

OmniRetarget: Interaction-Preserving Data Generation for Humanoid Whole-Body Loco-Manipulation and Scene Interaction

Lujie Yang^{*1,2}, Xiaoyu Huang^{*1,3}, Zhen Wu^{*1},
Angjoo Kanazawa^{†1,3}, Pieter Abbeel^{†1,3}, Carmelo Sferrazza^{†1}, C. Karen Liu^{†1,4}, Rocky Duan^{†1}, Guanya Shi^{†1,5}




Fig. 1: OMNIRETARGET enables reinforcement learning policies to learn complex, long-horizon loco-manipulation skills in challenging environments that transfer zero-shot from simulation to a Unitree G1 humanoid. Thanks to the high-quality interaction-preserving motion retargeting, these policies are trained and deployed in a *minimal and unified* way: it involves only 5 rewards, 4 robot domain randomization terms, and a purely proprioceptive observation space, shared by all tasks. Demonstrated behaviors include (a) 30-second parkour course involving chair moving, stepping & vault, and jump & roll, (b) object transportation, (c) crawling on a slope, and (d) fast platform climbing and sitting.

Abstract—A dominant paradigm for teaching humanoid robots complex skills is to retarget human motions as kinematic references to train reinforcement learning (RL) policies. However, existing retargeting pipelines often struggle with the significant embodiment gap between humans and robots, producing physically implausible artifacts like foot-skating and penetration. More importantly, common retargeting methods neglect the rich human-object and human-environment interactions essential for expressive locomotion and loco-manipulation. To address this, we introduce OMNIRETARGET, an interaction-preserving data generation engine based on an interaction mesh that explicitly models and preserves the crucial spatial and contact relationships between an agent, the terrain, and manipulated objects. By minimizing the Laplacian deformation between the human and robot meshes while enforcing kinematic constraints, OMNIRETARGET generates kinematically feasible trajectories. Moreover, preserving task-relevant interactions enables efficient data augmentation, from a single demonstration to different robot embodiments, terrains, and object configurations. We comprehensively evaluate OMNIRETARGET by retargeting motions from OMOMO [1], LAFAN1 [2], and our in-house MoCap datasets, generating over 8-hour trajectories that achieve better kinematic constraint satisfaction and contact preservation than widely used baselines. Such high-quality data enables proprioceptive RL policies to successfully execute long-horizon (up to 30 seconds) parkour and loco-manipulation skills on a Unitree G1 humanoid, trained with only 5 reward terms and simple domain randomization shared by all tasks, without

any learning curriculum. All code, retargeted datasets, and result videos can be found at <https://omniretarget.github.io>.

I. INTRODUCTION

The quest to enable humanoid robots to perform complex whole-body scene- and object-interaction tasks has long been constrained by a fundamental data bottleneck. While deep reinforcement learning (RL) has shown remarkable success in robot control, efficient exploration is highly sensitive to reward engineering [3]. This challenge is further amplified on humanoids, whose high-dimensional action spaces and complex dynamics make learning natural, expressive behaviors from scratch both difficult and inefficient.

To address these challenges, imitating human motions offers a powerful alternative for learning whole-body control, especially for complex scene interactions. Human demonstrations capture dynamic coordination, such as lifting objects while walking on uneven terrain, and have been used effectively in animation [4], [5], [6]. A critical challenge arises in robotics: unlike virtual characters, physical humanoids only approximate human morphology, with significant differences in shape, proportion and degrees of freedom. This embodiment gap means that simply adapting human motions is insufficient; it is essential to also adapt their scene interactions to the robot’s specific form to generate usable references.

To this end, researchers have pursued two main strategies. The first one is teleoperation [7], [8], [9], where only a

¹Amazon FAR (Frontier AI & Robotics), ²MIT, ³UC Berkeley, ⁴Stanford University, ⁵CMU. * Equal contribution, work done while interning at Amazon FAR. † Amazon FAR team co-lead.

human operator’s motions are retargeted to control the robot online. This approach leverages the human operator for real-time adaptation, which sidesteps the need for automatic interaction retargeting. However, despite the advantage of on-line feedback, the method remains labor-intensive and does not scale well for large-scale data generation. The second and more scalable strategy is offline interaction retargeting, which holistically adapts both the human’s motion and their scene interactions to the robot’s specific embodiment.

However, most existing retargeting methods [10], [9], [11] fall short in this regard. They predominantly rely on unconstrained or softly-penalized optimization, resulting in implausible motions with artifacts such as foot skating and penetration. More importantly, they do not explicitly consider interaction preservation—i.e., maintaining spatial and contact relationship—in the retargeting formulation, relying instead on simple keypoint matching. Consequently, the resulting references are of lower quality, which in turn complicates the downstream RL policy training [12], [8], [13].

In this work, we introduce OMNIRETARGET, an open-source data generation engine that transforms human demonstrations into diverse, high-quality kinematic references for humanoid whole-body control. By modeling spatial and contact relationships between robots, objects, and terrains via an interaction mesh [14], OMNIRETARGET preserves essential interactions and generates kinematically feasible variations. While existing methods require separate demonstrations for each variation—making data collection costly and limiting coverage—OMNIRETARGET addresses this bottleneck directly. Inspired by data augmentation frameworks for contact-rich manipulation [15], our framework automatically augments a single demonstration into a large number of training examples across object configurations, shapes, robot embodiments, and environments.

Our pipeline employs constrained optimization to enforce physical feasibility, including collision avoidance, joint limits, and foot contact stability, while minimizing interaction mesh deformation. The resulting motions are interaction-preserving and exhibit only minimal kinematic artifacts, providing dense learning signals that accelerate RL with minimal reward engineering. On a diverse suite of whole-body interaction tasks such as box lifting, platform climbing, and slope crawling, policies trained on OMNIRETARGET datasets outperform those from prior retargeting methods in both motion quality and robustness, with successful zero-shot sim-to-real transfer onto a physical humanoid robot.

Our contributions are fourfold:

- 1) The first interaction-preserving humanoid retargeting framework that handles rich robot-object-terrain interactions while enforcing hard physical constraints.
- 2) A systematic data augmentation pipeline that transforms a single human demonstration into a diverse, large-scale set of high-quality kinematic trajectories on various robot embodiments.
- 3) A large-scale, open-source dataset of retargeted, kinematically-feasible loco-manipulation trajectories.
- 4) Successful zero-shot sim-to-real transfer of proprioceptive RL policies on a physical humanoid, demonstrat-

ing a diverse set of scene-interaction tasks, including a long, agile sequence of object carrying, platform climbing, jumping, rolling and wall-flipping.

II. RELATED WORKS

A. Motion Retargeting

In computer graphics, transferring motions across characters has been extensively explored. Researchers have employed optimization-based methods to retarget human motions to avatars by preserving distances and orientations between keypoints [16], minimizing deformation energy [14], [17], or scaling the motions to satisfy hard constraints [18]. Others leverage data-driven methods that map diverse skeletons to a canonical representation [19], solve inverse kinematics with neural networks [20], or use reinforcement learning to preserve an interaction graph [21].

Retargeting motions to humanoid robots introduces challenges beyond character animation, particularly the need to enforce physical constraints. For example, PHC [10], a graphics method adopted in robotics [13], [8], uses keypoint matching with unconstrained optimization, often leading to penetration, foot skating, and lack of object or scene awareness. Similarly, GMR [9] extends keypoint matching to orientations but suffers the same issues. VideoMimic [11] improves realism with soft contact and collision penalties but offers no guarantees and requires careful tuning.

The closest method to ours is Interaction Mesh based Motion Adaptation (IMMA) [22], which also leverages an interaction mesh [14] to preserve the spatial relationship between body parts. However, it is not open-sourced and ignores kinematic limits and interactions with the environment or manipulated objects. In contrast, OMNIRETARGET unifies all hard constraints, including foot sticking, non-penetration, and joint and velocity limits, while explicitly preserving environment and object interactions.

B. Learning-Based Humanoid Whole-Body Control

Recent learning-based whole-body control has enabled humanoids to traverse dynamic scenes and manipulate objects [23], [24], [25], [26], [27], [28], [29]. These methods typically train with hand-crafted rewards or task interfaces (e.g., velocity tracking, contact schedules, end-effector targets) but depend on extensive reward engineering and mostly fail to yield natural, human-level motions.

Motion imitation offers a promising alternative. In graphics, DeepMimic [4] shows that using human references yields natural, human-like behaviors with agile, dynamic motions. However, applying this approach to humanoid robots remains difficult due to the lack of reliable open-source kinematic retargeting pipelines. With suboptimal reference motions, practitioners are forced to either manually clean the data [12] or re-introduce extensive reward engineering, such as ad-hoc regularizers for contact, slipping, and air time, to compensate for artifacts [9], [13], [30]. In contrast, trackers with minimal reward formulation like BeyondMimic [31] achieve state-of-the-art results on hardware with high-fidelity references [32], but those are scarce and robot-only, without interactions.

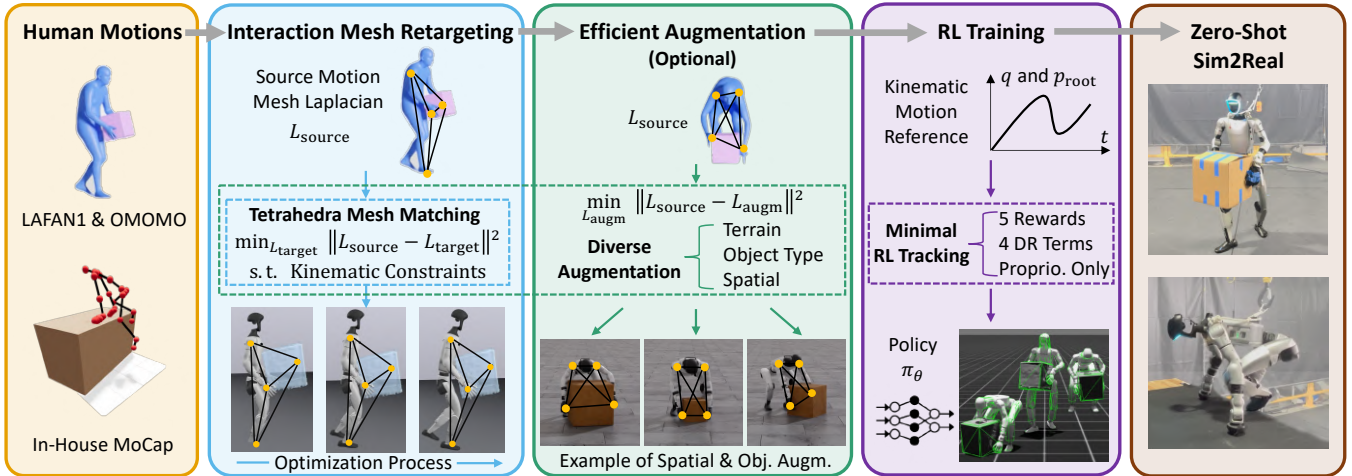


Fig. 2: **OMNIRETARGET overview.** Human demonstrations are retargeted to the robot via interaction-mesh-based constrained optimization. Each spatial and shape augmentation is solved as a new optimization, producing diverse trajectories that serve as references for RL training with minimal reward design and domain randomization, enabling zero-shot transfer to real-world humanoids.

Beyond single-character motion, human–scene interaction data has proven effective for terrain traversal and loco-manipulation in character animation [5], [6], but translating this to robotics remains challenging. VideoMimic [11] applies this idea to human–terrain traversal by reconstructing motions and terrains from video, but suffers from artifacts and is limited to static–scene interactions. To bridge this gap, OMNIRETARGET enables natural, agile robot-object-scene interactions with high-quality reference from retargeting without manual post-processing or reward engineering.

C. Data Generation for Humanoid Loco-Manipulation

The demand for whole-body interaction data has motivated many prior works on data generation. One approach is direct human teleoperation [7], [8], [9], [33]. While it provides online feedback, teleoperation is difficult to scale: it’s labor-intensive, prone to operator fatigue, and limited by the embodiment gap between human and robot kinematics. The lack of rich haptic feedback and difficulty stabilizing extreme motions (e.g., deep squats) further constrain its applicability. To address these scaling challenges, automated data augmentation has been explored, particularly for robotic manipulation. Many works leverage state-of-the-art generative models for visual [34], [35] and semantic [36], [37], [38] augmentations, while others rely on simple open-loop kinematic replay of base trajectories [39], [40] or trajectory optimization [15] in simulation. Despite the advances in manipulation, data augmentation for whole-body loco-manipulation remains largely unexplored. The closest prior work [41] interpolates keypoints to augment objects of different shape, but it cannot deal with varied object poses either. OMNIRETARGET directly addresses this gap.

III. INTERACTION-PRESERVING MOTION RETARGETING

A. Interaction Mesh with Hard Constraints

We leverage the interaction mesh [14] to preserve spatial relationships between body parts, objects, and the environment. The interaction mesh is defined as a volumetric

structure whose vertices consist of key robot or human joints together with points sampled from objects and the environment. By shrinking or stretching this mesh, we can warp human motion onto the robot while preserving relative spatial configurations and contact relationships.

Interaction Mesh Construction. We construct the interaction mesh by applying Delaunay tetrahedralization [42] to user-defined key joint positions and randomly sampled object and environment points. To more accurately maintain contact relationships, we sample the object and environment surfaces more densely than the body joints.

Optimization Objectives and Constraints. To preserve the spatial relationships between the body parts, objects and terrains, our primary objective is to minimize the Laplacian deformation energy of the interaction meshes [43], [44] constructed from two corresponding sets of keypoints. The source set at frame t , $\mathcal{P}_t^{\text{source}}$, is composed of user-defined anatomical points on the human, and points randomly sampled on the manipulated object and the environment. The target set for the retargeted motion, $\mathcal{P}_t^{\text{target}}$, consists of corresponding anatomical points on the robot, the same manipulated object and environment points. Our method is relatively robust to the precise placement of these keypoints, requiring only a semantically consistent correspondence between the human and robot (e.g., hand to hand).

The Laplacian coordinate of the i -th keypoint $p_{t,i} \in \mathcal{P}_t$ is defined as the difference between the point and the weighted average of its neighbors $j \in \mathcal{N}(i)$:

$$L(p_{t,i}) = p_{t,i} - \sum_{j \in \mathcal{N}(i)} w_{ij} \cdot p_{t,j}, \quad (1)$$

where w_{ij} is the normalized weight and we omit L ’s dependence on $\{p_{t,j}\}_{j \neq i}$ in the function definition for conciseness. For all our experiments, we use uniform weights, setting $w_{ij} = 1/|\mathcal{N}(i)|$. The deformation energy measures the change in these Laplacian coordinates between the source

demonstration mesh $\mathcal{P}_t^{\text{source}}$ and the retargeted mesh $\mathcal{P}_t^{\text{target}}$:

$$E_L = \sum_{p_{t,i}^{\text{source}} \in \mathcal{P}_t^{\text{source}}, p_{t,i}^{\text{target}} \in \mathcal{P}_t^{\text{target}}} \|L(p_{t,i}^{\text{source}}) - L(p_{t,i}^{\text{target}})\|^2. \quad (2)$$

We seek the robot configuration q_t , consisting of the floating base pose (quaternion and translation) and all joint angles, that minimizes this deformation energy while satisfying a set of hard kinematic constraints. The robot’s keypoints are determined by its configuration q_t via forward kinematics f_i as $p_{t,i}^{\text{robot}}(q_t) = f_i(q_t) \in \mathcal{P}_t^{\text{target}}$. At each time step, we solve the following constrained, nonconvex program:

$$q_t^* = \arg \min_{q_t} \sum_i \|L(p_{t,i}^{\text{source}}) - L(p_{t,i}^{\text{target}}(q_t))\|^2 + \|q_t - q_{t-1}\|_Q^2 \quad (3a)$$

$$\text{s.t. } \phi_j(q_t) \geq 0, \forall j \quad (3b)$$

$$q_{\min} \leq q_t \leq q_{\max} \quad (3c)$$

$$v_{\min} \cdot dt \leq q_t - q_{t-1} \leq v_{\max} \cdot dt \quad (3d)$$

$$p_t^F = p_{t-1}^F, \forall \text{stance foot}, \quad (3e)$$

where Q is a cost matrix that encourages temporal smoothness, ϕ_j denotes the signed distance function for the j -th collision pair, q_{\min}/q_{\max} and v_{\min}/v_{\max} are the configuration and velocity bounds, p_t^F denotes the foot position. A foot is considered to be in the stance phase if its horizontal velocity in the source motion (in the xy -plane) falls below a threshold of 1 cm/s. This optimization program solves for a temporally consistent robot trajectory that minimizes interaction mesh deformation, subject to hard constraints for collision avoidance (3b), joint and velocity limits (3c)–(3d), and preventing foot skating (3e).

We solve (3) sequentially for each timestep using a customized Sequential Quadratic Programming (SQP)-style solver. Within each iteration, the objective (3a) is quadratically approximated and the hard constraints (3b)–(3e) are linearized around the solution from the previous iteration. To ensure temporal consistency and accelerate convergence, the optimization at frame t is warm-started with the optimal solution from the previous frame q_{t-1}^* . A key challenge in this formulation is computing derivatives involving the quaternion-based floating base orientation; our implementation leverages the automatic differentiation framework in Drake [45], which correctly handles the differential geometry of rotations on the \mathbb{S}^3 manifold [46].

Our interaction-mesh-based kinematic pipeline is highly general. It adapts to different robot embodiments, including the Unitree G1, H1, and Booster T1 (Fig. 3), by modifying only the keypoint correspondences in the interaction mesh and the robot’s collision model. It also supports diverse interaction types: **robot-object** interactions from the OMOMO [1], **robot-terrain** interactions from in-house MoCap data, and **robot-only** motions on flat terrain from LAFANI [2].

B. Terrain, Object Shape and Spatial Augmentation

A key advantage of our framework is its capability for systematic data augmentation, which eliminates the need for collecting numerous, repetitive demonstrations with minor spatial variations. Our method can transform a single human

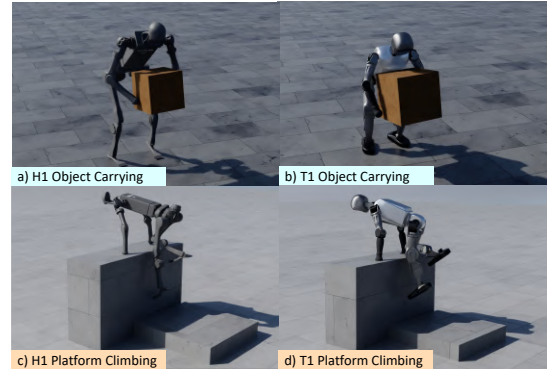


Fig. 3: Cross-embodiment robot-object-terrain interaction.

demonstration into a rich and diverse dataset by parametrically altering object configurations, shapes, or terrain features. For each new scenario, we re-solve the optimization problem with fixed $\mathcal{P}_t^{\text{source}}$ and augmented \mathcal{P}_t : minimizing the interaction mesh deformation finds a new, kinematically valid robot motion $\{q_t\}$ that preserves the essential spatial and contact relationships of the original interaction.

Robot-Object. We generate diverse interactions by augmenting both the object’s spatial configuration and its shape. We apply translations and rotations to modify the object’s initial pose (Fig. 4b) and blend the new initial pose with the original object motion via interpolation with an exponential schedule detailed in (14). In addition, we scale the three dimensions of the object (Fig. 4c). A critical component of this process is constructing the interaction mesh in the object’s local frame, which ensures that the robot’s interacting body parts naturally follow the object’s transformation (Sec. VI-C.2).

However, this alone can lead to trivial augmentations where the entire robot undergoes a rigid transformation along with the object. To generate more meaningful data diversity, we introduce cost terms and constraints that anchor parts of the robot’s body to the nominal trajectory $\{\bar{q}_t^*\}$. For instance, in a pick-up task, we encourage the robot to discover new upper-body coordination by penalizing lower-body deviations from the original motion:

$$\|q_t - \bar{q}_t^*\|_W, \quad (4)$$

where W heavily penalizes the lower-body entries, constraining the initial foot poses to match the nominal trajectory

$$p_0^F = \bar{p}_0^{F*} \quad \text{for left and right feet.} \quad (5)$$

These added objectives prevent the optimization from collapsing to a simple rigid transform of the nominal trajectory and instead produce genuinely new and diverse interactions.

Robot-Terrain. We generate diverse terrain scenarios by scaling environmental features, such as varying the platform height and depth (Fig. 4a), and introducing additional constraints. For instance, to encourage stable ground contact when the terrain is elevated, we uniformly sample a grid of points on the ground surface into the interaction mesh.

IV. RL TRAINING WITH MINIMAL FORMULATION

Having established our method for generating high-quality kinematic references, we use RL to bridge the gap to

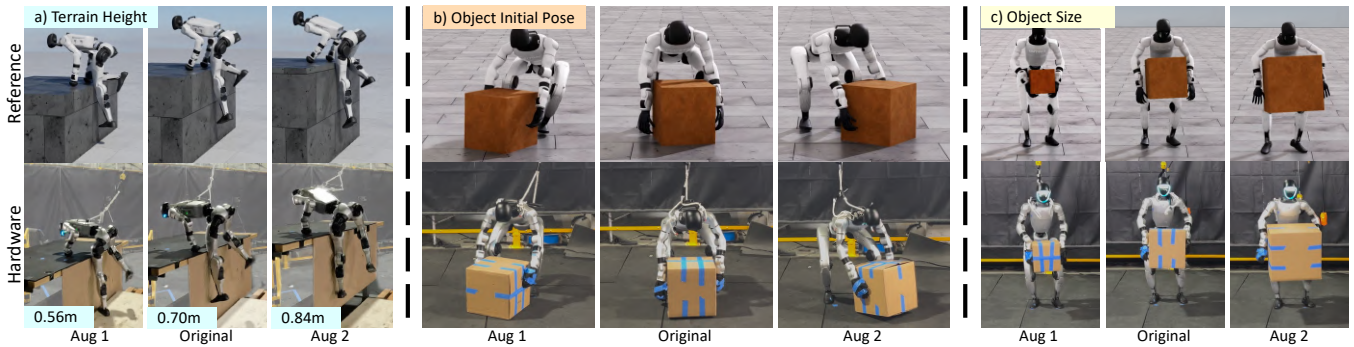


Fig. 4: OMNIRETARGET generates systematic variations of (a) terrain height, (b) object initial pose, and (c) object shape from a single human demonstration, with optimized motions in simulation (top) transferring consistently to hardware (bottom).

dynamics by training a low-level policy that converts these trajectories into physically realizable actions, enabling zero-shot transfer from simulation to hardware.

Reward engineering is often the main difficulty in humanoid RL: prior works [9], [13], [30] rely on many ad-hoc regularizers (e.g., foot flight and contact time) to compensate for artifacts in noisy references, but tuning these terms is tedious and fragile. In contrast, BeyondMimic [31] shows that when references are clean [32], a minimal reward is already sufficient for high-quality tracking. Since OmniRetarget produces such artifact-free, interaction-preserving references, we can follow this minimal formulation directly, achieving faithful tracking of dynamic interactions and zero-shot sim-to-real transfer *without any hyperparameter tuning*.

Observations. To show that high-quality reference motions provide a sufficient prior for complex tasks, we design a *minimal proprioceptive* observation space, as listed below, where the agent is blind to explicit scene and object information and must follow the reference trajectory precisely.

- *Reference Motion:* Reference Joint Position/Velocity, Reference Pelvis Position/Orientation Error;
- *Proprioception:* Joint Position/Velocity;
- *Previous Action:* Policy action from last timestep.

Rewards. To show the benefits of high-quality reference and avoid reward tuning, we use only five reward terms:

- *Body Tracking:* DeepMimic-style tracking term for body position, orientation, linear and angular velocity;
- *Object Tracking (where applicable):* DeepMimic-style tracking term for object position and orientation;
- *Action Rate:* Penalize rapid changes in action;
- *Soft Joint Limit:* Penalize robot joint limit violation;
- *Self-Collision:* Binary penalty on each body if its self-collision force exceeds 1 N.

We use the same weights and hyperparameters from [31] out of the box without tuning. For object tracking, we use the same hyperparameters as body tracking terms.

Termination. We terminate training episodes with large body tracking deviations [31]. For object loco-manipulation, episodes terminate when the object deviates more than 1.0m and 45° from the reference trajectory. We only apply this criterion after the policy achieves reasonable body tracking.

Domain Randomization. To improve generalization across object properties for a single reference motion, we randomize the object’s physical parameters: mass (0.1–2 kg),

center of mass (± 0.08 m), inertia (50–150%), and shape ($\pm 10\%$). For the robot specifically, in contrast to the many terms in prior works (e.g., random force injection (RFI), motor PD, action delay), we only apply four terms:

- *Torso COM Position:* ± 0.025 m in x , ± 0.05 m in y , ± 0.075 m in z ;
- *Joint default position:* ± 0.01 rad;
- *Random push:* 0.3 m/s, 0.78 rad/s for (1–3) s;
- *Observation noise:* ± 0.05 for orientation in Rot6D, ± 0.5 m/s and ± 0.2 rad/s for linear and angular velocity, ± 0.01 rad and ± 0.5 rad/s for joint position and velocity.

Policy Training. We group similar motions for faster training. All box-moving motions share a single multi-task policy, while platform climbing uses one policy per reference.

V. EXPERIMENTAL RESULTS

In this section, we present a comprehensive experimental validation of OMNIRETARGET. We first demonstrate the breadth of complex behaviors enabled by our approach, including natural object manipulation and terrain interaction. We then provide a quantitative benchmark against state-of-the-art baselines, evaluating performance across kinematic quality metrics and downstream policy performance.

A. Whole-Body Scene-Object-Interaction

a) Agile Loco-Manipulation: OMNIRETARGET enables RL policies to learn agile, whole-body motions for complex scene interactions and loco-manipulation in simulation, culminating in successful zero-shot sim-to-real transfer to hardware. Shown in Fig. 5, policies trained on our data reproduce a diverse range of expressive behaviors on a Unitree G1 humanoid, including natural box-carrying motions retargeted from the OMOMO dataset, dynamically climbing a 0.9m-high platform (70% of the robot’s height), and crawling over slopes, showing clean and accurate contact sequences.

To showcase the full capabilities of our framework, we present a long-horizon, dynamic sequence inspired by the Boston Dynamics Atlas tool-use demo [47]. Visualized in Fig. 1, our retargeted data enables the robot to carry a 4.6 kg chair to a platform, use it as a stepstone to climb up, and then leap off, performing a parkour-style roll to absorb the landing impact. This 30-second, complex, multi-stage task highlights OMNIRETARGET’s ability to produce precise and versatile




Fig. 5: Additional hardware results showing diverse, agile and human-like behaviors.



Fig. 6: Hardware results showing a high-dynamic wall-flip motion. The robot reaches a maximum linear velocity of 3.5 m/s and a peak angular velocity of 15 rad/s.

reference motions, pushing the boundaries of what is possible for humanoids learning agile, human-like behaviors.

We additionally showcase a high-dynamic wall-flip motion in Fig. 6. The robot completes the full flip in approximately 0.5 second, reaching a peak angular velocity of 15 rad/s. The trained policy is robust and achieves a 5/5 success rate in our real-world experiments.

b) Sim-to-real with Augmented Data: We show that the augmented motions from our pipeline can be used for training and deployment effectively. As shown in Fig. 4, the interaction mesh formulation allows OMNIRETARGET to generalize a single nominal motion into box-picking across shapes and positions, as well as platform climbing at different heights. Notably, these augmented motions transfer to hardware without reward tuning, effectively expanding the repertoire of scenes and behaviors we can achieve in real.

In comparison, relying solely on domain randomization—which perturbs object shapes and poses only during training—performs poorly under our RL formulation, as the policies struggle to explore far beyond the nominal reference. Policies trained on our augmented data instead yield reliable success (see video for comparison). Admittedly, additional reward engineering could help, but it contradicts our minimal design goal. Quantitatively, training and evaluating on the full augmented dataset achieves a 79.1% success rate, comparable to 82.2% when evaluating on nominal motions only, showing

that kinematics augmentation substantially enlarges coverage without significant performance degradation.

B. Benchmark Against Prior Retargeting Pipelines

We compare OMNIRETARGET against several widely-used open-source retargeting baselines: PHC [10], GMR [9] and VideoMimic [11]. The generated dataset including 2.78 hours of box carrying in OMOMO, 1 hour of in-house MoCap and 4.6 hours of LAFAN1 will be open-sourced.

a) Kinematics Quality: We evaluate the kinematic quality of retargeted motions on a Unitree G1 with three criteria:

- 1) **Penetration:** Measured by the time duration (normalized by the trajectory length) and maximum depth of intersections between the robot, objects, and terrain.
- 2) **Foot Skating:** Quantified by the time duration (normalized by the total desired foot sticking length) and maximum skating velocity of a stance foot.
- 3) **Contact Preservation:** Measured as the fraction of time during which the distance between relevant geometry pairs remains below a specified threshold, normalized by the desired contact duration. For *robot-object* tasks, we measure hand-object contact. For *robot-terrain* tasks, we measure contact between the robot’s hands, toes, and heels with the terrain surface.

As illustrated in Tab. I, OMNIRETARGET significantly outperforms all baselines across most kinematic metrics. While

Method	Penetration		Foot Skating		Contact Preservation	Downstream RL Policy
	Duration ↓	Max Depth (cm) ↓	Duration ↓	Max Vel. (cm/s) ↓	Duration ↑	Success Rate ↑
<i>Robot-Object Interaction (Retargeting from the OMOMO Dataset)</i>						
PHC [10]	0.68 ± 0.21	5.11 ± 3.09	0.05 ± 0.05	1.40 ± 0.80	0.96 ± 0.09	71.28% ± 22.55%
GMR [9]	0.83 ± 0.14	8.50 ± 3.94	0.02 ± 0.01	1.46 ± 0.45	0.99 ± 0.04	50.83% ± 23.89%
VideoMimic [11]	0.60 ± 0.27	7.48 ± 4.95	0.12 ± 0.07	1.50 ± 0.70	0.77 ± 0.25	3.85% ± 8.41%
OMNIRETARGET	0.00 ± 0.01	1.34 ± 0.34	0	0	0.96 ± 0.09	82.20% ± 9.74%
<i>Robot-Terrain Interaction (Retargeting from the In-House MoCap Dataset)</i>						
PHC	0.66 ± 0.36	7.74 ± 4.53	0.15 ± 0.04	2.03 ± 1.83	0.45 ± 0.28	52.63% ± 49.93%
GMR	0.91 ± 0.16	5.72 ± 3.84	0.04 ± 0.05	1.75 ± 3.01	0.67 ± 0.26	78.94% ± 40.77%
VideoMimic	0.83 ± 0.11	5.97 ± 3.58	0.14 ± 0.05	1.85 ± 1.38	0.47 ± 0.25	51.75% ± 49.23%
OMNIRETARGET	0.01 ± 0.02	1.37 ± 0.18	0	0	0.72 ± 0.19	94.73% ± 22.33%
<i>Robot-Only (Retargeting from the LAFANI Dataset)</i>						
Unitree [32]	0.09 ± 0.13	3.22 ± 2.64	0.06 ± 0.03	1.46 ± 0.01	N/A	100%
OMNIRETARGET	0.00 ± 0.00	1.07 ± 0.00	0	0	N/A	100%

TABLE I: Quantitative comparison of kinematic retargeting quality and downstream RL performances.

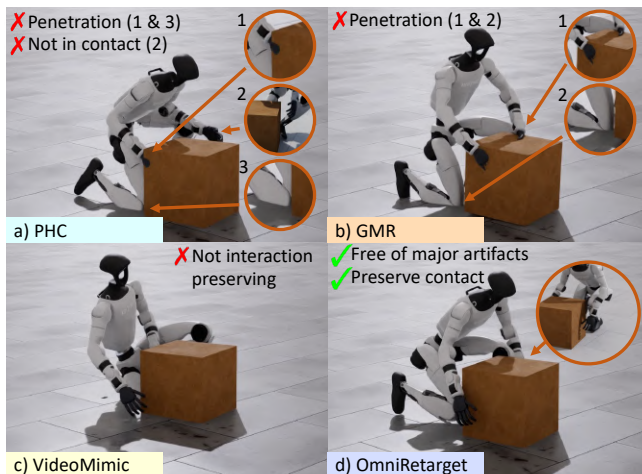


Fig. 7: Artifacts resulting from the retargeting baselines.

OMNIRETARGET occasionally incurs minor penetration due to the linearization of constraints (3b) in the sequential SOCP solver, the violations are minimal and can be efficiently fixed by RL. GMR achieves the highest contact preservation score for robot-object interaction tasks; however, this outcome largely reflects its keypoint-matching objective. In practice, scaling human hand keypoints to the robot’s size often drives the robot’s hands inside the object, leading to substantial penetration errors (Fig. 7b). Overall, all baselines exhibit significant penetration and foot skating (Fig. 7), degrading the downstream RL performance, as discussed next.

For a direct comparison on pure locomotion, we retarget motions from the LAFANI MoCap dataset [2] and benchmark them against the publicly available Unitree LAFANI retargeted dataset [32]. This serves as a strong baseline, as it is widely considered a high-quality data source for RL-based locomotion training [31]. Table I shows that OMNIRETARGET’s motions exhibit fewer physical artifacts, achieving better satisfaction of hard constraints.

b) Downstream RL Performance: A central observation from prior works [31], [12] is that the quality of retargeted motions strongly influences the performance of downstream RL. To verify this, we select 39 challenging motions for

OMNIRETARGET and baselines, and train RL policies using identical hyperparameters from [31] without manual tuning. We evaluate the policies in simulation, and success is measured by training termination criteria.

Shown in Tab. I, retargeting quality directly impacts RL success rates. OMNIRETARGET consistently achieves the highest performance across tasks, exceeding baselines by over 10% with lower variance, which indicates more stable learning across different motions. PHC performs better than GMR in object manipulation, likely due to lower penetration with sufficient contact preservation, but worse in terrain interaction, where its contact preservation drops by nearly 50%. Specifically for terrain interaction, we see that contact preservation is directly proportional to the success rate. These results suggest that both contact preservation and penetration reduction are critical for generalizing RL policies across diverse tasks, and OMNIRETARGET shows strength in both.

VideoMimic shows the weakest interaction preservation among all baselines (Fig. 7c), likely due to its collision avoidance soft cost conflicting with the keypoint matching cost. This is compounded by its coarse collision model originally designed for heightmaps, which is ill-suited for precise loco-manipulation. Consequently, while its terrain-interaction results are comparable to PHC, its performance on object manipulation is poor. Although this could be partially attributed to the tuning of its soft penalties, OMNIRETARGET demonstrates that a hard-constraint formulation avoids such sensitivities altogether.

VI. CONCLUSION

In this work, we tackled a key data bottleneck caused by a lack of high-quality, interaction-aware retargeting pipeline in humanoid whole-body loco-manipulation. We introduced OMNIRETARGET, a unified, interaction-preserving data generation engine that leverages an interaction mesh within a single constrained optimization. Our experiments showed that OMNIRETARGET significantly outperforms prior methods in kinematic quality, producing a diverse set of artifact-free trajectories from single demonstrations. This high-quality data enabled a proprioceptive RL policy, trained with

minimal formulation, to achieve long-horizon dynamic skills on a physical humanoid via zero-shot sim-to-real transfer.

Ultimately, OMNIRETARGET demonstrates a paradigm shift from patching lower-quality reference motions with complex reward engineering to solving the problem at its source with principled data generation. While our current frame-by-frame optimization is mostly effective, future work could explore jointly optimizing the entire trajectory to enhance the framework’s robustness to noisier motion sources, such as video data, or learning autonomous visuomotor policies. By open-sourcing our complete framework and the large-scale dataset of retargeted trajectories, we hope to accelerate progress towards more agile, capable, and versatile humanoid robots.

REFERENCES

- [1] J. Li, J. Wu, and C. K. Liu, “Object motion guided human motion synthesis,” *ACM Transactions on Graphics (TOG)*, 2023.
- [2] F. G. Harvey, M. Yurick, D. Nowrouzezahrai, and C. Pal, “Robust motion in-betweening,” vol. 39, no. 4, 2020.
- [3] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Science robotics*, 2020.
- [4] 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 (TOG)*, 2018.
- [5] M. Xu, Y. Shi, K. Yin, and X. B. Peng, “Parc: Physics-based augmentation with reinforcement learning for character controllers,” in *Proceedings of the SIGGRAPH Conference Papers*, 2025.
- [6] Z. Wu, J. Li, P. Xu, and C. K. Liu, “Human-object interaction from human-level instructions,” *arXiv preprint arXiv:2406.17840*, 2024.
- [7] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, “Humanplus: Humanoid shadowing and imitation from humans,” *CoRL*, 2024.
- [8] 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 *CoRL*, 2025.
- [9] Y. Ze, Z. Chen, J. P. AraÅsjo, Z.-a. Cao, X. B. Peng, J. Wu, and C. K. Liu, “Twist: Teleoperated whole-body imitation system,” *CoRL*, 2025.
- [10] Z. Luo, J. Cao, A. W. Winkler, K. Kitani, and W. Xu, “Perpetual humanoid control for real-time simulated avatars,” in *ICCV*, 2023.
- [11] A. Allshire, H. Choi, J. Zhang, D. McAllister, A. Zhang, C. M. Kim, T. Darrell, P. Abbeel, J. Malik, and A. Kanazawa, “Visual imitation enables contextual humanoid control,” *CoRL*, 2025.
- [12] T. Zhang, B. Zheng, R. Nai, Y. Hu, Y.-J. Wang, G. Chen, F. Lin, J. Li, C. Hong, K. Sreenath *et al.*, “Hub: Learning extreme humanoid balance,” *CoRL*, 2025.
- [13] T. He, J. Gao, W. Xiao, Y. Zhang, Z. Wang, J. Wang, Z. Luo, G. He, N. Sobanbabu, C. Pan, Z. Yi, G. Qu, K. Kitani, J. Hodgins, L. J. Fan, Y. Zhu, C. Liu, and G. Shi, “Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills,” *arXiv*, 2025.
- [14] E. S. L. Ho, T. Komura, and C.-L. Tai, “Spatial relationship preserving character motion adaptation,” *ACM Transactions on Graphics*, 2010.
- [15] L. Yang, H. Suh, T. Zhao, B. P. Graesdal, T. Kelestemur, J. Wang, T. Pang, and R. Tedrake, “Physics-driven data generation for contact-rich manipulation via trajectory optimization,” *RSS*, 2025.
- [16] T. Cheynel, T. Rossi, B. Bellot-Gurlet, D. Rohmer, and M.-P. Cani, “Sparse motion semantics for contact-aware retargeting,” in *ACM SIGGRAPH Conference on Motion, Interaction and Games*, 2023.
- [17] Y. Kim, H. Park, S. Bang, and S.-H. Lee, “Retargeting human-object interaction to virtual avatars,” *IEEE transactions on visualization and computer graphics*, vol. 22, no. 11, pp. 2405–2412, 2016.
- [18] M. Gleicher, “Retargeting motion to new characters,” in *Proceedings of the 25th annual conference on Computer graphics and interactive techniques*, 1998, pp. 33–42.
- [19] K. Aberman, P. Li, D. Lischinski, O. Sorkine-Hornung, D. Cohen-Or, and B. Chen, “Skeleton-aware networks for deep motion retargeting,” *ACM Transactions on Graphics (TOG)*, vol. 39, no. 4, pp. 62–1, 2020.
- [20] R. Villegas, J. Yang, D. Ceylan, and H. Lee, “Neural kinematic networks for unsupervised motion retargeting,” in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- [21] Y. Zhang, D. Gopinath, Y. Ye, J. Hodgins, G. Turk, and J. Won, “Simulation and retargeting of complex multi-character interactions,” in *ACM SIGGRAPH 2023 Conference Proceedings*, 2023.
- [22] S. Nakaoka and T. Komura, “Interaction mesh based motion adaptation for biped humanoid robots,” in *Humanoids*, 2012.
- [23] J. Dao, H. Duan, and A. Fern, “Sim-to-real learning for humanoid box loco-manipulation,” in *ICRA*. IEEE, 2024.
- [24] J. Long, J. Ren, M. Shi, Z. Wang, T. Huang, P. Luo, and J. Pang, “Learning humanoid locomotion with perceptive internal model,” *arXiv*, 2024.
- [25] J. He, C. Zhang, F. Jenelten, R. Grandia, M. Bächer, and M. Hutter, “Attention-based map encoding for learning generalized legged locomotion,” *Science Robotics*, vol. 10, no. 105, p. eadv3604, 2025.
- [26] X. He, R. Dong, Z. Chen, and S. Gupta, “Learning getting-up policies for real-world humanoid robots,” *RSS*, 2025.
- [27] Y. Kuang, H. Geng, A. Elhafi, T.-D. Do, P. Abbeel, J. Malik, M. Pavone, and Y. Wang, “Skillblender: Towards versatile humanoid whole-body loco-manipulation via skill blending,” *arXiv*, 2025.
- [28] Y. Xue, W. Dong, M. Liu, W. Zhang, and J. Pang, “A unified and general humanoid whole-body controller for versatile locomotion,” *RSS*, 2025.
- [29] Y. Zhang, Y. Yuan, P. Gurunath, T. He, S. Omidshafiei, A.-a. Aghamohammadi, M. Vazquez-Chanlatte, L. Pedersen, and G. Shi, “Falcon: Learning force-adaptive humanoid loco-manipulation,” *arXiv*, 2025.
- [30] Z. Li, X. B. Peng, P. Abbeel, S. Levine, G. Berseth, and K. Sreenath, “Reinforcement learning for versatile, dynamic, and robust bipedal locomotion control,” *IJRR*, 2025.
- [31] Q. Liao, T. E. Truong, X. Huang, G. Tevet, K. Sreenath, and C. K. Liu, “Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion,” *arXiv e-prints*, pp. arXiv–2508, 2025.
- [32] U. Robotics and Contributors, “Unitree lafan1 retargeting dataset,” https://huggingface.co/datasets/lvhaidong/LAFANI_Retargeting-Dataset, 2025.
- [33] Q. Ben, F. Jia, J. Zeng, J. Dong, D. Lin, and J. Pang, “Homie: Humanoid loco-manipulation with isomorphic exoskeleton cockpit,” *RSS*, 2025.
- [34] X. Zhang, M. Chang, P. Kumar, and S. Gupta, “Diffusion meets dagger: Supercharging eye-in-hand imitation learning,” in *RSS*, 2024.
- [35] L. Y. Chen, C. Xu, K. Dharmarajan, Z. Irshad, R. Cheng, K. Keutzer, M. Tomizuka, Q. Vuong, and K. Goldberg, “Rovi-aug: Robot and viewpoint augmentation for cross-embodiment robot learning,” in *Conference on Robot Learning (CoRL)*, 2024.
- [36] Z. Mandi, H. Bharadhwaj, V. Moens, S. Song, A. Rajeswaran, and V. Kumar, “Cacti: A framework for scalable multi-task multi-scene visual imitation learning,” *arXiv preprint arXiv:2212.05711*, 2022.
- [37] Z. Chen, S. Kiami, A. Gupta, and V. Kumar, “Genaug: Retargeting behaviors to unseen situations via generative augmentation,” *RSS*, 2023.
- [38] T. Yu, T. Xiao, A. Stone, J. Tompson, A. Brohan, S. Wang, J. Singh, C. Tan, J. Peralta, B. Ichter *et al.*, “Scaling robot learning with semantically imagined experience,” *RSS*, 2023.
- [39] A. Mandlekar, S. Nasiriany, B. Wen, I. Akinola, Y. Narang, L. Fan, Y. Zhu, and D. Fox, “Mimicgen: A data generation system for scalable robot learning using human demonstrations,” in *CoRL*, 2023.
- [40] Z. Jiang, Y. Xie, K. Lin, Z. Xu, W. Wan, A. Mandlekar, L. Fan, and Y. Zhu, “Dexmimicgen: Automated data generation for bimanual dexterous manipulation via imitation learning,” *ICRA*, 2025.
- [41] S. Starke, H. Zhang, T. Komura, and J. Saito, “Neural state machine for character-scene interactions,” *ACM Transactions on Graphics*, 2019.
- [42] H. Si and K. Gärtner, “Meshing piecewise linear complexes by constrained delaunay tetrahedralizations,” in *Proceedings of the 14th international meshing roundtable*. Springer, 2005, pp. 147–163.
- [43] M. Alexa, “Differential coordinates for local mesh morphing and deformation,” *The Visual Computer*, vol. 19, no. 2, pp. 105–114, 2003.
- [44] K. Zhou, J. Huang, J. Snyder, X. Liu, H. Bao, B. Guo, and H.-Y. Shum, “Large mesh deformation using the volumetric graph laplacian,” in *ACM SIGGRAPH 2005 Papers*. ACM, 2005, pp. 496–503.
- [45] R. Tedrake and the Drake Development Team, “Drake: Model-based design and verification for robotics,” 2019.
- [46] B. E. Jackson, K. Tracy, and Z. Manchester, “Planning with attitude,” *IEEE Robotics and Automation Letters*, 2021.
- [47] Boston Dynamics, “Atlas Gets a Grip.” YouTube, available: https://www.youtube.com/watch?v=e-1_QhJIEhQ.
- [48] K. Zakka, “Mink: Python inverse kinematics based on MuJoCo,” May 2025. [Online]. Available: <https://github.com/kevinzakka/mink>
- [49] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, “SMPL: A skinned multi-person linear model,” *ACM Trans. Graphics (Proc. SIGGRAPH Asia)*, vol. 34, no. 6, pp. 248:1–248:16, Oct. 2015.