotics: Science and Systems XIX,*
 581 *Daegu, Republic of Korea, July 10-14, 2023*, 2023. doi: 10.15607/RSS.2023.XIX.025. URL
 582 <https://doi.org/10.15607/RSS.2023.XIX.025>.

583
 584 Berk Çalli, Aaron Walsman, Arjun Singh, Siddhartha S. Srinivasa, Pieter Abbeel, and Aaron M. Dol-
 585 lar. Benchmarking in manipulation research: The YCB object and model set and benchmarking
 586 protocols. *CoRR*, abs/1502.03143, 2015. URL <http://arxiv.org/abs/1502.03143>.

587
 588 Yuanpei Chen, Chen Wang, Li Fei-Fei, and C. Karen Liu. Sequential dexterity: Chaining dexterous
 589 policies for long-horizon manipulation. *CoRR*, abs/2309.00987, 2023. doi: 10.48550/ARXIV.
 590 2309.00987. URL <https://doi.org/10.48550/arXiv.2309.00987>.

591
 592 Shreyansh Daftary, J. Andrew Bagnell, and Martial Hebert. Learning transferable policies for monoco-
 593 lular reactive MAV control. In Dana Kulic, Yoshihiko Nakamura, Oussama Khatib, and Gen-
 594 tiane Venture (eds.), *International Symposium on Experimental Robotics, ISER 2016, Tokyo,*
 595 *Japan, October 3-6, 2016*, volume 1 of *Springer Proceedings in Advanced Robotics*, pp. 3–11.
 596 Springer, 2016. doi: 10.1007/978-3-319-50115-4_1. URL https://doi.org/10.1007/978-3-319-50115-4_1.

- 594 Sudeep Dasari, Frederik Ebert, Stephen Tian, Suraj Nair, Bernadette Bucher, Karl Schmeckpeper,
 595 Siddharth Singh, Sergey Levine, and Chelsea Finn. Robonet: Large-scale multi-robot learning.
 596 In Leslie Pack Kaelbling, Danica Kragic, and Komei Sugiura (eds.), *3rd Annual Conference
 597 on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*,
 598 volume 100 of *Proceedings of Machine Learning Research*, pp. 885–897. PMLR, 2019. URL
 599 <http://proceedings.mlr.press/v100/dasari20a.html>.
- 600 Matt Deitke, Eli VanderBilt, Alvaro Herrasti, Luca Weihs, Kiana Ehsani, Jordi Salvador, Win-
 601 son Han, Eric Kolve, Aniruddha Kembhavi, and Roozbeh Mottaghi. 12796065039 pro-
 602 thor: Large-scale embodied AI using procedural generation. In Sanmi Koyejo, S. Mo-
 603 hamed, A. Agarwal, Danielle Belgrave, K. Cho, and A. Oh (eds.), *Advances in Neural
 604 Information Processing Systems 35: Annual Conference on Neural Information Process-
 605 ing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9,
 606 2022*. URL http://papers.nips.cc/paper_files/paper/2022/hash/27c546ab1e4f1d7d638e6a8dfbad9a07-Abstract-Conference.html.
- 607 Frederik Ebert, Yanlai Yang, Karl Schmeckpeper, Bernadette Bucher, Georgios Georgakis, Kostas
 608 Daniilidis, Chelsea Finn, and Sergey Levine. Bridge data: Boosting generalization of robotic
 609 skills with cross-domain datasets. In Kris Hauser, Dylan A. Shell, and Shoudong Huang (eds.),
 610 *Robotics: Science and Systems XVIII, New York City, NY, USA, June 27 - July 1, 2022*, 2022.
 611 doi: 10.15607/RSS.2022.XVIII.063. URL <https://doi.org/10.15607/RSS.2022.XVIII.063>.
- 612 Abby O'Neill et al. Open x-embodiment: Robotic learning datasets and RT-X models : Open
 613 x-embodiment collaboration. In *IEEE International Conference on Robotics and Auto-
 614 mation, ICRA 2024, Yokohama, Japan, May 13-17, 2024*, pp. 6892–6903. IEEE, 2024. doi:
 615 10.1109/ICRA57147.2024.10611477. URL <https://doi.org/10.1109/ICRA57147.2024.10611477>.
- 616 Justin Fu, Aviral Kumar, Ofir Nachum, George Tucker, and Sergey Levine. D4RL: datasets for deep
 617 data-driven reinforcement learning. *CoRR*, abs/2004.07219, 2020. URL <https://arxiv.org/abs/2004.07219>.
- 618 Chuang Gan, Siyuan Zhou, Jeremy Schwartz, Seth Alter, Abhishek Bhandwaldar, Dan Gutfreund,
 619 Daniel L. K. Yamins, James J. DiCarlo, Josh H. McDermott, Antonio Torralba, and Joshua B.
 620 Tenenbaum. The threedworld transport challenge: A visually guided task-and-motion plan-
 621 ning benchmark towards physically realistic embodied AI. In *2022 International Conference
 622 on Robotics and Automation, ICRA 2022, Philadelphia, PA, USA, May 23-27, 2022*, pp. 8847–
 623 8854. IEEE, 2022. doi: 10.1109/ICRA46639.2022.9812329. URL <https://doi.org/10.1109/ICRA46639.2022.9812329>.
- 624 Jiayuan Gu, Devendra Singh Chaplot, Hao Su, and Jitendra Malik. Multi-skill mobile manipu-
 625 lation for object rearrangement. In *The Eleventh International Conference on Learning Rep-
 626 resentations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net, 2023a. URL
 627 https://openreview.net/forum?id=Z3IC1M_bzvP.
- 628 Jiayuan Gu, Fanbo Xiang, Xuanlin Li, Zhan Ling, Xiqiang Liu, Tongzhou Mu, Yihe Tang, Stone
 629 Tao, Xinyue Wei, Yunchao Yao, Xiaodi Yuan, Pengwei Xie, Zhiao Huang, Rui Chen, and Hao Su.
 630 Maniskill2: A unified benchmark for generalizable manipulation skills. In *The Eleventh Inter-
 631 national Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
 632 OpenReview.net, 2023b. URL https://openreview.net/forum?id=b_CQDy9vrD1.
- 633 Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy
 634 maximum entropy deep reinforcement learning with a stochastic actor. In Jennifer G. Dy and
 635 Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning,
 636 ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings
 637 of Machine Learning Research*, pp. 1856–1865. PMLR, 2018. URL <http://proceedings.mlr.press/v80/haarnoja18b.html>.
- 638 Shengyi Huang, Rousslan Fernand Julien Dossa, Chang Ye, Jeff Braga, Dipam Chakraborty, Ki-
 639 nal Mehta, and João G.M. Araújo. Cleanrl: High-quality single-file implementations of deep

- 648 reinforcement learning algorithms. *Journal of Machine Learning Research*, 23(274):1–18, 2022.
 649 URL <http://jmlr.org/papers/v23/21-1342.html>.
 650
- 651 Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot
 652 learning benchmark & learning environment. *IEEE Robotics Autom. Lett.*, 5(2):3019–3026,
 653 2020. doi: 10.1109/LRA.2020.2974707. URL <https://doi.org/10.1109/LRA.2020.2974707>.
 654
- 655 Alexander Khazatsky, Karl Pertsch, Suraj Nair, Ashwin Balakrishna, Sudeep Dasari, Siddharth
 656 Karamcheti, Soroush Nasiriany, Mohan Kumar Srirama, Lawrence Yunliang Chen, Kirsty Ellis,
 657 Peter David Fagan, Joey Hejna, Masha Itkina, Marion Lepert, Yecheng Jason Ma, Patrick Tree
 658 Miller, Jimmy Wu, Suneel Belkhale, Shivin Dass, Huy Ha, Arhan Jain, Abraham Lee, Young-
 659 woon Lee, Marius Memmel, Sungjae Park, Ilija Radosavovic, Kaiyuan Wang, Albert Zhan, Kevin
 660 Black, Cheng Chi, Kyle Beltran Hatch, Shan Lin, Jingpei Lu, Jean Mercat, Abdul Rehman,
 661 Pannag R. Sanketi, Archit Sharma, Cody Simpson, Quan Vuong, Homer Rich Walke, Blake
 662 Wulfe, Ted Xiao, Jonathan Heewon Yang, Arefeh Yavary, Tony Z. Zhao, Christopher Agia, Ro-
 663 han Baijal, Mateo Guaman Castro, Daphne Chen, Qiuyu Chen, Trinity Chung, Jaimyn Drake,
 664 Ethan Paul Foster, and et al. DROID: A large-scale in-the-wild robot manipulation dataset. *CoRR*,
 665 abs/2403.12945, 2024. doi: 10.48550/ARXIV.2403.12945. URL <https://doi.org/10.48550/arXiv.2403.12945>.
 666
- 667 Eric Kolve, Roozbeh Mottaghi, Daniel Gordon, Yuke Zhu, Abhinav Gupta, and Ali Farhadi. AI2-
 668 THOR: an interactive 3d environment for visual AI. *CoRR*, abs/1712.05474, 2017. URL <http://arxiv.org/abs/1712.05474>.
 669
- 670 Youngwoon Lee, Joseph J. Lim, Anima Anandkumar, and Yuke Zhu. Adversarial skill chaining
 671 for long-horizon robot manipulation via terminal state regularization. In Aleksandra Faust, David
 672 Hsu, and Gerhard Neumann (eds.), *Conference on Robot Learning, 8-11 November 2021, London,
 673 UK*, volume 164 of *Proceedings of Machine Learning Research*, pp. 406–416. PMLR, 2021. URL
 674 <https://proceedings.mlr.press/v164/lee22a.html>.
 675
- 676 Chengshu Li, Fei Xia, Roberto Martín-Martín, Michael Lingelbach, Sanjana Srivastava, Bokui Shen,
 677 Kent Elliott Vainio, Cem Gokmen, Gokul Dharan, Tanish Jain, Andrey Kurenkov, C. Karen Liu,
 678 Hyowon Gweon, Jiajun Wu, Li Fei-Fei, and Silvio Savarese. igibson 2.0: Object-centric sim-
 679 ulation for robot learning of everyday household tasks. In Aleksandra Faust, David Hsu, and
 680 Gerhard Neumann (eds.), *Conference on Robot Learning, 8-11 November 2021, London, UK*,
 681 volume 164 of *Proceedings of Machine Learning Research*, pp. 455–465. PMLR, 2021. URL
 682 <https://proceedings.mlr.press/v164/li22b.html>.
 683
- 684 Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-
 685 Martín, Chen Wang, Gabrael Levine, Michael Lingelbach, Jiankai Sun, Mona Anvari, Minjune
 686 Hwang, Manasi Sharma, Arman Aydin, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan
 687 Lou, Caleb R. Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Silvio
 688 Savarese, Hyowon Gweon, C. Karen Liu, Jiajun Wu, and Li Fei-Fei. BEHAVIOR-1K: A bench-
 689 mark for embodied AI with 1,000 everyday activities and realistic simulation. In Karen Liu, Dana
 690 Kulic, and Jeffrey Ichnowski (eds.), *Conference on Robot Learning, CoRL 2022, 14-18 December
 691 2022, Auckland, New Zealand*, volume 205 of *Proceedings of Machine Learning Research*, pp.
 692 80–93. PMLR, 2022. URL <https://proceedings.mlr.press/v205/li23a.html>.
 693
- 694 Ajay Mandlekar, Yuke Zhu, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao,
 695 John Emmons, Anchit Gupta, Emre Orbay, Silvio Savarese, and Li Fei-Fei. ROBOTURK: A
 696 crowdsourcing platform for robotic skill learning through imitation. In *2nd Annual Conference on
 697 Robot Learning, CoRL 2018, Zürich, Switzerland, 29-31 October 2018, Proceedings*, volume 87
 698 of *Proceedings of Machine Learning Research*, pp. 879–893. PMLR, 2018. URL <http://proceedings.mlr.press/v87/mandlekar18a.html>.
 699
- 700 Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-
 701 Fei, Silvio Savarese, Yuke Zhu, and Roberto Martín-Martín. What matters in learning from offline
 702 human demonstrations for robot manipulation. In Aleksandra Faust, David Hsu, and Gerhard
 703 Neumann (eds.), *Conference on Robot Learning, 8-11 November 2021, London, UK*, volume 164
 704 of *Proceedings of Machine Learning Research*, pp. 1678–1690. PMLR, 2021. URL <https://proceedings.mlr.press/v164/mandlekar22a.html>.
 705

- 702 Ajay Mandlekar, Soroush Nasiriany, Bowen Wen, Iretiayo Akinola, Yashraj S. Narang, Linxi Fan,
 703 Yuke Zhu, and Dieter Fox. Mimicgen: A data generation system for scalable robot learning
 704 using human demonstrations. In Jie Tan, Marc Toussaint, and Kourosh Darvish (eds.), *Con-*
 705 *ference on Robot Learning, CoRL 2023, 6-9 November 2023, Atlanta, GA, USA*, volume 229
 706 of *Proceedings of Machine Learning Research*, pp. 1820–1864. PMLR, 2023. URL <https://proceedings.mlr.press/v229/mandlekar23a.html>.
- 707
- 708 Detlef Mewes and Fritz Mauser. Safeguarding crushing points by limitation of forces. *Inter-*
 709 *national journal of occupational safety and ergonomics : JOSE*, 9:177–91, 02 2003. doi:
 710 10.1080/10803548.2003.11076562.
- 711
- 712 Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A. Rusu, Joel Veness, Marc G.
 713 Bellemare, Alex Graves, Martin A. Riedmiller, Andreas Fidjeland, Georg Ostrovski, Stig Pe-
 714 tersen, Charles Beattie, Amir Sadik, Ioannis Antonoglou, Helen King, Dharshan Kumaran,
 715 Daan Wierstra, Shane Legg, and Demis Hassabis. Human-level control through deep rein-
 716 forcement learning. *Nat.*, 518(7540):529–533, 2015. doi: 10.1038/NATURE14236. URL
 717 <https://doi.org/10.1038/nature14236>.
- 718
- 719 Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi,
 720 Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for
 721 generalist robots. *CoRR*, abs/2406.02523, 2024. doi: 10.48550/ARXIV.2406.02523. URL
 722 <https://doi.org/10.48550/arXiv.2406.02523>.
- 723
- 724 Stéphane Ross, Geoffrey J. Gordon, and Drew Bagnell. A reduction of imitation learning and
 725 structured prediction to no-regret online learning. In Geoffrey J. Gordon, David B. Dunson,
 726 and Miroslav Dudík (eds.), *Proceedings of the Fourteenth International Conference on Artificial*
727 Intelligence and Statistics, AISTATS 2011, Fort Lauderdale, USA, April 11-13, 2011, volume 15
 728 of *JMLR Proceedings*, pp. 627–635. JMLR.org, 2011. URL <http://proceedings.mlr.press/v15/ross11a/ross11a.pdf>.
- 729
- 730 Manolis Savva, Jitendra Malik, Devi Parikh, Dhruv Batra, Abhishek Kadian, Oleksandr Maksymets,
 731 Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, and Vladlen Koltun. Habitat: A
 732 platform for embodied AI research. In *2019 IEEE/CVF International Conference on Computer*
 733 *Vision, ICCV 2019, Seoul, Korea (South), October 27 - November 2, 2019*, pp. 9338–9346. IEEE,
 734 2019. doi: 10.1109/ICCV.2019.00943. URL <https://doi.org/10.1109/ICCV.2019.00943>.
- 735
- 736 John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy
 737 optimization algorithms. *CoRR*, abs/1707.06347, 2017. URL <http://arxiv.org/abs/1707.06347>.
- 738
- 739 Andrew Szot, Alexander Clegg, Eric Undersander, Erik Wijmans, Yili Zhao, John M. Turner,
 740 Noah Maestre, Mustafa Mukadam, Devendra Singh Chaplot, Oleksandr Maksymets, Aaron
 741 Gokaslan, Vladimir Vondrus, Sameer Dharur, Franziska Meier, Wojciech Galuba, Angel X.
 742 Chang, Zsolt Kira, Vladlen Koltun, Jitendra Malik, Manolis Savva, and Dhruv Batra. Hab-
 743 itat 2.0: Training home assistants to rearrange their habitat. In Marc’Aurelio Ranzato, Alina
 744 Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (eds.), *Ad-*
745 vances in Neural Information Processing Systems 34: Annual Conference on Neural In-
746 formation Processing Systems 2021, NeurIPS 2021, December 6-14, 2021, virtual, pp.
 747 251–266, 2021. URL <https://proceedings.neurips.cc/paper/2021/hash/021bbc7ee20b71134d53e20206bd6feb-Abstract.html>.
- 748
- 749 Stone Tao, Fanbo Xiang, Arth Shukla, Yuzhe Qin, Xander Hinrichsen, Xiaodi Yuan, Chen Bao,
 750 Xinsong Lin, Yulin Liu, Tse kai Chan, Yuan Gao, Xuanlin Li, Tongzhou Mu, Nan Xiao, Ar-
 751 nav Gurha, Zhiao Huang, Roberto Calandra, Rui Chen, Shan Luo, and Hao Su. Maniskill3:
 752 Gpu parallelized robotics simulation and rendering for generalizable embodied ai. *arXiv preprint*
753 arXiv:2410.00425, 2024.
- 754
- 755 Fanbo Xiang, Yuzhe Qin, Kaichun Mo, Yikuan Xia, Hao Zhu, Fangchen Liu, Minghua Liu, Hanxiao
 Jiang, Yifu Yuan, He Wang, Li Yi, Angel X. Chang, Leonidas J. Guibas, and Hao Su. SAPIEN:
 A simulated part-based interactive environment. In *2020 IEEE/CVF Conference on Computer*

- 756 *Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pp. 11094–
 757 11104. Computer Vision Foundation / IEEE, 2020. doi: 10.1109/CVPR42600.2020.01111.
 758 URL https://openaccess.thecvf.com/content_CVPR_2020/html/Xiang_SAPIEN_A_SimulAted_Part-Based_Interactive_ENvironment_CVPR_2020_paper.html.
- 761 Jinwei Xing. Pytorch implementations of reinforcement learning algorithms for visual continuous
 762 control. <https://github.com/KarlXing/RL-Visual-Continuous-Control>,
 763 2022.
- 764 Sherry Yang, Yilun Du, Seyed Kamyar Ghasemipour, Jonathan Tompson, Leslie Pack Kael-
 765 bling, Dale Schuurmans, and Pieter Abbeel. Learning interactive real-world simulators. In
 766 *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Aus-
 767 tria, May 7-11, 2024*. OpenReview.net, 2024. URL <https://openreview.net/forum?id=sFyTZEqmUY>.
- 770 ZebraTechnologies. Autonomous mobile robots, 2024. URL <https://www.zebra.com/us/en/products/autonomous-mobile-robots.html>.
- 772 Yuke Zhu, Josiah Wong, Ajay Mandlekar, and Roberto Martín-Martín. robosuite: A modular
 773 simulation framework and benchmark for robot learning. *CoRR*, abs/2009.12293, 2020. URL
 774 <https://arxiv.org/abs/2009.12293>.
- 775

A APPENDIX

A.1 SUBTASK DEFINITIONS AND INITIALIZATION

A.1.1 GENERAL INITIALIZATION

We refer to the robot end-effector as ee , and its rest position as r_{pos} . The end-effector resting position is (0.5 m, 0 m, 1.25 m) relative to the base³. Let q_{arm} be the arm joint positions, r_{arm} be the arm resting joint positions, and \dot{q}_{arm} be the arm joint velocities. Similarly, for the torso define q_{tor} , r_{tor} , and \dot{q}_{tor} . Let v_{base} be the base linear velocity in $m s^{-1}$, and $v_{base,x}, v_{base,y}$ its x and y components respectively. Let ω_{base} be the base angular velocity in $rad s^{-1}$.

We initialize the robot to $r_{pos}, r_{arm}, r_{tor}$ with $\dot{q}_{arm} = (0 \dots 0), \dot{q}_{tor} = 0, v_{base} = 0, \omega_{base} = 0$. Then, q_{arm} is perturbed by clipped Gaussian noise $\text{clip}(\mathcal{N}(0, 0.1), -0.2, 0.2)$, the base position is perturbed by $\text{clip}(\mathcal{N}(0, 0.1), -0.2, 0.2)$, and the base rotation is perturbed by $\text{clip}(\mathcal{N}(0, 0.25), -0.5, 0.5)$.

A.2 SUBTASK DEFINITIONS

Let $d_b^a = \|a_{pos} - b_{pos}\|_2^2$, units in m. For example, d_{ee}^r is the distance in m between the end-effector and its rest position. Next, let $j_k = \max_{1 \leq i \leq |q_k|} |q_{k,i} - r_{k,i}|$. For example, j_{arm} is the maximum absolute difference in the arm joint positions and corresponding resting positions. Let $C_{[0:t]}$ be the cumulative robot collisions in N until step t . Finally, we define two commonly-used success conditions for the Fetch robot:

$$\mathbf{1}_{\text{grasped}(x)} = \mathbf{1}\{\text{x is grasped (computed by simulator)}\}$$

$$\mathbf{1}_{\text{is_static}} = \mathbf{1}\left\{\left(\max_{1 \leq i \leq |q_{arm}|} |\dot{q}_{arm,i}| \leq 0.2\right) \wedge v_{base,x} \leq 0.05 \wedge v_{base,y} \leq 0.05 \wedge \omega_{base} \leq 0.05\right\}$$

Pick[a , optional](x_{pose}): Pick object x (from articulation a , if provided).

- Initialization: Spawn robot facing x , within 2 m of x , with noise, and without collisions.
- Success:

$$\mathbf{1}_{\text{grasped}(x)} \wedge d_{ee}^r \leq 0.05 \wedge j_{arm} \leq 0.6 \wedge \mathbf{1}_{\text{is_static}} \wedge C_{[0:t]} \leq 5000$$

³The z-axis is ‘up’ in ManiSkill3.

- 810 • Failure: $C_{[0:t]} > 5000 \text{ N}$
 811

812 **Place**[a , optional](x_{pose} , g_{pos}): Place object x at goal g (in articulation a , if provided).

- 813
 814 • Initialization: Spawn with grasp pose sampled from $\text{Pick}(x_{pose})$ policy, robot facing g ,
 815 within 2 m of g , with noise, and without collisions.
 816 • Success:
 817
 818 $\neg \mathbf{1}_{\text{grasped}(x)} \wedge d_x^g \leq 0.15 \wedge d_{ee}^r \leq 0.05 \wedge j_{arm} \leq 0.2 \wedge j_{tor} \leq 0.01 \wedge \mathbf{1}_{\text{is_static}} \wedge C_{[0:t]} \leq 7500$
 819
 820 • Failure: $C_{[0:t]} > 7500 \text{ N}$

821 **Open**[a](a_{pos}): Open articulation a with handle at a_{pos} .

- 822
 823 • Initialization: Spawn facing a . If a is a fridge, spawn within $[0.933, -0.6] \times [1.833, 0.6]$
 824 region in front of a , otherwise within $[0.3, -0.6] \times [1.5, 0.6]$. With noise, without collisions.
 825
 826 • Success: Let a_q, a_{qmax}, a_{qmin} be the current, max, and min joint positions for the target
 827 articulation (drawer or fridge). Then, let $a_{ofrac} = \{0.75 \text{ if } a \text{ is a fridge else } 0.9\}$. We
 828 define

$$\mathbf{1}_{\text{open}(a)} = \mathbf{1} \{a_q \geq a_{ofrac} \cdot (a_{qmax} - a_{qmin}) + a_{qmin}\}$$

829 Hence, we have success condition

$$\mathbf{1}_{\text{open}(a)} \wedge d_{ee}^r \leq 0.05 \wedge j_{arm} \leq 0.2 \wedge j_{tor} \leq 0.01 \wedge \mathbf{1}_{\text{is_static}} \wedge C_{[0:t]} \leq 10000$$

- 830
 831 • Failure⁴: $C_{[0:t]} > 10000 \text{ N}$

832 **Close**[a](a_{pos}): Close articulation a with handle at a_{pos} .

- 833
 834 • Initialization: Spawn facing a . If a is a fridge, spawn within $[0.933, -0.6] \times [1.833, 0.6]$
 835 region in front of a , otherwise within $[0.3, -0.6] \times [1.5, 0.6]$. With noise, without collisions.
 836
 837 • Success: Let a_q, a_{qmax}, a_{qmin} be the current, max, and min joint positions for the target
 838 articulation (drawer or fridge). We define

$$\mathbf{1}_{\text{close}(a)} = \mathbf{1} \{a_q \leq 0.01 \cdot (a_{qmax} - a_{qmin}) + a_{qmin}\}$$

839 Hence, we have success condition

$$\mathbf{1}_{\text{close}(a)} \wedge d_{ee}^r \leq 0.05 \wedge j_{arm} \leq 0.2 \wedge j_{tor} \leq 0.01 \wedge \mathbf{1}_{\text{is_static}} \wedge C_{[0:t]} \leq 10000$$

- 840
 841 • Failure⁴: $C_{[0:t]} > 10000 \text{ N}$

842 A.3 RL SUBTASK EVALUATION CURVES

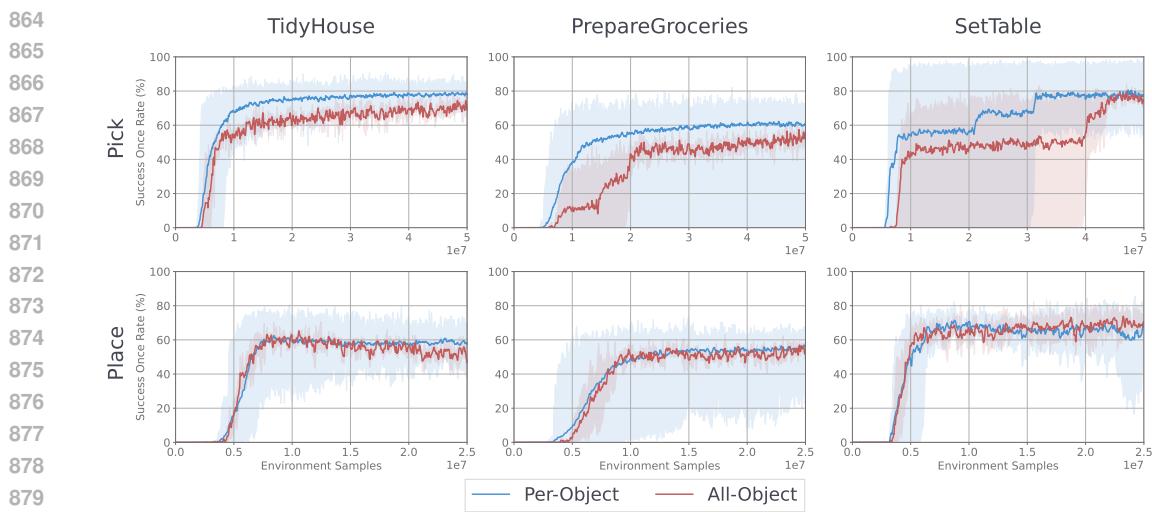
843
 844
 845 During training, we evaluate our policies every 10000 steps on 189 episodes. The per vs all-object
 846 training curves in Fig. 4 demonstrate a similar trend as seen in Sec. 6.2.1: per-object policies show
 847 the most significant improvements when grasping many objects with different geometries (Tidy-
 848 House Pick and PrepareGroceries Pick) or when manipulating objects in tight constraints where
 849 object geometry is important (PrepareGroceries Place).

850
 851 Per Tables 9-12, the performance limitations for Open and Close seen in Fig. 5 are caused primarily
 852 by the 10 000 N cumulative robot force limit we set, which is not used in the original implementation
 853 of the HAB (Szot et al., 2021).

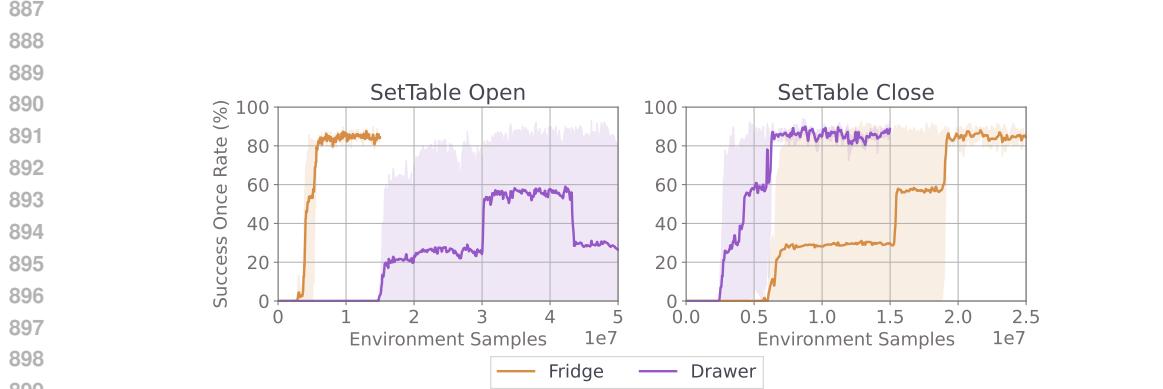
854 A.4 ADDITIONAL EXPERIMENTS

855 A.4.1 DATASET SIZE

856
 857 ⁴Originally, the HAB does not specify collision limits for Open or Close, but we add them to enforce safety.



882 Figure 4: Per-object vs all-object RL success once rate (%) evaluation curves for Pick and Place
883 policies across tasks. We run 3 seeds for each per-object policy and 3 seeds for the all-object policy.
884 TidyHouse and PrepareGroceries involve 9 objects, while SetTable involves 2 objects. Since we
885 group runs for different per-object policies into one curve, we use minimum and maximum for the
886 shaded region. Best viewed zoomed.



901 Figure 5: Open and Close training success once rate (%) curves for Drawer and Fridge. Since
902 success once rate jumps very quickly once the policy learns to solve the task, we use minimum and
903 maximum for the shaded region.

904
905
906
907 To highlight the importance of generating scal-
908 able datasets, we train IL policies for Tidy-
909 House/PrepareGroceries/SetTable Pick/Place subtasks
910 at 1, 10, 100, 500, and 1000 demonstrations per object.
911 In Table 4, we run 1000 evaluation episodes per policy,
912 and group results by demonstrations per object. We then
913 report average success once rate and 95% CIs for each
914 demonstrations per object value.

915 We find that 1000 demonstrations per object leads to the
916 most performant policies. Furthermore there are large
917 jumps in success rate as demonstrations per object in-
creases from 10 to 100 to 500.

Table 4: Success once rate (**SoR**) with
95% CIs depending on demos per object
(Demos).

DEMOS	SoR
1	0.00 ± 0.00
10	0.02 ± 0.03
100	0.27 ± 0.19
500	0.53 ± 0.13
1000	0.62 ± 0.09

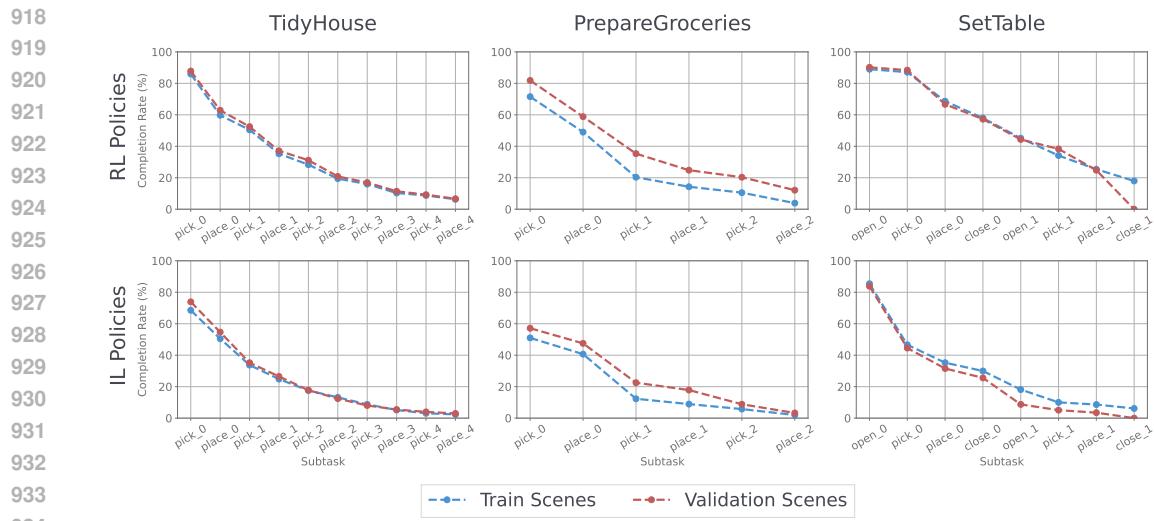


Figure 6: Progressive completion rate (%) on simplified long-horizon tasks with RL and IL policies. We remove all collision requirements, and allow placing on the full target receptacle surface. Best viewed zoomed.

A.4.2 PERFORMANCE WITH TASK SIMPLIFICATIONS

In Sec. 6.1, we find that improving subtask success rate is the most effective way to increase long-horizon task success rate. In Fig. 6, we simplify the long-horizon tasks by (1) removing all collision requirements, and (2) marking Place subtasks successful if the object remains anywhere on the target receptacle surface. We find that progressive completion rate increases for both our RL and IL policies, but we achieve no higher than 20% overall task success on any task or split. Hence, the largest challenge in training subtask policies is low-level whole-body control. Meanwhile, collision requirements and subtask success conditions pose some difficulty, but not as much.

A.4.3 SAC vs PPO FOR RL TRAINING

To justify our choices of RL algorithm for each policy, we compare SAC and PPO performance across tasks and subtasks. For Pick and Place, we compare all-object SAC and PPO subtask success once rate (%) on the train split. As described in Sec. 5.1, we train SAC with 25 million samples. Since PPO trains faster wall-time, we provide it 50 millions samples for a fair comparison.

Furthermore, for Open and Close, we compare per-object success once rate (%). We provide PPO with the same total samples as listed in Sec. 5.1, and we train SAC with 20 million samples per run.

For Pick and Place, despite training on double the total samples, PPO policies achieve notably lower performance compared to the SAC policies. We hypothesize this is because our large replay buffer can store a greater diversity of examples across objects, spawn locations, and obstructions, allowing SAC policies to better learn manipulation in diverse settings (Haarnoja et al., 2018). Hence, we use SAC for RL Pick and Place baselines.

Meanwhile, in Open Fridge and Close Drawer PPO policies perform better than SAC policies, in Close Fridge PPO policies perform marginally worse than SAC policies, and in Open Drawer PPO policies perform better only with many more samples. Since performance between PPO and SAC is generally comparable in Open and Close, we choose PPO for our baselines since it has faster wall-time training. For consistency, we use the same RL algorithm across all Open and Close variants.

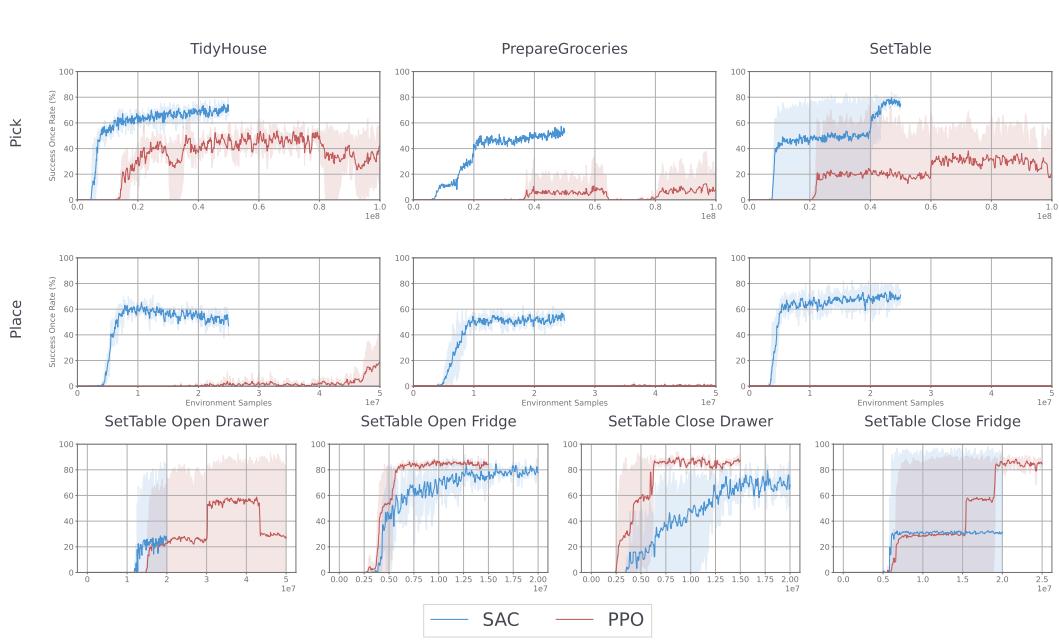


Figure 7: SAC vs PPO subtask success once rate (%) curves on the train split. Lines are averaged across 3 seeds; since success rate can jump rapidly, shaded regions represent min/max values. For Pick and Place, we compare all-object SAC and PPO policies, and for Open and Close, we compare per-object policies. Note that for PrepareGroceries and SetTable Place, lines are drawn but near-zero. Best viewed zoomed.

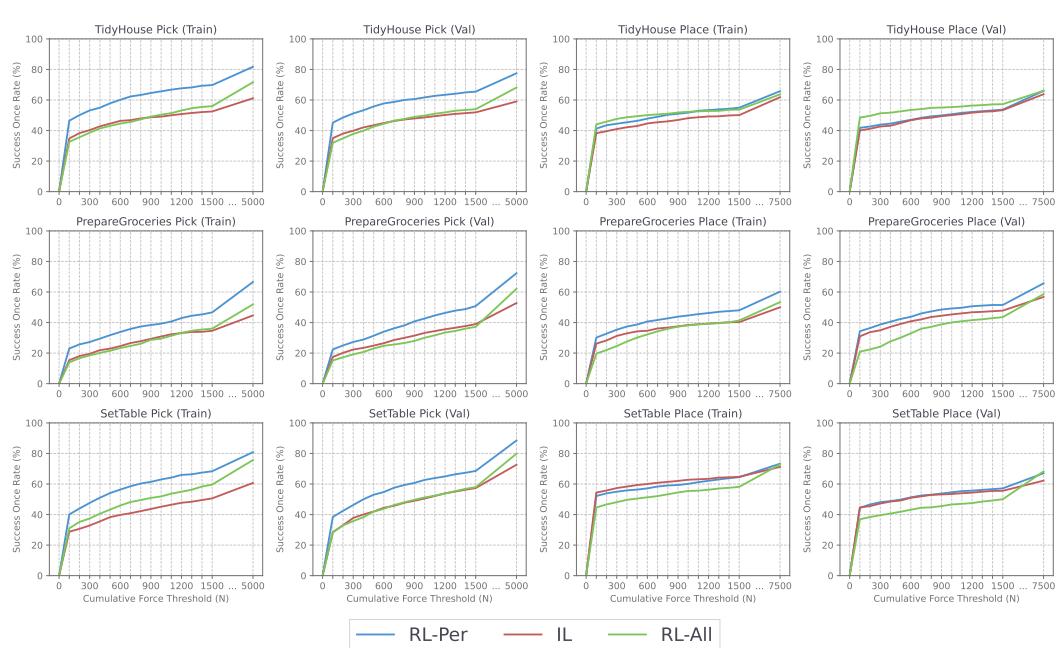


Figure 8: Success once rates (%) for RL-Per, RL-All, and IL policies in Pick and Place subtasks under varying cumulative collision thresholds. Best viewed zoomed.

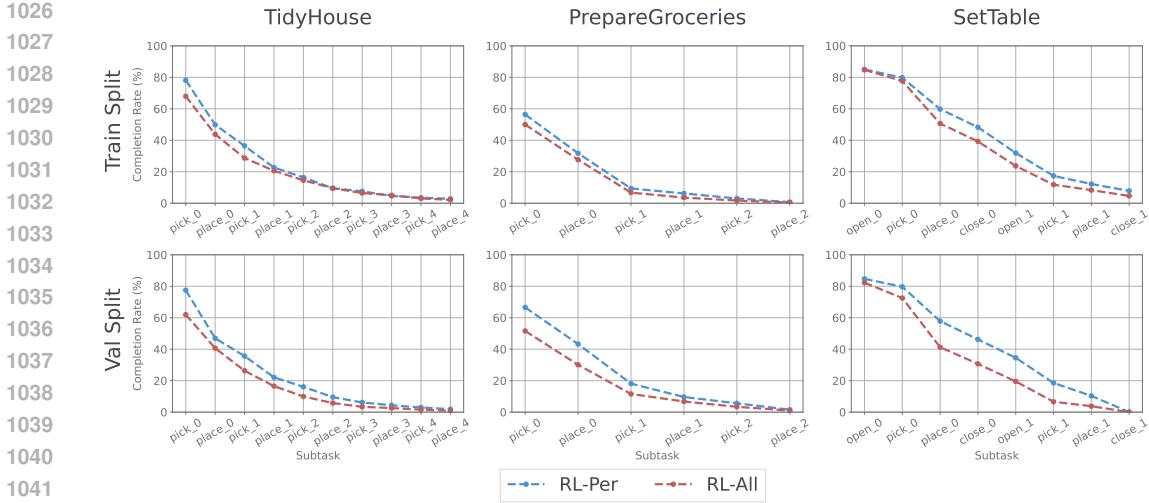


Figure 9: Success once rates (%) for RL-Per and RL-All policies on long-horizon tasks in train and validation split. Best viewed zoomed.

A.4.4 PERFORMANCE UNDER LOW COLLISION THRESHOLDS

To evaluate the safety of our policies in a real-world setting, we compare performance for RL-Per, RL-All, and IL policies on Pick and Place subtasks under low cumulative collision thresholds. Per industry safety standards, we use 1400 N for as the measure for safe execution (Mewes & Mauser, 2003).

As seen in Fig. 8, we observe a 5-20% drop in performance depending on subtask when using a cumulative collision threshold of 1400 N. Additionally, the per-object RL policies notably outperform all-object RL policies under lower collision thresholds.

A.4.5 PER VS ALL-OBJECT POLICY LONG-HORIZON PERFORMANCE

In addition to superior subtask success rates, as seen in Fig. 9, per-object RL policies outperform their all-object counterparts on full long horizon tasks in both train and validation splits.

A.5 TRAJECTORY CATEGORIZATION AND DATASET FILTERING

In this section, we provide definitions for our event labeling and trajectory categorization system. We additionally provide statistics on policy success and failure modes using our trajectory categorization system. Some example videos for Pick and Place failure modes are provided in the supplementary and project website.

A.5.1 DATASET FILTERING AND GENERATION

We generate 1000 demonstrations per object/articulation for each subtask using per-object RL policies on the train split. We use our trajectory labeling system to filter demonstrations (full definitions in Appendix A.5.2). For Pick, we require “straightforward success” demonstrations, where the agent successfully picks the object without dropping it while remaining within the cumulative collision threshold. For Place, we require “placed in goal success” demonstrations, where the agent releases the object within 15cm of the goal, the object stays in the goal without rolling or falling out, and the agent remains within the cumulative collision threshold. For Open and Close, we require “open success” and “closed success” demonstrations, where the agent opens/closes the articulation without excessive collisions, and the articulation remains within the open/close state.

1080 A.5.2 DEFINITIONS
1081

1082 For each success and failure mode definition, we provide a plain text description in addition to the
1083 boolean definitions. We heavily rely on notation used in Appendix A.1, in addition to those defined
1084 below.

1085 Using the criteria defined below, for each trajectory $\tau_{\text{subtask}} = (s_0, a_0, \dots, s_n, a_n)$, we create a
1086 chronologically ordered event list $E_{\text{subtask}} = (e_1, \dots, e_k)$. Using E_{subtask} , we categorize τ_{subtask}
1087 into a success/failure mode.

1088 Let $i_{\text{subtask}, \text{event}} = \{\text{index of } e_{\text{event}} \text{ in } E_{\text{subtask}} \text{ if } e_{\text{event}} \in E_{\text{subtask}} \text{ else } -1\}$. Also, let $F_{a,b,t}$ be
1089 the pairwise force between a and b in \mathbb{N} at timestep t .

1090 **Pick**[a , optional](x_{pose}): Pick object x (from articulation a , if provided).

1091

- 1093 Events: We define time-conditioned events at timestep t :

$$\begin{aligned} e_{\text{contact}} &= |F_{ee,x,t-1}| = 0 \wedge |F_{ee,x,t}| \geq 0 \\ e_{\text{grasped}} &= \neg \mathbf{1}_{\text{grasped}(x),t-1} \wedge \mathbf{1}_{\text{grasped}(x),t} \\ e_{\text{dropped}} &= \mathbf{1}_{\text{grasped}(x),t-1} \wedge \neg \mathbf{1}_{\text{grasped}(x),t} \\ e_{\text{success}} &= \neg \mathbf{1}_{\text{success},t-1} \wedge \mathbf{1}_{\text{success},t} \\ e_{\text{excessive_collisions}} &= C_{[0:t-1]} \leq 5000 \wedge C_{[0:t]} > 5000 \end{aligned}$$

1101 For $t \in \{1, \dots, n\}$ (in increasing order), we evaluate each e_{event} in the order shown above.
1102 If $e_{\text{event}} = 1$, we add it to E_{pick} .

- 1104 Success Modes: if $e_{\text{success}} \in E_{\text{pick}}$, then categorize using the following success modes:
1105
 - i Straightforward success: Agent successfully grasps x and returns to rest without dropping or excessive collisions.
1106 $E_{\text{pick}} = (e_{\text{contact}}, e_{\text{grasped}}, e_{\text{success}})$
 - ii Winding success: Agent (eventually) successfully grasps x (but drops x along the way) and returns to rest without excessive collisions.
1107 $E_{\text{pick}} = (e_{\text{contact}}, e_{\text{grasped}}, \dots, e_{\text{success}}) \wedge |E_{\text{pick}}| > 3 \wedge e_{\text{excessive_collisions}} \notin E_{\text{pick}}$
 - iii Success then drop: Agent successfully picks x and returns to rest without excessive collisions, but irrecoverably drops x after.
1108 $e_{\text{dropped}} \in E_{\text{pick}} \wedge i_{\text{subtask}, \text{dropped}} > i_{\text{subtask}, \text{grasped}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{pick}}$
 - iv Success then excessive collisions: Agent picks x and returns to rest, but exceeds collision threshold afterwards.
1109 $e_{\text{excessive_collisions}} \in E_{\text{pick}}$

- 1119 Failure Modes: if $e_{\text{success}} \notin E_{\text{pick}}$, then categorize using the following failure modes:
1120
 - v Excessive collision failure: Agent exceeds collision threshold.
1121 $e_{\text{excessive_collisions}} \in E_{\text{pick}}$
 - vi Mobility failure: Agent cannot reach x .
1122 $E_{\text{pick}} = ()$
 - vii Can't grasp failure: Agent reaches x , but cannot grasp it.
1123 $E_{\text{pick}} = (e_{\text{contact}})$
 - viii Drop failure: Agent grasps x , but drops it before returning to rest.
1124 $e_{\text{dropped}} \in E_{\text{pick}} \wedge i_{\text{subtask}, \text{dropped}} > i_{\text{subtask}, \text{grasped}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{pick}}$
 - ix Too slow failure: Agent (eventually) grasps x , but the episode truncates before it can reach success.
1125 $e_{\text{grasped}} \in E_{\text{pick}} \wedge i_{\text{subtask}, \text{grasped}} > i_{\text{subtask}, \text{dropped}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{pick}}$

1133 **Place**[a , optional]($x_{\text{pose}}, g_{\text{pos}}$): Place object x at goal g (in articulation a , if provided).

- 1134 • Events: We define time-conditioned events at timestep t :
- 1135
- $$e_{\text{grasped}} = \neg \mathbf{1}_{\text{grasped}(x), t-1} \wedge \mathbf{1}_{\text{grasped}(x), t}$$
- 1136
- $$e_{\text{obj_at_goal}} = d_{x,t-1}^g > 0.15 \wedge d_{x,t}^g \leq 0.15$$
- 1137
- $$e_{\text{released_at_goal}} = d_x^g \leq 0.15 \wedge \mathbf{1}_{\text{grasped}(x), t-1} \wedge \neg \mathbf{1}_{\text{grasped}(x), t}$$
- 1138
- $$e_{\text{released_outside_goal}} = d_x^g > 0.15 \wedge \mathbf{1}_{\text{grasped}(x), t-1} \wedge \neg \mathbf{1}_{\text{grasped}(x), t}$$
- 1139
- $$e_{\text{obj_left_goal}} = d_{x,t-1}^g \leq 0.15 \wedge d_{x,t}^g > 0.15$$
- 1140
- $$e_{\text{success}} = \neg \mathbf{1}_{\text{success}, t-1} \wedge \mathbf{1}_{\text{success}, t}$$
- 1141
- $$e_{\text{excessive_collisions}} = C_{[0:t-1]} \leq 7500 \wedge C_{[0:t]} > 7500$$
- 1142
- 1143
- 1144
- 1145 For $t \in \{1, \dots, n\}$ (in increasing order), we evaluate each e_{event} in the order shown above.
- 1146 If $e_{\text{event}} = 1$, we add it to E_{place} .
- 1147 • Success Modes: if $e_{\text{success}} \in E_{\text{place}}$, then categorize using the following success modes:
- 1148 i Place in goal success: Agent releases and successfully places x to within 15cm of g_{pos} , then returns to rest.
- 1149 $|E_{\text{place}}| \leq 4 \wedge (e_{\text{released_at_goal}} \in E_{\text{place}} \vee d_{x,0}^g \leq 0.15) \wedge i_{\text{place,obj_left_goal}} \leq i_{\text{place,obj_at_goal}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$
- 1150 ii Dropped to goal success: Agent releases x beyond 15cm of g_{pos} , x drops into the region within 15cm of g_{pos} , and the agent returns to rest.
- 1151 $|E_{\text{place}}| \leq 4 \wedge (e_{\text{released_outside_goal}} \in E_{\text{place}} \vee d_{x,0}^g > 0.15) \wedge i_{\text{place,obj_left_goal}} \leq i_{\text{place,obj_at_goal}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$
- 1152 iii Dubious success: x is manipulated to within 15cm of g_{pos} , and the robot returns to rest, but x leaves g before truncation.
- 1153 $i_{\text{place,obj_at_goal}} < i_{\text{place,obj_left_goal}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$
- 1154 iv Winding success: x leaves the goal at least once, but the agent (eventually) successfully places/drops x to within 15cm of g_{pos} , where it remains as the agent returns to rest.
- 1155 $|E_{\text{place}}| > 4 \wedge i_{\text{place,obj_at_goal}} > i_{\text{place,obj_left_goal}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$
- 1156 v Success then excessive collisions: The agent successfully places/drops x to within 15cm of g_{pos} and returns to rest, but exceeds collision threshold after.
- 1157 $e_{\text{excessive_collisions}} \in E_{\text{place}}$
- 1158 • Failure Modes: if $e_{\text{success}} \notin E_{\text{place}}$, then categorize using the following failure modes:
- 1159 vi Excessive collision failure: Agent exceeds collision threshold.
- 1160 $e_{\text{excessive_collisions}} \in E_{\text{place}}$
- 1161 vii Didn't grasp failure: Agent fails to grasp x at initialization.
- 1162 $E_{\text{place}} = () \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$
- 1163 viii Didn't reach goal failure: Agent grasps x , but cannot manipulate x to within 15cm of g_{pos} .
- 1164 $|E_{\text{place}}| > 0 \wedge e_{\text{obj_at_goal}} \notin E_{\text{place}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$
- 1165 ix Place in goal failure: Agent places x to within 15cm of g_{pos} , but x leaves this region (i.e., rolls or falls out) before the agent returns to rest.
- 1166 First, we define
- 1167 $\mathbf{1}_{\text{placed_is_latest_sequence}} = (|E_{\text{place}}| \leq 2 \wedge d_{x,0}^g \leq 0.15) \vee (i_{\text{place,released_at_goal}} > i_{\text{place,released_outside_goal}} \wedge i_{\text{place,released_at_goal}} > i_{\text{place,grasped}})$
- 1168 Hence, we have failure mode definition
- 1169 $e_{\text{obj_at_goal}} \in E_{\text{place}} \wedge \mathbf{1}_{\text{placed_is_latest_sequence}} \wedge i_{\text{place,obj_left_goal}} > i_{\text{place,obj_at_goal}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$
- 1170 x Dropped to goal failure: Agent drops x beyond 15cm away from g_{pos} , and x drops into the region within 15cm of g_{pos} , but leaves this region (i.e., rolls or falls out) before the agent returns to rest. First, we define
- 1171 $\mathbf{1}_{\text{dropped_is_latest_sequence}} = (|E_{\text{place}}| \leq 2 \wedge d_{x,0}^g > 0.15) \vee (i_{\text{place,released_outside_goal}} > i_{\text{place,released_at_goal}} \wedge i_{\text{place,released_at_goal}} > i_{\text{place,grasped}})$
- 1172 Hence, we have failure mode definition
- 1173 $e_{\text{obj_at_goal}} \in E_{\text{place}} \wedge \mathbf{1}_{\text{dropped_is_latest_sequence}} \wedge i_{\text{place,obj_left_goal}} > i_{\text{place,obj_at_goal}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{place}}$

- 1188 xi Won't let go failure: The agent is able to manipulate x to within 15cm of g_{pos} , but
 1189 does not release x .
 1190 $e_{obj_at_goal} \in E_{place} \wedge i_{place,grasped} > i_{place,released_at_goal} \wedge i_{place,grasped} >$
 1191 $i_{place,released,outside_goal} \wedge e_{excessive_collisions} \notin E_{place}$
- 1192 xii Too slow failure: The agent is able to manipulate x to within 15cm of the goal, releases
 1193 x , but is unable to return to rest before truncation.
 1194 The condition is no other failure mode is applicable. This also implies
 1195 $i_{place,obj_at_goal} > i_{place,obj_left_goal} \wedge e_{excessive_collisions} \notin E_{place}$

1196
 1197 **Open**[a](a_{pos}): Open articulation a with handle at a_{pos} .

- 1198 • Events: First, define indicator

1200 $\mathbf{1}_{\text{slightly_opened}(a),t} = \mathbf{1}\{a_{q,t} \geq 0.1 \cdot (a_{qmax} - a_{qmin}) + a_{qmin}\}$

1201 Now, we define time-conditioned events at timestep t :

1202 $e_{contact} = |F_{ee,a,t-1}| = 0 \wedge |F_{ee,a,t}| \geq 0$
 1203 $e_{opened} = \neg \mathbf{1}_{\text{open}(a),t-1} \wedge \mathbf{1}_{\text{open}(a),t}$
 1204 $e_{slightly_opened} = \neg \mathbf{1}_{\text{slightly_opened}(a),t-1} \wedge \mathbf{1}_{\text{slightly_opened}(a),t}$
 1205 $e_{closed} = \mathbf{1}_{\text{open}(a),t-1} \wedge \neg \mathbf{1}_{\text{open}(a),t}$
 1206 $e_{success} = \neg \mathbf{1}_{\text{success},t-1} \wedge \mathbf{1}_{\text{success},t}$
 1207 $e_{excessive_collisions} = C_{[0:t-1]} \leq 10000 \wedge C_{[0:t]} > 10000$

1208 For $t \in \{1, \dots, n\}$ (in increasing order), we evaluate each e_{event} in the order shown above.
 1209 If $e_{event} = 1$, we add it to E_{open} .

- 1210 • Success Modes: if $e_{success} \in E_{open}$, then categorize using the following success modes:
- 1211 i Open success: Agent successfully opens a and returns to rest without excessive collisions.
 1212 $e_{excessive_collisions} \notin E_{open} \wedge i_{open,opened} > i_{open,closed}$
- 1213 ii Dubious success: Agent successfully opens a and returns to rest without excessive collisions, but accidentally closes a after
 1214 $e_{excessive_collisions} \notin E_{open} \wedge i_{open,opened} < i_{open,closed}$
- 1215 iii Success then excessive collisions: Agent successfully opens a and returns to rest, but exceeds collision threshold after.
 1216 $e_{excessive_collisions} \notin E_{open}$
- 1217 • Failure Modes: if $e_{success} \notin E_{open}$, then categorize using the following failure modes:
- 1218 iv Excessive collision failure: Agent exceeds collision threshold.
 1219 $e_{excessive_collisions} \in E_{open}$
- 1220 v Can't reach articulation failure: Agent cannot reach a .
 1221 $e_{contact} \notin E_{open} \wedge e_{excessive_collisions} \notin E_{open}$
- 1222 vi Closed after open failure: Agent opens a , but closes it before returning to rest.
 1223 $e_{closed} \in E_{open} \wedge i_{open,closed} > i_{open,opened} \wedge i_{open,closed} > i_{open,slightly_opened} \wedge$
 1224 $e_{excessive_collisions} \notin E_{open}$
- 1225 vii Slightly opened failure: Agent at least slightly opens a , but cannot fully open it.
 1226 Previous failure modes are not applicable, and $i_{open,slightly_opened} > i_{open,opened} \wedge$
 1227 $i_{open,slightly_opened} > i_{open,closed} \wedge e_{excessive_collisions} \notin E_{open}$
- 1228 viii Too slow failure: Agent is able to open a , but cannot return to rest in time.
 1229 Previous failure modes are not applicable, and $e_{opened} \in E_{open}$
- 1230 ix Can't open failure: Agent reaches a , but cannot open it.
 1231 The condition is no other failure mode is applicable. This also implies $e_{contact} \in$
 1232 $E_{open} \wedge e_{opened} \notin E_{open}$

1233
 1234 **Close**[a](a_{pos}): Close articulation a with handle at a_{pos} .

- 1242 • Events: First, define indicator
 1243
 1244 $\mathbf{1}_{\text{slightly_closed}(a),t} = \mathbf{1}\{a_{q,t} < a_{q,0} - 0.05 \cdot (a_{q\max} - a_{q\min})\}$
 1245 Now, we define time-conditioned events at timestep t :
 1246
 1247 $e_{\text{contact}} = |F_{ee,a,t-1}| = 0 \wedge |F_{ee,a,t}| \geq 0$
 1248 $e_{\text{closed}} = \neg \mathbf{1}_{\text{closed}(a),t-1} \wedge \mathbf{1}_{\text{closed}(a),t}$
 1249 $e_{\text{slightly_closed}} = \neg \mathbf{1}_{\text{slightly_closed}(a),t-1} \wedge \mathbf{1}_{\text{slightly_closed}(a),t}$
 1250
 1251 $e_{\text{open}} = \mathbf{1}_{\text{closed}(a),t-1} \wedge \neg \mathbf{1}_{\text{closed}(a),t}$
 1252 $e_{\text{success}} = \neg \mathbf{1}_{\text{success},t-1} \wedge \mathbf{1}_{\text{success},t}$
 1253 $e_{\text{excessive_collisions}} = C_{[0:t-1]} \leq 10000 \wedge C_{[0:t]} > 10000$
 1254
 1255 For $t \in \{1, \dots, n\}$ (in increasing order), we evaluate each e_{event} in the order shown above.
 1256 If $e_{\text{event}} = 1$, we add it to E_{close} .
 1257 • Success Modes: if $e_{\text{success}} \in E_{\text{close}}$, then categorize using the following success modes:
 1258 i Close success: Agent successfully closes a and returns to rest without excessive collisions.
 1259 $e_{\text{excessive_collisions}} \notin E_{\text{close}} \wedge i_{\text{close,closed}} > i_{\text{close,opened}}$
 1260 ii Dubious success: Agent successfully closes a and returns to rest without excessive collisions, but accidentally opens a after.
 1261 $e_{\text{excessive_collisions}} \notin E_{\text{close}} \wedge i_{\text{close,closed}} < i_{\text{close,opened}}$
 1262 iii Success then excessive collisions: Agent successfully closes a and returns to rest, but exceeds collision threshold after.
 1263 $e_{\text{excessive_collisions}} \notin E_{\text{close}}$
 1264 • Failure Modes: if $e_{\text{success}} \notin E_{\text{close}}$, then categorize using the following failure modes:
 1265 iv Excessive collision failure: Agent exceeds collision threshold.
 1266 $e_{\text{excessive_collisions}} \in E_{\text{close}}$
 1267 v Can't reach articulation failure: Agent cannot reach a .
 1268 $e_{\text{contact}} \notin E_{\text{close}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{close}}$
 1269 vi Opened after closed failure: Agent closes a , but opens it before returning to rest.
 1270 $e_{\text{closed}} \in E_{\text{close}} \wedge i_{\text{close,opened}} > i_{\text{close,closed}} \wedge i_{\text{close,opened}} > i_{\text{close,slightly_closed}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{close}}$
 1271 vii Slightly closed failure: Agent at least slightly closes a , but cannot fully close it.
 1272 Previous failure modes are not applicable, and $i_{\text{close,slightly_closed}} > i_{\text{close,closed}} \wedge i_{\text{close,slightly_closed}} > i_{\text{close,opened}} \wedge e_{\text{excessive_collisions}} \notin E_{\text{close}}$
 1273 viii Too slow failure: Agent is able to close a , but cannot return to rest in time.
 1274 Previous failure modes are not applicable, and $e_{\text{closed}} \in E_{\text{close}}$
 1275 ix Can't close failure: Agent reaches a , but cannot close it.
 1276 The condition is no other failure mode is applicable. This also implies $e_{\text{contact}} \in E_{\text{close}} \wedge e_{\text{closed}} \notin E_{\text{close}}$

A.5.3 TRAJECTORY CATEGORIZATION STATISTICS

In Tables 5-12, we categorize trajectories with our automated event labeling method. We run 1000 episodes for every task/subtask/target combination using all relevant policies (RL-Per, RL-All, IL) for each object/articulation. We provide success once rate (**SoR**), success at end rate (**SaeR**), failure rate (**FR**), and proportions for each success and failure mode as percentages (labeled with roman numerals corresponding to those use in Appendix A.5.2).

Table 5: Train split Pick policy trajectory labeling on 1000 episodes per target object. All numbers are percentages. Best viewed zoomed.

TASK	OBJ	TYPE	SoR	SaeR	(i)	(ii)	(iii)	(iv)	FR	(v)	(vi)	(vii)	(viii)	(ix)
TH	002	RL-Per	82.34	72.52	70.63	1.88	0.00	9.82	17.66	13.79	3.37	0.40	0.10	0.00
		RL-All	78.97	60.81	58.43	2.38	0.00	18.15	21.03	13.69	5.65	0.89	0.50	0.30
		IL	72.62	70.63	67.76	2.88	0.20	1.79	27.38	21.03	0.69	1.79	1.29	2.58
	003	RL-Per	71.63	62.80	60.91	1.88	0.10	8.73	28.37	17.26	7.64	2.48	0.89	0.10
		RL-All	29.46	23.41	23.02	0.40	0.00	6.05	70.54	34.52	6.25	28.17	1.19	0.40
		IL	17.96	17.16	16.87	0.30	0.00	0.79	82.04	49.01	0.20	31.25	0.89	0.69
	004	RL-Per	78.47	68.15	65.48	2.68	0.00	10.32	21.53	13.29	6.25	1.19	0.60	0.20
		RL-All	74.60	57.64	55.16	2.48	0.00	16.96	25.40	16.27	6.25	1.88	0.40	0.60
		IL	60.81	57.44	54.27	3.17	0.00	3.37	39.19	27.18	0.50	6.05	1.59	3.87
	005	RL-Per	81.05	78.27	75.10	3.17	0.00	2.78	18.95	15.67	2.58	0.40	0.10	0.20
		RL-All	76.69	62.30	59.13	3.17	0.20	14.19	23.31	17.06	5.16	0.30	0.40	0.40
		IL	74.01	70.73	66.37	4.37	0.10	3.17	25.99	21.13	0.40	0.79	0.89	2.78
PG	007	RL-Per	83.23	80.36	78.57	1.79	0.20	2.68	16.77	10.52	5.95	0.20	0.10	0.00
		RL-All	77.98	70.44	69.44	0.99	0.00	7.54	22.02	15.97	5.56	0.10	0.10	0.30
		IL	76.59	75.20	72.42	2.78	0.10	1.29	23.41	18.95	1.59	1.69	0.20	0.99
	008	RL-Per	75.99	72.22	63.19	9.03	0.50	3.27	24.01	17.36	3.08	2.18	0.79	0.60
		RL-All	73.91	62.80	55.95	6.85	0.60	10.52	26.09	18.25	5.36	1.69	0.69	0.10
		IL	65.18	62.50	54.07	8.43	0.50	2.18	34.82	26.39	0.40	3.77	2.18	2.08
	009	RL-Per	81.25	76.69	71.63	5.06	0.40	4.17	18.75	12.40	5.46	0.50	0.30	0.10
		RL-All	71.43	60.42	50.10	10.32	0.40	10.62	28.57	20.54	4.86	2.28	0.60	0.30
		IL	67.16	64.29	55.85	8.43	0.79	2.08	32.84	25.00	0.79	3.08	2.38	1.59
	010	RL-Per	81.65	52.78	48.61	4.17	0.10	28.77	18.35	12.40	4.86	0.79	0.20	0.10
		RL-All	75.50	61.90	55.16	6.75	0.00	13.59	24.50	18.95	4.27	0.99	0.20	0.10
		IL	58.33	54.17	48.02	6.15	0.50	3.67	41.67	26.98	0.79	8.43	1.19	4.27
	024	RL-Per	82.74	78.47	68.35	10.12	0.20	4.07	17.26	12.90	3.67	0.30	0.20	0.20
		RL-All	73.91	65.08	56.15	8.93	0.00	8.83	26.09	19.05	6.15	0.69	0.10	0.10
		IL	62.20	58.63	50.40	8.23	0.69	2.88	37.80	25.20	0.20	7.74	2.18	2.48
	all	RL-Per	81.75	73.41	67.26	6.15	0.20	8.13	18.25	12.40	4.17	1.19	0.30	0.20
	all	RL-All	71.63	59.13	54.17	4.96	0.30	12.20	28.37	19.74	4.96	3.17	0.30	0.20
	all	IL	61.11	59.42	54.56	4.86	0.20	1.49	38.89	26.39	0.69	7.64	0.89	3.27
ST	002	RL-Per	69.05	55.16	50.89	4.27	0.40	13.49	30.95	14.58	16.17	0.10	0.00	0.10
		RL-All	62.70	49.40	38.10	11.31	1.49	11.81	37.30	24.60	11.31	0.89	0.50	0.00
		IL	63.10	59.82	55.16	4.66	0.20	3.08	36.90	31.55	2.38	1.09	0.99	0.89
	003	RL-Per	51.98	43.15	40.67	2.48	2.08	6.75	48.02	25.10	12.30	8.63	1.79	0.20
		RL-All	11.51	8.13	7.64	0.50	0.60	2.78	88.49	60.62	11.21	16.17	0.20	0.30
		IL	16.27	14.38	13.79	0.60	1.29	0.60	83.73	67.06	2.28	12.10	1.98	0.30
	004	RL-Per	64.48	52.88	50.20	2.68	0.20	11.41	35.52	19.44	13.99	0.69	0.10	1.29
		RL-All	59.82	46.03	37.10	8.93	1.39	12.40	40.18	26.98	11.41	0.99	0.40	0.40
		IL	51.09	47.32	44.84	2.48	0.50	3.27	48.91	39.98	2.38	2.88	1.98	1.69
	005	RL-Per	61.21	48.51	44.74	3.77	0.30	12.40	38.79	25.79	11.51	0.10	0.20	1.19
		RL-All	63.29	51.29	39.58	11.71	0.69	11.31	36.71	25.89	10.42	0.10	0.30	0.00
		IL	61.01	56.25	54.07	2.18	0.20	4.56	38.99	34.33	2.28	0.50	0.89	0.99
	007	RL-Per	66.57	52.58	47.02	5.56	0.10	13.89	33.43	18.55	13.29	1.19	0.30	0.10
		RL-All	64.38	45.63	34.42	11.21	0.89	17.86	35.62	23.71	10.81	0.99	0.10	0.00
		IL	61.61	59.82	57.04	2.78	0.00	1.79	38.39	32.14	3.08	1.59	0.89	0.69
	008	RL-Per	62.90	45.73	37.00	8.73	0.89	16.27	37.10	21.33	14.38	0.79	0.50	0.10
		RL-All	45.93	34.03	21.63	12.40	2.58	9.33	54.07	36.81	12.40	3.17	1.49	0.20
		IL	40.38	35.71	32.24	3.47	1.19	3.47	59.62	49.80	2.78	3.97	2.58	0.50
	009	RL-Per	63.59	46.73	41.67	5.06	1.09	15.77	36.41	18.75	16.47	0.30	0.50	0.40
		RL-All	49.01	38.49	26.19	12.30	1.69	8.83	50.99	32.14	13.00	3.67	1.69	0.50
		IL	43.25	39.78	36.01	3.77	1.19	2.28	56.75	45.73	3.77	4.56	2.18	0.50
	010	RL-Per	65.18	51.49	44.64	6.85	0.30	13.39	34.82	20.93	13.29	0.20	0.30	0.10
		RL-All	54.76	42.06	27.58	14.48	1.49	11.21	45.24	30.26	12.10	2.18	0.30	0.40
		IL	50.99	48.61	41.77	6.85	0.10	2.28	49.01	41.27	2.68	2.38	1.69	0.99
	024	RL-Per	74.60	56.45	39.58	16.87	0.30	17.86	25.40	15.87	8.73	0.30	0.20	0.30
		RL-All	51.09	35.02	22.92	12.10	0.79	15.28	48.91	37.00	10.12	1.09	0.60	0.10
		IL	20.73	18.25	14.68	3.57	0.60	1.88	79.27	66.37	1.88	7.54	1.49	1.98
	all	RL-Per	66.57	52.78	46.63	6.15	0.40	13.39	33.43	19.54	11.81	1.09	0.60	0.40
	all	RL-All	51.88	37.70	28.17	9.52	1.79	12.40	48.12	33.93	10.12	2.48	1.29	0.30
	all	IL	44.64	42.36	39.38	2.98	0.40	1.88	55.36	47.22	1.49	4.27	1.49	0.89
1343	013	RL-Per	65.87	54.07	47.82	6.25	0.30	11.51	34.13	11.11	21.73	0.00	0.40	0.89
		RL-All	59.03	35.71	33.04	2.68	0.10	23.21	40.97	16.57	24.01	0.10	0.00	0.30
	024	RL-Per	94.35	85.81	75.30	10.52	0.20	8.33	5.65	4.86	0.50	0.10	0.20	0.00
		RL-All	93.95	79.37	64.68	14.68	0.20	14.38	6.05	5.06	0.30	0.00	0.50	0.20
	all	RL-Per	80.85	69.64	60.62	9.03	0.20	11.01	19.15	8.04	10.12	0.00	0.30	0.69
	all	RL-All	75.69	56.55	47.52	9.03	0.10	19.05	24.31	11.01	12.40	0.30	0.50	0.10
	all	IL	60.71	50.40	44.74	5.65	1.09	9.23	39.29	31.25	1.88	3.87	1.29	0.99

Table 6: Val split Pick policy trajectory labeling on 1000 episodes per target object. All numbers are percentages. Best viewed zoomed.

	TASK	OBJ	TYPE	SoR	SaeR	(i)	(ii)	(iii)	(iv)	FR	(v)	(vi)	(vii)	(viii)	(ix)
TH	002	RL-Per	80.06	70.63	68.95	1.69	0.30	9.13	19.94	15.77	3.67	0.30	0.00	0.20	
		RL-All	78.57	61.71	59.82	1.88	0.00	16.87	21.43	15.18	5.36	0.69	0.10	0.10	
	003	IL	73.61	72.22	69.64	2.58	0.20	1.19	26.39	21.03	0.20	1.69	1.09	2.38	
		RL-Per	73.41	64.38	61.61	2.78	0.10	8.93	26.59	16.67	6.94	1.98	0.99	0.00	
	004	RL-All	33.73	26.29	25.79	0.50	0.10	7.34	66.27	33.13	7.34	24.50	1.09	0.20	
		IL	16.67	16.17	15.58	0.60	0.10	0.40	83.33	47.52	0.10	33.93	0.69	1.09	
	005	RL-Per	77.28	69.44	66.96	2.48	0.00	7.84	22.72	13.39	6.85	1.29	0.79	0.40	
		RL-All	75.20	57.44	54.07	3.37	0.00	17.76	24.80	15.97	6.35	1.88	0.20	0.40	
	007	IL	59.62	57.34	54.76	2.58	0.20	2.08	40.38	27.78	0.40	7.24	1.19	3.77	
		RL-Per	81.15	78.08	75.79	2.28	0.10	2.98	18.85	15.87	2.78	0.20	0.00	0.00	
	008	RL-All	75.50	61.90	58.43	3.47	0.10	13.49	24.50	18.55	5.06	0.20	0.40	0.30	
		IL	71.03	69.44	65.77	3.67	0.20	1.39	28.97	21.92	0.99	2.28	0.79	2.98	
	009	RL-Per	82.74	77.98	76.69	1.29	0.00	4.76	17.26	11.61	5.46	0.20	0.00	0.00	
		RL-All	77.78	69.44	68.75	0.69	0.00	8.33	22.22	15.18	6.55	0.40	0.00	0.10	
	010	IL	72.72	71.23	67.66	3.57	0.10	1.39	27.28	22.22	1.29	1.79	0.10	1.88	
		RL-Per	75.40	70.34	60.62	9.72	0.69	4.37	24.60	18.45	2.58	2.48	0.79	0.30	
	024	RL-All	72.02	63.19	55.95	7.24	0.40	8.43	27.98	19.54	5.26	2.38	0.79	0.00	
		IL	61.61	60.32	54.76	5.56	0.50	0.79	38.39	29.37	0.89	3.87	2.58	1.69	
	all	RL-Per	78.57	74.70	69.15	5.56	0.30	3.57	21.43	14.98	5.26	0.89	0.20	0.10	
		RL-All	71.63	60.62	52.78	7.84	0.30	10.71	28.37	20.34	6.35	0.60	0.99	0.10	
	PG	IL	64.48	62.20	54.96	7.24	1.09	1.19	35.52	28.57	1.09	2.78	1.88	1.19	
		RL-Per	82.24	53.37	50.00	3.37	0.00	28.87	17.76	13.19	4.07	0.50	0.00	0.00	
		RL-All	74.90	63.99	56.94	7.04	0.10	10.81	25.10	19.25	4.46	0.79	0.30	0.30	
		IL	55.75	52.68	46.83	5.85	0.10	2.98	44.25	27.88	0.30	9.42	1.69	4.96	
		RL-Per	79.56	74.31	63.79	10.52	0.50	4.76	20.44	15.58	4.46	0.20	0.20	0.00	
		RL-All	69.74	62.90	53.37	9.52	0.00	6.85	30.26	22.32	7.14	0.50	0.20	0.10	
		IL	58.04	54.76	44.94	9.82	0.40	2.88	41.96	29.96	0.30	7.04	1.39	3.27	
		RL-Per	77.48	69.44	65.77	3.67	0.30	7.74	22.52	16.57	4.86	0.89	0.00	0.20	
		RL-All	68.15	57.34	53.97	3.37	0.00	10.81	31.85	21.73	5.36	4.27	0.10	0.40	
		IL	59.03	57.04	53.87	3.17	0.20	1.79	40.97	28.47	0.30	8.43	1.39	2.38	
	002	RL-Per	84.62	67.56	63.69	3.87	0.89	16.17	15.38	12.10	2.88	0.10	0.10	0.20	
		RL-All	76.98	54.27	41.47	12.80	2.68	20.04	23.02	20.34	0.99	1.29	0.30	0.10	
	003	IL	74.11	68.45	64.58	3.87	1.19	4.46	25.89	20.93	0.50	1.29	2.18	0.99	
		RL-Per	56.15	42.56	39.68	2.88	2.38	11.21	43.85	30.46	2.08	9.72	1.39	0.20	
	004	RL-All	14.19	11.41	10.02	1.39	0.60	2.18	85.81	57.24	0.69	26.88	0.99	0.00	
		IL	18.25	16.77	16.37	0.40	0.69	0.79	81.75	61.90	0.20	16.67	2.58	0.40	
	005	RL-Per	69.74	60.52	57.24	3.27	0.30	8.93	30.26	25.00	2.08	1.29	0.40	1.49	
		RL-All	71.23	47.82	38.99	8.83	4.17	19.25	28.77	23.61	1.29	2.78	0.89	0.20	
	007	IL	58.43	53.57	51.19	2.38	0.40	4.46	41.57	34.33	0.30	4.17	1.49	1.29	
		RL-Per	76.69	64.88	60.12	4.76	0.20	11.61	23.31	21.43	0.69	0.60	0.30	0.30	
	008	RL-All	75.60	53.97	40.87	13.10	1.19	20.44	24.40	21.03	1.09	1.79	0.30	0.20	
		IL	75.00	68.35	65.67	2.68	0.30	6.35	25.00	20.54	0.79	1.19	1.29	1.19	
	009	RL-Per	76.98	56.25	50.69	5.56	0.30	20.44	23.02	19.35	2.58	0.60	0.40	0.10	
		RL-All	76.98	56.15	41.96	14.19	0.40	20.44	23.02	19.64	1.59	1.79	0.00	0.00	
	010	IL	72.02	69.84	65.48	4.37	0.10	2.08	27.98	21.83	1.69	2.78	0.79	0.89	
		RL-Per	73.21	48.91	37.60	11.31	0.69	23.61	26.79	22.52	1.59	1.98	0.30	0.40	
	013	RL-All	58.93	42.86	24.70	18.15	3.67	12.40	41.07	29.07	2.38	6.65	2.98	0.00	
		IL	44.35	37.40	32.04	5.36	1.88	5.06	55.65	44.15	1.59	5.65	3.47	0.79	
	010	RL-Per	70.73	38.00	32.24	5.75	1.88	30.85	29.27	26.19	1.69	0.99	0.20	0.20	
		RL-All	62.90	46.92	29.37	17.56	2.88	13.10	37.10	26.09	2.38	5.36	2.98	0.30	
	024	IL	44.54	38.79	33.73	5.06	1.79	3.97	55.46	42.56	1.19	7.64	3.67	0.40	
		RL-Per	74.50	55.26	48.61	6.65	0.60	18.65	25.50	23.91	0.50	0.79	0.30	0.00	
	024	RL-All	70.04	49.40	34.42	14.98	2.88	17.76	29.96	23.41	2.38	3.57	0.60	0.00	
		IL	59.23	57.04	52.18	4.86	0.50	1.69	40.77	33.93	0.69	2.98	2.18	0.99	
	all	RL-Per	78.17	55.65	39.09	16.57	0.10	22.42	21.83	15.08	5.75	0.50	0.10	0.40	
		RL-All	61.21	37.00	25.79	11.21	1.19	23.02	38.79	35.81	0.69	1.49	0.79	0.00	
	024	IL	22.82	20.14	16.27	3.87	0.30	2.38	77.18	65.28	0.79	7.64	0.89	2.58	
		RL-Per	72.32	53.47	45.93	7.54	0.50	18.35	27.68	23.02	1.39	2.38	0.79	0.10	
	all	RL-All	62.10	43.75	31.25	12.50	1.79	16.57	37.90	28.47	0.89	7.54	0.89	0.10	
		IL	52.78	47.52	43.75	3.77	0.50	4.76	47.22	36.90	0.69	6.55	1.79	1.29	
ST	013	RL-Per	80.36	58.43	48.02	10.42	0.20	21.73	19.64	17.96	0.30	0.00	0.20	1.19	
		RL-All	66.67	26.29	19.74	6.55	0.10	40.28	33.33	31.94	0.89	0.30	0.20	0.00	
	024	IL	78.08	55.56	51.09	4.46	1.49	21.03	21.92	18.06	1.49	0.30	0.69	1.39	
		RL-Per	93.65	83.33	71.83	11.51	0.20	10.12	6.35	5.85	0.10	0.20	0.20	0.00	
	all	RL-All	89.98	77.38	62.00	15.38	0.40	12.20	10.02	9.33	0.30	0.20	0.20	0.00	
		IL	62.20	54.86	48.61	6.25	0.20	7.14	37.80	26.59	2.28	5.95	1.39	1.59	

1512
 1513
 1514
 1515
 1516 Table 9: Train split Open policy trajectory labeling on 1000 episodes per target articulation. All
 1517 numbers are percentages. Best viewed zoomed.
 1518

TASK	OBJ	TYPE	SoR	SaeR	(i)	(ii)	(iii)	FR	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
ST	drawer	RL-Per	84.92	84.92	84.92	0.00	0.00	15.08	11.41	1.69	0.00	1.49	0.20	0.30
		IL	79.86	79.86	79.86	0.00	0.00	20.14	13.59	0.89	1.09	3.67	0.10	0.79
ST	fridge	RL-Per	83.43	82.24	82.24	0.00	1.19	16.57	14.58	0.20	0.00	1.49	0.20	0.10
		IL	74.01	68.85	68.85	0.00	5.16	25.99	22.12	0.99	0.10	1.69	0.40	0.69

1523
 1524
 1525
 1526
 1527
 1528
 1529 Table 10: Val split Open policy trajectory labeling on 1000 episodes per target articulation. All
 1530 numbers are percentages. Best viewed zoomed.
 1531

TASK	OBJ	TYPE	SoR	SaeR	(i)	(ii)	(iii)	FR	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
ST	drawer	RL-Per	84.52	84.52	84.52	0.00	0.00	15.48	11.41	1.88	0.00	1.88	0.00	0.30
		IL	78.57	78.37	78.37	0.00	0.20	21.43	13.69	1.29	0.20	4.17	1.19	0.89
ST	fridge	RL-Per	88.10	37.30	37.30	0.00	50.79	11.90	6.94	0.00	0.00	4.07	0.89	0.00
		IL	53.67	1.49	1.49	0.00	52.18	46.33	43.85	0.00	0.00	1.29	0.89	0.30

1536
 1537
 1538
 1539
 1540
 1541
 1542 Table 11: Train split Close policy trajectory labeling on 1000 episodes per target articulation. All
 1543 numbers are percentages. Best viewed zoomed.
 1544

TASK	OBJ	TYPE	SoR	SaeR	(i)	(ii)	(iii)	FR	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
ST	drawer	RL-Per	88.79	88.79	88.79	0.00	0.00	11.21	3.17	2.68	0.30	0.89	4.17	0.00
		IL	88.39	88.39	88.39	0.00	0.00	11.61	3.27	3.08	0.00	0.79	4.46	0.00
ST	fridge	RL-Per	86.81	25.00	25.00	0.00	61.81	13.19	13.19	0.00	0.00	0.00	0.00	0.00
		IL	86.90	29.96	29.96	0.00	56.94	13.10	12.90	0.00	0.00	0.20	0.00	0.00

1550
 1551
 1552
 1553
 1554
 1555
 1556 Table 12: Val split Close policy trajectory labeling on 1000 episodes per target articulation. All
 1557 numbers are percentages. Best viewed zoomed.
 1558

TASK	OBJ	TYPE	SoR	SaeR	(i)	(ii)	(iii)	FR	(iv)	(v)	(vi)	(vii)	(viii)	(ix)
ST	drawer	RL-Per	89.29	89.29	89.29	0.00	0.00	10.71	4.56	2.18	0.20	0.20	3.37	0.20
		IL	87.60	87.60	87.60	0.00	0.00	12.40	4.66	2.18	0.00	0.30	5.16	0.10
ST	fridge	RL-Per	0.00	0.00	0.00	0.00	0.00	100.00	81.15	18.85	0.00	0.00	0.00	0.00
		IL	0.00	0.00	0.00	0.00	0.00	100.00	95.93	4.07	0.00	0.00	0.00	0.00