



CLONE: Closed-Loop Whole-Body Humanoid Teleoperation for Long-Horizon Tasks

Yixuan Li^{*,1,2}

lyx@bit.edu.cn

Yutang Lin^{*,3,4,5,6}

yutang.lin@stu.pku.edu.cn

Jieming Cui^{2,3,4,6}

cuijieming@stu.pku.edu.cn

Tengyu Liu^{2,7}

liutengyu@bigai.ai

Wei Liang^{✉,1}

liangwei@bit.edu.cn

Yixin Zhu^{✉,3,4,6,8}

yixin.zhu@pku.edu.cn

Siyuan Huang^{✉,2,7}

syhuang@bigai.ai

* Equal contributors¹ School of Computer Science and Technology, Beijing Institute of Technology

² State Key Laboratory of General Artificial Intelligence, BIGAI

³ School of Psychological and Cognitive Sciences, Peking University

⁴ Institute for Artificial Intelligence, Peking University⁵ Yuanpei College, Peking University

⁶ Beijing Key Laboratory of Behavior and Mental Health, Peking University

⁷ Joint Laboratory of Embodied AI and Humanoid Robots, BIGAI & UniTree Robotics

⁸ Embodied Intelligence Lab, PKU-Wuhan Institute for Artificial Intelligence

<https://humanoid-clone.github.io/>



Figure 1: CLONE employs an MoE-based policy with **closed-loop** error correction for humanoid teleoperation, enabling precise **whole-body coordination** and **long-horizon** task execution.

Abstract: Humanoid teleoperation plays a vital role in demonstrating and collecting data for complex humanoid-scene interactions. However, current teleoperation systems face critical limitations: they decouple upper- and lower-body control to maintain stability, restricting natural coordination, and operate open-loop without real-time position feedback, leading to accumulated drift. The fundamental challenge is achieving precise, coordinated whole-body teleoperation over extended durations while maintaining accurate global positioning. Here we show that an MoE-based teleoperation system, **CLONE**, with closed-loop error correction enables unprecedented whole-body teleoperation fidelity, maintaining minimal positional drift over long-range trajectories using only head and hand tracking from an MR headset. Unlike previous methods that either sacrifice coordination for stability or suffer from unbounded drift, **CLONE** learns diverse motion skills while preventing tracking error accumulation through real-time feedback, enabling complex coordinated movements such as “picking up objects from the ground.” These results establish a new milestone for whole-body humanoid teleoperation for long-horizon humanoid-scene interaction tasks.

Keywords: Humanoid; Whole-body teleoperation; Humanoid-scene interaction

1 Introduction

The ability to seamlessly coordinate whole-body movements while navigating complex environments represents one of humanity’s most remarkable capabilities [1, 2]. From squatting to retrieve objects from the ground to walking across rooms while carrying items, humans effortlessly integrate locomotion and manipulation in ways that remain challenging for robots [3–7]. Humanoid robots (humanoids henceforth), with their human-like morphology, offer the promise of replicating these capabilities—potentially enabling applications from household assistance to operations in hazardous environments where human-like dexterity and mobility are essential [8–12].

However, realizing this potential requires solving a fundamental challenge: enabling intuitive and precise teleoperation that maintains coordination across the entire body over extended periods. Long-horizon tasks, such as navigating to distant locations while manipulating objects, demand not only moment-to-moment stability but also sustained accuracy in both movement execution and global positioning. Current teleoperation approaches fall short of these requirements, creating a significant capability gap between human operators and humanoids.

Recent advances in humanoid teleoperation and loco-manipulation [13–20] have made notable progress. Nevertheless, existing methods struggle with precise teleoperation over extended durations and fall short of the whole-body coordination necessary for humanoid-scene interaction. Two fundamental challenges persist in bridging this capability gap.

The first challenge centers on achieving **coordinated whole-body coordination**. Many systems decouple upper- and lower-body control for stability [18, 21], sacrificing the natural synergies required for fluid motion. While this separation provides safety, it fundamentally limits integrated actions such as reaching while walking or adjusting posture during manipulation. Alternative approaches that rely on motion capture data [13, 15, 19, 22–27] often emphasize stability at the cost of expressiveness, yielding conservative motions constrained by training data distributions. Moreover, these methods consistently overlook key factors like hand orientation that are critical for dexterous tasks, further restricting humanoids’ potential for sophisticated whole-body movements.

The second challenge involves **accumulated positional drift** over time due to the absence of real-time feedback about the robot’s actual position in the environment. Unlike wheeled robots with straightforward odometry, humanoids exhibit complex foot-ground interactions and non-holonomic dynamics that complicate accurate state estimation. Without closed-loop correction, small pose errors compound with each step, progressively degrading the operator’s spatial awareness and control authority, eventually leading to complete task failure. This drift becomes particularly acute during manipulation tasks that require precise positioning relative to environmental objects.

To tackle the above challenges in humanoid long-horizon tasks requiring whole-body coordination and accurate positioning, we present **CLONE**. As illustrated in Fig. 2, **CLONE** is a **closed-loop** whole-body teleoperation system combining learning-based coordination and real-time feedback correction. Our system employs a Mixture-of-Experts (MoE) architecture that learns to coordinate diverse motion skills while a LiDAR-based error correction mechanism prevents the accumulation of positional drift. Critically, **CLONE** requires only head and hand tracking from a single commercial Mixed Reality (MR) headset, making it practical for real-world deployment while achieving unprecedented fidelity in long-horizon tasks.

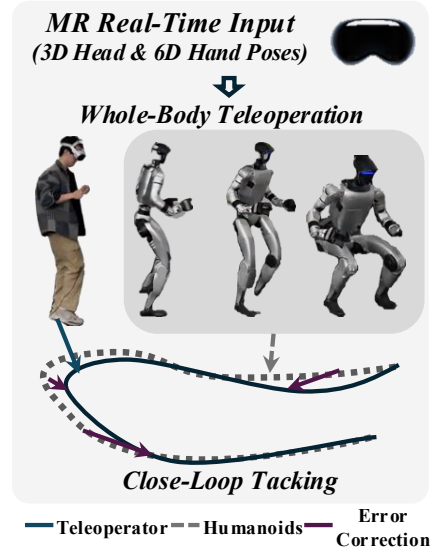


Figure 2: **Whole-body humanoid teleoperation from minimal input.** Our approach enables intuitive control of a humanoid robot using only head and hand poses from mixed reality input, generating coordinated whole-body motions including natural locomotion. Through closed-loop tracking, the system maintains accurate correspondence between operator and robot over extended operation periods, enabling complex long-horizon tasks that require sustained precision.

Our approach integrates three key components: (i) We develop an MoE framework that enables unified learning of diverse motion skills while maintaining natural upper- and lower-body coordination. (ii) We implement closed-loop error correction using LiDAR odometry [28] and Apple Vision Pro (AVP) tracking to provide continuous global pose feedback and prevent drift accumulation. (iii) We curate a comprehensive dataset, **CLONED**, that augments AMASS [29] with additional motion-captured sequences and an online hand orientation generation method, ensuring robust generalization to complex manipulation scenarios involving coordinated whole-body movements.

Our experiments demonstrate that **CLONE** enables capabilities previously unattainable with existing systems: whole-body coordination over long trajectories with minimal positional drift, complex coordinated movements like object retrieval from ground level, and robust performance across diverse operator configurations and environmental conditions. Using only minimal input from a commercial MR headset, **CLONE** achieves improved tracking precision over existing open-loop approaches, opening new possibilities for practical humanoid applications in unstructured environments.

Our contributions are four-fold: (i) the first MoE-based framework for coordinated whole-body teleoperation that maintains natural movement synergies; (ii) a closed-loop system that solves the fundamental position drift problem in long-horizon tasks through real-time pose correction; (iii) a comprehensive dataset, **CLONED**, enabling robust learning of dexterous whole-body motions with proper hand orientation coverage; and (iv) extensive validation demonstrating substantial improvements in real-world humanoid-scene interaction capabilities across diverse scenarios.

2 Related Work

Whole-Body Humanoid Teleoperation Humanoid teleoperation enables robots to replicate human movements for complex tasks using motion capture systems [24, 17, 30], haptic devices [31–33], or virtual reality interfaces [15, 19, 34–36]. A key challenge is developing whole-body control¹ policies that balance robot stability with motion tracking fidelity. Current methods struggle to reproduce the full diversity and fluidity of human motions [41], primarily due to monolithic MLP-based architectures that inadequately handle conflicting objectives across different motion types (*e.g.*, walking *vs.* crouching) [42–44]. Although mixture-based models have shown promise in other domains [45–48], their application to humanoid teleoperation remains underexplored. To address these limitations, we leverage an MoE framework for adaptive learning and unified representation of diverse motion patterns within a single policy.

Long-Horizon Loco-Manipulation Long-horizon task execution [49] has been extensively studied for fixed-base arms [50–53], mobile manipulators [54–58], and aerial manipulators [59], typically in structured settings. In contrast, humanoid teleoperation remains limited to short-horizon motion replication [19, 23, 24], operating open-loop due to difficulties in real-time global state estimation for bipedal systems. Although recent advances in odometry have improved state tracking for legged robots [60, 61], their application to long-horizon humanoid control remains largely unexplored. To bridge this gap, we integrate LiDAR odometry into our teleoperation framework to enable closed-loop error correction and significantly reduce accumulated drift.

Datasets for Training Humanoids Large-scale motion capture (MoCap) datasets [29, 62] have been instrumental in training humanoid control policies [15, 19, 23, 63]. Even after augmenting the datasets with generative models [24], these datasets were still confined primarily to animation and graphics [64, 25] rather than robotics applications. While they contain semantically distinct actions (*e.g.*, waving, hugging, drinking), they underrepresent the kinematic configurations and dynamic transitions required for robust, generalizable controller training in real-world scenarios. To address

¹Whole-Body Control (WBC) in robotics literature [37–39] was traditionally formulated as an optimization problem [3, 40], coordinating multiple competing tasks (such as balance and reaching) through hierarchical control objectives. Learning-based approaches [13, 19, 24] have recently extended this concept by formulating whole-body control as a reinforcement learning problem. We refer to our approach as a whole-body control policy, as it similarly coordinates all degrees of freedom of the humanoid in a unified manner.

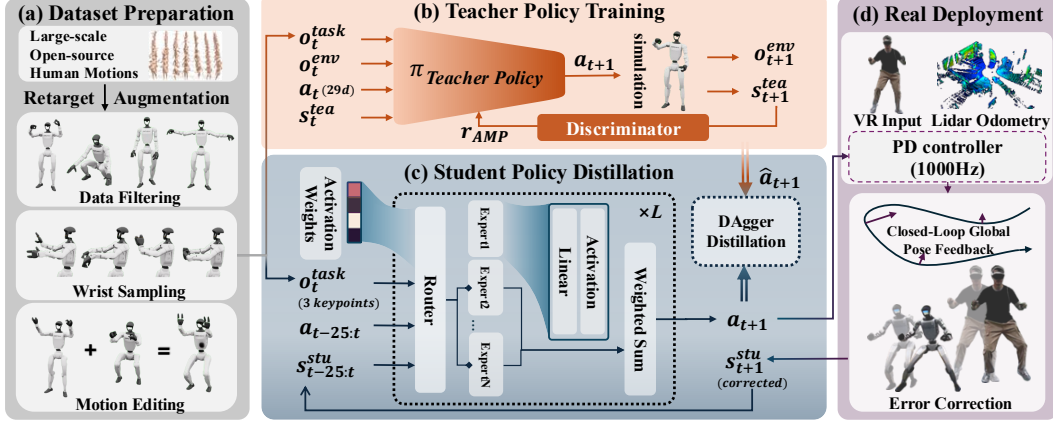


Figure 3: **The CLONE framework.** (a) **CLONED** curates and augments retargeted AMASS [29] data through motion editing to introduce diverse humanoid motions and detailed hand movements. (b) A teacher policy is trained using privileged information, including full robot state and environmental context. (c) An MoE network serves as the student policy, distilled from the teacher to operate with real-world observations only. (d) For real-world deployment, we integrate LiDAR odometry to obtain real-time humanoid states, enabling closed-loop error correction during teleoperation.

these limitations, we introduce **CLONED** by augmenting AMASS [29] through motion editing and collecting additional human MoCap data, specifically tailored for humanoid controllers. This expansion increases coverage of motions and transitions relevant to humanoid control tasks.

3 The CLONE Framework

Our teleoperation framework, **CLONE**, captures a minimal set of control signals from the teleoperator, consisting solely of the 6D poses (position and orientation) of both wrists and the 3D position of the head, tracked using an AVP headset. These three points (see also Fig. 3) serve as the complete control interface, providing an intuitive yet powerful means of directing the humanoid’s whole-body motion while maintaining a simple setup that requires no additional hardware or complex calibration procedures.

CLONE addresses two fundamental challenges through complementary components. First, we develop a teacher-student policy learning approach that transforms these sparse control signals into coordinated whole-body movements (Sec. 3.1). Second, we implement a closed-loop error correction mechanism that maintains positional accuracy during extended operation (Sec. 3.2). The system is supported by carefully designed reward structures and randomization techniques (Sec. 3.3) and trained on a newly curated dataset, **CLONED**, that ensures robust generalization (Sec. 3.4). Additional details of the implementation are provided in Appendix C.

3.1 Policy Learning

We employ a teacher-student training strategy for the teleoperation policy, following the overall framework of OmniH2O [19] (see Appendix A.1 for the problem formulation). This approach first trains a teacher policy with privileged information, then distills this knowledge into a student policy that operates using only real-world observations.

Teacher Policy Training The teacher policy π_{tea} is implemented as an Multi-Layer Perceptrons (MLP) that leverages comprehensive state information unavailable on real robots. At each timestep t , it processes observations $\mathbf{o}_t^{\text{tea}} = [\mathbf{s}_t^{\text{tea}}, \mathbf{o}_t^{\text{task}}, \mathbf{a}_t, \mathbf{o}_t^{\text{env}}]$ and outputs target joint positions $\mathbf{a}_{t+1} \in \mathbb{R}^{29}$ for PD control. The privileged states $\mathbf{s}_t^{\text{tea}} = [\mathbf{p}_t, \theta_t, \mathbf{v}_t, \omega_t]$ include joint angular positions \mathbf{p}_t and the 6D poses, linear velocities, and angular velocities $\theta_t, \mathbf{v}_t, \omega_t$ of all robot links. Task observations $\mathbf{o}_t^{\text{task}} = [\hat{\mathbf{p}}_{t+1} - \mathbf{p}_t, \hat{\theta}_{t+1} - \theta_t, \hat{\mathbf{v}}_{t+1} - \mathbf{v}_t, \hat{\omega}_{t+1} - \omega_t, \hat{\mathbf{p}}_{t+1}, \hat{\theta}_{t+1}]$ capture both reference motion

(denoted by $\hat{\cdot}$) and tracking errors between reference and current states. Environmental observations $\mathbf{o}_t^{\text{env}}$ provide context including ground friction coefficient and robot mass distribution.

Student Policy Distillation The student policy must operate without privileged information, following $a_{t+1} = \pi_{\text{stu}}(s_{t-25:t}^{\text{stu}}, a_{t-25:t}, o_t^{\text{task}})$. The robot state sequence $s_{t-25:t}^{\text{stu}}$ contains joint positions q , joint velocities \dot{q} , root angular velocity ω^{root} , and root gravity vector g obtained from on-device IMU over the past 25 frames. Task observations o_t^{task} consist of $\hat{p}_{t+1} - p_t$, \hat{p}_{t+1} , \hat{p}_{t+1} , h_t , and \hat{h}_{t+1} , where p_t represents the 3D positions of head and two wrists obtained from LiDAR odometry and forward kinematics, \hat{p}_{t+1} and \hat{p}_{t+1} are target positions and velocities from reference motion, and h_t , \hat{h}_{t+1} represent current and target wrist orientations.

The key challenge lies in handling diverse motion patterns within a single policy. Walking requires different control strategies than crouching or reaching, yet traditional monolithic architectures struggle with these conflicting objectives. We address this through an MoE architecture, as shown in Fig. 3, which allows specialized processing for different motion types.

The MoE design consists of L layers, each comprising N experts that function as independent feed-forward sub-layers with distinct parameters. At each layer, a router dynamically selects which experts are activated based on the input, generating weight distributions over all experts. The layer output combines the top- k experts with highest routing weights: $f = \sum_i^k w_i \cdot E_i(\cdot)$, where w_i is the routing weight for the i -th selected expert and $E_i(\cdot)$ is the output of the i -th expert. This design enables different experts to focus on distinct motion patterns. To prevent model collapse to only a few experts, we introduce a balancing loss that encourages uniform expert selection:

$$\mathcal{L}_{\text{balance}} = \sum_{l=1}^L \sum_{e=1}^N [\max(p_e - \frac{1+\epsilon}{N}, 0) + \min(\frac{1-\epsilon}{N} - p_e, 0)], \quad (1)$$

where $p_e = \mathbb{E}[w_e]$ represents the expected activation probability of expert e , and ϵ is a slack constant that allows slight deviations from perfect uniformity.

3.2 Closed-Loop Error Correction

Traditional humanoid teleoperation systems operate in an open-loop configuration, where small errors in position tracking accumulate over time, leading to significant drift during extended operations. This fundamental limitation becomes particularly problematic during long-horizon tasks that require sustained positional accuracy. To address this challenge, we implement a closed-loop error correction mechanism that continuously monitors and compensates for positional discrepancies between the teleoperator and the humanoid.

Our approach utilizes LiDAR odometry to maintain accurate global position estimates for both the humanoid and the teleoperator. We employ FAST-LIO2 [28], an algorithm that tightly couples IMU and LiDAR data through an iterated Kalman filter to provide robust real-time state estimation even during dynamic movements (see more details in Appendix A.2). This choice ensures reliable tracking performance across diverse motion patterns, from walking to complex manipulation tasks.

The system tracks global positions for both agents: the humanoid’s position $p \in \mathbb{R}^3$ is computed from onboard sensors, while the teleoperator’s position $\hat{p} \in \mathbb{R}^3$ is similarly tracked through a MR hardware equipped with a comparable odometry pipeline. The student teleoperation policy directly consumes the difference between p and \hat{p} , enabling it to generate actions that systematically reduce positional drift and maintain accurate correspondence between the operator and the humanoid.

3.3 Reward Design and Domain Randomization

We build upon the reward terms and domain randomizations from OmniH2O [19] as the foundation of our approach, with specific enhancements to address the challenges of real-world teleoperation. Detailed reward functions and domain randomization settings are provided in Appendix B.

To enhance robustness against LiDAR odometry errors, we introduce a velocity-dependent Stochastic Differential Equation (SDE) noise model during training that reflects real-world error characteristics. For the head position \vec{p}_{head} , we define the randomized position \vec{P}_{head} as:

$$d\vec{P}_{\text{head}} = \dot{\vec{p}}_{\text{head}}dt + \left(\frac{\|\dot{\vec{p}}_{\text{head}}\|}{c_{\text{vel}}} + c_{\text{min}} \right) d\vec{W}, \quad (2)$$

where \vec{W} is a standard Wiener process, and c_{vel} and c_{min} are constants that scale the noise proportionally to movement speed and establish a minimum randomization level. This formulation mirrors real-world dynamics, where faster movements tend to produce greater odometry errors. We use forward kinematics to compute other body positions based on the randomized head position, while periodically resetting and constraining the maximum deviation to avoid unrealistic drift.

Since **CLONE** provides only upper-body references (head and wrists), we must generate appropriate lower-body behaviors without explicit guidance. To tackle this challenge, we employ an Adversarial Motion Priors (AMP) reward [65] to regularize lower-body movements and encourage natural, stable behavior. Through this combination of specialized domain randomization and reward design, **CLONE** learns to generate robust lower-body behaviors while maintaining precise upper-body control aligned with operator commands.

3.4 The CLONED Dataset

The training dataset **CLONED** comprises three complementary components to support robust whole-body teleoperation. These include: (i) an augmented AMASS [29] subset of 149 curated sequences featuring diverse pairings of upper- and lower-body movements, enhanced via targeted motion editing to increase compositional diversity and policy generalization; (ii) 14 custom sequences captured with an IMU-based Xsens MoCap system to fill coverage gaps, emphasizing continuous transitions and diverse upper-body poses critical for manipulation; and (iii) systematic hand orientation augmentation through procedurally generated 6D wrist targets, smoothed via Spherical Linear Interpolation (SLERP) to ensure coherent and natural hand motions for teleoperation.

4 Real-World Experiments

We evaluated **CLONE** on a physical Unitree G1 humanoid through comprehensive experiments demonstrating exceptional whole-body motion fidelity and precise position tracking. Our experiments focus on two key capabilities: (i) global position tracking accuracy during extended teleoperation to validate our closed-loop error correction mechanism, and (ii) whole-body motion tracking fidelity across diverse skills to demonstrate coordination capabilities. These experiments collectively validate both the technical performance and practical applicability of our approach to real-world humanoid teleoperation.




Figure 4: **Global position tracking accuracy in real-world experiments.** **CLONE** achieves mean tracking errors of 5.1cm across distances up to 8.9m, demonstrating effective closed-loop error correction in extended teleoperation.

Global Position Tracking To evaluate global positioning accuracy over extended distances, we designed a controlled path-following experiment. We marked fixed initial positions for both operator and robot, then had the operator walk along straight paths to target positions at 3m, 6m, and 8.9m distances while teleoperating the robot. We measured the discrepancy between the robot’s final position and expected position as tracking error, repeating each condition ten times.

CLONE achieved a mean tracking error of 5.1cm with a maximum deviation of 12.0cm at 8.9m, as shown in Fig. 4. This exceptional positional accuracy validates that **CLONE**’s closed-loop error correction effectively mitigates accumulated positional errors during extended teleoperation.

Statistical analysis confirms consistent tracking performance across all tested distances. Independent samples t-tests comparing distance groups yielded: 3m versus 6m ($t = 0.165, p = 0.871$) and 6m




Figure 5: **Whole-body motion tracking on Unitree G1.** CLONE successfully tracks diverse skills including (a) waving, (b)(d) squatting, and (c) jumping, showcasing comprehensive whole-body coordination capabilities.

versus 8.9m ($t = 0.048, p = 0.963$). With all p-values > 0.05 , **CLONE** shows no significant performance degradation across the tested range, demonstrating that our approach maintains high accuracy over extended distances without drift accumulation.

Whole-Body Motion Tracking As shown in Fig. 5, **CLONE** successfully enables real-time teleoperation across a diverse range of whole-body skills. The robot accurately tracks complex motions including waving, squatting, standing up from squatted positions, and jumping. These results demonstrate high whole-body motion fidelity for real-time humanoid teleoperation, particularly for dynamic skills like jumping requiring precise coordination of balance control and force application.

Long-Horizon Mixed Navigation To validate system performance in complex scenarios, we conducted extended teleoperation sessions incorporating multiple movement types. As visualized in Fig. 6, we recorded a continuous teleoperation sequence where the operator traversed a complex path spanning over 15m, incorporating diverse locomotion patterns including forward walking, turning, side-stepping, and returning to the original position.

The robot consistently tracked the operator’s movements throughout this extended sequence and returned to its starting position with minimal drift. This result validates **CLONE**’s robustness for extended teleoperation sessions that combine locomotion and whole-body motion control—a critical capability for practical humanoid applications requiring sustained coordination over long horizons.

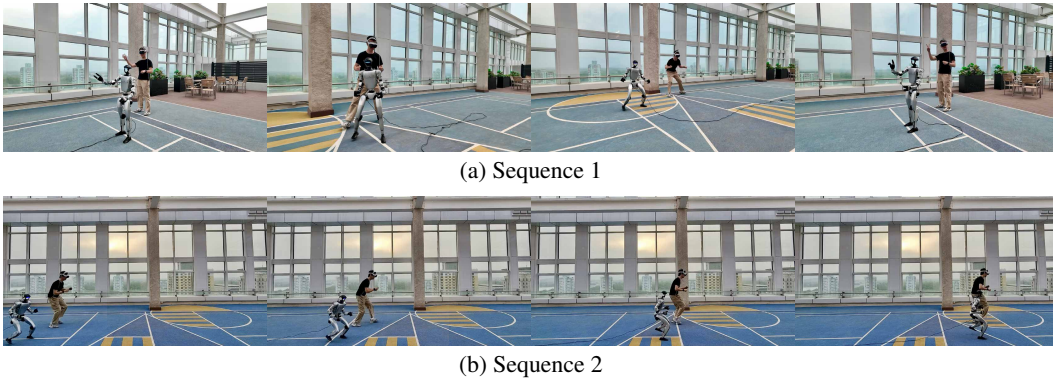


Figure 6: **Long-horizon teleoperation.** The humanoid accurately tracks both the operator’s local pose and global translation throughout a complex navigation sequence, demonstrating robust performance.

5 Simulations

We present comprehensive evaluations of **CLONE** in simulation across four key settings: reference motion tracking, diverse stance tracking, ablation studies (Appendix D.2), and expert activation analysis (Appendix D.4). These experiments are designed to: (i) quantify motion tracking accuracy in controlled simulation environments, and (ii) assess robustness across diverse stance configurations. The evaluation metrics are detailed in Appendix D.1.

Motion Tracking We compared **CLONE** against two ablated baselines: **CLONE**[†] and **CLONE**^{*}. **CLONE**[†] employs an MLP as the student policy, resembling the OmniH2O baseline trained on our data and task. **CLONE**^{*} represents our **CLONE** trained on OmniH2O data instead of **CLONED**. Quantitative results in Tab. 1 reveal that both the MoE architecture and **CLONED** contribute significantly to accurate reference motion tracking. Qualitative comparisons are provided in Appendix D.3.

Tracking Diverse Stances To assess **CLONE**’s robustness across varying postures, we evaluated performance in tracking motions with head heights from 1.2m (standing) to 0.6m (deep squatting) in 0.1m decrements. We generated these challenging motions by systematically editing sequences from the **CLONED** dataset, creating unseen poses that test teleoperation system limits.

Fig. 7 reveals an interesting performance trade-off: while **CLONE** underperforms baselines in absolute position accuracy (MPKPE), it consistently outperforms them in local motion metrics (R-MPKPE, velocity error, and hand orientation). This pattern indicates that **CLONE** prioritizes faithful reproduction of reference motions—particularly for challenging postures—sometimes at the expense of global positioning accuracy. All methods exhibit increased tracking errors at lower heights, confirming the inherent difficulty of teleoperating robots in squatting postures.

Table 1: **Motion tracking evaluation on CLONED dataset.** Comparison of **CLONE** against ablations: **CLONE**[†] uses an MLP instead of MoE architecture, while **CLONE**^{*} trains on OmniH2O data instead of **CLONED**.

Method	SR \uparrow	$E_{\text{mpkpe}} \downarrow$	$E_{\text{r-mpkpe}} \downarrow$	$E_{\text{vel}} \downarrow$	$E_{\text{hand-rot}} \downarrow$
CLONE [†]	100%	113.97	35.55	245.11	4.73
CLONE [*]	100%	102.20	41.07	309.65	4.61
CLONE	100%	87.84	33.30	227.17	3.61




Figure 7: **Motion tracking performance across stance heights.** Comparison between **CLONE** (blue solid), **CLONE**^{*} (green dashed), and **CLONE**[†] (red dashed) across different postures from standing to deep squatting. Lower values indicate better performance.

6 Conclusion and Discussions

We present **CLONE**, a closed-loop MoE-based teleoperation system that enables comprehensive humanoid control while addressing accumulated tracking errors in long-horizon tasks. Our approach integrates three key innovations: an MoE architecture coordinating diverse motion patterns, LiDAR-based closed-loop error correction preventing positional drift, and the **CLONED** dataset providing robust whole-body coordination training. Experimental validation demonstrates exceptional performance: precise global position tracking with 5.1cm mean error over 8.9m distances, accurate whole-body motion coordination across diverse skills, and robust long-horizon teleoperation using only minimal MR headset input.

While demonstrating significant capabilities, important limitations require future investigation. The minimal input configuration, though enabling intuitive control with head and hand tracking, inherently limits fine-grained stability control in certain scenarios. Additionally, the system exhibits reduced performance during highly dynamic movements like jumping, stemming from training data constraints and balance control challenges.

Future work should explore additional sensing modalities to enhance stability while preserving interface simplicity, and expand motion datasets with specialized reward functions for dynamic behaviors. **CLONE** establishes new benchmarks for practical humanoid teleoperation, achieving unprecedented fidelity in complex tasks while maintaining essential simplicity.

Acknowledgments

The authors gratefully acknowledge Unitree Robotics for their support with hardware. This work is supported in part by the National Science and Technology Major Project (2022ZD0114900), the National Natural Science Foundation of China (62376031), the Beijing Nova Program, the State Key Lab of General AI at Peking University, the PKU-Bingji Joint Laboratory for Artificial Intelligence, and the National Comprehensive Experimental Base for Governance of Intelligent Society, Wuhan East Lake High-Tech Development Zone.

References

- [1] B. Henze, M. A. Roa, and C. Ott. Passivity-based whole-body balancing for torque-controlled humanoid robots in multi-contact scenarios. *International Journal of Robotics Research (IJRR)*, 35(12):1522–1543, 2016.
- [2] P. M. Wensing and D. E. Orin. Improved computation of the humanoid centroidal dynamics and application for whole-body control. *International Journal of Humanoid Robotics*, 13(01):1550039, 2016.
- [3] L. Sentis and O. Khatib. Synthesis of whole-body behaviors through hierarchical control of behavioral primitives. *International Journal of Humanoid Robotics*, 2(04):505–518, 2005.
- [4] O. Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal on Robotics and Automation*, 3(1):43–53, 2003.
- [5] T. Fukuda, P. Dario, and G.-Z. Yang. Humanoid robotics—history, current state of the art, and challenges. *Science Robotics*, 2(13):eaar4043, 2017.
- [6] A. Hereid, C. M. Hubicki, E. A. Cousineau, and A. D. Ames. Dynamic humanoid locomotion: A scalable formulation for hzd gait optimization. *IEEE Transactions on Robotics (T-RO)*, 34(2):370–387, 2018.
- [7] J. Cui, T. Liu, Z. Meng, J. Yu, R. Song, W. Zhang, Y. Zhu, and S. Huang. Grove: A generalized reward for learning open-vocabulary physical skill. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [8] Y. Tong, H. Liu, and Z. Zhang. Advancements in humanoid robots: A comprehensive review and future prospects. *IEEE/CAA Journal of Automatica Sinica*, 11(2):301–328, 2024.
- [9] Vecna Robotics. Battlefield Extraction-Assist Robot (BEAR). <https://robotsguide.com/robots/bear>, 2006.
- [10] U.S. Naval Research Laboratory. Autonomous Shipboard Humanoid (ASH). <https://www.navy.mil/Resources/Fact-Files/Display-FactFiles/Article/2160601/shipboard-autonomous-firefighting-robot-saffir/>, 2012.
- [11] AGIBOT Robotics. AGIBOT A2 Humanoid Robot. <https://www.agibot.com/products/A2>, 2024.
- [12] Z. Zhao, W. Li, Y. Li, T. Liu, B. Li, M. Wang, K. Du, H. Liu, Y. Zhu, Q. Wang, K. Althoefer, and S.-C. Zhu. Embedding high-resolution touch across robotic hands enables adaptive human-like grasping, 2024.

- [13] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn. Humanplus: Humanoid shadowing and imitation from humans. In *Conference on Robot Learning (CoRL)*, 2024.
- [14] Y. Ze, Z. Chen, W. Wang, T. Chen, X. He, Y. Yuan, X. B. Peng, and J. Wu. Generalizable humanoid manipulation with improved 3d diffusion policies, 2024.
- [15] T. He, W. Xiao, T. Lin, Z. Luo, Z. Xu, Z. Jiang, C. Liu, G. Shi, X. Wang, L. Fan, and Y. Zhu. Hover: Versatile neural whole-body controller for humanoid robots, 2024.
- [16] J. Cui, T. Liu, N. Liu, Y. Yang, Y. Zhu, and S. Huang. Anyskill: Learning open-vocabulary physical skill for interactive agents. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [17] S. Daffarra, K. Darvish, R. Grieco, G. Milani, U. Pattacini, L. Rapetti, G. Romualdi, M. Salvi, A. Scalzo, I. Sorrentino, et al. icub3 avatar system: Enabling remote fully immersive embodiment of humanoid robots. *Science Robotics*, 9(86):eadh3834, 2024.
- [18] Q. Ben, F. Jia, J. Zeng, J. Dong, D. Lin, and J. Pang. Homie: Humanoid loco-manipulation with isomorphic exoskeleton cockpit. In *Robotics: Science and Systems (RSS)*, 2025.
- [19] T. He, Z. Luo, X. He, W. Xiao, C. Zhang, W. Zhang, K. Kitani, C. Liu, and G. Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. In *Conference on Robot Learning (CoRL)*, 2024.
- [20] T. Haarnoja, B. Moran, G. Lever, S. H. Huang, D. Tirumala, J. Humplik, M. Wulfmeier, S. Tulyasuvunakool, N. Y. Siegel, R. Hafner, et al. Learning agile soccer skills for a bipedal robot with deep reinforcement learning. *Science Robotics*, 9(89):eadi8022, 2024.
- [21] A. Matsiko. Humanoid robot learning of complex behaviors with llms. *Science Robotics*, 10(98):eadv4627, 2025.
- [22] T. He, Z. Luo, W. Xiao, C. Zhang, K. Kitani, C. Liu, and G. Shi. Learning human-to-humanoid real-time whole-body teleoperation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [23] X. Cheng, Y. Ji, J. Chen, R. Yang, G. Yang, and X. Wang. Expressive whole-body control for humanoid robots. In *Robotics: Science and Systems (RSS)*, 2024.
- [24] M. Ji, X. Peng, F. Liu, J. Li, G. Yang, X. Cheng, and X. Wang. Exbody2: Advanced expressive humanoid whole-body control, 2024.
- [25] N. Jiang, Z. Zhang, H. Li, X. Ma, Z. Wang, Y. Chen, T. Liu, Y. Zhu, and S. Huang. Scaling up dynamic human-scene interaction modeling. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2024.
- [26] N. Jiang, Z. He, H. Li, Y. Chen, S. Huang, and Y. Zhu. Autonomous character-scene interaction synthesis from text instruction. In *SIGGRAPH Asia Conference Papers*, 2024.
- [27] N. Jiang, H. Li, Z. Yuan, Z. He, Y. Chen, T. Liu, Y. Zhu, and S. Huang. Dynamic motion blending for versatile motion editing. In *Proceedings of Conference on Computer Vision and Pattern Recognition (CVPR)*, 2025.
- [28] W. Xu, Y. Cai, D. He, J. Lin, and F. Zhang. Fast-lid2: Fast direct lidar-inertial odometry. *IEEE Transactions on Robotics (T-RO)*, 38(4):2053–2073, 2022.
- [29] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black. AMASS: Archive of motion capture as surface shapes. In *Proceedings of International Conference on Computer Vision (ICCV)*, 2019.

- [30] K. Darvish, Y. Tirupachuri, G. Romualdi, L. Rapetti, D. Ferigo, F. J. A. Chavez, and D. Pucci. Whole-body geometric retargeting for humanoid robots. In *International Conference on Humanoid Robots (Humanoids)*, 2019.
- [31] A. Brygo, I. Sarakoglou, N. Garcia-Hernandez, and N. Tsagarakis. Humanoid robot teleoperation with vibrotactile based balancing feedback. In *Haptics: Neuroscience, Devices, Modeling, and Applications*, 2014.
- [32] L. Peternel and J. Babič. Learning of compliant human–robot interaction using full-body haptic interface. *Advanced Robotics*, 27(13):1003–1012, 2013.
- [33] J. Ramos and S. Kim. Dynamic locomotion synchronization of bipedal robot and human operator via bilateral feedback teleoperation. *Science Robotics*, 4(35):eaav4282, 2019.
- [34] J. Chagas Vaz, D. Wallace, and P. Y. Oh. Humanoid loco-manipulation of pushed carts utilizing virtual reality teleoperation. In *International Mechanical Engineering Congress and Exposition*, 2021.
- [35] L. Penco, N. Scianca, V. Modugno, L. Lanari, G. Oriolo, and S. Ivaldi. A multimode teleoperation framework for humanoid loco-manipulation: An application for the icub robot. *IEEE Robotics & Automation Magazine*, 26(4):73–82, 2019.
- [36] S. Tachi, Y. Inoue, and F. Kato. Telesar vi: Telexistence surrogate anthropomorphic robot vi. *International Journal of Humanoid Robotics*, 17(05):2050019, 2020.
- [37] Y. Nakamura, H. Hanafusa, and T. Yoshikawa. Task-priority based redundancy control of robot manipulators. *International Journal of Robotics Research (IJRR)*, 6(2):3–15, 1987.
- [38] O. Khatib, L. Sentis, J. Park, and J. Warren. Whole-body dynamic behavior and control of human-like robots. *International Journal of Humanoid Robotics*, 1(01):29–43, 2004.
- [39] A. Dietrich, C. Ott, and A. Albu-Schäffer. An overview of null space projections for redundant, torque-controlled robots. *International Journal of Robotics Research (IJRR)*, 34(11):1385–1400, 2015.
- [40] F. L. Moro and L. Sentis. Whole-body control of humanoid robots. *Humanoid robotics: a reference*, pages 1161–1183, 2019.
- [41] M. Moniruzzaman, A. Rassau, D. Chai, and S. M. S. Islam. Teleoperation methods and enhancement techniques for mobile robots: A comprehensive survey. *Robotics and Autonomous Systems*, 150:103973, 2022.
- [42] R. Huang, S. Zhu, Y. Du, and H. Zhao. Moe-loco: Mixture of experts for multitask locomotion, 2025.
- [43] S. Zhou, W. Zhang, J. Jiang, W. Zhong, J. GU, and W. Zhu. On the convergence of stochastic multi-objective gradient manipulation and beyond. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [44] K. Darvish, L. Penco, J. Ramos, R. Cisneros, J. Pratt, E. Yoshida, S. Ivaldi, and D. Pucci. Teleoperation of humanoid robots: A survey. *IEEE Transactions on Robotics (T-RO)*, 39(3): 1706–1727, 2023.
- [45] C. Yang, K. Yuan, Q. Zhu, W. Yu, and Z. Li. Multi-expert learning of adaptive legged locomotion. *Science Robotics*, 5(49):eabb2174, 2020.
- [46] Z. Xie, S. Starke, H. Y. Ling, and M. van de Panne. Learning soccer juggling skills with layer-wise mixture-of-experts. In *ACM SIGGRAPH / Eurographics Symposium on Computer Animation (SCA)*, 2022.

- [47] W. Song, H. Zhao, P. Ding, C. Cui, S. Lyu, Y. Fan, and D. Wang. Germ: A generalist robotic model with mixture-of-experts for quadruped robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.
- [48] G. Cheng, L. Dong, W. Cai, and C. Sun. Multi-task reinforcement learning with attention-based mixture of experts. *IEEE Robotics and Automation Letters (RA-L)*, 8(6):3812–3819, 2023.
- [49] C. R. Garrett, R. Chitnis, R. Holladay, B. Kim, T. Silver, L. P. Kaelbling, and T. Lozano-Pérez. Integrated task and motion planning. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):265–293, 2021.
- [50] C. Wang, L. Fan, J. Sun, R. Zhang, L. Fei-Fei, D. Xu, Y. Zhu, and A. Anandkumar. Mimicplay: Long-horizon imitation learning by watching human play. In *Conference on Robot Learning (CoRL)*, 2023.
- [51] Z. Lin, Y. Chen, and Z. Liu. Hierarchical human-to-robot imitation learning for long-horizon tasks via cross-domain skill alignment. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [52] H. Shi, H. Xu, S. Clarke, Y. Li, and J. Wu. Robocook: Long-horizon elasto-plastic object manipulation with diverse tools. *Conference on Robot Learning (CoRL)*, 2023.
- [53] Z. Zhao, Y. Li, W. Li, Z. Qi, L. Ruan, Y. Zhu, and K. Althoefer. Tac-man: Tactile-informed prior-free manipulation of articulated objects. *IEEE Transactions on Robotics (T-RO)*, 41: 538–557, 2024.
- [54] Z. Jiao, Z. Zeyu, W. Wang, D. Han, S.-C. Zhu, Y. Zhu, and H. Liu. Efficient task planning for mobile manipulation: a virtual kinematic chain perspective. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [55] Z. Jiao, Z. Zeyu, X. Jiang, D. Han, S.-C. Zhu, Y. Zhu, and H. Liu. Consolidating kinematic models to promote coordinated mobile manipulations. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021.
- [56] Z. Jiao, Y. Niu, Z. Zhang, S.-C. Zhu, Y. Zhu, and H. Liu. Planning sequential tasks on contact graph. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2022.
- [57] P. Zhi, Z. Zhang, M. Han, Z. Zhang, Z. Li, Z. Jiao, B. Jia, and S. Huang. Closed-loop open-vocabulary mobile manipulation with gpt-4v. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2024.
- [58] P. Zhi, P. Li, J. Yin, B. Jia, and S. Huang. Learning unified force and position control for legged loco-manipulation, 2025.
- [59] Y. Su, J. Li, Z. Jiao, M. Wang, C. Chu, H. Li, Y. Zhu, and H. Liu. Sequential manipulation planning for over-actuated unmanned aerial manipulators. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2023.
- [60] D. Wisth, M. Camurri, and M. Fallon. Vilens: Visual, inertial, lidar, and leg odometry for all-terrain legged robots. *IEEE Transactions on Robotics (T-RO)*, 39(1):309–326, 2022.
- [61] G. Ou, D. Li, and H. Li. Leg-kilo: Robust kinematic-inertial-lidar odometry for dynamic legged robots. *IEEE Robotics and Automation Letters (RA-L)*, 2024.
- [62] F. G. Harvey, M. Yurick, D. Nowrouzezahrai, and C. Pal. Robust motion in-betweening. *ACM Transactions on Graphics (TOG)*, 39(4):60–1, 2020.
- [63] L. Ma, Z. Meng, T. Liu, Y. Li, R. Song, W. Zhang, and S. Huang. Styleloco: Generative adversarial distillation for natural humanoid robot locomotion, 2025.

- [64] Z. Wang, Y. Chen, T. Liu, Y. Zhu, W. Liang, and S. Huang. Humanise: Language-conditioned human motion generation in 3d scenes. In *Proceedings of Advances in Neural Information Processing Systems (NeurIPS)*, 2022.
- [65] X. B. Peng, Z. Ma, P. Abbeel, S. Levine, and A. Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (TOG)*, 40(4): 1–20, 2021.

A Preliminaries

A.1 Formulation

We formulate the humanoid teleoperation as a Markov Decision Process (MDP) $\mathcal{M} = \{\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}\}$. \mathcal{S} includes proprioceptive states s and task-oriented observations o^{task} . The action space $\mathcal{A} \in \mathbb{R}^{29}$ represents the humanoid’s joint angles in our method. \mathcal{T} is the transition function conditioned on \mathcal{S} and \mathcal{A} . The reward functions \mathcal{R} is defined based on \mathcal{S}, \mathcal{A} . A policy π is proposed to maximize the overall reward \mathcal{R} using the Proximal Policy Optimization (PPO) algorithm.

A.2 LiDAR Odometry and Closed-Loop Error Correction

LiDAR odometry is designed to accurately determine the robot’s state, including its orientation and position. In this paper, we adopt FAST-LIO2 [28], which utilizes onboard LiDAR and IMU to estimate the humanoid’s global position. FAST-LIO2 [28] leverages IMU data and LiDAR point clouds to construct and update a 3D map in real time. It then registers the current LiDAR point clouds with the map to estimate the robot’s current position.

Previous teleoperation systems [19, 23, 24] often operate in an open-loop manner, primarily due to the absence of the humanoid’s global position. Consequently, stepwise tracking errors accumulate over time, leading to significant drift during long-horizon tasks. In this work, we leverage LiDAR odometry to determine the robot’s global position. Similarly, we obtain the operator’s global position with the odometry from Apple Vision Pro. We integrate the difference between the two positions into the task observation o^{task} , and design a reward function for our teleoperation policy that minimizes this difference. Notably, the LiDAR operates at 10 Hz, and our policy runs at 50 Hz. Our policy uses the latest available odometry position at each timestep.

B Reward Functions and Domain Randomization

Tab. A1 provides a detailed overview of the reward structure utilized in this study, while Tab. A2 outlines the domain randomization scheme employed.

C Implement Details

We train our policy in IsaacGym using a single A800 GPU. The teacher policy is trained for $1M$ iterations with 8192 parallel environments, while the student policy is trained for $600K$ iterations with 4096 parallel environments. In our implementation, the student policy consists of $L = 3$ MoE layers, each containing $N = 4$ experts. It uses a history length of $H = 25$ frames and activates the top $k = 2$ experts based on the highest weights determined by the router.

D Experiments

D.1 Evaluation Metrics

We evaluated CLONE on motion tracking tasks from CLONED using five metrics: success rate **SR** (%), mean per-keybody position error (MPKPE) E_{mpkpe} (mm), root-relative mean per-keybody position error (R-MPKPE) $E_{\text{r-mpkpe}}$ (mm), average joint velocity error E_{vel} (mm/s), and hand orientation tracking error E_{hand} . Success rate (**SR**) represents the proportion of episodes where: (i) the robot maintains balance without falling, and (ii) the average per-keybody distance between the robot and reference motion remains below 1.5m across the three controlled joints. We defined the hand orientation tracking error as $E_{\text{hand}} = 1 - \langle \hat{\mathbf{q}}, \mathbf{q} \rangle^2$, where $\hat{\mathbf{q}}$ and \mathbf{q} represent the reference and robot hand quaternions.

D.2 Ablation Study

Table A1: **Reward functions.** The details of the primary reward function used in our training process.

Term	Expression	Weight
Torque	$\ \boldsymbol{\tau}\ _2^2$	-0.0001
Torque limits	$[\tau \notin [\tau_{\min}, \tau_{\max}]]_1$	-2
DoF position limits	$[\mathbf{p}_t \notin [\mathbf{p}_{\min}, \mathbf{p}_{\max}]]_1$	-625
DoF velocity limits	$[\dot{\mathbf{p}}_t \notin [\dot{\mathbf{p}}_{\min}, \dot{\mathbf{p}}_{\max}]]_1$	-50
Termination	termination_1	$-e^4$
DoF acceleration	$\ \ddot{\mathbf{q}}_t\ _2^2$	$-1.1e^{-5}$
DoF velocity	$\ \dot{\mathbf{q}}_t\ _2^2$	-0.004
Lower-body action rate	$\ \mathbf{a}_t^{\text{lower}} - \mathbf{a}_{t-1}^{\text{lower}}\ _2^2$	-1.0
Upper-body action rate	$\ \mathbf{a}_t^{\text{upper}} - \mathbf{a}_{t-1}^{\text{upper}}\ _2^2$	-0.3
Feet air time	$T_{\text{air}} - 0.3$	2500
Stumble	$[(\mathbf{F}_{\text{feet}}^{xy} > 5 \times \mathbf{F}_{\text{feet}}^z)]_1$	$-1.25e^4$
Slippage	$\ v_t^{\text{feet}}\ _2^2 \cdot [(\mathbf{F}_{\text{feet}} \geq 1)]_1$	-80
Feet orientation	$\ \mathbf{R}_{\text{feet}}^z\ $	-62.5
In the air	$[(\mathbf{F}_{\text{feet}}^{\text{left}}, \mathbf{F}_{\text{feet}}^{\text{right}} < 1)]_1$	-750
Orientation	$\ \mathbf{R}_{\text{root}}^z\ $	-50
DoF position	$\exp(-0.25\ \hat{\mathbf{p}} - \mathbf{p}\ _2)$	100
DoF velocity	$\exp(-0.25\ \hat{\dot{\mathbf{p}}} - \dot{\mathbf{p}}\ _2^2)$	10
Extend body position	$\exp(-0.5\ \hat{\mathbf{q}} - \mathbf{q}\ _2^2)$	250
Body position (MR)	$\exp(-0.5\ \mathbf{q}_{\text{vr}} - \hat{\mathbf{q}}_{\text{vr}}\ _2^2)$	150
Body rotation	$\exp(-0.1\ \theta \ominus \hat{\theta}\)$	400
Body velocity	$\exp(-10.0\ \mathbf{v} - \hat{\mathbf{v}}\ _2)$	80
Body angular velocity	$\exp(-0.01\ \boldsymbol{\omega} - \hat{\boldsymbol{\omega}}\ _2)$	8
Body hand rotation	$(\theta_{\text{hand}} - \hat{\theta}_{\text{hand}})^2$	500
AMP	Sec. 3.3	500

Table A2: **Domain Randomization.** The details of the primary domain randomization used in our training process.

Term	Value
Friction	$\mathcal{U}(0.6, 2.0)$
Base CoM offset	$\mathcal{U}(-0.04, 0.04)\text{m}$
Link mass	$\mathcal{U}(0.7, 1.25) \times \text{default kg}$
P Gain	$\mathcal{U}(0.85, 1.15) \times \text{default}$
D Gain	$\mathcal{U}(0.85, 1.15) \times \text{default}$
Torque RFI	$0.05 \times \text{torque limit N} \cdot \text{m}$
Control delay	$\mathcal{U}(0.0, 20)\text{ms}$
Global step delay	$\mathcal{U}(0.0, 80)\text{ms}$
Rand born distance	$\mathcal{U}(0.0, 2.0)\text{m}$
Rand heading degree	$\mathcal{U}(-20.0, 20.0)\text{degree}$
Push robot	interval = 5s, $v_{xy} = 1.5\text{m/s}$
Terrain type	flat, rough, low obstacles [19]

We investigated the impact of key design choices, specifically history length and MoE parameters, through systematic ablation experiments reported in [Tab. A3](#). Our results indicate that a configuration using 25 timesteps of history, three MoE layers, and four experts per layer yields optimal performance across most evaluation metrics. We observed that shorter history lengths and increased expert counts can produce marginally lower R-MPKPE values and larger global tracking errors, suggesting a trade-off between local and global motion fidelity.

Table A3: Ablation study on history length and architecture components.

Method	E_{mpkpe}	$E_{\text{r-mpkpe}}$	E_{vel}	$E_{\text{hand-rot}}$
(a) History Length Analysis				
History5	93.97	31.99	236.12	3.80
History50	135.60	41.33	286.66	12.23
History25(CLONE)	87.84	33.30	227.17	3.61
(b) Architecture Ablation				
CLONE ($L = 1$)	134.06	37.56	270.14	7.22
CLONE ($N = 8$)	89.21	30.90	251.10	4.26
CLONE	87.84	33.30	227.17	3.61




Figure A1: Qualitative Results of **CLONE** and **CLONE***. (a) and (b) show the “crouch” tracking results of **CLONE***, while (c) and (d) present the results of **CLONE**.

		Layer 1				Layer 2				Layer 3			
Category	squat	0.08	0.08	0.42	0.42	0.15	0.14	0.36	0.35	0.15	0.36	0.15	0.35
	crouchwalk	0.06	0.07	0.44	0.43	0.15	0.14	0.36	0.36	0.14	0.36	0.14	0.35
	turn	0.19	0.20	0.31	0.30	0.16	0.15	0.35	0.34	0.16	0.35	0.15	0.33
	punch	0.23	0.24	0.27	0.26	0.21	0.18	0.32	0.30	0.20	0.32	0.18	0.29
	jump	0.29	0.30	0.21	0.20	0.23	0.22	0.28	0.27	0.23	0.28	0.22	0.26
	stepside	0.31	0.32	0.19	0.18	0.18	0.17	0.33	0.32	0.18	0.33	0.17	0.31
	forward	0.38	0.39	0.12	0.11	0.15	0.14	0.36	0.35	0.15	0.36	0.14	0.34
	backward	0.45	0.46	0.05	0.04	0.12	0.11	0.39	0.38	0.12	0.39	0.11	0.38
	stand	0.44	0.45	0.06	0.05	0.10	0.07	0.42	0.40	0.10	0.43	0.08	0.40
		E1	E2	E3	E4	E1	E2	E3	E4	E1	E2	E3	E4

Figure A2: The activation status of each expert.

D.3 Qualitative Results Comparison

We analyze the qualitative results of **CLONE** and **CLONE*** in Fig. A1. Subfigures (a) and (b) show that **CLONE***, trained on OmniH2O [19], fails to track motions like “crouch” or “squat to pick up an object” and falls down. In contrast, subfigures (c) and (d) present the results of **CLONE**, which tracks these motions accurately and robustly. Although **CLONE*** is trained on a larger dataset (more than $8k$ motions, compared to **CLONE**’s 345 motions), it struggles with these tasks. Meanwhile, our model effectively tracks these motions and performs manipulation skills using only about 20% of the data. Since the OmniH2O [19] dataset also includes motions like “squat,” this result suggests that a smaller dataset can still yield excellent tracking performance, as large-scale training data may cause the policy to overly generalize and compromise certain skills.

Expert Activation Analysis To better understand the specialization within our mixture-of-experts architecture, we visualized expert activation weights across nine distinct motion types in Fig. A2. Results reveal clear specialization patterns where motions requiring similar skills activate specific experts. In the first layer, experts 1 and 2 are predominantly activated during standing motions, while experts 3 and 4 show stronger activation during squatting motions. Notably, all four experts in the first layer become activated during dynamic motions such as jumping and punching, suggesting col-

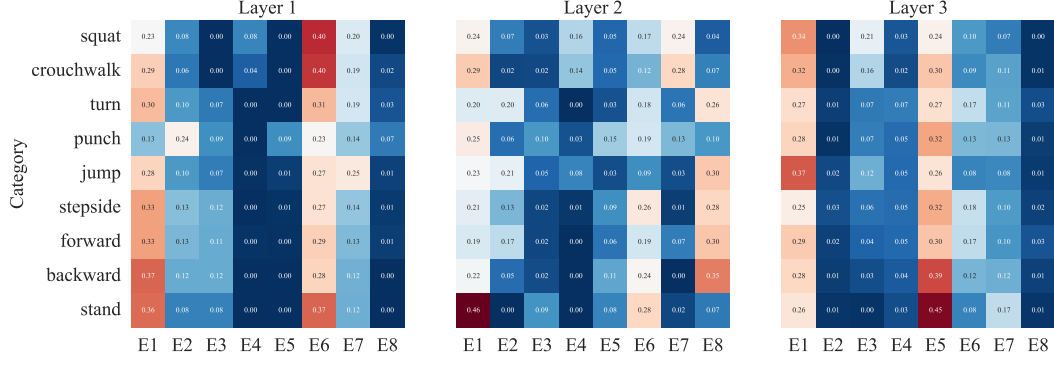


Figure A3: Experts activation when $N = 8$

laborative processing of complex movements. Similar specialization patterns emerged in subsequent layers, albeit with reduced variance across different motion categories.

D.4 The Choice of the Number of MoE Layers and Number of Experts




Figure A4: Experts activation when $L = 4$




Figure A5: Experts activation when $L = 5$

We visualize the activation patterns of experts in Fig. A2 to A6. Fig. A3 shows that MoE layers with $N = 8$ experts activate only half of the experts in each layer, revealing that 8 experts are redundant for the current training data distribution, while 4 experts are sufficient. Fig. A6 demonstrates that **CLONE*** ($L = 1$), which uses only one MoE layer, is still capable of activating different experts. However, as shown in Tab. A3, its tracking performance is inferior to that of **CLONE**. This is primarily attributed to the model's parameters being too limited to effectively learn such diverse motions. Though 4 MoE layers and 5 MoE layers also has same activation patterns, like shown in




Figure A6: Experts activation when $L = 1$

Fig. A4 and A5, we choose 3 MoE layers for a balance of training cost and policy performance. Therefore, we select the MoE policy with three MoE layers and four experts as our final model.