2026 IEEE/SICE International Symposium on System Integration (SII)
January 11-14, 2026. Cancún, México

AIREC-Basic: Consistent Demonstration Data Collection for Imitation Learning with Redundant Robot Arms

Hiroshi Ito[1], Yoshiki Kanai[1], Akira Kanazawa[1], Hideyuki Ichiwara[1], Takahiro Yoshida[1], Naoaki Noguchi[1], Tetsuya Ogata[2]

Abstract— **The performance of robotic imitation learning (IL) largely depends on the quality of human demonstrations. To address this challenge, we present AIREC-Basic, a leader–follower teleoperation system equipped with a dual-arm mobile manipulator that enables efficient data collection. The system employs 8-DoF redundant arms to realize diverse task postures, but such redundancy can reduce the consistency of demonstrations. To overcome this issue, we propose a novel control strategy, Soft Homing Control (SHC), which mitigates redundancy while preserving intuitive operator control, thereby improving dataset consistency. We validate our approach on three household tasks using state-of-the-art IL algorithms (ACT, Diffusion Policy, and HSARNN). Experimental results show that SHC significantly reduces joint trajectory variance and improves task success rates, particularly in scenarios with strong trajectory constraints and frequent contacts.**

I. INTRODUCTION

Recent breakthroughs in deep learning have enabled robots to autonomously execute complex tasks that were previously beyond the reach of traditional methods. Among these, deep reinforcement learning (DRL) has drawn significant attention for its ability to discover novel behaviors and efficient motion strategies through trial-and-error interaction with the environment, often surpassing human-designed policies [1], [2], [3]. While simulation accelerates DRL training, direct application in the real world remains impractical due to the vast number of required episodes, high operational costs, and risks of hardware damage [4]. Furthermore, the persistent reality gap between simulated and physical environments continues to be a fundamental barrier [4].

Imitation learning (IL) offers a promising alternative [5], [6], [7]. By leveraging demonstrations collected directly in real environments, IL mitigates the challenges of reward engineering and helps bridge the simulation-to-reality gap. Teleoperation enables human operators to provide rich sensorimotor data, allowing robots to adapt more effectively to real-world constraints. However, the success of IL hinges on the quality and diversity of demonstrations. Current teleoperation interfaces—such as VR devices, haptic systems, and 3D mice—suffer from limitations in intuitiveness, force feedback, and simultaneous support for both manipulation and navigation. Consequently, collecting diverse, low-noise, high-quality demonstrations remains a key bottleneck.

To address these limitations, we introduce AIREC-Basic, a leader–follower teleoperation system built around a dual-arm mobile manipulator designed for efficient demonstration collection. The system integrates five core features: (i) 8-DoF redundant arms enabling dexterous manipulation in constrained environments; (ii) a lifting torso mechanism extending the workspace from floor-level to elevated tasks; (iii) a compact omnidirectional base with independently driven wheels; (iv) an intuitive teleoperation interface that combines a scaled-down manipulator with a 3D mouse for simultaneous manipulation and navigation control; and (v) an AR-based visual feedback system aligning the operator's viewpoint with the robot's egocentric perspective, ensuring consistency between demonstrations and execution. Together, these components facilitate the collection of high-quality datasets tailored for imitation learning.

Yet, redundancy in 8-DoF manipulators introduces new challenges. Multiple feasible joint configurations can achieve the same end-effector pose, often leading to inconsistent and less reliable datasets. Existing redundancy management approaches include mechanical designs (e.g., compliant elements [8]) and algorithmic techniques based on null-space projection or regularization [9]. However, these methods typically generate restoring forces that resist deviations from reference postures, which in turn reduces teleoperation intuitiveness by opposing operator intent.

In this work, we propose Soft Homing Control (SHC), a control strategy designed to improve both intuitiveness and consistency in demonstration data collection. SHC applies small, constant torque limits toward predefined neutral joint postures. This lightweight mechanism preserves operator freedom for task-relevant motions while stabilizing redundant configurations and reducing posture variance. Importantly, SHC can be implemented directly using the current-limiting functions of standard commercial actuators, without requiring additional hardware modifications. By balancing stability with usability, SHC ensures that teleoperated demonstrations remain natural, consistent, and well-suited for imitation learning. We validate SHC through experiments on three household tasks, using datasets collected with our teleoperation system and evaluated across three state-of-the-art IL algorithms. Results show that SHC significantly reduces joint trajectory variance and improves task success rates, with the greatest benefits observed in tasks demanding strict trajectory adherence and frequent contact interactions.

[1] Hiroshi Ito, Yoshiki Kanai, Akira Kanazawa, Hideyuki Ichiwara, Takahiro Yoshida and Naoaki Noguchi are with the Research & Development Group, Hitachi, Ltd., Japan.

[2] Tetsuya Ogata is with the Faculty of Science and Engineering, Waseda University. ogata@waseda.jp

978-1-6654-5784-2/26/$31.00 ©2026 IEEE

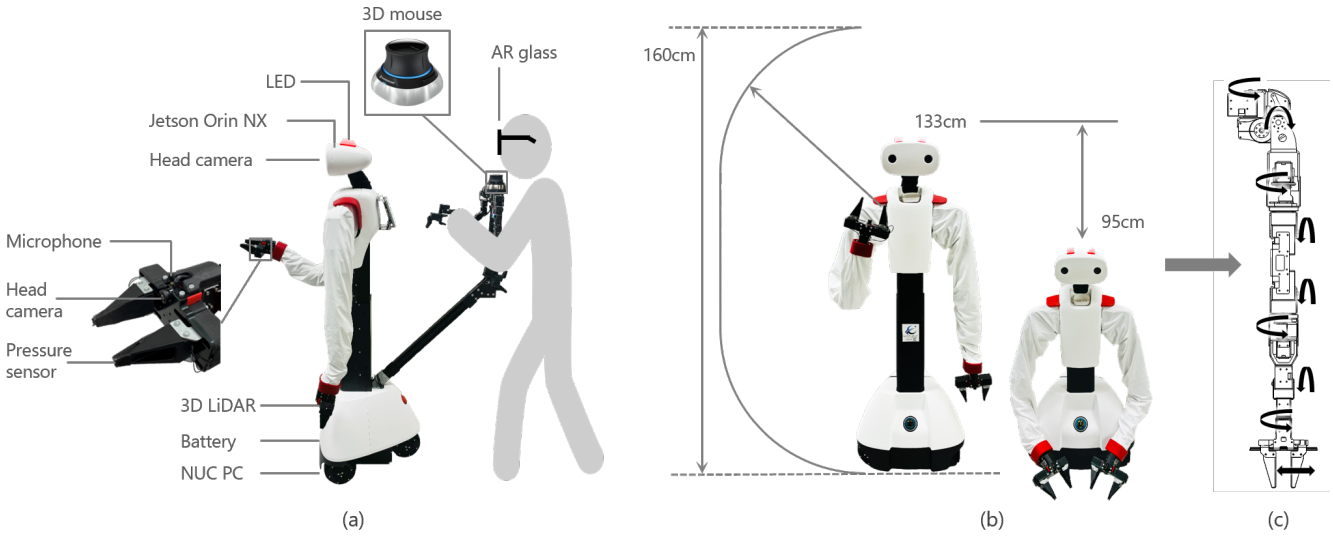

Fig. 1. Robot overview: (a) side view of the robot, and the operator can teach the robot to move from a robot-centric view; (b) front view with arms raised/lowered for low-position manipulation; (c) axis arrangement of the robot arm.

II. RELATED WORK

A variety of teleoperation modalities have been explored. VR-based interfaces provide immersive control but often suffer from mismatches between real-world perception and virtual rendering, which can compromise the naturalness of demonstrations [10]. Haptic devices enable precise control with force feedback but remain costly and largely confined to laboratory use [11]. Low-cost devices such as 3D mice offer affordability but lack intuitive feedback, limiting their utility for collecting diverse demonstrations. Moreover, many existing interfaces focus exclusively on manipulation and do not support the simultaneous teaching of mobility. More recently, leader–follower teleoperation schemes have been introduced to facilitate naturalistic demonstrations [12], yet the design of cost-effective and intuitive teleoperation platforms continues to be an open challenge.

Redundant manipulators further complicate IL. While redundancy provides multiple joint configurations for a single end-effector pose—enhancing flexibility in cluttered or constrained environments [13]—it also introduces challenges for teleoperation. Small variations in operator input can lead to large differences in null-space postures, reducing dataset consistency and hindering policy generalization. Existing redundancy management strategies span mechanical approaches (e.g., springs or elastic components that bias manipulators toward neutral postures, as in the GELLO robot [8]) and algorithmic techniques (e.g., null-space projection or regularization, as in FACTR [9]). Although effective, these methods typically impose restoring forces proportional to deviations from reference postures, which may conflict with operator intent and diminish teleoperation intuitiveness. Thus, redundancy management strategies that stabilize null-space configurations without compromising operator control remain an important open problem.

III. ROBOT OVERVIEW

A. Design Concept

The AIREC (AI-driven Robot for Embrace and Care) project envisions accessible lifelong companion robots for all individuals[14]. Dry-AIREC, developed by Tokyo Robotics, is a high-performance system that combines strong load-bearing capability with flexible motion within this concept. Research toward deploying it in application domains such as healthcare and nursing care is currently underway. However, Dry-AIREC is not well suited for imitation learning applications due to its large body size and the difficulty of intuitive motion teaching. Efficient and user-friendly demonstration collection is essential for such applications.

In contrast, the AIREC-Basic, developed by Hitachi, Ltd., was designed primarily with affordability and teaching flexibility in mind. Although its output torque is lower than that of the Dry-AIREC, the AIREC-Basic enables the teaching and learning of diverse and complex motions. It is intended for daily household assistance tasks, such as cleaning and laundry, where versatility and ease of demonstration are more important than raw load capacity.

Fig. 1 shows an overview of the AIREC-Basic. Table I summarises the key specifications. The system is designed to support diverse data collection and complex task execution in real-world environments. It integrates two processors and five categories of sensors, including stereo cameras for depth estimation and a 3D LiDAR for autonomous navigation, as well as a total of 24 degrees of freedom (DoFs). To ensure uniformity, reliability and maintainability of the hardware, the actuation system employs standardised ROBOTIS Dynamixel XM430 and XM540 series actuators. Current-based motor actuation enables both bilateral control and torque estimation. The composition and functionality of each unit are detailed in the following sections.

TABLE I

SPECIFICATIONS OF THE DEVELOPED ROBOT

DoF	16 (Arms), 2 (Gripper), 1 (Torso), 2 (Head), 3 (Base)
Weight	Approx. 30 kg
Size	42 x 35 x (95-133) cm
Payload	1.5 kg/arm (Worst case holding posture)
Speed	1.3 km/h
Battery life	2.5h (216Wh Lithium-ferrite battery)
Sensors	Stereo camera (head), Fish-eye camera (gripper), Microphone (gripper), Pressure sensor (gripper), 3D LiDAR (base)
Processors	ASUS NUC 14 Pro NVIDIA Jetson Orin NX 16GB

B. Specifications and Units

Head Unit: The head integrates a ZED 2 stereo camera mounted on a 2-DoF pan–tilt mechanism, enabling wide-range visual perception. Visual processing and head motor control are handled by an NVIDIA Jetson Orin NX, which communicates with the main processor over LAN. A full-color LED on the head provides visual feedback of the robot's operational state.

Arm Unit: The robot is equipped with dual 8-DoF manipulators designed for versatile operation in diverse environments. A tandem mechanism with two actuators at each shoulder base enhances payload capacity. Each end-effector incorporates a 1-DoF parallel gripper augmented with a fisheye camera, microphone, and pressure sensor, enabling both grasping and multimodal environmental perception.

Torso Unit: The torso includes a lifting mechanism that extends the operational range from floor-level tasks to elevated workspaces. Powered by a Dynamixel XM540-W150 motor driving a ball screw mechanism (lead: 30 mm, reduction ratio: 0.5), the system achieves an average lift time of 6.7 seconds between its lowest and highest positions.

Base Unit: The base houses the main processor (ASUS NUC 14 Pro), battery, and locomotion system, serving as the central hub for inference and actuator control of both torso and base. A fully Dynamixel-based drive system simplifies control and enhances maintainability. The locomotion system employs independent three-wheel steering, providing omni-directional mobility with a compact and stable footprint.

Teleoperation Unit: This unit consists of a scaled-down manipulator and a 3D mouse. The miniaturized manipulator mirrors the DoF structure of the main robot arms, with joint angles directly mapped for intuitive motion teaching. The 3D mouse controls torso lifting and mobile base movement.

AR Unit: An XREAL Air 2 device provides an ego-centric visual perspective from the robot to the operator. Accelerometer data from the AR headset are transmitted to synchronize the operator's head movements with the robot's head orientation. This alignment ensures consistency between the operator's demonstrations and the robot's viewpoint, which is critical for imitation learning. Without this alignment, human reliance on peripheral vision beyond the robot's field of view could degrade demonstration quality. By enforcing robot-centered teleoperation, the system ensures consistent perspectives and enables the collection of high-quality datasets for imitation learning.

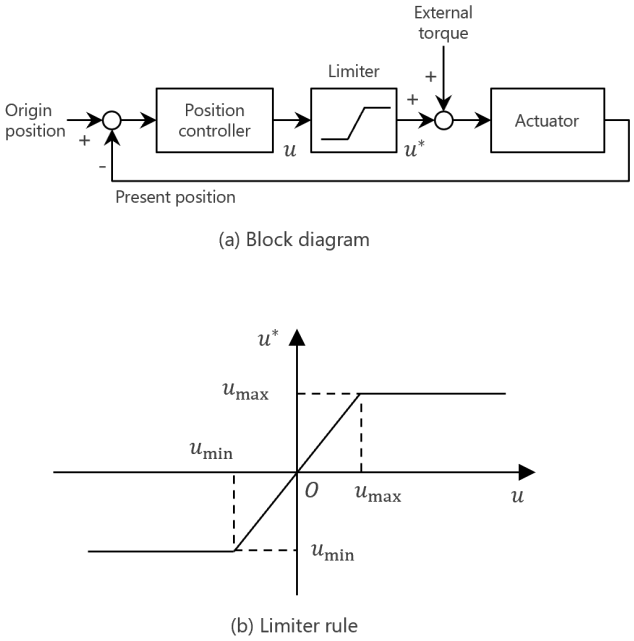

Fig. 2. Block Diagram of Soft Homing Controller

C. Soft Homing Control

AIREC-Basic is equipped with an 8-DoF redundant manipulator, which allows for a wide variety of possible postures. Even for simple tasks such as reaching toward an object in front, there exist countless trajectories and joint angle combinations to achieve the goal. As a result, demonstration data collected through teleoperation often exhibit significant variability, which degrades the quality of datasets for imitation learning. While training on end-effector position data could be considered, this approach faces challenges such as singularities and the computational burden of inverse kinematics.

To address this issue, we propose Soft Homing Control (SHC), a method that applies weak constant torques to gently restore the manipulator toward predefined neutral postures, without interfering with the operator's input. A block diagram of the proposed controller is shown in Fig. 2. Each joint is regulated by a position controller targeting its neutral posture, and the control input u is bounded by upper and lower limits $(u_{\max}, u_{\min})$ to avoid excessive restoring forces, as illustrated in Fig. 2(b). These limits can be implemented as actuator current constraints, and may be further adjusted based on the gravity-compensation torque at the neutral posture. For simplicity, we defined neutral postures as configurations with zero gravitational torque, set $|u_{\max}| = |u_{\min}|$, and assigned approximately 5% of

the actuator's maximum torque as the limit for the scaled-down manipulator. With this configuration, restoring forces remain constant and saturated even under large deviations, effectively preventing posture drift without obstructing user operations. Moreover, SHC requires only conventional position control with actuator current limiting, and thus can be implemented without any hardware modifications.

IV. EXPERIMENTS

To evaluate the effectiveness of the proposed redundant 8-DoF manipulators and the Soft Homing Control (SHC) method, we designed three manipulation tasks with varying environmental constraints and motion complexity. These tasks were chosen to represent a broad spectrum of scenarios commonly encountered in household environments.

(a) Pickup Task (low constraint, simple trajectory): The robot grasps a target object (cup) placed on a table and lifts it vertically. The object is unconstrained, and the end-effector trajectory is simple.

(b) Laundry Task (position + orientation in confined space): The robot inserts clothing into the opening of a washing machine. While the object itself is not highly constrained, the operation must be executed within a limited workspace. This requires precise control of both end-effector position and orientation to avoid collisions with the washing machine.

(c) Toaster Task (position + orientation with contact): The robot opens a toaster door, retrieves a slice of bread, and places it on a plate. This task involves continuous control of end-effector position and orientation under kinematic constraints imposed by the toaster door, making it a contact-rich and sequential manipulation scenario.

To evaluate the impact of SHC on imitation learning generalization performance, we collected demonstrations under two conditions: with and without SHC enabled. For each task, the target object was placed in three different locations. Ten demonstrations were recorded at each location, for a total of 30 demonstrations. These datasets were then used to train and evaluate three representative imitation learning algorithms: ACT (Action Chunking Transformer)[12], DP (Diffusion Policy)[15], and HSARNN (Hierarchical Spatial Attention Recurrent Neural Network) [16].

V. RESULTS AND DISCUSSION

A. Evaluation of Soft Homing Control

To assess the effectiveness of SHC, we analyzed the variability of redundant arm trajectories using 10 demonstrations collected at the same target position. To account for temporal misalignment between trials, trajectory variability was quantified using Dynamic Time Warping (DTW)[17].

The results are summarized in Table II. Overall, applying SHC substantially reduced trajectory variance across all joints (q1–q8) and tasks. For instance, in the Object Pickup Task (a), the maximum variance without SHC was approximately 10, whereas with SHC it was suppressed to around 2. These results demonstrate that SHC effectively mitigates null-space drift in redundant manipulators, thereby enhancing the reproducibility of demonstrations. A task-specific analysis revealed the following trends:

TABLE II
JOINT TRAJECTORY SIMILARITY USING DYNAMIC TIME WARPING

		q1	q2	q3	q4	q5	q6	q7	q8
(a)	w/o SHC	10.7	4.1	9.2	6.3	5.9	8.9	5.5	7.4
	w/ SHC	1.9	2.5	1.7	6.6	5.0	2.9	4.5	4.2
(b)	w/o SHC	15.3	5.4	12.1	5.3	7.2	9.6	10.5	8.0
	w/ SHC	4.3	2.1	3.2	1.8	3.8	3.1	5.2	5.3
(c)	w/o SHC	13.3	6.4	7.1	8.3	7.2	8.6	10.5	2.1
	w/ SHC	3.3	2.3	2.6	3.8	2.1	2.6	3.4	1.3

(a) Pickup Task: Even without SHC, errors remained relatively small and non-critical. Nevertheless, SHC further reduced the error by nearly half, confirming improved reproducibility even in simple, position-only tasks.

(b) Laundry Task: Without SHC, errors reached up to 15.3, indicating instability in posture maintenance. With SHC, errors decreased to within 2–5, significantly improving reproducibility and ensuring reliable collision avoidance in confined spaces.

(c) Toaster Task: Although contact introduced additional sources of error, SHC markedly suppressed trajectory variance. Notably, the variance of q7 decreased from 10.5 to 3.4, while that of q8 decreased from 13.1 to 2.1, underscoring SHC's contribution to maintaining stability during contact-rich interactions.

Figure 4 shows the trajectory variance of joint q3 in the Laundry Insertion Task, which exhibited the largest variance. The blue curve represents trials without SHC, and the red curve represents trials with SHC; shaded regions denote trial-to-trial standard deviations. Without SHC, significant variance appeared between steps 40–80, corresponding to the phase where the robot inserted clothing into the washing machine—a posture-sensitive operation. In contrast, with SHC, this variance was effectively suppressed, converging to a stable value around 3. These findings confirm that SHC reduces unnecessary null-space variations, thereby maintaining consistency in both end-effector position and orientation.

B. Task Success Rates

The task success rates over 10 trials using three imitation learning policies are summarized in Table III. Object positions were randomized in each trial, and success was defined as accurate reproduction of the demonstrated behavior: grasping a cup (Pickup Task), inserting clothes into the washing machine (Laundry Task), or opening the toaster door (Toaster Task). Overall, the results show that applying SHC improved success rates across most tasks:

(a) Object Pickup Task: As this task requires only end-effector position control without environmental constraints, all policies achieved high success rates. ACT and HSARNN consistently reached 100% regardless of SHC, indicating that SHC contributes little in simple, position-only scenarios.

(b) Laundry Task: Because this task required simultaneous control of end-effector position and orientation, SHC

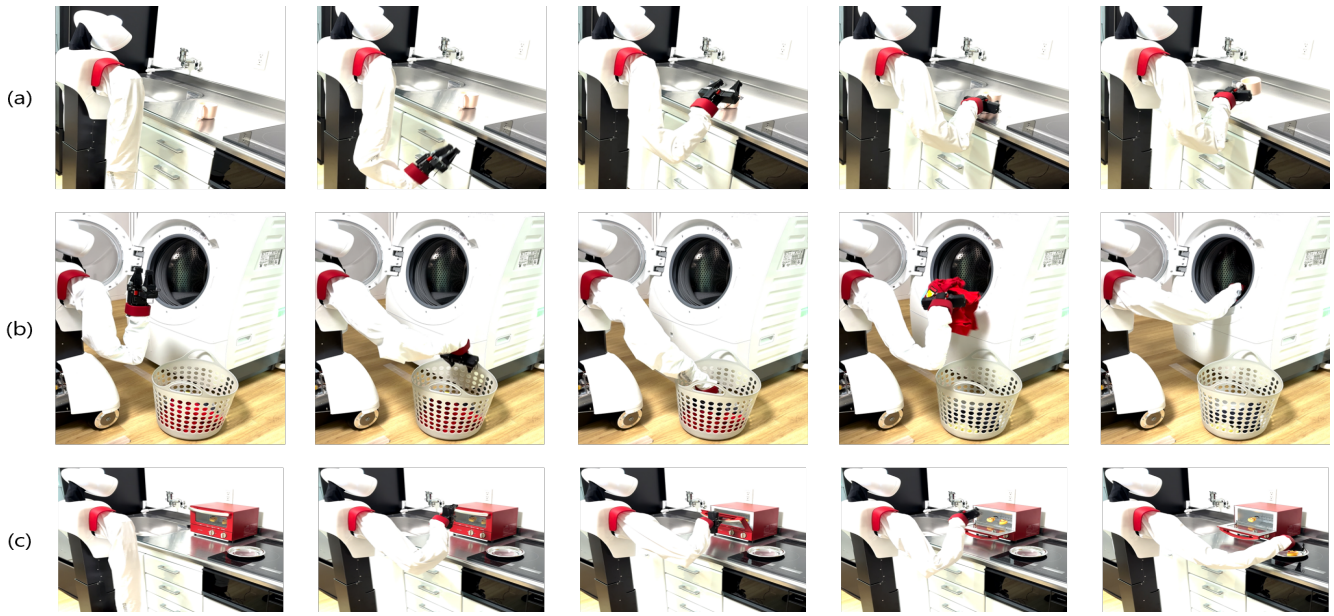

Fig. 3. Experimental task

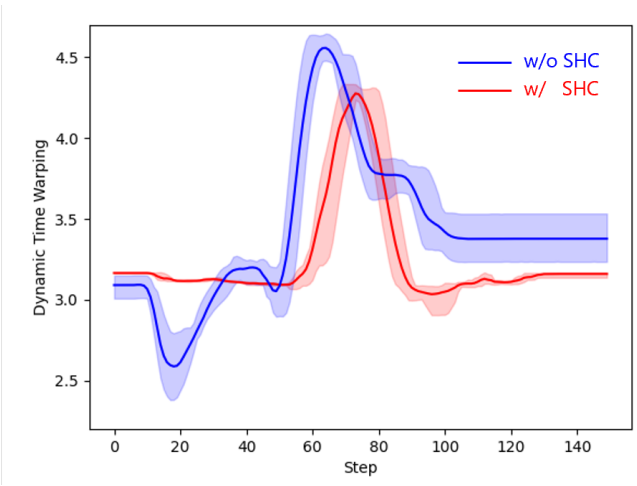

Fig. 4. Trajectory variance of joint q3 during the laundry task. The solid lines represent the mean trajectory barycenter calculated via soft-DTW over 10 trials, while the shaded areas indicate the standard deviation.

had a marked impact. HSARNN improved from 60% to 90%, while both ACT and Diffusion Policy also showed gains. Without SHC, null-space drift frequently caused arm collisions with the washing machine, leading to failures. By suppressing posture drift, SHC significantly enhanced reproducibility and collision-avoidance reliability in confined workspaces.

(c) Toaster Task: As a contact-rich, sequential manipulation scenario, this task was particularly error-prone. With SHC, ACT improved from 70% to 90% and HSARNN from 70% to 80%. These improvements suggest that SHC stabilized joint trajectories and preserved posture consistency at contact points. Diffusion Policy, however, showed little

TABLE III

TASK SUCCESS RATES OVER 10 TRIALS FOR EACH POLICY [%]

		ACT	DP	HSARNN
(a)	w/o SHC	100	80	100
	w/ SHC	100	70	100
(b)	w/o SHC	70	40	60
	w/ SHC	80	60	90
(c)	w/o SHC	70	70	70
	w/ SHC	90	70	80

improvement, likely due to intrinsic limitations in handling contact dynamics.

C. Limitation

Our results demonstrate the effectiveness of the proposed Soft Homing Control (SHC), but several limitations remain. The SHC parameters—such as the neutral postures and torque limits—still need to be manually tuned for each task, and the torque limits are set empirically. Establishing a general and theoretically grounded method that can be applied across different robot structures and actuator types, as well as conducting comparative experiments with prior methods such as FACTR, are important directions for future work.

VI. CONCLUSION

In this paper, we proposed the humanoid robot AIREC-Basic, which has a dual 8-DoF redundant arms. While this redundancy allows for a variety of task postures, it often decreases the consistency of demonstrations, which is essential for imitation learning. To address this issue, we proposed Soft Homing Control (SHC), a lightweight strategy that stabilizes joints with weak restoring torques

while preserving operator freedom. Experiments on household tasks demonstrated that SHC reduces trajectory variance and improves success rates, especially in constrained and contact-rich scenarios.

ACKNOWLEDGEMENT

This work was supported by JST [Moonshot R&D] Grant-Number [JPMJMS2031].

REFERENCES

[1] S. Levine, C. Finn, T. Darrell, and P. Abbeel, "End-to-end training of deep visuomotor policies," *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 1334–1373, 2016.

[2] S. Gu, E. Holly, T. Lillicrap, and S. Levine, "Deep reinforcement learning for robotic manipulation with asynchronous off-policy updates," in *2017 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2017, pp. 3389–3396.

[3] S. Levine, P. Pastor, A. Krizhevsky, J. Ibarz, and D. Quillen, "Learning hand-eye coordination for robotic grasping with deep learning and large-scale data collection," *The International Journal of Robotics Research*, vol. 37, no. 4-5, pp. 421–436, 2018.

[4] 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," *The International Journal of Robotics Research*, vol. 39, no. 1, pp. 3–20, 2020.

[5] T. Osa, J. Pajarinen, G. Neumann, J. A. Bagnell, P. Abbeel, J. Peters *et al.*, "An algorithmic perspective on imitation learning," *Foundations and Trends® in Robotics*, vol. 7, no. 1-2, pp. 1–179, 2018.

[6] H. Ito, K. Yamamoto, H. Mori, and T. Ogata, "Efficient multitask learning with an embodied predictive model for door opening and entry with whole-body control," *Science Robotics*, vol. 7, no. 65, p. eaax8177, 2022.

[7] T. Tsuji, Y. Kato, G. Solak, H. Zhang, T. Petrič, F. Nori, and A. Ajoudani, "A survey on imitation learning for contact-rich tasks in robotics," *arXiv preprint arXiv:2506.13498*, 2025.

[8] P. Wu, Y. Shentu, Z. Yi, X. Lin, and P. Abbeel, "Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators," *arXiv preprint arXiv:2309.13037*, 2023.

[9] J. J. Liu, Y. Li, K. Shaw, T. Tao, R. Salakhutdinov, and D. Pathak, "Factr: Force-attending curriculum training for contact-rich policy learning," *arXiv preprint arXiv:2502.17432*, 2025.

[10] T. Zhang, Z. McCarthy, O. Jow, D. Lee, X. Chen, K. Goldberg, and P. Abbeel, "Deep imitation learning for complex manipulation tasks from virtual reality teleoperation," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 5628–5635.

[11] P. Pastor, H. Hoffmann, T. Asfour, and S. Schaal, "Learning and generalization of motor skills by observation," in *2009 IEEE international conference on robotics and automation*. IEEE, 2009, pp. 763–768.

[12] T. Z. Zhao, V. Kumar, S. Levine, and C. Finn, "Learning fine-grained bimanual manipulation with low-cost hardware," *arXiv preprint arXiv:2304.13705*, 2023.

[13] T. Yoshikawa, "Analysis and control of robot manipulators with redundancy," in *Robotics Research*, H. Hanafusa and H. Inoue, Eds. MIT Press, 1985, pp. 735–747.

[14] "Moonshot r&d program goal 3: Ai-driven robot for embrace and care (airec) project outline." https://airec-waseda.jp/en/about_en/, accessed 1 September 2025.

[15] C. Chi, S. Feng, Y. Du, Z. Xu, E. Cousineau, B. Burchfiel, and S. Song, "Diffusion policy: Visuomotor policy learning via action diffusion," *arXiv preprint arXiv:2303.04137*, 2023.

[16] H. Ichiwara, H. Ito, K. Yamamoto, H. Mori, and T. Ogata, "Multimodal time series learning of robots based on distributed and integrated modalities: Verification with a simulator and actual robots," in *2023 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2023, pp. 9551–9557.

[17] H. Sakoe and S. Chiba, "Dynamic programming algorithm optimization for spoken word recognition," *IEEE transactions on acoustics, speech, and signal processing*, vol. 26, no. 1, pp. 43–49, 2003.