DartBot: Overhand Throwing of Deformable Objects with Tactile Sensing and Reinforcement Learning

Shoaib Aslam*, Krish Kumar*, Pokuang Zhou*, Hongyu Yu, Michael Yu Wang, *Fellow, IEEE*, and Yu She

***Abstract*—Object transfer through throwing is a classic dynamic manipulation task that necessitates precise control and perception capabilities. However, developing dynamic models for unstructured environments using analytical methods presents challenges. In this study, we present DartBot, a robot that integrates tactile exploration and reinforcement learning to achieve robust throwing skills for nonrigid relatively small objects under the influence of moment of inertia which cause the object to spin in the air. Unlike traditional sim-to-real transfer methods, our approach involves direct training of the agent on a real hardware robot equipped with a high-resolution tactile sensor, enabling reinforced learning in a realistic and dynamic environment. By leveraging tactile perception, we incorporate pseudo-embeddings of the physical properties of objects into the learning process through tilting actions at two distinct angles. Furthermore, we demonstrate that the quality of a grasp significantly impacts the success rate of the throwing task. We evaluate the effectiveness of our method through extensive experiments, demonstrating superior performance and generalization capabilities in real-world throwing scenarios. We achieved a success rate of 95% for unseen objects with a mean error of 3.15 cm from the goal.**

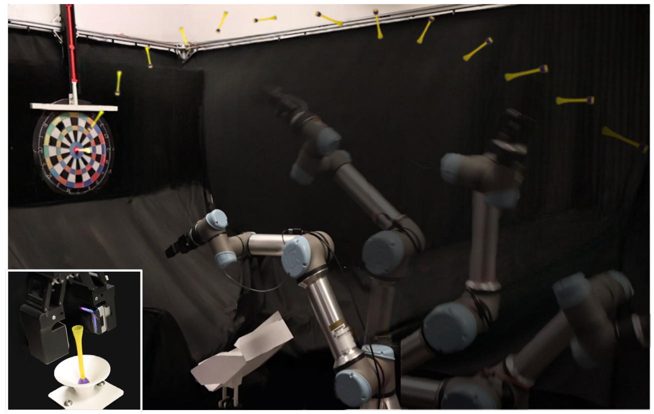

Fig. 1: Throwing object with precision: an example scenario of throwing a nonrigid cylindrical object beyond a robot's workspace by executing an overhand throw trajectory while the object spins in the air due to its moment of inertia. We develop a reinforcement learning-based method using a GelSight Mini tactile sensor, allowing the robot to utilize tactile feedback to perceive object properties and determine optimal throw parameters. The bottom left image shows the gripper with the tactile sensor grasping the nonrigid object's deformable shaft.

I. Introduction

Throwing manipulation is a fundamental human motor skill [1] with applications in sports and industrial automation [2]. It enables robots to transport objects beyond their kinematic limits by leveraging dynamic extrinsic dexterity. However, achieving accurate object throws is challenging due to multiple factors, including mass, center-of-mass, friction, softness, shape, and aerodynamics. The difficulty increases with non-rigid objects, which are common in daily life, making precise control and manipulation crucial for successful throwing tasks.

Despite growing research on throwing manipulation [3], [4], most efforts focus on underhand throws of rigid objects. To our knowledge, no prior work has explored overhand throwing tasks. Moreover, integrating tactile perception and reinforcement learning (RL) for enhancing robotic throwing remains underexplored.

∗ Authors with equal contribution.

Shoaib Aslam, Pokuang Zhou, and Yu She are with the School of Industrial Engineering, Purdue University, West Lafayette, USA (email: saslam@purdue.edu; zhou1458@purdue.edu; shey@purdue.edu). Shoaib Aslam is also with the Department of Mechanical and Aerosapce Engineering, The Hong Kong University of Science and Technology, Hong Kong (email: saslamaa@connect.ust.hk). *(Corresponding author: Yu She)*

Krish Kumar is with the Department of Computer Engineering, Purdue University, West Lafayette, USA (email: kumar585@purdue.edu).

Hongyu Yu is with the Department of Mechanical and Aerosapce Engineering, The Hong Kong University of Science and Technology, Hong Kong (email: hongyuyu@ust.hk).

Michael Yu Wang is with the School of Engineering, Great Bay University, Dongguan, China (email: mywang@gbu.edu.cn).

Unlike vision-based methods, which have shown promise in dynamic object representation [5], our approach exclusively utilizes tactile feedback to incorporate object-centric information. Vision-based methods remain constrained to structured environments and face challenges in real-world deployment, whereas tactile sensing provides a robust alternative for object interaction.

To the best of our knowledge, this study is the first to address high-speed overhand throwing manipulation on a real robot by integrating tactile sensing and reinforcement learning. We propose **TT-RL**, an RL framework that leverages high-resolution tactile feedback to develop precise throwing skills. Our method enables the transfer of small nonrigid objects beyond the robot's workspace, achieving high accuracy with mean errors of 2.20 cm for seen objects and 3.15 cm for unseen objects.

II. Methodology

The aim of DartBot is to make a robot arm capable of throwing arbitrary nonrigid objects, under the influence of moment of inertia, to a target location with a single subsequent exploratory action, illustrated in Fig. 1.

1) Preliminaries: We aim for the robot to complete the task by first grasping the dart and obtaining its physical properties through tactile sensors using two different grasping poses. Upon reaching the designated pose, the gripper's opening width is adjusted empirically, enabling in-hand manipulation of the dart. This imparts an initial rotational speed, allowing

TABLE I: Experiment results of learned robot overhand throwing skills with generalization to varying robot-to-target distance (mean score 0.79, mean distance from goal 3.15cm).

Object #	**Distance between target and robot (cm)**											
	155		**162**		**170**		**173**		**176**		**182**	
	Score	Success times	Score	Success times	Score	Success times	Score	Success times	Score	Success times	Score	Success times
1	0.80	18/20	0.77	18/20	0.89	20/20	0.86	20/20	0.82	20/20	0.84	19/20
2	0.72	16/20	0.71	17/20	0.83	18/20	0.81	18/20	0.76	19/20	0.78	19/20
3	0.77	17/20	0.75	19/20	0.85	20/20	0.84	20/20	0.83	20/20	0.83	20/20
4	0.76	18/20	0.74	17/20	0.86	20/20	0.83	19/20	0.81	20/20	0.82	20/20
5	0.71	16/20	0.69	16/20	0.78	17/20	0.77	19/20	0.72	19/20	0.75	18/20

it to travel further, as shown in Fig. 3. The camera records the dart's position only when it hits the target at the correct angle. We employ the TD3 method for online reinforcement learning training, where the MDP process consists of a series of dart-throwing actions, as shown in Fig. 4. The final reward is defined by the dart's landing position, while the actions include the end-effector's grasping motion, the robot base joint angles, and the timing of the dart's release by the gripper. These actions influence whether the dart hits the target with the correct pose, its flight distance, and its final landing position.

2) Experiments and benchmarks: Our experiment utilizes online reinforcement learning, with a platform enabling the robot to cycle through grasping, throwing, and retrieving the dart, improving training efficiency. We tested various darts and different distances between the robot and the target. As shown in Figure 2, by replacing the three modular components, we obtained darts with varying lengths, centers of mass, and weights. These darts have a magnetic head and a deformable mesh tail. We also evaluated the performance of other grippers. Common UR5 robotic grippers, equipped with force sensors, failed the task due to the minimal gripping force required for the dart. Additionally, since the mass variation among different darts is small (only 0.3g–0.5g), we compared a pre-trained ResNeXt network in the CNN method with our CNN-based improved approach. Our method increased classification accuracy by more than 10%.

3) Results: Experiments show that our method achieves excellent results compared to several mainstream approaches, as shown in Table I. The five template objects alongside the six distinct robot locations relative to the dartboard yield a total of 30 cases. The 24 unseen scenarios assess the learned policy's performance. The average success rate is 95%, with a mean distance from the goal of 3.15 cm.

III. Conclusion

We developed and implemented the TT-RL framework to train a real robot in overhand throwing, focusing on small objects with a deformable shaft and rigid head. The key challenge was achieving precise throws despite unpredictable post-release spins and variations in object properties. Using high-resolution tactile sensing, our method demonstrated strong generalization across different objects and target distances, highlighting its real-world potential. Currently, TT-RL focuses on dart-throwing to study overhand dynamics, laying the foundation for broader applications.

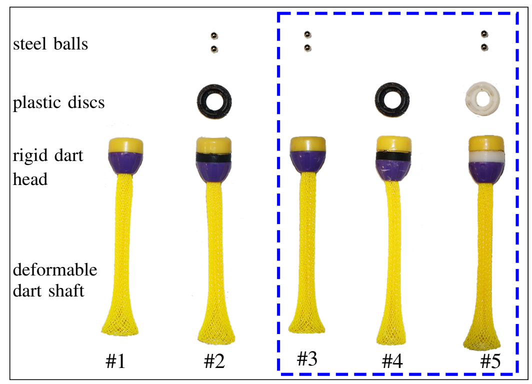

Fig. 2: For our experiments, we used template objects with three parts: a rigid purple head, a magnetic face with a magnet in the yellow top, and a flexible yellow mesh shaft. The objects within the blue dashed box are unseen to the agent.

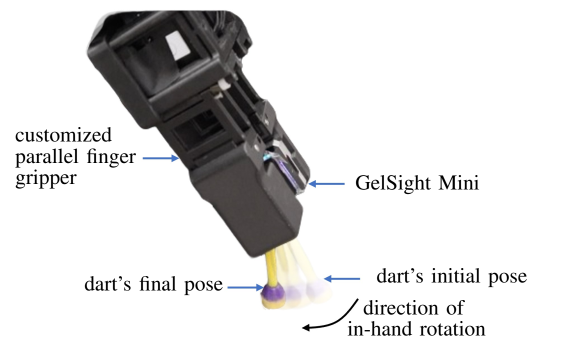

Fig. 3: The robot grasps the dart with a two-fingered gripper and a GelSight Mini sensor, maintaining the hold from pickup to the throwing pose. The gripper adjusts its width to manipulate the dart, allowing it to rotate left under gravity.

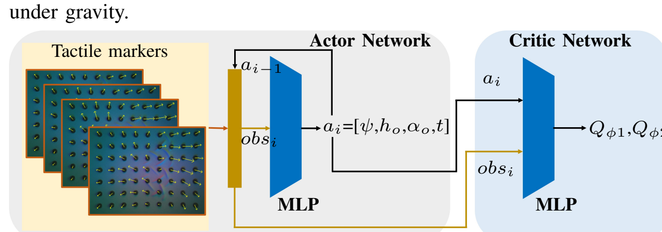

Fig. 4: The Actor and Critic neural network models in Twin Delayed Deep Deterministic Policy Gradient (TD3). The actor network takes a 24×1 *obs* vector consisting of tactile features obtained after two tilt motions and the previous timestep's action space parameters as input, producing action space elements as output. The critic network takes the *obs* vector and action space as inputs and outputs learned policies Q_{ϕ_1} and Q_{ϕ_2}.

REFERENCES

[1] K. John E, "Throwing, the shoulder, and human evolution," *American journal of orthopedics (Belle Mead, N.J.)*, vol. 45, pp. 110–114, 2016.

[2] F. Raptopoulos, M. Koskinopoulou, and M. Maniadakis, "Robotic pick-and-toss facilitates urban waste sorting," in *2020 IEEE 16th International Conference on Automation Science and Engineering (CASE)*, pp. 1149–1154, 2020.

[3] A. Zeng, S. Song, J. Lee, A. Rodriguez, and T. Funkhouser, "Tossingbot: Learning to throw arbitrary objects with residual physics," *IEEE Transactions on Robotics*, vol. 36, no. 4, pp. 1307–1319, 2020.

[4] M. Monastirsky, O. Azulay, and A. Sintov, "Learning to throw with a handful of samples using decision transformers," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 576–583, 2023.

[5] Z. Xu, J. Wu, A. Zeng, J. B. Tenenbaum, and S. Song, "Densephysnet: Learning dense physical object representations via multi-step dynamic interactions," 2019.