

TeleGate: Whole-Body Humanoid Teleoperation via Gated Expert Selection with Motion Prior

Jie Li[†], Bing Tang[‡], Feng Wu^{†*}

[†]School of Computer Science and Technology, University of Science and Technology of China.

[‡]AnyWit Robotics Co., Ltd., Shushan District, Hefei, Anhui, China.

Abstract—Real-time whole-body teleoperation is a critical method for humanoid robots to perform complex tasks in unstructured environments. However, developing a unified controller that robustly supports diverse human motions remains a significant challenge. Existing methods typically distill multiple expert policies into a single general policy, which often inevitably leads to performance degradation, particularly on highly dynamic motions. This paper presents TeleGate, a unified whole-body teleoperation framework for humanoid robots that achieves high-precision tracking across various motions while avoiding the performance loss inherent in knowledge distillation. Our key idea is to preserve the full capability of domain-specific expert policies by training a lightweight gating network, which dynamically activates experts in real-time based on proprioceptive states and reference trajectories. Furthermore, to compensate for the absence of future reference trajectories in real-time teleoperation, we introduce a VAE-based motion prior module that extracts implicit future motion intent from historical observations, enabling anticipatory control for motions requiring prediction such as jumping and standing up. We conducted empirical evaluations in simulation and also deployed our technique on the Unitree G1 humanoid robot. Using only 2.5 hours of motion capture data for training, our TeleGate achieves high-precision real-time teleoperation across diverse dynamic motions (e.g., running, fall recovery, and jumping), achieving the best overall performance compared to the baseline methods in both tracking accuracy and success rate. Our project website: <https://anywitresearch.github.io/TeleGate/>

I. INTRODUCTION

Whole-body teleoperation is a critical technique for real-time mapping of human motion control to humanoid robots. Compared to pre-programmed control, it not only endows robots with onsite decision-making capabilities in unstructured scenarios (such as disaster rescue and industrial inspection), but also provides high-quality motion data for building embodied intelligence [2, 18, 64]. However, as shown in Fig. 1, building a robust whole-body teleoperation system still faces severe challenges: on one hand, unlike offline motion tracking, real-time teleoperation cannot obtain future reference trajectories and needs to generalize to unseen motion distributions, making it difficult for robots to anticipate action switches such as from jumping to standing up; on the other hand, human motions exhibit high diversity, with significantly different dynamic characteristics across different movements, posing extremely high requirements for a single controller to maintain whole-body balance while accurately tracking diverse motions.




Fig. 1. Whole-body teleoperation of the Unitree G1 humanoid robot using inertial motion capture equipment: (a) Grasping and placing toys into a basket; (b) Standing long jump; (c) Prone position stand-up; (d) Kicking a ball.

Although Reinforcement Learning (RL) based whole-body control has recently shown excellent performance on single tasks [8, 20–22, 40, 44, 45], training directly on multiple tasks unfortunately leads to severe “catastrophic forgetting” due to significant distribution differences between different motion types [40]. To address this, many efforts have been made. For example, PHC [40] validated the effectiveness of increasing network capacity for multi-motion tracking, but its validation was limited to simulation environments. In the field of whole-body teleoperation, TWIST [64] attempted to train a unified policy on large-scale datasets (42h), but in real-world deployment, its lower-body control still cannot support highly dynamic motions. Most recently, SONIC [42] improves performance by scaling data to 700 hours, but still has limited generalizability due to extremely high training costs.

Another class of approaches adopts a two-stage pipeline of first training multiple expert policies and then performing behavioral cloning (distillation) [66]. This approach clusters motion data and trains expert policies independently for different motion types, and then integrates them into a unified controller through distillation. However, the distillation process is essentially viewed as a lossy compression, where the general policy inevitably suffers capability degradation when fitting multiple experts’ decision boundaries within a single network. Overall, the existing methods still face trade-offs

*Corresponding author: wufeng02@ustc.edu.cn.

between data efficiency and tracking accuracy. This motivates us to consider a more efficient framework that can achieve high-precision tracking of diverse motions with limited data.

To address these challenges, we propose *TeleGate* — a novel whole-body teleoperation framework for humanoid robots based on gated expert selection. Specifically, it achieves high-precision real-time tracking of diverse motions while avoiding the performance loss inherent in knowledge distillation. The core idea is to preserve the full capability of domain-specific experts and utilize a gating network for dynamic expert routing at runtime. To this end, we first train multiple domain experts based on dynamic similarity. Then, with expert parameters frozen, we train a lightweight gating network to enable seamless switching of expert capabilities. Furthermore, to compensate for the absence of future information in real-time scenarios, we introduce a VAE-based motion prediction prior that extracts implicit future motion intent, providing the policy with anticipatory control capabilities.

We conducted systematic comparisons with various baseline methods in simulation and complete physical deployment validation on the Unitree G1 humanoid robot. With only 2.5 hours of self-collected motion capture data, *TeleGate* achieves high-precision real-time tracking of diverse human motions including running, fall recovery, and jumping, achieving the best overall performance in success rates and tracking accuracy compared to baseline methods.

Our main contributions are summarized as follows:

- We design a *VAE-based motion prediction prior* module that endows robots with anticipatory control capabilities in real-time teleoperation.
- We propose a unified control framework based on *gated expert selection* for multiple motion categories, avoiding the performance loss caused by knowledge distillation and enabling a single policy to maintain expert-level tracking accuracy across various motion types.
- We implement a *real-time whole-body teleoperation* system supporting multiple complex motions on the Unitree G1 humanoid robot, showing the effectiveness of the proposed method in sim-to-real transfer and its generalization capability to out-of-distribution motions.

To put them together, our *TeleGate* framework provides a system paradigm that combines high data efficiency with high precision for humanoid robot whole-body teleoperation.

II. RELATED WORK

A. Whole-Body Motion Control for Humanoid Robots

Reinforcement learning-based whole-body control for humanoid robots has achieved rapid development in recent years. Compared to traditional model-based control, learning-based methods can handle complex contact dynamics and generate more natural movements. Early works such as DeepMimic [44] and AMP [45] achieved high-quality tracking of reference motions in physics simulation environments. Building upon this, researchers have made significant progress on single tasks including walking [6, 13, 17, 30, 31, 33, 37, 38,

47–49], jumping [61, 65, 69], parkour [69], dancing [8, 24], and fall recovery [21, 22, 68]. These methods have demonstrated successful sim-to-real transfer through techniques such as domain randomization and privileged learning. Adversarial motion priors [41, 45] provide effective motion priors for policy learning, guiding the policy to generate natural and human-like movements through a discriminator that distinguishes between generated and reference motions.

However, when training objectives extend from single skills to multiple motion categories, PHC [40] points out that catastrophic forgetting is the core bottleneck—state-action distributions of different motion types differ significantly, making it difficult for a single policy to master all skills simultaneously. For instance, the contact patterns and balance strategies for walking versus jumping are fundamentally different, and optimizing for one often degrades performance on the other. To address this, SONIC [42] improves performance by scaling training data to 700 hours, but the cost is extremely high; Any2Track [66] and BumbleBee [59] adopt a strategy of first training experts then distilling, but the distillation process inevitably suffers capability degradation; GMT [7] borrows the Mixture-of-Experts (MoE) [53] architecture for joint training, but its superiority has not been fully validated; the MoE paradigm has also been employed for imitation learning of legged locomotion in earlier work such as MPC-Net [5]. VMP [51] encodes trajectory segments via VAE [26] to enhance policy perception of motion patterns, providing inspiration for our motion prediction prior design. More recently, BeyondMimic [32] presents a unified motion tracking formulation that scales to diverse agile human motions with shared hyperparameters and no per-motion tuning, and further composes the learned skills at test time via a guided diffusion model. Overall, existing methods still face trade-offs between data efficiency and tracking accuracy, motivating us to explore a more efficient framework.

B. Humanoid Robot Teleoperation Systems

Teleoperation is an effective way to collect real-world expert demonstrations for imitation learning [3, 9, 16, 34, 63]. Prior work has explored various input modalities, including VR keypoints [9, 18, 23, 39, 63], exoskeletons [2, 14, 25, 52, 60, 62], motion capture suits [4, 10–12, 27, 35, 36, 56, 58, 64], visual pose estimation [15, 19, 29, 46, 54], and robot arm joint matching [14, 16, 60, 67]. Each modality presents different trade-offs between tracking accuracy, deployment cost, and operator comfort.

Based on lower-body control strategies, existing systems can be divided into two categories: base velocity tracking and leg joint tracking. This distinction reflects a fundamental trade-off between ease of control and motion expressiveness. Early work primarily focused on upper-body control [9, 23, 46, 62, 63], with the lower body in fixed states, severely limiting the operational workspace. Mobile-TeleVision [39], AMO [28], and Homie [2] decouple upper and lower body control [8, 15, 18, 19, 39], with the lower-body gait controller receiving velocity commands, but restricting the lower body to preset

TABLE I
COMPARISON OF HUMANOID ROBOT TELEOPERATION SYSTEMS

Method	Leg Joint Tracking	Dynamic Motions	Training Data	MoCap Device
Mobile-TeleVision [39]	✗	✗	–	VR
AMO [28]	✗	✗	–	VR / Vision
ULC [55]	✗	✗	–	VR
Homie [2]	✗	✗	–	Exoskeleton
OmniH2O [18]	✓	✗	>30h	VR / Vision
TWIST [64]	✓	✗	42h	Optical
SONIC [42]	✓	✓	700h	VR / Vision
TeleGate (Ours)	✓	✓	2.5h	Inertial

- *Leg Joint Tracking*: Policy supports leg joint-level control.
- *Dynamic Motions*: Supports highly dynamic whole-body motions such as running, jumping, squatting, and fall recovery.

gait spaces, unable to execute movements requiring fine leg control.

Leg joint tracking methods directly output whole-body joint commands, offering the potential to track arbitrary whole-body motions. However, existing approaches have various limitations: HumanPlus [15] relies on visual pose estimation but has insufficient root position accuracy; OmniH2O [18, 19] can only capture sparse upper-body information; ExBody [8] can only track low-speed smooth motions; TWIST [64] trains a unified policy on 42h of data, but actual deployment cannot support highly dynamic motions such as running, jumping, or fall recovery.

The proposed TeleGate uses inertial motion capture to balance accuracy and portability, achieving whole-body joint-level real-time teleoperation of highly dynamic motions on only 2.5h of data through a gated expert selection mechanism. Unlike distillation-based approaches, our method preserves the full capability of domain experts while enabling seamless switching between motion types (Table I).

III. PROBLEM FORMULATION

We model the whole-body teleoperation of humanoid robots as a *Markov Decision Process* (MDP): $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{P}, r, \gamma)$, where \mathcal{S} is the state space, \mathcal{A} is the action space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \times \mathcal{S} \rightarrow [0, 1]$ is the state transition probability, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, and $\gamma \in (0, 1)$ is the discount factor.

Our goal is to learn a policy $\pi : \mathcal{S} \rightarrow \mathcal{A}$ that outputs joint actions a_t based on current observation s_t at each timestep t , enabling the robot to track the human operator’s reference motion trajectory in real-time. Unlike offline motion tracking tasks, in real-time teleoperation scenarios, the policy cannot access future reference trajectory information M_t^+ and can only rely on *historical reference trajectories* M_t^- and *current proprioceptive state* for motion control.

A. State Space with Reference Trajectory and Motion Prior

The policy observation $s_t = \{o_t, m_t, z_t\} \in \mathcal{S}$ consists of three components: proprioceptive state o_t , reference trajectory information m_t , and motion prediction prior z_t . Specifically, the proprioceptive state o_t is defined as:

$$o_t = \{g_t^{\text{proj}}, v_t, \omega_t, q_t, \dot{q}_t, q_{t-1}^{\text{target}}\} \quad (1)$$

where $g_t^{\text{proj}} \in \mathbb{R}^3$ is the projection of gravity direction in the robot base coordinate frame, $v_t \in \mathbb{R}^3$ and $\omega_t \in \mathbb{R}^3$ are the linear and angular velocities in the base frame, $q_t \in \mathbb{R}^n$ and $\dot{q}_t \in \mathbb{R}^n$ are the joint positions and velocities, $q_{t-1}^{\text{target}} \in \mathbb{R}^n$ is the target motor position at the previous timestep, and n is the number of joint Degrees of Freedom (DoF).

A single-frame reference trajectory m_t is defined as:

$$m_t = \{\hat{g}_t^{\text{proj}}, \hat{v}_t, \hat{\omega}_t, \hat{q}_t, \hat{\dot{q}}_t, \hat{h}_t^{\text{pelvis}}, \hat{h}_t^{\text{feet}}\} \quad (2)$$

where \hat{q}_t and $\hat{\dot{q}}_t$ are the reference joint angles and angular velocities, \hat{g}_t^{proj} is the gravity projection in the pelvis frame, \hat{v}_t and $\hat{\omega}_t$ are the linear and angular velocities in the pelvis frame, and $\hat{h}_t^{\text{pelvis}}$ and \hat{h}_t^{feet} are the pelvis height and feet heights, respectively. The historical reference trajectory window M_t^- considered in real-time teleoperation is defined as:

$$M_t^- = \{m_{t-20}, m_{t-10}, m_{t-5}, m_{t-1}, m_t\} \quad (3)$$

We employ non-uniform sparse sampling rather than consecutive frames: recent frames $(t-1, t)$ capture instantaneous dynamics, while distant frames $(t-20, t-10)$ provide motion trend context, covering a longer time span with the same computational cost.

The future reference trajectory window M_t^+ is defined as:

$$M_t^+ = \{m_{t+5}, m_{t+10}, m_{t+20}\} \quad (4)$$

Note that the motion prediction prior z_t is a low-dimensional latent vector that encodes future motion intent information extracted from historical reference trajectories. The detailed construction method of our motion prediction prior will be introduced later in Section IV-B.

B. Action Space for Residual Correction of Whole-Body Joints

The action of the MDP $a_t \in \mathbb{R}^n$ is defined as the residual correction for all n joints of the whole body. The target joint angle for the PD controller is determined jointly by the reference trajectory \hat{q}_t and the residual action a_t as:

$$q_t^{\text{target}} = \hat{q}_t + a_t \cdot \alpha \quad (5)$$

where α is the action scaling coefficient. Joint torques are computed through the following PD control law as:

$$\tau_t = K_p(q_t^{\text{target}} - q_t) - K_d\dot{q}_t \quad (6)$$

where K_p and K_d are the proportional and derivative gain matrices, respectively. The policy outputs actions at 50Hz.

C. Reward Design and Domain Randomization

In our settings, the reward function r_t consists of: tracking reward r^{track} , survival reward r^{alive} , and regularization reward r^{reg} . The tracking reward encourages the robot to accurately reproduce the reference trajectory; the survival reward provides positive reinforcement when tracking error is below a preset threshold, and terminates the current episode when error exceeds the threshold; the regularization reward penalizes abrupt action changes and foot sliding.

To improve the sim-to-real transfer capability and disturbance robustness of the policy, we randomize external push

disturbances, state observation noise, PD gain coefficients, and contact friction coefficients during training. More details are provided in the appendix.

IV. TELEGATE METHOD

We propose the TeleGate method for real-time whole-body teleoperation of humanoid robots. As shown in Fig. 2, the framework consists of three stages: I) data collection and preprocessing, II) expert policy training with motion prediction prior, and III) gating network training. We first collect whole-body motion data from operators using inertial motion capture equipment and retarget it to the humanoid robot as reference trajectories. Then, we partition the data into several categories based on motion characteristics and train corresponding expert tracking policies. To further enhance the real-time tracking capability of expert policies, we introduce a *Variational Autoencoder* (VAE) to extract motion prediction priors from historical trajectories. Finally, we design a gating network that dynamically selects the optimal expert based on the current state during inference, enabling seamless switching and high-quality tracking across arbitrary motion patterns.

A. Whole-Body Motion Data Collection and Preprocessing

As shown in Stage I of Fig. 2, we use *inertial motion capture* equipment to collect whole-body motion data from human operators, and map human motion to the robot joint space in real-time through online retargeting algorithms [1] to generate reference trajectories. During data collection, the reference trajectories are recorded in the form of the robot’s global base pose (including position and orientation) and joint angles. In the training phase, we compute the position and velocity information of each link from the recorded reference trajectories through forward kinematics and numerical differentiation, which serve as the target tracking signals for the policy network.

Compared to optical motion capture systems, inertial motion capture equipment is convenient to wear and carry, and is not limited by lighting conditions or venue space; compared to VR devices, inertial motion capture can provide more refined whole-body joint information, making it more suitable for diverse whole-body teleoperation scenarios. To eliminate distribution mismatch between training and deployment, the entire data processing pipeline remains completely consistent with the final physical deployment. Furthermore, we partition the collected data into several categories based on the dynamic characteristics of the motions, providing clear task boundaries for subsequent expert policy training.

B. Expert Policy Training with Motion Prediction Prior

As shown in Stage II of Fig. 2, the expert policy training stage involves joint optimization of the policy network and VAE module. In our implementation, we use the Proximal Policy Optimization (PPO) [50] algorithm for RL training in the MuJoCo physics simulator, employing an asymmetric actor-critic architecture: the actor only receives deployable observation information, while the critic can access privileged

information from the simulation environment (such as ground-truth robot states, future reference trajectories, etc.), thereby providing more accurate value estimates without affecting policy deployability.

1) *Motion Prediction Prior*: For offline trajectory tracking tasks (e.g., behavior clone), the policy can utilize future reference trajectory information to prepare in advance for actions requiring anticipation. As mentioned, in real-time teleoperation scenarios, the controller cannot access the operator’s future motion intent. To compensate for this information gap, we propose using a VAE to extract implicit motion prediction priors from historical trajectories.

The choice of VAE as the motion prior extractor is based on the following considerations: from an information bottleneck perspective, the KL divergence regularization forces the latent variable z_t to retain only information most relevant for predicting the future. Compared to directly concatenating historical trajectories, VAE guides the latent space to learn predictive representations through explicit reconstruction supervision. Furthermore, unlike VMP [51] which encodes the current state, our VAE adopts an asymmetric design—the encoder receives historical M_t^- while the decoder reconstructs future M_t^+ —making z_t inherently contain implicit predictions of future motion.

Specifically, we design a VAE module based on a small Transformer [57]. The encoder E_ϕ takes the historical reference trajectory window M_t^- as input and outputs distribution parameters (μ_t, σ_t) of the latent variable z_t as below:

$$\mu_t, \sigma_t = E_\phi(M_t^-) \quad (7)$$

Then, the latent variable z_t is sampled through the reparameterization trick as follow:

$$z_t = \mu_t + \sigma_t \odot \epsilon, \quad \epsilon \sim \mathcal{N}(0, I) \quad (8)$$

The decoder D_ψ takes the latent variable z_t as condition and predicts the future reference trajectory window:

$$\tilde{M}_t^+ = D_\psi(z_t) \quad (9)$$

Given this, the latent variable z_t is then concatenated to the policy network’s observation input, enabling the policy to make anticipatory control decisions even in the absence of explicit future information.

2) *Joint Training Loss*: The VAE module and expert policy are trained jointly, rather than in separate stages. This design is crucial: if the VAE is first trained independently and then frozen while training the policy, the semantics of latent variable z_t may not align with the policy’s decision-making requirements, causing the policy to treat z_t as noise and ignore it. Joint training allows the VAE’s latent space representation to be shaped by policy gradients while learning to predict future trajectories, thereby producing motion priors that are truly useful for control decisions. The overall training loss \mathcal{L} consists of three components:

$$\mathcal{L} = \mathcal{L}_{\text{PPO}} + \lambda_{\text{recon}} \mathcal{L}_{\text{recon}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}} \quad (10)$$

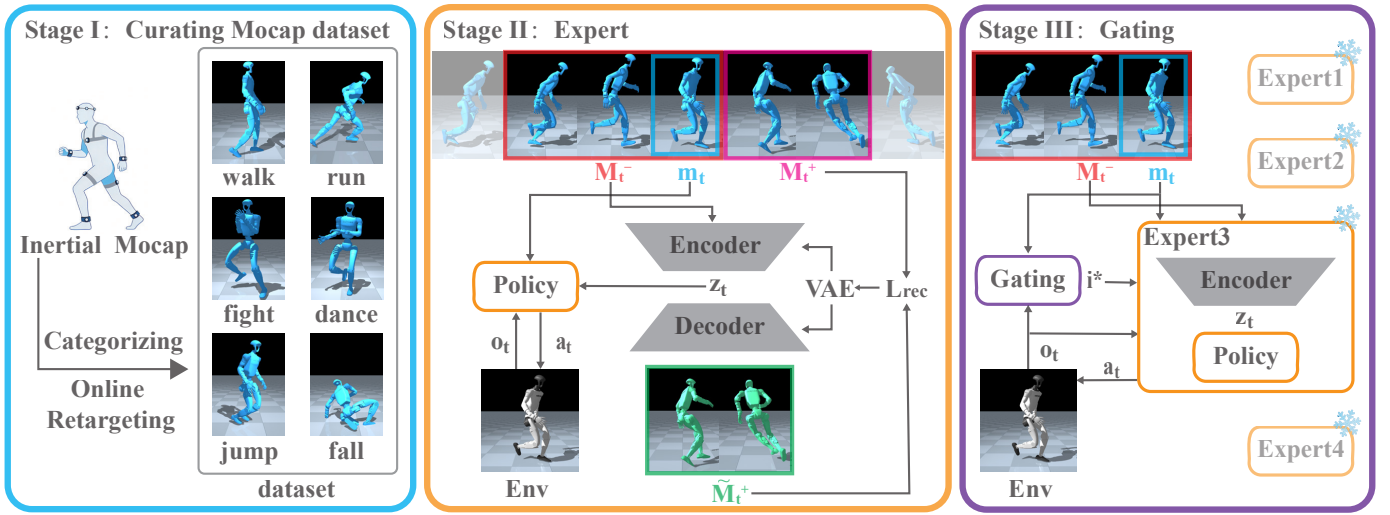


Fig. 2. Framework overview. Our method consists of three stages: (I) Data collection and preprocessing using inertial motion capture; (II) Expert policy training with VAE-based motion prediction prior; (III) Gating network training for dynamic expert selection during inference.

where \mathcal{L}_{PPO} is the policy gradient loss of *Proximal Policy Optimization* (PPO); $\mathcal{L}_{\text{recon}}$ is the mean squared error between predicted and true future trajectories:

$$\mathcal{L}_{\text{recon}} = \|\tilde{M}_t^+ - M_t^+\|_2^2 \quad (11)$$

\mathcal{L}_{KL} is the KL divergence between the latent variable distribution and a standard Gaussian prior, used to constrain the latent space structure:

$$\mathcal{L}_{\text{KL}} = D_{\text{KL}}(\mathcal{N}(\mu_t, \sigma_t^2) \|\mathcal{N}(0, I)) \quad (12)$$

where $\mu_t \in \mathbb{R}^d$ and $\sigma_t \in \mathbb{R}^d$ are the latent variable mean and standard deviation output by the encoder, respectively, and d is the latent variable dimension.

The weight coefficient λ_{recon} controls the strength of trajectory reconstruction loss, while λ_{KL} controls the strength of latent space regularization; a larger λ_{recon} helps improve trajectory prediction accuracy, while an appropriate λ_{KL} ensures the latent space structure is compact and has good generalization.

3) *Failure-Rate Based Curriculum Sampling*: Motion clips in the dataset vary in difficulty, and uniform sampling leads to insufficient training on difficult clips. To address this, we propose an adaptive trajectory sampling strategy: first, motion clips are segmented to within certain time period (e.g., 10 seconds) with some overlap preserved, and the tracking failure rate of each trajectory is dynamically tracked during training. Specifically, the sampling weight w_i for the i -th trajectory is calculated as below:

$$w_i = T_i \cdot (1 + \min(\gamma \cdot f_i, \beta)) \quad (13)$$

where T_i is the trajectory duration, γ is the failure rate amplification coefficient, and β is the weight gain upper limit. The failure rate f_i is updated by exponential moving average.

The sampling probability of each trajectory is positively correlated with its duration and failure rate, automatically directing training resources toward difficult samples, effectively improving the policy’s success rate on high-difficulty motions.

C. Expert Policy Selection via Gating Network

1) *Expert Partitioning Strategy*: Human motions exhibit high diversity and heterogeneity, with significant differences in dynamic characteristics across different motion types. Previous work has shown that using a single policy to learn all motion types leads to poor policy performance. Therefore, we partition the six motion data categories into four groups based on dynamic similarity and train corresponding expert policies: walking and running, dancing and martial arts, fall and recovery, and jumping. Each expert policy π_i ($i \in 1..K$) is trained only on its corresponding motion subset, achieving higher tracking accuracy and robustness.

2) *Gating Network Design*: Traditional methods often use distillation to merge multiple experts into a single policy, but when different experts have decision conflicts for similar states (e.g., in a squatting posture, a jumping expert accumulates power while a recovery expert stands up steadily), a single network struggles to fit all decision boundaries, and supervision signals from different experts may produce gradient conflicts.

To address this, we propose a gating-based expert selection mechanism (as shown in Stage III of Fig. 2): keeping all expert parameters frozen, we train a lightweight gating network to dynamically route to the optimal expert, transforming the problem from “executing all motions” to “selecting which expert”.

The gating network $G_\theta : \mathcal{S}' \rightarrow \mathbb{R}^K$ takes proprioceptive state s_t and reference trajectory information m_t as input, and outputs selection scores for K experts. At each control cycle, the expert i^* with the highest score is selected to execute the action a_t as below:

$$a_t = \pi_{i^*}(o_t, m_t, z_t^{i^*}), \text{ and } i^* = \arg \max_i G_\theta(o_t, m_t)_i \quad (14)$$

3) *Gating Network Training*: During gating network training, all expert policies and their corresponding VAE encoder parameters remain frozen, with only the gating network pa-

parameters optimized. Training uses exactly the same settings as the expert stage, including simulation environment, reward function, domain randomization configuration, and data sampling strategy, with the loss function containing only PPO’s policy gradient loss.

4) *Expert Switching Mechanism*: In actual operation, frequent expert switching may occur, requiring the system to maintain control continuity after switching. Our design ensures this from two aspects: first, each expert policy’s input includes the previous timestep’s target motor position and current proprioceptive state, enabling the newly activated expert to sense the system’s instantaneous dynamics; second, both expert and gating training stages include action regularization rewards to encourage output action smoothness, suppressing action jumps at switching instants.

D. Implementation of Physical Whole-Body Teleoperation

During physical deployment, the operator wears inertial motion capture equipment for movement. The operator’s motion data is transmitted in real-time to the robot via wireless link. The robot first converts the motion capture data through skeleton retargeting to the humanoid robot’s reference trajectory, then performs policy inference and outputs joint actions, with the policy output sent to the low-level PD controller after low-pass filtering.

V. EXPERIMENTS

In this section, we aim to answer the following questions through our simulation and real-world experiments: *E1*: How does TeleGate perform compared to existing whole-body motion tracking methods? *E2*: What advantages does gated expert selection have over other expert integration approaches? *E3*: Can introducing motion prediction prior effectively improve tracking quality? *E4*: Can TeleGate achieve diverse whole-body teleoperation on real robots?

A. Experimental Setup

1) *Simulation Environment*: We conduct training and evaluation in the MuJoCo physics simulator, using the Unitree G1 humanoid robot (29 DoF). The policy outputs actions at 50Hz, and the low-level PD controller runs at 500Hz. The latent variable dimension of the VAE module is set to 32.

2) *Motion Dataset*: Existing whole-body motion control research commonly relies on large-scale public motion capture datasets (such as AMASS[43]). However, these datasets have two limitations: first, the collection equipment and processing pipeline often differ from actual deployment scenarios, leading to distribution shift during transfer—especially for the inertial motion capture equipment used in this work, whose data distribution differs significantly from optical motion capture; second, the datasets contain diverse and ambiguously labeled motion types, making fine-grained categorization difficult. Therefore, we independently collected approximately 2.5 hours of motion data using the same inertial motion capture equipment as deployment, covering six motion categories:

TABLE II
COMPARISON WITH EXISTING METHODS (4-SEED MEAN±STD).

Method	SR↑	E_{plve} ↓	E_{mpjpe} ↓	E_{mpjae} ↓
TWIST [64]	68.9±2.4%	3.88±0.10	53.55±1.45	0.158±0.007
Any2track [66]	91.2±0.1%	8.93±0.04	18.40±0.07	0.090±0.001
GMT [7]	92.0±0.4%	9.65±0.09	29.14±0.12	0.181±0.002
TeleGate (Ours)	97.3±0.1%	8.37±0.04	17.22±0.05	0.085±0.001

walking (40 min), running (24 min), dancing (24 min), martial arts (20 min), fall recovery (26 min), and jumping (16 min).

All sequences are segmented into motion clips less than 10 seconds for precise tracking of failure rates and accuracy. Although the data scale is much smaller than existing work, the consistency between collection and deployment equipment effectively eliminates data domain shift, and the precise categorization of motion types lays the foundation for expert policy training, enabling high-quality sim-to-real transfer with a relatively small amount of data.

3) *Evaluation Metrics*: We adopt the following metrics to evaluate tracking performance: Success Rate (SR): proportion of trajectories where tracking error remains below threshold throughout; pelvis linear velocity error in local coordinate frame (E_{plve} , mm/frame); mean per-joint position error in base coordinate frame (E_{mpjpe} , mm); mean per-joint angle error (E_{mpjae} , rad). Tracking is considered failed when any joint position error in the local coordinate frame exceeds 0.5m or shoulder height error exceeds 0.3m.

4) *Baseline Methods*: We compare with three current state-of-the-art motion tracking methods: TWIST [64] adopts a teacher-student distillation framework to train a unified policy; Any2track [66] first trains multiple experts then distills through DAGger; GMT [7] adopts joint training with Mixture-of-Experts (MoE). Among them, TWIST uses official open-source code; for Any2track, the expert training stage is implemented based on the OpenTrack [66] open-source code, while the distillation stage is reproduced according to the original paper using KL divergence as the distillation loss; GMT is reproduced according to the original paper description. Our expert policy training code is also based on modifications to OpenTrack [66]. For fairness, all methods are trained and tested on the same dataset mentioned above.

B. How does TeleGate perform compared to existing whole-body motion tracking methods? (E1)

As shown in Table II, TeleGate achieves the best overall performance in both success rate (97.3%) and tracking accuracy (E_{mpjpe} =17.22mm, E_{mpjae} =0.085rad). Specifically: TWIST achieves the best pelvis velocity tracking (E_{plve} =3.88) but has only 68.9% success rate. This phenomenon is because TWIST often fails early on difficult samples, while successfully completed simple samples have easier velocity tracking, leading to lower average velocity error. Any2track achieves 91.2% success rate through expert distillation, but information compression during distillation reduces performance. GMT adopts end-to-end MoE training, with E_{mpjpe} (29.14mm) and




Fig. 3. Success rate under different error thresholds. TeleGate is optimal across all thresholds; shading indicates 4-seed standard deviation.

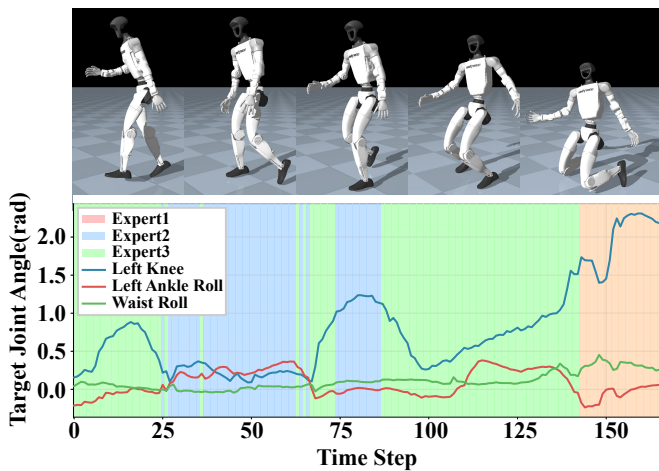


Fig. 4. Expert switching analysis during continuous motion. Top: Key frame sequence; Bottom: Target joint position curves for several joints, with different background colors representing different experts.

E_{mpjae} (0.181rad) both significantly worse than expert distillation methods. In contrast, TeleGate completely preserves each expert’s capability through the gated expert selection mechanism, achieving the best success rate and tracking accuracy.

To further verify the robustness of the success rate improvement to the choice of error threshold, we sweep over a range of thresholds and plot the resulting success rate curves in Fig. 3. TeleGate is optimal across all thresholds, demonstrating that the improvement is not sensitive to the particular threshold used.

1) *Expert Switching Smoothness Analysis*: A potential issue with the gating mechanism is whether expert switching causes control command jumps. Fig. 4 shows a motion sequence containing motion pattern transitions. It can be observed that joint commands remain continuous at switching moments without severe jumps. To quantitatively validate switching smoothness, we compute the average target joint position change (degree/frame) across all timesteps in the dataset: 1.38°

TABLE III
COMPARISON OF DIFFERENT EXPERT INTEGRATION APPROACHES (W/O VAE, 4-SEED MEAN \pm STD).

Method	SR \uparrow	E_{plve} \downarrow	E_{mpjpe} \downarrow	E_{mpjae} \downarrow
Expert (Oracle)	96.3%	8.35	18.03	0.094
Single Policy	94.7 \pm 0.1%	9.61 \pm 0.02	24.25 \pm 0.40	0.121 \pm 0.002
Dagger Distillation	91.2 \pm 0.1%	8.93 \pm 0.04	18.40 \pm 0.07	0.090 \pm 0.001
Top-2 Gating	93.7 \pm 0.1%	8.75 \pm 0.04	17.72\pm0.08	0.086\pm0.001
Random Partition	95.0 \pm 0.5%	8.96 \pm 0.07	19.89 \pm 0.15	0.094 \pm 0.002
Top-1 Gating (Ours)	96.7\pm0.1%	8.44\pm0.03	18.03 \pm 0.06	0.092 \pm 0.001

without expert switching versus 3.49° with expert switching. Although the angle change increases at switching moments, the magnitude remains within an acceptable range and does not significantly impact control stability. This smoothness stems from our state design—each expert’s input includes the previous timestep’s target position q_{t-1}^{target} , enabling the newly activated expert to sense the system’s instantaneous dynamics and output continuous commands.

C. What advantages does gated expert selection have over other expert integration approaches? (E2)

To validate the effectiveness of the gated expert selection mechanism, we compare with multiple expert integration approaches: Single Policy directly trains a single generalist policy; Dagger Distillation distills pre-trained experts into a unified policy; Top-2 Gating performs weighted fusion of outputs from the two highest-scoring experts; Random Partition randomly divides data into 4 groups to train experts. Additionally, we introduce Expert (Oracle) as a performance upper bound, which directly selects the optimal expert using ground-truth motion category labels. To exclude VAE influence, all methods here do not use motion prediction prior.

The results in Table III reveal the following key findings:

- **Single policy struggles to master all motions.** Single Policy’s success rate (94.7%) is not the lowest, but its tracking accuracy is worst (E_{mpjpe} =24.25mm), indicating that a single network has difficulty mastering multiple motion types with significantly different dynamic characteristics simultaneously.
- **Distillation causes performance loss.** Dagger Distillation’s success rate (91.2%) is significantly lower than the Oracle upper bound, validating our core hypothesis—distillation is essentially lossy compression, where the generalist policy inevitably suffers capability degradation when fitting multiple experts’ decision boundaries.
- **Multi-expert fusion introduces interference.** Top-2 Gating is slightly better in joint accuracy but has lower success rate (93.7%). This shows that suboptimal experts’ outputs interfere with the optimal expert’s decisions; weighted fusion is not an effective integration strategy.
- **Expert partitioning depends on dynamic similarity.** Random Partition’s E_{mpjpe} (19.89mm) is significantly worse than methods based on dynamic similarity partitioning, indicating that reasonable expert partitioning is a prerequisite for gating mechanism effectiveness.

TABLE IV
 ABLATION STUDY ON MOTION PREDICTION PRIOR (4-SEED MEAN \pm STD).
 JUMP E_{mpjpe} REDUCED BY 13%, FALL RECOVERY BY 8%.

Category	SR \uparrow		E_{mpjpe} \downarrow	
	w/o VAE	w/ VAE	w/o VAE	w/ VAE
Walk & Run	99.5 \pm 0.2%	99.6 \pm 0.3%	16.46 \pm 0.05	16.48 \pm 0.03
Dance & Martial	100.0 \pm 0.0%	100.0 \pm 0.0%	12.66 \pm 0.12	12.31\pm0.08
Fall Recovery	89.3 \pm 0.4%	91.8\pm0.3%	28.23 \pm 0.24	25.94\pm0.20
Jump	91.7 \pm 0.4%	93.2\pm0.2%	22.27 \pm 0.15	19.40\pm0.12
Overall	96.7 \pm 0.1%	97.3\pm0.1%	18.03 \pm 0.06	17.22\pm0.05

- **Gating selection approaches Oracle upper bound.**

Our Top-1 Gating is comparable or even slightly better than Oracle in both success rate (96.7% vs 96.3%) and tracking error. This phenomenon stems from motion clip labels not being frame-level accurate—for example, a trajectory labeled “fall recovery” may contain walking segments, while the gating network can select the optimal expert frame-by-frame, making it reasonable to exceed the segment-label-based Oracle in certain metrics.

D. Can introducing motion prediction prior effectively improve tracking quality? (E3)

To address the inability to obtain future reference trajectories in real-time teleoperation, we introduce a VAE-based motion prediction prior. This section analyzes the module’s impact on overall performance through ablation experiments.

As shown in Table IV, the VAE prediction yields the most significant gains on motions requiring anticipation: Jump E_{mpjpe} is reduced by 13% (from 22.27mm to 19.40mm) and Fall Recovery E_{mpjpe} by 8% (from 28.23mm to 25.94mm), with success rates of these two categories also improved by 1.5–2.5 percentage points. For periodic and quasi-static motions (Walk & Run, Dance & Martial), where future intent is largely deducible from the current state, the VAE brings only marginal changes. Overall, the success rate increases from 96.7% to 97.3% and E_{mpjpe} decreases from 18.03mm to 17.22mm, confirming that the motion prediction prior is most beneficial precisely for the highly dynamic motions it is designed to assist.

Fig. 5 illustrates the trajectory prediction performance of the VAE module across different motion types. To validate that the VAE can effectively predict future motion intent from historical reference trajectories, we visualized its reconstruction results. In the visualization, the base height and all joint angles of the predicted trajectory are obtained from VAE reconstruction, while the horizontal position and orientation of the base are set to ground truth values for comparison since they are not reconstruction targets. It can be observed that for motions requiring anticipation such as standing up and lying down, the VAE accurately captures motion trends from historical trajectories and predicts reasonable future poses; for periodic motions like running and walking, the predicted trajectories also closely match the ground truth. This demonstrates that the VAE successfully learns the temporal dynamics of different motion patterns, providing the policy network with effective




Fig. 5. Trajectory reconstruction performance of VAE motion prediction prior. Left: Historical reference trajectory sequence; Right: Ground truth future trajectory (blue) and VAE-predicted future trajectory (green). From top to bottom: stand-up, lie-down, running, and walking.




Fig. 6. More real-world teleoperation skills: (a) sitting; (b) walking; (c) running; (d) lying prone; (e) lateral slide; (f) single-leg hop.

future motion intent information.

E. Can TeleGate achieve diverse whole-body teleoperation on real robots? (E4)

We deploy TeleGate on the Unitree G1 humanoid robot for real-world teleoperation validation. During physical deployment, policy inference runs at 50Hz, while the PD controller and online retargeting module run at 500Hz, with a low-pass filter with cutoff frequency of 37.5Hz connecting the policy output and PD controller.

As shown in Fig. 1, TeleGate successfully achieves real-time teleoperation of multiple highly dynamic motions: (a) grasping toys and placing them into a storage basket; (b) standing long jump; (c) prone position stand-up; (d) ball kicking. Fig. 6 demonstrates more whole-body teleoperation skills: (a) sitting; (b) walking; (c) running; (d) lying prone; (e) lateral slide; (f) single-leg hop.

To quantify sim-to-real transfer, we randomly selected 36 motion clips (6 per category, 20 s each) for real-robot teleoperation evaluation. TeleGate achieves an overall success rate of 97.2% with E_{mpjpe} =18.57 mm on the real robot, exhibiting only a small sim-to-real gap from the simulation result of 100% / 17.14 mm. All categories except Fall Recovery (83.3%) reach a 100% success rate; the slight drop on Fall Recovery is mainly attributable to sim-to-real gaps in contact dynamics during whole-body landing and ground interaction.

Together with the qualitative demonstrations above, these results validate TeleGate’s effectiveness in sim-to-real transfer and its generalization capability to out-of-distribution motions, spanning diverse scenarios from fine manipulation to highly dynamic motions. More demonstrations can be found in the supplementary video.

VI. CONCLUSION

This paper presents TeleGate, a unified real-time teleoperation framework for diverse whole-body motions of humanoid robots. By dynamically selecting domain expert policies through a gating network, TeleGate achieves seamless switching and high-precision tracking across multiple motion patterns while avoiding the performance loss inherent in knowledge distillation. Through VAE-based motion prediction prior, TeleGate compensates for the information bottleneck of missing future reference trajectories in real-time teleoperation, endowing the policy with anticipatory control capabilities for dynamic motions such as jumping and standing up. Simulation and real-world experiments demonstrate that TeleGate achieves the best overall performance in tracking accuracy and success rate compared to existing methods, and successfully achieves real-time teleoperation of highly dynamic motions including running, jumping, and fall recovery on the Unitree G1 humanoid robot. This work provides an efficient and scalable solution for agile whole-body teleoperation of humanoid robots, laying the foundation for their practical applications in unstructured environments.

Future work will explore the following directions: investigating how to leverage high-quality motion data collected through teleoperation to train autonomous policies, achieving transfer from teleoperation to autonomous control.

ACKNOWLEDGMENTS

We thank all reviewers and ACs for their insightful comments and valuable suggestions. We thank Cong Liang and Rongyun Cao for the fruitful discussions, Lifan Li and Wenhao Tang from AnyWit Robotics for their support in setting up the real-world system and video recording, and Wanying Chen for the help with motion capture data collection. This work was supported in part by Major Research Plan of the National Natural Science Foundation of China (92048301), Anhui Provincial Major Research and Development Plan (202004H07020008), and Anhui Provincial Natural Science Foundation (2208085MF172).

REFERENCES

- [1] João Pedro Araújo, Yanjie Ze, Pei Xu, Jiajun Wu, and C. Karen Liu. Retargeting matters: General motion retargeting for humanoid motion tracking. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2026.
- [2] Qingwei Ben, Feiyu Jia, Jia Zeng, Junting Dong, Dahua Lin, and Jiangmiao Pang. Homie: Humanoid locomanipulation with isomorphic exoskeleton cockpit. In *Robotics: Science and Systems (RSS)*, 2025.
- [3] Kevin Black, Noah Brown, Danny Driess, Adnan Esmail, Michael Equi, Chelsea Finn, Niccolo Fusai, Lachy Groom, Karol Hausman, Brian Ichter, et al. π 0: A vision-language-action flow model for general robot control. *arXiv preprint arXiv:2410.24164*, 2024.
- [4] Manuel Caeiro-Rodríguez, Iván Otero-González, Fernando A Mikic-Fonte, and Martín Llamas-Nistal. A systematic review of commercial smart gloves: Current status and applications. *Sensors*, 21(8):2667, 2021.
- [5] Jan Carius, Farbod Farshidian, and Marco Hutter. MPC-Net: A first principles guided policy search. *IEEE Robotics and Automation Letters*, 5(2):2897–2904, 2020.
- [6] Zixuan Chen, Xialin He, Yen-Jen Wang, Qiayuan Liao, Yanjie Ze, Zhongyu Li, S Shankar Sastry, Jiajun Wu, Koushil Sreenath, Saurabh Gupta, et al. Learning smooth humanoid locomotion through lipschitz-constrained policies. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.
- [7] Zixuan Chen, Mazeyu Ji, Xuxin Cheng, Xuanbin Peng, Xue Bin Peng, and Xiaolong Wang. Gmt: General motion tracking for humanoid whole-body control. *arXiv preprint arXiv:2506.14770*, 2025.
- [8] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. In *Robotics: Science and Systems (RSS)*, 2024.
- [9] Xuxin Cheng, Jialong Li, Shiqi Yang, Ge Yang, and Xiaolong Wang. Open-television: Teleoperation with immersive active visual feedback. In *Conference on Robot Learning (CoRL)*, 2024.
- [10] Stefano Dafarra, Ugo Pattacini, Giulio Romualdi, Lorenzo Rapetti, Riccardo Grieco, Kourosh Darvish, Gianluca Milani, Enrico Valli, Ines Sorrentino, Paolo Maria Viceconte, et al. icub3 avatar system: Enabling remote fully immersive embodiment of humanoid robots. *Science Robotics*, 9(86):eadh3834, 2024.
- [11] Kourosh Darvish, Yeshasvi Tirupachuri, Giulio Romualdi, Lorenzo Rapetti, Diego Ferigo, Francisco Javier Andrade Chavez, and Daniele Pucci. Whole-body geometric retargeting for humanoid robots. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, pages 679–686, 2019.
- [12] Anca D Dragan, Kenton CT Lee, and Siddhartha S Srinivasa. Legibility and predictability of robot motion. In *ACM/IEEE International Conference on Human-Robot*

- Interaction (HRI)*, pages 301–308, 2013.
- [13] Pranay Dugar, Aayam Shrestha, Fangzhou Yu, Bart van Marum, and Alan Fern. Learning multi-modal whole-body control for real-world humanoid robots. *arXiv preprint arXiv:2408.07295*, 2024.
- [14] Hongjie Fang, Hao-Shu Fang, Yiming Wang, Jieji Ren, Jingjing Chen, Ruo Zhang, Weiming Wang, and Cewu Lu. Airexo: Low-cost exoskeletons for learning whole-arm manipulation in the wild. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 15031–15038, 2024.
- [15] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. In *Conference on Robot Learning (CoRL)*, 2024.
- [16] Zipeng Fu, Tony Z Zhao, and Chelsea Finn. Mobile aloha: Learning bimanual mobile manipulation with low-cost whole-body teleoperation. In *Conference on Robot Learning (CoRL)*, 2024.
- [17] Xinyang Gu, Yen-Jen Wang, Xiang Zhu, Chengming Shi, Yanjiang Guo, Yichen Liu, and Jianyu Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. *arXiv preprint arXiv:2408.14472*, 2024.
- [18] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. In *Conference on Robot Learning (CoRL)*, 2024.
- [19] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024. Oral Presentation.
- [20] Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbabu, Chaoyi Pan, Zeji Yi, Guannan Qu, Kris Kitani, Jessica K. Hodgins, Linxi Fan, Yuke Zhu, Changliu Liu, and Guanya Shi. Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. In *Robotics: Science and Systems (RSS)*, 2025.
- [21] Xialin He, Runpei Dong, Zixuan Chen, and Saurabh Gupta. Humanup: Learning getting-up policies for real-world humanoid robots. In *Robotics: Science and Systems (RSS)*, 2025.
- [22] Tao Huang, Junli Ren, Huayi Wang, Zirui Wang, Qingwei Ben, Muning Wen, Xiao Chen, Jianan Li, and Jiangmiao Pang. Host: Learning humanoid standing-up control across diverse postures. In *Robotics: Science and Systems (RSS)*, 2025. Best Systems Paper Finalist.
- [23] Aadithya Iyer, Zhuoran Peng, Yinlong Dai, Irmak Guzey, Siddhant Haldar, Soumith Chintala, and Lerrel Pinto. Open teach: A versatile teleoperation system for robotic manipulation. *arXiv preprint arXiv:2403.07870*, 2024.
- [24] Mazeyu Ji, Xuanbin Peng, Fangchen Liu, Jialong Li, Ge Yang, Xuxin Cheng, and Xiaolong Wang. Exbody2: Advanced expressive humanoid whole-body control. *arXiv preprint arXiv:2412.13196*, 2024.
- [25] Yizhou Jiang, Ruihai Zhang, Josiah Wong, Chris Wang, Yanjie Ze, Hang Yin, Celso Gokmen, Shuran Song, Jiajun Wu, and Li Fei-Fei. Behavior robot suite: Streamlining real-world whole-body manipulation for everyday household activities. *arXiv preprint arXiv:2503.05652*, 2025.
- [26] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2014. ICLR 2014.
- [27] Johannes Koenemann, Felix Burget, and Maren Bennewitz. Real-time imitation of human whole-body motions by humanoids. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2806–2812, 2014.
- [28] Jialong Li, Xuxin Cheng, Tianshu Huang, Shiqi Yang, Ri-Zhao Qiu, and Xiaolong Wang. Amo: Adaptive motion optimization for hyper-dexterous humanoid whole-body control. In *Robotics: Science and Systems (RSS)*, 2025.
- [29] Jinhan Li, Yifeng Zhu, Yuqi Xie, Zhenyu Jiang, Mingyo Seo, Georgios Pavlakos, and Yuke Zhu. Okami: Teaching humanoid robots manipulation skills through single video imitation. In *Conference on Robot Learning (CoRL)*, 2024. Oral Presentation.
- [30] Zhongyu Li, Xuxin Cheng, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for robust parameterized locomotion control of bipedal robots. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 2811–2817, 2021.
- [31] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control. *The International Journal of Robotics Research*, page 02783649241285161, 2024.
- [32] Qiayuan Liao, Takara E. Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C. Karen Liu. Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion. *arXiv preprint arXiv:2508.08241*, 2025.
- [33] Qiayuan Liao, Bike Zhang, Xuanyu Huang, Xiaoyu Huang, Zhongyu Li, and Koushil Sreenath. Berkeley humanoid: A research platform for learning-based control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [34] Toru Lin, Yu Zhang, Qiyang Li, Haozhi Qi, Brent Yi, Sergey Levine, and Jitendra Malik. Learning visuotactile skills with two multifingered hands. *arXiv preprint arXiv:2404.16823*, 2024.
- [35] Hangxin Liu, Xu Xie, Matt Millar, Mark Edmonds, Feng Gao, Yixin Zhu, Veronica J Santos, Brandon Rothrock, and Song-Chun Zhu. A glove-based system for studying

- hand-object manipulation via joint pose and force sensing. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6617–6624, 2017.
- [36] Hangxin Liu, Zhenliang Zhang, Xu Xie, Yixin Zhu, Yue Liu, Yongtian Wang, and Song-Chun Zhu. High-fidelity grasping in virtual reality using a glove-based system. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 5180–5186, 2019.
- [37] Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. *arXiv preprint arXiv:2411.14386*, 2024.
- [38] Junfeng Long, Wenhan Yu, Quanyi Li, Zirui Wang, Dahua Lin, and Jiangmiao Pang. Learning h-infinity locomotion control. *arXiv preprint*, 2024.
- [39] Chenhao Lu, Xuxin Cheng, Jialong Li, Shiqi Yang, Mazeyu Ji, Chengjing Yuan, Ge Yang, Sha Yi, and Xiaolong Wang. Mobile-television: Predictive motion priors for humanoid whole-body control. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2025.
- [40] Zhengyi Luo, Jinkun Cao, Alexander W Winkler, Kris Kitani, and Weipeng Xu. Perpetual humanoid control for real-time simulated avatars. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10895–10904, 2023.
- [41] Zhengyi Luo, Jinkun Cao, Josh Merel, Alexander Winkler, Jing Huang, Kris Kitani, and Weipeng Xu. Universal humanoid motion representations for physics-based control. In *International Conference on Learning Representations (ICLR)*, 2024. Spotlight.
- [42] Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Sirui Chen, Fernando Castañeda, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, Xingye Da, Runyu Ding, Cyrus Hogg, Lina Song, Edy Lim, Eugene Jeong, Tairan He, Haoru Xue, Wenli Xiao, Zi Wang, Simon Yuen, Jan Kautz, Yan Chang, Umar Iqbal, Linxi Fan, and Yuke Zhu. Sonic: Supersizing motion tracking for natural humanoid whole-body control. *arXiv preprint arXiv:2511.07820*, 2025.
- [43] Naureen Mahmood, Nima Ghorbani, Nikolaus F Troje, Gerard Pons-Moll, and Michael J Black. Amass: Archive of motion capture as surface shapes. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5442–5451, 2019.
- [44] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. In *ACM Transactions on Graphics (TOG)*, volume 37, pages 1–14, 2018.
- [45] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (TOG)*, 40(4):1–20, 2021. SIGGRAPH 2021.
- [46] Yuzhe Qin, Wei Yang, Binghao Huang, Karl Van Wyk, Hao Su, Xiaolong Wang, Yu-Wei Chao, and Dieter Fox. Anyteleop: A general vision-based dexterous robot arm-hand teleoperation system. *arXiv preprint arXiv:2307.04577*, 2023.
- [47] Ilija Radosavovic, Sarthak Kamat, Trevor Darrell, and Jitendra Malik. Learning humanoid locomotion over challenging terrain. *arXiv preprint arXiv:2410.03654*, 2024.
- [48] Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath. Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 9(89):eadi9579, 2024.
- [49] Ilija Radosavovic, Bike Zhang, Baifeng Shi, Jathushan Rajasegaran, Sarthak Kamat, Trevor Darrell, Koushil Sreenath, and Jitendra Malik. Humanoid locomotion as next token prediction. 2024.
- [50] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [51] Agon Serifi, Ruben Grandia, Espen Knoop, Markus Gross, and Moritz Bächer. Vmp: Versatile motion priors for robustly tracking motion on physical characters. *Computer Graphics Forum (ACM SIGGRAPH / Eurographics Symposium on Computer Animation)*, 43(8), 2024.
- [52] Kenneth Shaw, Yulong Li, Jiahui Yang, Mohan Kumar Srirama, Ray Liu, Haoyu Xiong, Russell Mendonca, and Deepak Pathak. Bimanual dexterity for complex tasks. In *8th Annual Conference on Robot Learning*, 2024.
- [53] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations (ICLR)*, 2017.
- [54] Aravind Sivakumar, Kenneth Shaw, and Deepak Pathak. Robotic telekinesis: Learning a robotic hand imitator by watching humans on youtube. In *Proceedings of Robotics: Science and Systems*, New York City, NY, USA, 2022.
- [55] Wandong Sun, Luying Feng, Baoshi Cao, Yang Liu, Yaochu Jin, and Zongwu Xie. Unified loco-manipulation controller for humanoid robots. *arXiv preprint arXiv:2507.06905*, 2025.
- [56] Bing Tang, Rongyun Cao, Rongya Chen, Xiaoping Chen, Bei Hua, and Feng Wu. Automatic generation of robot facial expressions with preferences. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7606–7613, London, UK, May 2023. doi: 10.1109/ICRA48891.2023.10160409.
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 30, 2017.
- [58] Chen Wang, Haochen Shi, Weizhuo Wang, Ruohan Zhang, Li Fei-Fei, and C Karen Liu. Dexcap: Scalable and portable mocap data collection system for dexterous

- manipulation. *arXiv preprint arXiv:2403.07788*, 2024.
- [59] Yuxuan Wang, Ming Yang, Weishuai Zeng, Yu Zhang, Xinrun Xu, Haobin Jiang, Ziluo Ding, and Zongqing Lu. From experts to a generalist: Toward general whole-body control for humanoid robots. *arXiv preprint arXiv:2506.12779*, 2025.
- [60] Philipp Wu, Yide Shentu, Zhongke Yi, Xingyu Lin, and Pieter Abbeel. Gello: A general, low-cost, and intuitive teleoperation framework for robot manipulators. 2023.
- [61] Yufei Xue, Wentao Dong, Minghuan Liu, Weinan Zhang, and Jiangmiao Pang. Hugwbc: A unified and general humanoid whole-body controller for versatile locomotion. In *Robotics: Science and Systems (RSS)*, 2025.
- [62] Shiqi Yang, Minghuan Liu, Yuzhe Qin, Runyu Ding, Jialong Li, Xuxin Cheng, Ruihan Yang, Sha Yi, and Xiaolong Wang. Ace: A cross-platform visual-exoskeletons system for low-cost dexterous teleoperation. In *Conference on Robot Learning (CoRL)*, 2024.
- [63] Yanjie Ze, Zixuan Chen, João Pedro Araújo, Zi-ang Cao, Xue Bin Peng, Jiajun Wu, and C Karen Liu. Generalizable humanoid manipulation with improved 3d diffusion policies. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.
- [64] Yanjie Ze, Zixuan Chen, João Pedro Araújo, Zi-ang Cao, Xue Bin Peng, Jiajun Wu, and C. Karen Liu. Twist: Teleoperated whole-body imitation system. In *Conference on Robot Learning (CoRL)*, 2025.
- [65] Chong Zhang, Wenli Xiao, Tairan He, and Guanya Shi. Wococo: Learning whole-body humanoid control with sequential contacts. In *Conference on Robot Learning (CoRL)*, 2024. Oral Presentation.
- [66] Zhikai Zhang, Jun Guo, Chao Chen, Jilong Wang, Chenghuai Lin, Yunrui Lian, Han Xue, Zhenrong Wang, Maoqi Liu, Jiangran Lyu, Huaping Liu, He Wang, and Li Yi. Track any motions under any disturbances, 2025. URL <https://arxiv.org/abs/2509.13833>.
- [67] Tony Z Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. *arXiv preprint arXiv:2304.13705*, 2023.
- [68] Ziwen Zhuang and Hang Zhao. Embrace collisions: Humanoid shadowing for deployable contact-agnostics motions. *arXiv preprint arXiv:2502.01465*, 2025.
- [69] Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. In *Conference on Robot Learning (CoRL)*, pages 1975–1991. PMLR, 2024.

APPENDIX

A. Hyperparameters and Training Settings

1) *PPO Hyperparameters*: Proximal Policy Optimization (PPO) is adopted for policy gradient training of both expert policies and the gating network. The learning rate is set to 3×10^{-4} , with a clip range of $\epsilon_{\text{clip}} = 0.2$. The Generalized Advantage Estimation (GAE) parameter is $\lambda = 0.95$, and the discount factor is $\gamma = 0.97$. Each batch of data is updated 4 times, with 32 mini-batches, a batch size of 1024, and an unroll length of 20 steps. The maximum gradient norm is set to 1.0, and the entropy loss coefficient is 0.01.

2) *VAE, Curriculum Sampling, and Action Scaling*: The motion prediction prior based on Variational Autoencoder (VAE) is jointly trained with expert policies, with future trajectory reconstruction loss weight $\lambda_{\text{recon}} = 0.5$, KL divergence weight $\lambda_{\text{KL}} = 0.0005$, and latent dimension $d = 32$. The trajectory sampling weight is computed as $w_i = T_i \cdot (1 + \min(\gamma \cdot f_i, \beta))$, where T_i is the length of trajectory i , f_i is the failure rate of trajectory i , $\gamma = 0.6$ is the failure rate amplification coefficient, and $\beta = 19.0$ is the weight gain upper bound (i.e., the maximum weight multiplier is $1 + 19 = 20$ times). The action scaling formula is $q_t^{\text{target}} = \hat{q}_t + a_t \cdot \alpha$, where \hat{q}_t is the reference joint position and $\alpha = 1.0$ is the action scaling coefficient.

3) *Training Environment and Scale*: All policies are trained in the MuJoCo physics simulator with NVIDIA RTX A6000 PRO GPUs, and implemented based on the mjlab framework. The number of parallel environments is 32768, with a maximum episode length of 500 steps. During the expert policy phase, four expert groups (Walk/Run, Dance/Fight, Fall/Getup, Jump) are trained. Each expert is trained for approximately 1.5B environment steps over 30 hours. During the gating network phase, the expert parameters are frozen, and the gating network is trained alone for approximately 150M environment steps over 1.5 hours.

B. Network Architecture

1) *VAE with Transformer-based Encoder/Decoder*: The motion prediction prior adopts a Transformer-based VAE architecture: the encoder E_ϕ takes as input the historical reference trajectory M_t^- (5 frames) and outputs latent distribution parameters (μ_t, σ_t) ; the decoder D_ψ predicts the future window \tilde{M}_t^+ (3 frames) conditioned on z_t . The architecture is shown in Table V.

The encoder concatenates or projects frame features from M_t^- and feeds them into the Transformer, takes the last layer’s [CLS] or pooled output, and passes it through an MLP to obtain $\mu_t \in \mathbb{R}^d$ and $\sigma_t \in \mathbb{R}^d$. The decoder, conditioned on z_t , expands it through an MLP and feeds it into Transformer decoder layers to output trajectory features for the next three frames.

2) *Expert Policy Network (Actor)*: The Actor takes as input $s_t = (o_t, m_t, z_t)$ and outputs action $a_t \in \mathbb{R}^{29}$. It adopts a 5-layer MLP with hidden layer dimensions of (512, 512, 256, 256, 128) and ReLU activation. The output

TABLE V
VAE / TRANSFORMER ARCHITECTURE

Component/Hyperparameter	Value
<i>Encoder</i>	
Transformer Layers	3
Attention Heads	8
Hidden Dimension (d_model)	256
Feedforward Dimension	512
<i>Decoder</i>	
Transformer Layers	3
Attention Heads	8
Hidden Dimension (d_model)	256
Feedforward Dimension	512
Latent Dimension d	32

TABLE VI
DOMAIN RANDOMIZATION CONFIGURATION

Item	Range/Scale
External Push	$\mathcal{U}(0.1, 1.0)$ m/s
Push Interval	$\mathcal{U}(8.0, 15.0)$ s
Noise—Joint Position	0.02 rad
Noise—Joint Velocity	1.5 rad/s
Noise—Gravity Vector	0.03
Noise—Gyroscope	0.1 rad/s
Noise—Ref. Root Lin. Vel.	0.5 m/s
Noise—Ref. Root Ang. Vel.	0.5 rad/s
Contact Friction	$\times \mathcal{U}(0.5, 1.5)$
Joint Friction Loss	$\times \mathcal{U}(0.7, 1.5)$
Joint Armature	$\times \mathcal{U}(1.0, 1.05)$
Torso COM Position	$+\mathcal{U}(-0.10, 0.10)$ m
Torso Mass	$+\mathcal{U}(-1.5, 2.0)$ kg
Initial Joint Position	$+\mathcal{U}(-0.02, 0.02)$ rad

layer uses Tanh activation followed by multiplication with action scaling coefficient $\alpha = 1.0$, producing a 29-dimensional action (corresponding to the number of controllable joints).

3) *Critic Network*: The Critic takes as input privileged observations (e.g., true state and future reference trajectories) and outputs a scalar state value $V(s_t) \in \mathbb{R}$. The network architecture is the same as the Actor, adopting a 5-layer MLP with hidden layer dimensions of (512, 512, 256, 256, 128), ReLU activation, and a 1-dimensional output layer.

4) *Gating Network*: The gating network $G_\theta : (o_t, m_t) \mapsto \mathbb{R}^K$ outputs scores for $K = 4$ experts, and takes $\arg \max$ to obtain the current expert index. The network adopts a 5-layer MLP with hidden layer dimensions of (512, 512, 256, 256, 128), ReLU activation, and outputs a 4-dimensional vector corresponding to four expert groups (Walk/Run, Dance/Fight, Fall/Getup, Jump).

C. Domain Randomization

To improve sim-to-real transfer and robustness, the following quantities are randomized during training, where $\mathcal{U}(a, b)$ denotes a uniform distribution. For physical properties with nominal values (contact friction, joint friction loss, and joint armature), the randomization scales are multiplied by their respective nominal values. Noise on the gravity vector is applied element-wise as $\mathcal{U}(1 - 0.03, 1 + 0.03)$ per component.

TABLE VII
REWARD FUNCTION FORMULATION AND WEIGHTS

Reward Term	Formula	Weight
<i>Tracking Rewards</i>		
Upper Body Position	$\exp(-10.0 \cdot \sum x_t^{\text{upper}} - \hat{x}_t^{\text{upper}})$	1.0
Lower Body Position	$\exp(-10.0 \cdot \sum x_t^{\text{lower}} - \hat{x}_t^{\text{lower}})$	0.5
Foot Position	$\exp(-10.0 \cdot \sum x_t^{\text{feet}} - \hat{x}_t^{\text{feet}})$	2.1
Rigid Body Rotation	$\exp(-10.0 \cdot \sum 2 \arccos(q^w))$	0.5
Rigid Body Linear Velocity	$\exp(-0.2 \cdot \sum \ v_t - \hat{v}_t\ ^2)$	0.5
Rigid Body Angular Velocity	$\exp(-0.02 \cdot \sum \ \omega_t - \hat{\omega}_t\ ^2)$	0.5
Joint Position	$\exp(-0.1 \cdot \sum q_t - \hat{q}_t)$	0.75
Joint Velocity	$\exp(-1.0 \cdot \sum \dot{q}_t - \hat{\dot{q}}_t)$	0.5
Root Linear Velocity	$\exp(-1.0 \cdot \sum v_t^{\text{root}} - \hat{v}_t^{\text{root}})$	1.0
Root Angular Velocity	$\exp(-0.1 \cdot \sum \omega_t^{\text{root}} - \hat{\omega}_t^{\text{root}})$	1.0
Roll-Pitch Angle	$\exp(-5.0 \cdot \sum \text{rp}_t - \hat{\text{rp}}_t)$	1.0
Root Height	$\exp(-10.0 \cdot h_t^{\text{root}} - \hat{h}_t^{\text{root}})$	1.0
Foot Height	$\exp(-10.0 \cdot \sum h_t^{\text{feet}} - \hat{h}_t^{\text{feet}})$	1.0
<i>Regularization Penalties</i>		
Action Rate	$-\ a_t - a_{t-1}\ ^2$	-0.5
Torque	$-\sum \tau_t^2$	-2×10^{-5}
Joint Position Limit	$-\sum \max(0, q_t - q_{\text{soft}})$	-10
Joint Velocity Limit	$-\sum \max(0, \dot{q}_t - \dot{q}_{\text{limit}})$	-5
Self-Collision	$-\sum \mathbb{1}(\text{collision})$	-10
Termination	$-\mathbb{1}(\text{terminated})$	-200
Foot Slip	$-\sum \ v_{\text{xy}}^{\text{foot}}\ \cdot \mathbb{1}(\text{contact})$	-0.5
Foot Contact Angular Velocity	$-\sum \ \omega^{\text{foot}}\ \cdot \mathbb{1}(\text{contact})$	-0.3
Joint Smoothness	$-\sum (0.02 \dot{q}_t^2 + \ddot{q}_t^2)$	-10^{-6}

D. Complete Reward Function Definition

The total reward r_t is a weighted combination of multiple tracking rewards and regularization penalties. The definition and weight of each reward term are detailed in Table VII. For quaternion-based rigid body rotation tracking, q^w denotes the scalar component of the quaternion. In regularization penalties, q_{soft} denotes soft joint position limits, \dot{q}_{limit} is the joint velocity limit, \ddot{q}_t is joint acceleration, and $\mathbb{1}(\cdot)$ is the indicator function.

All tracking rewards are summed with their respective weights, and the regularization penalty terms are negative. The total reward is $r_t = \sum_i w_i r_i^{\text{track}} + \sum_j w_j r_j^{\text{reg}}$, where all terms are weighted by their respective coefficients and then scaled by the timestep $\Delta t = 0.02$ s.