Predicting Object Interactions with Behavior Primitives: An Application in Stowing Tasks

Anonymous Author(s)
Affiliation
Address
email

Abstract: Stowing, the task of placing objects in cluttered shelves or bins, is a common task in warehouse and manufacturing operations. However, this task is still predominantly carried out by human workers as stowing is challenging to automate due to the complex multi-object interactions and long-horizon nature of the task. Previous works typically involve extensive data collection and costly human labeling of semantic priors across diverse object categories. This paper presents a method to learn a generalizable robot stowing policy from predictive model of object interactions and a single demonstration with behavior primitives. We propose a novel framework that utilizes Graph Neural Networks (GNNs) to predict object interactions within the parameter space of behavioral primitives. We further employ primitive-augmented trajectory optimization to search the parameters of a predefined library of heterogeneous behavioral primitives to instantiate the control action. Our framework enables robots to proficiently execute long-horizon stowing tasks with a few keyframes (3-4) from a single demonstration. Despite being solely trained in a simulation, our framework demonstrates remarkable generalization capabilities. It efficiently adapts to a broad spectrum of real-world conditions, including various shelf widths, fluctuating quantities of objects, and objects with diverse attributes such as sizes and shapes.

Keywords: Robotic Manipulation, Model Learning, Graph-Based Neural Dynamics, Multi-Object Interactions

1 Introduction

Stowing, defined as relocating an object from a table to a cluttered shelf, is one of the dominating warehouse activities. In stowing, an agent is required to pick up an object from a table. The agent must then actively create free space within the shelf before inserting the object from the table. A successful stow execution is characterized by the placement of all objects with poses in some predefined thresholds. While stowing can be performed effortlessly by humans, it remains challenging when automated with robots. The difficulty stems from the long-horizon nature, multi-object interactions, and the variety of objects and configurations involved in stowing tasks.

The challenge of long-horizon stowing task is not only due to the nature and variety of objects involved but also due to several inherent constraints in existing methods. First, the nature of these tasks requires determining the characteristics of contacts and frictions, which is a task that presents considerable difficulties. Conventional first-order models fall short in capturing the physical effects, and identifying the specific parameters of contact and friction is challenging [1]. Thus, designing a controller for such tasks becomes a tedious and laborious process. Second, the existing methodologies, including manually pre-programmed systems and recent advancements in category-level object manipulation, exhibit notable limitations. Classical pre-programmed systems struggle with adaptability, unable to efficiently handle variations introduced by different arrangements of objects on the shelf. Meanwhile, recent strategies for category-level object manipulation are curbed by the need

Submitted to the 7th Conference on Robot Learning (CoRL 2023). Do not distribute.

for expensive data collection and human labelling, thus failing to provide a scalable solution. Additionally, pure learning-based methods, such as Deep Reinforcement Learning (DRL), also present drawbacks in terms of extensive training times and poor data efficiency [2, 3], making learning from scratch on real robots impractical for long-horizon tasks.

To address these challenges, we propose a novel framework that uses Graph Neural Networks (GNNs) to predict object interactions within the parameter space of behavior primitives. When trained with various situations in the simulator, GNNs can learn to model the forward dynamics associated with interactions of rigid objects. Our framework also applies primitive-augmented trajectory optimization to search the parameters of a predefined library of heterogeneous skills or behavior primitives. The incorporation of behavior primitives enables our policy to handle tasks of significant complexity with improved efficiency.

We make three key contributions: (1) We introduce a novel model-based imitation learning framework to learn from minimal demonstrations, which enables robots to acquire complex skills. (2) We create a comprehensive stowing benchmark for long-horizon manipulations, which is highly prevalent in various both the industrial and household applications. (3) We demonstrate the effectiveness and generalization of our framework across a wide range of real-world experiments in stowing tasks.

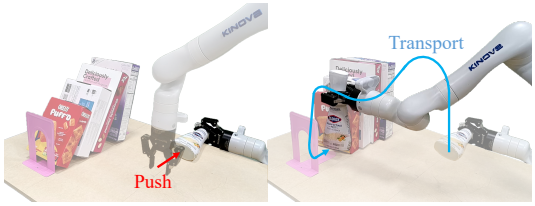

Figure 1: The robot places an object into a cluttered shelf by exploiting object interactions. It uses the grasped object to manipulate other objects within the shelf through pushing and sliding actions and finally places all the objects in the appropriate location.

2 Related Works

One-shot Imitation Learning in Manipulation: Recent advancements in imitation learning aim to solve unseen new tasks with as few demonstrations as possible [4, 5, 6, 7, 8]. Typical approaches use a two-phase pipeline consisting of a meta-learning phase to train a general policy from a large number of tasks, followed by a fine-tuning phase to optimize the policy for a specific task using a single demonstration [8]. Certain replay-based methods employ a strategy of estimating a 'bottleneck pose', splitting the trajectories into two segments, and then replaying the demonstration trajectory from the bottleneck [5]. Other techniques emphasize learning object representations and identifying correspondence between specific objects [4, 9] or objects within the same category [7]. However, these methodologies primarily handle relatively short-horizon tasks and face difficulties in modeling object dynamics.

Model Learning in Robotic Manipulation: Dynamics models have emerged as an important component in robotic systems, capable of predicting future world states based on current states and actions [10, 11, 12, 13, 14, 15]. The model can process different representations of inputs, ranging from images and latent vectors, to point sets and graph representations. Among these, graph representations have demonstrated a remarkable ability to capture interactions between multiple objects [16, 17, 18, 19, 20, 21, 22]. The capability to model interactions between various objects and the robot end effector is crucial for our work, leading us to use a graph to model interactions between objects in our system.

Long-Horizon Planning in Robotic Manipulation: Addressing long-horizon manipulation problems presents substantial complexity. Hierarchical reinforcement learning (HRL) approaches aim to solve this issue by employing a high-level policy to select optimal subtasks and low-level policies to execute them [23]. However, these methods face the challenges of the sim2real gap and the difficulty of real-world data collection, hampering their real-world transferability. An alternative approach by Di Palo et al. [24] enhances replay-based methodologies by integrating primitive skills for everyday long-horizon tasks, but the task configuration lacks versatility. Recent efforts by Lin et al. [25] have explored sequential manipulation of deformable objects through a differential simulator and trajectory optimization. However, this work is only validated in simulation, and real-world deployment

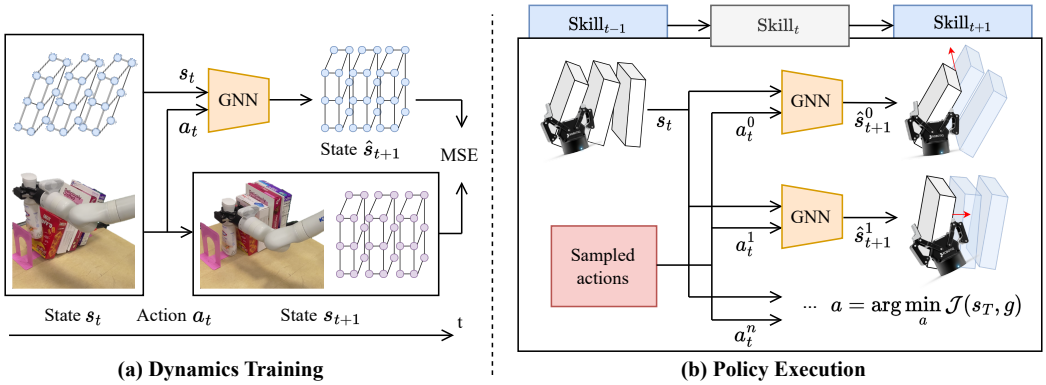

Figure 2: **Overview of the proposed framework.** (a) We employ a particle-based representation to characterize the object state. The object state's predicted outcome, following the execution of robot actions, is computed alongside the ground truth object state using the MSE loss function to train the GNN. (b) For each skill execution, we apply random shooting to sample parameters within the action parameter space, utilizing the GNN to predict object movement. We then select the action that brings us closest to the desired state. Each skill is executed in sequence.

is non-trivial due to the difficulty of obtaining gradients of state changes. To this end, we propose to apply trajectory optimization to GNN-based forward dynamics prediction modules, incorporating heterogeneous behavior primitives.

3 Approach

In our proposed system, a GNN first predicts system dynamics. Then, a primitive-augmented trajectory optimization method achieves subgoals from a single demonstration, which is shown in Figure 2. Initially, the object state is represented as particles. We then train the GNN with the MSE loss between the predicted outcome following robot actions and the ground truth state. We use the random shooting to explore the action parameter space and use the GNN's predicted results to select the optimal action that aligns most closely with our desired state. The skills are executed in a sequential manner, guiding our system to accomplish its tasks efficiently.

3.1 Learning Forward Dynamics via GNN

Graph Construction: We define a dynamics function that describes a rigid-body system state $\mathcal{S}$ with M objects and N particles. We model each rigid object as a collection of uniformly distributed particles, offering a representation that is flexible and adaptable to variations in object shape, size, and geometry. The dynamics function is expressed as $\Phi : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{T}$. $\mathcal{A}$ represents the skill type and its associated parameters, and $\mathcal{T}$ denotes M rigid transformations containing the translation and rotation for each dynamic object. The future state of an object can be determined by applying a sequence of these rigid transformations. We represent each rigid objects as uniformly distributed particles as such representations are versatile to object shape, object size, and object geometry.

Each object in the system is represented by its respective particles. We define the graph state $s_t = (\mathcal{O}_t, \mathcal{E}_t)$, where the graph's vertices $\mathcal{O}_t$ represent an object's particles, and edges $\mathcal{E}_t$ denotes the relations between particles. Each vertex $o_{i,t} = \langle x_{i,t}, c_{i,t} \rangle$ consists of the particle's position $x_{i,t}$ and object's attributes $c_{i,t}$ including its dynamism (dynamic or static), the object it belongs to, its offset from the rigid body's center, and gravity. Edges $\mathcal{E}_t$ are formed when the distance between two vertices is less than a specified threshold. The edge relations are represented as $e_k = \langle i_k, j_k, c_k \rangle$, in which i_k, j_k denote the indices of the receiver and sender particles, respectively. The edge index is denoted by k, and nature of the relationship such as intra-object relation, or gripper-to-object relation) is denoted by c_k. Since our focus is predicting the motions of dynamic objects, we restrict node building to vertices associated with these dynamic objects. However, to effectively model the

interactions between dynamic and static objects (i.e, shelf and table), we choose to incorporate the particles of static objects during the construction of edges.

Message Passing: The features of all vertices and edges are processed through encoder multi-layer perceptrons (MLPs), resulting in latent vertex representations denoted as h_i^O and latent edge representations represented as $h_{i,j}^E$. At each time step, we apply the following message function to the vertex feature and edge feature for several iterations to handle signal propagation.

$$h_{i,j}^E \leftarrow \rho^E(h_{i,j}^E, h_i^O, h_j^O), \quad h_i^O \leftarrow \rho^O(h_i^O, \sum_j h_{i,j}^E). \tag{1}$$

where the message passing functions ρ^E and ρ^O are MLPs. Subsequently, an MLP decoder is applied to the output from the final layer of the vertex feature processor. The rigid transformation for each individual object is determined by calculating the mean of the decoder outputs corresponding to that object.

Representing Action as Particles: The control action is represented through the planned position and motion of the gripper, further described by particles as $o_{i,t} = \langle x_{i,t}, v_{i,t}, c_{i,t} \rangle$, where $x_{i,t}$ denotes the current position of gripper, $v_{i,t}$ denotes the planned motion of the gripper. The particles associated with the gripper are subsequently encoded by the encoder. Additionally, we predict the future positions of the gripper. The discrepancy between these predicted positions and the actual achieved positions serves as an auxiliary loss, which helps GNNs better understand the inherent physical laws and constraints.

3.2 Control with the Learned Dynamics

In this section, we discuss the design of behavior primitives, and trajectory optimization algorithms used to generate parameters for different skills.

Behavior Primitives: We introduce the behavior primitives as a higher level of temporal abstraction to facilitate efficient planning. The behavior primitives simplify the task space by generating key poses for the system, which subsequently executes actions using Operational Space Control (OSC) [26] as a lower-level controller. The GNN is used only to predict the system's state at the key poses of behavior primitives, which significantly reduces the number of forward passes and cumulative prediction error over time. Behavior primitives function serve as building blocks and can be easily extended for various manipulation tasks. We further specify a maximum execution time T_{skl} for each behavior primitive. Our system integrates a collection of three primitives, encompassing both prehensile and non-prehensile motions. The primitives and their associated parameters are described as follows:

1. **Sweeping:** The robot sweeps objects on the shelf using its end-effector, aiming to stand them upright. Sweeping is parameterized by the starting offset y in the shelf direction, sweeping height h, sliding distance d, and the angle of gripper rotation θ during the sweep.
2. **Pushing:** Pushing involves the robot nudging the object to establish potential grasp poses. The starting push position (x, y) and the distance of the push d are the parameters for pushing.
3. **Transporting:** The robot picks up the object from the table, places it in the shelf, and, if necessary, adjusts its position through sliding. This skill is defined by parameters such as the starting offset in the shelf direction y, the height of insertion h, the sliding distance d, and the gripper rotation angle θ during the insertion process.

Goal-conditioned Trajectory Optimization: We optimize the parameters of a given skill by minimizing the mean square error (MSE) loss between the predicted particle positions and their corresponding positions as demonstrated. Keyframes collected during demonstrations can be denoted by g. We search for the skill parameter a_p that minimizes the cost function $\mathcal{J}$ representing the MSE between the predicted and target positions of the object particles. Mathematically, this optimization problem is represented as $a_p = \arg\min_{a_p} \mathcal{J}(s_T, g)$. The resulting low-level control actions are

generated from the the skill which is parameterized by the skill parameters a_p. Our dynamics network Φ makes forward predictions of the future state of the system after the execution of each skill. We employ trajectory optimization to find the skill parameters that yield the lowest cost.

4 Experiment Setup

Our experimental setup consists of both simulated and real-world environments. The simulated environment is built using Robosuite [27] and the MuJoCo physics engine [28], operating with a 7-DOF (Degrees of Freedom) Kinova Gen3 Robot. The real-world counterpart consists the same Kinova Gen3 Robot and a Robotiq 2F-85 gripper.

Data collection: We collect a dataset of 300 episodes of action execution for each behavior primitive, where the actions are executed based on randomly sampled skill parameters. For each episode, the robot randomly sampled parameters within its parameter space and executed the skill in the simulator. We gather the key poses for each skill, subsequently training a GNN to take the state at these key poses and the robot action and predict the future state at the subsequent key poses.

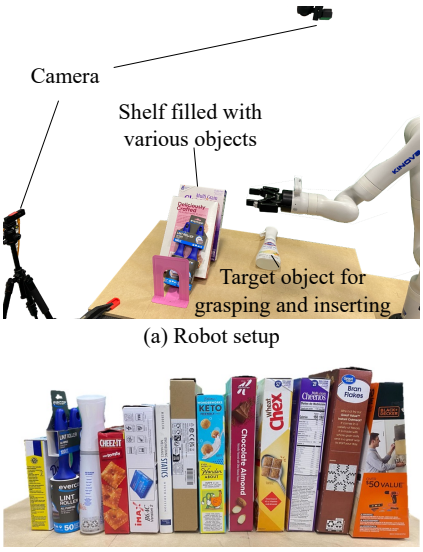

(a) Robot setup

(b) Objects used in this work

Figure 3: Overview of the robot setup.

Simulated environment: In our simulation, the shelf width is initialized to randomly vary within the range from 0.18 to 0.35 meters. We also randomly select the number of objects placed on the shelf, varying between two and four. The properties of each object, including size and density, were randomly generated. All the objects are created with a box shape. We use a SpaceMouse to teleoperate the robot to complete the task, and the ending poses of each skill were collected. These poses are then used as subgoals for each skill during execution.

Real-world environment: Figure 3 illustrates our real-world environment setup. We use OAK-D Pro cameras, with a top-down camera to estimate the pose of the object for grasping and inserting, and a side-view camera to estimate the object's pose in the shelf. Objects placed on the table are always oriented perpendicular to the table edge and positioned adjacent to it. Shelf sizes of 0.18m and 0.35m are tested, and objects in the shelf are placed at randomized positions and orientations. A point cloud representation of each object is created, including their sizes, positions, and orientations as the state representation.

Evaluation metrics: We evaluated our dynamics model and manipulation outcomes in simulation based on the final prediction error, applying metrics such as Mean Squared Error (MSE), Earth Mover's Distance (EMD) [29], and Chamfer distance (CD) [30]. In real-world scenarios, success rates were computed for each setup. A success was defined as all boxes ending up within the shelf with their orientations falling within a predefined threshold θ.

5 Experimental Results

In this section, we analyze the performance of our forward dynamics prediction model and conduct a comprehensive evaluation of our proposed framework in both simulated and real-world environments. We use a diverse range of objects and shelf dimensions to provide a broad and challenging test. The results highlight the potential of our framework for zero-shot sim2real transfer capabilities, demonstrating its applicability in handling real-world conditions without the necessity for prior real-world data.

Table 1: **Dynamics Model Prediction Quantitative Results and Ablations.** Our model consistently outperforms RoboCraft in complex primitives such as 'Sweep' and 'Transport', evident in the lower MSE, EMD, and CD values. In the simpler 'Push' primitive, our model retains comparable MSE, EMD, and CD values within one standard deviation, indicating robust performance even in simpler scenarios.

Primitive	**Method**	**MSE** (mm) (↓)	**EMD** (mm) (↓)	**CD** (mm) (↓)
Sweep	RoboCraft [13]	0.351 (± 0.222)	24.510 (± 6.434)	33.277 (± 5.064)
	Ours	**0.287** (± **0.185**)	**21.792** (± **5.533**)	**30.017** (± **3.259**)
Push	RoboCraft	**0.216** (± **0.148**)	**12.509** (± **3.494**)	15.43 (± **2.902**)
	Ours	0.292 (± 0.179)	14.569 (± 3.752)	**15.046** (± 3.085)
Transport	RoboCraft	1.091 (± 0.512)	42.162 (± 10.317)	55.615 (± 9.997)
	Ours	**0.666** (± **0.41**)	**31.232** (± **9.108**)	**38.068** (± **6.605**)

Table 2: **Quantitative Results of Control in Simulation.** Our method consistently outperforms PPO and SAC, exhibiting markedly lower execution errors, illustrating the critical role of incorporating a model for long-horizon stowing tasks.

Method	**MSE** (mm) (↓)	**EMD** (mm) (↓)	**CD** (mm) (↓)
PPO	87.479	266.554	173.522
SAC	265.411	465.639	368.571
Ours	**0.905**	**29.697**	**39.914**

5.1 Dynamics Model Learning and Ablations

We trained our dynamics model utilizing MSE loss. As part of our ablation study, we incorporated a version of the GNN presented in [13], which we will refer to as "RoboCraft". It's important to note that this model does not encompass dynamic-static object interactions, nor does it utilize the auxiliary loss derived from gripper movements.

The quantitative results from the model learning are shown in Table 1. The relatively small prediction errors suggest that our model is able to accurately predict the interactions involved in the task; these include object/object interactions, environment/object interactions, and robot gripper/object interactions. The results further demonstrate that the model can effectively learn the rigid motions of the objects, resulting in only minor errors.

Compared with RoboCraft, the introduction of gripper movement information and dynamic-static object edge information improves the prediction accuracy of our GNN model, particularly in complex behavior primitives such as sweeping and transporting. Even in simpler behavior primitives like pushing, our GNN maintains a performance comparable to the RoboCraft. The final prediction error stays within one standard deviation, indicating the model's continued effectiveness even in simpler scenarios.

5.2 Manipulation Results

Results in Simulation: We collect six demonstrations in the simulation. The keyframes from these demonstrations serve as the goal state for trajectory optimization, implemented with Random Shooting (RS). In our analysis, we evaluated Random Shooting, the Cross-Entropy Method, and Gradient Descent - all of which exhibited similar performance. We present the results obtained from RS and denote them as "Ours" in the subsequent discussion. We use Proximal Policy Optimization (PPO) [31] and Soft Actor-Critic (SAC) [32], both state-of-the-art, model-free RL algorithms, as baselines.

Setups	**Success** (↑)		
	Heuristic	**3 skills**	**2 skills**
(a)	1/10	10/10	10/10
(b)	3/10	9/10	9/10
(c)	3/10	9/10	9/10
(d)	1/10	10/10	10/10
(e)	2/10	10/10	10/10
(f)	1/10	9/10	9/10
Average	18%	95%	95%

Table 3: **Success Rates of Control Actions in Real-World Environments.** The performance of our proposed method, using two different skill sets (2-skill and 3-skill) and a heuristic-based control action selection, is evaluated across six different setups.

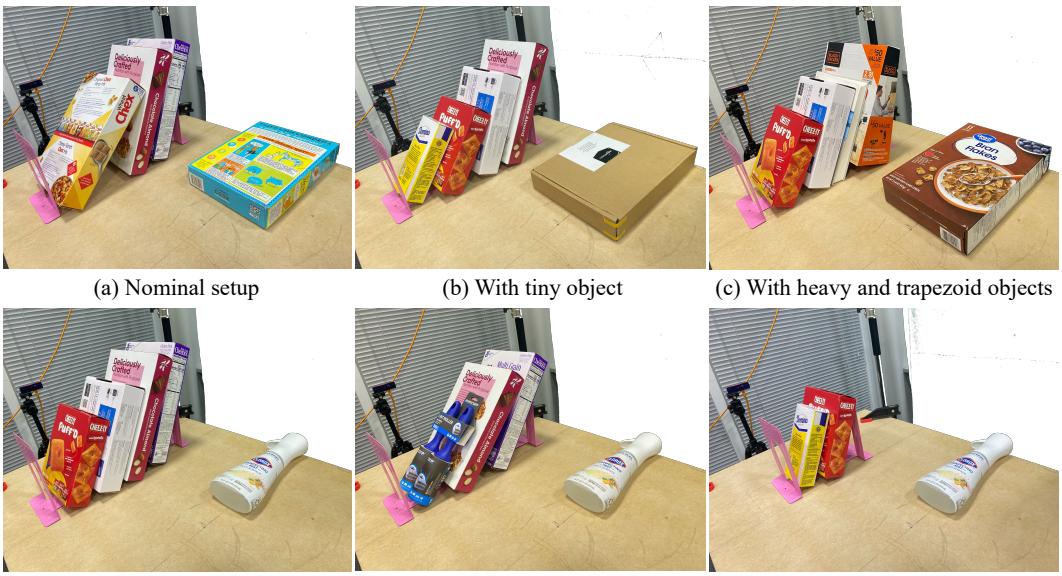

(a) Nominal setup (b) With tiny object (c) With heavy and trapezoid objects

(d) With bottle to be grasped (e) With slippery roller (f) Small shelf setup

Figure 4: Different Setups Used in Real-World Manipulation Experiments. Six setups represent a wide range of conditions including different combinations of objects, shelf sizes, object dimensions, and shapes.

The choice of these algorithms aims to highlight the necessity of a model for long-horizon stowing tasks. The simulation results are presented in Table 2. Each method's performance is assessed using MSE, EMD, and CD as metrics. PPO and SAC utilize the negative MSE between the current and goal object states as the reward function. In comparison to these methods, RS consistently yields lower execution errors across all metrics, which indicates its superior capability in minimizing the gap between actual and desired states. Model-free RL methods such as PPO and SAC struggle to perform effectively due to the large exploration space and the long-horizon nature of the task. Their effectiveness is further hampered by their lack of knowledge regarding the model of the environment.

Results in the Real World: Our framework is tested in six different real-world setups, with each setup executed for ten test trails. We manually randomize the initial orientations and positions of the objects within the shelf in each trail. The object poses in these scenarios are identified using Scale-Invariant Feature Transform (SIFT) on the captured images. The various setups represent a broad spectrum of conditions including different object combinations, shelf sizes, object dimensions, and shapes. This wide range of conditions is depicted in Figure 4.

In our experiments, we implement two distinct skill combinations for each setup: a 3-skill set comprising sweeping, pushing, and transporting, and a 2-skill set, which only included pushing and transporting. The term "heuristic" refers to a process where humans fine-tune a relatively small parameter space and assign the tuned parameters to the skills. The heuristic-based approach is likewise conducted utilizing all three skills: sweeping, pushing, and transporting.

Figure 4 presents the success rates of the different control strategies. Our 3-skill method significantly outperforms the heuristic-based approach with a success rate of 95%, indicating the effective handling of varied setups by our dynamics prediction module. Interestingly, the 2-skill set also achieves the same 95% success rate, indicating that the robot's ability to understand the interactions between the gripper-held object and the objects within the shelf enables it to determine the optimal position for insertion and placement. These high success rates demonstrate the robustness and efficacy of our method. In contrast, the heuristic-based strategy yielded an average success rate of only 18%. Despite being trained solely with box-shaped objects, our method generalizes effectively to out-of-distribution objects, showing its versatility in a variety of real-world conditions.

Figure 5 provides a qualitative comparison of the sweeping skill execution between the heuristic-based method and our approach. Our method, equipped with the ability to anticipate future states

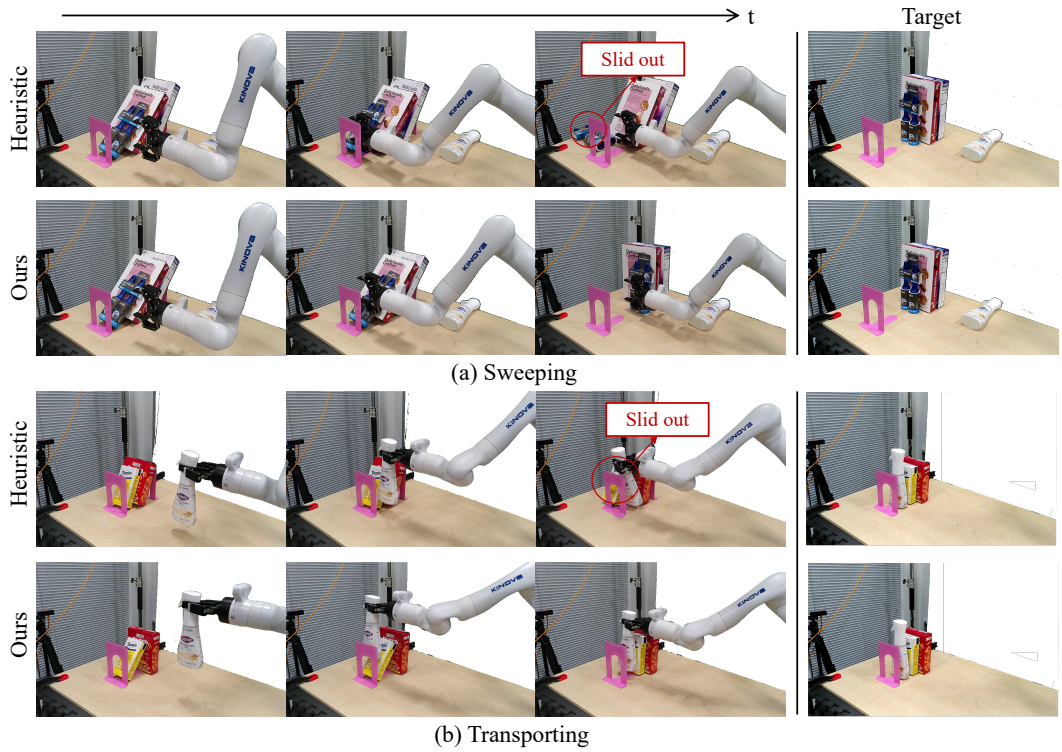

Figure 5: Comparison of the heuristic-based method and our approach during the execution of the sweeping skill and transporting skill. Our method anticipates future states and arranges objects into upright positions within the shelf, unlike the heuristic-based method which pushes objects out of the shelf.

based on specific robot actions, is capable of sweeping and transporting objects into upright positions within the shelf. In contrast, the heuristic-based method tends to push objects out of the shelf.

6 Limitations

Our system currently has a few limitations. Firstly, it relies on manual human labeling of ordered keyframes from demonstrations, which could potentially restrict scalability and deployment in larger and more complex scenarios. Secondly, we use box-shaped point clouds to represent objects during training and inference. This simplistic representation may not accurately reflect the geometrical properties of objects, especially in scenarios involving more complex interactions and contacts. Addressing these limitations, particularly improving object representation, presents a promising direction for future research.

7 Conclusion

In this work, we focus on stowing tasks wherein a robot must manipulate a large, ungraspable flat object and subsequently stow it within a cluttered shelf. The robot's task involves creating sufficient space within the cluttered shelf to accommodate the object appropriately. We introduce a system that utilizes behavior primitives in combination with forward dynamics prediction to achieve the task. We discuss the design choices for our system and demonstrate its effectiveness in real robot scenarios. Moreover, we show the system's ability to generalize to various stowing conditions.

Our work opens several potential avenues for future research. One promising direction involves developing the ability to composite skills and further reduce the sub-goals presented in the demonstrations. Another is that the design and definition of the behavior primitives library need additional exploration and research, which can enhance the adaptability and versatility of robotics systems in performing complex manipulation tasks.

References

[1] F. R. Hogan and A. Rodriguez. Feedback control of the pusher-slider system: A story of hybrid and underactuated contact dynamics. In *Workshop on the Algorithmic Foundations of Robotics*, 2016.

[2] J. Luo, O. Sushkov, R. Pevceviciute, W. Lian, C. Su, M. Vecerik, N. Ye, S. Schaal, and J. Scholz. Robust Multi-Modal Policies for Industrial Assembly via Reinforcement Learning and Demonstrations: A Large-Scale Study. In *Proceedings of Robotics: Science and Systems*, Virtual, July 2021. doi:10.15607/RSS.2021.XVII.088.

[3] T. Zhao, J. Luo, O. O. Sushkov, R. Pevceviciute, N. M. O. Heess, J. Scholz, S. Schaal, and S. Levine. Offline meta-reinforcement learning for industrial insertion. *2022 International Conference on Robotics and Automation (ICRA)*, pages 6386–6393, 2022.

[4] C.-Y. Chai, K.-F. Hsu, and S.-L. Tsao. Multi-step pick-and-place tasks using object-centric dense correspondences. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4004–4011. IEEE, 2019.

[5] E. Johns. Coarse-to-fine imitation learning: Robot manipulation from a single demonstration. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 4613–4619. IEEE, 2021.

[6] S. Nair, A. Rajeswaran, V. Kumar, C. Finn, and A. Gupta. R3m: A universal visual representation for robot manipulation. In *Conference on Robot Learning (CoRL)*, 2022.

[7] B. Wen, W. Lian, K. Bekris, and S. Schaal. You only demonstrate once: Category-level manipulation from single visual demonstration. In *Robotics: Science and Systems (RSS)*, 2022.

[8] Y. Duan, M. Andrychowicz, B. Stadie, O. Jonathan Ho, J. Schneider, I. Sutskever, P. Abbeel, and W. Zaremba. One-shot imitation learning. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL https://proceedings.neurips.cc/paper_files/paper/2017/file/ba3866600c3540f67c1e9575e213be0a-Paper.pdf.

[9] A. Ganapathi, P. Sundaresan, B. Thananjeyan, A. Balakrishna, D. Seita, J. Grannen, M. Hwang, R. Hoque, J. E. Gonzalez, N. Jamali, et al. Learning dense visual correspondences in simulation to smooth and fold real fabrics. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 11515–11522. IEEE, 2021.

[10] H. T. Suh and R. Tedrake. The surprising effectiveness of linear models for visual foresight in object pile manipulation. In *Algorithmic Foundations of Robotics XIV: Proceedings of the Fourteenth Workshop on the Algorithmic Foundations of Robotics 14*, pages 347–363. Springer, 2021.

[11] W. Yan, A. Vangipuram, P. Abbeel, and L. Pinto. Learning predictive representations for deformable objects using contrastive estimation. In *Conference on Robot Learning*, pages 564–574. PMLR, 2021.

[12] P. Mitrano, D. McConachie, and D. Berenson. Learning where to trust unreliable models in an unstructured world for deformable object manipulation. *Science Robotics*, 6(54):eabd8170, 2021.

[13] H. Shi, H. Xu, Z. Huang, Y. Li, and J. Wu. Robocraft: Learning to see, simulate, and shape elasto-plastic objects with graph networks. In *Robotics: Science and Systems (RSS)*, 2022.

[14] X. Lin, Y. Wang, Z. Huang, and D. Held. Learning visible connectivity dynamics for cloth smoothing. In *Conference on Robot Learning*, pages 256–266. PMLR, 2022.

[15] Z. Huang, X. Lin, and D. Held. Mesh-based dynamics with occlusion reasoning for cloth manipulation. In *Robotics: Science and Systems (RSS)*, 2022.

[16] Y. Li, J. Wu, R. Tedrake, J. B. Tenenbaum, and A. Torralba. Learning particle dynamics for manipulating rigid bodies, deformable objects, and fluids. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net, 2019. URL https://openreview.net/forum?id=rJgbSn09Ym.

[17] Y. Li, J. Wu, J.-Y. Zhu, J. B. Tenenbaum, A. Torralba, and R. Tedrake. Propagation networks for model-based control under partial observation. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 1205–1211. IEEE, 2019.

[18] P. Battaglia, R. Pascanu, M. Lai, D. Jimenez Rezende, et al. Interaction networks for learning about objects, relations and physics. *Advances in neural information processing systems*, 29, 2016.

[19] M. Chang, T. D. Ullman, A. Torralba, and J. B. Tenenbaum. A compositional object-based approach to learning physical dynamics. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net, 2017. URL https://openreview.net/forum?id=Bkab5dqxe.

[20] A. Sanchez-Gonzalez, N. Heess, J. T. Springenberg, J. Merel, M. Riedmiller, R. Hadsell, and P. Battaglia. Graph networks as learnable physics engines for inference and control. In *International Conference on Machine Learning*, pages 4470–4479. PMLR, 2018.

[21] N. Funk, G. Chalvatzaki, B. Belousov, and J. Peters. Learn2assemble with structured representations and search for robotic architectural construction. In *Conference on Robot Learning*, pages 1401–1411. PMLR, 2022.

[22] T. Silver, R. Chitnis, A. Curtis, J. B. Tenenbaum, T. Lozano-Pérez, and L. P. Kaelbling. Planning with learned object importance in large problem instances using graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35, pages 11962–11971, 2021.

[23] A. Gupta, V. Kumar, C. Lynch, S. Levine, and K. Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. In L. P. Kaelbling, D. Kragic, and K. Sugiura, editors, *3rd Annual Conference on Robot Learning, CoRL 2019, Osaka, Japan, October 30 - November 1, 2019, Proceedings*, volume 100 of *Proceedings of Machine Learning Research*, pages 1025–1037. PMLR, 2019. URL http://proceedings.mlr.press/v100/gupta20a.html.

[24] N. Di Palo and E. Johns. Learning multi-stage tasks with one demonstration via self-replay. In *Conference on Robot Learning*, pages 1180–1189. PMLR, 2022.

[25] X. Lin, Z. Huang, Y. Li, J. B. Tenenbaum, D. Held, and C. Gan. Diffskill: Skill abstraction from differentiable physics for deformable object manipulations with tools. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net, 2022. URL https://openreview.net/forum?id=Kef8cKdHWpP.

[26] O. Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53, 1987. doi:10.1109/JRA.1987.1087068.

[27] Y. Zhu, J. Wong, A. Mandlekar, R. Martín-Martín, A. Joshi, S. Nasiriany, and Y. Zhu. robosuite: A modular simulation framework and benchmark for robot learning. In *arXiv preprint arXiv:2009.12293*, 2020.

[28] E. Todorov, T. Erez, and Y. Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033, 2012. doi:10.1109/IROS.2012.6386109.

[29] Y. Rubner, C. Tomasi, and L. Guibas. A metric for distributions with applications to image databases. In *Sixth International Conference on Computer Vision (IEEE Cat. No.98CH36271)*, pages 59–66, 1998. doi:10.1109/ICCV.1998.710701.

[30] H. G. Barrow, J. M. Tenenbaum, R. C. Bolles, and H. C. Wolf. Parametric correspondence and chamfer matching: Two new techniques for image matching. In *Proceedings of the 5th International Joint Conference on Artificial Intelligence - Volume 2*, IJCAI'77, page 659–663, San Francisco, CA, USA, 1977. Morgan Kaufmann Publishers Inc.

[31] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov. Proximal policy optimization algorithms, 2017.

[32] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In J. G. Dy and A. Krause, editors, *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 1856–1865. PMLR, 2018. URL http://proceedings.mlr.press/v80/haarnoja18b.html.