

Embracing Bulky Objects with Humanoid Robots: Whole-Body Manipulation with Reinforcement Learning

Chunxin Zheng, Kai Chen, Zhihai Bi, Yulin Li, Liang Pan, Jinni Zhou,
Haoang Li, and Jun Ma, *Senior Member, IEEE*

Abstract—Whole-body manipulation (WBM) for humanoid robots presents a promising approach for executing embracing tasks involving bulky objects, where traditional grasping relying on end-effectors only remains limited in such scenarios due to inherent stability and payload constraints. This paper introduces a reinforcement learning framework that integrates a pre-trained human motion prior with a neural signed distance field (NSDF) representation to achieve robust whole-body embracing. Our method leverages a teacher-student architecture to distill large-scale human motion data, generating kinematically natural and physically feasible whole-body motion patterns. This facilitates coordinated control across the arms and torso, enabling stable multi-contact interactions that enhance the robustness in manipulation and also the load capacity. The embedded NSDF further provides accurate and continuous geometric perception, improving contact awareness throughout long-horizon tasks. We thoroughly evaluate the approach through comprehensive simulations and real-world experiments. The results demonstrate improved adaptability to diverse shapes and sizes of objects and also successful sim-to-real transfer. These indicate that the proposed framework offers an effective and practical solution for multi-contact and long-horizon WBM tasks of humanoid robots. The open-source project can be found at <https://github.com/ChunxinZHENG/Embracing-Bulky-Objects-with-Humanoid-Robots>.

I. INTRODUCTION

Recent advances in humanoid robots have facilitated their deployment in diverse domains, such as industrial applications and home services, where they are increasingly required to perform whole-body manipulation (WBM) tasks [1]–[3]. In these applications, humanoid robots are unavoidably required to grasp and transport bulky objects. Unlike the manipulation of relatively small objects with an end-effector, humanoid robots can leverage the whole body (including arms and torso) to envelop and embrace objects in a human-like manner, which could distribute contact forces across the body to stably manipulate bulky objects. However, there are several key considerations for successful completion of such

This work was supported by the National Key Research and Development Program of China under Grant 2024YFB4707903. (*Corresponding Author: Jun Ma.*)

Chunxin Zheng, Kai Chen, Zhihai Bi, Yulin Li, Jinni Zhou, and Haoang Li are with the Robotics and Autonomous Systems Thrust, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou 511453, China (e-mail: czheng739@connect.hkust-gz.edu.cn).

Liang Pan is with the School of Computing and Data Science, The University of Hong Kong, Hong Kong SAR, China (e-mail: liangpan1005@connect.hku.hk).

Jun Ma is with the Robotics and Autonomous Systems Thrust, The Hong Kong University of Science and Technology (Guangzhou), Guangzhou 511453, China, and also with the Cheng Kar-Shun Robotics Institute, The Hong Kong University of Science and Technology, Hong Kong SAR, China (e-mail: jun.ma@ust.hk).




Fig. 1. Whole-body manipulation task for a humanoid robot. The task mainly consists of three stages: (i) Approaching: the robot moves toward the object; (ii) Embracing: the robot establishes a whole-body enveloping grasp around the object; (iii) Transporting: the robot carries the object to a target location. The red dot denotes the object's geometric center, and the yellow dot indicates the robot's center of mass.

challenging tasks. First, it requires an advanced perception system that can accurately model the spatial relationships between the robot and objects in its environment. Also, it necessitates a highly effective method that empowers humanoid robots to acquire anthropomorphic behaviors through a whole-body control policy.

Essentially, WBM has witnessed significant success, which shows noteworthy performance through multi-contact strategies [4]. It leverages the entire body and even environmental supports to execute more forceful, stable, and dexterous manipulation of bulky objects [5]. This makes it particularly effective for tasks that surpass the force or kinematic constraints of traditional grasping strategies with the end effector. Existing WBM approaches typically rely on the simplified robot model and kinematic relationships to estimate the spatial positions of the robot's links (usually simplified as multiple line segments) [6]. However, such methods are unable to accurately model or perceive the precise geometry of the surfaces of the links, leading to significant errors in position estimation during contact-rich manipulation tasks. As a result, these inaccuracies could ultimately result in task failures.

In general, control methods for WBM can be categorized into model-based [7], [8] and learning-based approaches [9]–[11]. Model-based WBM methods need to specify the contact points and determine the contact sequence, but the computation is highly expensive due to the high degrees of freedom of humanoid robots. Learning-based methods, which generally depend on carefully engineered rewards, struggle to produce robust and coherent behaviors in these contact-intensive, dynamically unstable scenarios. In the existing literature, the developments on WBM for humanoid robots are still severely limited. Especially for tasks like carrying bulky

objects with WBM, it demands a sequence of coordinated actions as illustrated in Fig. 1, which include approaching the object, embracing the object (by establishing multi-contact grasps and lifting with whole-body stabilization), and transporting objects to a designated location. In essence, acquiring the anthropomorphic skill of manipulating bulky objects with the whole body is extremely challenging due to: 1) the *inherent high complexity* of humanoid robot systems with their numerous degrees of freedom, 2) the *difficulty in constructing stable contact* through precise perception and control, and 3) the *long-horizon nature* of manipulation tasks that require sustained coordination and planning.

To overcome these limitations, we propose a reinforcement learning (RL) framework that integrates a pre-trained human motion prior with a neural signed distance field (NSDF) representation, tailored to the whole-body embracing manipulation task of humanoid robots. The NSDF provides precise and efficient self-modeling capability that enhances the robot’s spatial self-awareness and robustness of interactions. The human motion prior offers a biologically plausible kinematic distribution that stabilizes policy training and facilitates convergence in complex manipulation tasks. Unlike conventional grasping with an end-effector only, our method dynamically coordinates the arms and torso of the humanoid robot to embrace bulky objects. In this sense, it overcomes the limitations of end-effector-only manipulation and significantly enhances payload capacity. Crucially, the system maintains stability through distributed multi-contact interactions, ensuring balance while executing whole-body motions.

This work makes several significant contributions to the field, which are outlined as follows:

- We propose an RL framework for humanoid robots to manipulate bulky objects using their entire bodies, including arms and torso. To the best knowledge of the authors, this is the first RL framework that enables the humanoid robots to actively engage with bulky objects through whole-body manipulation.
- We introduce the motion prior into the embracing policy training pipeline, which speeds up the convergence in policy training toward such multi-contact and long-horizon tasks. Also, it empowers the humanoid robots to acquire anthropomorphic skills through the proposed framework.
- We construct an NSDF representation of the humanoid robot that precisely perceives the robot-object interactions. It is further utilized to design both the observation space and reward function that guides the upper body to maintain sustained contact with the object, which significantly enhances the robustness of manipulation during the task.
- We thoroughly validate the proposed approach through simulations and real-world experiments on a humanoid robot, and the results attained demonstrate noteworthy performance in complex whole-body embracing tasks.

II. RELATED WORK

A. Whole-Body Manipulation


Considering the manipulation methods using only end-effectors as the interactive part, such as dexterous grippers and end sticks, WBM introduces the ability to manipulate objects using all possible effective parts of the robot. This approach represents the ultimate form of anthropomorphic manipulation by coordinating the whole body of a robot [12]. Despite this potential, current research largely focuses on single or bimanual manipulation platforms. Early works, such as [13], employ model-based techniques to achieve whole-arm grasping on single-arm systems. Subsequently, research utilizing fixed-base humanoid platforms begin incorporating contact models involving dual arms, objects, and the torso to enable whole-arm manipulation [5]. However, the inherent complexity of contact-rich scenarios poses significant challenges for model-based methods. Consequently, recent approaches have turned to learning-based strategies. Notably, an RL method is presented in [9] to achieve WBM on a bimanual platform. Also, some researchers have sought to enhance contact perception in WBM by equipping robot links with soft pneumatic sensors [8]. These tactile enhancements provide richer contact feedback and facilitate more effective whole-body interaction. In the most recent state-of-the-art approaches, the integration of RL with whole-body tactile sensing has enabled more compliant and human-like motion generation for WBM on humanoid robots, but it is currently limited to upper-body platforms [11].

B. Reinforcement Learning for Humanoid Robot Control

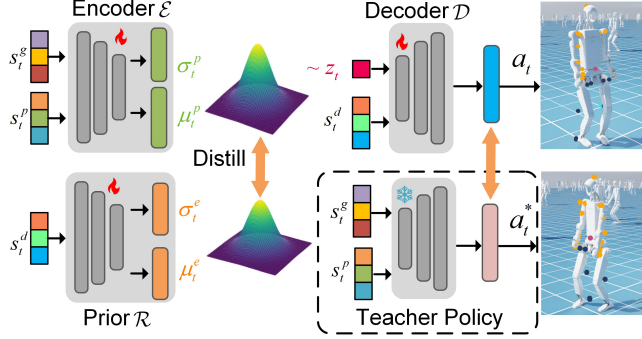
Significant progress has been made in humanoid robot control using RL, with extensive prior research dedicated to task-driven controllers for specialized behaviors, such as agile running [14] and locomotion across challenging terrains [15]. While these methods achieve high performance in their respective tasks, they often struggle to generalize to complex, long-horizon scenarios. Moreover, the bionicity of actions cannot be well guaranteed, or additional gait rewards need to be added during the training phase to guide the strategy. To address this limitation, an emerging trend leverages behavior cloning (BC) from physics-based animation [16], [17], where control policies are trained on large-scale human motion datasets (e.g., AMASS). This approach enables robots to imitate natural, human-like movements while maintaining physical feasibility. Although BC policy has great performance on whole-body tracking, it typically requires an additional action generator [3], [18] or policy distillation [19] to adapt to specific tasks (e.g., locomotion, manipulation, or teleoperation). Moreover, they remain limited in interactive tasks due to the absence of critical environmental information in human motion datasets, such as contact forces, terrain geometry, and dynamic interactions, which are essential for robust real-world deployment.

Recent advancements in humanoid robotics have focused on developing universal motion controllers that can generalize across various tasks. These approaches often utilize

(a) Training Framework of Whole-Body Embracing Policy



(b) Motion Prior Distillation



(c) Embracing Policy Training

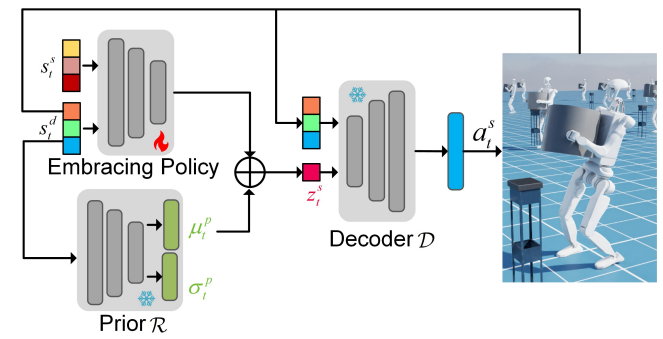


Fig. 2. System architecture for training of embracing task with whole-body manipulation. (a) **Training Framework:** The framework consists of two components, data processing and policy training. In data processing, we start from the AMASS mocap dataset, use MaskedMimic to discard sequences that violate kinematic constraints (e.g., self-collision and foot slip), and retarget the remaining motions to the robot morphology to collect robot-reference trajectories. In the policy training process, we first distill a human motion prior from the collected trajectories with a teacher policy and a student VAE policy, and then use the learned motion prior to guide the training of the high-level embracing policy. (b) **Motion Prior Distillation:** A pretrained teacher policy guides robotic motion imitation by providing reference actions. The student policy, structured as a VAE-based model, learns to reconstruct the teacher’s action distribution through imitation learning. This process constructs a learnable humanoid motion prior \mathcal{R} , which facilitates efficient training for downstream tasks. (c) **Embracing Policy Training:** Based on the learned humanoid motion prior \mathcal{R} , we train the whole-body manipulation policy for the embracing task.

pre-trained BC policies to construct a comprehensive action space. For example, Luo et al. [20] propose the use of human motion priors, distilled from BC policies, to create universal controllers. These controllers provide a robust foundation for further task-specific fine-tuning, facilitating the efficient adaptation of humanoid robots to diverse and complex tasks. In [21], it combines text description and pre-trained BC policy to build a text-motion mapping action space. Also, it has other methods by distilling different robotics skill policies to build action space servicing for universal humanoid robot controller [22].

III. METHODOLOGY

In this work, we propose an RL framework for humanoid robots to perform bulky-object embracing tasks, and the complete system architecture is illustrated in Fig. 2. In the following, we first present how to convert the AMASS mocap dataset to the kinematically feasible references of the robot (Section III-A). Then, with these references, we detail the learning process of the humanoid motion prior through a teacher-student structure (Section III-B), which provides a compact latent action space for downstream tasks. Next, in Section III-C, we introduce the embracing policy training approach based on the learned motion prior. Pertinent notations used in this work are listed in Table I for convenience.

TABLE I
NOMENCLATURE

Symbol	Dimension	Description
\mathbf{a}_t	\mathbb{R}^{27}	Target joint position
\mathbf{v}_t	\mathbb{R}^3	Root linear velocity
$\boldsymbol{\omega}_t$	\mathbb{R}^3	Root angular velocity
\mathbf{g}_t	\mathbb{R}^3	Projected gravity vector
\mathbf{q}_t	\mathbb{R}^{27}	Joint position
$\dot{\mathbf{q}}_t$	\mathbb{R}^{27}	Joint velocity
\mathbf{a}_{t-1}	\mathbb{R}^{27}	Previous action
$\hat{\mathbf{p}}_{t+1}$	$\mathbb{R}^{3 \times 27}$	Reference body position
$\hat{\mathbf{p}}_{t+1} - \mathbf{p}_t$	$\mathbb{R}^{3 \times 27}$	Positional offset

A. Data Processing

To enable humanoid robots to acquire fundamental movement skills from human demonstrations, the motion learning pipeline typically begins with the evaluation of raw motion capture data, represented using the SMPL human model. For example, MaskedMimic [23] processes and filters these input motions to identify kinematically feasible trajectories. The process leads to a cleaned human motion dataset, denoted as \mathcal{D}_h , which contains only physically plausible movement patterns. We then adopt the retargeting framework from H2O [24] to transfer these motions onto our target humanoid platform, explicitly accounting for differences in body proportions and joint ranges of motion. This produces the robot motion dataset, represented by \mathcal{D}_r , which provides validated trajectories to serve as training targets for the mimic teacher

policy.

B. Motion Prior Distillation

The goal of this module is to distill the motion priors based on the robot motion dataset, which retain human-like dexterity while ensuring physical feasibility for humanoid robots. Following well-established techniques [19], [20], we implement a teacher-student learning scheme. First, the retargeted robot motion dataset is utilized to train a teacher policy capable of precise trajectory tracking. Then, this teacher policy is distilled through a VAE-based student policy to construct the motion prior. The learned motion prior captures the essential characteristics of natural human movement, while remaining strictly aligned with the robot’s physical constraints.

1) *Teacher Policy*: Building upon the retargeted robot motion dataset, we develop a motion-tracking teacher policy using the PHC framework [17]. The policy π_{teacher} maps the current robot state and reference motion to actions:

$$\mathbf{a}_t^* = \pi_{\text{teacher}}(\mathbf{s}_t^p, \mathbf{s}_t^g), \quad (1)$$

where \mathbf{a}_t^* is the action generated by teacher policy, $\mathbf{s}_t^p = \{\mathbf{v}_t, \boldsymbol{\omega}_t, \mathbf{g}_t, \mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{a}_{t-1}^*\} \in \mathbb{R}^{90}$, which encodes the proprioceptive state through root linear velocity \mathbf{v}_t , angular velocity $\boldsymbol{\omega}_t$, projected gravity vector \mathbf{g}_t , joint positions \mathbf{q}_t , joint velocities $\dot{\mathbf{q}}_t$, and previous action \mathbf{a}_{t-1}^* . The goal state is defined as $\mathbf{s}_t^g = \{\hat{\mathbf{p}}_{t+1} - \mathbf{p}_t, \hat{\mathbf{p}}_{t+1}\} \in \mathbb{R}^{2 \times 3 \times 27}$, which incorporates reference motion trajectories from \mathcal{D}_r , specified through both positional offsets ($\hat{\mathbf{p}}_{t+1} - \mathbf{p}_t$) and target configurations $\hat{\mathbf{p}}_{t+1}$ for each rigid body. The complete specifications of all variables, including their respective dimensions and detailed descriptions, are summarized in Table I. The policy is implemented as a multi-layer perceptron (MLP) and optimized using Proximal Policy Optimization (PPO) [25].

2) *Humanoid Motion Prior Distillation*: To enable efficient and transferable control of humanoid motion, we aim to distill the complex action space learned by the teacher policy into a compact latent representation that naturally induces human-like behavior. To this end, we adopt a VAE-based distillation approach, following PULSE [20], which encodes diverse motion behaviors into a latent space that can be effectively leveraged for downstream tasks.

Specifically, we construct a variational encoder-decoder framework to represent the humanoid motion space. The encoder $\mathcal{E}(\mathbf{z}_t | \mathbf{s}_t^p, \mathbf{s}_t^g)$ infers a distribution over the latent motion variable \mathbf{z}_t given the current proprioceptive state \mathbf{s}_t^p and the goal state \mathbf{s}_t^g . The decoder $\mathcal{D}(\mathbf{z}_t | \mathbf{s}_t^d, \mathbf{z}_t)$ maps the latent code back to the action space. To address sim-to-real transfer, we design $\mathbf{s}_t^d = \{\boldsymbol{\omega}_t, \mathbf{g}_t, \mathbf{q}_t, \dot{\mathbf{q}}_t, \mathbf{a}_{t-1}\} \in \mathbb{R}^{87}$, which excludes the root linear velocity component \mathbf{v}_t from \mathbf{s}_t^p , making the representation feasible in real-world scenarios. Here, we use \mathbf{a}_t to represent the action reconstructed by the student policy, and \mathbf{a}_{t-1} to represent the previous student action at the last time. Furthermore, a learnable prior network $\mathcal{R}(\mathbf{z}_t | \mathbf{s}_t^d)$ is used to approximate the latent code distribution

produced by the encoder. Formally, we have the following:

$$\mathcal{E}(\mathbf{z}_t | \mathbf{s}_t^p, \mathbf{s}_t^g) = \mathcal{N}(\mathbf{z}_t | \boldsymbol{\mu}_t^e, \boldsymbol{\sigma}_t^e), \quad (2)$$

$$\mathcal{R}(\mathbf{z}_t | \mathbf{s}_t^d) = \mathcal{N}(\mathbf{z}_t | \boldsymbol{\mu}_t^p, \boldsymbol{\sigma}_t^p), \quad (3)$$

$$\mathcal{D}(\mathbf{a}_t | \mathbf{s}_t^d, \mathbf{z}_t) = \mathcal{N}(\mathbf{a}_t | \boldsymbol{\mu}_t^d, \hat{\boldsymbol{\sigma}}_t^d), \quad (4)$$

where \mathcal{D} outputs a Gaussian distribution over actions with a fixed diagonal covariance.

The overall training objective for the encoder, decoder, and prior is:

$$\mathcal{L}_{\text{all}} = \mathcal{L}_{\text{action}} + \alpha \mathcal{L}_{\text{regu}} + \beta \mathcal{L}_{\text{KL}}, \quad (5)$$

where $\mathcal{L}_{\text{action}} = \|\mathbf{a}_t^* - \mathbf{a}_t\|_2^2$ is the action reconstruction loss. The regularization term, $\mathcal{L}_{\text{regu}} = \|\boldsymbol{\mu}_t^e - \boldsymbol{\mu}_{t-1}^e\|_2^2$, penalizes abrupt changes in the latent trajectory by enforcing temporal consistency between consecutive latent means, where $\boldsymbol{\mu}_t^e$ and $\boldsymbol{\mu}_{t-1}^e$ are the mean of the encoder’s output distribution at time t and $t-1$, respectively. The KL divergence term, \mathcal{L}_{KL} , aligns the encoder’s latent distribution with the learned prior. The coefficients α and β coordinate the relative importance of the smoothness and KL regularization terms, respectively.

Once the distillation is complete, we freeze the parameters of the decoder \mathcal{D} and the prior \mathcal{R} , thus defining a new low-dimensional action space for downstream task learning.

C. Embracing Policy Training

In this subsection, we formulate the task policy π_{task} for the WBM task, especially for bulky objects. This task can be generally divided into three stages: approaching the object, embracing the object, and transporting the object to the target location. To address this long-horizon manipulation task with bulky objects, we develop a tailored PPO algorithm incorporating three key components: stage randomized initialization, specific rewards, and humanoid motion prior distribution.

1) *Observations and Actions*: The task observation at time step t is defined as $\mathbf{s}_t^{\text{task}} = \{\mathbf{s}_t^d, \mathbf{s}_t^s\}$, which is the concatenation of the decoder input \mathbf{s}_t^d and a set of task-specific features \mathbf{s}_t^s . Specifically, we define

$$\mathbf{s}_t^s = \{\hat{\rho}_t^{\text{box}}, \hat{\theta}_t^{\text{box}}, \hat{\rho}_t^{\text{target}}, \hat{\theta}_t^{\text{target}}, \mathbf{d}_t\} \in \mathbb{R}^{19},$$

where $\hat{\rho}_t^{\text{box}}$ and $\hat{\theta}_t^{\text{box}}$ denote the distance and direction angle differences between the robot torso and the center of the box; $\hat{\rho}_t^{\text{target}}$ and $\hat{\theta}_t^{\text{target}}$ denote the distance and direction angle differences between the robot torso and the target position. Notably, the term $\mathbf{d}_t \in \mathbb{R}^{15}$ comprises NSDF features, which are computed by a pre-trained network f_θ evaluating the shortest distances between a selected target point and a set of mesh segments [26], as shown in Fig. 3. In our implementation, we choose the center of the box as the target point $\mathbf{p}_t^{\text{target}}$ and the robot’s upper links as the mesh collection, so that the nearest distances to the robot links are defined as:

$$\mathbf{d}_t = f_\theta(\mathbf{p}_t^{\text{target}}). \quad (6)$$




Fig. 3. Illustration of NSDF between the target object and the robot's upper body. The dots distributed along each link of the robot represent the closest point from the corresponding joint to the object's center of mass. The dashed lines indicate the direction of the NSDF between the object's center of mass and the robot's links.

The task policy π_{task} operates in this latent space, producing a latent action \mathbf{z}_t^s given the current task state \mathbf{s}_t^s :

$$\mathbf{z}_t^s = \pi_{\text{task}}(\mathbf{s}_t^s), \quad (7)$$

$$\mathbf{a}_t^s = \mathcal{D}(\mathbf{z}_t^s + \mu_t^p). \quad (8)$$

As shown in Fig. 2(c), the learnable prior \mathcal{R} generates a motion prior μ_t^p . The output of π_{task} is then combined with this motion prior to form a new latent token, which is decoded into the final robot actions \mathbf{a}_t^s .

2) *Reward Design*: We design multiple reward functions to train the humanoid robot for the WBM task. The detailed descriptions of rewards on the smoothness and physical limitation are presented in Table II.

Smoothness Reward: The smoothness of the policy is encouraged by the reward term $r_{\text{smooth}} = r_{\text{torque}} + r_{\text{acc}} + r_{\text{action}}$, where r_{torque} penalizes high joint torques, r_{acc} penalizes large actuator accelerations, and r_{action} penalizes rapid changes in action values. Collectively, these terms help guide the policy towards generating smooth and physically plausible actions.

Physical Limitation Reward: To protect the robot during deployment, we introduce $r_{\text{limit}} = r_{\text{dof}} + r_{\text{slippage}} + r_{\text{feet}}$, where r_{dof} encourages the joints to stay within their physical limits and discourages configurations near joint boundaries, r_{slippage} enhances the stability of the robot. By limiting the contact force on the robot's foot, r_{feet} prevents the robot from stepping too heavily on the ground, making its movements smoother.

Task Reward: The total task reward is given by $r_{\text{task}} = r_{\text{walk}} + r_{\text{carry}} + r_{\text{arm}} + r_{\text{NSDF}}$. The WBM task is decomposed into three distinct stages: approaching the object, embracing it, and transporting it to the target location. To support stage-specific reward allocation, a circular *pick-up zone* of radius $d_p = 0.35$ m is centered around the object, acting as a transition boundary between approaching and manipulation. An annular region with inner radius $d_i = 3.5$ m and outer radius $d_o = 4.0$ m serves dual roles: it is used as the initial spawning area for the robot during the approaching phase and as the target delivery region for the object upon task completion.

TABLE II
DESIGN OF REWARD TERMS

Reward Terms	Definition	Weight
Torque	$\ \boldsymbol{\tau}\ _2^2$	-1e-7
Joint Acceleration	$\ \ddot{\mathbf{q}}\ _2^2$	-2.5e-8
Action Rate	$\ \mathbf{a}_{t-1} - \mathbf{a}_t\ _2^2$	-0.5
Feet Slippage	$\sum \mathbf{v}_{\text{foot}} \cdot (\mathbf{f}_{\text{contact}} > 1)$	-0.05
Feet Contact Force	$\sum [\max(0, \mathbf{f}_{\text{contact}} - F_{\text{max}})]$	-1e-5
Joint Angle Limitation	$\sum \left[-\min(0, \mathbf{q} - \mathbf{q}_{\text{lowerlimit}}) + \max(0, \mathbf{q} - \mathbf{q}_{\text{upperlimit}}) \right]$	-1e-3

In the first stage (when the robot approaches the object), the robot is not in the *pick-up zone*. In this sense, r_{walk} encourages the robot to move closer to the object, r_{carry} and r_{arm} are set to zero. The reward function is formulated as:

$$r_{\text{walk}} = \begin{cases} 1, & \text{In pick-up zone,} \\ \exp(-\|\sigma \hat{p}_t^{\text{box}}\|_2) + \exp(-\|\sigma \hat{\theta}_t^{\text{box}}\|_2) + \exp(-\|\sigma \hat{v}_t^{\text{body}}\|_2), & \text{Out of pick-up zone.} \end{cases} \quad (9)$$

In the second and third stages when robot embraces and transports the object, it stays in the *pick-up zone*, r_{carry} encourages the box to move closer to the final target. In this case, the reward function is defined as:

$$r_{\text{carry}} = \begin{cases} 0, & \text{Out of pick-up zone,} \\ \exp(-\|\sigma \hat{p}_t^{\text{target}}\|_2) + \exp(-\|\sigma \hat{\theta}_t^{\text{target}}\|_2) + \exp(-\|\sigma \hat{v}_t^{\text{box}}\|_2), & \text{In pick-up zone.} \end{cases} \quad (10)$$

Additionally, r_{arm} is used to adjust the arm position:

$$r_{\text{arm}} = \exp(-\sigma(\|\hat{p}_t^{\text{lh}}\|_2 + \|\hat{p}_t^{\text{rh}}\|_2)) + \exp(-\sigma(\|\hat{h}_t^{\text{lh}}\|_2 + \|\hat{h}_t^{\text{rh}}\|_2)), \quad (11)$$

where \hat{p}_t^{lh} and \hat{p}_t^{rh} are the distances from the left and right end-effectors to the object, respectively. Meanwhile, \hat{h}_t^{lh} and \hat{h}_t^{rh} denote the distance between the height of the end-effectors and height of the target object.

Additionally, we evaluate the proximity between the robot's upper body and the object:

$$r_{\text{nsdf}} = \exp(-\sigma \|\mathbf{d}_t\|_1). \quad (12)$$

3) *Random Initialization*: To address long-horizon tasks, we randomly generate initial scenarios that are categorized into three stages. The first scenario type involves placing the target object on a table while randomly initializing the robot's position in the initial spawning area. This configuration allows the training to focus primarily on the first stage of the task. In the second scenario, the object remains placed on the table, but the robot's initial state is manually configured such that its upper joints are positioned in a pre-hugging posture, and the robot's position is initialized directly in the pick-up zone. This setup is designed to optimize training for the object pickup phase. The third scenario initializes the robot's




Fig. 4. Effectiveness of the NSDF module. (a) Results with the NSDF module. (b) Results without the NSDF module. The red circles indicate contact regions. As shown, the model with NSDF generates more contact points during manipulation, significantly enhancing grasping stability and success rate.

joints directly in a hugging posture, with the object generated in its arms. This configuration facilitates training for the final stage of the task. During training, one of these three scenarios is randomly selected upon resetting the environment.

IV. EXPERIMENTS

In this section, we evaluate the proposed framework through both simulation and real-world experiments. The training environment is constructed using the Isaac Sim simulator, with 4096 parallel agents deployed for efficient data collection. The target object to be embraced is represented as a cylinder measuring $\Phi 42 \text{ cm} \times 40 \text{ cm}$. Initially, the object is placed 4 meters away from the robot. The robot is then required to embrace the object and transport it to a target location situated 4 meters away from the object's original position, thereby accomplishing an integrated whole-body manipulation task. Due to the absence of suitable open-source baselines in this domain, we mainly perform comparative experiments to validate the effectiveness and performance of the proposed modules. All training and experiments are conducted on a single workstation equipped with an NVIDIA RTX 4080 GPU.

A. Validation of Our Framework

To thoroughly evaluate the effectiveness of our proposed WBM framework, we design a set of experiments to analyze its core components. Specifically, we investigate two key aspects: the importance of the NSDF module in guiding the robot's whole-body motion toward the target object, and the impact of randomization during early training on improving convergence efficiency.

1) *Effectiveness of the NSDF Module:* To evaluate the impact of this design, we adopt identical experimental settings without the NSDF-related observations from the input space and disable the NSDF reward term. As illustrated in Fig. 4(a), after training for the same number of epochs, the robot trained with our proposed method successfully establishes multi-point contact using multiple upper-body joints during object transportation. This ability results in

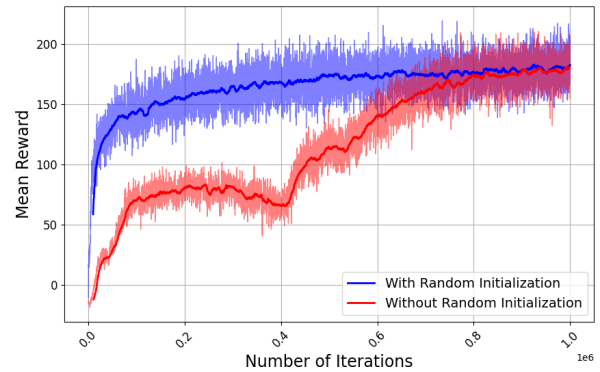


Fig. 5. Comparison of mean rewards in the ablation study on multi-stage random initialization. The red curve represents the policy with the proposed multi-stage random initialization, while the blue curve denotes the baseline without this mechanism. The policy with initialization (blue) demonstrates smoother convergence and higher stability throughout training. In contrast, the baseline (red) exhibits slower reward growth during the initial phase and greater instability.

TABLE III

SUCCESS RATE OF OUR WHOLE-BODY MANIPULATION POLICY FOR OBJECTS WITH DIFFERENT SIZES, MASSES, AND SHAPES IN MUJoCo.

Object	Size [cm ³]	Mass [kg]	Success rate [%]	NSDF
Cylinder	$\Phi 42 \times 40$	3	0	w/o
Cuboid	$42 \times 42 \times 42$	3	0	w/o
Sphere	$\Phi 42$	3	0	w/o
Cylinder	$\Phi 42 \times 40$	1	100	w
Cylinder	$\Phi 42 \times 40$	3	100	w
Cylinder	$\Phi 42 \times 40$	7	93	w
Cylinder	$\Phi 50 \times 40$	3	100	w
Cuboid	$42 \times 42 \times 42$	1	100	w
Cuboid	$42 \times 42 \times 42$	3	100	w
Cuboid	$42 \times 42 \times 42$	7	80	w
Cuboid	$30 \times 30 \times 30$	3	100	w
Sphere	$\Phi 42$	1	100	w
Sphere	$\Phi 42$	3	100	w
Sphere	$\Phi 42$	7	87	w
Sphere	$\Phi 30$	3	100	w

increased contact points with the object and significantly enhances manipulation stability, particularly when handling bulky items. In contrast, Fig. 4(b) depicts the trained policy without the NSDF module. During task execution, only the torso moves close to the object, while the arms fail to embrace properly. This leads to inferior policy performance and an inability to complete the transportation task.

2) *Evaluation of Randomized Initialization:* The proposed randomized initialization strategy mainly enhances state space exploration by resetting the robot to varied initial states at the beginning of each episode. To evaluate its effectiveness, we perform the comparison against a baseline that starts each episode deterministically from a fixed initial state. As shown in Fig. 5, the proposed method (blue curve) demonstrates significantly improved convergence speed and stability compared to the fixed initialization baseline (red curve). Specifically, our method promotes balanced learning across all task phases, effectively overcoming the limitations of slow convergence and unstable reward progression observed in the baseline.




Fig. 6. Illustration of sim-to-sim experiments in Mujoco. The illustration compares the manipulation performance across three common object primitives: a cylinder, a cuboid, and a sphere. Objects are varied in both shape and weight to evaluate generalization. The red objects match the dimensions used during training, while those in blue represent unseen shapes for generalization testing. The green platforms serve as pedestals for initial object placement.

B. Adaptability to Different Object Properties

To comprehensively evaluate the robustness and generalization of the policy, we perform a sim-to-sim transfer assessment in MuJoCo, which enables standardized and reproducible testing across a wide range of physical scenarios, as illustrated in Fig. 6. In this environment, the policy is tested on objects with diverse sizes, masses, and geometries to evaluate its success rate in WBM tasks, even though it was trained solely on a fixed-size cylindrical object. All target objects are placed on the green platforms initially. For each trial, success is defined as the object’s center of mass remaining within 0.1 m of the target location without falling. To ensure statistical reliability, the success rates are computed over 30 independent trials for each test condition.

The results summarized in Table III reveal several key aspects of the policy’s generalization capability. First, when the NSDF module is incorporated, the policy achieves and maintains a 100% success rate across multiple object types, including cylinders, cuboids, and spheres, under standard mass (3 kg) and size conditions. This demonstrates strong sim-to-sim transfer performance. In contrast, the policy exhibits complete failure (0% success) across all shape categories without NSDF, highlighting the essential role of this module in capturing geometric and semantic attributes for stable manipulation. Furthermore, the policy shows notable robustness to mass variations, retaining high success rates even at 7 kg, with 93% for cylinders, 87% for spheres, and 80% for cuboids. The slightly lower performance observed




Fig. 7. Sim-to-real transfer of whole-body manipulation in H1-2 humanoid robot. The sequence illustrates the sim-to-real transfer of a whole-body manipulation task. From 0 to 2 s, the H1 robot approaches the target object. At 4 s, it positions itself to initiate contact and prepare for lifting. By 6 s, the robot successfully lifts the object and begins locomotion to complete the task.

for heavier cuboidal objects can be attributed to their flat surface morphology. Additionally, variations in object size, such as larger objects like the $\Phi 50$ cm cylinder and smaller ones like the $\Phi 30$ cm sphere, do not degrade performance. This further confirms the policy’s adaptability beyond the training distribution.

C. Real-World Experiment

In our real-world experiment, we deploy the trained policy on the Unitree H1-2 platform equipped with an onboard Intel i7 computer that executes all policy inference computations with the trained policy running at 50 Hz, ensuring real-time control performance. The global poses of the robot and the embracing target are tracked in real time via a high-precision motion capture system. For the target object, we selected a cylinder with dimensions $\Phi 40\text{cm} \times 60\text{cm}$ to evaluate the policy’s performance when handling objects of realistic size and structural properties. As shown in Fig. 7, the experiment showcases the robot’s noteworthy capability in manipulation of a bulky object. The trial begins at 0 s, with the robot initiating motion toward the object. It successfully enters the pick-up zone at 4 s, then achieves a stable embracing of the object and moves into the transport phase by 6 s. From the real-world experiments, they showcase the successful embracing performance toward bulky objects, while using the whole body of the humanoid robot including arms and torso. In this sense, the effectiveness of our proposed method is clearly demonstrated.

V. CONCLUSIONS

In this study, we propose an RL framework that integrates a pre-trained human motion prior and an NSDF, enabling humanoid robots to perform whole-body embracing tasks for bulky objects. By constructing a hierarchical control architecture, the method effectively coordinates the movements of robot arms and torso, to distribute contact forces across multiple contact points, which significantly enhances operational stability and load capacity. Experimental validation demonstrates that the proposed approach not only adapts to objects of various shapes, sizes, and weights in simulation but also transfers effectively to a real-world robot platform toward WBM tasks with complex contacts.

REFERENCES

- [1] M. Murooka, K. Chappellet, A. Tanguy, M. Benallegue, I. Kumagai, M. Morisawa, F. Kanehiro, and A. Kheddar, "Humanoid locomotion pattern generation and stabilization control," *IEEE Robotics and Automation Letters*, vol. 6, no. 3, pp. 5597–5604, 2021.
- [2] D. J. Agravante, A. Cherubini, A. Sherikov, P.-B. Wieber, and A. Kheddar, "Human-humanoid collaborative carrying," *IEEE Transactions on Robotics*, vol. 35, no. 4, pp. 833–846, 2019.
- [3] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, "Humanplus: Humanoid shadowing and imitation from humans," in *8th Annual Conference on Robot Learning (CoRL)*, 2024.
- [4] Y. Li, Y. Zhang, A. Razmjoo, and S. Calinon, "Representing robot geometry as distance fields: Applications to whole-body manipulation," in *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 2024, pp. 15 351–15 357.
- [5] M. Florek-Jasińska, T. Wimböck, and C. Ott, "Humanoid compliant whole arm dexterous manipulation: control design and experiments," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2014, pp. 1616–1621.
- [6] F. Liu, Z. Gu, Y. Cai, Z. Zhou, H. Jung, J. Jang, S. Zhao, S. Ha, Y. Chen, D. Xu, and Y. Zhao, "Opt2skill: Imitating dynamically-feasible whole-body trajectories for versatile humanoid loco-manipulation," *arXiv:2409.20514*, 2025.
- [7] M. Murooka, S. Nozawa, Y. Kakiuchi, K. Okada, and M. Inaba, "Whole-body pushing manipulation with contact posture planning of large and heavy object for humanoid robot," in *2015 IEEE International Conference on Robotics and Automation (ICRA)*, 2015, pp. 5682–5689.
- [8] A. Goncalves, N. Kuppaswamy, A. Beaulieu, A. Uttamchandani, K. M. Tsui, and A. Alspach, "Punyo-1: Soft tactile-sensing upper-body robot for large object manipulation and physical human interaction," in *2022 IEEE 5th International Conference on Soft Robotics (RoboSoft)*, 2022, pp. 844–851.
- [9] W. Yuan, K. Hang, H. Song, D. Kragic, M. Y. Wang, and J. A. Stork, "Reinforcement learning in topology-based representation for human body movement with whole arm manipulation," in *2019 International Conference on Robotics and Automation (ICRA)*, 2019, pp. 2153–2160.
- [10] M. Zhang, J. Barreiros, and A. O. Onol, "Plan-guided reinforcement learning for whole-body manipulation," *arXiv:2310.12263*, 2023.
- [11] J. A. Barreiros, A. Özgün Onol, M. Zhang, S. Creasey, A. Goncalves, A. Beaulieu, A. Bhat, K. M. Tsui, and A. Alspach, "Learning contact-rich whole-body manipulation with example-guided reinforcement learning," *Science Robotics*, vol. 10, no. 105, p. eads6790, 2025.
- [12] Z. Gu, J. Li, W. Shen, W. Yu, Z. Xie, S. McCrory, X. Cheng, A. Shamsah, R. Griffin, C. K. Liu *et al.*, "Humanoid locomotion and manipulation: Current progress and challenges in control, planning, and learning," *arXiv preprint arXiv:2501.02116*, 2025.
- [13] P. Song, M. Yashima, and V. Kumar, "Dynamics and control of whole arm grasps," in *2001 IEEE International Conference on Robotics and Automation (ICRA)*, vol. 3, 2001, pp. 2229–2234.
- [14] Z. Zhuang, S. Yao, and H. Zhao, "Humanoid parkour learning," in *8th Annual Conference on Robot Learning (CoRL)*, 2024.
- [15] W. Xie, C. Bai, J. Shi, J. Yang, Y. Ge, W. Zhang, and X. Li, "Humanoid whole-body locomotion on narrow terrain via dynamic balance and reinforcement learning," *arXiv:2502.17219*, 2025.
- [16] X. B. Peng, P. Abbeel, S. Levine, and M. Van de Panne, "Deepmimic: Example-guided deep reinforcement learning of physics-based character skills," *ACM Transactions on Graphics*, vol. 37, no. 4, pp. 1–14, 2018.
- [17] Z. Luo, J. Cao, A. Winkler, K. Kitani, and W. Xu, "Perpetual humanoid control for real-time simulated avatars," in *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2023, pp. 10 861–10 870.
- [18] M. Ji, X. Peng, F. Liu, J. Li, G. Yang, X. Cheng, and X. Wang, "Exbody2: Advanced expressive humanoid whole-body control," in *RSS 2025 Workshop on Whole-body Control and Bimanual Manipulation: Applications in Humanoids and Beyond*, 2025.
- [19] T. He, Z. Luo, X. He, W. Xiao, C. Zhang, W. Zhang, K. M. Kitani, C. Liu, and G. Shi, "Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning," in *8th Annual Conference on Robot Learning (CoRL)*, 2024.
- [20] Z. Luo, J. Cao, J. Merel, A. Winkler, J. Huang, K. M. Kitani, and W. Xu, "Universal humanoid motion representations for physics-based control," in *The Twelfth International Conference on Learning Representations*, 2024.
- [21] Y. Shao, X. Huang, B. Zhang, Q. Liao, Y. Gao, Y. Chi, Z. Li, S. Shao, and K. Sreenath, "Langwbc: Language-directed humanoid whole-body control via end-to-end learning," in *Robotics: Science and Systems*, 2025.
- [22] Z. Zhang, C. Chen, H. Xue, J. Wang, S. Liang, Y. Liu, Z. Zhang, H. Wang, and L. Yi, "Unleashing humanoid reaching potential via real-world-ready skill space," *arXiv:2505.10918*, 2025.
- [23] C. Tessler, Y. Guo, O. Nabati, G. Chechik, and X. B. Peng, "Masked-mimic: Unified physics-based character control through masked motion inpainting," *ACM Transactions on Graphics*, vol. 43, no. 6, pp. 1–21, 2024.
- [24] T. He, Z. Luo, W. Xiao, C. Zhang, K. Kitani, C. Liu, and G. Shi, "Learning human-to-humanoid real-time whole-body teleoperation," in *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024, pp. 8944–8951.
- [25] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv:1707.06347*, 2017.
- [26] K. Chen, Z. Bi, G. Zhao, C. Zheng, Y. Li, H. Zhao, and J. Ma, "SAMP: Spatial anchor-based motion policy for collision-aware robotic manipulators," *arXiv: 2509.11185*, 2025.