SMR: LEARNING STREAMING MOTION RETARGETING FOR HUMANOIDS FROM HUMAN SEMANTIC

Chun Gu[1,2,3] **Juntong Zhang**[2,3] **Jingyu Li**[1,2,3] **Li Zhang**[1,2,3]*

[1]School of Data Science, Fudan University [2]Shanghai Innovation Institute
[3]Logos Intelligence Lab

logosintelligencelab.github.io/SMR

ABSTRACT

Retargeting human motion to humanoids requires preserving motion semantics while accounting for the robot's physical constraints. Kinematic pipelines prescribe a robot trajectory through geometric matching, without establishing whether it can be executed under the robot's dynamics. Prior physics-supervised neural methods train bidirectional full-sequence models on physics-refined kinematic trajectories, remaining tied to prescribed kinematic solutions and requiring noncausal context. We propose SMR, a streaming neural motion retargeter learned from semantic-conditioned physics experts. We convert human motion into robot-aware task-space targets for body landmark positions, orientations, and velocities. Physics experts track these targets directly in simulation under contact, dynamics, and actuator constraints. We train a single strictly causal Transformer on their executed state trajectories that complete the motion and pass semantic and physical validity checks. At each incoming frame, the student uses only current and past human observations to generate a robot motion reference for a downstream whole-body controller. Experiments demonstrate that semantic-conditioned physics experts generate motion with higher semantic fidelity and physical validity than direct optimization-based retargeting, while the resulting streaming retargeter achieves improved closed-loop performance over kinematic and physics-refined baselines under a common whole-body controller.

1 INTRODUCTION

Large-scale human motion datasets such as AMASS, LAFAN1, and BONES-SEED provide a rich source of whole-body motion for humanoid learning and control (Mahmood et al., 2019; Harvey et al., 2020; Bones Studio, 2026). Transferring these motions to robots requires more than geometric matching because humans and robots differ in morphology, joint structure, mass distribution, contact mechanics, and actuation. For streaming applications such as live teleoperation, retargeting must satisfy two requirements: preserving human motion semantics in physically executable robot references, and generating those references from current and past human observations.

The first requirement depends on how robot motion is generated for supervision. Kinematic retargeting pipelines use inverse kinematics or geometric optimization to construct robot trajectories (Penco et al., 2018; Darvish et al., 2019; Araujo et al., 2025). Although these methods can incorporate contact and collision constraints, geometric agreement alone does not establish dynamic executability. Using the resulting trajectories as neural supervision can tie the learned mapping to solutions selected by the kinematic solver. Physics-supervised methods such as NMR (Zhao et al., 2026) refine these trajectories through simulated tracking, but still use the prescribed kinematic motion as the expert's tracking target. This motivates specifying the motion to preserve directly through robot-aware task-space semantics and using the resulting physical execution to generate robot supervision.

The second requirement is causal inference from a live human stream. NMR's learned retargeter (Zhao et al., 2026) uses bidirectional modeling over a complete motion sequence and therefore

*Corresponding author (lizhangfd@fudan.edu.cn).

Table 1: Representative comparison of retargeting methods across kinematics-aware (GMR (Araujo et al., 2025), PHUMA (Lee et al., 2025), UMR (Cao et al., 2026)), dynamics-grounded (IKMR (Chen et al., 2025a), ReActor (Müller et al., 2026), NMR (Zhao et al., 2026)), and real-time teleoperation (GMR (Araujo et al., 2025)) capabilities.

Capability	GMR	PHUMA	UMR	IKMR	ReActor	NMR	SMR (Ours)
Kinematics-aware	✓	✓	✓	✓	✓	✓	✓
Dynamics-grounded	×	×	×	✓	✓	✓	✓
Real-time teleoperation	✓	×	×	×	×	×	✓

depends on future observations. This prevents use in a streaming setting where the reference at time t must be generated from observations available up to t. The challenge is therefore to retain the benefits of physics-based supervision while learning a causal retargeter.

To address these challenges, we propose SMR, a streaming neural motion retargeter learned from semantic-conditioned physics experts. We represent human motion as robot-aware task-space targets for body landmark positions, orientations, and velocities. The experts track these targets directly in closed-loop simulation, using robot state and limited future semantic context to generate motion under contact, dynamics, and actuator constraints. IK is used only for initialization, rather than to prescribe a tracking trajectory. We retain executed state rollouts that complete the motion and pass semantic and physical validity checks, and use them to train a single strictly causal Transformer. At each incoming frame, the student generates a robot motion reference for a downstream whole-body controller using only current and past human observations. As summarized in Table 1, SMR is the only compared method combining dynamics-grounded supervision with real-time teleoperation.

Experiments show that semantic-conditioned physics experts generate motion with higher semantic fidelity and physical validity than direct optimization-based retargeting. Under the same frozen whole-body controller, SMR achieves improved closed-loop performance over kinematic and physics-refined baselines. We further demonstrate whole-body motion execution on a physical Unitree G1 using SMR-generated references.

Our contributions are threefold: (1) we formulate physics-based retargeting around robot-aware human semantics, allowing experts to determine the robot trajectory instead of refining a prescribed kinematic solution; (2) we distill privileged expert rollouts into a single causal retargeter that generates streaming motion references from human observations, without robot-state inputs or future frames; and (3) we demonstrate improved closed-loop performance over kinematic and physics-refined baselines under a common frozen controller, alongside whole-body motion demonstrations on a physical Unitree G1.

2 Related Work

2.1 Optimization-based motion retargeting

Early motion retargeting methods transferred motion between articulated characters through constrained optimization (Gleicher, 1998; Popović & Witkin, 1999; Tak & Ko, 2005). For humanoids, geometric methods use Cartesian objectives to accommodate differences in morphology and joint structure (Penco et al., 2018; Darvish et al., 2019). Later approaches improve morphology adaptation and geometric matching (Araujo et al., 2025), incorporate physical constraints into data generation (Lee et al., 2025), and preserve human–object or human–environment interactions (Yang et al., 2025). UMR further uses learned surface correspondences to guide constrained pose matching and contact transfer (Cao et al., 2026). Dynamics-aware methods incorporate force equilibrium (Rouxel et al., 2022), optimal control (Grandia et al., 2023), or kinodynamic trajectory optimization (Zhang et al., 2026; Wang et al., 2026; Dhedin et al., 2026). Direct Dynamic Retargeting uses sampling-based model predictive control to optimize task-space imitation directly in simulation, without an intermediate kinematic trajectory (Roux et al., 2026). These methods generate robot motion by optimizing input poses or trajectories. SMR instead learns from accepted physics-expert rollouts and generates streaming references with a causal neural model, without solving a retargeting optimization problem at deployment.

2.2 Learning-based and physics-supervised retargeting

Learning-based retargeters learn cross-skeleton mappings through shared latent spaces, skeleton-aware architectures, and paired or unpaired motion translation (Villegas et al., 2018; Aberman et al., 2020; Lee et al., 2023; Choi et al., 2020; Yan et al., 2023). Online upper-body retargeting has also been demonstrated using unpaired adversarial learning (Yagi et al., 2024). Physics enters these models through feasibility constraints in Human2Humanoid (Huang et al., 2026) and simulated tracking feedback in IKMR (Chen et al., 2025a). ReActor jointly adapts reference motions and trains a tracking policy through bilevel optimization (Müller et al., 2026), while OmniTrack uses privileged simulation rollouts to generate references for training a downstream tracking policy (Li et al., 2026). NMR is most closely related to our approach: it trains clustered RL experts to track GMR-generated trajectories and uses their rollouts to supervise a full-sequence neural retargeter (Zhao et al., 2026). In SMR, experts directly track robot-aware human semantics, with IK used only for initialization. Accepted state rollouts then supervise a human-only, strictly causal motion-reference generator.

2.3 Physics-based motion tracking

Physics-based motion tracking studies how robots and simulated characters execute reference motions under dynamics and contact constraints. DeepMimic and PHC address this problem through reinforcement learning in simulation (Peng et al., 2018; Luo et al., 2023). Subsequent humanoid systems improve sim-to-real transfer (He et al., 2025) and motion diversity (Liao et al., 2026; Chen et al., 2025b; Han et al., 2025), introduce online dynamics adaptation (Zhang et al., 2025), and scale motion tracking through larger datasets and models (Luo et al., 2026; Qi et al., 2026). MimicLite focuses on efficient training of general motion-tracking policies (Robotparty Lab, 2026). These approaches use robot-state feedback to produce control actions, whereas SMR generates motion references from human observations. We evaluate these references with a frozen whole-body controller to isolate the effect of retargeting on downstream execution.

3 Method

3.1 Overview

SMR learns streaming motion retargeting from physics experts that execute human semantic commands. In Section 3.2, we construct robot-aware task-space commands and train physics experts to track them in simulation. We retain the resulting state rollouts that pass semantic and physical validity checks. In Section 3.3, we use these rollouts to train a single strictly causal Transformer that generates streaming motion references. Figure 1 illustrates the complete SMR pipeline.

3.2 Semantic-conditioned physics experts

Human semantic commands. We assemble a corpus from AMASS (Mahmood et al., 2019), LAFAN1 (Harvey et al., 2020), BONES-SEED (Bones Studio, 2026), 100STYLE (Mason et al., 2022), and several additional smaller sources. We use the source human motions rather than converted robot trajectories and represent each motion as a canonical SMPL-H sequence at 50 Hz (Loper et al., 2015; Romero et al., 2017). Let H_t denote the SMPL-H state at frame t, including the root transform and articulated body pose. For a target robot $\mathcal{R}$, a deterministic morphology-aware transport maps the human motion to robot-shaped body landmarks using a fixed neutral-leg scale and the robot's neutral geometry. The resulting landmarks form human semantics for the target robot. The transport only produces task-space targets; it neither solves IK nor produces a robot joint trajectory. From the transported positions and orientations, we form a predefined set $\mathcal{K}$ of 15 body landmarks and construct the human semantic command:

$$c_t = \Psi_{\mathcal{R}}(H_t) = \{p_{k,t}, R_{k,t}, v_{k,t}, \omega_{k,t}\}_{k\in\mathcal{K}}\,. \tag{1}$$

Here each landmark $k \in \mathcal{K}$ contributes a transported position $p_{k,t}$ and orientation $R_{k,t}$, together with their finite-difference linear and angular velocities $v_{k,t}$ and $\omega_{k,t}$. These fields define the human semantic command at 50 Hz. The archived command fields are world-frame quantities. For expert control, the root displacement and orientation residuals in the actor observation are expressed in the

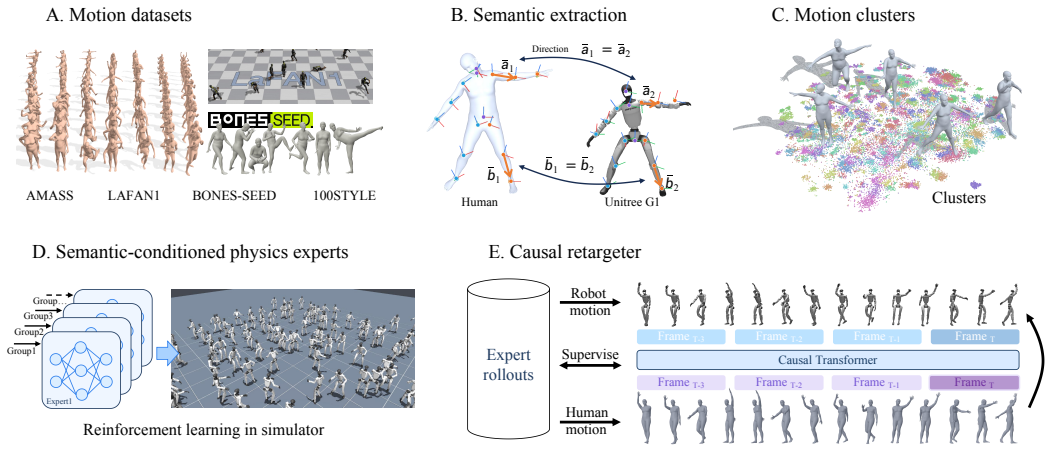


Figure 1: Overview of SMR. (A–C) Human motions from multiple datasets are converted into robot-aware task-space semantics and grouped for expert training. (D) Physics experts track these targets directly in simulation. (E) Accepted expert state rollouts supervise a single strictly causal transformer that generates robot motion references from current and past human observations.

current robot heading-local frame, while root tracking errors use world-frame quantities and body-landmark errors use the heading-local frame. Thus the human semantic command specifies the task-space motion to preserve without prescribing a robot joint trajectory. Figure 2 visualizes these transported semantics together with the robot tracking them in simulation. See the supplementary material for details.

Motion clustering. To organize the heterogeneous motion patterns in the corpus, we partition the motions into training pools before fitting the physics experts. For the large-scale sources, we follow Humanoid-GPT (Qi et al., 2026) and apply Harmonic Motion Embedding (HME) directly to 50-Hz human semantic sequences, whereas Humanoid-GPT embeds retargeted robot motion. A Periodic Autoencoder (Starke et al., 2022) extracts harmonic amplitude and frequency features from two-second windows. We concatenate the mean and standard deviation of these features for each motion into a 256-dimensional descriptor, cluster the descriptors, and use each nonempty cluster as the training pool for one semantic-conditioned expert. Smaller or specialized motion sets bypass clustering and form one single-motion expert per motion. This procedure produces more than 400 experts that share the same human semantic command definition and training objective; the clustering only organizes expert training and rollout generation.

Semantic-conditioned expert training. For each resulting motion pool, we train a closed-loop policy in the target-robot simulator with proximal policy optimization (PPO) (Schulman et al., 2017). At policy step t, its observation is the concatenation

$$o_t = \mathrm{concat}\big(\tilde{c}_t, x_t^{\mathrm{hist}}, a_{t-3:t-1}\big), \qquad a_t \sim \pi_e(\cdot \mid o_t), \tag{2}$$

where $\tilde{c}_t$ is the heading-local window of the human semantic command, x_t^{hist} is the sparse robot proprioceptive history, and $a_{t-3:t-1}$ stacks the three most recent policy actions. The command window contains the current positional and orientation components of the human semantic command, a short history, and a fixed finite future context. Landmark velocities are used only in the tracking reward and the privileged critic observation, and are not policy inputs. The policy outputs a 29-dimensional joint-position target. The reward combines human semantic tracking with regularization terms. Tracking terms compare the simulated root and selected body landmarks with the positions, orientations, and velocities specified by the human semantic command. Regularization terms discourage abrupt action changes and excessive joint velocity, penalize soft joint-limit violations and self-collision, and encourage survival. The actor observation fields, dimensions, and reward weights are summarized in Table 2. Please refer to the supplement for details. Episodes are initialized from a robot pose obtained by offline per-frame IK to provide a suitable starting state and avoid premature termination.

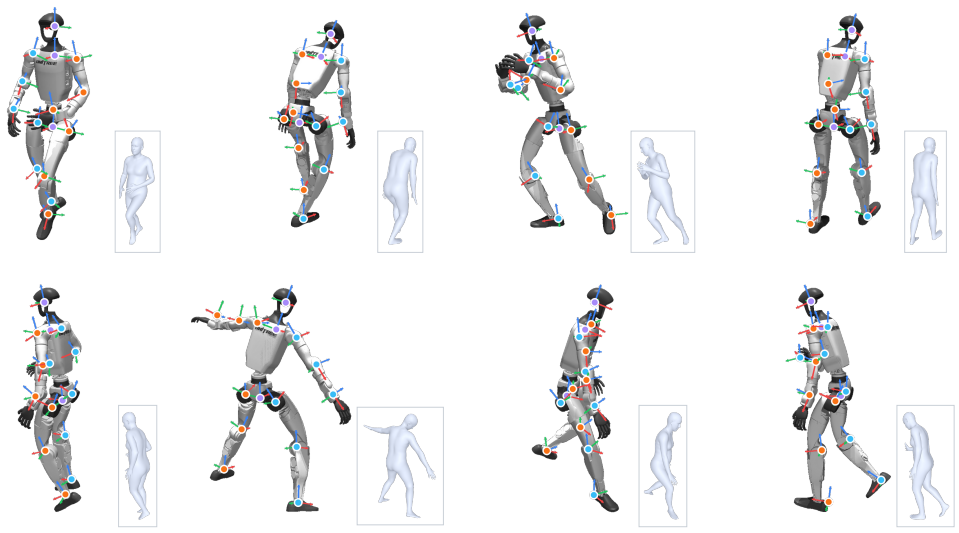


Figure 2: Qualitative examples of robot-aware semantic tracking across eight motions. Each example pairs a source SMPL-H pose (inset) with the corresponding Unitree G1 pose produced by a physics expert in simulation. Colored markers and axes depict the transported landmark targets, while the Unitree G1 poses preserve distinctive features of the source postures despite differences in morphology.

Unlike NMR, which first constructs a robot trajectory with GMR and uses it as the expert tracking target, our expert directly uses the human semantic command as its tracking target. This avoids introducing the kinematic bias of an optimization-based GMR solution into the physics stage.

Expert rollout supervision. After training, each expert deterministically replays its assigned motions in simulation and records the joint, root, and body states at 50 Hz. A rollout is accepted only if it reaches the final frame without termination and passes the fixed semantic-fidelity and physical-validity gates. Accepted state sequences are paired frame by frame with their source human motions and corresponding human semantic commands, where m_t^e denotes the simulator state at timestamp t. The resulting supervision is *dynamics-grounded* because each rollout is executed and accepted under the specified robot, contact, actuator, and simulator models.

3.3 Neural Retargeting

Input and output representation. SMR distills the accepted expert rollouts into a single retargeter that uses only human observations. The student receives no simulator state, expert action, expert identity, or future human frame. At time t, its input is an 85-D feature consisting of a 79-D human semantic pose and a 6-D human-root twist. The pose features contain (i) the positions of 14 non-root landmarks relative to the pelvis (42-D), (ii) pelvis height (1-D), (iii) a 6-D heading-removed pelvis-tilt representation, and (iv) 6-D orientations of the torso, wrists, and ankles (30-D). The twist (v) contains linear and angular backward differences of the human root, expressed in the preceding human-root frame at the native observation rate. Landmark positions and orientations use the heading frame of the current human root. The pose features contain no absolute horizontal position or yaw. Given the causal history $u_{0:t}$, the student predicts the robot motion reference

$$\widehat{m}_t = f_\theta(u_{0:t}) = \mathrm{concat}\Big(\widehat{q}_t, \widehat{\Delta p}_t, \widehat{\Delta r}_t\Big) \in \mathbb{R}^{35}, \tag{3}$$

where $\widehat{q}_t \in \mathbb{R}^{29}$ contains Unitree G1 joint positions, $\widehat{\Delta p}_t \in \mathbb{R}^3$ is a pelvis translation offset, and $\widehat{\Delta r}_t \in \mathbb{R}^3$ parameterizes the relative pelvis rotation. Both offsets are expressed in the current human-pelvis frame, and the root reference is reconstructed by composing them with the current human pelvis pose. This anchored representation avoids integrating a predicted root velocity, while

Table 2: Actor observations and expert rewards for semantic-conditioned expert training.

Actor observation	Dim.
Semantic command window (8 samples)	
Root displacement	$8 \times 3 = 24$
Root orientation	$8 \times 6 = 48$
Tracked body position	$8 \times 13 \times 3 = 312$
Selected body orientation	$8 \times 5 \times 6 = 240$
Robot proprioception history (7 samples)	
Root angular velocity	$7 \times 3 = 21$
Projected gravity	$7 \times 3 = 21$
Joint position offset	$7 \times 29 = 203$
Joint velocity	$7 \times 29 = 203$
Policy action history (3 actions)	
Joint-position action	$3 \times 29 = 87$
Total	1159

(a) Actor observations

Expert reward	Weight
Command tracking	
Root position	1.0
Root orientation	0.5
Root linear velocity	0.5
Root angular velocity	0.5
Body position	1.0
Body orientation	0.5
Body linear velocity	0.5
Body angular velocity	0.25
Regularization	
Action rate	0.5
Joint velocity	10^{-4}
Soft joint limits	10.0
Self-collision	-1.0
Survival	5.0

(b) Expert reward

the reconstructed root can still inherit drift, heading noise, and ground-height error from the human input.

Causal transformer student. To support streaming retargeting, we instantiate f_θ as a strictly causal Transformer (Vaswani et al., 2017) with four pre-normalized blocks of width 512, multi-head self-attention with 16 heads, and rotary position embeddings (Su et al., 2024). A final normalization and linear head produce $\widehat{m}_t$. Each position attends only to itself and the 127 preceding tokens, so each arriving frame is processed with bounded computation independent of stream length. Training uses the same causal mask. Rotary positions are scaled by $50/f$ at input rate f so the 128-token context covers the same physical duration across rates. During streaming inference, each layer maintains a ring KV cache and processes one arriving frame per forward pass. Both the attention path and the feature derivatives use only observations available at time t, enforcing zero lookahead.

Training. We train the retargeter on the accepted expert rollouts with

$$\mathcal{L} = \mathcal{L}_{\text{dist}} + 0.1\mathcal{L}_v + 0.01\mathcal{L}_{\text{pen}} + 0.1\mathcal{L}_{\text{slip}}. \tag{4}$$

where m_t^e is the accepted simulator state expressed in the student's 35-D output contract. The distillation term $\mathcal{L}_{\text{dist}}$ is a per-dimension normalized mean-squared error between $\widehat{m}_t$ and m_t^e. The velocity term $\mathcal{L}_v$ matches joint and reconstructed-root motion, $\mathcal{L}_{\text{pen}}$ penalizes foot geometry below the ground, and $\mathcal{L}_{\text{slip}}$ matches horizontal foot speeds on support frames inferred from the expert rollout. To support real-time teleoperation at different downstream input rates, we sample causal windows at $f \in \{30, 50, 90\}$ Hz and align human observations with expert labels at the same physical timestamps. Only the input history is corrupted with temporally correlated pose noise and one- or two-frame causal holds to model imperfect or delayed observations. The prediction is passed directly to the downstream whole-body controller as its motion reference.

4 Experiments

To assess whether SMR preserves human motion semantics while producing references the robot can execute, we evaluate it on the Unitree G1. We examine reference quality and closed-loop execution, alongside further analyses and qualitative streaming teleoperation on hardware. Implementation details are provided in the supplementary material.

4.1 Expert and student reference quality

We ask whether tracking robot-aware task-space semantics yields more faithful references than tracking a prescribed GMR trajectory, and whether SMR retains that fidelity after causal distillation (Table 3). The GMR-conditioned RL expert tracks GMR-derived trajectories for each motion cluster, and its accepted state rollouts train the GMR-Expert Student. The GMR-supervised Student

Table 3: Reference quality on the test set. FK and Sem. are world-frame and root-relative landmark errors. Root is mean pelvis position error, and End. is final pelvis error. Limit and Coll. are interval percentages with joint-limit violations and self-collisions. Slip is support-foot sliding speed.

Reference source	FK (m) ↓	Sem. (m) ↓	Root (m) ↓	End. (m) ↓	Limit (%) ↓	Coll. (%) ↓	Slip (m/s) ↓
GMR	0.077	0.029	0.064	0.075	**0.00**	3.53	0.045
NMR	0.163	0.032	0.156	0.387	**0.00**	2.75	0.050
GMR-supervised Student	0.063	0.030	0.051	0.044	29.59	2.90	0.040
GMR-Expert Student	0.068	0.034	0.053	0.045	0.08	**1.78**	0.033
SMR (Ours)	**0.027**	**0.020**	**0.016**	**0.019**	0.01	1.79	**0.030**
GMR-conditioned Expert	0.082	0.038	0.064	0.075	**0.02**	**0.01**	0.024
Semantic Expert (Ours)	**0.028**	**0.023**	**0.013**	**0.011**	**0.02**	**0.01**	**0.017**

Table 4: Closed-loop evaluation on 400 motions using the same frozen whole-body controller. SMR achieves the highest success rate and lowest tracking and human-target errors, demonstrating more reliable execution and better preservation of human motion semantics.

	GMR	NMR	GMR-supervised Student	GMR-Expert Student	SMR (Ours)
Success (%) ↑	90.0	92.0	91.5	93.0	**97.0**
MPJPE (m) ↓	0.060	0.057	0.055	0.051	**0.039**
E-MPJPE (m) ↓	0.115	0.111	0.111	0.112	**0.080**

learns from GMR references, while SMR learns from accepted semantic-expert rollouts. All sources use the same test motions. Physics-based tracking of GMR-derived trajectories greatly reduces self-collision, but leaves landmark error close to GMR's. The semantic expert achieves similarly low collision with much lower landmark error, supporting our choice to let the expert determine the robot trajectory from the semantic task. SMR slightly improves landmark accuracy over the semantic expert and has lower error than the GMR-based students. Learning from many accepted rollouts may let it generalize to motions that individual experts track less accurately. At inference, SMR has no simulator feedback to explicitly correct collisions or slipping contacts, which may explain its worse self-collision and foot-slip metrics.

4.2 Closed-loop execution

We next evaluate whether the generated references can be executed while preserving the intended human motion. Table 4 compares five methods on the same 400 motions with a frozen MimicLite controller in mjlab (Zakka et al., 2026). The Unitree G1 asset, scene, actuators, initialization, and reference offsets are fixed across methods. Success requires completion without absolute pelvis-height error above 0.30 m or an absolute projected-gravity vertical-component difference above 0.8. Non-finite or incomplete rollouts fail. MPJPE measures tracking of the robot reference, and E-MPJPE measures how closely the executed wrists and ankles follow the human targets. Both errors include all 400 motions, including failures. Our final model completes 97.0% of motions, compared with 90.0% for GMR, 92.0% for NMR, and 91.5–93.0% for the GMR-based students. It also achieves the lowest MPJPE (0.039 m) and E-MPJPE (0.080 m). Together, higher completion and lower execution errors show that SMR's streaming references, distilled from dynamics-grounded expert rollouts, serve as effective targets for the frozen WBC to drive Unitree G1 in simulation while retaining the intended human motion. Figure 3 shows synchronized snapshots of different SMPL motion streams retargeted by GMR, NMR, and SMR. In the selected crouching, kneeling, and floor-level poses, SMR more closely matches the human motion, complementing the reference quality and closed-loop evidence for semantic fidelity and executability.

4.3 Ablations

To assess the effects of data and model scale, we train the Base student with different fractions of the training set and different widths (Figure 4). Held-out error decreases markedly up to 50% of the training data, but doubling the data to 100% brings little further gain. Under the same training budget, increasing the model beyond 12.7M parameters slightly worsens held-out error despite lower training loss, suggesting that the added capacity may overfit at this data scale. We therefore use

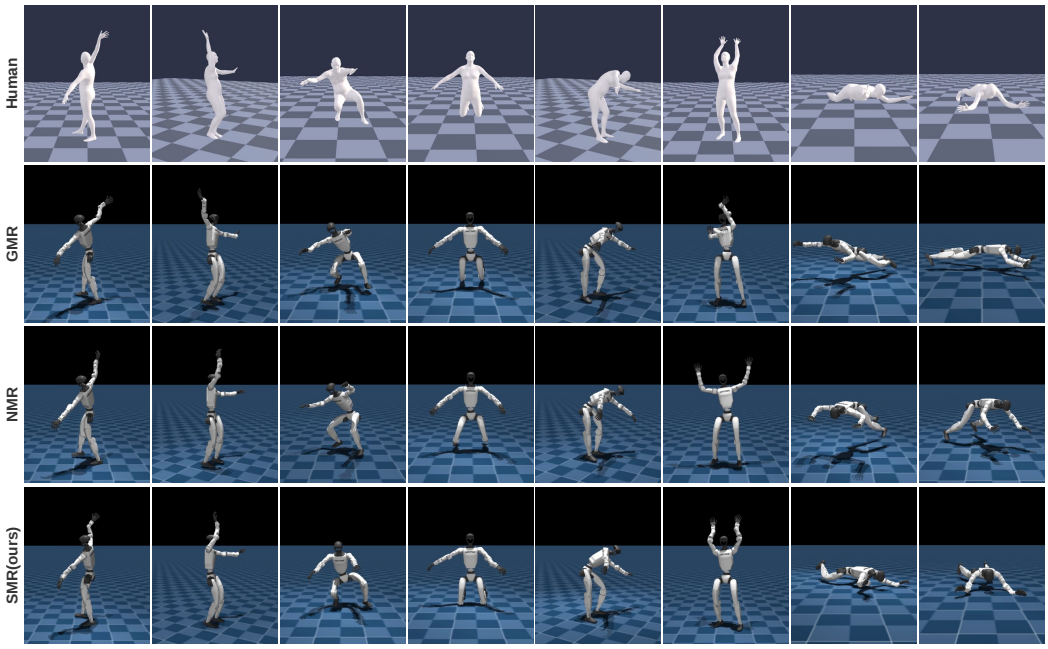


Figure 3: Qualitative comparison of Unitree G1 reference poses at matched timestamps from eight human motions. The rows show source SMPL poses and the corresponding references produced by GMR, NMR, and SMR. Compared with GMR and NMR, SMR better preserves the source whole-body posture, especially in the crouching, kneeling, and floor-level examples.

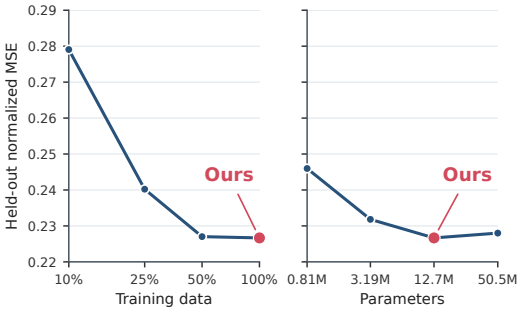


Figure 4: Data and model scaling. Gains diminish beyond 50% training data, while the 12.7M model achieves the lowest held-out error.

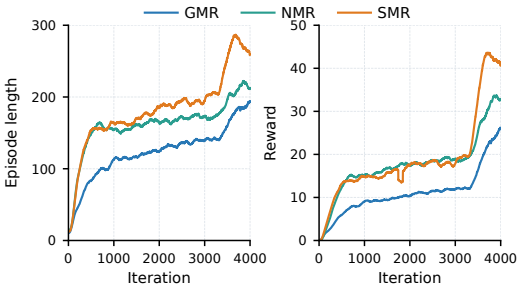


Figure 5: MimicLite training with different motion references. SMR achieves longer episodes and higher rewards by 4,000 iterations.

Table 5: Matched ablation of auxiliary losses (Aux.: velocity, foot penetration, support-foot slip) and input augmentation (Aug.: correlated pose noise and 1–2-frame causal holds). Sem. and Slip use the metrics defined in Table 3. MPJPE, E-MPJPE, and Success are measured with the frozen WBC. Lower is better for errors, higher for Success. Best values are bold.

Variant	Aux.	Aug.	MSE (norm.) ↓	Sem. (m) ↓	Slip (m/s) ↓	MPJPE (m) ↓	E-MPJPE (m) ↓	Success (%) ↑
Base	—	—	0.24580	0.0288	0.0512	0.0428	0.0839	96.75
Aux-only	✓	—	**0.24521**	**0.0276**	**0.0325**	0.0412	0.0814	96.75
Aug-only	—	✓	0.24642	0.0287	0.0501	0.0411	0.0819	96.75
SMR (Ours)	✓	✓	0.24863	0.0283	0.0329	**0.0392**	**0.0795**	**97.00**

the 12.7M-parameter architecture for our final model. Figure 5 compares all 4,000 iterations of MimicLite training on GMR, NMR, and neural references. All runs use 8,192 environments per GPU on eight GPUs. The neural run reaches higher reward and longer episodes by iteration 4,000,

Table 6: Warm per-frame latency of the GMR (Araujo et al., 2025) and four SMR model sizes. Compared with Small and Medium, Large has substantially more parameters but only slightly higher latency. Together with its lowest held-out error in Figure 4, this makes Large a favorable balance of accuracy and speed among the tested sizes. Bold marks the lowest latency.

		SMR			
	GMR	Small	Medium	Large (Ours)	XLarge
Parameters (M)	–	0.81	3.19	12.7	50.5
Median (ms/frame) ↓	4.87	**0.34**	0.35	0.37	0.58



Figure 6: Streaming Unitree G1 teleoperation driven by GMR (top two rows) and SMR (bottom two rows) through MimicLite. Each row shows four frames from an independent recording in temporal order, covering arm motion or kneeling.

although the reward curves cross earlier in training. Reference-processing protocols differ across methods, so these curves describe training on each reference distribution rather than a controlled comparison of retargeters alone. Table 5 compares students trained with and without auxiliary losses and input augmentation. Adding auxiliary losses reduces reference foot slip and lowers both frozen-WBC tracking errors in each matched pair, while success changes little. Augmentation also lowers both WBC errors, although its effects on clean-input reference metrics are mixed.

4.4 Streaming Real-Robot Deployment

At 50 Hz, SMR's model step takes 0.37 ms per frame, compared with 4.87 ms for GMR's iterative solve (Table 6). The values average three warm-run medians. XLarge was measured on a separate GPU. SMR also achieves higher completion and lower execution errors under the frozen WBC (Table 4), combining faster reference generation with better closed-loop simulation performance. Figure 6 shows streaming deployment on a physical Unitree G1. In both cases, a MimicLite policy controls the robot using motion references from the retargeter. GMR's kinematic optimization does not enforce dynamically stable execution: the GMR-driven robot becomes unstable while kneeling, and its raised-arm motion falls out of phase with the operator. With SMR references, MimicLite produces smooth, stable motion throughout the demonstrated teleoperation sequences.

5 Conclusion

We presented SMR, which uses semantic-conditioned physics experts to produce supervision for a strictly causal neural retargeter. The experts track robot-aware task-space commands directly, with IK used only to initialize episodes. Accepted simulator state rollouts, executed under the Unitree G1 contact and actuator models, train the student to generate Unitree G1 motion references from current and past human observations. This separates the semantic task from its physical realization while preserving a streaming input contract.

In the shared 400-motion Unitree G1 evaluation, SMR completes 97.0% of motions under a frozen whole-body controller, compared with 90.0% for GMR and 92.0% for NMR, while lowering MPJPE and E-MPJPE (Table 4). The semantic expert also achieves lower landmark error and fewer self-collisions than direct optimization-based GMR (Table 3). These results support using accepted expert rollouts to train streaming references under the tested simulator and controller.

6 Limitations and future work

Our evaluation is centered on the Unitree G1 under a single simulation and control setting. The semantic transport, expert policies, and closed-loop evaluation are tied to this embodiment and controller, so transfer to other morphologies, actuation, contact conditions, controllers, and hardware remains untested. SMR also faces a causal distillation gap: the student sees only current and past human observations, whereas the physics expert can access robot state and limited future semantic context. Its performance ceiling may therefore be lower than that of a bidirectional model with access to the complete motion sequence, which is the cost of zero-lookahead streaming inference. Future work should evaluate broader embodiments, controllers, and hardware, develop uncertainty-aware causal predictors, and extend the semantic interface to robot-to-human retargeting.

Ethics statement

This work uses existing human-motion datasets for human-to-humanoid motion retargeting. Source motions and derived assets are used and redistributed under their respective data and model licenses. Human-motion and teleoperation recordings are treated as potentially identifiable data, and their collection, storage, and sharing follow participant consent and privacy protections. Generated robot motion references can lead to falls or collisions during execution. Acceptance of an expert rollout in simulation does not establish hardware safety, and the qualitative demonstration does not establish safety across robots, environments, or human interactions. Physical deployment therefore requires separate risk assessment, supervised testing, and appropriate safeguards.

Reproducibility statement

Section 3 defines the semantic interface, physics-expert objective, and causal student's input, output, and training objective. Appendix A specifies semantic construction, motion grouping, expert training, and rollout acceptance. Appendix B specifies the student architecture, training recipe, loss terms, and streaming inference procedure. Section 4 describes the reported experimental comparisons under the specified robot and simulation conditions.

References

Kfir Aberman, Peizhuo Li, Dani Lischinski, Olga Sorkine-Hornung, Daniel Cohen-Or, and Baoquan Chen. Skeleton-aware networks for deep motion retargeting. *ACM Transactions on Graphics*, 39 (4), 2020. doi: 10.1145/3386569.3392462.

Joao Pedro Araujo, Yanjie Ze, Pei Xu, Jiajun Wu, and C. Karen Liu. Retargeting matters: General motion retargeting for humanoid motion tracking. *arXiv preprint arXiv:2510.02252*, 2025.

Bones Studio. BONES-SEED: Skeletal everyday embodiment dataset. Hugging Face dataset release, 2026. URL https://huggingface.co/datasets/bones-studio/seed.

Hanyang Cao, Yuetong Fang, Taesoo Kwon, Runyi Yu, Ji Ma, Jing Tan, Yangchen Zhou, Baoze Du, Yi Gu, Yukang Gao, Ruoli Dai, Lei Han, and Renjing Xu. Unified motion retargeting for humanoids with learned point cloud correspondence. *arXiv preprint arXiv:2609.02134*, 2026.

Xingyu Chen, Hanyu Wu, Sikai Wu, Mingliang Zhou, Diyun Xiang, Haodong Zhang, Yangchen Zhou, Yukang Gao, Yi Gu, and Renjing Xu. A scalable whole-body motion transfer via implicit kinodynamic motion retargeting. *arXiv preprint arXiv:2509.15443*, 2025a.

Zixuan Chen, Mazeyu Ji, Xuxin Cheng, Xuanbin Peng, Xue Bin Peng, and Xiaolong Wang. GMT: General motion tracking for humanoid whole-body control. *arXiv preprint arXiv:2506.14770*, 2025b.

Sungjoon Choi, Matthew K. X. J. Pan, and Joohyung Kim. Nonparametric motion retargeting for humanoid robots on shared latent space. In *Robotics: Science and Systems*, 2020. doi: 10.15607/RSS.2020.XVI.071.

Kourosh Darvish, Yeshasvi Tirupachuri, Giulio Romualdi, Lorenzo Rapetti, Diego Ferigo, Francisco Javier Andrade Chavez, and Daniele Pucci. Whole-body geometric retargeting for humanoid robots. In *IEEE-RAS 19th International Conference on Humanoid Robots*, pp. 679–686, 2019. doi: 10.1109/HUMANOIDS43949.2019.9035059.

Victor Dhedin, Ilyass Taouil, Shafeef Omar, Dian Yu, Kun Tao, Angela Dai, and Majid Khadiv. DynaRetarget: Dynamically-feasible retargeting using sampling-based trajectory optimization. *IEEE Robotics and Automation Letters*, pp. 1–8, 2026. doi: 10.1109/LRA.2026.3730386.

Michael Gleicher. Retargetting motion to new characters. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 33–42, 1998. doi: 10.1145/280814.280820.

Ruben Grandia, Farbod Farshidian, Espen Knoop, Christian Schumacher, Marco Hutter, and Moritz Bächer. DOC: Differentiable optimal control for retargeting motions onto legged robots. *ACM Transactions on Graphics*, 42(4):1–14, 2023. doi: 10.1145/3592454.

Jinrui Han, Weiji Xie, Jiakun Zheng, Jiyuan Shi, Weinan Zhang, Ting Xiao, and Chenjia Bai. KungfuBot2: Learning versatile motion skills for humanoid whole-body control. *arXiv preprint arXiv:2509.16638*, 2025.

Félix G. Harvey, Mike Yurick, Derek Nowrouzezahrai, and Christopher Pal. Robust motion inbetweening. *ACM Transactions on Graphics*, 39(4), 2020. doi: 10.1145/3386569.3392480.

Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbab, Chaoyi Pan, Zeji Yi, Guannan Qu, Kris Kitani, Jessica Hodgins, Linxi Fan, Yuke Zhu, Changliu Liu, and Guanya Shi. ASAP: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. *arXiv preprint arXiv:2502.01143*, 2025.

Tianchen Huang, Feiyang Yuan, Junchi Gu, Shurui Fang, Xiaohu Zhang, Yu Wang, Wei Gao, and Shiwu Zhang. Human2Humanoid: Physics-aware cross-morphology motion retargeting for humanoid robots. *arXiv preprint arXiv:2606.03476*, 2026.

Kyungmin Lee, Sibeen Kim, Youngdo Lee, Minho Park, Hyunseung Kim, Dongyoon Hwang, Donghu Kim, Hojoon Lee, and Jaegul Choo. PHUMA: Physically reliable humanoid locomotion dataset. *arXiv preprint arXiv:2510.26236*, 2025.

Sunmin Lee, Taeho Kang, Jungnam Park, Jehee Lee, and Jungdam Won. SAME: Skeleton-agnostic motion embedding for character animation. In *SIGGRAPH Asia Conference Papers*, pp. 1–11, 2023. doi: 10.1145/3610548.3618206.

Yuhan Li, Peiyuan Zhi, Yunshen Wang, Tengyu Liu, Sixu Yan, Wenyu Liu, Xinggang Wang, Baoxiong Jia, and Siyuan Huang. OmniTrack: General motion tracking via physics-consistent reference. *arXiv preprint arXiv:2602.23832*, 2026.

Qiayuan Liao, Takara E. Truong, Xiaoyu Huang, Yuman Gao, Guy Tevet, Koushil Sreenath, and C. Karen Liu. BeyondMimic: From motion tracking to versatile humanoid control via guided diffusion. *Science Robotics*, 11(117), 2026. doi: 10.1126/scirobotics.adx8924.

Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J. Black. SMPL: A skinned multi-person linear model. *ACM Transactions on Graphics*, 34(6):1–16, 2015. doi: 10.1145/2816795.2818013.

Zhengyi Luo, Jinkun Cao, Alexander Winkler, Kris Kitani, and Weipeng Xu. Perpetual humanoid control for real-time simulated avatars. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10895–10904, 2023. doi: 10.1109/ICCV51070.2023.01000.

Zhengyi Luo, Ye Yuan, Tingwu Wang, Chenran Li, Fernando Castañeda, Sirui Chen, Zi-Ang Cao, Jiefeng Li, David Minor, Qingwei Ben, Jinhyung Park, David Sami, Zi Wang, Xingye Da, Runyu Ding, Cyrus Hogg, Lina Song, Edy Lim, Eugene Jeong, Tairan He, Haoru Xue, Wenli Xiao, Simon Yuen, Jan Kautz, Yan Chang, Umar Iqbal, Linxi Fan, and Yuke Zhu. SONIC: Supersizing motion tracking for natural humanoid whole-body control. *Science Robotics*, 11(117), 2026. doi: 10.1126/scirobotics.aed4592.

Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. AMASS: Archive of motion capture as surface shapes. In *IEEE/CVF International Conference on Computer Vision*, pp. 5441–5450, 2019. doi: 10.1109/ICCV.2019.00554.

Ian Mason, Sebastian Starke, and Taku Komura. Real-time style modelling of human locomotion via feature-wise transformations and local motion phases. *arXiv preprint arXiv:2201.04439*, 2022. doi: 10.48550/arXiv.2201.04439.

David Müller, Agon Serifi, Sammy Christen, Ruben Grandia, Espen Knoop, and Moritz Bächer. ReActor: Reinforcement learning for physics-aware motion retargeting. *ACM Transactions on Graphics*, 45(4):1–11, 2026. doi: 10.1145/3811378.

Luigi Penco, Brice Clément, Valerio Modugno, Enrico Mingo Hoffman, Gabriele Nava, Daniele Pucci, Nikos G. Tsagarakis, Jean-Baptiste Mouret, and Serena Ivaldi. Robust real-time whole-body motion retargeting from human to humanoid. In *IEEE-RAS 18th International Conference on Humanoid Robots*, pp. 425–432, 2018. doi: 10.1109/HUMANOIDS.2018.8624943.

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. Deepmimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions on Graphics*, 37(4), 2018. doi: 10.1145/3197517.3201311.

Zoran Popović and Andrew Witkin. Physically based motion transformation. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques*, pp. 11–20, 1999. doi: 10.1145/311535.311536.

Zekun Qi, Xuchuan Chen, Dairu Liu, Chenghuai Lin, Yunrui Lian, Sikai Liang, Zhikai Zhang, Yu Guan, Jilong Wang, Wenyao Zhang, Xinqiang Yu, He Wang, and Li Yi. Humanoid-GPT: Scaling data and structure for zero-shot motion tracking. *arXiv preprint arXiv:2606.03985*, 2026.

Robotparty Lab. MimicLite: Efficient and effective general humanoid motion tracking. Technical report, 2026. URL https://github.com/Roboparty/MimicLite.

Javier Romero, Dimitrios Tzionas, and Michael J. Black. Embodied hands: Modeling and capturing hands and bodies together. *ACM Transactions on Graphics*, 36(6):245:1–245:17, 2017. doi: 10.1145/3130800.3130883.

Constant Roux, Ludovic De Matteïs, Armand Jordana, Valentin Guillet, Nicolas Mansard, Olivier Stasse, and Philippe Souères. Direct dynamic retargeting for humanoid imitation learning from videos, 2026.

Quentin Rouxel, Kai Yuan, Ruoshi Wen, and Zhibin Li. Multicontact motion retargeting using whole-body optimization of full kinematics and sequential force equilibrium. *IEEE/ASME Transactions on Mechatronics*, 27(5):4188–4198, 2022. doi: 10.1109/TMECH.2022.3152844.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Sebastian Starke, Ian Mason, and Taku Komura. DeepPhase: Periodic autoencoders for learning motion phase manifolds. *ACM Transactions on Graphics*, 41(4):1–13, 2022. doi: 10.1145/3528223.3530178.

Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. RoFormer: Enhanced transformer with rotary position embedding. *Neurocomputing*, 568:127063, 2024. doi: 10.1016/j.neucom.2023.127063.

Seyoon Tak and Hyeong-Seok Ko. A physically-based motion retargeting filter. *ACM Transactions on Graphics*, 24(1):98–117, 2005. doi: 10.1145/1037957.1037963.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems*, volume 30, 2017. URL https://papers.nips.cc/paper/2017/hash/3f5ee243547dee91fbd053c1c4a845aa-Abstract.html.

Ruben Villegas, Jimei Yang, Duygu Ceylan, and Honglak Lee. Neural kinematic networks for unsupervised motion retargetting. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8639–8648, 2018. doi: 10.1109/CVPR.2018.00901.

Hanwen Wang, Qiayuan Liao, Bike Zhang, Kunzhao Ren, Koushil Sreenath, and Xiaobin Xiong. SPARK: Skeleton-parameter aligned retargeting on humanoid robots with kinodynamic trajectory optimization. *arXiv preprint arXiv:2603.11480*, 2026.

Satoshi Yagi, Mitsunori Tada, Eiji Uchibe, Suguru Kanoga, Takamitsu Matsubara, and Jun Morimoto. Unsupervised neural motion retargeting for humanoid teleoperation. *arXiv preprint arXiv:2406.00727*, 2024.

Yashuai Yan, Esteve Valls Mascaro, and Dongheui Lee. ImitationNet: Unsupervised human-to-robot motion retargeting via shared latent space. In *IEEE-RAS 22nd International Conference on Humanoid Robots*, pp. 1–8, 2023. doi: 10.1109/Humanoids57100.2023.10375150.

Lujie Yang, Xiaoyu Huang, Zhen Wu, Angjoo Kanazawa, Pieter Abbeel, Carmelo Sferrazza, C. Karen Liu, Rocky Duan, and Guanya Shi. OmniRetarget: Interaction-preserving data generation for humanoid whole-body loco-manipulation and scene interaction. *arXiv preprint arXiv:2509.26633*, 2025.

Kevin Zakka, Qiayuan Liao, Brent Yi, Louis Le Lay, Koushil Sreenath, and Pieter Abbeel. mjlab: A lightweight framework for GPU-accelerated robot learning. *arXiv preprint arXiv:2601.22074*, 2026. doi: 10.48550/arXiv.2601.22074.

Xiaoyu Zhang, Steven Haener, Varun Madabushi, and Maegan Tucker. Kinodynamic motion retargeting for humanoid locomotion via multi-contact whole-body trajectory optimization. *arXiv preprint arXiv:2603.09956*, 2026.

Zhikai Zhang, Jun Guo, Chao Chen, Jilong Wang, Chenghuai Lin, Yunrui Lian, Han Xue, Zhenrong Wang, Maoqi Liu, Jiangran Lyu, Huaping Liu, He Wang, and Li Yi. Track any motions under any disturbances. *arXiv preprint arXiv:2509.13833*, 2025.

Qingrui Zhao, Kaiyue Yang, Xiyu Wang, Shiqi Zhao, Yi Lu, Xinfang Zhang, Qiu Shen, Xiao-Xiao Long, and Xun Cao. Make tracking easy: Neural motion retargeting for humanoid whole-body control. *arXiv preprint arXiv:2603.22201*, 2026.

APPENDIX

A Semantic-conditioned physics experts 14

A.1 Robot-aware semantic command details 14

A.2 Motion grouping and expert assignment 14

A.3 Physics expert training 15

A.4 Expert rollout supervision 16

B Neural retargeting 16

A SEMANTIC-CONDITIONED PHYSICS EXPERTS

This section specifies the construction of the robot-aware semantic command, the offline expert population, and the accepted simulator rollouts used as student supervision.

A.1 ROBOT-AWARE SEMANTIC COMMAND DETAILS

We construct the semantic command $c_t = \Psi_{\mathcal{R}}(H_t)$ from a canonical SMPL-H motion H_t at 50 Hz. Forward kinematics uses the neutral SMPL-H skeleton after conversion from the source Y-up convention to the robot Z-up convention. The resulting sparse skeleton contains 17 joints, including both toes. The toes are used only to determine the vertical alignment and are excluded from the landmark set $\mathcal{K}$. The 15 landmarks in $\mathcal{K}$ are the pelvis, neck, head, and the hips, knees, ankles, shoulders, elbows, and wrists on both sides. Each landmark has a position, orientation, linear velocity, and angular velocity.

Table 7: The predefined landmark set $\mathcal{K}$.

Trunk	Leg (per side)	Arm (per side)
Pelvis, neck, head	Hip, knee, ankle	Shoulder, elbow, wrist

Morphology-aware transport. All transport parameters are fixed by the neutral SMPL-H rest pose and the Unitree G1 neutral configuration. Let $\ell_{\mathcal{R}}$ and ℓ_{S} be the sums of the hip–knee and knee–ankle rest lengths over both legs of the Unitree G1 and SMPL-H skeleton, respectively. We scale the three coordinates of the human root by $s_{\mathcal{R}} = \ell_{\mathcal{R}}/\ell_{\mathrm{S}}$ and anchor this scaling at the world origin. A constant vertical offset then aligns the lowest toe of the scaled human rest pose with the corresponding robot rest pose. Cropping a sequence therefore does not change the transported pose at an unchanged timestamp. The remaining landmarks are assembled along the human segment directions with lengths from the robot rest pose rather than by pointwise scaling. The neck is attached to a fixed waist anchor, and the head is placed using the neutral Unitree G1 head geometry. For every landmark, including the root, we map the SMPL orientation into a reference frame fixed at the Unitree G1 neutral configuration. This transport produces task-space landmark targets only. It performs no inverse-kinematics solve and produces no robot joint trajectory.

Command fields. For each landmark, the command contains the transported position and orientation and their linear and angular velocities. All four fields are computed after transport and archived in the world frame at 50 Hz. The velocities are finite differences of the transported quantities and are used by the expert critic and tracking rewards, not as additional student inputs. Section A.3 specifies the frame conventions under which the expert observes this command.

A.2 MOTION GROUPING AND EXPERT ASSIGNMENT

We partition the heterogeneous motion corpus into offline training pools before fitting the physics experts. This subsection details the Harmonic Motion Embedding (HME) grouping summarized in Section 3.2. The grouping input is the robot-aware semantic command sequence at 50 Hz. Motions

shorter than the 101-frame, two-second window are padded by holding the terminal pose, so no motion is removed by the windowing step.

Periodic autoencoder. A Periodic Autoencoder (PAE) (Starke et al., 2022) encodes each 101-frame window with a convolutional encoder and 64 phase channels. A Fourier decomposition of each phase channel provides a harmonic amplitude and a power-weighted frequency. We train one PAE per clustered source with reconstruction error as the only objective, using 50,000 AdamW steps, learning rate 3×10^{-4}, weight decay 10^{-4}, batch size 512, 16 windows per motion, a 2% validation split, and random seed 0. We slide the same window with a stride of 0.2 s over each complete motion. Concatenating the mean and standard deviation of the 64 amplitudes and 64 frequencies over these windows yields one 256-dimensional descriptor per motion. We do not segment a motion into separate descriptors.

Clustering and expert assignment. We standardize each descriptor dimension and apply Mini-Batch K-means separately within each large source, with $K = 256$ for BONES-SEED and $K = 32$ for 100STYLE. Each nonempty cluster forms the training pool of one cluster expert. The smaller specialized sources bypass clustering and form one single-motion expert per motion. Table 8 summarizes the resulting Unitree G1 roster, which contains 465 experts for 145,970 candidate motions. This assignment determines only the offline expert training pools. The student does not receive cluster or expert identity.

Table 8: The expert assignment on Unitree G1, with total motion duration per source.

Source	Assignment	Experts	Motions	Hours
BONES-SEED (Bones Studio, 2026)	Clustered	256	142,220	288.29
100STYLE (Mason et al., 2022)	Clustered	32	3,573	18.65
LAFAN1 (Harvey et al., 2020)	Single-motion	77	77	4.60
AMASS (Mahmood et al., 2019)	Single-motion	22	22	0.05
Five smaller specialized sets	Single-motion	78	78	0.81
Total		465	145,970	312.40

A.3 PHYSICS EXPERT TRAINING

We train each expert as a closed-loop PPO policy in the target-robot simulator. The semantic command window contains eight samples at 50 Hz: the current sample, three samples from the preceding 160 ms, and four samples from the following 80 ms. The proprioceptive history uses offsets $\{0, 1, 2, 3, 4, 8, 16\}$ control steps in the past. Within the command window, the actor observes only position and orientation fields. Its full observation also includes the sparse proprioceptive history and the three most recent policy actions. Landmark velocities are reserved for the privileged critic and for tracking rewards. The policy outputs a 29-dimensional joint-position target at 50 Hz. Each of the eight tracking terms uses the kernel $\exp(-\bar{e}/\sigma)$ with its own width. Position tracking covers the root and every landmark of $\mathcal{K}$ except the head, and the five selected orientation bodies are the torso, both wrists, and both ankles. Root displacement and orientation residuals are encoded in the current robot heading-local frame for the actor, while root tracking errors are evaluated in the world frame and body-landmark errors in the current heading-local frame. A curriculum changes the tracking widths from 1.25 to 0.5 times their base values during the first 1,050 iterations, while the weights in Table 2 remain fixed. An episode terminates when a root or landmark error exceeds its threshold for consecutive control steps or when the motion ends.

Optimization, episodes, and reset states. The actor and critic are MLPs with Mish activations and running observation normalization. The actor has three 1,024-unit hidden layers, and the critic has hidden layers of 1,024, 512, and 256 units. PPO uses a clip ratio of 0.2, entropy coefficient 0.012, AdamW learning rate 5×10^{-5} with an adaptive KL schedule, discount 0.99, and $\lambda = 0.95$. We train with 16,384 parallel environments for 1,500 iterations for single-motion experts and 4,000 iterations for cluster experts. After a failure, an episode restarts with probability 0.8 from 25–125 frames earlier, with the start frame selected by an adaptive difficulty sampler. The initial robot state is obtained from offline per-frame differential IK using landmark positions as primary targets. The

IK solve uses neither contact reasoning nor cross-frame continuity. It supplies only the simulator reset state and is not a tracking target for the expert.

A.4 Expert rollout supervision

After training, each expert deterministically replays every assigned motion once at 50 Hz with its mean policy, starting from the same offline IK reset state used during training. The recorded simulator state at frame t is paired with the semantic command at frame t, which is the frame used by the tracking reward. A rollout is accepted only if it reaches the final frame without termination and passes every check in Table 9. We pair each accepted rollout with its source human motion and assign motions to train, validation, or test splits in a 90/5/5 ratio by motion group. No motion contributes frames to more than one split. Experts are trained before this split, so the holdout evaluates student distillation on unseen motion groups rather than expert generalization. Rejected candidates remain in the candidate-level acceptance denominator and are excluded from positive student labels. On Unitree G1, the semantic expert release accepts 105,634 of 145,970 candidate motions, while the GMR-conditioned release accepts 106,019.

Table 9: Label-production acceptance gates.

Check	Threshold
Completion	full target length, no termination
Joint limits	violation fraction ≤ 0.001, margin ≥ -0.02 rad
Joint speed	≤ 40 rad/s
Effort saturation	fraction ≤ 0.05 at $0.95\times$ limit
Self-collision	contact fraction ≤ 0.01 above 1 N
Root position	p95 error ≤ 0.20 m
Body position	max p95 error ≤ 0.20 m
Root orientation	p95 error $\leq 40^\circ$
Selected orientations	p95 error $\leq 40^\circ$
Persistent orientation	runs above 60° shorter than 25 frames

B Neural retargeting

This section specifies the causal student's input and output contracts, streaming state, and training recipe.

Input and output contract. At time t, the input is the 79-D human pose representation from items (i)–(iv) and the 6-D human-root twist from item (v) in Section 3.3. The twist is a backward difference computed at the native observation rate and is zero at the first frame of each session. The student receives no simulator state, expert action or identity, or future human observation. The output is the 35-D robot motion reference $\widehat{m}_t = (\widehat{q}_t, \widehat{\Delta p}_t, \widehat{\Delta r}_t)$: 29 Unitree G1 joint positions, a 3-D pelvis translation offset, and a 3-D relative pelvis rotation. The two offsets are expressed in the current human-pelvis frame and are composed with the current human pelvis pose to reconstruct the robot root. This representation avoids integrating a predicted root velocity, but the reconstructed root can still inherit human-root drift, heading noise, and ground-height error. The student produces a motion reference, not a torque or action. The downstream whole-body controller executes this reference.

Architecture and streaming. The deployed model has four pre-normalized Transformer blocks of width 512 and 16 attention heads. Each head has dimension 32, and the feed-forward layer is four times the model width with GELU activations. Rotary embeddings use base 10,000, dropout is disabled, and the model has approximately 12.7 M parameters. RoPE positions are scaled by $50/f$ at input rate f, so the 128-token context spans the same physical duration across rates. The causal mask gives each output access only to the current token and the 127 preceding tokens. During streaming inference, each layer maintains a bounded ring KV cache and processes one arriving frame per call. The cache is reset at a session boundary. Input normalization is folded into the deployed graph, and outputs are reported in physical units. Strict causality describes this retargeter interface. The

downstream WBC may use an 80-ms future-reference buffer. This buffer is not retargeter lookahead and is reported separately.

Training recipe. Student-Final uses a single-stage run of 15,000 optimization steps with batch size 32, AdamW, and a cosine learning-rate schedule from 3×10^{-4} to 3×10^{-5}. Each training window contains 256 frames and is sampled at the input rates used in the main text. Positions are linearly interpolated and orientations use shortest-arc interpolation at the resampled timestamps. Human features and root offset labels are recomputed at those timestamps under the causal alignment, so the input at time t does not use an observation after t. With probability 0.5, only the input history is corrupted. The corruption combines temporally correlated pose noise of 2 mm for positions and 0.01 rad for orientations with one- or two-frame causal holds, applied with a per-frame drop probability of 0.06. Offset labels are recomputed from the uncorrupted expert rollout, whereas the root twist is recomputed from the corrupted input stream.