

# Make Tracking Easy: Neural Motion Retargeting for Humanoid Whole-body Control

Qingrui Zhao<sup>1</sup>, Kaiyue Yang<sup>1</sup>, Xiyu Wang<sup>1,2</sup>, Shiqi Zhao<sup>1</sup>, Yi Lu<sup>1</sup>, Xinfang Zhang<sup>2</sup>, Wei Yin<sup>3</sup>,  
Qiu Shen<sup>1</sup>, Xiao-Xiao Long<sup>1\*</sup>, Xun Cao<sup>1\*</sup>

\*Corresponding authors.

<sup>1</sup>Nanjing University, <sup>2</sup>Huawei Technologies, <sup>3</sup>Horizon Robotics

Project page: <https://3dv-humanoidgroup.github.io/nmr.github.io/>

*Abstract*—Humanoid robots require diverse motor skills to integrate into complex environments, but bridging the kinematic and dynamic embodiment gap from human data remains a major bottleneck. We demonstrate through Hessian analysis that traditional optimization-based retargeting is inherently non-convex and prone to local optima, leading to physical artifacts like joint jumps and self-penetration. To address this, we reformulate the targeting problem as learning data distribution rather than optimizing optimal solutions, where we propose NMR, a Neural Motion Retargeting framework that transforms static geometric mapping into a dynamics-aware learned process. We first propose Clustered-Expert Physics Refinement (CEPR), a hierarchical data pipeline that leverages VAE-based motion clustering to group heterogeneous movements into latent motifs. This strategy significantly reduces the computational overhead of massively parallel reinforcement learning experts, which project and repair noisy human demonstrations onto the robot’s feasible motion manifold. The resulting high-fidelity data supervises a non-autoregressive CNN-Transformer architecture that reasons over global temporal context to suppress reconstruction noise and bypass geometric traps. Experiments on the Unitree G1 humanoid across diverse dynamic tasks (e.g., martial arts, dancing) show that NMR eliminates joint jumps and significantly reduces self-collisions compared to state-of-the-art baselines. Furthermore, NMR-generated references accelerate the convergence of downstream whole-body control policies, establishing a scalable path for bridging the human-robot embodiment gap.

## I. INTRODUCTION

Humanoid robots are at a critical stage in their transition from laboratory settings to complex human environments, and the acquisition of diverse motor skills is fundamental to this progression. At present, the dominant research paradigm relies on large-scale human motion data, such as video recordings or motion-capture databases, as prior guidance for training robot motor control policies through imitation learning or reinforcement learning (RL)[1, 2, 3]. Within this pipeline, motion retargeting serves as a critical bridge between human demonstrations and robotic execution. Conventional retargeting methods, including inverse kinematics (IK)-based approaches and differential optimization schemes such as GMR[4], primarily seek optimal joint configurations at the geometric level.

However, this conventional decoupled architecture of “retargeting first, tracking later” suffers from two major bottlenecks: **i) mathematical non-convexity**[5, 6, 7]. Motion retargeting is inherently a highly non-convex optimization problem and is therefore prone to becoming trapped in local optima. As a

result, such methods are highly sensitive to initialization and require tedious parameter tuning. When poorly initialized, they often produce physically infeasible artifacts, such as abrupt joint jerks, self-interpenetration, and foot sliding, thereby forcing downstream controllers to learn compensatory behaviors or lower stability. **ii) noise propagation.** Human motion data at the source side, such as SMPL-based estimations, often contain noise in the form of ground penetration or temporal jitter. Geometric optimization methods lack awareness of physical plausibility and therefore merely propagate these errors mechanically, leading to a classic “garbage in, garbage out” dilemma.

To address these limitations, we propose a Neural Motion Retargeting (NMR) framework. The central idea is to reformulate retargeting from static optimization over frame-wise states into dynamic mapping between motion distributions. In a data-driven manner, the model can directly learn a mapping from the human motion space to the robot’s feasible motion manifold. However, realizing this vision entails a chicken-and-egg dilemma: training a highly robust neural retargeter requires large-scale, high-quality robot motion data, yet such data are extremely difficult to obtain without an efficient retargeting tool.

To obtain physical plausible motion data, we design a carefully structured hierarchical data pipeline, termed Clustering-Expert Physical Refinement (CEPR). We first use a variational autoencoder (VAE) to extract motion features and cluster heterogeneous human motion data accordingly. Subsequently, we train parallel reinforcement learning expert policies to drive the robot to track these clustered motion sets in a physics simulator, thereby automatically correcting physical defects in the original data and generating “ground-truth” motions that satisfy dynamic constraints. In summary, the contribution of this work includes:

- Proposed a Neural Motion Retargeting (NMR) framework for human to humanoid motion retargeting. By reformulating the retargeting problem as a distribution mapping from the human motion space to the robot’s motion manifold, our method alleviated issues in optimization-based methods such as local minima, joint discontinuities, and self-collisions.
- Introduced a hierarchical data construction pipeline, termed Clustered-Expert Physics Refinement (CEPR).

Through motion clustering, parallel reinforcement learning expert policies, and physics-based simulation refinement, large-scale, high-fidelity, and physically consistent human–robot paired data are automatically generated, providing reliable supervision for neural retargeting model training.

- Developed a transformer-based motion retargeting network and corresponding two-stage training strategy to enable broad motion coverage and bake physical feasibility.
- The effectiveness of the proposed method is validated through diverse motion experiments on the Unitree G1 robot. The results show that joint discontinuities, self-collisions, and joint-limit violations are significantly reduced by NMR, while the training efficiency and tracking performance of downstream whole-body control policies are improved.

## II. RELATED WORK

Motion retargeting aims to transfer human motion data to humanoid robots while accounting for their distinct kinematic structures and physical constraints. We review related work from three perspectives: optimization-based retargeting methods, data-driven retargeting methods, and physics-based motion imitation.

### A. Optimization-based Retargeting Methods

Motion retargeting originated from character animation research in computer graphics. Classical methods[8, 9, 10] employed optimization-based spacetime constraint solvers, formulating motion retargeting as optimization problems with kinematic constraints. While these approaches perform well for single-frame poses, they struggle to guarantee temporal consistency and physical feasibility.

In robotics, simple approaches[3, 11] directly copy joint rotations from human motion to the robot joint space. However, topological and morphological differences between humans and humanoid robots cause such direct mapping to produce artifacts including floating, foot penetration, and end-effector drift. To address joint space misalignment, Whole-Body Geometric Retargeting (WBGR) methods[12, 13] IK to match Cartesian positions and orientations. However, these methods ignore human-robot scale differences and contact states, leading to floating, foot sliding, and ground penetration artifacts.

Recent advances incorporate parametric human body models. The PHC[14] leverages the SMPL model[15] to fit robot skeleton shape parameters and solves IK through gradient descent, an approach widely adopted by H2O[16], HOVER[1], and OmniH2[17]. However, this method is computationally expensive and neglects contact constraints. GMR[4] addresses these issues through non-uniform local scaling and two-stage IK optimization, significantly reducing foot sliding and self-penetration artifacts. PHUMA[18] further introduces multiple physical constraints and jointly optimizes entire motion sequences.

Despite these advances, optimization-based methods remain fundamentally constrained by the inherent non-convexity of IK problems[5, 6, 7, 19], leading to initialization sensitivity and frequent convergence to suboptimal solutions. This mathematical limitation motivates a shift toward data-driven paradigms.

### B. Data-driven Retargeting Methods

To circumvent the local optima problem inherent in optimization methods, data-driven approaches in character animation achieve motion transfer by learning shared latent spaces across different skeletal structures, enabling cross-skeleton retargeting without paired data or 3D reconstruction [20, 21, 22, 23]. These methods leverage cycle-consistency constraints, adversarial training, or contrastive learning to discover correspondences between different embodiments without explicit supervision[24, 25].

In humanoid robotics, similar latent space alignment approaches have been explored for bridging the embodiment gap. Early efforts focused on learning shared representations for translating motions between humans and robots using manually collected paired datasets[26]. To overcome the data bottleneck, recent works introduced self-supervised techniques for automating correspondence discovery: ImitationNet[27] and its variants[28, 29] employ GAN-based or cycle-consistency approaches for teleoperation without paired data. More recent work explores contrastive learning to enhance the expressiveness and smoothness of motion retargeting[30, 31]. However, these methods primarily focus on upper-body manipulation or simple arm motions due to the difficulty of acquiring paired whole-body data and satisfying dynamic constraints for locomotion.

Existing data-driven methods face two limitations: they inherit local optima artifacts from optimization-based supervision, and lack physical reasoning to filter out source motion noise such as ground penetration and temporal jitter.

## III. METHOD

We propose a neural motion retargeting framework that learns a direct mapping from human SMPL motion sequences to feasible humanoid robot motion, bypassing the local-optima failures commonly observed in conventional optimisation-based approaches.

In Section III-A, we introduce typical optimization-based retargeting methods to provide a non-convexity analysis and explained why these methods frequently stall in local optima. This theoretical limitation motivates a shift toward our data-driven paradigm that can circumvent these geometric traps. In Section III-B, we propose a three-stage "Clustered-Expert Physics Refinement" pipeline, which yields approximately 30,000 physics preserved SMPL-robot motion dataset. In Section III-C, we propose a motion retargeting network that learns to directly map human SMPL sequences to physically feasible humanoid motion, and can suppress artifacts caused by upstream pose. In Section III-D, we introduce our "Detail to physical" training scheme that can guarantee both retargeted motion quality and physical fidelity.

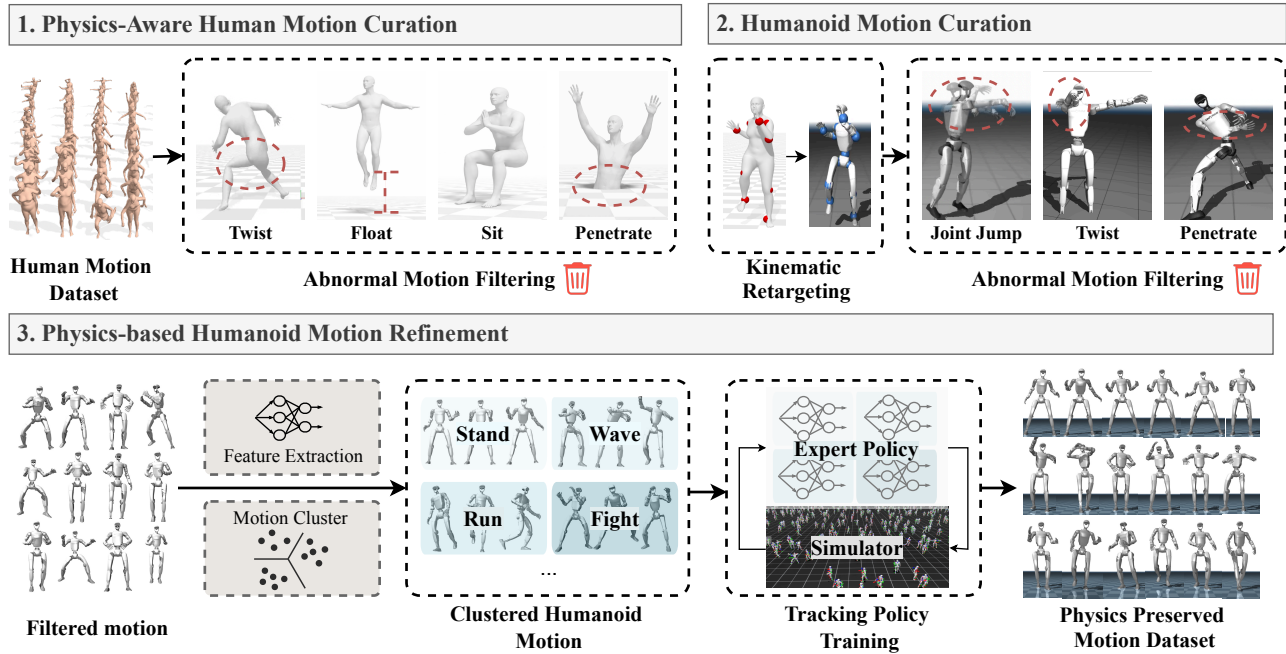


Fig. 1: Data Construction Pipeline. We obtain high-quality human–humanoid motion pairs through three processing stages.

#### A. Preliminary: Optimization-based Human to Humanoid Retargeting

##### Optimization-based Retargeting Methods.

Motion retargeting aims to convert a human SMPL motion sequence  $\{\mathbf{s}_t\}_{t=1}^T$  into a robot-executable sequence  $\{\mathbf{q}_t\}_{t=1}^T$ , where  $\mathbf{q}$  is robot generalized coordinates (root translation, root rotation, and joint values), while preserving motion semantics and satisfying the robot’s physical constraints. Conventional approaches such as GMR[4] formulate this as a per-frame optimization problem:

$$f(\mathbf{q}) = \sum_{(i,j) \in \mathcal{M}} w_{i,j}^R \|R_i^h \ominus R_j(\mathbf{q})\|^2 + \sum_{(i,j) \in \mathcal{M}_{ee}} w_{i,j}^p \|p_i^{\text{target}} - p_j(\mathbf{q})\|^2 \quad (1)$$

where  $R_i^h \in SO(3)$  is the orientation of the human body  $i$ ,  $\mathbf{p}_j(\mathbf{q})$  and  $R_j(\mathbf{q}) \in SO(3)$  are the Cartesian position and orientation of the robot body  $j$ ,  $\mathbf{q}$  is robot generalized coordinates (root translation, root rotation, and joint values),  $R_i \ominus R_j$  is the exponential map representation of the orientation difference between  $R_i$  and  $R_j$ . Although this formulation is formally concise, the underlying optimization is highly non-convex, owing to two mutually coupled geometric challenges.

##### Non-convexity Analysis.

For notational clarity in the following analysis, we denote the joint configuration as  $\boldsymbol{\theta} \in \mathbb{R}^n$ , which corresponds to the joint-angle components of  $\mathbf{q}$ . To expose the geometric source of non-convexity, we consider a single end-effector pose  $T(\boldsymbol{\theta}) \in SE(3)$  and a target pose  $T^* \in SE(3)$ , and

define the relative pose error  $E(\boldsymbol{\theta}) = (T^*)^{-1}T(\boldsymbol{\theta}) \in SE(3)$ . We lift this error to the Lie algebra via

$$\boldsymbol{\xi}(\boldsymbol{\theta}) = \text{Log}(E(\boldsymbol{\theta})) = \begin{pmatrix} \boldsymbol{\omega}(\boldsymbol{\theta}) \\ \mathbf{v}(\boldsymbol{\theta}) \end{pmatrix} \in \mathbb{R}^6, \quad (2)$$

where  $\boldsymbol{\omega} \in \mathbb{R}^3$  is the rotational log-coordinate and  $\mathbf{v} \in \mathbb{R}^3$  is the translational log-coordinate. We then study the weighted surrogate cost

$$\tilde{f}(\boldsymbol{\theta}) = \frac{1}{2} \boldsymbol{\xi}(\boldsymbol{\theta})^\top W \boldsymbol{\xi}(\boldsymbol{\theta}), \quad W = \text{diag}(w_R I_3, w_p I_3), \quad (3)$$

with weights  $w_R, w_p > 0$ . By direct differentiation, the Hessian of this surrogate objective decomposes as

$$\nabla^2 \tilde{f}(\boldsymbol{\theta}) = \underbrace{J_\xi^\top W J_\xi}_{\text{Gauss-Newton term, PSD}} + \underbrace{\sum_{a=1}^6 (W \boldsymbol{\xi})_a \nabla^2 \xi_a}_{\text{curvature correction}}, \quad (4)$$

where  $J_\xi = \frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\theta}} \in \mathbb{R}^{6 \times n}$  is the Jacobian of the log-coordinate error. The first term is always positive semi-definite; the negative curvature arise from the second term, which captures the second-order geometry of the error map  $\boldsymbol{\theta} \mapsto \text{Log}((T^*)^{-1}T(\boldsymbol{\theta}))$ .

Detailed analysis (see Appendix VI-A) reveals two geometrically distinct sources of non-convexity: (i) the curvature of forward kinematics due to recursive trigonometric composition, and (ii) the nonlinearity of the logarithmic map on  $SO(3)/SE(3)$ . Together they imply:

**Proposition 1 (existence of negative curvature).** *Let  $n \geq 2$  and  $w_R, w_p > 0$ . For the surrogate objective  $\tilde{f}$  in (3), there exist feasible target poses  $T^*$  and robot configurations  $\boldsymbol{\theta}$  such*

that the Hessian  $\nabla^2 \tilde{f}(\theta)$  has a strictly negative directional curvature; namely, there exists a direction  $\mathbf{u} \neq \mathbf{0}$  satisfying

$$\mathbf{u}^\top \nabla^2 \tilde{f}(\theta) \mathbf{u} < 0. \quad (5)$$

Hence,  $\tilde{f}$  is generally non-convex.

Proposition 1 establishes that the retargeting landscape admits configurations with strictly negative curvature generated by either forward-kinematics curvature or logarithmic-map curvature. This explains why gradient-based retargeting can be sensitive to initialization and may stall in poor local minima even for seemingly simple target motions.

### Our Core Idea.

The non-convexity of retargeting creates practical limitations. Differential Inverse Kinematics (Differential IK) linearizes the objective at the current iterate, reducing the problem to a convex quadratic program. However, this approximation tends to be valid only within a limited basin of attraction around the true solution, offers no assurance of global convergence, and can be sensitive to initialization and the choice of weighting parameters.

We reframe retargeting as a supervised learning task. But this shift is difficult for two reasons: (i) data: naive use of kinematic retargeting outputs as supervision would inherit the same local-optima failures, making high-quality training pairs difficult to obtain; and (ii) architecture: the network must generalize across diverse motion styles while respecting physical feasibility. To address these challenges, we introduce physics-simulation-based data generation pipeline (III-B) and motion retargeting network (III-C).

### B. Clustered-Expert Physics Refinement.

The systematic transformation from raw human motion to robot-feasible trajectories is summarized in Figure 1. This pipeline acts as a hierarchical filter: Step 1 ensures semantic relevance, Step 2 enforces kinematic validity, and Step 3 leverages the physics engine to resolve dynamic inconsistencies, eventually yielding a high-fidelity dataset for neural network supervision.

#### Step 1 Physics-Aware Human Motion Curation.

The raw SMPL dataset contains a large number of motions that are semantically incompatible with robotic applications that would introduce substantial spurious noise. Following PHUMA[18], we apply the same filtering method to remove physically inconsistent motion including (i) excessive jerk, (ii) a CoM position far outside its support base, or (iii) insufficient foot-ground contact (float and penetration).

After this stage, all retained sequences correspond semantically to motions that the robot can in principle execute, thereby providing a foundation for the subsequent fine-grained filtering steps.

#### Step 2 Kinematic Retargeting and Quality Filtering.

Then, we employ optimization-based kinematic retargeting method (GMR[4]) on curated SMPL sequences to obtain the initial corresponding robot motion dataset. While GMR

effectively mitigates common artifacts via human-robot rest-pose alignment and multi-stage inverse-kinematics (IK) optimization, it remains fundamentally constrained by the non-convex nature of the optimization problem (as analyzed in Section III-A). Consequently, the output quality cannot be universally guaranteed, particularly for complex or highly dynamic motions.

To address potential optimization failures or IK singularities, we apply a hard-threshold filtering pipeline. A motion segment is preserved only if it satisfies the following joint-space and geometric constraints:

- **Joint Continuity and Stability:** We compute the inter-frame joint velocity  $\dot{\mathbf{q}}_t = (\mathbf{q}_t - \mathbf{q}_{t-1})/\Delta t$ . Any segment containing a peak velocity  $|\dot{\mathbf{q}}_t| > \dot{\mathbf{q}}_{max}$  is discarded, where  $\dot{\mathbf{q}}_{max}$  represents the hardware-specific velocity saturation limit. This prunes abrupt, non-smooth jumps caused by IK singularities.
- **Geometric Self-Intersection Detection:** We leverage the MuJoCo collision detection engine to identify interpenetrations of the robot’s geometric links. By loading the robot’s URDF into a simulation context, we check for contacts between all geoms. A sequence is rejected if the fraction of self-intersecting frames exceeds a tolerance `cross_ratio = 0.05`.
- **Floating Foot Rectification:** To ensure the motion is physically grounded, we calculate the average foot clearance relative to the estimated ground plane. If the mean elevation of the lowest foot point across the sequence exceeds `float_threshold = 0.10` m, the motion is classified as "floating" (e.g., sitting or lying poses in the air) and pruned from the training set.

### Step 3 Physics-based Humanoid Motion Refinement

This stage constitutes the core of the entire data pipeline, with the objective of transforming kinematic references into physically consistent robot motion trajectories, using the physics simulator and RL policy as the pseudo ground-truth source.

#### Motion Clustering.

Training a single RL tracking policy over the full motion dataset suffers from distributional conflict [14], leading to unstable performance and degraded tracking accuracy. Conversely, training a separate policy on individual motion sequences is prohibitively expensive in both computation and training time. To address this, we partition the motion library into behaviorally related clusters, allowing each expert policy to specialize over a homogeneous motion distribution.

Specifically, we leverage TMR [32] to train a motion-text retrieval model via contrastive and reconstruction losses, which establishes a well-structured cross-modal latent space. The resulting motion encoder is then used to extract latent representations for all motion sequences. Our goal is to cluster motions by semantic type, for instance, grouping jumps into one cluster and in-place motions into another. The semantically-aligned motion-text features ensure that motions sharing similar semantics are embedded in proximity, even

TABLE I: Policy Observation Space

Observation term	Description	Dimension
<i>Reference motion state <math>s_t^g</math></i>		
$q^g$	Reference joint positions	29
$\dot{q}^g$	Reference joint velocities	29
$p_b^g$	Reference body positions (world)	$14 \times 3 = 42$
$v_b^g$	Reference body linear velocities	$14 \times 3 = 42$
$\sigma_b^g$	Reference body orientations (quat)	$14 \times 4 = 56$
$\omega_b^g$	Reference body angular velocities	$14 \times 3 = 42$
<i>Robot proprioception <math>s_t^p</math></i>		
$p_b^p$	Robot body positions (world)	$14 \times 3 = 42$
$v_b^p$	Robot body linear velocities	$14 \times 3 = 42$
$\sigma_b^p$	Robot body orientations (quat)	$14 \times 4 = 56$
$\omega_b^p$	Robot body angular velocities	$14 \times 3 = 42$
$q^p$	Robot joint positions (relative)	29
$\dot{q}^p$	Robot joint velocities (relative)	29
<i>Action history</i>		
$a_{t-1}$	Previous action	29
<b>Total</b>		<b>509</b>

if their kinematic patterns differ. We then apply K-Means algorithm with cosine similarity as the distance metric to partition all motion sequences in the latent space.

#### Expert Policy Training.

For each cluster, we train an RL tracking policy in a massively parallelized physics simulation environment to make the policy to replicate the target motion in simulator as close as possible. We leveraged symmetric actor-critic framework and PPO algorithm for policy training to maximize the data efficiency.

As shown in Table I, we provide the policy with comprehensive state information including both the reference motion state  $s_t^g$  and the robot’s proprioceptive state  $s_t^p$ . This rich observation space enables the policy to accurately perceive the tracking error between the current robot configuration and the target motion, facilitating precise whole-body motion tracking.

As shown in Table II, we mainly add tracking rewards and only minimal regulation rewards to guarantee accurate motion tracking. Moreover, we employ an adaptive standard deviation schedule for the tracking rewards. When training on large-scale motion datasets, the effective sample count per motion decreases compared to single-motion training. To compensate for this reduced sample efficiency and achieve lower tracking errors, we progressively tighten the reward function by decreasing  $\sigma$  from  $\sigma_{\text{start}}$  to  $\sigma_{\text{end}}$  over the course of training:

$$\sigma(i) = \sigma_{\text{start}} + (\sigma_{\text{end}} - \sigma_{\text{start}}) \cdot \frac{i - i_0}{i_{\text{max}} - i_0} \quad (6)$$

where  $i$  denotes the current training iteration. This curriculum learning strategy allows the policy to first learn coarse motion patterns with relaxed reward tolerances, then gradually refine the tracking precision as training progresses.

#### Generating Physically Faithful Data Pairs.

Once the expert policies have converged, we roll out each

TABLE II: Reward Terms for Expert policy training

Reward term	Weight
<i>Tracking rewards <math>\mathcal{R}(s_t^p, s_t^g)</math></i>	
Anchor position	1.0
Anchor orientation	1.0
Anchor velocity	1.0
Body link position (rel.)	1.0
Body link orientation (rel.)	1.0
Body link linear velocity	1.0
Body link angular velocity	1.0
<i>Penalty terms <math>\mathcal{P}(s_t^p, a_t)</math></i>	
Action rate	-0.1
Undesired contacts	-0.1

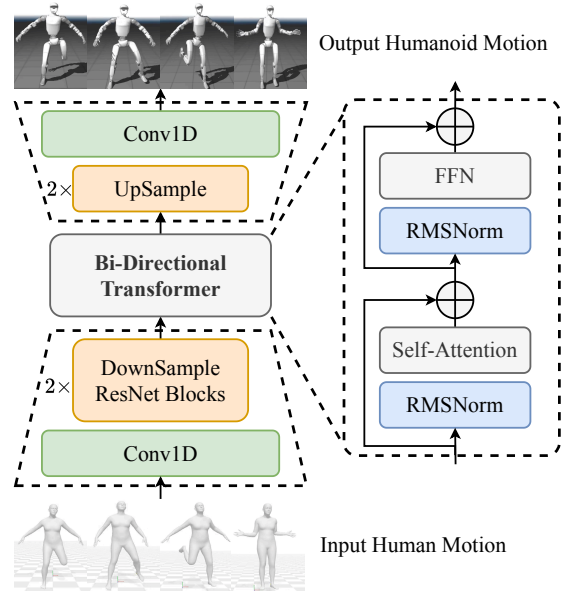


Fig. 2: Overview of our neural motion retargeting network, which maps human motion to humanoid motion.

policy in the simulator over its corresponding reference sequences, record the resulting robot state trajectories, and pair them with the corresponding input SMPL sequences, forming (SMPL sequence, physically consistent robot motion) data pairs. This process yields approximately 30K paired sequences in total, each implicitly validated by the physics simulation, constituting a compact yet high-quality subset of our overall dataset.

#### C. Motion Retargeting Network

##### Motion Representation.

Human and humanoid motion differ in joint structure, including the number of joints and the fact that humanoid robots are actuated by DoFs rather than full joint rotations. To accommodate these differences while maintaining alignment, we adopt distinct motion representations for humans and humanoids.

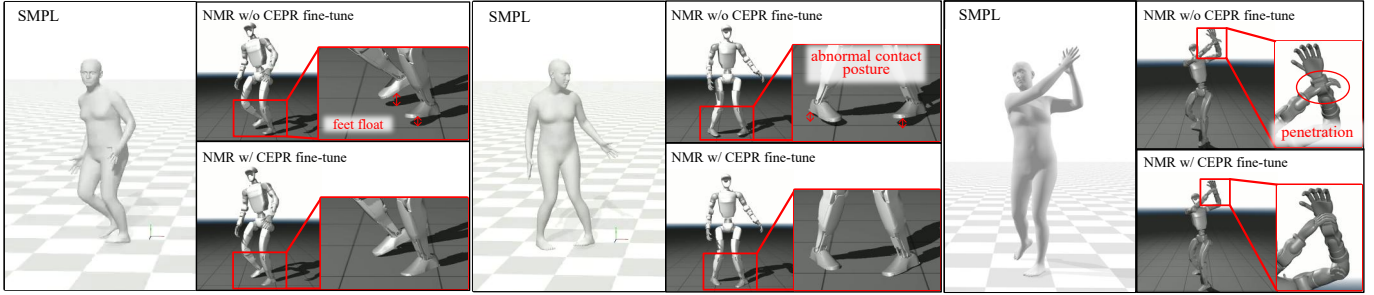


Fig. 3: Visualization of NMR retargeting results with and without CEPR data fine-tuning

For human motion, we build upon the 272-dimensional representation used in MotionMillion [33] and reformulate it as:

$$m^i = \{r^x, r^z, r, j^p, j^v\}, \quad (7)$$

where  $r^x, r^z \in \mathbb{R}$  denote the root linear velocities on the XZ-plane,  $r \in \mathbb{R}^6$  represents the root orientation in 6D rotation representation, and  $j^p, j^v \in \mathbb{R}^{3k}$  denote local joint positions and velocities, respectively. Humanoid motion is represented similarly, with the addition of the robot’s joint DoFs:

$$m_{bot}^i = \{r_{bot}^x, r_{bot}^z, r_{bot}, j_{bot}^p, j_{bot}^v, q\}, \quad (8)$$

where  $r_{bot}^x$  and  $r_{bot}^z$  denote the humanoid root linear velocities on the XZ-plane,  $r_{bot}$  denotes the root orientation,  $j_{bot}^p, j_{bot}^v \in \mathbb{R}^{3d}$  denote local joint positions and velocities, and  $q \in \mathbb{R}^n$  denotes the joint DoFs.

#### Motion Retargeting Network.

As illustrated in Fig. 2, our neural motion retargeting network directly maps human motion to corresponding humanoid sequences. Given a human motion sequence, we first extract latent features via a 1D ResNet-based encoder. The features are then passed into a Transformer-based network following the architectural design of LLaMA [34]. Since human and humanoid motions are temporally aligned with a strict one-to-one correspondence, we replace causal attention with full self-attention, enabling parallel timestep-wise prediction conditioned on the entire input sequence. The output of the Transformer is subsequently decoded through upsampling and 1D-Conv layers to produce the final humanoid sequences. The network is optimized by minimizing an L1 loss:

$$\mathcal{L} = \sum_{t=1}^T \|m_{bot}^t - \hat{m}_{bot}^t\|_1, \quad (9)$$

where T denotes the length of the motion sequence.

#### D. Two-Stage Training Scheme

Corresponding to the two-tier data hierarchy described in Section III-B, we train the retargeting network in two sequential stages: large-scale kinematic pre-training followed by physics-guided fine-tuning. This strategy is necessary because neither dataset alone is sufficient: the kinematic dataset provides breadth but lacks physical guarantees, while the physics

dataset provides feasibility signals but is too small to support generalization.

#### Step 1: Kinematic Alignment with large-scale data.

We first pre-train the network on the large-scale kinematic retargeting dataset, minimizing the regression loss between the predicted G1 joint angles and the kinematic reference targets. Despite the residual physical artifacts in kinematic data (e.g., foot skating, ground penetration), the dataset’s scale and diversity endow the network with a foundational embodiment mapping across a broad range of motion categories, including locomotion, upper-limb manipulation, and martial-arts motions.

#### Step 2: Physical Grounding with CEPR data.

Building on the pre-trained checkpoint, we fine-tune the network using approximately 30,000 physically consistent motion pairs validated by physics simulation (Section III-B). Although this dataset is smaller than the kinematic set by roughly an order of magnitude, each sample carries a strong physical feasibility signal that has been verified through RL policy rollouts in simulation. Fine-tuning shifts the output distribution toward the robot’s dynamically feasible motion manifold, enabling the network to implicitly suppress physically infeasible components in upstream SMPL-X noise rather than propagating them to the output.

The necessity of both stages can be argued from two directions. Pre-training without physics fine-tuning (the *NMR w/o RL* ablation) produces kinematically plausible but physically unconstrained outputs. Conversely, training on physics data alone without pre-training leads to overfitting to the limited set of motion patterns covered by the RL training corpus, failing to generalize to unseen motion types. Together, the two stages equip the network with both broad motion coverage and improved physical feasibility.

## IV. EXPERIMENT

### A. Experiment Setup

For training the motion retargeting network, we adopt a two-stage optimization scheme. During the kinematic-alignment setp, we use the AdamW optimizer with a batch size of 128 and an initial learning rate of  $2 \times 10^{-4}$ , scheduled via cosine annealing. The network is trained for 500 epochs. During the physical grounding setp, we use the same optimizer and batch

TABLE III: Quantitative comparison of different methods in terms of joint jump, self collision and reaching joint limit

Method	Joint Jump ↓	Self Collision ↓	Joint Limit ↓
GMR	56 (0.11%)	947 (1.91%)	21443 (43.21%)
PHUMA	12 (0.02%)	2456 (4.95%)	10580 (21.32%)
NMR w/o RL	2 (0.00%)	779 (1.57%)	19324 (38.94%)
<b>NMR</b>	<b>0 (0.00%)</b>	<b>431 (0.87%)</b>	<b>8339 (16.80%)</b>

size, while reducing the learning rate to  $1 \times 10^{-5}$  for a warm start, and train for an additional 50 epochs.

### B. Datasets and Baselines

To evaluate the proposed model’s performance in human-to-robot retargeting, we curated a test suite from the AMASS[35] dataset containing 82 motion sequences (totaling 119K frames at 120Hz). All test data were strictly excluded from the training phase to ensure unbiased assessment. The dataset is categorized by Motion Complexity and Sequence Length:

1) *Motion Complexity*: Sequences are classified into three levels based on kinematic mapping and hardware constraints:

- Upper-limb-only (ULOM): Involves stationary lower bodies; evaluates workspace mapping and self-collision avoidance.
- Whole-body primitive (WBPM): Includes basic locomotion (e.g., walking, running); assesses coordinated multi-joint retargeting and Center of Mass (CoM) stability.
- Whole-body complex (WBCM): Covers high-dynamic motions (e.g., acrobatics, martial arts); tests robustness against mechanical joint limits and singularities.

2) *Sequence Length*: To measure temporal stability and error accumulation, data is partitioned by frame count:

- Short (< 250 frames): Evaluates initialization speed and transient response to sudden kinematic changes.
- Medium (250–1000 frames): Assesses motion smoothness and the consistency of cyclic gait patterns.
- Long (> 1000 frames): Tests the suppression of cumulative errors and drift over extended operations.

We adopt GMR[4] as baseline methods. GMR employs a unique non-uniform scaling strategy and a two-stage optimization scheme to maximize the likelihood of obtaining a good initialization for the inverse kinematics optimization problem, thereby enabling accurate motion mapping. PHUMA improves the physical plausibility of the retargeted motions by introducing multiple physical constraints and jointly optimizing the entire motion sequence.

We use BeyondMimic[36] for RL-based tracking policy training. We keep reward, policy observation, and domain randomization settings identical to the Beyondmimic, only set `num_envs = 16384` to improve value sampling.

### C. Retargeting Quality and Policy Tracking

We evaluate retargeted motions on three physical plausibility metrics. **Joint Jump** counts frames where max single-step joint angle change exceeds 0.5 rad. **Self-Collision** flags frames

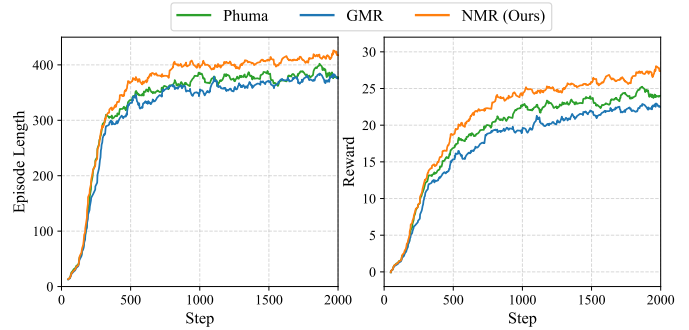


Fig. 4: Comparison of training episode length and reward of motion retargeted with different methods.

with non-hand body-segment contacts via MuJoCo forward kinematics. **Joint Limit** counts frames where any joint comes within 0.05 rad of its hardware boundary.

As shown in Tab. III, NMR achieves the best results across all metrics: zero joint jumps, 54% fewer self-collisions than GMR, and joint limit violations reduced to 16.80%—roughly half of PHUMA. The ablation (NMR w/o RL) confirms that physics refinement is essential for hardware-feasible output.

As shown in Figure 5, The GMR method (green curve) encounters the lower joint limit of the right shoulder roll joint during the initial phase of the motion, causing the optimizer to converge to a local optimum due to improper initialization. During the interval from  $t=0.4s$  to  $t=0.8s$ , the joint angle of GMR remains stagnant near the limit, resulting in substantial deviation of the robot motion from the original SMPL reference, manifested as distorted upper-limb postures. Once the accumulated error exceeds a certain threshold, GMR abruptly escapes from the local minimum at approximately  $t=0.8s$ , with the joint angle undergoing a sudden change of approximately 1.5 rad within 0.2s. The corresponding angular velocity reaches 7.5 rad/s. By contrast, both NMR (blue curve) and PHUMA (orange curve) generate smooth and continuous joint trajectories, with NMR achieving the optimal motion smoothness while maintaining motion similarity.

The downstream effect of motion quality is further reflected in the RL tracking policy training curves (Fig. 4). Policies trained on NMR references reach longer episodes and higher rewards, indicating that cleaner reference motions yield more efficient policy learning.

Table IV reports the end-to-end tracking performance of policies trained on each method’s retargeted motions, evaluated by success rate, mean per-joint pose error (MPJPE), and mean per-joint pose error under world frame (W-MPJPE) across short, medium, and long sequences. NMR achieves the highest success rate and lowest MPJPE/W-MPJPE in all settings. Notably, PHUMA’s high W-MPJPE on short sequences (0.660m vs. 0.237m for NMR) suggests its sequence-level optimization distorts short motions.

TABLE IV: Comparison of tracking accuracy.

Method	Short			Medium			Long		
	Success Rate $\uparrow$	MPJPE $\downarrow$	W-MPJPE $\downarrow$	Success Rate $\uparrow$	MPJPE $\downarrow$	W-MPJPE $\downarrow$	Success Rate $\uparrow$	MPJPE $\downarrow$	W-MPJPE $\downarrow$
PHUMA	26/33	0.058	0.660	<b>41/46</b>	0.042	0.444	9/13	0.0429	0.460
GMR	25/33	0.048	0.244	40/46	0.046	0.450	9/13	0.043	0.319
NMR (Ours)	<b>31/33</b>	<b>0.040</b>	<b>0.237</b>	<b>41/46</b>	<b>0.035</b>	<b>0.308</b>	<b>10/13</b>	<b>0.040</b>	<b>0.291</b>

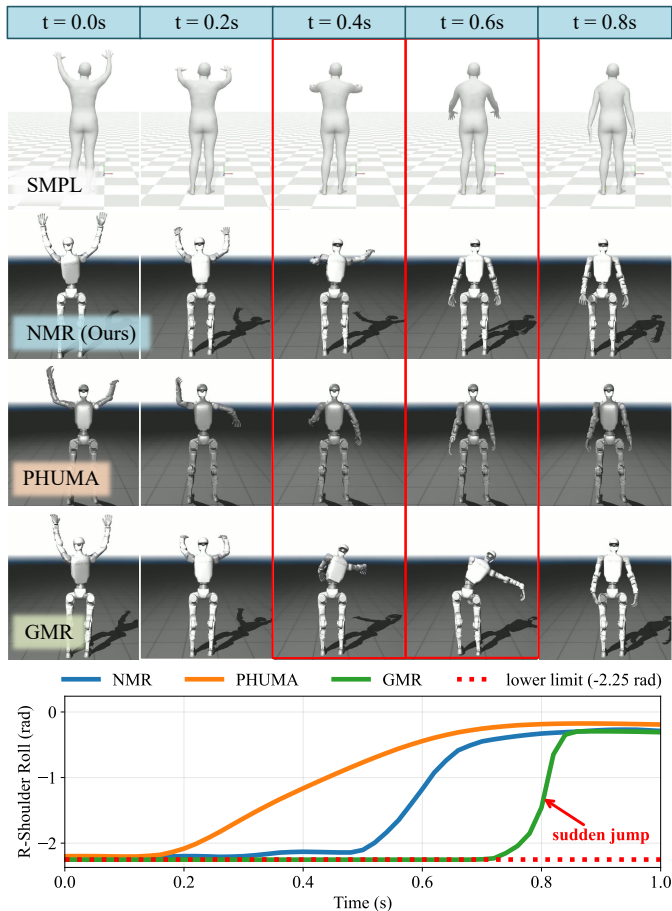


Fig. 5: Visual comparative of different motion retargeting methods when processing a human “arm raise-lower” motion sequence. The top row displays the reference motion sequence from the SMPL human body model, while the subsequent three rows illustrate the retargeted robot motions generated by NMR (ours), PHUMA, and GMR, respectively. The red bounding boxes highlight the key frames where GMR exhibits significant motion anomalies. The line graph at the bottom illustrated “right shoulder roll angle” of the retargeted motion.

#### D. Correcting Upstream Errors

As shown in Figure 6, when the original SMPL sequence contains abnormal jitters caused by pose estimation errors, NMR can implicitly filter out these artifacts and generate smooth, physically feasible robot motion trajectories.

This robustness arises from two key design choices. First, because our training data have undergone physical-consistency

filtering (see Section III-B), motion jitters caused by upstream estimation errors are treated as out-of-distribution samples with respect to the training distribution. Unlike optimization-based methods that solve the problem independently on a frame-by-frame basis, neural networks naturally exhibit a smoothing generalization effect on out-of-distribution inputs: when the input deviates from the training distribution, the model tends to predict continuous and plausible motion sequences rather than mechanically reproducing the input noise.

Second, our network employs a bidirectional self-attention mechanism, enabling the prediction at each frame to access the full temporal context. This global modeling capability allows the network to leverage the normal motion information before and after an anomalous frame for implicit interpolation and correction, thereby suppressing the propagation of local noise. In contrast, optimization-based methods such as GMR can process only one frame at a time and lack cross-frame constraints on physical plausibility.

As a result, they tend to pass upstream errors directly to the output, making them difficult for downstream tracking policies to compensate for and thereby causing unstable motion on real robots.

#### V. CONCLUSION

This paper addresses the physical-feasibility gap in human-to-humanoid motion retargeting by proposing NMR, a Neural Motion Retargeting framework. The central insight is to treat retargeting as a learned distribution mapping rather than a frame-wise geometric optimization.

To enable this data-driven approach, we propose CEPR, a pipeline that leverages VAE-based clustering and massively parallel RL expert policies to generate approximately 30K physically validated human-robot motion pairs. A Transformer-based network, pretrained on large-scale kinematic data and fine-tuned with CEPR-refined pairs, performs efficient inference without requiring a physics simulator. Experiments on the Unitree G1 demonstrate that NMR achieves zero joint jumps, reduces self-collision frames by 54%, and cuts joint-limit violations by 61%, while also accelerating the convergence of downstream whole-body control policies. Furthermore, global temporal attention enables NMR to implicitly suppress upstream SMPL estimation errors rather than propagating them.

A current limitation is that CEPR is morphology-specific; extending to other platforms requires regenerating the data pipeline, and developing morphology-conditioned architectures remains a direction for future work.

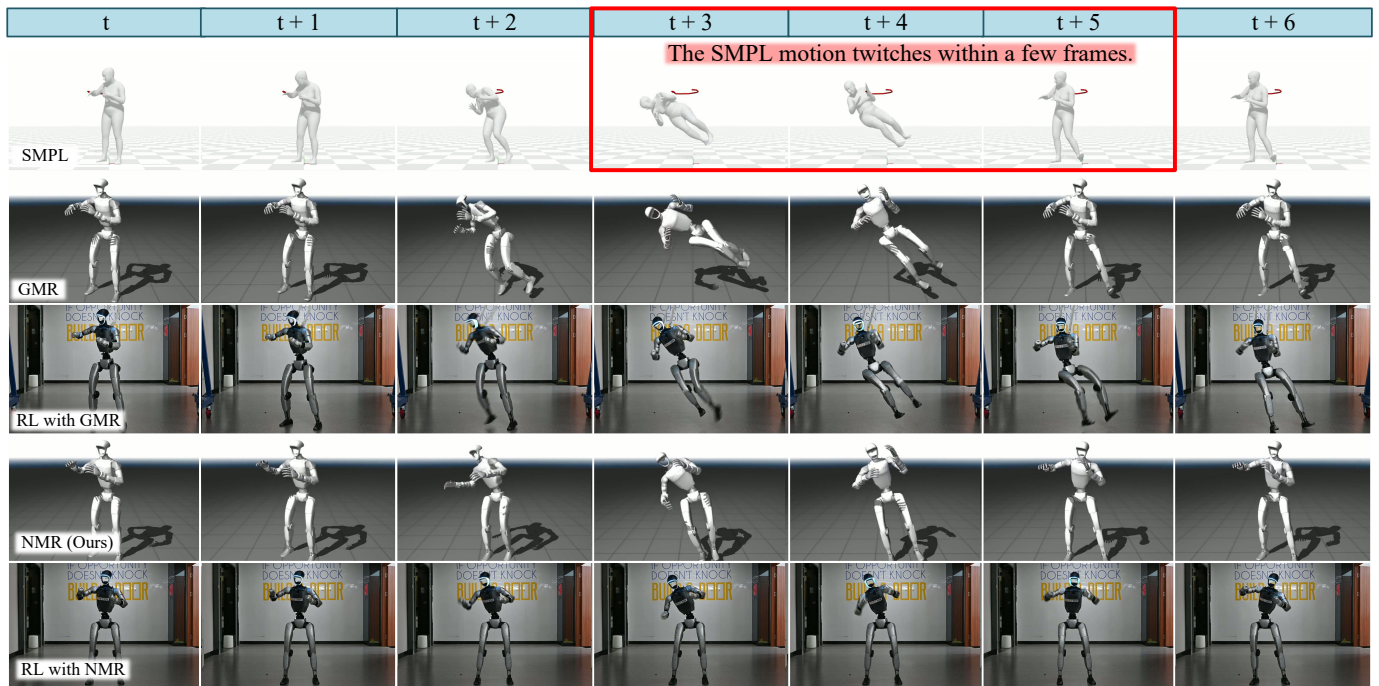


Fig. 6: Comparison under abnormal SMPL motions, frame interval is around 0.06 s. When abnormal poses appear in the original SMPL sequence between t+2 and t+5, the proposed NMR method implicitly filter out the abnormal, producing smoother and more feasible robot motion. In contrast, GMR retain the error from source motion, making it difficult for the learned policy to compensate and causing instability on the real robot.

#### ACKNOWLEDGMENTS

This research is supported by HUAWEI’s AI Hundred Schools Program and was carried out using the Ascend AI technology stack. Thank Tianhao Jiang for helping with hardware experiment.

#### REFERENCES

- [1] T. He, W. Xiao, T. Lin, Z. Luo, Z. Xu, Z. Jiang, J. Kautz, C. Liu, G. Shi, X. Wang *et al.*, “Hover: Versatile neural whole-body controller for humanoid robots,” in *2025 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2025, pp. 9989–9996.
- [2] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, “Humanplus: Humanoid shadowing and imitation from humans,” *arXiv preprint arXiv:2406.10454*, 2024.
- [3] X. Cheng, Y. Ji, J. Chen, R. Yang, G. Yang, and X. Wang, “Expressive whole-body control for humanoid robots,” *arXiv preprint arXiv:2402.16796*, 2024.
- [4] J. P. Araujo, Y. Ze, P. Xu, J. Wu, and C. K. Liu, “Retargeting Matters: General Motion Retargeting for Humanoid Motion Tracking,” *arXiv preprint arXiv:2510.02252*, 2025.
- [5] H. Dai, G. Izatt, and R. Tedrake, “Global inverse kinematics via mixed-integer convex optimization,” *The International Journal of Robotics Research*, vol. 38, no. 12-13, pp. 1420–1441, 2019.
- [6] J. Haviland and P. Corke, “Manipulator differential kinematics: Part i: Kinematics, velocity, and applications [tutorial],” *IEEE Robotics & Automation Magazine*, vol. 31, no. 4, pp. 149–158, 2023.
- [7] J. J. Craig, *Introduction to robotics: mechanics and control, 3/E*. Pearson Education India, 2009.
- [8] Z. Popović and A. Witkin, “Physically based motion transformation,” in *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*, 1999, pp. 11–20.
- [9] S. Tak and H.-S. Ko, “A physically-based motion re-targeting filter,” *ACM Transactions on Graphics (ToG)*, vol. 24, no. 1, pp. 98–117, 2005.
- [10] E. Lyard and N. Magnenat-Thalmann, “Motion adaptation based on character shape,” *Computer Animation and Virtual Worlds*, vol. 19, no. 3-4, pp. 189–198, 2008.
- [11] Z. Fu, Q. Zhao, Q. Wu, G. Wetzstein, and C. Finn, “Humanplus: Humanoid shadowing and imitation from humans,” in *Conference on Robot Learning*. PMLR, 2025, pp. 2828–2844.
- [12] K. Darvish, Y. Tirupachuri, G. Romualdi, L. Rapetti, D. Ferigo, F. J. A. Chavez, and D. Pucci, “Whole-body geometric retargeting for humanoid robots,” in *2019 IEEE-RAS 19th International Conference on Humanoid Robots (Humanoids)*. IEEE, 2019, pp. 679–686.
- [13] L. Penco, B. Clément, V. Modugno, E. M. Hoffman, G. Nava, D. Pucci, N. G. Tsagarakis, J.-B. Mouret, and S. Ivaldi, “Robust real-time whole-body motion retargeting from human to humanoid,” in *2018 IEEE-RAS 18th International Conference on Humanoid Robots*

- (Humanoids). IEEE, 2018, pp. 425–432.
- [14] Z. Luo, J. Cao, K. Kitani, W. Xu *et al.*, “Perpetual humanoid control for real-time simulated avatars,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 10 895–10 904.
- [15] M. Loper, N. Mahmood, J. Romero, G. Pons-Moll, and M. J. Black, “Smpl: A skinned multi-person linear model,” in *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 2023, pp. 851–866.
- [16] 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)*. IEEE, 2024, pp. 8944–8951.
- [17] T. He, Z. Luo, X. He, W. Xiao, C. Zhang, W. Zhang, K. Kitani, C. Liu, and G. Shi, “Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning,” *arXiv preprint arXiv:2406.08858*, 2024.
- [18] K. Lee, S. Kim, M. Park, H. Kim, D. Hwang, H. Lee, and J. Choo, “Phuma: Physically-grounded humanoid locomotion dataset,” *arXiv*, 2025.
- [19] J. Nocedal and S. J. Wright, *Numerical optimization*. Springer, 2006.
- [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, pp. 8639–8648.
- [21] 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.
- [22] S. Lee, T. Kang, J. Park, J. Lee, and J. Won, “Same: Skeleton-agnostic motion embedding for character animation,” in *SIGGRAPH Asia 2023 Conference Papers*, 2023, pp. 1–11.
- [23] Z. Yang, W. Zhu, W. Wu, C. Qian, Q. Zhou, B. Zhou, and C. C. Loy, “Transmomo: Invariance-driven unsupervised video motion retargeting,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 5306–5315.
- [24] L. Hu, Z. Zhang, C. Zhong, B. Jiang, and S. Xia, “Pose-aware attention network for flexible motion retargeting by body part,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 30, no. 8, pp. 4792–4808, 2023.
- [25] H. Zhang, Z. Chen, H. Xu, L. Hao, X. Wu, S. Xu, Z. Zhang, Y. Wang, and R. Xiong, “Semantics-aware motion retargeting with vision-language models,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024, pp. 2155–2164.
- [26] S. Choi, M. K. Pan, and J. Kim, “Nonparametric motion retargeting for humanoid robots on shared latent space,” in *Robotics: science and systems*, 2020.
- [27] Y. Yan, E. V. Mascaro, and D. Lee, “Imitationnet: Unsupervised human-to-robot motion retargeting via shared latent space,” in *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*. IEEE, 2023, pp. 1–8.
- [28] S. Yagi, M. Tada, E. Uchibe, S. Kanoga, T. Matsubara, and J. Morimoto, “Unsupervised neural motion retargeting for humanoid teleoperation,” *arXiv preprint arXiv:2406.00727*, 2024.
- [29] M. Stanley, L. Tao, and X. Zhang, “Robust motion mapping between human and humanoids using cycleautoencoder,” in *2021 IEEE International Conference on Robotics and Biomimetics (ROBIO)*. IEEE, 2021, pp. 93–98.
- [30] T. Wang, D. Bhatt, X. Wang, and N. Atanasov, “Cross-embodiment robot manipulation skill transfer using latent space alignment,” *arXiv preprint arXiv:2406.01968*, 2024.
- [31] Y. Yan and D. Lee, “Learning a unified latent space for cross-embodiment robot control,” *arXiv preprint arXiv:2601.15419*, 2026.
- [32] M. Petrovich, M. J. Black, and G. Varol, “Tmr: Text-to-motion retrieval using contrastive 3d human motion synthesis,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 9488–9497.
- [33] K. Fan, S. Lu, M. Dai, R. Yu, L. Xiao, Z. Dou, J. Dong, L. Ma, and J. Wang, “Go to zero: Towards zero-shot motion generation with million-scale data,” in *CVPR*, 2025, pp. 13 336–13 348.
- [34] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar *et al.*, “Llama: Open and efficient foundation language models,” *arXiv preprint arXiv:2302.13971*, 2023.
- [35] N. Mahmood, N. Ghorbani, N. F. Troje, G. Pons-Moll, and M. J. Black, “AMASS: Archive of motion capture as surface shapes,” in *International Conference on Computer Vision*, Oct. 2019, pp. 5442–5451.
- [36] Q. Liao, T. E. Truong, X. Huang, Y. Gao, G. Tevet, K. Sreenath, and C. K. Liu, “Beyondmimic: From motion tracking to versatile humanoid control via guided diffusion,” *arXiv preprint arXiv:2508.08241*, 2025.

## VI. APPENDIX

### A. Non-convexity Analysis of Retargeting Optimization

This appendix provides the detailed derivation for the non-convexity analysis summarized in Section III-A.

#### Geometric Preliminaries

We briefly introduce the notation on the special Euclidean group  $SE(3) = SO(3) \times \mathbb{R}^3$  and its Lie algebra  $\mathfrak{se}(3) \cong \mathbb{R}^6$ . For any vector  $\mathbf{v} \in \mathbb{R}^3$ , let  $[\mathbf{v}]_{\times}$  denote the  $3 \times 3$  skew-symmetric matrix such that  $[\mathbf{v}]_{\times} \mathbf{u} = \mathbf{v} \times \mathbf{u}$ . For a rotation matrix  $R \in SO(3)$ , its logarithm is written as

$$\text{Log}(R) = \boldsymbol{\omega} = \phi \hat{\mathbf{n}},$$

where  $\phi \in [0, \pi)$  is the rotation angle and  $\hat{\mathbf{n}}$  is the unit rotation axis.

The left Jacobian of  $SO(3)$  and its inverse are

$$\begin{aligned}\mathcal{J}_{SO(3)}(\boldsymbol{\omega}) &= \frac{\sin \phi}{\phi} I + \left(1 - \frac{\sin \phi}{\phi}\right) \hat{\mathbf{n}}\hat{\mathbf{n}}^\top + \frac{1 - \cos \phi}{\phi} [\hat{\mathbf{n}}]_\times, \\ \mathcal{J}_{SO(3)}^{-1}(\boldsymbol{\omega}) &= \frac{\phi/2}{\tan(\phi/2)} I + \left(1 - \frac{\phi/2}{\tan(\phi/2)}\right) \hat{\mathbf{n}}\hat{\mathbf{n}}^\top - \frac{\phi}{2} [\hat{\mathbf{n}}]_\times.\end{aligned}\quad (10)$$

### A Unified Surrogate Objective

To expose the geometric source of non-convexity, we consider a single end-effector pose  $T(\boldsymbol{\theta}) \in SE(3)$  and a target pose  $T^* \in SE(3)$ , and define the relative pose error

$$E(\boldsymbol{\theta}) = (T^*)^{-1}T(\boldsymbol{\theta}) \in SE(3). \quad (11)$$

We lift this error to the Lie algebra via

$$\boldsymbol{\xi}(\boldsymbol{\theta}) = \text{Log}(E(\boldsymbol{\theta})) = \begin{pmatrix} \boldsymbol{\omega}(\boldsymbol{\theta}) \\ \mathbf{v}(\boldsymbol{\theta}) \end{pmatrix} \in \mathbb{R}^6. \quad (12)$$

Here  $\boldsymbol{\omega} \in \mathbb{R}^3$  is the rotational log-coordinate and  $\mathbf{v} \in \mathbb{R}^3$  is the translational log-coordinate. Note that  $\mathbf{v}$  is the translational component of the  $SE(3)$  logarithm and, in general, is not identical to the Euclidean position error; rather, it provides a geometrically consistent coupled pose representation.

We then study the weighted surrogate cost

$$\tilde{f}(\boldsymbol{\theta}) = \frac{1}{2} \boldsymbol{\xi}(\boldsymbol{\theta})^\top W \boldsymbol{\xi}(\boldsymbol{\theta}), \quad W = \text{diag}(w_R I_3, w_p I_3), \quad (13)$$

with weights  $w_R, w_p > 0$ . This objective is not globally identical to the original retargeting loss, but serves as a geometrically meaningful surrogate for analyzing second-order structure.

### Gradient and Hessian Decomposition

Let

$$J_\xi(\boldsymbol{\theta}) = \frac{\partial \boldsymbol{\xi}}{\partial \boldsymbol{\theta}} \in \mathbb{R}^{6 \times n}$$

denote the Jacobian of the log-coordinate error. By direct differentiation,

$$\nabla \tilde{f}(\boldsymbol{\theta}) = J_\xi(\boldsymbol{\theta})^\top W \boldsymbol{\xi}(\boldsymbol{\theta}), \quad (14)$$

and the Hessian is

$$\nabla^2 \tilde{f}(\boldsymbol{\theta}) = \underbrace{J_\xi^\top W J_\xi}_{\text{Gauss-Newton term, PSD}} + \underbrace{\sum_{a=1}^6 (W \boldsymbol{\xi})_a \nabla^2 \xi_a}_{\text{curvature correction}}. \quad (15)$$

The first term is always positive semi-definite. Therefore, any negative curvature must arise from the second term, which captures the second-order geometry of the error map  $\boldsymbol{\theta} \mapsto \text{Log}((T^*)^{-1}T(\boldsymbol{\theta}))$ .

### Source I: Curvature Induced by Forward Kinematics

Even if one ignores the logarithmic-map nonlinearity and focuses only on translational kinematics, the mapping from joint angles to end-effector position is nonlinear due to recursive trigonometric composition. Consider a planar or spatial serial chain with at least two revolute joints ( $n \geq 2$ ), and let the target position be chosen so that the translational error is aligned with the fully extended direction of the chain. Around

the canonical extended configuration  $\boldsymbol{\theta} = \mathbf{0}$ , the end-effector position admits the standard second-order expansion

$$p(\boldsymbol{\theta}) = p(\mathbf{0}) + J_p(\mathbf{0}) \boldsymbol{\theta} + \frac{1}{2} \mathcal{H}_p(\mathbf{0})[\boldsymbol{\theta}, \boldsymbol{\theta}] + o(\|\boldsymbol{\theta}\|^2), \quad (16)$$

where  $J_p$  is the translational Jacobian and  $\mathcal{H}_p$  denotes the second derivative tensor of forward kinematics.

For a perturbation direction that bends an interior joint away from the extended pose, the second-order positional variation points opposite to the extension direction. Hence, if the target is placed further along the extension direction, the translational component of the curvature correction contributes negatively along that perturbation. Equivalently, there exists a direction  $\mathbf{u} \in \mathbb{R}^n$  and a target pose  $T^*$  such that

$$\mathbf{u}^\top \left( \sum_{a=4}^6 (W \boldsymbol{\xi})_a \nabla^2 \xi_a \right) \mathbf{u} < 0 \quad (17)$$

at  $\boldsymbol{\theta} = \mathbf{0}$ . Thus, the nonlinear forward-kinematics map alone can induce negative second-order curvature.

### Source II: Curvature Induced by the Logarithmic Map on $SO(3)$

A second and geometrically distinct source of non-convexity comes from the differential of the logarithm map itself. Consider the pure rotational component

$$\boldsymbol{\omega}(\boldsymbol{\theta}) = \text{Log}(R^{*\top} R(\boldsymbol{\theta})).$$

Its Jacobian involves the inverse left Jacobian  $\mathcal{J}_{SO(3)}^{-1}(\boldsymbol{\omega})$ , whose coefficients depend nonlinearly on the rotation angle  $\phi = \|\boldsymbol{\omega}\|$ . In particular, the scalar coefficient

$$\alpha(\phi) = \frac{\phi/2}{\tan(\phi/2)} \quad (18)$$

satisfies

$$\alpha'(\phi) < 0, \quad \phi \in (0, \pi), \quad (19)$$

which shows that the differential of Log varies nonlinearly and increasingly sharply as the rotation error approaches  $\pi$ .

Now choose a target orientation  $R^*$  such that at some configuration  $\boldsymbol{\theta}_0$  the relative rotation

$$R^{*\top} R(\boldsymbol{\theta}_0)$$

has angle  $\phi(\boldsymbol{\theta}_0) \in (\pi/2, \pi)$ , and choose a perturbation direction  $\mathbf{u}$  such that  $d\phi(\boldsymbol{\theta}_0)[\mathbf{u}] \neq 0$ . Then the second derivative of the rotational log error contains a term proportional to the variation of  $\mathcal{J}_{SO(3)}^{-1}$  with respect to  $\phi$ , which contributes negatively along  $\mathbf{u}$  for a suitable choice of target and local motion direction. Therefore, there exist  $\boldsymbol{\theta}_0$ ,  $R^*$ , and  $\mathbf{u}$  such that

$$\mathbf{u}^\top \left( \sum_{a=1}^3 (W \boldsymbol{\xi})_a \nabla^2 \xi_a \right) \mathbf{u} < 0. \quad (20)$$

This shows that negative curvature can arise even when the kinematic map is locally regular, purely due to the geometry of the rotational logarithm.

### Proof of Proposition 1

The two mechanisms above are geometrically distinct: the first originates from the second-order curvature of forward

kinematics, while the second comes from the nonlinearity of the logarithmic chart on  $SO(3)/SE(3)$ . Together they imply the following result.

**Proposition 1 (existence of negative curvature).** *Let  $n \geq 2$  and  $w_R, w_p > 0$ . For the surrogate objective  $\tilde{f}$  in (13), there exist feasible target poses  $T^*$  and robot configurations  $\theta$  such that the Hessian  $\nabla^2 \tilde{f}(\theta)$  has a strictly negative directional curvature; namely, there exists a direction  $\mathbf{u} \neq \mathbf{0}$  satisfying*

$$\mathbf{u}^\top \nabla^2 \tilde{f}(\theta) \mathbf{u} < 0. \quad (21)$$

Hence,  $\tilde{f}$  is generally non-convex.

*Proof:* By the analysis in Sections VI-A.4 and VI-A.5, there exist configurations where either the forward-kinematics curvature or the logarithmic-map curvature dominates, making the curvature correction term in (15) sufficiently negative to overcome the positive semi-definite Gauss–Newton term along certain directions. This establishes the existence of negative directional curvature. ■

### B. Test Motion Names Used in Evaluation

TABLE V: List of Test Motion Files from AMASS (1/2)

Type	Source	File Name
ULOM/L	ACCAD (Female1Gestures_c3d)	D2
ULOM/L	ACCAD (Female1Gestures_c3d)	D3
ULOM/L	CMU	02_07
ULOM/L	CMU	02_10
ULOM/L	CMU	13_20
ULOM/M	ACCAD (Female1Gestures_c3d)	D1
ULOM/M	ACCAD (Male1General_c3d)	General_A2
ULOM/M	ACCAD (Male1General_c3d)	General_A4
ULOM/M	ACCAD (Male1General_c3d)	General_A5
ULOM/M	ACCAD (Male1General_c3d)	General_A6
ULOM/M	CMU	02_05
ULOM/M	CMU	02_08
ULOM/M	CMU	02_09
ULOM/M	CMU	79_46
ULOM/M	CMU	79_90
ULOM/S	ACCAD (Male2MartialArtsPunches_c3d)	E1
ULOM/S	ACCAD (Male2MartialArtsPunches_c3d)	E2
ULOM/S	ACCAD (Male2MartialArtsPunches_c3d)	E3
ULOM/S	ACCAD (Male2MartialArtsPunches_c3d)	E5
ULOM/S	ACCAD (Male2MartialArtsPunches_c3d)	E6
WBCM/L	ACCAD (Male2MartialArtsExtended_c3d)	Extended_1
WBCM/L	ACCAD (Male2MartialArtsExtended_c3d)	Extended_2
WBCM/L	ACCAD (Male2MartialArtsExtended_c3d)	Extended_3
WBCM/L	CMU	01_01
WBCM/L	CMU	06_13
WBCM/L	CMU	122_22
WBCM/M	ACCAD (Female1Gestures_c3d)	D5
WBCM/M	ACCAD (Male2MartialArtsStances_c3d)	D10
WBCM/M	ACCAD (Male2MartialArtsStances_c3d)	D6
WBCM/M	ACCAD (Male2MartialArtsStances_c3d)	D7
WBCM/M	ACCAD (Male2MartialArtsStances_c3d)	D9_-_f2
WBCM/M	ACCAD (Male2MartialArtsStances_c3d)	D9_-_victory
WBCM/M	ACCAD (Male2MartialArtsStances_c3d)	D9_-_warm

TABLE VI: List of Test Motion Files from AMASS (2/2, continue)

Type	Source	File Name
WBCM/M	CMU	05_02
WBCM/M	CMU	05_04
WBCM/M	CMU	05_18
WBCM/M	CMU	106_15
WBCM/S	ACCAD (Male2MartialArtsKicks_c3d)	G11
WBCM/S	ACCAD (Male2MartialArtsKicks_c3d)	G16
WBCM/S	ACCAD (Male2MartialArtsKicks_c3d)	G17
WBCM/S	ACCAD (Male2MartialArtsKicks_c3d)	G19
WBCM/S	ACCAD (Male2MartialArtsStances_c3d)	D13
WBCM/S	ACCAD (Male2MartialArtsStances_c3d)	D14
WBCM/S	ACCAD (Male2MartialArtsStances_c3d)	D1
WBCM/S	ACCAD (Male2MartialArtsStances_c3d)	D3
WBCM/S	CMU	02_04
WBPM/L	CMU	02_06
WBPM/L	CMU	137_03
WBPM/M	ACCAD (Female1Running_c3d)	C10
WBPM/M	ACCAD (Female1Running_c3d)	C14
WBPM/M	ACCAD (Female1Running_c3d)	C21
WBPM/M	ACCAD (Female1Running_c3d)	C25
WBPM/M	ACCAD (Female1Running_c3d)	C2
WBPM/M	ACCAD (Female1Running_c3d)	C5
WBPM/M	ACCAD (Female1Running_c3d)	C6
WBPM/M	ACCAD (Female1Running_c3d)	C8
WBPM/M	ACCAD (Female1Walking_c3d)	B11
WBPM/M	ACCAD (Female1Walking_c3d)	B15
WBPM/M	ACCAD (Female1Walking_c3d)	B16
WBPM/M	ACCAD (Female1Walking_c3d)	B20
WBPM/M	ACCAD (Female1Walking_c3d)	B21
WBPM/M	ACCAD (Female1Walking_c3d)	B21_s2
WBPM/M	ACCAD (Female1Walking_c3d)	B21_s3
WBPM/M	ACCAD (Female1Walking_c3d)	B22
WBPM/M	ACCAD (Female1Walking_c3d)	B5
WBPM/M	ACCAD (Male1General_c3d)	General_A12
WBPM/M	CMU	06_10
WBPM/M	CMU	06_12
WBPM/M	CMU	09_12
WBPM/M	CMU	122_33
WBPM/M	CMU	122_36
WBPM/M	CMU	137_04
WBPM/M	CMU	143_38
WBPM/S	ACCAD (Female1Running_c3d)	C13
WBPM/S	ACCAD (Female1Running_c3d)	C21_s2
WBPM/S	ACCAD (Female1Running_c3d)	C26
WBPM/S	ACCAD (Female1Running_c3d)	C27
WBPM/S	ACCAD (Female1Running_c3d)	C4
WBPM/S	ACCAD (Female1Running_c3d)	C7
WBPM/S	CMU	02_01
WBPM/S	CMU	02_02
WBPM/S	CMU	02_03
WBPM/S	CMU	07_01
WBPM/S	CMU	07_02
WBPM/S	CMU	07_03
WBPM/S	CMU	07_06
WBPM/S	CMU	07_07
WBPM/S	CMU	07_08
WBPM/S	CMU	07_09
WBPM/S	CMU	07_10
WBPM/S	CMU	07_11
WBPM/S	CMU	07_12