

Force-Aware Robotic Manipulation via Haptic-Guided Inverse Reinforcement Learning

Do-Gyeong Yuk

*Dept. of Aeronautics, Mechanical and Electronic Convergence
Engineering, Kumoh National Institute of Technology
Gumi, Republic of Korea
dgyuk@kumoh.ac.kr*

Jung Woo Sohn

*School of Mechanical Engineering
Kumoh National Institute of Technology
Gumi, Republic of Korea
jwsohn@kumoh.ac.kr*

Abstract—Robotic manipulation in unstructured environments requires not only precise motion control but also fine regulation of contact forces. Excessive force can break fragile objects, making reward design for reinforcement learning difficult. To address this, we propose a simulation and learning framework that integrates haptic feedback with adversarial inverse reinforcement learning. Human demonstrations are collected through haptic teleoperation, where operators control the manipulator and perceive contact forces. These demonstrations encode implicit strategies for safe force regulation. Joint torque-based force estimates are incorporated into AIRL states, enabling policies that balance accuracy with force control. Simple breakage conditions in simulation model fragile interactions. The framework allows agents to acquire robust, force-aware manipulation strategies generalizing to delicate and uncertain dynamics.

Keywords—Inverse Reinforcement Learning, Haptic Feedback, Robot Manipulation, Contact-rich Task

I. INTRODUCTION

Robotic manipulation in unstructured environments requires far more than simple position accuracy; it demands adaptive regulation of contact forces under uncertainty. Unlike structured industrial scenarios where objects are rigid, geometries are fixed, and trajectories can be pre-programmed, unstructured settings involve fragile, deformable, or variably constrained objects that can easily be damaged or mishandled if the robot applies inappropriate force. For example, inserting a delicate component, manipulating soft biological tissue, or writing with chalk on a surface all require the robot not only to reach a desired position but also to continuously adapt applied forces at fine granularity. Failure to regulate these forces can result in breakage, slippage, or task failure, even when positional accuracy is high. Consequently, developing policies that balance force precision and trajectory execution has emerged as a critical challenge in robotic learning. However, manually designing reinforcement learning (RL) reward functions that capture this balance remains extremely difficult: too much emphasis on trajectory can encourage brittle strategies that ignore contact safety, while excessive focus on force regulation may hinder task completion. This makes force-adaptive policy learning in unstructured

environments a central open problem for advancing robust robotic autonomy.

Several recent works have investigated aspects of integrating haptic signals into robot learning and control, yet important gaps remain. In Haptic-ACT (Haptic Attention over Contact Trajectories), force-torque signals are treated as attention cues to highlight salient phases of manipulation and to guide imitation learning [1]. While this approach demonstrates that haptic cues can improve generalization across tasks, it remains centered on augmenting demonstration efficiency rather than enabling robots to autonomously acquire novel contact strategies. In contrast, Learning Variable Impedance Control via IRL successfully recovers both impedance policies and latent reward functions from expert demonstrations, enabling robots to adaptively adjust stiffness and damping during contact [2]. Nevertheless, these methods are typically task-specific, optimized for one manipulation skill (e.g., surface wiping or peg insertion), and therefore struggle to generalize across the diverse and unpredictable contact conditions encountered in unstructured environments. Furthermore, research such as TACTO-based tactile representation learning employs vision-based tactile sensors to infer latent states and enhance robot perception in contact-rich scenarios [3]. While such work advances state estimation under partial observability, it primarily improves perception rather than teaching the robot how to strategically adapt its forces during execution.

Despite these advances, none of these approaches fully leverage haptic feedback as a primary learning signal for autonomous policy acquisition in unstructured environments. Teleoperation methods tend to prioritize assisting human operators rather than enabling robots to learn from the feedback itself. IRL applications for force control often remain constrained to pre-specified tasks. And tactile representation learning focuses on perception instead of strategic force adaptation. This study aims to close this gap by integrating haptic feedback directly into the learning loop via Adversarial Inverse Reinforcement Learning (AIRL). Our goal is to enable a robot to autonomously develop generalizable force regulation strategies in environments where visual or positional cues are insufficient, such as unpredictable or fragile object interactions. Unlike prior work, which emphasizes human-centered assistance or narrow imitation, our approach promotes autonomous robotic learning of robust contact strategies through skilled demonstration, enabling better adaptability to real-world variability

This work was supported by a National Research Foundation of Korea (NRF) grant funded by the Korean Government (Ministry of Science and ICT, MSIT) (No. RS-2023-NR076925, 50%) and in part by the Ministry of Science and ICT (MSIT), South Korea, through the Information Technology Research Center (ITRC) Support Program under Grant (IITP-2025-RS-2024-00437190, 50%) supervised by the Institute for Information and Communications Technology Planning and Evaluation (IITP).

II. METHODOLOGY

A. System Setup

The experimental platform (Fig. 1) consists of a UR3e robotic manipulator equipped with a Robotiq 2F-85 parallel gripper. A piece of chalk is rigidly grasped by the gripper and interacts directly with a virtual blackboard, where the chalk is modeled as a rigid body with a simple breakage mechanism triggered by a force–torque threshold. The simulation environment was constructed using Isaac Sim, which provides high-fidelity physics and contact dynamics, while the learning framework was implemented with Isaac Lab to enable inverse reinforcement learning.

In addition, a Force Dimension Omega.6 haptic device (Fig. 2) was used to teleoperate the robot. The user’s hand motion was mapped to the end-effector pose, and contact forces from the robot–blackboard interaction were fed back to the haptic device, allowing the user to feel the forces in real time.




Fig. 1. UR3e manipulator holding a piece of chalk



Fig. 2. Haptic teleoperation setup with the Omega.6 device

B. Force-Aware Learning Framework

During the interaction, external torques at the joints of the UR3e are measured and used as indirect indicators of contact forces at the end-effector. These signals are incorporated into the state representation of the learning framework, enabling the policy to develop force-aware behaviors.

Instead of conventional reinforcement learning, we employ Adversarial Inverse Reinforcement Learning (AIRL)

implemented in Isaac Lab. In this setup, expert demonstrations collected via the Omega.6 haptic device provide trajectories where the human operator controls the end-effector pose and perceives contact forces through haptic feedback. The AIRL framework learns a reward function that captures the implicit strategies of human demonstrations, while simultaneously training a policy that replicates these behaviors. The observation space consists of end-effector position, joint states, and measured external torques, while the action space is defined as incremental target poses of the end-effector. This design encourages smooth and stable chalk–blackboard interaction while penalizing excessive contact forces that may lead to chalk breakage.

III. CONCLUSIONS

This work presented a simulation and learning framework for robot–environment interaction using both force/torque feedback and inverse reinforcement learning. A UR3e robotic manipulator with a Robotiq 2F-85 gripper was teleoperated through a Force Dimension Omega.6 device, enabling intuitive control and real-time haptic feedback during contact with a virtual blackboard. Chalk breakage was incorporated into the environment to realistically model fragile interactions. External joint torques measured during contact were used as informative signals for inverse reinforcement learning, implemented in Isaac Lab. Through this setup, the system enables robots to acquire policies that are guided not only by positional trajectories but also by force-based cues, paving the way for more natural and robust manipulation in tasks involving delicate contact dynamics. In the next step of this work, the possibility of applicability to real environments will be secured through experiments.

ACKNOWLEDGEMENT

This work was supported by the IITP(Institute of Information & Communications Technology Planning & Evaluation)-ITRC(Information Technology Research Center)(IITP-2025- RS-2024-004371905) grant funded by the Korea government(Ministry of Science and ICT).

REFERENCES

- [1] K. Li, S. M. Wagh, N. Sharma, S. Bhadani, W. Chen, C. Liu, and P. Kormushev, “Haptic-ACT: Bridging Human Intuition with Compliant Robotic Manipulation via Immersive VR,” 2024, arXiv:2409.11925. [Online]. Available: <https://arxiv.org/abs/2409.11925>
- [2] X. Zhang, L. Sun, Z. Kuang, and M. Tomizuka, “Learning variable Impedance Control via Inverse Reinforcement Learning for Force-Related Tasks,” *IEEE Robotics and Automation Letters*, vol. 6, No. 2, pp. 2225–2232.
- [3] S. Wang, M. Lambeta, P. W. Chou, and R. Calandra, “TACTO: A Fast, Flexible, and Open-Source Simulator for High-Resolution Vision-Based Tactile Sensors,” *IEEE Robotics and Automation Letters*, vol. 7, No. 2, pp. 3930–3937.