Effectiveness of Kinesthetic Sensing in In-Hand Rotation of Objects with an Eccentric Center of Mass

Chanyoung Ahn, Sungwoo Park, and Donghyun Hwang

***Abstract*— In-hand manipulation is a key capability for dexterous control, yet it becomes challenging when the mass or center of mass (CoM) of an object is not well known. Such intrinsic properties are difficult to infer precisely through visual sensing alone, which limits the reliability of manipulation strategies. This study investigates how kinesthetic sensing can support in-hand rotational tasks by enabling reinforcement learning (RL) agents to adjust to variations in object dynamics, particularly weight and center of mass. Our method incorporates both proprioceptive signals, such as joint angles, and kinesthetic data, joint forces and torques, captured from sensors embedded in a four-finger robotic hand. To reduce the dimensionality of the input space while retaining the relevant dynamics, we applied Principal Component Analysis (PCA). The resulting policy demonstrates improved adaptability. In the simulation, the manipulation success rates increased by 2.09 and 2.40 times on six and twelve previously unseen CoM configurations, respectively. In addition, kinesthetic detection improves performance 1.52 times in ten known configurations. These findings indicate that kinesthetic feedback contributes substantially to robust and generalizable in-hand manipulation. To access our data and video, please visit our project page: `https://cold-young.github.io/kinesthetic_rotation/`.**

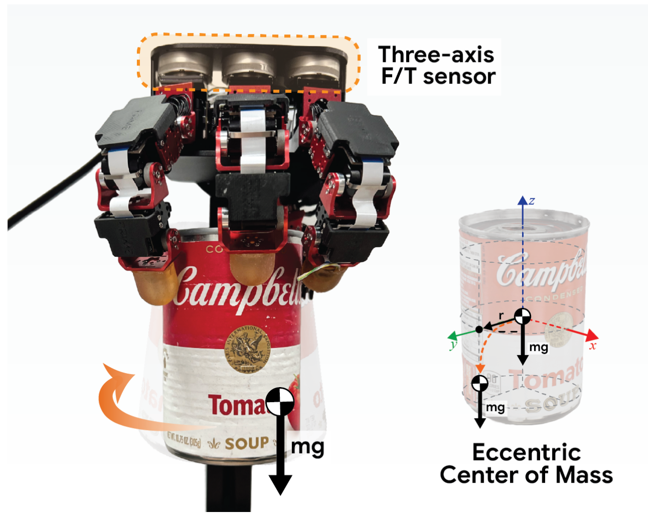

Fig. 1: In-hand rotation tasks of objects with varying weight and an eccentric CoM. The robotic hand manipulates cylindrical objects and estimates center of mass using kinesthetic feedback from force/torque (F/T) sensors embedded in the fingers.

I. Introduction

Robotic hands, having higher degrees of freedom than typical grippers, can perform sophisticated in-hand manipulation tasks such as precisely reorienting objects and adaptively grasping them in complex environments [1], [2]. To automate these tasks, recent studies increasingly utilize reinforcement learning (RL) and imitation learning (IL), along with sensory inputs such as vision and tactile feedback [3], [4], [5]. Despite these efforts, reliably manipulating diverse objects remains challenging, particularly due to variations in intrinsic properties [1].

Relying solely on visual sensing often leads to poor performance, especially when handling internal object properties like an eccentric CoM configuration. To address this limitation, we propose using proprioceptive and kinesthetic feedback, enabling the robot to physically perceive and adaptively respond to object variations [6], [7]. Although prior work has briefly explored kinesthetic feedback for object recognition and limited manipulation tasks [6], our research specifically investigates its effectiveness for adapting to changes in CoM during in-hand rotation.

All authors are with the Center for Humanoid Research, KIST, Seoul 02792, South Korea, Sungwoo Park is also with the Department of Electrical and Computer Engineering, Korea University, Seoul 02841, South Korea. ({`chanyoung.ahn`, `sungwoo.park`, `donghyun`}`@kist.re.kr`).

To effectively handle the high-dimensional nature of sensory data, we apply Principal Component Analysis (PCA) for state representation, motivated by previous successful studies [8], [9]. In this study, we systematically analyze how kinesthetic feedback affects policy learning and performance in RL-based in-hand rotation tasks with objects having varying intrinsic properties.

In this study, we leverage kinesthetic feedback to facilitate in-hand object rotation, particularly for objects with varying weight and an eccentric CoM. We analyze the role of kinesthetic feedback in RL-based approaches of in-hand rotation tasks, demonstrating how it influences task performance and policy learning under changing object properties.

The main contributions are as follows.

- We demonstrate how kinesthetic feedback enables perception and adaptation to intrinsic object properties, such as weight and CoM, for in-hand rotation tasks.
- We evaluate in-hand rotation performance using kinesthetic feedback under 15 unseen weight and CoM conditions to assess adaptability.
- We analyze the effectiveness of state representation techniques in improving RL performance and increasing success rates in manipulation tasks.

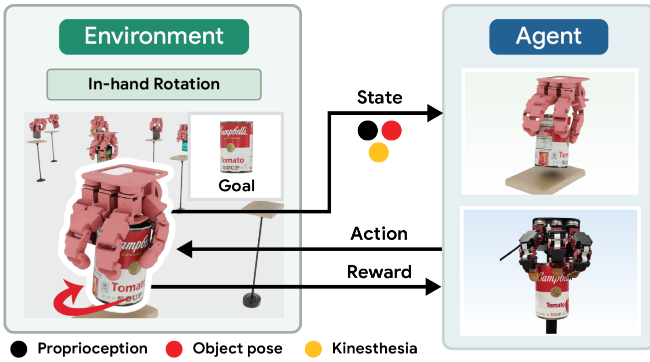

Fig. 2: Overall learning system architecture. To achieve adaptive in-hand rotation of cylindrical objects with an eccentric CoM, the robotic agent receives proprioceptive data, object pose, kinesthetic feedback, and a specified goal orientation. Based on these inputs, the agent outputs a 16-dimensional action corresponding to target joint angles. Further details are provided in Section II.

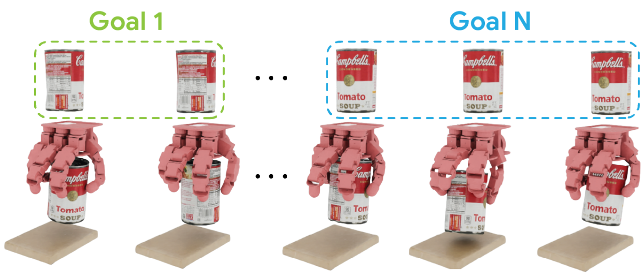

Fig. 3: In-hand rotation with eccentric objects. During inference, the robotic hand rotates cylindrical objects toward the goal orientation within a rollout period of 5 seconds. The agent receives higher rewards for achieving more successful goal orientations.

II. Kinesthetic Sensing for In-Hand Rotation with Eccentric CoM

This study aims to validate the role of kinesthetic sensing, derived from F/T feedback from sensors located at the base joints of the fingers, in in-hand rotation tasks, as illustrated in the concept image in Fig. 1. An overview of our RL framework is illustrated in Fig. 2. To achieve this, we design the framework to acquire downward-facing rotation skills, as shown in Fig. 3.

Our approach leverages privileged observations, incorporating object pose data, proprioceptive signals, and kinesthetic feedback from four force/torque (F/T) sensors embedded in the finger joints. These privileged observations improve state estimation accuracy and facilitate effective policy learning, enabling adaptation to real-world scenarios.

A. Problem Formulation

We define the in-hand rotation problem as a Markov Decision Process, denoted as $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{R}, \mathcal{P})$, where $\mathcal{S}$ represents the state space, $\mathcal{A}$ the action space, $\mathcal{R}$ the reward function, and $\mathcal{P}$ the transition dynamics. Since $\mathcal{R}$ and $\mathcal{P}$ are

TABLE I: **Observations.** The three-axis kinesthetic observations are collected at three time steps. The sensors are located at the joints near the base of the fingers, as indicated by the orange dotted circles in Fig. 5 (b).

Joint angles	q_t
Object pose	(x_t, R_t)
Goal orientation	R_{goal}
Delta rotation	$R_{goal}^{-1} R_t$
Fingertip pose	$(x_t^{finger}, R_t^{finger})$
Previous target	$\bar{q}_t$
F/T sensing	$(\Delta o_t, \Delta o_{t-1}, \Delta o_{t-2})$

unknown to the robot, The robotic hand agent perceives a state s_t at each step t and generates an action $a_t = \pi(s_t)$ based on the current policy π. The agent then receives a reward $r_t = R(s_t, a_t, s_{t+1})$. Its objective is to maximize the discount return $\sum_{t=0}^{T} \gamma^t r_t$, where γ is the discount factor. The elements of this MDP as defined as follows.

1) State: the state of the system consists of the joint position of the robotic hand platform $q_t \in \mathbb{R}^{16}$ with difference of force feedback sensor as observation $o_t \in \{-1.5, 1.5\}^{12}$. and the previous position target $\bar{q}_t \in \mathbb{R}^{16}$. Since the F/T sensing data at one step may not be sufficient for control, we also stack it with other three historical states as the input when we use a multilayer perception (MLP) as the policy network. The detail observation o_t passed to the policy is shown in Table I.

2) Action: at each timestep, the policy network outputs a relative command $a_t \in \mathbb{R}^{16}$, which the PD controller uses to adjust the hand's joint positions, updating the target as $\tilde{q}_{t+1} = \tilde{q}_t + a_t$. To ensure smooth finger movements and prevent conflicts between consecutive actions, we apply an exponential moving average for target updates: $\tilde{q}_{t+1} = \tilde{q}_t + a_t$, where $\tilde{a}_t = \eta a_t + (1-\eta)\tilde{a}_{t-1}$, $t \geq 1$ and $\tilde{a}_0 = 0$. Our experiments show that setting $\eta = 0.035$ provides stable performance. The PD controller runs at a control frequency of 30 Hz.

3) Reward: we design a reward function that is able to rotate the object in a smooth and transferable way. The reward function used in this work is a weighted mixture of several components.

$$r = \omega_1 r_{rot} + \omega_2 r_{fall} + \omega_3 r_{cont} + \omega_4 r_{vel} + \omega_5 r_{dist} + \omega_6 r_{goal}. \tag{1}$$

The reward function comprises six components, each designed to guide the agent's behavior toward successful in-hand rotation. The first term, r_{rot}, provides a positive signal based on the alignment between the object's current orientation and the target pose. The second component, r_{fall}, imposes a penalty when the object falls from the table, as illustrated in Fig. 3. The third term, r_{cont}, discourages contact between the object and the table surface, promoting free-space rotation during rollouts. The fourth reward, r_{vel}, penalizes overly fast rotations in simulation, encouraging

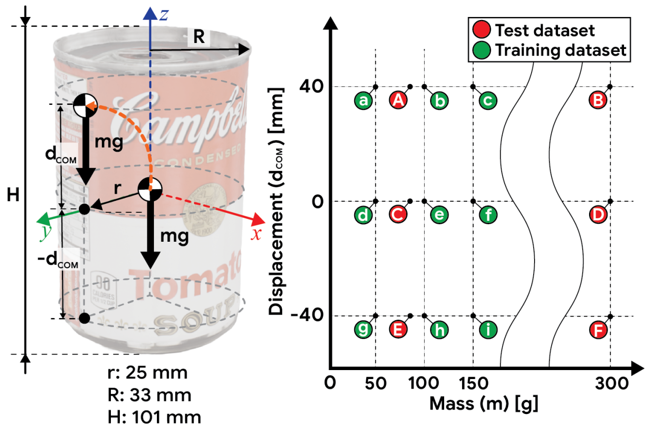

Fig. 4: Training and test objects with eccentric CoM dataset. *(Left)* Illustration of the cylindrical object, where we vary intrinsic properties such as CoM position and mass. *(Right)* The dataset consists of nine training objects (green circles) with three different mass values: 50 g, 100 g, and 150 g. The six test objects (red circles) have unknown intrinsic attributes, particularly mass, with three objects weighing 80 g and the other three weighing 300 g.

smoother trajectories that can be reliably transferred to the real world. The fifth component, r_{dist}, offers a positive reward based on the normalized distance between the current and goal poses, increasing as the object approaches the target. Finally, r_{goal} is a sparse reward granted when the agent successfully completes the rotation task.

III. Experimental Setup

s

A. Experimental Setup

Our goal is to evaluate the influence of kinesthetic sensing on in-hand rotation tasks. To achieve this, we design a learning environment for a downward-facing rotation task using Isaac Lab [10]. We use 25 cylindrical objects with varying mass and CoM positions to investigate the role of kinesthetic sensing in in-hand manipulation. To analyze its effectiveness, we train three policies with different observation settings: (1) a baseline policy using only privileged information, such as object pose and robot joint angles; (2) a policy that incorporates privileged information along with kinesthetic sensing differences over three timesteps; and (3) a policy similar to the second but utilizing PCA as an encoder to represent kinesthetic data in a lower-dimensional space while preserving meaningful correlations.

Hardware Setup. We use a four-fingered robotic hand platform [11]. The robotic hand platform is capable of kinesthetic feedback using a three-axis F/T sensor on each finger, as shown in Fig. 5. It weighs approximately 550 g and has a maximum payload capacity of around 3 kg. Each joint of the robotic hand is controlled by a PD position controller with a control frequency of 120 Hz. The target position commands are converted to torque using a PD controller (K_p = 1.0, K_d = 0.1). Each episode consists of 150 control steps, corresponding to a duration of 5 senconds.

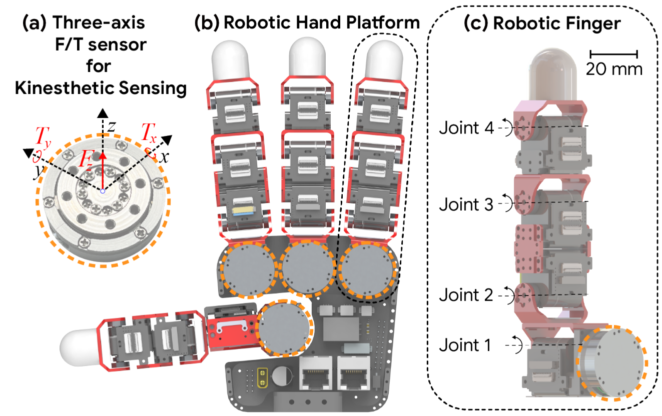

Fig. 5: Illustration of robotic hand platform capable of kinesthetic feedback at each finger. (a), (b), and (c) show the three-axis F/T sensor embedded into the robotic hand platform for kinesthetic feedback, the robotic hand platform, and the robotic finger, respectively.

B. Training Procedure

We train our control policy using the proximal policy optimization (PPO) algorithm [12] with a multilayer perceptron (MLP) for both the policy and value networks. For simulation, we use Isaac Sim with Isaac Lab [10], setting the timestep to $dt = 1/120$ s with four simulation sub-steps. The simulation runs 4096 parallel environments, and the policy executes actions every four steps, corresponding to a 30 Hz control frequency.

Policies are trained for 40K steps using the skrl RL library [13]. All results are obtained using five random seeds, tested across 500 rollouts with 100 parallel environments. All experiments are conducted on single RTX 4090 GPU.

C. Baselines

We evaluate three baselines to assess the role of kinesthetic sensing in learning in-hand manipulation skills. All baselines share the same environment, goal, and reward settings, differing only in their observation modalities.

Proprioception. The policy uses only privileged observations, including robot joint angles, object pose, and target goal. The total state dimension is 75.

Proprioception+Kinesthesia. This policy extends the first baseline by incorporating kinesthetic sensing. Kinesthetic feedback consists of three-axis F/T information from the robotic hand's phalanges. To emphasize dynamic changes, we use the difference between the current and previous kinesthetic readings and apply a three-step time stack. This approach helps capture transient signals such as slip and contact by leveraging correlations in sensory data. The total state dimension is 111.

Proprioception+Kinesthesia with PCA. This policy builds upon the second baseline but incorporates PCA as an encoder to compress kinesthetic sensing data. By reducing

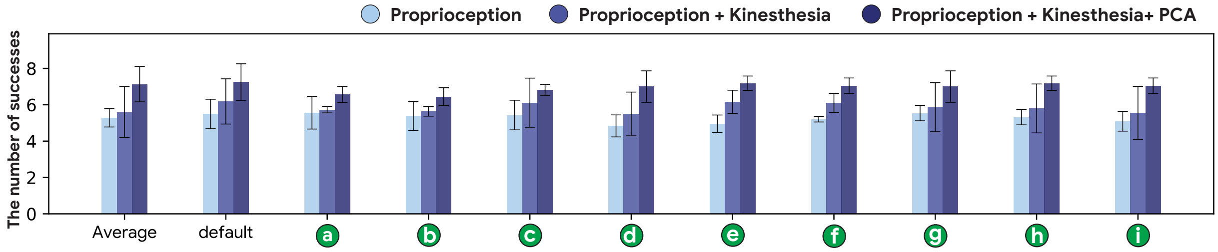

Fig. 6: Performance Evaluation on Pre-trained Samples. We evaluate our policies on 10 objects with varying mass and CoM positions, measuring the number of successful rotations to the target orientation over 500 inferences. We test five policy seeds trained on these 10 objects. The results indicate that incorporating kinesthetic sensing improves performance compared to policies without it. Notably, using PCA to represent and compress kinesthetic data further enhances performance, achieving even higher performance.

dimensionality while preserving key correlations, we aim to demonstrate that structuring sensory data enhances manipulation performance by capturing intrinsic object properties such as mass and CoM. The total state dimension is 93, as PCA reduces the force feedback data from 36 to 18 dimensions.

D. Performance Evaluation on Pre-trained Samples

First, we evaluate three baseline policies using ten objects that were included in the training process to compare their performance. We train the policies with five random seeds and 40K steps. As shown in Fig. 6, in the evaluation of trained objects, the policy incorporating kinesthetic sensing outperforms the policy without sensor input. This suggests that kinesthetic sensing enhances performance by providing additional state information, similar to real-world scenarios where only proprioceptive and F/T data from the robotic hand are available.

When combining PCA with kinesthetic data, we observe an overall improvement in performance. However, for objects already encountered during training, the difference in performance between this approach and other baselines is not substantial. In particular, the Proprioception + Kinesthesia with PCA configuration increased the number of successful manipulations by 1.52-fold improvement compared to using proprioception alone.

E. Performance Evaluation on Novel Samples

We evaluate the same three policies on two datasets of previously unseen objects. These objects share the same geometric shape but differ in intrinsic properties such as mass and CoM. The policies are tested under conditions where the objects have different weights—80 g and 300 g—compared to the training set.

A key observation is that kinesthetic sensing becomes increasingly beneficial as the object's weight increases. For the 80 g objects (A, C, and E in Fig. 4 *(right)*), all three baselines perform relatively well, even though the objects are novel. However, when manipulating the 300 g objects (B, D, and F in Fig. 4 *(right)*), there is a sharp drop in performance for all policies except the one that incorporates PCA with kinesthetic data. This likely occurs because, as object weight increases, correctly estimating the CoM becomes crucial for

TABLE II: Results of the In-Hand Rotation Experiment: The number of successful manipulations within a fixed time limit (5 seconds) for 10 pre-trained and 15 novel samples.

	Max Reward (mean)	Pre-trained Samples Fig. 6	Unknown Mass Fig. 8 (a)	Unknown CoM Positions Fig. 8 (b)
Proprioception	419.05 ± 92.41	4.47 ± 1.14	3.29 ± 0.95	2.90 ± 0.77
Proprioception + Kinesthesia	477.30 ± 143.31	5.01 ± 1.75	3.73 ± 1.26	3.35 ± 1.19
Proprioception + Kinesthesia with PCA	**569.71 ± 104.59**	**6.78 ± 1.24**	**6.87 ± 1.35**	**6.95 ± 1.36**

effective manipulation. Unlike other approaches, the PCA-based representation accounts for sensor correlations, rather than just including raw kinesthetic data, thereby mitigating performance degradation. The Proprioception + Kinesthesia with PCA configuration improved the number of successful manipulations by 2.09 times compared to proprioception alone.

In addition, we conducted experiments with 12 novel CoM configurations to examine performance variations when the CoM changes while keeping the object's weight constant at 100 g (as in pre-training). The absolute radius r remained fixed (Fig. 4 *(left)*), but the CoM configuration was systematically varied along different axes, from $[-r, 0, -d_{com}]$ to $[r, 0, d_{com}]$, and from $[0, -r, -d_{com}]$ to $[0, r, d_{com}]$, resulting in 12 different configurations.

IV. Results & Conclusion

Our experiments demonstrate that kinesthetic sensing significantly enhances in-hand manipulation performance, particularly when handling objects with varying mass and CoM. On pre-trained objects, the Proprioception + Kinesthesia policy outperformed the proprioception-only baseline, achieving a 1.52-fold increase in successful manipulations when combined with PCA-based data encoding as shown in Table II.

Our findings highlight the critical role of kinesthetic sensing in RL-based in-hand manipulation. Kinesthetic feedback enables more stable and adaptable manipulation by providing essential information about an object's dynamic properties.

Furthermore, applying PCA to structure kinesthetic data improves learning efficiency and performance by preserving key correlations while reducing dimensionality.

REFERENCES

[1] A. Billard and D. Kragic, “Trends and challenges in robot manipulation,” *Science*, vol. 364, no. 6446, p. eaat8414, 2019.

[2] T. Feix, J. Romero, C. H. Ek, H.-B. Schmiedmayer, and D. Kragic, “A metric for comparing the anthropomorphic motion capability of artificial hands,” *Transactions on Robotics*, vol. 29, no. 1, pp. 82–93, 2012.

[3] H. Van Hoof, T. Hermans, G. Neumann, and J. Peters, “Learning robot in-hand manipulation with tactile features,” in *2015 IEEE-RAS 15th International Conference on Humanoid Robots (Humanoids)*, pp. 121–127, IEEE, 2015.

[4] H. Qi, A. Kumar, R. Calandra, Y. Ma, and J. Malik, “In-hand object rotation via rapid motor adaptation,” in *Conference on Robot Learning (CoRL)*, Proceedings of Machine Learning Research, pp. 1722–1732, PMLR, 14–18 Dec 2023.

[5] Z.-H. Yin, B. Huang, Y. Qin, Q. Chen, and X. Wang, “Rotating without seeing: Towards in-hand dexterity through touch,” *Proceedings of Robotics: Science and Systems (RSS)*, 2023.

[6] J. Arolovitch, O. Azulay, and A. Sintov, “Kinesthetic-based in-hand object recognition with an underactuated robotic hand,” in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pp. 18179–18185, IEEE, 2024.

[7] M. Emami, A. Bayat, R. Tafazolli, and A. Quddus, “A survey on haptics: Communication, sensing and feedback,” *IEEE Communications Surveys & Tutorials*, 2024.

[8] O. M. Andrychowicz, B. Baker, M. Chociej, R. Jozefowicz, B. McGrew, J. Pachocki, A. Petron, M. Plappert, G. Powell, A. Ray, *et al.*, “Learning dexterous in-hand manipulation,” *International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.

[9] R. Antonova, P. Shi, H. Yin, Z. Weng, and D. Kragic, “Dynamic environments with deformable objects,” in *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.

[10] M. Mittal, C. Yu, Q. Yu, J. Liu, N. Rudin, D. Hoeller, J. L. Yuan, R. Singh, Y. Guo, H. Mazhar, A. Mandlekar, B. Babich, G. State, M. Hutter, and A. Garg, “Orbit: A unified simulation framework for interactive robot learning environments,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 8, no. 6, pp. 3740–3747, 2023.

[11] S. Park and D. Hwang, “Three-axis flat and lightweight force/torque sensor for enhancing kinesthetic sensing capability of robotic hand,” *IEEE Transactions on Industrial Electronics*, vol. 71, no. 10, pp. 12707–12717, 2024.

[12] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.

[13] A. Serrano-Muñoz, D. Chrysostomou, S. Bøgh, and N. Arana-Arexolaleiba, “skrl: Modular and flexible library for reinforcement learning,” *Journal of Machine Learning Research*, vol. 24, no. 254, pp. 1–9, 2023.

[14] Y. Lin, J. Lloyd, A. Church, and N. F. Lepora, “Tactile gym 2.0: Sim-to-real deep reinforcement learning for comparing low-cost high-resolution robot touch,” *IEEE Robotics and Automation Letters (RA-L)*, vol. 7, no. 4, pp. 10754–10761, 2022.

APPENDIX

A. Limitations & Future Work

This research still has several limitations. First, we did not deploy the policy on a real robotic platform due to the sim-to-real gap. Kinesthetic sensing data can differ significantly between simulation and real-world. To partially address this, we used differential features of tactile sensing. However, bridging this gap requires mapping real sensor data to simulation sensing. In future work, we plan to develop a sensing transfer model, following approaches similar to [14]. Additionally, the current policy assumes access to the object’s pose as part of the state. To remove this assumption, we intend to incorporate a vision sensor and a perception model to estimate the object pose directly from visual input.

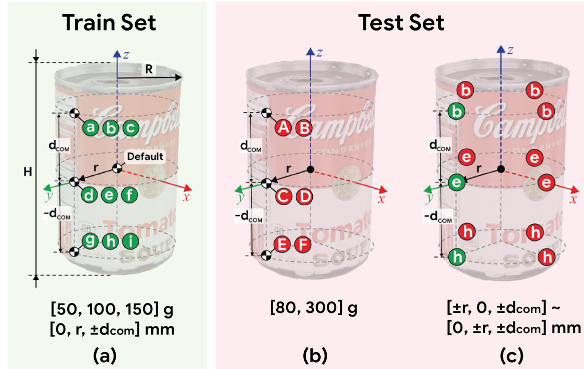

Fig. 7: Test objects with eccentric CoM. (a) Ten training objects with varying center of mass locations; result are shown in Fig. 6. (b) Six test objects with unseen masses (80g and 300g); corresponding results are shown in Fig. 8. (c) Twelve test objects, including three known objects, used to evaluate performance under unseen CoM positions; results are shown in Fig. 8.

Secondly, while we used PCA to encode kinesthetic feedback, we did not explore other representation learning methods like autoencoder-based models. Our focus in this study was to analyze the effectiveness of kinesthetic feedback. However, we are also interested in applying more advanced representation learning techniques. In particular, we observed that PCA is limited in handling multi-modal sensor data, as it assumes linear and orthogonal relationships between modalities. To address this, we aim to develop a multi-modal representation model that can jointly encode kinesthetic feedback, joint angles, joint torques, and other sensing modalities.

B. Performance Evaluation on Novel Samples

We evaluated the three policies on two test sets to validate their generality (see Fig. 7). The results are shown in Table II and Fig. 8.

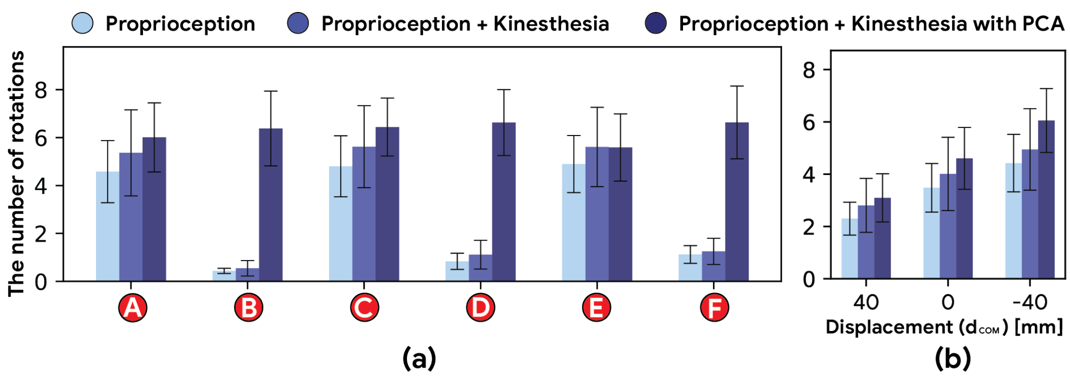

Fig. 8: Simulation Performance of Our Policy with 15 unknown objects. (a) Performance on six novel objects with identical shapes but different intrinsic properties, evaluated at two masses (80 g, 300 g). Performance declines for heavier objects when kinesthetic data is not combined with PCA. (b) Performance on 12 novel objects with different CoM positions. Kinesthetic sensing improves manipulation success, while performance decreases as the CoM shifts higher.