

LangWBC: Language-directed Humanoid Whole-Body Control via End-to-end Learning

Yiyang Shao Xiaoyu Huang Bike Zhang Qiayuan Liao Yuman Gao
Yufeng Chi Zhongyu Li Sophia Shao Koushil Sreenath
University of California, Berkeley




Fig. 1: We propose a language-directed humanoid whole-body control framework that translates natural language commands into continuous robot actions through a Conditional Variational Autoencoder (CVAE). The structured latent space brought by the CVAE enables smooth transitions between diverse and agile behaviors, as shown in the sequence where the robot seamlessly transitions from walking to running, concluding with a hand-waving motion prompted by the corresponding text commands. See more experiments at <https://youtu.be/9AN0GulqWwc>

Abstract—General-purpose humanoid robots are expected to interact intuitively with humans, enabling seamless integration into daily life. Natural language provides the most accessible medium for this purpose. However, translating language into humanoid whole-body motion remains a significant challenge, primarily due to the gap between linguistic understanding and physical actions. In this work, we present an end-to-end, language-directed policy for real-world humanoid whole-body control. Our approach combines reinforcement learning with policy distillation, allowing a single neural network to interpret language commands and execute corresponding physical actions directly. To enhance motion diversity and compositionality, we incorporate a Conditional Variational Autoencoder (CVAE) structure. The resulting policy achieves agile and versatile whole-body behaviors conditioned on language inputs, with smooth transitions between various motions, enabling adaptation to linguistic variations and the emergence of novel motions. We validate the efficacy and generalizability of our method through extensive simulations and real-world experiments, demonstrating robust whole-body control. Please see our website at LangWBC.github.io for more information.

I. INTRODUCTION

Humanoid robots hold immense potential for integration into human environments due to their anthropomorphic design, particularly in areas such as healthcare, personal assistance, and interactive services. For such robots to be truly effective, especially for users with limited technical proficiency such as the elderly, intuitive interaction modalities are essential. Natural language stands out as the most accessible and natural medium for human-robot communication, enabling users to

convey complex instructions effortlessly.

However, translating natural language commands into dynamic, agile, and robust whole-body motions for humanoid robots remains a significant challenge. The fundamental challenge stems from two interconnected aspects: First, as a motion generation problem, the system needs to produce diverse and generalizable movements that accurately reflect the intent of varied natural language commands. Second, as a real-world control problem, it must ensure these generated motions are physically executable while maintaining balance and stability under environmental uncertainties and disturbances. These two aspects are tightly coupled – generated motions must be physically realizable, while the control strategy must be flexible enough to accommodate the diversity of language-commanded behaviors. While prior works on language-directed real-world humanoid control have shown success by decoupling the problem into kinematic motion generation and whole-body tracking control [34, 10, 25], this hierarchical approach has key limitations. The generated motions are often physically implausible – e.g., lower-body floating in the air or upper-body exceeding stability margins – forcing the tracking policy to trade off between accurately tracking these motions and maintaining balance. Moreover, these methods are restricted to fixed-duration motions, limiting their ability to handle disturbances or ensure smooth transitions between motions.

In this work, we introduce LangWBC, a framework that addresses these dual challenges through a single end-to-end

model, eliminating the inherent conflicts between motion generation and physical feasibility. This approach enables humanoid robots to execute agile and diverse whole-body motions from natural language commands with flexible duration, smoothly transition between sequential actions, and synthesize diverse and novel motions through latent space interpolation.

LangWBC uses a two-stage training process. First, a teacher policy is trained via reinforcement learning to track retargeted motion capture (MoCap) data, acquiring a rich repertoire of dynamic and physically plausible behaviors. Then, a student policy based on a Conditional Variational Autoencoder (CVAE) [36] is trained via behavior cloning to learn the mapping from natural language commands and proprioceptive history *directly* to control actions, forming a joint distribution of language and actions within a structured and unified latent space.

We demonstrate the capabilities of LangWBC through extensive simulations and real-world experiments. The robot successfully executes agile motions, including running and quickly turning around, as well as expressive motions like waving and clapping. It also exhibits robustness to disturbances, such as recovering from kicks while executing textual commands. Furthermore, our framework enables smooth transitions between motion clips and generates novel motions through interpolation, demonstrating generalization beyond the training data.

Our key contributions are summarized as follows:

- We propose a novel framework that maps natural language commands *directly* to whole-body robot actions in a closed-loop control setup, achieving agile and robust performance suitable for real-world deployment.
- Our method enables the generation of diverse motions, smooth transitions, and adaptability to a wide range of textual inputs, including the synthesis of novel behaviors through latent space interpolation using the CVAE architecture.
- We validate our approach extensively on a physical humanoid robot, demonstrating its practical applicability, robustness to disturbances, and ability to execute complex whole-body motions from natural language commands.

II. RELATED WORK

This work explores the intersection of learning-based humanoid whole-body control and generative action modeling. Here, we review prior research in both fields.

A. Learning-based Humanoid Whole-body Control

Learning-based controllers have demonstrated the ability to perform complex whole-body control for humanoid robots. In physics-based animation, robots have learned various dynamic tasks [28, 29, 38, 16], object interactions [43, 24, 7], and even full-body motions from the AMASS dataset [23, 39, 14].

However, transferring these controllers to real-world hardware faces challenges due to the sim-to-real gap. As a result, prior work has largely focused on specialized controllers for a

limited set of agile motions on bipedal robots, such as walking [4, 21, 31, 8, 22], jumping [19, 42], and running [20, 37]. More recent efforts aim to scale up the range of feasible humanoid motions. For example, [1, 15, 5] incorporate dozens to a few hundred motions, while [11, 10, 6, 12] focus on a curated subset of the AMASS dataset. Unlike policies trained for specific skills, these approaches treat motion generation as a tracking problem: the policy learns to track kinematic trajectories but does not inherently compose them into downstream tasks. Instead, an additional high-level planner is required to execute specific motions at test time. In contrast, our work directly learns a text-conditioned motion generation policy, enabling both robust sim-to-real transfer and the ability to generate and execute diverse motions as a downstream task within a single framework.

B. Generative Action Modeling

Prior approaches to model actions with a generative model can be broadly categorized into two approaches: hierarchical kinematics-based tracking and end-to-end action generation.

1) *Hierarchical Kinematics-based Tracking*: A common approach is to use a hierarchical framework, where a high-level generative model produces diverse kinematic motions conditioned on inputs such as text [40, 35], keyframes [3], obstacles [17], etc., while a low-level tracking controller learns to follow these trajectories. Most prior works adopt this scheme. For example, OmniH2O [10] uses a pre-trained fixed-length MDM model [40] for text-conditioned motion generation, followed by a tracking controller. HumanPlus [6] uses behavior cloning to learn a high-level policy from human tele-operation to output target poses for downstream tasks. Exbody2 [15] separately trains a CVAE to generate kinematic motions autoregressively, but lacks text conditioning.

While hierarchical methods have proven effective, they require complex frameworks and often suffer from artifacts or physically infeasible motions in generated trajectories, such as floating bodies, foot sliding, and penetration. To address these challenges, Robot MDM [34] incorporates a learned Q-function to refine motion generation and enhance feasibility. However, this adds additional training overhead.

2) *End-to-end Action Generation*: An alternative is to model control actions directly using a generative approach, eliminating the gap between kinematics generation and tracking. While this method offers advantages in continuity and coherence, it remains under-explored due to its complexity, particularly in high-dimensional dynamic control.

In robotic manipulation, diffusion-based policies [2, 41] have demonstrated this approach for quasi-static manipulators. For legged robots, DiffuseLoco [13] extends it to dynamic control but is limited to state-based commands and quadrupedal robots. A more recent work, UH-1 [25], introduces text-to-action generation but only supports open-loop control, making it less robust to real-world disturbances. In this work, we present a fully functional text-conditioned end-to-end generative controller for humanoid robots. Our model not only

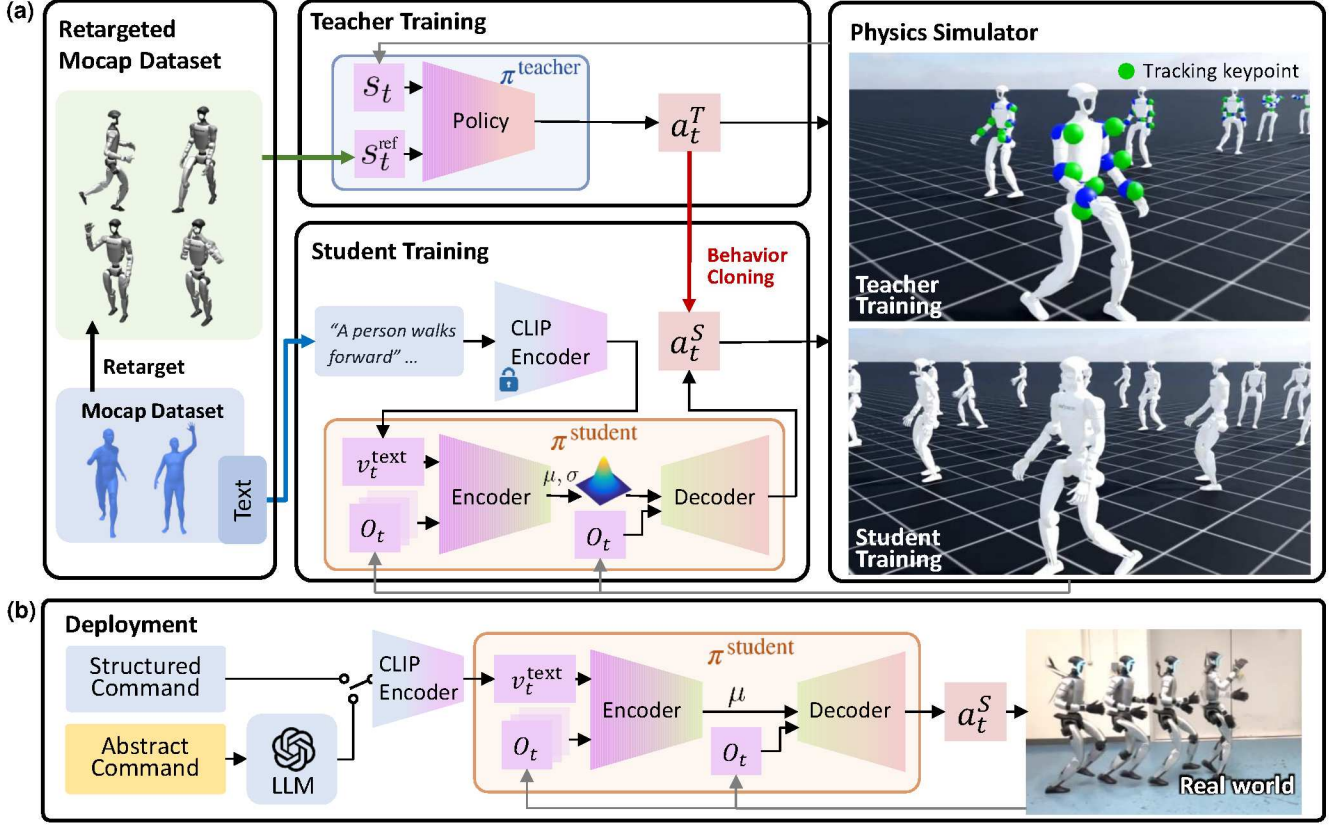


Fig. 2. **The Overview of the Training Framework.** The training process includes a motion-tracking teacher training phase and a language-directed student training phase. We first retarget the MoCap dataset and train a teacher policy via reinforcement learning. Then, a student policy, leveraging a CVAE architecture, jointly models high-level linguistic instructions and low-level physical actions of the teacher policy in a unified latent space. During deployment, we use the student policy for zero-shot sim-to-real transfer on hardware, demonstrating diverse behaviors.

enables robust real-world deployment but also generates novel, unseen motions while generalizing to similar text commands.

III. METHODS

In this section, we present LangWBC, an end-to-end framework that jointly models high-level linguistic instructions and low-level physical actions, enabling robots to execute complex whole-body motions directly from language commands.

We begin by training a language-agnostic teacher policy to learn and track a diverse set of human motions. A CVAE student policy is then used to align these physically-plausible motions with language inputs, forming a unified latent space that captures the joint distribution of language and actions. This latent space facilitates generalization, smooth interpolation, and seamless behavior transitions. Simultaneously, by training the CVAE with behavior cloning, we transfer the privileged teacher policy to the student policy that operates solely on proprioceptive inputs, enabling zero-shot sim-to-real transfer using onboard sensors without additional training. An overview of this framework is illustrated in Fig. 2.

A. Motion-Tracking Teacher Policy

The teacher policy is designed solely as a motion-tracking policy to track complex human motions without language understanding. The training process of the teacher policy involves two stages, motion retargeting and motion tracking.

1) *Motion Retargeting*: To ensure the MoCap trajectories are kinematically feasible for the teacher policy to track, we perform motion retargeting by applying inverse kinematics (IK) based on the Levenberg–Marquardt (LM) algorithm [26]. We formulate the retargeting as a nonlinear least squares optimization problem that minimizes the position and orientation errors between the robot and MoCap keypoints, while incorporating smoothness constraints to ensure natural transitions between frames. The optimization is solved using the LM algorithm with joint limit constraints, yielding kinematically feasible motions that closely match the original MoCap data. The detailed formulation and implementation are provided in Appendix B.

2) *Motion Tracking*: The primary objective of the teacher policy is to accurately track the retargeted MoCap trajectories without language information. Therefore, we employ a simple neural network architecture consisting of a multi-layer percep-

tron (MLP) with layer sizes of 512, 256, and 128 units.

The teacher policy can be formulated as:

$$a_t^T = \pi^{\text{teacher}}(s_t, s_t^{\text{ref}}), \quad (1)$$

where $s_t \in \mathbb{R}^{175}$ represents robot states, including both proprioceptive states and privileged information (friction, mass, external perturbations, and motor properties) available only in simulation [18], and $s_t^{\text{ref}} \in \mathbb{R}^{141}$ is the reference motion, specifically, the future five-frame keypoint positions in body frame and reference joint positions from the retargeted motion. The inclusion of privileged information (in s_t) enhances the policy's ability to master complex dynamic skills by providing additional context about the environment and the robot's physical properties. The definition of each specific input state can be found in Appendix A. The action output $a_t^T \in \mathbb{R}^{27}$ corresponds to the desired joint positions for the low-level PD controllers. We apply domain randomization for the teacher policy, with details provided in Appendix D.

Since MoCap datasets contain highly agile motions that are difficult to track in the early stages of training, including the entire datasets often leads to high gradient variance and slow convergence. To improve training efficiency, we design a motion curriculum that gradually increases motion complexity, allowing the policy to adapt progressively to more challenging motions.

We categorize the motions into two levels of difficulty:

- a) Easy motions: static or quasi-static movements, typically characterized by low-speed motions.
- b) Hard motions: agile motions that require more dynamic whole-body coordination, including actions such as sudden turns or rapid running.

Training begins with easy motions, and gradually we add hard motions as tracking performance improves. With this curriculum, the teacher policy learns a wide range of physical skills required for executing diverse motions.

The teacher policy is trained using Proximal Policy Optimization (PPO) [33] to minimize the discrepancy between the robot's movements and the reference motions. To encourage symmetry in the learned policy, we also incorporate symmetry-based data augmentation and an additional symmetry loss. Specifically, for each state-action pair (s_t, a_t^T) , we generate its mirrored counterpart $(s_t^m, a_t^{T,m})$ through left-right reflection. The augmented training objective is formulated as

$$\mathcal{L}_{\text{teacher}} = \mathcal{L}_{\text{PPO}} + \lambda_{\text{sym}} \mathcal{L}_{\text{sym}}, \quad (2)$$

where λ_{sym} is a weighting coefficient, and \mathcal{L}_{sym} encourages consistent policy outputs for mirrored states, i.e.,

$$\mathcal{L}_{\text{sym}} = \mathbb{E}_{s_t \sim \mathcal{D}} [\|\pi^{\text{teacher}}(s_t) - \mathcal{M}(\pi^{\text{teacher}}(s_t^m))\|^2]. \quad (3)$$

Here, $\mathcal{M}(\cdot)$ denotes the mirroring operation for actions. This symmetry constraint helps the policy learn more balanced and natural movements while reducing the sample complexity of training. The tracking reward formulation is summarized in Table I. The teacher policy runs at 50 Hz.

TABLE I
REWARD FUNCTION COMPONENTS FOR TEACHER POLICY

Term	Expression	Weight
Z linear vel penalty	$\ v_z^{\text{root}}\ ^2$	-0.2
XY angular vel penalty	$\ \omega_{xy}^{\text{root}}\ ^2$	-0.05
Joint torque penalty	$\sum \tau_j^2$	-2×10^{-6}
Joint acc penalty	$\sum \ddot{\theta}_i^2$	-1×10^{-7}
Joint action rate penalty	$\sum (\Delta a_i^T)^2$	-0.05
Energy cost	$\sum \ \tau_i \cdot \theta_i\ ^2$	-1×10^{-6}
Termination penalty	$\mathbb{I}_{\text{terminated}}$	-200
Joint limit penalty	$\mathbb{I}_{\theta_i \notin [\theta_{\min}, \theta_{\max}]}$	-1
Orientation penalty	$\ \theta_{zy}^{\text{proj}}\ ^2$	-10.0
Feet slide penalty	$\mathbb{I}(F_{\text{feet}} > 100N) \cdot \sum \ v_{xy}^{\text{feet}}\ ^2$	-0.1
Hip joint deviation	$\sum \theta_{\text{hip}} - \theta_{\text{default}} $	-0.03
Leg joint deviation	$\sum \theta_{\text{leg}} - \theta_{\text{default}} $	-0.01
Keypoint tracking	$\exp\left(-\frac{\ p_{\text{key}} - p_{\text{ref}}\ ^2}{2}\right)$	1.0
Joint tracking	$\exp\left(-\frac{\ \theta - \theta_{\text{ref}}\ ^2}{4}\right)$	1.0
Single stance reward	$\mathbb{I}_{\substack{\Delta h_{\text{ref}}^{\text{feet}} > 0.05 \\ t_{\text{stance}} \in [0.1, 0.5]}} \cdot t_{\text{stance}}$	1.5

Notations: $\mathbb{I}(\cdot)$ denotes indicator function, $\|\cdot\|$ is Euclidean norm, τ : joint torque, θ : joint angle, v : velocity, ω : angular velocity, F_{feet} : ground reaction force, t_{stance} : single-foot stance duration.

B. Language-Directed Student Policy

To enable the robot to interpret and act on natural language commands, we design a CVAE-based student policy that encodes textual instructions and physical actions into a unified latent space, using only language inputs and proprioceptive readings.

The input of the student policy consists of two parts:

- 1) **Text Caption Embedding:** We utilize the CLIP text encoder [30] to convert the input natural language command c_t^{text} into a fixed-length embedding vector

$$v_t^{\text{text}} = f_{\text{CLIP}}(c_t^{\text{text}}) \in \mathbb{R}^{512}. \quad (4)$$

This embedding captures the semantic meaning of the text command.

- 2) **History of Proprioceptive Observations:** Instead of providing the full privileged state used in the teacher policy, we provide only proprioceptive observations $o_t \in \mathbb{R}^{90}$ to the student policy, which encapsulate joint positions, joint velocities, base linear velocities, base angular velocities, and projected gravity. We input a sequence of historical observations and actions, sampled at 10 Hz over a 2-second window, yielding a 20-step trajectory of input-output pairs.

The encoder processes the concatenated textual and observational inputs to produce the parameters of a latent Gaussian distribution, outputting a mean vector $\mu \in \mathbb{R}^{128}$ and a diagonal covariance matrix represented by $\sigma \in \mathbb{R}^{128}$. This architecture models the conditional distribution of robot motions given text commands through the latent space, where the text embedding serves as a conditioning signal that shapes the latent distribution. During training, we sample the latent vector z using the

standard reparameterization trick:

$$\mu_t, \sigma_t = \pi_{\text{enc.}}^{\text{student}}(o_t, \dots, o_{t-20}, v_t^{\text{text}}), \quad (5)$$

$$z_t = \mu_t + \sigma_t \odot \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, I), \quad (6)$$

$$a_t^S = \pi_{\text{dec.}}^{\text{student}}(z_t, o_t), \quad (7)$$

where \odot denotes element-wise multiplication, $\pi_{\text{enc.}}^{\text{student}}$ is the student encoder, and $\pi_{\text{dec.}}^{\text{student}}$ denotes the decoder. This reparameterization allows gradients to flow through the sampling process. The decoder then takes the sampled latent vector z_t along with the latest state observation to output the action. We use an MLP with layer sizes of 2048, 1024, and 512 units for the encoder, and an MLP with layer sizes of 512, 1024, and 2048 units for the decoder. During inference, we simply use the mean μ_t of the encoded distribution as the latent vector, eliminating the sampling step to ensure deterministic behavior. The student policy is applied with the same domain randomization as the teacher.

We employ the Dataset Aggregation (Dagger) algorithm [32] to train the student policy from the teacher policy with language labels. The training objective follows the variational lower bound

$$\mathcal{L}_{\text{student}} = \|a_t^T - a_t^S\|_2^2 + \lambda_{\text{KL}} D_{\text{KL}}(q_\phi(z_t | o_{t-20:t}, v_t^{\text{text}}) \| p(z_t)), \quad (8)$$

where D_{KL} is the KL-Divergence operator, and λ_{KL} balances reconstruction quality in behavior cloning with the structural regularization of the latent space.

The training process consists of five steps:

- 1) **Data Collection:** We simulate 1,024 parallel environments. At each time step, the student is given the language command and its history observation.
- 2) **Teacher Action Query:** For each state encountered by the student, the corresponding optimal action is obtained by querying the teacher policy.
- 3) **Experience Buffer Construction:** We insert the collected student’s observations and the teacher’s actions to a buffer of 1024×512 (500,000) state-action pairs.
- 4) **Loss Computation:** In the early stage of training, the student policy results in large accumulated errors that push the teacher policy out of its training distribution. To mitigate this, instead of tracking absolute positions, the student policy tracks displacement relative to past positions. Let p_t be the robot’s root position at time t , with $\Delta p_t = p_t - p_{t-\Delta t}$ representing its displacement over interval Δt , and $\Delta p_{\text{ref},t} = p_{\text{ref},t} - p_{\text{ref},t-\Delta t}$ denoting the reference displacement. The robot’s tracking objective then becomes minimizing the error between its own displacement and the reference displacement

$$\min \|\Delta p_t - \Delta p_{\text{ref},t}\|^2. \quad (9)$$

This mitigates deviations from the reference motion and preserves the quality of teacher demonstrations.

- 5) **Policy Update:** We update the student with the loss in (8). We use a batch size of 1024×64 and a learning rate of 1×10^{-5} , with one epoch per iteration. We then use the student’s actions to step the environment.




Fig. 3. **Robustness to External Disturbances.** The humanoid robot demonstrates robust stability while executing a hand-waving motion under external perturbations. When subjected to kicks (top row) and pushes (bottom row), the robot maintains balance and continues the commanded motion, showcasing effective disturbance rejection capabilities without interrupting the primary task.

We repeat the iterative process, where the student progressively learns to replicate the teacher’s behavior while understanding both language inputs and its own observation history. The student policy also runs at 50 Hz.

IV. EXPERIMENTS

We conduct extensive experiments to evaluate our framework for language-directed humanoid whole-body control with a Unitree G1 humanoid robot. We begin with an overview and demonstrate diverse motions enabled by our approach. We then analyze the learned latent space and its contribution to the policy’s generalization to unseen commands, highlight key features such as smooth transitions and latent interpolation, and follow up with an ablation study on core design choices. Finally, we showcase a complex LLM-guided compositional task, illustrating the full capabilities of LangWBC. We use Isaac Lab [27] for training.

A. Diverse Humanoid Motions

In order to learn a diverse set of motions, we utilize the HumanML3D dataset [9] in training the teacher policy, which provides human MoCap data annotated with textual descriptions. For deployment, we use an AMD Ryzen 9 CPU for inference. As shown in Fig. 3 and 4, the robot successfully executes a diverse range of upper- and lower-body motions in response to natural language commands, including walking in different directions, turning, performing hand gestures, and executing more complex whole-body movements, while remaining robust to external perturbations such as heavy kicks and pushes. For a diverse set of whole-body motions, we provide more results in Appendix C. Through these demonstrations, our framework exhibits zero-shot sim-to-real transfer capabilities, effectively addressing both core challenges through a unified network – generating diverse,




Fig. 4. **Real World Demonstration.** Conditioned on text commands, our framework is able to learn a diverse distribution of whole-body motions in action generation directly, and can be zero-shot deployed on real-world robots. More results are shown in the accompanying video.

language-aligned motions while maintaining robust control under real-world conditions and disturbances.

B. Latent Space Analysis

One key advantage of using a CVAE as the student policy – rather than a simple MLP network as in previous works [10, 15] – is that it provides a structured latent space. This structured latent space aligns language inputs and motion actions to the same latent codes, enabling the model to learn disentangled representations where each latent variable captures both semantic meaning and motion dynamics. As a result, variations in latent space correspond not only to distinct motion styles and transitions but also to meaningful differences in language instructions. This allows for better generalization to unseen commands, smoother motion interpolation, and more coherent transitions between behaviors.

To verify this property, we apply the t-SNE algorithm to embed the high-dimensional latent codes of various motions into a 2D plane. As shown in Fig. 5, we plot nine different motions from four categories (walking, raising the left or right hand, and clapping), with each category color-coded.

From the visualization, we see that the latent space has several interpretable features. First, motions in each category form distinct clusters, reflecting clear separation by motion type. Second, there is a striking symmetry: raising the left hand versus the right hand appears mirrored about the center of the latent space, near which more symmetric motions (e.g., walking or clapping with both hands) are located. Third, all motions exist in a common region near the origin, which we interpret as the “standing” latent code – every motion begins and ends in a standing pose. Overall, this analysis confirms that the CVAE learns a meaningful latent manifold, which we




Fig. 5. **t-SNE Analysis of Latent Space.** The plot shows 9 motions from 4 categories of motion, as shown in the legend. We see that similar motions (in the same color band) are closer than dissimilar ones. The axes suggest an interpretable structure: lateral symmetry (left/right motions mirrored across the x-axis) and vertical hierarchy (upper-body motions cluster at higher y-values, lower-body motions at lower y-values). We observe that all motions share a common region near the origin (0,0), likely representing a typical standing posture.

can leverage to perform smooth transitions between motions and generate novel, unseen motions.

C. Generalization to Unseen Texts

Understanding language variations is key to human communication, and essential for flexible communication with robots as well. Since our CVAE’s latent space is inherently structured, where semantically similar commands cluster together while remaining distinct from dissimilar ones, we hypothesize it provides robustness to unseen yet contextually similar commands.

To verify this hypothesis, we qualitatively evaluate the policy’s response to three semantically similar text commands, as illustrated in Fig. 6. We find the policy performs forward motion in a consistent speed and style despite phrasing differences like “move” vs. “walk,” demonstrating robustness to linguistic variation.

However, this generalization capability may also stem from the CLIP text encoder itself. To isolate the contribution of the CVAE architecture, we compare our approach (**CLIP+CVAE**) to a baseline that pairs an MLP with a CLIP encoder alone (**CLIP+MLP**). We evaluate the generated motion quality of both models on a test set of 15 unseen commands spanning three categories:

- 1) *Similar* (e.g., *Walk slowly*),
- 2) *Moderately different* (e.g., *Walk into store*),
- 3) *Semantically distant* (e.g., *Jump*).

Shown in Table II, both variants perform similarly on commands close to the training data. However, as commands become more unfamiliar, **CLIP+CVAE** consistently produces higher-quality motions. We posit that it is because, while the




Fig. 6. **Rollouts of Unseen Text Commands.** Our method can generalize to unseen text commands with similar semantical meanings. Of the three, only one command “A person walks forward” (a) is included in the training dataset.

TABLE II
MOTION QUALITY METRIC ON UNSEEN COMMANDS.

Model	Similar	Moderate	Different
CLIP+CVAE	80.92%	69.58%	54.62%
CLIP+MLP	80.62%	64.20%	50.28%

The Motion Quality Metric is a weighted sum of keypoint and joint errors, normalized by an exponential function to (0, 1] (see Table I).

CLIP encoder handles minor linguistic variations well, it produces significantly different encodings for out-of-distribution commands, which the MLP policy struggles to generalize from. In contrast, the CVAE’s structured joint latent space reduces the effective distance between novel and seen commands, mitigating overfitting to the training set and enabling more plausible motions for out-of-distribution commands.

Summarizing the results, we conclude that our CVAE-based approach generalizes more effectively to language variation, compared to non-structured MLP baselines. This makes it well-suited for integration with large language models (LLMs) in more complex reasoning tasks, as shown in Section IV-G.

D. Smooth Transitions Between Agile Motions

Furthermore, a major benefit of the structured latent space is the ability to transition smoothly between motions. We demonstrate this capability through examples of agile motion transitions, a challenging task in humanoid control. As shown in Fig. 1, the robot seamlessly transitions from walking to running, then gradually slows to a stop before waving its hand. In another rollout shown in Fig. 7, the robot performs an upper-body movement (waving its hand), then begins running, stops, and waves again. Given the high-dimensional




Fig. 7. **Smooth Transitions between Different Text Commands.** The humanoid robot seamlessly executes a sequence of actions: waving its right hand, transitioning into running, coming to a stop, and concluding with another hand wave. The policy demonstrates the ability to handle diverse motion transitions within a single execution, without requiring resets between different actions.

dynamics of humanoid motion, achieving smooth and coherent transitions – such as running, stopping, and switching to limb movements – within a single policy, without requiring resets, is a significant challenge that has not been demonstrated in prior works. These results underscore the diversity and robustness of our method, as well as its ability to generalize to transition motions that were underrepresented in the training dataset.

Additionally, we observe that the policy can autonomously transition from the end of a motion back to its beginning, enabling seamless looping. As shown in Fig. 4a, where the reference consists of a single turn, the policy can perform two consecutive turns without interruption. Although the starting and ending poses can be quite different, the policy successfully infers the correct timing for transitioning to ensure a continuous motion.

E. Interpolation in Latent Space

Since the CVAE creates a smooth and continuous latent space, we find that it also enables the generation of novel motions via a meaningful interpolation between latent codes. This is achieved by encoding the current observation and CLIP-encoded text into latent codes via the CVAE encoder, and then interpolating them in the CVAE latent space and decoding to generate the corresponding action.

To illustrate, we show an example of interpolation between the two distinct commands: “a man walks forward briskly” and “a person shuffles from left to right, then shuffles back to the left”. As shown in Fig. 8, the policy generates a novel yet intuitive motion - blending forward walking with lateral shuffling - despite the absence of such behavior in the training data. This interpolated movement emerges naturally as a result of latent space blending, demonstrating the capacity to synthesize new behaviors from learned patterns. Moreover, the robot’s movement stays agile and stable, demonstrating the framework’s robustness to unseen latent codes.

TABLE III
ABLATION STUDY ON THE PROPOSED ARCHITECTURE, SYMMETRY LOSS, AND STUDENT TRACKING OBJECTIVES

Metric	CVAE	No-Symm	No-Rel	MLP
Motion Quality \uparrow	96.2%	91.6%	93.8%	91.9%
Stability \uparrow	99.10%	98.50%	96.60%	96.81%
Imitation Loss \downarrow	0.085	0.107	0.098	0.091

All metrics are reported after 10k iterations. The Motion Quality metric is defined in Table II, the Stability metric indicates success rate over 1000 steps under perturbations, while the Imitation Loss is the student’s final loss in training.

Similar to the ablation in Section IV-C, since both the CLIP encoder and CVAE architecture have interpolatable latent spaces, this raises a natural question: Is CLIP alone sufficient for interpolating diverse motion commands, or does the CVAE provide essential capabilities? To answer this question, we again isolate the effectiveness of the CVAE latent space by comparing its interpolation performance against a CLIP-only MLP baseline, assessing which latent mixing generates meaningful novel behaviors.

- Interpolating *directly* in the CLIP text-embedding space (**CLIP Interpolation**).
- Interpolating in the *CVAE latent space* learned by the student (**CVAE Interpolation, Ours**).

As shown in Fig. 9, we see that, with the help of the CVAE architecture, the interpolated motion results in a smooth diagonal walk, whereas the baseline with the CLIP encoder alone produces noticeable jitter and difficulty in walking (as seen by the orange CoM trajectory). This result highlights that the CVAE induces a smoother and more structured latent space than the CLIP encoder alone, enabling better generalization to unseen motions through meaningful motion interpolation.

These results highlight the strength of our CVAE-based architecture in capturing motion space structure, enabling novel motions from dataset diversity. This allows our proposed method to generate flexible, diverse motions and generalize beyond explicit training examples.

F. Ablation Study

In the ablation study, we comprehensively evaluate the effectiveness of the three core design choices in our proposed framework: (i) the symmetry loss applied during teacher training, (ii) the relative-tracking objective used for the student, and (iii) the CVAE architecture of the student policy. Below, we summarize our full framework alongside the ablation baselines used to assess the contribution of each component. We conduct the ablations in simulation.

- 1) **LangWBC (CVAE, Ours)**: Our proposed method, in which the teacher is trained with a symmetry loss, and the student leverages a CVAE architecture trained using the relative-tracking objective.
- 2) **No Teacher Symmetry (No-Symm)**: A variant of LangWBC in which the symmetry loss is removed from the teacher’s training, while all other components remain unchanged.




Fig. 8. **Interpolation in the Latent Space.** The CVAE gives a structured latent space, enabling the policy to generalize to interpolated command. Here, interpolating between walking (Command 1) and side stepping (Command 2) produces walking to the side, a whole-body motion that does not exist in the training distribution.




Fig. 9. **Latent Space Interpolation: CLIP+CVAE vs. CLIP Alone** Comparison of motion quality when interpolating between forward and side-ways walking. The CLIP+CVAE model (left) produces smooth and coherent diagonal walking, while the CLIP-only baseline (right) results in jittery motions.

- 3) **No Relative Tracking (No-Rel):** A variant where the student is trained *without* the relative-tracking objective, while all other components remain unchanged.
- 4) **MLP Student (MLP):** A variant that replaces the student’s CVAE architecture with an MLP, while all other components remain unchanged.

The quantitative results, summarized in Table III, indicate that our full framework outperforms all ablation baselines, confirming the contribution of each proposed component to efficient learning and high-quality motion generation. Notably, the CVAE architecture not only enhances generalization, as discussed earlier, but also improves motion tracking accuracy and robustness compared to the MLP baseline.

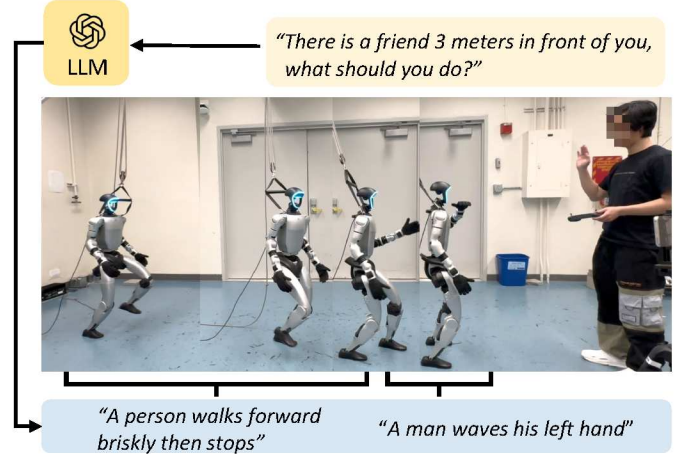


Fig. 10. **LLM-guided Humanoid Motion Sequence.** Given the social scenario “There is a friend 3 meters in front”, the LLM decomposes this high-level instruction into primitive motion commands, which the robot executes by walking forward, stopping, and greeting with a hand wave.

G. Integrating LLMs for Complex Tasks

To handle abstract instructions and social scenarios, we integrate our framework with a large language model (LLM) as a high-level planner, as shown in Fig. 2(b). The LLM translates unstructured natural language inputs into structured command sequences executable by our end-to-end policy.

However, LLMs tend to provide commands that are semantically similar to, but not identical to the training data. Thus, the generalization to similar text commands provided by the CVAE becomes critical in this task.

We prompt the LLM to decompose abstract tasks into motion primitives similar to the training dataset, using the structure “(Text Command): (Duration in Seconds).” Given the social scenario “There is a friend 3 meters in front of you, what should you do?”, the LLM generates an intuitive sequence of timed commands:

- “A person walks forward briskly then stops: 4.0”
- “A man waves his right hand: 5.0”

During deployment, we run the CLIP encoder at a lower frequency and reuse its embedding across multiple control steps. As shown in Fig. 10, the robot successfully performs this socially appropriate sequence, demonstrating our framework’s ability to execute high-level, context-aware instructions.

V. LIMITATIONS

We have demonstrated a set of humanoid whole-body motions directed by natural language commands. However, the number of language-conditioned motions remains limited to only several dozens of motions due to compute constraints. Expanding the range of language-directed humanoid motions on a larger scale could further validate and enhance our approach. Additionally, the current examples predominantly focus on locomotion-oriented whole-body control due to the absence of a vision module. In future work, we aim to incorporate more whole-body motions that can be leveraged in agile locomanipulation tasks critical for the deployment of humanoid robots. Due to the limited expressiveness of the variational autoencoder model, we experience some sim-to-real gap which could have been improved via incorporating a more expressive generative model, such as diffusion models, which could capture more subtle diversity in domain randomizations.

VI. CONCLUSION

In this work, we presented an end-to-end language-directed humanoid whole-body control framework. It combines learning-based whole-body control with generative action modeling, allowing a single neural network to interpret language commands and execute corresponding physical actions robustly on humanoid hardware in a zero-shot manner. Our framework demonstrates the following characteristics:

Diversity, Generalization, Smooth Transition: Our framework enables diverse humanoid whole-body motions and composes novel behaviors through latent space interpolation. Leveraging the structured latent space which models the distribution of actions conditioned on language semantics and proprioceptive observations, the policy generates a wide range of motion sequences, including agile behaviors like running, conditioned on language inputs. Furthermore, our method generalizes to contextually similar commands and generates novel motion compositions beyond the training distribution.

Moreover, it enables smooth transitions between agile motions, such as from walking to running, with simple language commands, a capability not demonstrated in prior work.

End-to-End Learning: Unlike prior works, our language-directed humanoid whole-body control framework is trained end-to-end, mapping language commands directly to actions without intermediate representations or complex planning modules. By aligning text embeddings with the motion latent space, our approach effectively bridges the gap between language understanding and physical robot skills.

We believe that learning a language-conditioned generative action model for complex humanoid behaviors marks a significant step toward developing a foundation model for humanoid control, paving the way for more intuitive, adaptable, and generalizable deployment of humanoid robots in real-world environments.

VII. ACKNOWLEDGMENT

This work was supported in part by The Robotics and AI Institute. We thank Jiaze Cai and Kaixuan Wang for help with hardware experiment. We thank Vishnu Sangli for helpful discussions. K. Sreenath has financial interests in The Robotics and AI Institute. He and the company may benefit from the commercialization of the results of this research.

REFERENCES

- [1] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. *arXiv preprint arXiv:2402.16796*, 2024.
- [2] Cheng Chi, Zhenjia Xu, Siyuan Feng, Eric Cousineau, Yilun Du, Benjamin Burchfiel, Russ Tedrake, and Shuran Song. Diffusion policy: Visuomotor policy learning via action diffusion. *The International Journal of Robotics Research*, 0(0), 2024. doi: 10.1177/02783649241273668.
- [3] Setareh Cohan, Guy Tevet, Daniele Reda, Xue Bin Peng, and Michiel van de Panne. Flexible motion in-betweening with diffusion models. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–9, 2024.
- [4] Helei Duan, Bikram Pandit, Mohitvishnu S Gadde, Bart Van Marum, Jeremy Dao, Chanh Kim, and Alan Fern. Learning vision-based bipedal locomotion for challenging terrain. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 56–62. IEEE, 2024.
- [5] 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.
- [6] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. *arXiv preprint arXiv:2406.10454*, 2024.
- [7] Jiawei Gao, Ziqin Wang, Zeqi Xiao, Jingbo Wang, Tai Wang, Jinkun Cao, Xiaolin Hu, Si Liu, Jifeng Dai, and

- Jiangmiao Pang. Coohoi: Learning cooperative human-object interaction with manipulated object dynamics. *arXiv preprint arXiv:2406.14558*, 2024.
- [8] 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.
- [9] Chuan Guo, Shihao Zou, Xinxin Zuo, Sen Wang, Wei Ji, Xingyu Li, and Li Cheng. Generating diverse and natural 3d human motions from text. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5152–5161, June 2022.
- [10] 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. *arXiv preprint arXiv:2406.08858*, 2024.
- [11] 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 *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8944–8951. IEEE, 2024.
- [12] Tairan He, Wenli Xiao, Toru Lin, Zhengyi Luo, Zhenjia Xu, Zhenyu Jiang, Jan Kautz, Changliu Liu, Guanya Shi, Xiaolong Wang, et al. Hover: Versatile neural whole-body controller for humanoid robots. *arXiv preprint arXiv:2410.21229*, 2024.
- [13] Xiaoyu Huang, Yufeng Chi, Ruofeng Wang, Zhongyu Li, Xue Bin Peng, Sophia Shao, Borivoje Nikolic, and Koushil Sreenath. Diffuseloco: Real-time legged locomotion control with diffusion from offline datasets. *arXiv preprint arXiv:2404.19264*, 2024.
- [14] Xiaoyu Huang, Takara Truong, Yunbo Zhang, Fangzhou Yu, Jean Pierre Sleiman, Jessica Hodgins, Koushil Sreenath, and Farbod Farshidian. Diffuse-cloc: Guided diffusion for physics-based character look-ahead control. *arXiv preprint arXiv:2503.11801*, 2025.
- [15] 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.
- [16] Jordan Juravsky, Yunrong Guo, Sanja Fidler, and Xue Bin Peng. Superpadl: Scaling language-directed physics-based control with progressive supervised distillation. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–11, 2024.
- [17] Korrawe Karunratanakul, Konpat Preechakul, Supasorn Suwajanakorn, and Siyu Tang. Guided motion diffusion for controllable human motion synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2151–2162, 2023.
- [18] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47):eabc5986, 2020.
- [19] 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 *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2811–2817. IEEE, 2021.
- [20] 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*, 44(5):840–888, 2024. doi: 10.1177/02783649241285161.
- [21] Qiayuan Liao, Bike Zhang, Xuanyu Huang, Xiaoyu Huang, Zhongyu Li, and Koushil Sreenath. Berkeley humanoid: A research platform for learning-based control. *arXiv preprint arXiv:2407.21781*, 2024.
- [22] 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.
- [23] Zhengyi Luo, Jinkun Cao, Kris Kitani, Weipeng Xu, et al. Perpetual humanoid control for real-time simulated avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10895–10904, 2023.
- [24] Zhengyi Luo, Jinkun Cao, Sammy Christen, Alexander Winkler, Kris Kitani, and Weipeng Xu. Grasping diverse objects with simulated humanoids. *arXiv preprint arXiv:2407.11385*, 2024.
- [25] Jiageng Mao, Siheng Zhao, Siqi Song, Tianheng Shi, Junjie Ye, Mingtong Zhang, Haoran Geng, Jitendra Malik, Vitor Guizilini, and Yue Wang. Learning from massive human videos for universal humanoid pose control. *arXiv preprint arXiv:2412.14172*, 2024.
- [26] Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [27] Mayank Mittal, Calvin Yu, Qinxi Yu, Jingzhou Liu, Nikita Rudin, David Hoeller, Jia Lin Yuan, Ritvik Singh, Yunrong Guo, Hammad Mazhar, Ajay Mandlekar, Buck Babich, Gavriel State, Marco Hutter, and Animesh Garg. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023. doi: 10.1109/LRA.2023.3270034.
- [28] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018.
- [29] 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.
- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya

- Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Aspell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [31] 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. doi: 10.1126/scirobotics.adi9579.
- [32] Stéphane Ross, Geoffrey Gordon, and Drew Bagnell. A reduction of imitation learning and structured prediction to no-regret online learning. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 627–635. JMLR Workshop and Conference Proceedings, 2011.
- [33] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [34] Agon Serifi, Ruben Grandia, Espen Knoop, Markus Gross, and Moritz Bächer. Robot motion diffusion model: Motion generation for robotic characters. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–9, 2024.
- [35] Yi Shi, Jingbo Wang, Xuekun Jiang, Bingkun Lin, Bo Dai, and Xue Bin Peng. Interactive character control with auto-regressive motion diffusion models. *ACM Transactions on Graphics (TOG)*, 43(4):1–14, 2024.
- [36] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. *Advances in neural information processing systems*, 28, 2015.
- [37] Annan Tang, Takuma Hiraoka, Naoki Hiraoka, Fan Shi, Kento Kawaharazuka, Kunio Kojima, Kei Okada, and Masayuki Inaba. Humanmimic: Learning natural locomotion and transitions for humanoid robot via wasserstein adversarial imitation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 13107–13114. IEEE, 2024.
- [38] Chen Tessler, Yoni Kasten, Yunrong Guo, Shie Mannor, Gal Chechik, and Xue Bin Peng. Calm: Conditional adversarial latent models for directable virtual characters. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–9, 2023.
- [39] Chen Tessler, Yunrong Guo, Ofir Nabati, Gal Chechik, and Xue Bin Peng. Maskedmimic: Unified physics-based character control through masked motion inpainting. *ACM Transactions on Graphics (TOG)*, 43(6):1–21, 2024.
- [40] Guy Tevet, Sigal Raab, Brian Gordon, Yoni Shafir, Daniel Cohen-or, and Amit Haim Bermano. Human motion diffusion model. In *The Eleventh International Conference on Learning Representations*, 2023.
- [41] Dian Wang, Stephen Hart, David Surovik, Tarik Kestemur, Haojie Huang, Haibo Zhao, Mark Yeatman, Jiuguang Wang, Robin Walters, and Robert Platt. Equivariant diffusion policy. In *8th Annual Conference on Robot Learning*, 2024.
- [42] Chong Zhang, Wenli Xiao, Tairan He, and Guanya Shi. Wococo: Learning whole-body humanoid control with sequential contacts. *arXiv preprint arXiv:2406.06005*, 2024.
- [43] Yunbo Zhang, Deepak Gopinath, Yuting Ye, Jessica Hodgins, Greg Turk, and Jungdam Won. Simulation and retargeting of complex multi-character interactions. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023.

APPENDIX

A. Teacher Policy Input State

The input of the teacher policy is defined as follows.

- **Base Linear Velocity** (3 dimensions): The robot's current linear velocity in the base frame.
- **Base Angular Velocity** (3 dimensions): The robot's current angular velocity in the base frame.
- **Projected Gravity** (3 dimensions): The gravity vector projected onto the robot's coordinate frame.
- **Actual Keypoint Positions** (7 links \times 3 dimensions = 21): The positions of key body links (e.g., limbs and torso) in the robot's body frame that need to be tracked.
- **Joint Positions** (27 dimensions): The current positions of the robot's joints.
- **Joint Velocities** (27 dimensions): The current velocities of the robot's joints.
- **Last Action** (27 dimensions): The last action executed by the robot.
- **Privileged Information** (64 dimensions): Simulation-exclusive information encompasses: 1) physical properties including static friction, dynamic friction, and restitution coefficients (3 dims); 2) joint armature (27 dims); 3) body link mass (28 dims); 4) external force (3 dims); and 5) external torque (3 dims).
- **Target Keypoint Positions** (6 frames \times 7 links \times 3 dimensions = 126): Future desired positions of keypoints over six time steps, providing a trajectory for the robot to follow.
- **Degree-of-Freedom (DoF) Commands** (15 dimensions): Commands specifying desired joint configurations.

B. Motion Retargeting Details

We formulate the retargeting problem as a nonlinear least square optimization over the entire motion sequence. Let T be the number of frames, and $q_t \in \mathbb{R}^n$ be the vector of robot joint angles at frame t . Our objective is to find the set of joint angles that minimize the following cost function:

$$\min_{\{q_t\}} \sum_{t=1}^T \left(\|x_{\text{robot},t}(q_t) - x_{\text{mocap},t}\|^2 + w_{\text{ori}} \|\Delta r_t(q_t)\|^2 \right) + w_{\text{smooth}} \sum_{t=2}^T \|q_t - q_{t-1}\|^2, \quad (10)$$

subject to:

$$q_{\min} \leq q_t \leq q_{\max}, \quad \forall t = 1, \dots, T. \quad (11)$$

To solve this optimization problem, we employ the Levenberg-Marquardt algorithm, which is suitable for solving nonlinear least squares problems in IK. At each iteration, we linearize the residuals around the current estimate $q_t^{(k)}$ and compute the update $\Delta q_t = q_t^{(k+1)} - q_t^{(k)}$ by solving:

$$(J_k^\top J_k + \lambda I + w_{\text{smooth}} S^\top S) \Delta q = -J_k^\top f_k. \quad (12)$$

Where:

- $x_{\text{robot},t}$ denotes the positions of the robot's keypoints at frame t as a function of the joint angles q_t .
- $x_{\text{mocap},t}$ represents the target keypoint positions from the mocap data.
- $\Delta r_t(q_t)$ represents the orientation error between the robot's end-effectors and the mocap data at frame t , calculated using the Lie-algebra error.
- w_{ori} is the weighting factor for the orientation error term.
- w_{smooth} is the weighting factor for the smoothness term.
- q_{\min} and q_{\max} define the joint limit boundaries.
- $\Delta q = [\Delta q_1^\top, \Delta q_2^\top, \dots, \Delta q_T^\top]^\top$ is the stacked update vector for all frames.
- J_k is the Jacobian matrix of residuals at iteration k , combining position and orientation terms across all frames.
- f_k is the stacked residual vector at iteration k , defined as:

$$f_k = \begin{bmatrix} x_{\text{robot},1}(q_1^{(k)}) - x_{\text{mocap},1} \\ \sqrt{w_{\text{ori}}} \Delta r_1(q_1^{(k)}) \\ \vdots \\ x_{\text{robot},T}(q_T^{(k)}) - x_{\text{mocap},T} \\ \sqrt{w_{\text{ori}}} \Delta r_T(q_T^{(k)}) \\ \sqrt{w_{\text{smooth}}}(q_2^{(k)} - q_1^{(k)}) \\ \sqrt{w_{\text{smooth}}}(q_3^{(k)} - q_2^{(k)}) \\ \vdots \\ \sqrt{w_{\text{smooth}}}(q_T^{(k)} - q_{T-1}^{(k)}) \end{bmatrix} \quad (13)$$

- λ is the damping parameter that balances the Gauss-Newton and gradient descent methods, ensuring convergence and stability in the LM algorithm.
- I is the identity matrix.
- S is the sparse matrix representing the smoothness constraints, defined as:

$$S = \begin{bmatrix} -I & I & & & \\ & -I & I & & \\ & & \ddots & \ddots & \\ & & & -I & I \end{bmatrix} \quad (14)$$

The term $w_{\text{smooth}} S^\top S$ in Eq. (12) adds regularization that enforces smooth transitions between consecutive frames, promoting natural motion. By solving this equation, we obtain the updates Δq_t for all frames, which are used to update the joint angles:

$$q_t^{(k+1)} = q_t^{(k)} + \Delta q_t. \quad (15)$$

After each update, we enforce the joint limits by projecting $q_t^{(k+1)}$ onto the feasible set defined in Eq. (11).

Iteratively applying the LM algorithm under these constraints yields kinematically feasible motions that closely match the original mocap data, thus providing suitable reference trajectories for the teacher policy.

C. Diverse Whole-body Motions on Real Robot

We demonstrate our framework's capability to generate and execute diverse whole-body motions on real hardware. As

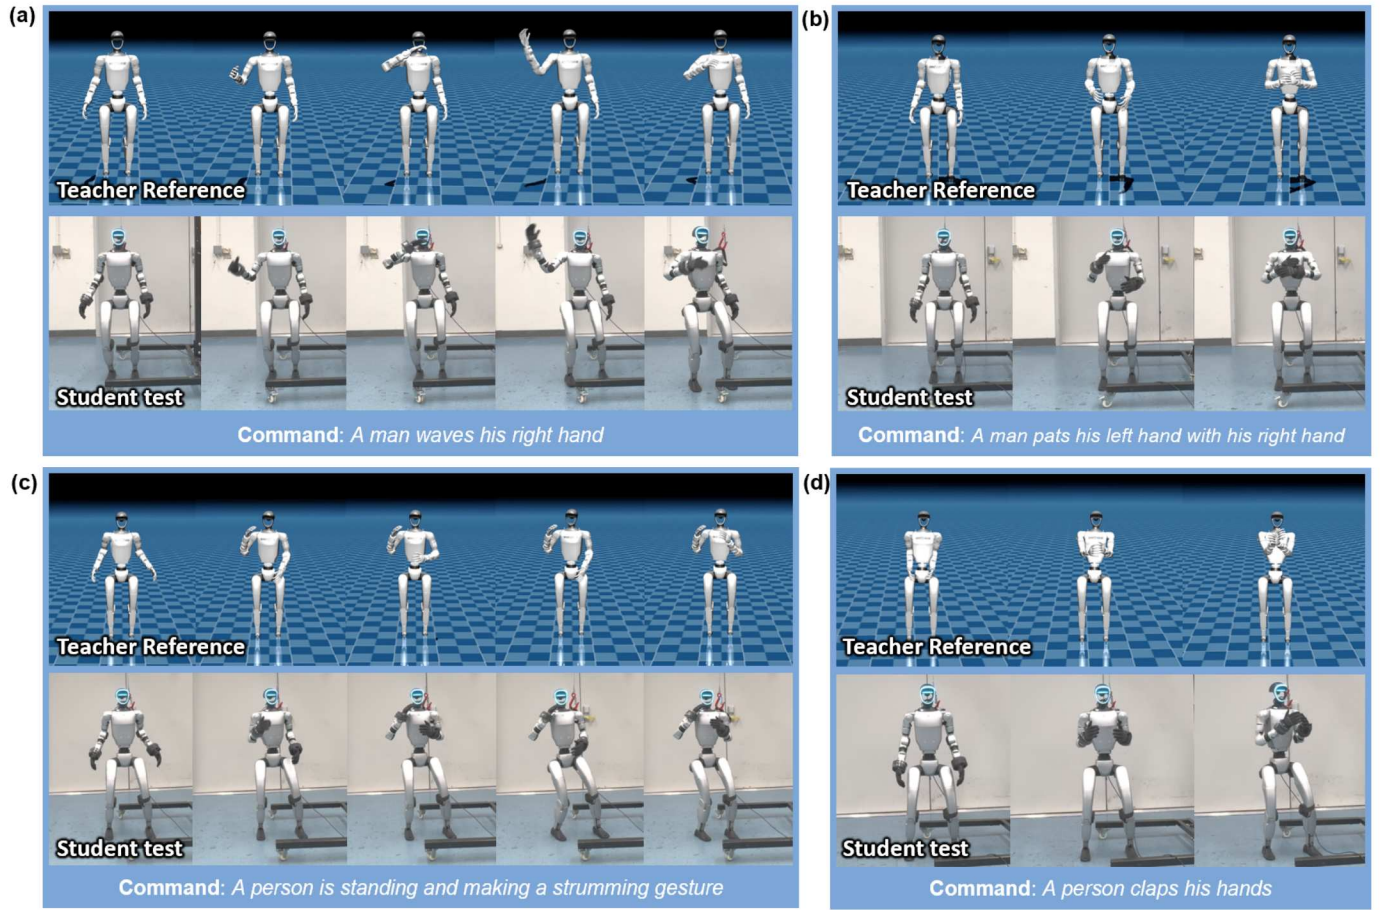


Fig. 11. **Upper-body Motion Examples.** Our framework generates diverse upper-body movements including reaching, lifting, and manipulation tasks. The learned motions can be directly deployed on the real robot.




Fig. 12. **Lower-body Motion Examples.** The framework also enables various lower-body movements such as stepping, squatting and balancing. These motions are also successfully transferred to the real robot without additional training.

TABLE IV
DOMAIN RANDOMIZATION PARAMETERS

Mode	Component	Range
Reset	Initial velocity (linear)	$[-0.5, 0.5]$ m/s
	Initial velocity (angular)	$[-0.5, 0.5]$ rad/s
Reset	Joint positions	$[0.5, 1.5] \times \text{default}$
	Joint velocities	$[0.0, 0.0]$ rad/s
Startup	Ankle friction (static)	$[0.2, 0.6]$
	Ankle friction (dynamic)	$[0.2, 0.6]$
	Ankle restitution	$[0.0, 0.4]$
Startup	Link masses	$[0.9, 1.1] \times \text{default}$
	Base mass modification	$[-1.0, 1.0]$ kg (additive)
	Joint armature	$[0.8, 1.2] \times \text{default}$
Startup	Joint default positions	$[-0.05, 0.05]$ rad (additive)
Reset	Base torque	$[-5.0, 5.0]$ Nm
Interval	Push robot (every 10-15s)	$[-1.0, 1.0]$ m/s (x,y)

shown in Fig. 11 and Fig. 12, our method can learn a rich distribution of both upper-body and lower-body movements, and successfully transfer them to the physical robot with zero-shot. The accompanying video provides more comprehensive demonstrations of these motions.

D. Domain Randomization Configuration

Table IV describes the domain randomization parameters applied to the robot. "Reset" parameters are applied at environment reset, "Startup" parameters are applied once during initialization, and "Interval" parameters are applied periodically during simulation.