

# Static Is Not Enough: A Comparative Study of VR and SpaceMouse in Static and Dynamic Teleoperation Tasks

Yijun Zhou

Vrije Universiteit Amsterdam  
Department of Computer Science  
Amsterdam the Netherlands  
y.zhou7@student.vu.nl

Muhan Hou

Vrije Universiteit Amsterdam  
Department of Computer Science  
Amsterdam the Netherlands  
m.hou@vu.nl

Kim Baraka

Vrije Universiteit Amsterdam  
Department of Computer Science  
Amsterdam the Netherlands  
k.baraka@vu.nl

## Abstract

Imitation learning relies on high-quality demonstrations, and teleoperation is a primary way to collect them, making teleoperation interface choice crucial for the data. Prior work mainly focused on static tasks, i.e., discrete, segmented motions, yet demonstrations also include dynamic tasks requiring reactive control. As dynamic tasks impose fundamentally different interface demands, insights from static-task evaluations cannot generalize. To address this gap, we conduct a within-subjects study comparing a VR controller and a SpaceMouse across two static and two dynamic tasks ( $N = 25$ ). We assess success rate, task duration, cumulative success, alongside NASA-TLX, SUS, and open-ended feedback. Results show statistically significant advantages for VR: higher success rates, particularly on dynamic tasks, shorter successful execution times across tasks, and earlier successes across attempts, with significantly lower workload and higher usability. As existing VR teleoperation systems are rarely open-source or suited for dynamic tasks, we release our VR interface to fill this gap.

## CCS Concepts

• **Human-centered computing** → **User studies**; • **Computer systems organization** → **Robotic control**.

## Keywords

human-robot interaction; teleoperation; interfaces; dynamic manipulation; learning from demonstrations

### ACM Reference Format:

Yijun Zhou, Muhan Hou, and Kim Baraka. 2026. Static Is Not Enough: A Comparative Study of VR and SpaceMouse in Static and Dynamic Teleoperation Tasks. In *Companion Proceedings of the 21st ACM/IEEE International Conference on Human-Robot Interaction (HRI Companion '26)*, March 16–19, 2026, Edinburgh, Scotland, UK. ACM, New York, NY, USA, 5 pages. <https://doi.org/10.1145/3776734.3794405>

## 1 INTRODUCTION

Imitation learning is an effective approach for learning robot policies from user-provided data, and teleoperation is one of the most common ways to collect such demonstrations [1, 17, 27]. The choice

of teleoperation interfaces therefore plays an important role in influencing the resulting data [12, 22], as it impacts how operators express their actions during control.

The demonstrations collected via teleoperation span a range of tasks, which can be grouped into static and dynamic tasks. Static tasks involve discrete, segmented motions, such as pick-and-place or reaching. Dynamic tasks, by contrast, require reactive, continuous control with time-sensitive adjustments [19, 23], such as flipping an object in a pan. As these two task types place fundamentally different demands on human control, an interface effective for static tasks may not necessarily perform well on dynamic ones.

Despite this distinction, prior evaluations of teleoperation interfaces have focused primarily on static tasks [5, 15, 18, 25]. As a result, existing insights do not directly clarify how well commonly used interfaces, such as VR controllers and SpaceMouse [15, 21], support dynamic tasks. In this work, VR refers solely to the use of a VR controller for 6-DoF tracking rather than a virtual environment. This gap makes it important to examine how different interfaces affect user task performance across both task types.

While several VR teleoperation systems have been used to collect high-quality demonstrations [6, 18, 24, 25], many are not publicly available, limiting systematic comparison across interfaces. To address this, we introduce an open-source VR teleoperation interface and conduct a within-subjects comparison of VR and SpaceMouse across static and dynamic tasks. Our contributions are:

- A systematic user study comparing VR and SpaceMouse (both 6 DoF) across static and dynamic tasks, evaluating user task performance and user experience.
- An open-source VR teleoperation interface that supports demonstration collection in dynamic settings.

## 2 RELATED WORK

### 2.1 Teleoperation Interface Comparisons

VR controllers [6, 18, 24, 25] and SpaceMouse devices [4, 30] have become the dominant teleoperation interfaces for collecting demonstrations in imitation learning. Prior comparisons across input devices [5, 18] and direct VR–SpaceMouse studies [15, 25] report differences in usability, workload, and policy learning outcomes. However, these works focus almost exclusively on static or quasi-static manipulation, where motion is discrete and segmented.

This creates a key gap: static tasks cannot reveal how interfaces behave under reactive, continuous motion, nor how pose-based and velocity-based control differ when timing and momentum matter. Our work extends these comparisons to dynamic tasks, where



This work is licensed under a Creative Commons Attribution 4.0 International License. *HRI Companion '26, Edinburgh, Scotland, UK*  
© 2026 Copyright held by the owner/author(s).  
ACM ISBN 979-8-4007-2321-6/2026/03  
<https://doi.org/10.1145/3776734.3794405>

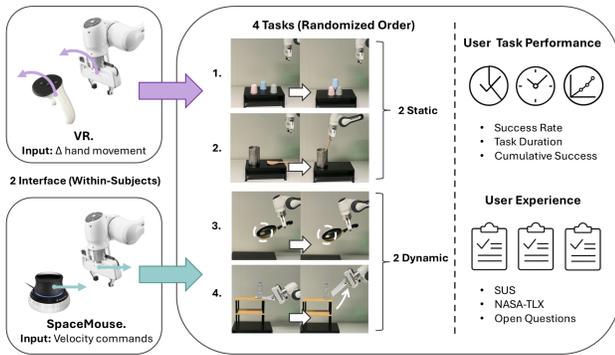


Figure 1: Overview of the study design, including the two teleoperation interfaces (VR, SpaceMouse), four randomized tasks (two static and two dynamic), and evaluated measures of user task performance and user experience.

these distinctions become clearer in how they influence user task performance and experience.

## 2.2 Dynamic Manipulation and Data Collection

Dynamic manipulation involves exploiting inertia and rapid motion rather than purely quasi-static control [19]. Recent systems such as ALOHA [29] and HiLSERL [16] successfully collect high-quality demonstrations on dynamic tasks through teleoperation, yet they provide limited analysis of how interface choice influences user experience. Moreover, these systems typically rely on a single input device, which makes cross-interface comparisons difficult.

A separate line of work learns dynamic manipulation directly from human videos [3, 9, 10, 28]. These approaches bypass teleoperation altogether, but they lack real-time human feedback, leaving open questions about how different teleoperation interfaces support data collection for dynamic tasks.

Another direction studies dexterous in-hand manipulation [7, 13, 20, 26], which centers on fine-grained finger motions rather than the larger, faster movements in arm-level dynamic tasks. Our study complements these efforts by examining interface effects in a more accessible setting.

## 3 METHODOLOGY

For VR teleoperation, we use the Oculus Quest 3 controller for 6-DoF tracking, while the headset is only used to capture motion and button inputs. Hand motion is mapped to robot end-effector commands via pose tracking [14]. At each control cycle, the controller pose is transformed into the robot frame to obtain the hand rotation and position  $(R_t^{hand}, p_t^{hand})$ .

Teleoperation is activated when the user presses any button, which records the current hand pose and the robot end-effector pose as references. End-effector motion is then generated from incremental hand motion. The positional command is computed from frame-to-frame displacement,

$$a_t^{pos} = \alpha (p_t^{hand} - p_{t-1}^{hand}), \quad (1)$$

and applied to update the end-effector position.

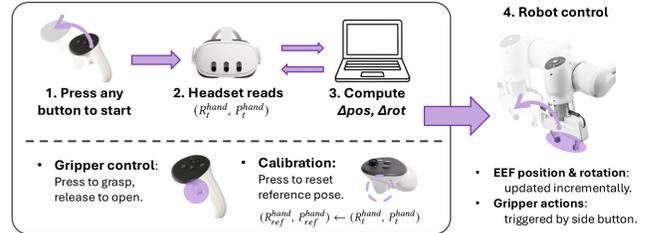


Figure 2: Workflow of our VR teleoperation interface. Controller pose is tracked and mapped to incremental end-effector commands, with button inputs handling gripper control and calibration.

Rotational motion is obtained from the relative hand rotation,

$$a_t^{rot} = \beta \cdot \text{Log}(R_t^{hand}(R_{ref}^{hand})^{-1}), \quad (2)$$

and applied through a quaternion update for numerical stability.

A calibration step is included to prevent VR drift: pressing the trigger resets the hand reference pose and its corresponding robot reference. This allows users to re-center their hand at any time without affecting the robot’s current pose.

## 4 EXPERIMENTAL SETUP

### 4.1 Ethics Statement

The study was approved by the Research Ethics Review Committee of the Faculty of Science, Vrije Universiteit Amsterdam (Ref. 25-045).

### 4.2 Hardware and Tasks

We use a 7-DoF Franka Emika Panda robotic arm controlled via the Cartesian impedance controller in panda-py [8]. We use the Oculus Quest 3 controller for VR teleoperation. For SpaceMouse, we use a 3Dconnexion SpaceMouse Compact, and its six-axis velocity input is mapped to end-effector pose deltas via a scaling factor.

We choose stacking three cubes as the training task, which is a classic manipulation task that allows participants to practice the basic skills required by both interfaces: moving, rotating, grasping, and controller calibration (for VR).

We select four tasks as our testbed (Figure 3), covering both static and dynamic tasks. For dynamic tasks, the action must be completed in a single continuous motion; slow or stop-and-go movements do not count as success. This structure allows us to compare user performance across task complexity and generalization.

The task order is in a constrained counterbalancing scheme. The two static tasks (T1, T2) and the two dynamic tasks (T3, T4) are randomly permuted separately. All participants complete the static pair before the dynamic pair, maintaining counterbalancing while providing a smooth transition from discrete to reactive actions.

### 4.3 Participants

We recruited 25 participants for a within-subjects study (17 male, 8 female, 0 other; mean age 24). Their prior experience with robots was mostly low, with 7 participants reporting no experience, 10 limited, 7 moderate, and 1 extensive. Prior gaming experience was generally higher, with 2 reporting no experience, 4 limited, 8 moderate, and 11 extensive. All participants completed both conditions.

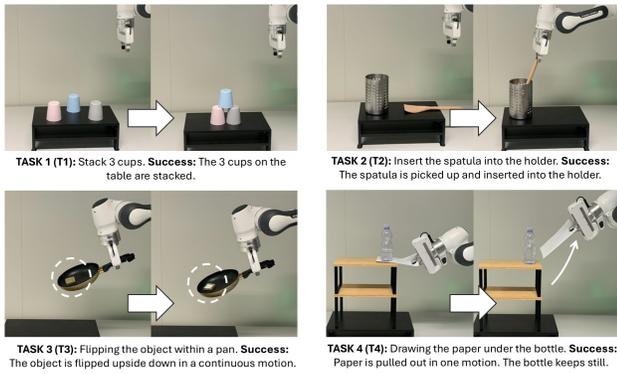


Figure 3: 4 Tasks used in our study, including success criteria. T1 and T2 are static tasks. T3 and T4 are dynamic tasks requiring continuous motion.

Participants were recruited via on-campus posters. The experiment followed the ethical guidelines of our faculty’s research ethics board and received its approval. Informed consent for participation and data collection was obtained before participation. Each participant received a €10 gift card as compensation upon finishing.

#### 4.4 Procedure

The experiment began with the VR controller condition. Participants first completed a training task that required one successful demonstration, with up to three minutes allowed. Participants who could not complete the task within this period were dismissed.

After training, participants performed four tasks. The order followed the constrained counterbalancing scheme: two static tasks followed by two dynamic ones, each internally randomized. For each task, participants had one minute to explore and up to five attempts, with each attempt capped at two minutes. A task ended once a single successful attempt was achieved. Attempts exceeded the time limit or failed the success criterion were recorded as failures. Object placements remained identical across all participants.

After the VR condition, participants filled out the NASA-TLX [11] and the System Usability Scale (SUS) [2]. After a one-minute break, they repeated the same procedure using the SpaceMouse interface. The fixed VR-SpaceMouse order was adopted for safety, as pilot observations showed that starting with VR would reduce early-stop issues by providing initial task familiarity.

At the end of the study, participants completed an open-ended questionnaire covering demographics, prior experience with robots and gaming, and subjective feedback.

#### 4.5 Metrics

**4.5.1 User Task Performance.** We measure user task performance using success rate, task duration, and cumulative success. For each participant, we record the number of attempts until their first successful trial, whether they succeeded at all, and the duration of each demonstration. Success is labeled manually based on predefined criteria. The detailed definitions of these metrics are as follows:

- **Success rate:** ratio of successful attempts to total attempts.

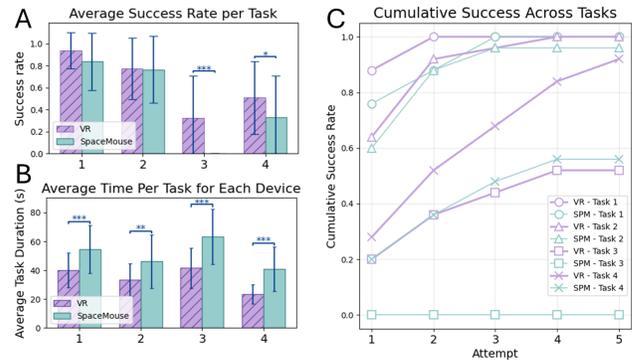


Figure 4: User task performance across four tasks. A: Average success rate for VR and SpaceMouse across tasks. B: Average completion time across tasks under both interfaces. C: Cumulative success rate over the five allowed attempts for each task. VR consistently achieves higher success rates, yields significantly shorter completion times, and reaches successful completion earlier across static and dynamic tasks. Significance levels are indicated as: \*  $p < .05$ , \*\*  $p < .01$ , \*\*\*  $p < .001$ .

- **Task duration:** duration of the successful attempt, or the average duration across attempts for participants who do not succeed.
- **Cumulative success:** at each attempt number, the proportion of participants who have achieved at least one successful completion.

**4.5.2 User Experience.** User experience is assessed through workload, usability, and open-ended feedback. We use an unweighted version of NASA-TLX to measure mental demand, physical demand, temporal demand, perceived performance, effort, and frustration, on a 1–7 scale. Perceived usability is evaluated with the standard 10-item SUS, scored on a 0–100 scale. Participants also give open-ended feedback on interface preference, perceived strengths, difficulties, and suggestions for improvement.

## 5 RESULTS AND DISCUSSION

### 5.1 User Task Performance

Success rates decrease across tasks, with T3 being the most difficult (Figure 4A). In the two static tasks, VR and SpaceMouse show no significant difference. Static manipulation relies on slow, incremental adjustments, and under such conditions both interfaces support similar success behavior (T1: 94.0% vs. 84.0%; T2: 77.7% vs. 76.7%).

Dynamic tasks present a sharply different pattern. Statistical tests show clear differences in both T3 and T4: VR outperforms SpaceMouse with a large gap in T3 (31.2% vs. 0%,  $p < .001$ ) and a smaller but significant difference in T4 (50.9% vs. 33.3%,  $p < .05$ ). The shift from static to dynamic tasks thus reveals where the two interfaces begin to diverge in success behavior.

T3 shows the strongest contrast. Flipping requires an impulsive wrist rotation that changes momentum abruptly. VR’s pose-delta input turns this motion into a sharp, high-acceleration command, while SpaceMouse outputs smoothed velocity, flattening impulsive

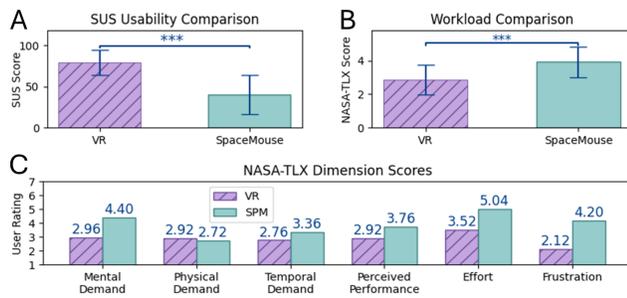


Figure 5: User subjective feedback. A: SUS scores. B: NASA-TLX scores. C: NASA-TLX scores for each individual dimension. VR shows higher usability and significantly lower workload across both static and dynamic tasks.

input into uniform motion. This makes the required impulse hard to produce, explaining the low SpaceMouse success rate.

T4 results in low success for both interfaces, but for different reasons: participants often manage to pull the paper but trigger the robot’s velocity-limit safety via VR, while they frequently activate self-collision protection before reaching the paper with SpaceMouse. These distinct failure modes further demonstrate how the two interfaces behave differently under dynamic demands.

Task duration (Figure 4B) shows significant differences across all tasks (T2:  $p < .01$ , T1/T3/T4:  $p < .001$ ), with VR completing tasks consistently faster. These results indicate that once participants identify a solution, VR allows more direct execution, while SpaceMouse requires longer, segmented motion.

Cumulative success curves (Figure 4C) suggest that our VR interface supports more accessible onboarding for users. For instance, T4 has a low overall success rate, yet 23 out of 25 participants eventually succeed by the fifth attempt via VR. This pattern highlights a clear trend: VR enables earlier success with fewer attempts, especially for dynamic tasks.

## 5.2 User Experience

The SUS results (Figure 5A) show a clear difference between the two interfaces. VR receives a substantially higher usability score ( $M = 79.1$ ) compared to SpaceMouse ( $M = 40.5$ ). A paired t-test confirms that this difference is significant ( $p < .001$ ).

Similar patterns appear in the raw NASA-TLX scores (Figure 5B, 5C). VR received lower workload ratings across all six dimensions, except for a slight increase in physical demand, likely due to the need to hold the VR controller throughout use. VR showed consistently lower perceived workload ( $M = 2.87$ ) compared to SpaceMouse ( $M = 3.91$ ), with the difference highly significant ( $p < .001$ ).

For interface preference, 92% participants preferred VR across tasks. They frequently described it as “intuitive”. Many felt that its spatial mapping matched their hand movements, e.g.,

“VR felt more natural and intuitive. Like I was using my own hand.” – P21

A few positive comments were made about SpaceMouse: three participants liked its low physical effort and minimal calibration, consistent with its common use in teleoperation.

Participants also pointed out limitations for each interface. With the VR controller, three participants felt that impulsive hand movements transferred too much momentum to the robot, sometimes pushing it to its velocity limit. This reflects both the responsiveness of VR and its proneness to momentum-driven errors.

Feedback for SpaceMouse was more critical. Participants found its control logic unclear and often struggled to map intended motions to the robot, especially for rotation, e.g.,:

“The SpaceMouse is limited in how much it can tilt, rotate and move up and down.” – P13

“For spacemouse, rotation was very unpredictable.” – P5

indicating the velocity-based input makes rotation cognitively heavier and less transparent to participants.

## 5.3 Limitations and Future works

This study has several limitations. Some participants noted that the transition from static to dynamic tasks was not fully smooth, and the difficulty gap may have influenced learning and performance. The study also has methodological constraints. The interfaces were not counterbalanced, as VR always preceded SpaceMouse, due to safety concerns. This fixed order may have introduced limited fatigue effects on the SpaceMouse results, despite the inclusion of a rest period. Additionally, trajectory data were not logged, which would prevent deeper analysis of movement quality, control strategies, and error patterns across interfaces.

Future work can address these issues by designing additional intermediate dynamic tasks to create a more gradual difficulty transition. This would allow a more thorough comparison between VR controllers and SpaceMouse on tasks with varying dynamism. Logging full end-effector and joint trajectories could also enable deeper analysis and support building a dataset for future research.

## 6 CONCLUSION

In this work, we study how teleoperation via a VR controller and a SpaceMouse differs in supporting demonstration collection for static and dynamic tasks. Results show that our VR interface yields higher success rates, shorter successful durations, and fewer attempts across all tasks, with participants reporting lower workload and higher usability. These findings indicate that pose-based control helps users produce fast or continuous motions, whereas the velocity-based SpaceMouse makes such actions harder to perform.

Overall, demonstrations collected with our VR interface are of higher quality, especially for dynamic tasks. As existing VR teleoperation systems are rarely open-source or suited for dynamic tasks, we release our VR interface to fill this gap. A promising direction for future work is to add end-effector trajectory logging and use the interface to build a dynamic-task dataset for future research.

## CODE AND VIDEO

Open-source code and documentation can be found at <https://github.com/ZoExOr/VR2Arm-Proj>. A video summarizing the contribution of the work and showing teleoperation examples with both interfaces is available through the supplementary material.

## References

- [1] Suneel Belkhale, Yuchen Cui, and Dorsa Sadigh. 2023. Data quality in imitation learning. *Advances in neural information processing systems* 36 (2023), 80375–80395.
- [2] John Brooke. 1996. *"SUS-A quick and dirty usability scale."* *Usability evaluation in industry*. CRC Press. <https://www.crcpress.com/product/isbn/9780748404605> ISBN: 9780748404605.
- [3] Haonan Chen, Cheng Zhu, Shuijing Liu, Yunzhu Li, and Katherine Rose Driggs-Campbell. 2025. Tool-as-Interface: Learning Robot Policies from Observing Human Tool Use. In *Proceedings of Robotics: Conference on Robot Learning (CoRL)*.
- [4] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. 2025. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research* 44, 10-11 (2025), 1684–1704.
- [5] Xiangyu Chu, Yunxi Tang, Lam Him Kwok, Yuanpei Cai, and Kwok Wai Samuel Au. 2023. Bootstrapping robotic skill learning with intuitive teleoperation: Initial feasibility study. In *International Symposium on Experimental Robotics*. Springer, 42–52.
- [6] Catie Cuan, Allison Okamura, and Mohi Khansari. 2024. Leveraging haptic feedback to improve data quality and quantity for deep imitation learning models. *IEEE Transactions on Haptics* 17, 4 (2024), 984–991.
- [7] Nikhil Chavan Dafle, Alberto Rodriguez, Robert Paolini, Bowei Tang, Siddhartha S. Srinivasa, Michael Erdmann, Matthew T. Mason, Ivan Lundberg, Harald Staab, and Thomas Fuhlbrügge. 2014. Extrinsic dexterity: In-hand manipulation with external forces. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*. 1578–1585. doi:10.1109/ICRA.2014.6907062
- [8] Jean Elsner. 2023. Taming the Panda with Python: A powerful duo for seamless robotics programming and integration. *SoftwareX* 24 (2023), 101532. doi:10.1016/j.softx.2023.101532
- [9] Huy Ha and Shuran Song. 2022. Flingbot: The unreasonable effectiveness of dynamic manipulation for cloth unfolding. In *Conference on Robot Learning*. PMLR, 24–33.
- [10] Siddhant Halder, Jyothish Pari, Anant Rai, and Lrel Pinto. 2023. Teach a Robot to FISH: Versatile Imitation from One Minute of Demonstrations. In *Robotics: Science and Systems*.
- [11] Sandra G. Hart and Lowell E. Staveland. 1988. Development of NASA-TLX (Task Load Index): Results of Empirical and Theoretical Research. In *Human Mental Workload*, Peter A. Hancock and Najmedin Meshkati (Eds.). Advances in Psychology, Vol. 52. North-Holland, 139–183. doi:10.1016/S0166-4115(08)62386-9
- [12] Rebecca Hetrick, Nicholas Amerson, Boyoung Kim, Eric Rosen, Ewart J de Visser, and Elizabeth Phillips. 2020. Comparing virtual reality interfaces for the teleoperation of robots. In *2020 Systems and Information Engineering Design Symposium (SIEDS)*. IEEE, 1–7.
- [13] Binghao Huang, Yuanpei Chen, Tianyu Wang, Yuzhe Qin, Yaodong Yang, Nikolay Atanasov, and Xiaolong Wang. 2023. Dynamic Handover: Throw and Catch with Bimanual Hands. In *Conference on Robot Learning*. PMLR, 1887–1902.
- [14] Frederik Ebert Jdrzej Orbik. 2021. Oculus Reader: Robotic Teleoperation Interface. [https://github.com/rail-berkeley/oculus\\_reader](https://github.com/rail-berkeley/oculus_reader) Accessed: YYYY-MM-DD.
- [15] Haozhuo Li, Yuchen Cui, and Dorsa Sadigh. 2025. How to Train Your Robots? The Impact of Demonstration Modality on Imitation Learning. *arXiv preprint arXiv:2503.07017* (2025).
- [16] Jianlan Luo, Charles Xu, Jeffrey Wu, and Sergey Levine. 2025. Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning. *Science Robotics* 10, 105 (2025), eads5033.
- [17] Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, and Roberto Martin-Martín. 2022. What Matters in Learning from Offline Human Demonstrations for Robot Manipulation. In *Conference on Robot Learning*. PMLR, 1678–1690.
- [18] Ajay Mandlekar, Yuke Zhu, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao, John Emmons, Ancht Gupta, Emre Orbay, et al. 2018. Roboturk: A crowdsourcing platform for robotic skill learning through imitation. In *Conference on Robot Learning*. PMLR, 879–893.
- [19] M.T. Mason and K.M. Lynch. 1993. Dynamic manipulation. In *Proceedings of 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '93)*, Vol. 1. 152–159 vol.1. doi:10.1109/IROS.1993.583093
- [20] A. Namiki, Y. Imai, M. Ishikawa, and M. Kaneko. 2003. Development of a high-speed multifingered hand system and its application to catching. In *Proceedings 2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2003) (Cat. No.03CH37453)*, Vol. 3. 2666–2671 vol.3. doi:10.1109/IROS.2003.1249273
- [21] Abby O'Neill, Abdul Rehman, Abhiram Maddukuri, Abhishek Gupta, Abhishek Padalkar, Abraham Lee, Acorn Pooley, Agrim Gupta, Ajay Mandlekar, Ajinkya Jain, et al. 2024. Open x-embodiment: Robotic learning datasets and rt-x models: Open x-embodiment collaboration 0. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 6892–6903.
- [22] Harish Ravichandar, Athanasios S Polydoros, Sonia Chernova, and Aude Billard. 2020. Recent advances in robot learning from demonstration. *Annual review of control, robotics, and autonomous systems* 3, 1 (2020), 297–330.
- [23] Fabio Ruggiero, Vincenzo Lippello, and Bruno Siciliano. 2018. Nonprehensile Dynamic Manipulation: A Survey. *IEEE Robotics and Automation Letters* 3, 3 (2018), 1711–1718. doi:10.1109/LRA.2018.2801939
- [24] Patrick Stotko, Stefan Krumpen, Max Schwarz, Christian Lenz, Sven Behnke, Reinhard Klein, and Michael Weinmann. 2019. A VR system for immersive teleoperation and live exploration with a mobile robot. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 3630–3637.
- [25] Yun-Peng Su, Xiao-Qi Chen, Tony Zhou, Christopher Pretty, and Geoffrey Chase. 2022. Mixed-reality-enhanced human-robot interaction with an imitation-based mapping approach for intuitive teleoperation of a robotic arm-hand system. *Applied Sciences* 12, 9 (2022), 4740.
- [26] Jun Wang, Ying Yuan, Haichuan Che, Haozhi Qi, Yi Ma, Jitendra Malik, and Xiaolong Wang. 2025. Lessons from Learning to Spin "Pens". In *Conference on Robot Learning*. PMLR, 3124–3138.
- [27] Maryam Zare, Parham M. Kebria, Abbas Khosravi, and Saied Nahavandi. 2024. A Survey of Imitation Learning: Algorithms, Recent Developments, and Challenges. *IEEE Transactions on Cybernetics* 54, 12 (2024), 7173–7186. doi:10.1109/TCYB.2024.3395626
- [28] Andy Zeng, Shuran Song, Johnny Lee, Alberto Rodriguez, and Thomas Funkhouser. 2020. Tossingbot: Learning to throw arbitrary objects with residual physics. *IEEE Transactions on Robotics* 36, 4 (2020), 1307–1319.
- [29] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. 2023. Learning Fine-Grained Bimanual Manipulation with Low-Cost Hardware. In *Proceedings of Robotics: Science and Systems*. Daegu, Republic of Korea. doi:10.15607/RSS.2023.XIX.016
- [30] Yifeng Zhu, Peter Stone, and Yuke Zhu. 2022. Bottom-up skill discovery from unsegmented demonstrations for long-horizon robot manipulation. *IEEE Robotics and Automation Letters* 7, 2 (2022), 4126–4133.

Received 2025-12-08; accepted 2026-01-12