

# Harnessing the Synergy between Pushing, Grasping, and Throwing to Enhance Object Manipulation in Cluttered Scenarios

Hamidreza Kasaei<sup>1\*</sup> and Mohammadreza Kasaei<sup>2\*</sup>

**Abstract**—In this work, we delve into the intricate synergy among non-prehensile actions like pushing, and prehensile actions such as grasping and throwing, within the domain of robotic manipulation. We introduce an innovative approach to learning these synergies by leveraging model-free deep reinforcement learning. The robot’s workflow involves detecting the pose of the target object and the basket at each time step, predicting the optimal push configuration to isolate the target object, determining the appropriate grasp configuration, and inferring the necessary parameters for an accurate throw into the basket. This empowers robots to skillfully reconfigure cluttered scenarios through pushing, creating space for collision-free grasping actions. Simultaneously, we integrate throwing behavior, showcasing how this action significantly extends the robot’s operational reach. Ensuring safety, we developed a simulation environment in Gazebo for robot training, applying the learned policy directly to our real robot. Notably, this work represents a pioneering effort to learn the synergy between pushing, grasping, and throwing actions. Extensive experimentation in both simulated and real-robot scenarios substantiates the effectiveness of our approach across diverse settings. Our approach achieves a success rate exceeding 80% in both simulated and real-world scenarios. A video showcasing our experiments is available online at: <https://youtu.be/q114BJVDbRw>

## I. INTRODUCTION

The field of robotic manipulation has evolved significantly, focusing on the synergy between non-prehensile actions like pushing and prehensile actions such as grasping and throwing [1], [2]. The importance of learning such a synergy became clear in cluttered scenarios. Pushing, grasping, and throwing are not just individual actions; they are interconnected skills that, when combined, unlock a world of possibilities for service robots. Our approach integrates these actions into a unified framework to enable robots to autonomously manage clutter and execute complex tasks. This involves a sequence of training steps, starting with an object-agnostic grasping foundation, followed by strategic push, and then throw policies. We adopt a modular training approach, allowing easier troubleshooting and updates without impacting the whole system. Moreover, this approach reduces the need for extensive data collection for end-to-end training, as each policy is optimized separately but coherently. Particularly, the push policy is conditioned on the output of an object-agnostic grasping policy, ensuring a structured framework essential for complex manipulation tasks.

<sup>1</sup>Hamidreza Kasaei is with the Department of Artificial Intelligence, University of Groningen, The Netherlands. Email: hamidreza.kasaei@rug.nl

<sup>2</sup> Mohammadreza Kasaei is with the School of Informatics, University of Edinburgh, UK. Email: m.kasaei@ed.ac.uk

\* Equal contribution - This work is partially supported by Google DeepMind through the Research Scholar Program for the “Continual Robot Learning in Human-Centered Environments” project.

To mitigate the safety risks and time-intensive nature of real-world training, we developed a simulation environment in Gazebo that closely mirrors our real robot. An overview of the proposed method is depicted in Fig. 1. Through simulation training and empirical evaluation, our framework proves to be both adaptable and effective, achieving superior performance compared to baseline alternatives. Our key contributions are groundbreaking in three main areas. Firstly, we are pioneers in integrating pushing, grasping, and throwing actions to expand the capabilities of robots in complex environments. Secondly, our approach, although trained on simulation data, shows remarkable adaptability and generalization in real-world scenarios. Lastly, our framework achieves a success rate exceeding 80% in both simulated and real-world scenarios, confirming its practical viability.

## II. RELATED WORK

In this section, we provide a brief overview of the related work, organizing it into several distinct subsections.

**Pushing techniques:** Numerous studies have contributed to the advancement of pushing strategies for efficiently reorganizing cluttered objects. Previous research has encompassed analytical techniques [3], [4] as well as learning-based approaches [5], [6], [7], [8], [9] for push planning. These investigations have significantly improved object rearrangement in intricate settings. In contrast to these approaches, our method conditions the push policy on the grasp quality map of the scene. This implies that the robot learns to execute a push action that renders the target object graspable.

**Grasping methods:** Grasping, as a cornerstone of robotic manipulation, has seen remarkable progress, with a spectrum of techniques ranging from analytical grasp synthesis to data-driven learning methods. Notably, object-agnostic grasping approaches have garnered attention for their adaptability across diverse object types and scenarios [10], [11], [10], [12]. We encourage readers to explore a comprehensive survey on object grasping [13] for a more in-depth understanding of this field. Among all possible object-grasping approaches, we adopted the MVGRASP [12] approach as our grasp policy network in this work.

**Synergy between pushing and grasping:** The domains of pushing and grasping have been the focal points of extensive research within the field of robotic manipulation [1], [14], [15], [2]. The synergy between these two fundamental actions has led to significant advancements in the realm of object manipulation. These studies have explored how pushing can pave the way for successful grasps, thereby enhancing efficiency and mitigating collisions in cluttered

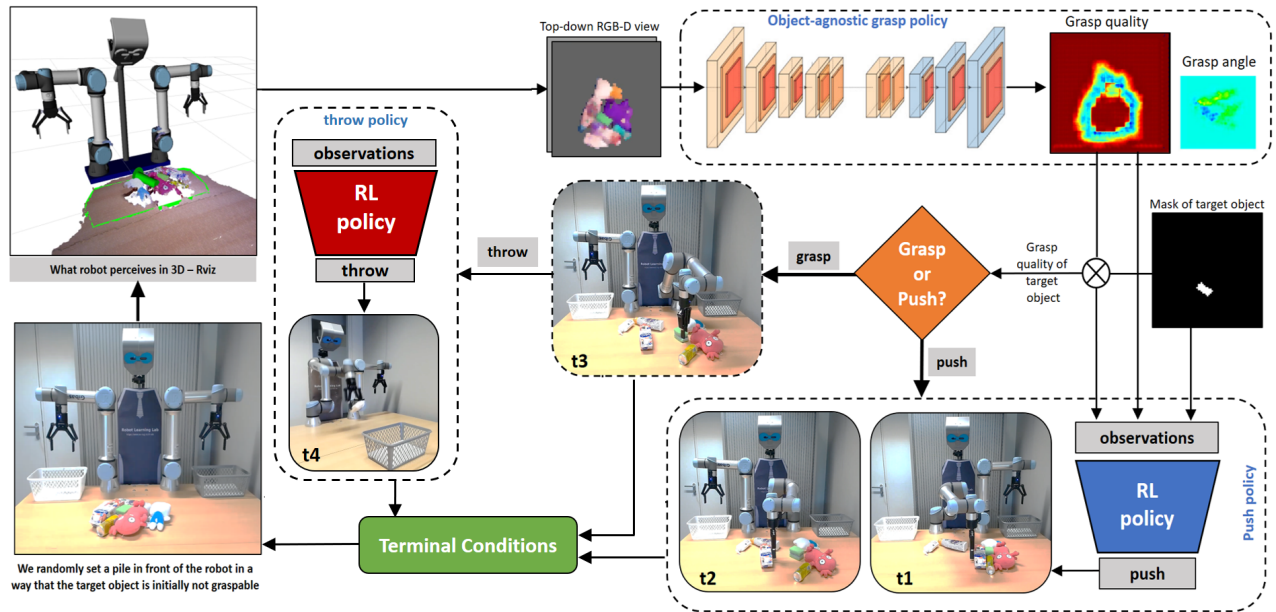


Fig. 1: Overview: The perception system provides essential inputs, including top-down RGB-D views of the workspace, and a mask highlighting the target object. These inputs are then processed through an object-agnostic grasp policy, resulting in pixel-wise grasp synthesis for the scene. Based on the grasp quality of the target object, the system makes a decision between executing a push action or a grasp action. Specifically, if the grasp quality surpasses a predefined threshold, the robot initiates a grasp; otherwise, it proceeds with a push action. After successfully grasping the target object, the robot leverages throwing actions when the target basket is out of its immediate reach.

environments [1], [14], [15], [2], [15]. In contrast to these approaches, our method prioritizes learning to grasp objects in cluttered scenarios before acquiring a push policy to render the target object graspable. Moreover, our approach stands out in its ability to simultaneously learn the length and direction of the push in a single pass. This differs from most state-of-the-art approaches, which typically require passing a batch of 16 images to the network and subsequently identifying the image with the highest q-value [1], [2]. The direction of the push is determined based on the index of the image, while the push’s length remains a fixed value.

**Throwing object:** Early works in the field of robotic throwing primarily focused on precise object placement. These studies often involved analytical approaches and dedicated throwing kernels [16], [17], [18]. The challenge lay in accurately predicting throwing parameters to achieve desired placements, a task further complicated by varying object shapes and environmental constraints [19], [20]. Recent advancements in deep reinforcement learning have offered new avenues for teaching robots the art of throwing [19], [20]. Similar to these approaches, we used an RL-based method to enable robots to adapt and generalize their throwing capabilities. Integrating throwing with pushing and grasping actions represents a novel and promising direction.

### III. METHOD

#### A. Preliminaries

**Markov Decision Process (MDP):** An MDP is a fundamental framework defined by a tuple comprising four essential elements:  $(s_t, a_t, p(s_{t+1}|s_t, a_t), r(s_{t+1}|s_t, a_t))$ . Here,

$s_t$  and  $a_t$  represent the continuous state and action at time step  $t$ , respectively. The function  $p(s_{t+1}|s_t, a_t)$  signifies the transition probability, indicating the likelihood of transitioning from the current state  $s_t$  to the next state  $s_{t+1}$  given the action  $a_t$ . The function  $r(s_{t+1}|s_t, a_t)$  denotes the immediate reward received from the environment.

**Off-policy RL:** In online RL, an agent engages in continuous interactions with the environment to gather experiences and learn the optimal policy  $\pi^*$ . The objective is to maximize the expected future return  $R_t = \mathbb{E}[\sum_{i=t}^{\infty} \gamma^{i-t} r_{i+1}]$ , with a discount factor  $\gamma \in [0, 1]$  accounting for the importance of future rewards. The expected return under a policy  $\pi$  after taking action  $a$  in the state  $s$  is computed by the corresponding action-value function  $Q^\pi(s, a) = \mathbb{E}[R_t | s_t = s, a_t = a]$ . This can be computed using the Bellman equation:

$$Q^\pi(s_t, a_t) = \mathbb{E}_{s_{t+1} \sim p} [r(s_t, a_t) + \gamma \mathbb{E}_{a_{t+1} \sim A} [Q_\pi(s_{t+1}, a_{t+1})]], \quad (1)$$

where  $A$  denotes the action space. The ultimate goal of RL algorithms is to discover an optimal policy  $\pi^*$  such that  $Q^{\pi^*}(s, a) = Q^*(s, a)$  for all states and actions.

#### B. Problem Formulation

Initially, the target object remains ungraspable due to surrounding objects. Given the initial pose of the target object,  $p_s$ , the push policy aims to determine the proper parameters (i.e., initial push point, direction, and length) to manipulate the scene, making the target object graspable (i.e., the grasp quality of the target object exceeds a predefined threshold). In essence, the push policy is strategically conditioned upon the output of an object-agnostic grasping policy,

which has undergone unsupervised training. Similarly, the throwing policy seeks to determine the appropriate residual parameters of the initial throwing kernel, previously taught to the robot through Kinesthetic teaching, while considering the goal pose  $p_g$  (i.e., the pose of the basket) in the environment (see



Fig. 2: Kinesthetic teaching of throwing kernel.

Fig. 2). These tasks can be represented using a fully observable MDP framework, and an off-policy reinforcement learning framework is employed to solve them. It is important to highlight that each of these policies is trained independently. This decision is grounded in the understanding that achieving end-to-end training for such multifaceted tasks would necessitate extensive training samples. Notably, the push policy is strategically reliant on the output of an object-agnostic grasping policy. This modular approach enhances the system's decomposability, thereby enabling efficient troubleshooting, fine-tuning, and updates for each policy, all without causing effects throughout the system.

1) **States:** In both the push and throw policies, we utilize a feature vector to characterize the continuous state. In the context of the push and grasp policy, the observation space includes a flattened grasp quality map consisting of  $50 \times 50$  pixels. Additionally, it encompasses eleven supplementary parameters, comprising the initial push point coordinates (X, Y, Z), push direction, and push length, all represented by five floating-point numbers. Further, it includes the coordinates of the target object (X, Y, Z), and three float numbers representing the initial grasp quality of the target object, the grasp quality of the target object after a push, and the difference between these qualities.

Within the throwing policy, the observation state encompasses several critical elements. These include the robot's proprioception, information regarding the target goal, releasing time, duration of trajectory execution, and the distance between the thrown object and the goal. The robot autonomously learns the initial and final values for the shoulder joint, the duration of trajectory execution (speed), and the optimal releasing time during the learning process.

At each time step  $t$ , we record essential parameters. This includes the initial and final shoulder joint values ( $j_i$  and  $j_f$ ) in radians, representing proprioception as  $\text{proprio} = (j_i, j_f) \in \mathbb{R}^2$ . The position of the goal (center of the box) in the task space is described as  $\text{goal} = (x^g, y^g, z^g) \in \mathbb{R}^3$ . Additionally, we consider the absolute distance of the thrown object relative to the goal, including distances in the X and Y axes. These distances are represented as  $\text{dist} = (d^g, d_x^g, d_y^g) \in \mathbb{R}^3$ . Furthermore, we record two timing profiles: the duration of executing the throwing trajectory  $\tau$ , and the time for releasing the object,  $t_r$ , where  $t_r < \tau$ . These timings are captured as  $\text{time} = (t_r, \tau) \in \mathbb{R}^2$ . In summary, the state space of throwing policy can be represented as a vector:  $s = (\text{proprio}, \text{goal}, \text{dist}, \text{time}) \in \mathbb{R}^{10}$ .

2) **Actions:** In the case of the push and grasp policy, the action space is defined as follows: A grasp action is taken when the maximum grasp quality of the target object exceeds a threshold, where the threshold is set at 0.7, otherwise, a push action is executed. For the throwing policy, each action is denoted by a vector  $a \in ([-1, 1])^4$ , which represents (i) the initial and (ii) the final shoulder joint values, (iii) the duration of trajectory execution, and (iv) the releasing time.

3) **Transition function:** During each training episode, we introduce random variations to the positions of the basket, the target object, and the surrounding objects in a manner that renders the target object initially ungraspable. Consequently, the dynamics of the transition function are influenced by the execution of push actions or throwing trajectories, based on the current state of the scene and sampled parameters. To be specific, we compute the subsequent state, denoted as  $s_{i+1}$ , following the execution of the action  $a_i$ , as defined by the transition function  $f_s$ ; in mathematical terms, this relationship is expressed as  $s_{i+1} = f_s(s_i, a_i)$ , where  $s_i$  represents the current state, and  $a_i$  signifies the action undertaken. It should be noted that, given the unknown nature of the transition function, our off-policy reinforcement learning framework remains model-free, adapting to various scenarios without requiring a predefined model.

4) **Rewards:** Our objective for the push and grasp policy is to equip the robot with the capability to efficiently manipulate the scene so that the target object becomes graspable following a push action. This is indicated by the grasp quality of the target object surpassing a predefined threshold, denoted as  $\tau_g = 0.7$ , for successful grasping. For successful grasping, we assign a maximum reward of 1.0. Additionally, if the grasp quality of the target object improves after the execution of a push action, we compute the rewards as follows:  $R = \alpha \cdot e^{(-d^2/0.001)} + (1 - \alpha) \cdot e^{(-d^2/0.05)}$ . Here,  $\alpha$  is set to 0.9, and  $d$  is calculated as  $1 - \beta$ , with  $\beta$  representing the maximum grasp quality of the target object. On the other hand, actions deemed as ineffective or useless are penalized with a reward of  $-0.1$ . Terminal conditions for the push and grasp policy include:

- The target object goes beyond the robot's workspace.
- The number of consecutive push actions exceeds a predefined budget (i.e., set at five consecutive pushes).
- The robot executes a grasp action.

In the case of the throwing action, a successful outcome is achieved if the thrown object lands inside the target basket ( $R = 1.0$ ). We determine success by calculating the absolute distance between the object and the goal, where the distance is compared to the radius ( $d$ ) of a cylindrical space fitted inside the basket. If the next state results in success (i.e., the thrown object lands inside the basket), we incentivize this behavior by setting the reward  $R = 1.0$ . Conversely, if the throwing action results in the object landing outside the target basket, the reward is penalized based on the distance, with  $R = -\text{dis}(o, g)$ . A throwing episode is terminated upon the execution of a throwing action.




Fig. 3: Visualizing the output of our perception system: world model information is provided through a top-down view, a grasp quality map, and a mask of the target object. The robot’s workspace is outlined by the green rectangle, and the predicted push action is shown by the green arrow.

### C. Perception

In the interest of safety, our push, grasp, and throw policies were initially trained in a simulated environment before being tested on our real-robot platform. To facilitate this transition, we developed a versatile interface capable of processing RGB-D sensory data, which was utilized seamlessly in both our simulation and real robot experiments [21][22][12]. Specifically, we employed an RGB-D Asus Xtion camera, capturing point cloud data at a rate of 30 Hz. For tracking the position of the target object, we adopted a straightforward object detection method that takes into account both shape and color data [23]. It is important to note that, we constrained the system to grasp the object from above and near its center of mass [12]. Our perception system serves as a world model service, enabling the agent to access real-time environmental information at each time step. Figure 3 provides a visual representation of the outcomes generated by our perception system. For a comprehensive understanding of our perception and grasping pipelines, we encourage referring to our earlier research for more detailed insights [23][12].

## IV. EXPERIMENTS AND RESULTS

Our experimental evaluation involves assessing the proposed approach based on several key performance metrics: (i) *Success Rate* calculating by dividing the number of instances in which the robot successfully places the target object into the basket by the total number of trials. (ii) *Required Number of Actions* needed to successfully singulate the target object.

### A. Experimental setup and tasks settings

Our experimental setups in both simulation and real-robot environments are represented in Fig. 4. Specifically, we developed a simulation environment in Gazebo, employing the ODE physics engine to closely emulate our real robot. Our setup includes an Asus Xtion camera, two Universal Robots (UR5e) equipped with Robotiq 2F-140 grippers, and an interface for initiating and concluding experiments. To evaluate the proposed approach, we devised three distinct tasks, each progressively more challenging:

- **Task 1:** The objective here is to singulate the target object from a cluttered scene and place it into a reachable basket. An illustrative example of this task in a real-robot setup can be seen in Fig. 4 (*center*).

- **Task 2:** The robot is tasked with throwing an object into a basket positioned randomly in front of it. Fig. 4 (*left*) offers a glimpse of this task in a simulation environment.
- **Task 3:** This multifaceted task requires the robot to first singulate the target object, followed by grasping and throwing it into a basket situated beyond the robot’s maximum kinematic range. A visual depiction of this task is shown in Fig. 4 (*right*).

For each of the proposed tasks, we trained the model for 100,000 steps in the simulation. Specifically, we trained the throwing policy using a cubic object as shown in Fig. 4 (*left*). Moreover, we conducted the training of the push and grasp policy in a cubic scenario, where the target object initially resides in a configuration that renders it ungraspable, as illustrated in Fig. 4 (*center*). Subsequently, we evaluated the learned policy on real robots across ten distinct scenarios, as depicted in Fig. 5.

### B. Baseline Methods

We harnessed two state-of-the-art, sample-efficient off-policy RL algorithms to train our robotic system: Deep Deterministic Policy Gradient (DDPG) [24], [25], and Soft Actor-Critic (SAC) [26], [27] via the stable baseline [28]. The neural network architectures

TABLE I: SAC hyper-parameters

Parameter	Value
#hidden layers (all networks)	2
#hidden units per layer	256
#samples per minibatch	256
optimizer	Adam
learning rate	$3 \times 10^{-4}$
batch size	256
#epochs	50K
discount ( $\gamma$ )	0.99
replay buffer size	$10^6$
nonlinearity	ReLU
target update rate ( $\tau$ )	0.005
target update interval	1
gradient steps	1

for both SAC and DDPG consisted of two hidden layers, each comprising 256 neurons, activated by Rectified Linear Units (ReLU) activation functions. The hyper-parameters of SAC are listed in Table. I, and the hyper-parameters of DDPG that are not listed in Table I are reported in the Experiments section of our previous work [29].

### C. Results

In the context of **Task 1**, which involves singulating the target object from a cluttered scene and depositing it into a reachable basket, we conducted a comprehensive evaluation comprising 1000 experiments with a cubic scenario, 500 simulation experiments with the SAC policy, 500 simulation experiments with the DDPG policy. Furthermore, we executed 200 real-world experiments, spanning across 10 distinct cubic scenarios (see Fig. 5). Notably, the first scenario closely resembled the simulated environment (#1), while the remaining nine scenarios presented entirely novel challenges that had not been encountered during training. For each scenario, we executed 10 real robot tests with the SAC policy and 10 real robot experiments with the DDPG policy. This round of experiments provides key insights into the comparative efficacy and efficiency of the two algorithms. Table II summarizes the outcomes.




Fig. 4: Our experimental setups: (from left to right) Training the throwing policy, the push-and-grasp policy, integrating all policies into a unified robotic system, and the real dual-arm robot setup.




Fig. 5: We created 10 cubic scenarios to assess the efficacy of the acquired policies in real robot experiments.

By comparing the obtained results, it is clear that under the SAC policy, the task success rate (percentage of successful trials) ranged from 80% to 100%, while under the DDPG policy, it varied from 70% to 100%. The average number of actions required to accomplish the task (lower is better) ranged from 1.6 to 3.8 for SAC and 2.7 to 4.8 for DDPG. Overall, the results indicate that the SAC policy outperformed DDPG in most scenarios, achieving a higher task success rate and requiring fewer actions on average to complete the task. On average, SAC achieved a task success rate of 94% with a standard deviation of approximately 6.63%, while DDPG achieved an average success rate of 86% with a standard deviation of about 10.19%. Additionally, SAC exhibited an average number of actions of approximately 2.66 with a standard deviation of 0.71, whereas DDPG had an average of about 3.75 actions with a standard deviation of approximately 0.66. These results suggest that the SAC policy is more robust and effective in handling the complexity of real-world scenarios, as it consistently outperformed DDPG across various scenarios in **Task 1**. Furthermore, we evaluated the learned policies using various piles of objects scenarios. We

TABLE II: Performance of the SAC and DDPG policies on 10 real scenarios in Task 1.

Scenario	Task Success Rate ( $\uparrow$ )		Avg. Number of Actions ( $\downarrow$ )	
	SAC	DDPG	SAC	DDPG
#1	100	80	1.6	2.7
#2	100	90	1.8	3.4
#3	90	90	2.4	4.3
#4	90	100	2.9	3.8
#5	100	90	2.8	3.3
#6	100	100	1.7	2.8
#7	100	80	3.2	4.2
#8	90	90	3.3	4.4
#9	80	70	3.8	4.8
#10	90	70	3.1	3.8
avg $\pm$ std	<b>94 <math>\pm</math> 6.63</b>	86 $\pm$ 10.19	<b>2.66 <math>\pm</math> 0.71</b>	3.75 $\pm$ 0.66

used 10 daily-life objects with different materials, shapes, sizes, and weights. In these scenarios, the initial visibility of the target object was manipulated to create three distinct conditions: (i) the target object was fully visible, (ii) it was partially visible, and (iii) it was entirely concealed. To provide clarity, Fig. 6 showcases a representative example of each of these scenarios. Similar to the previous round of experiments, for each scenario we performed 20 real robot experiments including 10 experiments with SAC and 10 experiments with DDPG. Results are reported in Table III. Fig. 7 shows an example of a fully visible scenario.

By comparing all results, it is clear that in the “fully visible” scenario, where the target object was completely visible initially, the SAC policy achieved a 100% task success rate, while the DDPG policy achieved an 80% success rate. The SAC policy required an average of 2.9 actions, whereas DDPG needed an average of 3.3 actions to complete the task. In the “partially visible” scenario, SAC achieved a 100% task success rate, while DDPG achieved an 80% success rate. SAC had an average action count of 3.1, while DDPG averaged 3.6 actions. In the “not visible” scenario, where the target object was entirely hidden from view initially, the SAC policy achieved an 80% task success rate, while DDPG reached 60%. SAC required an average of 3.5 actions, while DDPG needed an average of 4.7 actions to accomplish the task. These results demonstrate the robustness of the SAC policy across different initial visibility conditions. In scenarios where the target object was partially or fully visible, both policies achieved high success rates. However, the SAC policy consistently required fewer actions to complete the task. In the challenging “not visible” scenario, SAC outperformed DDPG, highlighting its effectiveness in handling complex real-world situations.

In the second round of experiments (**Task2**), which involves the task of throwing a target object into a basket, the SAC policy demonstrated superior performance when compared to the DDPG policy. In the simulated environment,



Fig. 6: Visibility of the target object (green sponge) in Task 1 : (left) fully visible, (center) partially visible, (right) completely concealed.




Fig. 7: A real robot experiment demonstrating the singulation and grasping of a fully visible target object (a green sponge) from a pile of household objects.

the robot achieved a success rate of 95% with the SAC policy and 92% with the DDPG policy. However, in the real-robot experiments, the differences in performance became more pronounced. The robot equipped with the SAC policy achieved remarkable results, boasting a success rate of 92%. In contrast, when utilizing the DDPG policy, the robot’s throwing performance experienced a decline, with a success rate of 81%. This observed disparity between SAC and DDPG in real-robot experiments, as opposed to the relatively similar performance observed in simulation, underscores the adaptability and robustness of the SAC in addressing the complexities inherent in real-world scenarios.

**Task 3** represents a significantly more complex challenge compared to the previous tasks, requiring the robot to sequentially singulate the target object, grasp it, and then execute a throw into a basket positioned beyond the robot’s maximum kinematic range. As anticipated, the robot equipped with the SAC policy consistently outperformed its counterpart using the DDPG policy in both simulation and real-robot experiments. Specifically, the robot with the SAC policy achieved an impressive 88% success rate while requiring an average of 1.8 push actions per episode. Conversely, the robot using the DDPG policy achieved a success rate of 79% but required an average of 2.7 push actions per episode. Transitioning to the real-robot experiments, the performance gap between the two policies persisted, albeit with some modifications. The robot equipped with the SAC policy maintained a relatively high success rate of 85% while slightly increasing the average number of push actions to 2.1. Similarly, the robot utilizing the DDPG policy experienced a drop in success rate to 70% and a more significant increase in the average number of push actions to 4.4.

Comparing the performance across all tasks, it becomes evident that the difference in performance between the SAC policy and alternative policies becomes more pronounced as the task complexity increases. This disparity is particularly evident in the real-robot experiments.

#### D. Failure Cases

Throughout our experimental trials, we encountered four distinct categories of failures, they are including:

**Out of workspace failure:** In some real-robot experiments,

TABLE III: Performcane of the robot on piles of real objects.

Scenario	Task Success Rate (↑)		Avg. Number of Actions (↓)	
	SAC	DDPG	SAC	DDPG
fully visible	100	80	2.9	3.3
partially visible	100	80	3.1	3.6
not visible	80	60	3.5	4.7

the robot unintentionally moved the target object out of the workspace. This type of failure underscores the importance of maintaining objects within the robot’s operational boundaries. Addressing such occurrences can lead to improved control strategies to prevent objects from being pushed out of the workspace. It is worth noting that such failures were infrequent in the simulation environment, primarily because we imposed spatial constraints by encapsulating objects within a box-like workspace.

**Inaccurate grasp pose failure:** Another failure scenario arose when the predicted grasp pose was inaccurate. This led to instances where the robot either failed to grasp the target object or mistakenly grasped multiple objects simultaneously. Addressing this challenge calls for enhanced perception systems and grasp pose estimation techniques.

**Pushing limitation failure:** There were cases where the robot could not make the target object graspable even after the maximum allowable number of pushes. This indicates that the policy might need refinement, particularly in scenarios with complex object arrangements. Developing strategies to effectively manipulate cluttered scenes, such as learning more sophisticated non-prehensile pushing behaviors, can be essential for overcoming this limitation.

**Throwing failure:** The primary factors contributing to the failures observed in the throwing policy were the imprecise prediction of parameters and the latency in executing the gripper commands. Specifically, the gripper’s control process entails an intermediate step, where commands are relayed through the robot’s controller before reaching the gripper itself. This intermediary layer introduces a potential delay, influenced by network status and the robot’s conditions.

## V. CONCLUSIONS

In this work, we addressed the intricate challenges of harnessing the synergy between pushing, grasping, and throwing actions to enhance object manipulation in cluttered scenarios. We decoupled the learning of pushing, grasping, and throwing policies, recognizing that end-to-end training for such multifaceted tasks demands an extensive corpus of training samples. The push policy strategically relies on the output of an object-agnostic grasping policy. This modular approach enhances the system’s decomposability, enabling streamlined troubleshooting, fine-tuning, and updates to each policy without cascading effects on the entire robotic system. Through extensive experiments in both simulation and real-robot settings, we evaluated the performance of our approach across tasks of ascending complexity. We achieved success rates of over 80% in both simulation and real-robot environments. However, we also encountered challenges and observed failures. These included unintentional movements of the target object out of the robot’s workspace, inaccuracies in grasp pose predictions leading to failed grasps or grasping multiple objects, and difficulties in making the target object graspable with a limited number of pushes. In future work, we will leverage these insights, in combination with the integration of common-sense knowledge through Large Language Models, to handle sophisticated household tasks.

## REFERENCES

- [1] W. Zhou and D. Held, "Learning to grasp the ungraspable with emergent extrinsic dexterity," in *Conference on Robot Learning*. PMLR, 2023, pp. 150–160.
- [2] A. Zeng, S. Song, S. Welker, J. Lee, A. Rodriguez, and T. Funkhouser, "Learning synergies between pushing and grasping with self-supervised deep reinforcement learning," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2018, pp. 4238–4245.
- [3] K.-T. Yu, M. Bauza, N. Fazeli, and A. Rodriguez, "More than a million ways to be pushed: a high-fidelity experimental dataset of planar pushing," in *2016 IEEE/RSJ international conference on intelligent robots and systems (IROS)*. IEEE, 2016, pp. 30–37.
- [4] J. Moura, T. Stouraitis, and S. Vijayakumar, "Non-prehensile planar manipulation via trajectory optimization with complementarity constraints," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 970–976.
- [5] M. Bauza, F. R. Hogan, and A. Rodriguez, "A data-efficient approach to precise and controlled pushing," in *Conference on Robot Learning*. PMLR, 2018, pp. 336–345.
- [6] K. Lowrey, S. Kolev, J. Dao, A. Rajeswaran, and E. Todorov, "Reinforcement learning for non-prehensile manipulation: Transfer from simulation to physical system," in *2018 IEEE International Conference on Simulation, Modeling, and Programming for Autonomous Robots (SIMPAP)*. IEEE, 2018, pp. 35–42.
- [7] N. Dengler, D. Großklaus, and M. Bennis, "Learning goal-oriented non-prehensile pushing in cluttered scenes," in *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2022, pp. 1116–1122.
- [8] P. Florence, C. Lynch, A. Zeng, O. A. Ramirez, A. Wahid, L. Downs, A. Wong, J. Lee, I. Mordatch, and J. Tompson, "Implicit behavioral cloning," in *Conference on Robot Learning*. PMLR, 2022, pp. 158–168.
- [9] J. D. A. Ferrandis, J. Moura, and S. Vijayakumar, "Nonprehensile planar manipulation through reinforcement learning with multimodal categorical exploration," *arXiv preprint arXiv:2308.02459*, 2023.
- [10] S. Kumra, S. Joshi, and F. Sahin, "Antipodal robotic grasping using generative residual convolutional neural network," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2020, pp. 9626–9633.
- [11] D. Morrison, P. Corke, and J. Leitner, "Closing the loop for robotic grasping: A real-time, generative grasp synthesis approach," *arXiv preprint arXiv:1804.05172*, 2018.
- [12] H. Kasaei and M. Kasaei, "Mvgrasp: Real-time multi-view 3d object grasping in highly cluttered environments," *Robotics and Autonomous Systems*, vol. 160, p. 104313, 2023.
- [13] R. Newbury, M. Gu, L. Chumbley, A. Mousavian, C. Eppner, J. Leitner, J. Bohg, A. Morales, T. Asfour, D. Kragic *et al.*, "Deep learning approaches to grasp synthesis: A review," *IEEE Transactions on Robotics*, 2023.
- [14] H. Zhang, H. Liang, L. Cong, J. Lyu, L. Zeng, P. Feng, and J. Zhang, "Reinforcement learning based pushing and grasping objects from ungraspable poses," *arXiv preprint arXiv:2302.13328*, 2023.
- [15] G. Zuo, J. Tong, Z. Wang, and D. Gong, "A graph-based deep reinforcement learning approach to grasping fully occluded objects," *Cognitive Computation*, vol. 15, no. 1, pp. 36–49, 2023.
- [16] J.-S. Hu, M.-C. Chien, Y.-J. Chang, S.-H. Su, and C.-Y. Kai, "A ball-throwing robot with visual feedback," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2010, pp. 2511–2512.
- [17] Y. Gai, Y. Kobayashi, Y. Hoshino, and T. Emaru, "Motion control of a ball throwing robot with a flexible robotic arm," *International Journal of Computer and Information Engineering*, vol. 7, no. 7, pp. 937–945, 2013.
- [18] D. M. Lofaro, R. Ellenberg, P. Oh, and J.-H. Oh, "Humanoid throwing: Design of collision-free trajectories with sparse reachable maps," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2012, pp. 1519–1524.
- [19] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, "Tossing-bot: Learning to throw arbitrary objects with residual physics," *IEEE Transactions on Robotics*, vol. 36, no. 4, pp. 1307–1319, 2020.
- [20] H. Kasaei and M. Kasaei, "Throwing objects into a moving basket while avoiding obstacles," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 3051–3057.
- [21] S. H. Kasaei, N. Shafii, L. S. Lopes, and A. M. Tomé, "Interactive open-ended object, affordance and grasp learning for robotic manipulation," in *2019 International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 3747–3753.
- [22] H. Kasaei, S. Luo, R. Sasso, and M. Kasaei, "Simultaneous multi-view object recognition and grasping in open-ended domains," *arXiv preprint arXiv:2106.01866*, 2021.
- [23] S. H. Kasaei, M. Oliveira, G. H. Lim, L. S. Lopes, and A. M. Tomé, "Towards lifelong assistive robotics: A tight coupling between object perception and manipulation," *Neurocomputing*, vol. 291, pp. 151–166, 2018.
- [24] T. P. Lillicrap, J. J. Hunt, A. Pritzel, N. Heess, T. Erez, Y. Tassa, D. Silver, and D. Wierstra, "Continuous control with deep reinforcement learning," *arXiv preprint arXiv:1509.02971*, 2015.
- [25] D. Silver, G. Lever, N. Heess, T. Degris, D. Wierstra, and M. Riedmiller, "Deterministic policy gradient algorithms," in *International conference on machine learning*. PMLR, 2014, pp. 387–395.
- [26] T. Haarnoja, A. Zhou, P. Abbeel, and S. Levine, "Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor," in *International conference on machine learning*. PMLR, 2018, pp. 1861–1870.
- [27] T. Haarnoja, A. Zhou, K. Hartikainen, G. Tucker, S. Ha, J. Tan, V. Kumar, H. Zhu, A. Gupta, P. Abbeel *et al.*, "Soft actor-critic algorithms and applications," *arXiv preprint arXiv:1812.05905*, 2018.
- [28] A. Raffin, A. Hill, A. Gleave, A. Kanervisto, M. Ernestus, and N. Dormann, "Stable-baselines3: Reliable reinforcement learning implementations," *Journal of Machine Learning Research*, vol. 22, no. 268, pp. 1–8, 2021. [Online]. Available: <http://jmlr.org/papers/v22/20-1364.html>
- [29] S. Luo, H. Kasaei, and L. Schomaker, "Accelerating reinforcement learning for reaching using continuous curriculum learning," in *2020 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2020, pp. 1–8.