On the Role of Haptic Feedback in Demonstration Collection for Visuomotor Policy Learning

Joanna Peng[1,2*], Aidan Schreder[1,2*], Yiran Liu[1,2*], Manjot Dola[1,2], Benjamin Tran[1,2], Yue Hu[2], Pengcheng Xi[1]
[1]National Research Council Canada [2]University of Waterloo
{j85peng, aschrede, y3275liu}@uwaterloo.ca; pengcheng.xi@nrc-cnrc.gc.ca

Abstract*— Precise physical interaction remains a major challenge for visual imitation learning methods such as diffusion policies. While recent approaches introduce more complex architectures, they often overlook a key limitation during data collection: standard teleoperation lacks haptic feedback, preventing operators from sensing contact. We propose a haptic-fused imitation learning framework that enables force-aware demonstration collection using a wrist-mounted force-torque sensor and a haptic controller. This setup allows operators to perceive contact through real-time force feedback. We conduct a comparative study on contact-rich tasks to evaluate the impact of haptic feedback on demonstration quality and downstream policy performance. Results show that haptic feedback improves force regulation and contact maintenance during data collection. However, this does not directly translate to improved task success in the learned policy, highlighting limitations of task selection and open-loop policy execution. Our findings suggest that closing the haptic loop during data collection is important, but its benefits are task-dependent. Code, video, dataset and model are available at our project page.

I. INTRODUCTION

Contact-rich manipulation tasks require precise and continuous regulation of interaction forces and it remains a fundamental challenge in robotics. Examples include insertion tasks, surface-following operations, and interactions with deformable objects. Unlike pick-and-place, these tasks demand not only spatial accuracy but also real-time adaptation to contact dynamics that are difficult to model.

Imitation learning (IL) provides a promising approach by enabling robots to learn directly from human demonstrations. Recent methods such as Diffusion Policy [1] and Action Chunking Transformer [2] capture the temporal structure of demonstrations and achieve strong performance across manipulation tasks, while reinforcement learning (RL) methods [3], [4], [5] further refine policies through interaction. However, these approaches rely primarily on visual observations, which do not capture interaction forces, limiting their effectiveness in contact-rich scenarios.

A key limitation persists: even when force data are recorded by the robot [6], [7], [8], this information is not fed back to the human operator during data collection. As a result, demonstrations are typically collected in a force-blind manner, with operators relying solely on visual cues to infer contact. Prior efforts [2], [9], [10] provide partial solutions, but real-time haptic feedback, where operators directly feel interaction forces through a dedicated device, remains largely absent from imitation learning pipelines. This is consequential for contact-rich tasks, where humans rely on tactile feedback to regulate force, detect slip, and adjust actions during interaction. Without such feedback, demonstrations may contain suboptimal force patterns, limiting downstream policy performance.

In this work, we propose a haptic-fused imitation learning framework that incorporates force feedback into demonstration collection. Using a wrist-mounted force–torque sensor and a Haply Inverse3 haptic controller, we stream contact forces from the robot to the operator in real time. We compare two teleoperation conditions: a force-blind baseline and a haptic-fused setup with real-time force feedback. Experiments on two contact-rich tasks show that haptic feedback leads to improved demonstration quality, including reduced excessive force and more stable contact. These improvements are reflected in the learned policy, although they do not translate into immediate gains in task success.

Our contributions are as followed: First, we present a haptic teleoperation framework for demonstration collection. Second, we provide quantitative evidence that haptic feedback improves force regulation during demonstrations and influences learned policy behavior. Third, we highlight the importance of task selection in contact-rich imitation learning, and suggest that tasks where force information is not visually inferable are more likely to benefit from haptic feedback.

II. METHODOLOGY

We propose a haptic-fused imitation learning framework for contact-rich manipulation, combining force-aware teleoperation with diffusion-based policy learning (Fig. 1). During data collection, we compare two modalities: force-blind teleoperation, using standard interfaces, and force-streamed teleoperation, where a haptic device (Haply controller) provides bidirectional interaction by transmitting user commands to the robot and returning force/torque feedback to the operator.

Demonstrations are used to train a policy that maps multimodal observations, including camera views, robot pose, and force/torque signals, to action sequences. A diffusion policy is trained on these observation sequences to generate robot trajectories at inference time. The framework is evaluated on

*Equal contribution.

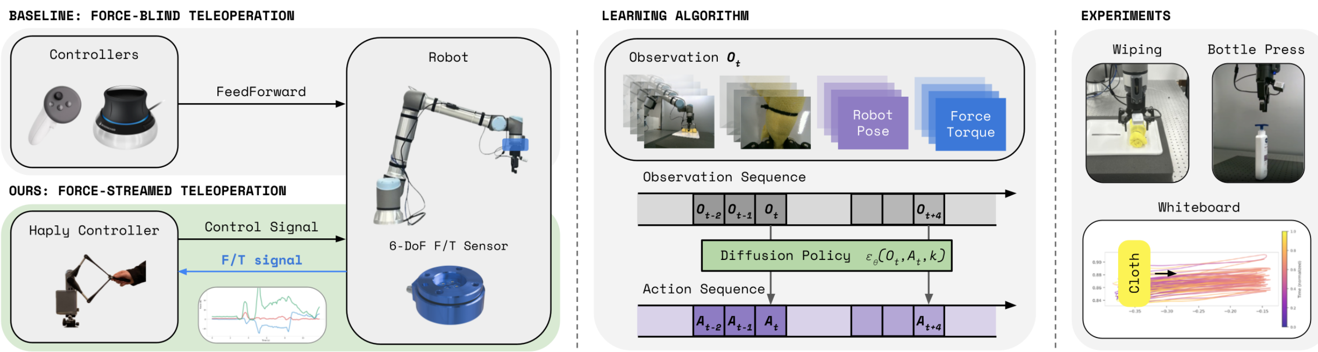

Fig. 1: Overview of the haptic-fused imitation learning framework.

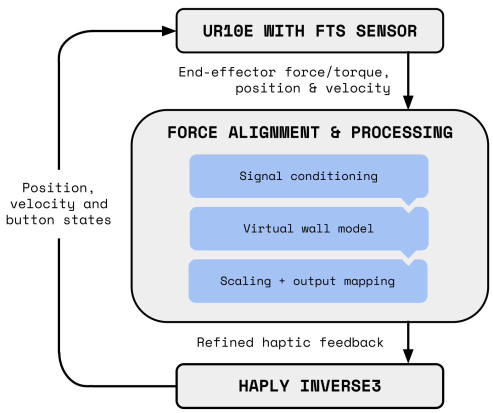

Fig. 2: Overview of the proposed haptic feedback loop.

contact-rich tasks, including wiping and pressing, to study the impact of force feedback.

A. Hardware

A UR10e robot is controlled via RTDE for low-latency operation. Teleoperation uses a Haply Inverse3 with VerseGrip (6-DoF input) and a Robotiq gripper. Two RealSense RGB cameras (wrist and external) provide visual input, while a SensuReal 6-DoF force–torque sensor measures interaction forces.

B. Teleoperation Pipeline

At each control cycle, the system synchronously acquires robot state, operator input, and visual observations to ensure temporally consistent multimodal data. Incremental pose updates from the haptic interface are integrated to generate target 6-DoF end-effector motions, while corresponding proprioceptive and force–torque measurements are recorded. Visual streams from multiple cameras are captured in parallel and aligned via timestamp-based resampling onto a unified time grid.

The haptic feedback pipeline (Fig. 2) establishes bidirectional coupling between the robot and the operator. Measured end-effector forces are processed to generate real-time feedback signals rendered through the haptic device, conveying contact dynamics during interaction. To ensure stability, the force signal is conditioned using low-pass filtering and rate limiting. A virtual contact model based on a unilateral spring–damper formulation regulates interaction along the contact axis, with threshold-based engagement and smooth transitions between free-space and contact. The resulting force is scaled and rendered to the operator, closing the loop between motion input and force feedback.

C. Policy Architecture

The policy is based on the Diffusion Policy framework [1]. End-effector poses in both observations and actions are represented relative to the start of each episode. The model is trained as a conditional diffusion policy that predicts future action sequences from multi-view RGB observations and proprioceptive inputs.

A global conditioning vector is constructed from an observation encoder over a four-step horizon. Each camera view is encoded using a ResNet-18 backbone[11], and the resulting features are concatenated with low-dimensional inputs, including proprioception and force–torque measurements. The combined representation conditions a CNN-based diffusion model, where the denoising network is implemented as a 1D U-Net [12] over the temporal action sequence.

During inference, Real-Time Chunking (RTC) [13] is used to improve execution of predicted trajectories in real-world settings.

D. Task Design

We evaluate the proposed approach on two manipulation tasks with distinct contact interaction requirements. The **bottle-cap press** task requires sensitivity to object compliance, where sufficient normal force must be applied to fully depress the cap without exerting excessive force. The **whiteboard wiping** task involves sustained surface contact, requiring the operator to maintain appropriate normal force while moving across the board to remove marker traces.

To study the effect of haptic feedback on both dataset quality and downstream policy performance, demonstrations are collected under two teleoperation modes: unilateral control (without haptic feedback) and bilateral control (with haptic

TABLE I: Peak Normal Force F_z

Condition	Median $\pm$ Std Dev (N)	Reduction
w/o Haptic Feedback	71.3 $\pm$ 18.0	64.0%
w/ Haptic Feedback	25.7 $\pm$ 17.4	

feedback). Dataset quality is evaluated using interaction-based metrics to better understand the operator and environment interaction:

- **Temporal profile of normal force:** The evolution of normal force F_z over an episode is analyzed to compare force application patterns across both tasks.
- **Peak normal force (bottle-cap press):** The maximum normal force reached in each episode is used to quantify excessive force during pressing. A reference range of 15–30 N, empirically determined from typical human usage, is used to assess whether demonstrations apply appropriate force magnitudes.
- **End-effector displacement along the surface-normal direction (whiteboard wipe):** Motion along the surface-normal axis is analyzed to assess whether consistent contact with the wiping surface is maintained.

To evaluate downstream policy performance, the same interaction-based metrics are applied. In addition, task-level performance is measured using functional success metrics:

- **Bottle-cap press:** A binary metric indicating whether the cap is fully depressed.
- **Whiteboard wipe:** A discrete score from 0 to 4 based on the proportion of the marked line removed, where each successfully erased segment contributes one point.

III. EXPERIMENTS AND RESULTS

We evaluate the proposed approach by examining whether haptic feedback during teleoperated data collection improves demonstration quality, and whether these improvements translate to downstream policy performance.

A. Demonstration Quality Comparison

Bottle-Cap Press. Demonstrations collected with haptic feedback exhibit substantially improved force regulation relative to demonstrations collected without feedback. As shown in Fig. 3, the temporal profile of the normal force F_z remains consistently lower throughout the contact phase for the bilateral teleoperation condition. Furthermore, the comparable spread of the standard deviation bands across both conditions indicates that this reduction in applied force is not achieved at the expense of increased variability, suggesting more controlled and repeatable interaction.

This trend is further supported by the per-episode peak force statistics shown in Table I, which show a reduction in median peak normal force from 71.3 N (w/o haptic feedback) to 25.7 N (w/ haptic feedback), corresponding to a 64% decrease. Moreover, the median force under haptic feedback falls within the predetermined reference range of 15–30 N, whereas in the unilateral condition, it substantially exceeds this range. These results indicate that haptic feedback

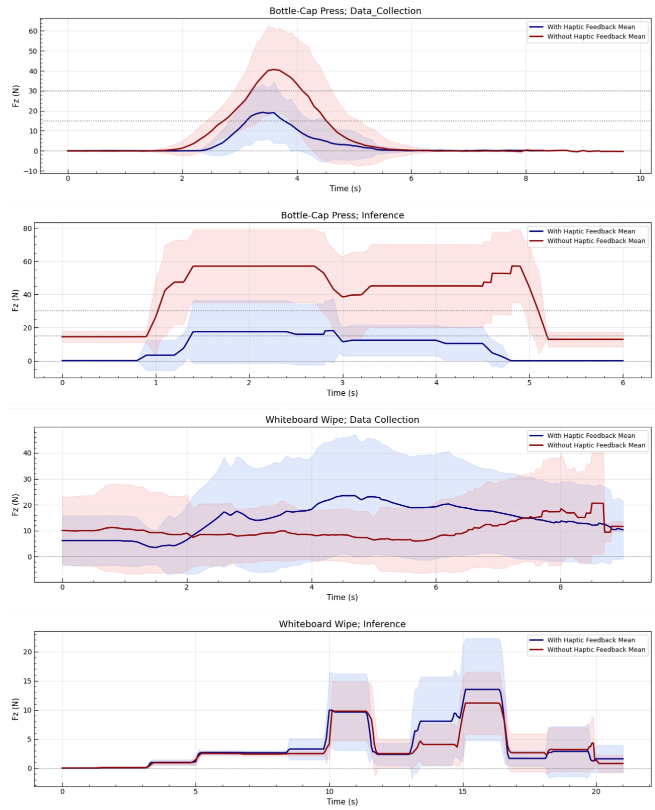

Fig. 3: Temporal evolution of mean normal contact force F_z during data collection and policy execution for both tasks. Shaded regions denote ± 1 standard deviation across episodes.

TABLE II: Contact Maintenance Ratio (ratio over each episode)

Condition	Mean $\pm$ Std	Increase
w/o Haptic Feedback	0.626 $\pm$ 0.311	45.3%
w/ Haptic Feedback	0.910 $\pm$ 0.146	

improves dataset quality for transient contact interactions by enabling more task-appropriate force application during demonstration collection.

Whiteboard Wipe. The desired interaction for this task differs from that of the bottle-cap press—successful execution requires sustained contact with the surface and sufficiently large maintained normal force throughout the wiping motion. Demonstrations collected with haptic feedback for this task also show improved interaction quality. The temporal force profile in Fig. 3 shows that the haptic-feedback condition produces higher maintained normal force during the main contact phase. Quantitatively, the mean maintained normal force is 14.5 $\pm$ 13.8 N with haptic feedback, compared to 8.30 $\pm$ 8.72 N without feedback. Although both conditions exhibit substantial variability, the higher maintained force under haptic feedback is more consistent with the task requirement of preserving contact with the whiteboard surface during motion.

TABLE III: Task Success Rate

Task	Wipe	Press
w/o Haptic Feedback	90%	65%
w/ Haptic Feedback	100%	68%

To further assess contact stability, we analyze end-effector displacement along the surface-normal direction. For each episode, the initial end-effector position is taken as the zero-reference threshold, and contact maintenance ratio is computed as the fraction of the episode during which the end-effector remains below this threshold. As shown in Table II, the mean contact maintenance ratio increases by 45.3% for demonstrations collected with feedback compared to those without. Additionally, a larger amount of residual whiteboard marker was observed after the completion of unilateral teleoperation, as opposed to teleoperation with feedback, because of more inconsistent whiteboard and cloth contact.

These results show that haptic feedback improves dataset quality across both tasks by allowing teleoperators to produce more task-appropriate interaction behavior.

B. Policy Performance Comparison

We next examine whether the improvements in demonstration quality will lead to improved downstream policy performance.

Bottle-Cap Press. Policies trained on datasets collected with or without haptic feedback achieve similar task completion, with success rates of 100% and 90%, respectively (Table III). We additionally analyze the force profiles from Fig. 3 during execution to determine whether the improved demonstration quality is reflected in learned interaction behavior. While the policies trained on unilateral and bilateral demonstrations exhibit different force profiles, the mean normal force in both cases remains outside the 15–30 N reference range. This indicates that the haptic condition does not improve task-appropriate force regulation during execution. Furthermore, due to the use of real-time chunking (RTC) during execution, the force profiles are smoothed and do not exhibit the same transient peaks observed in the demonstrations which may attenuate differences in force behavior learning from the demonstrations.

Whiteboard Wipe. A similar trend is observed for the whiteboard wipe task. Task-level performance is evaluated using the normalized mean of a 0–4 completion score. Table III shows that policies trained with or without haptic feedback achieve comparable performance, with normalized mean scores of 68% and 65%, respectively.

Overall, these results indicate that haptic feedback provides a clear advantage during data collection by improving force regulation and contact maintenance in a task-dependent manner. However, for the tasks considered, these improvements in dataset quality do not yield corresponding improvements in downstream policy performance.

IV. DISCUSSION

Our results demonstrate that haptic force feedback enables context-sensitive force regulation during teleoperation, rather than simply reducing the overall applied force. Qualitative analysis of force profiles confirms that during the bottle-cap press task, operators with haptic feedback eased pressure earlier, avoiding the prolonged over-pressing seen in the force-blind baseline; during the wiping task, operators maintained more stable contact with the surface. These yield an important quality improvement of the collected demonstrations.

However, despite reflecting the force patterns found in the demonstrations, the policy *did not* demonstrate statistically significant performance improvement during inference. We identify task selection as the major bottleneck: the two evaluated tasks impose relatively loose constraints on force regulation, and critically, the visual observation stream contains sufficient positional and deformation cues to infer contact state, leaving little room for force-enriched demonstrations to confer additional benefit.

We introduce *force-observability* as a criterion for task selection in contact-rich benchmarks—a task is force-observable if the contact state and force demand cannot be reliably inferred from vision alone. Candidate tasks satisfying this criterion include peg-in-hole insertion, cap-tightening, and multi-material slip detection, where we expect haptic-augmented demonstrations to yield direct policy performance gains.

V. CONCLUSION & FUTURE WORK

We presented a haptic-fused imitation learning framework that incorporates force feedback during demonstration collection, and conducted a controlled study to examine its impact on data quality and downstream policy behavior. We showed that enabling haptic feedback leads to clear force-aware behavior, including reduced peak contact force (up to 64%) and more stable contact during surface-following tasks. These differences are reflected in the learned policy, although they do not translate into improved task success under the current open-loop execution.

Our results highlight the importance of task characteristics in evaluating force-aware learning. In particular, tasks where contact state cannot be reliably inferred from visual observations are more likely to benefit from force-enriched demonstrations. We refer to this property as *force-observability*, and propose it as a useful consideration for future benchmark design.

Several directions remain for improving force integration into policy learning. Beyond incorporating force as an input, a promising approach is to combine trajectory prediction with compliance control, where the policy outputs motion together with a stiffness signal that is executed via an admittance controller [14], [15]. This separation of motion planning and force regulation may enable more robust execution in contact-rich scenarios.

REFERENCES

[1] C. Chi, Z. Xu, S. Feng, E. Cousineau, Y. Du, B. Burchfiel, R. Tedrake, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *The International Journal of Robotics Research*, 2024.

[2] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," in *Proceedings of Robotics: Science and Systems (RSS)*, 2023.

[3] J. Luo, Z. Hu, C. Xu, Y. L. Tan, J. Berg, A. Sharma, S. Schaal, C. Finn, A. Gupta, and S. Levine, "SERL: A software suite for sample-efficient robotic reinforcement learning," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2024.

[4] J. Luo, C. Xu, X. Geng, G. Feng, K. Fang, L. Tan, S. Schaal, and S. Levine, "Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning," *Science Robotics*, 2024.

[5] A. Rajeswaran, V. Kumar, A. Gupta, G. Vezzani, J. Schulman, E. Todorov, and S. Levine, "Learning complex dexterous manipulation with deep reinforcement learning and demonstrations," in *Proceedings of Robotics: Science and Systems (RSS)*, 2018.

[6] R. Shukla, S. Moode, R. Talan, and S. K. Gupta, "Learning force-conditioned visuomotor diffusion policy from human demonstrations for complex robotic assembly tasks," *Manufacturing Letters*, vol. 44, pp. 1513–1524, 2025.

[7] Z. He, H. Fang, J. Chen, H.-S. Fang, and C. Lu, "FoAR: Force-aware reactive policy for contact-rich robotic manipulation," 2025.

[8] E. Helmut, N. Funk, T. Schneider, C. de Farias, and J. Peters, "Tactile-conditioned diffusion policy for force-aware robotic manipulation," 2025.

[9] W. Liu, J. Wang, Y. Wang, W. Wang, and C. Lu, "ForceMimic: Force-centric imitation learning with force-motion capture system for contact-rich manipulation," 2025.

[10] H. Xue, J. Ren, W. Chen, G. Zhang, Y. Fang, G. Gu, H. Xu, and C. Lu, "Reactive diffusion policy: Slow-fast visual-tactile policy learning for contact-rich manipulation," 2025.

[11] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.

[12] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, 2015.

[13] K. Black, M. Y. Galliker, and S. Levine, "Real-time execution of action chunking flow policies," 2025. [Online]. Available: https://arxiv.org/abs/2506.07339

[14] J. A. Maples and J. J. Becker, "Experiments in force control of robotic manipulators," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 1986, pp. 695–702.

[15] Y. Hou and M. T. Mason, "Robust execution of contact-rich motion plans by hybrid force-velocity control," in *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2019, pp. 1933–1939.