Rhythm: Learning Interactive Whole-Body Control for Dual Humanoids

Hongjin Chen[1,2,*] Wei Zhang[2,*] Pengfei Li[2,3,*] Shihao Ma[1,2] Ke Ma[1,2] Yujie Jin[2] Zijun Xu[1,5]
Xiaohui Wang[1,2] Yupeng Zheng[6] Zining Wang[2] Jieru Zhao[4] Yilun Chen[2] Wenchao Ding[1,2,†]

[1]Fudan University [2]TARS Robotics [3]Tsinghua University [4]Shanghai Jiao Tong University
[5]Shanghai Innovation Institute [6]Institute of Automation, Chinese Academy of Sciences

[*]Equal contribution [†]Corresponding Author

Project Page: https://hoshi-no-ai.github.io/Rhythm/

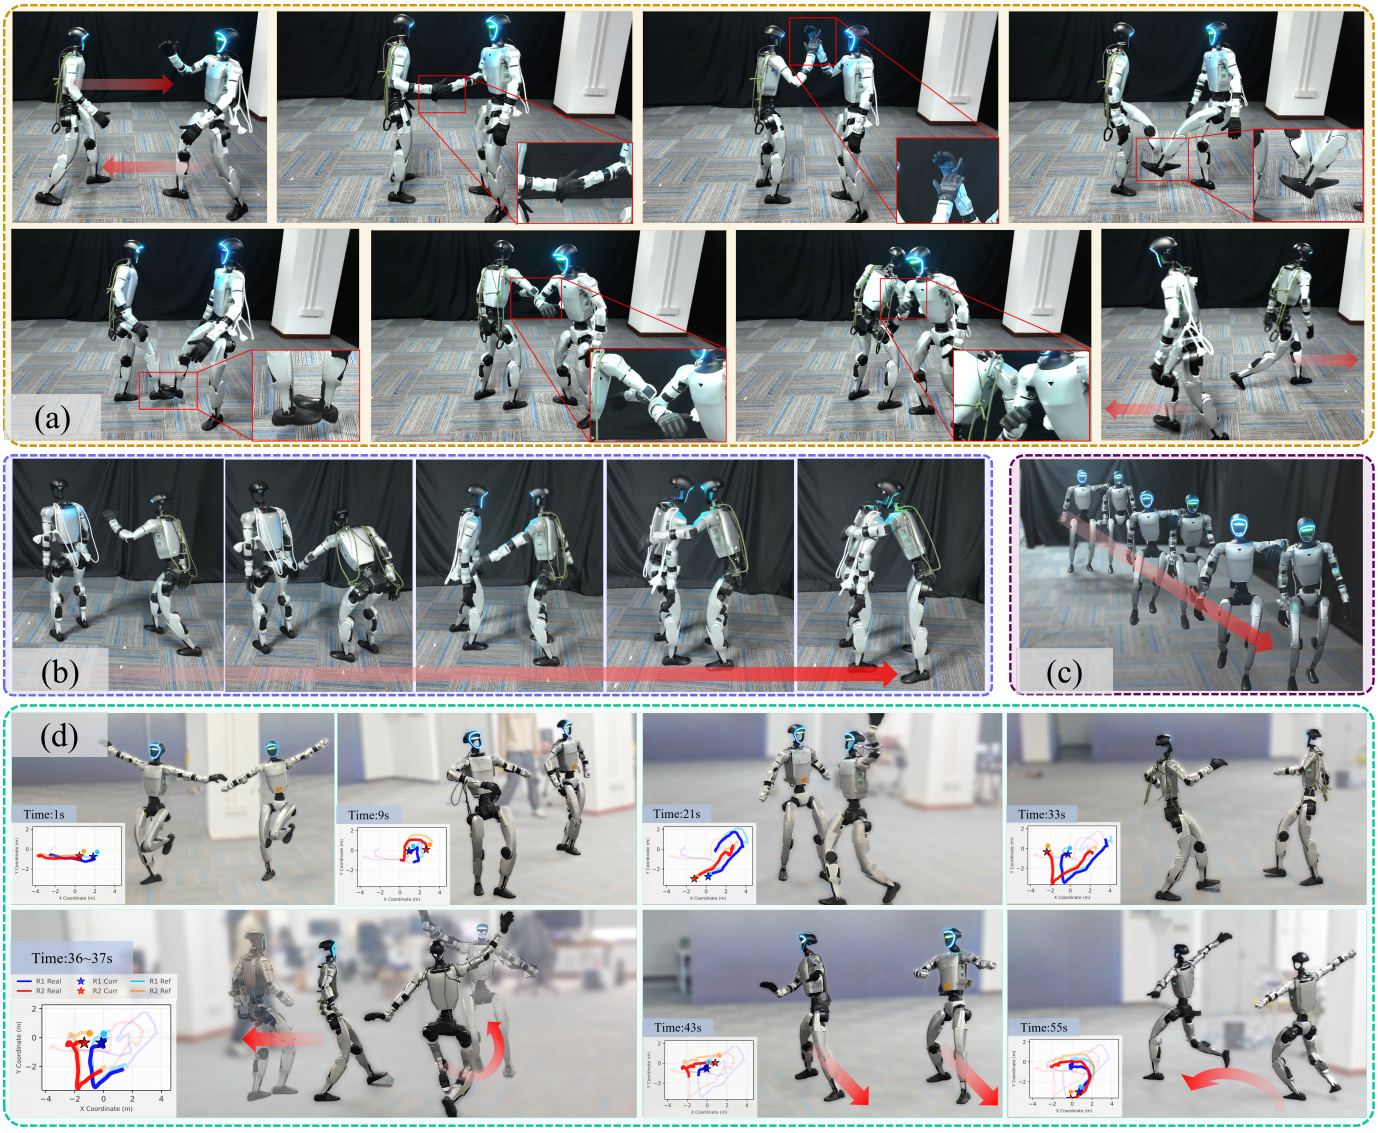

Fig. 1: The proposed framework, **Rhythm**, facilitates a spectrum of humanoid–humanoid interactions. **(a–c) Contact-Rich Interaction:** The method handles interactions ranging from light contact (Greeting) to intensive contact (Hug, Shoulder-to-Shoulder), maintaining fine-grained contact geometry without penetration (shown in the zoomed-in views). **(d) Coordinated Interaction:** The humanoids perform synchronized long-horizon dance (*La La Land*), with trajectories showing consistent spatiotemporal alignment and stable relative positioning over time.

***Abstract*—Realizing interactive whole-body control for multi-humanoid systems is critical for unlocking complex collaborative capabilities in shared environments. Although recent advancements have significantly enhanced the agility of individual robots, bridging the gap to physically coupled multi-humanoid interaction remains challenging, primarily due to**

severe kinematic mismatches and complex contact dynamics. To address this, we introduce Rhythm, the first unified framework enabling real-world deployment of dual-humanoid systems for complex, physically plausible interactions. Our framework integrates three core components: (1) an Interaction-Aware Motion Retargeting (IAMR) module that generates feasible humanoid interaction references from human data; (2) an Interaction-Guided Reinforcement Learning (IGRL) policy that masters coupled dynamics via graph-based rewards; and (3) a real-world deployment system that enables robust transfer of dual-humanoid interaction. Extensive experiments on physical Unitree G1 robots demonstrate that our framework achieves robust interactive whole-body control, successfully transferring diverse behaviors such as hugging and dancing from simulation to reality.

I. Introduction

Humanoid robotics has witnessed rapid evolution, achieving remarkable milestones in single-agent capabilities. Recent research has established a strong foundation in dynamic locomotion [12, 19, 23, 40, 48] and general whole-body control [8, 17, 22, 26, 36, 55], significantly enhancing the agility and robustness of individual robots. However, the broader vision of embodied intelligence necessitates agents that can operate beyond isolation [5, 37]. Realizing multi-agent systems capable of physical collaboration represents a critical next step. Yet, research in this domain remains disproportionately focused on single-robot tasks.

Despite growing interest in multi-agent interaction, current solutions are largely confined to virtual environments or simplified humanoid-object interactions. In computer graphics, physics-based animation has achieved realistic simulations of multi-character interactions [10, 25, 45, 50, 54], yet these methods often prioritize visual fidelity over the strict physical constraints necessary for real-world robotic deployment. In robotics, although human–robot collaboration [9, 20] and competitive sports in structured environments [28, 39, 47] have been explored, these typically involve a compliant human partner or passive objects. Research explicitly targeting multi-humanoid interaction remains scarce, and existing works are predominantly validated only in simulation [29]. Consequently, achieving robust, physically coupled whole-body control on real multi-humanoid hardware remains an unbridged gap in the field.

Two fundamental challenges hinder the Sim-to-Real transition for dual-humanoid systems: (1) the scarcity of feasible interaction data; and (2) the complexity of the training-to-deployment paradigm. First, acquiring high-quality interaction references is non-trivial. Directly transferring *Human-Human Interaction* [38, 41, 45, 46] data to robots introduces severe kinematic conflicts due to morphological differences between humans and humanoids (see Sec. III-A). Standard retargeting methods [3, 30, 49] struggle to preserve both individual motion style and precise interaction geometry, yielding suboptimal motion references. Second, the learning and deployment pipeline presents significant hurdles. Existing tracking policies [26, 44, 52] typically treat agents as isolated entities, failing to model the intricate coupled dynamics essential for close interaction. Meanwhile, a significant disparity exists between the global observability available in simulation and the asynchronous, ego-centric, partially observable reality of real hardware, making the deployment unstable.

To address these challenges, we introduce **Rhythm** (Interactive Whole-Body Control for Dual Humanoids), a unified framework designed to empower dual humanoids to execute complex, physically plausible interactions in real-world scenarios. The framework explicitly models interaction geometry and physical contact to achieve high-fidelity coupled behaviors. First, to resolve kinematic conflicts in reference generation, we introduce **Interaction-Aware Motion Retargeting (IAMR)**. By explicitly modeling interaction geometry and utilizing distance-aware dynamic weighting, this module adaptively balances self-motion fidelity with interaction geometry. The resulting geometrically consistent references serve as a crucial prior, laying the physical foundation for the subsequent training phase. Building upon this, we develop **Interaction-Guided Reinforcement Learning (IGRL)** to master the complex dual-agent dynamics. This module directly leverages the interaction geometry and contact preserved by IAMR through explicit graph-based rewards, guiding agents to learn synchronized and robust behaviors. Finally, to realize dual-humanoid interactive whole-body control in the real world, we implement a relative state estimation and inter-agent synchronization scheme, effectively bridging the gap between global simulation and ego-centric reality.

Fig. 1 illustrates the robustness of our framework in real-world scenarios, successfully handling tasks ranging from intensive contacts to synchronized long-horizon dancing. Our main contributions are summarized as follows:

- We propose **Rhythm**, a unified framework for whole-body dual-humanoid interaction that, to our knowledge, achieves the first successful **robust transfer** of complex interactive behaviors to physical hardware.
- We develop **IAMR** to resolve kinematic conflicts and generate humanoid-humanoid interaction motion references. Furthermore, we release **MAGIC**, the Multi-Humanoid Geometric Interaction Dataset, offering paired raw and retargeted interaction data.
- We introduce **IGRL**, a multi-agent learning module that masters coupled interaction dynamics. By incorporating graph-based rewards, it enables agents to learn robust, physically consistent interactive behaviors.
- We conduct extensive experiments on Unitree G1 humanoids in simulation and on physical hardware. Quantitative and qualitative evaluations across diverse interaction tasks demonstrate the superior performance and robustness of our framework.

II. Related Work

A. Learning-Based Humanoid Control and Interaction

Learning-based whole-body control methods primarily leverage Reinforcement Learning (RL) to endow humanoid robots with diverse motor skills [2, 4, 15–19, 24, 33, 40, 48, 52, 53]. Existing research largely focuses on single-agent

motion tracking [17, 44, 53], progressing from mimicking specific references to general motion tracking frameworks [7, 13, 26, 52, 55]. Beyond isolation, recent works have addressed the interaction between humanoids and environmental factors in sports such as table tennis [39], badminton [28], and soccer [34, 47], as well as human-robot collaboration [6, 9, 20]. However, these scenarios typically model the partner (human or object) as a passive entity or an external disturbance, ignoring the complex coupled dynamics inherent in multi-agent systems. Achieving interactive whole-body control on dual-humanoid hardware requires moving beyond these assumptions to explicitly model the mutual physical influence and geometric topology between active agents—a capability that remains absent in current robotics literature.

B. Physics-Based Animation of Human-Human Interaction

While robotics research focuses on hardware feasibility, the computer graphics community has expanded the scope of humanoid animation from single-agent generation [30–32] to the complex domain of *Human-Human Interaction* [38, 41, 46]. Addressing the inevitable physical artifacts caused by sensor inaccuracies in existing datasets [11, 25, 45, 51], recent works [10, 43, 50, 54] employ RL within physics-based simulations to guarantee human-like behavior and physical plausibility. However, these methods prioritize visual fidelity over dynamic feasibility, often relaxing rigorous constraints essential for real robots. Policies trained in these idealized environments struggle with the Sim-to-Real gap. To bridge this gap, we propose **Rhythm**, which, to our knowledge, represents the first demonstration of physically coupled dual-humanoid interaction on real-world hardware.

C. Motion Retargeting for Humanoid

Motion retargeting serves as a fundamental bridge for transferring human skills to humanoids. However, standard approaches like PHC [30] and GMR [3] typically treat agents in isolation, neglecting interaction behaviors. While OmniRetarget [49] advances this by employing *Interaction Meshes* [1, 56] to enforce human-object interaction constraints, its applicability is restricted to single-humanoid settings. Extending this to human-human interaction presents unique challenges. Recent attempts, such as PAIR [20] and Harmanoid [29], adopt a coupled formulation that intertwines self-motion with interaction constraints. Consequently, in the presence of inherent kinematic conflicts, they often compromise self-motion fidelity (e.g., causing unnatural distortions) to strictly enforce interaction geometry. To address this, we propose IAMR that effectively mitigates these kinematic conflicts to produce high-fidelity coupled motion references.

III. Methods

As illustrated in Fig. 2, **Rhythm** addresses the challenge of dual-humanoid interaction through three tightly integrated components: 1) **Interaction-Aware Motion Retargeting (IAMR)**, which synthesizes physically feasible priors by decoupling interaction geometry from self-motion fidelity (Sec. III-A); 2) **Interaction-Guided Reinforcement Learning (IGRL)**, which captures coupled dynamics via topology-aware rewards that enforce interaction consistency (Sec. III-B); and 3) **Real-World Deployment**, which overcomes the limitations of noisy global observability and asynchronous execution to bridge the Sim-to-Real gap (Sec. III-C).

A. Interaction-Aware Motion Retargeting (IAMR)

1) Problem Formulation: Given a source motion sequence of two human demonstrators with different anthropometry, our goal is to synthesize kinematically feasible trajectories for two humanoids. The core challenge lies in synthesizing motions that simultaneously preserve the individual motion style and the dense interaction geometry.

Interaction Mesh and Laplacian Coordinates. To mathematically model the coupled multi-agent system, we adopt the volumetric Interaction Mesh formalism [1, 56]. We represent the system as a connected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, comprising a vertex set $\mathcal{V}$ of the agents' key joints and an edge set $\mathcal{E}$ encoding their structural connections. To encode local geometric details, we utilize Laplacian coordinates. For a vertex p_i, the Laplacian operator $\mathcal{L}(\cdot)$ computes the deviation from the weighted average of its neighbors $\mathcal{N}(i)$:

$$\mathcal{L}(p_i) = p_i - \sum_{j \in \mathcal{N}(i)} c_{ij} p_j,$$

where c_{ij} denotes the normalized weights. In mesh-based motion retargeting, the objective is to find a target configuration q such that the local geometry of the retargeted vertices $p_i(q)$ matches the source reference. This is achieved by minimizing the deformation energy $\sum \|\mathcal{L}(p_i(q)) - \mathcal{L}(p_i^{src})\|^2$.

The Kinematic Conflict. While the standard formulation $\sum \|\mathcal{L}(p_i(q)) - \mathcal{L}(p_i^{src})\|^2$ is effective for single-agent editing, applying it to the transfer from heterogeneous humans to homogeneous robots creates a fundamental ambiguity in defining the source reference $\mathcal{V}^{src}$. Due to the embodiment mismatch, no single reference manifold can simultaneously satisfy both self-motion and interaction constraints:

- **Individual Manifold ($\mathcal{M}_{ind}$):** Constructed by scaling each human with individual ratios to match the robot's height. Using this as $\mathcal{V}^{src}$ preserves valid self-motion Laplacian coordinates but structurally disrupts the relative interaction geometry (e.g., causing "air handshakes").
- **Unified Manifold ($\mathcal{M}_{uni}$):** Constructed by applying a single global scale to the entire scene. Using this as $\mathcal{V}^{src}$ strictly preserves relative interaction edges but forces the robots to adopt kinematic constraints incompatible with their morphology (e.g., foot floating).

Topological Partitioning. To resolve this conflict, we propose to relax the monolithic structure of the standard Interaction Mesh. Instead of treating the system as a single deformable body, we explicitly partition $\mathcal{E}$ into two disjoint functional groups, as visualized in the left part of Fig. 2: (1) **Intra-Agent Edges ($\mathcal{E}_{self}$):** Edges connecting joints within a single

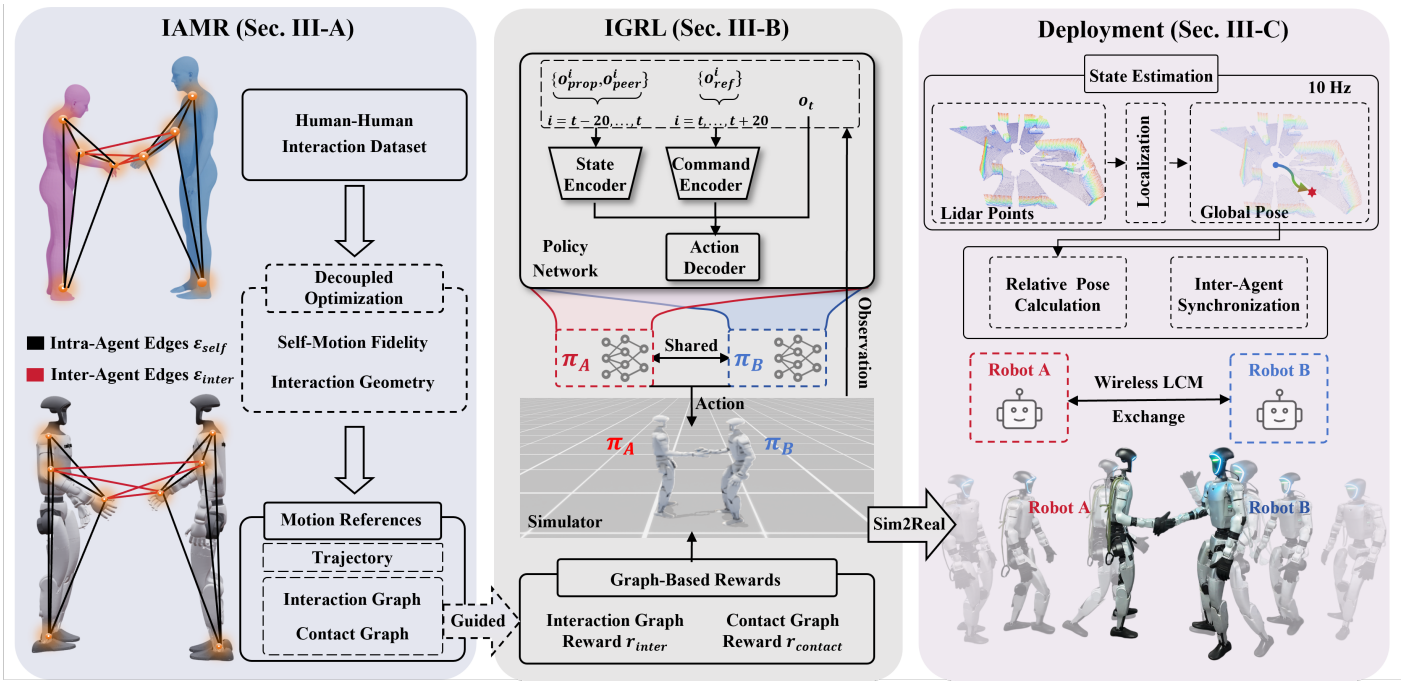

Fig. 2: Overview of Rhythm. IAMR utilizes decoupled optimization to generate high-quality humanoid-humanoid motion interaction references from human demonstrations. Guided by these references, IGRL employs MAPPO and graph-based rewards to learn robust coupled dynamics. Finally, the deployment module facilitates Sim-to-Real transfer via Lidar-fused state estimation and inter-agent synchronization.

robot, encoding the local self-motion topology. (2) **Inter-Agent Edges ($\mathcal{E}_{inter}$):** Edges connecting key joints between the two robots, encoding the relative interaction topology.

This topological decomposition allows us to assign distinct geometric references to different semantic parts of the graph, forming the basis for our decoupled optimization scheme.

2) Decoupled Optimization: Leveraging the topological partitioning defined above, we resolve the kinematic conflict by assigning distinct geometric references to the partitioned subgraphs. We formulate the retargeting as a dynamic spring system, where a self-motion term E_{self} ensures intra-agent edges $\mathcal{E}_{self}$ track the Independent Manifold $\mathcal{M}_{ind}$, while an interaction term E_{inter} constrains inter-agent edges $\mathcal{E}_{inter}$ to the Unified Manifold $\mathcal{M}_{uni}$.

We solve for the optimal joint configuration $q^* = \{q_1, q_2\}$ by minimizing a hybrid energy function:

$$q^* = \arg\min_q \left(E_{self}(q) + E_{inter}(q)\right) \quad \text{s.t.} \quad q \in \mathcal{C}_{phy},$$

where $\mathcal{C}_{phy}$ represents the set of feasible configurations satisfying joint limits and collision constraints.

Self-Motion Objective (E_{self}). To preserve individual motion quality, we align the Laplacian geometry of each robot to its Independent Reference $\mathcal{M}_{ind}$, strictly confining the operator to the local subgraph:

$$E_{self}(q) = \sum_{a\in\{1,2\}} \sum_{p_i\in\mathcal{V}_a} \|\mathcal{L}(p_i) - \mathcal{L}(p_i^{ind})\|^2 + \lambda_{rot} \sum_{a\in\{1,2\}} \sum_{k\in\mathcal{B}_a} \|\theta_k \ominus \hat{\theta}_k^{src}\|^2,$$

where $\mathcal{V}_a$ denotes the vertex set of robot a, and p_i^{ind} represents the corresponding vertex position in the reference $\mathcal{M}_{ind}$. For rotation, $\mathcal{B}_a$ denotes the key links, where $\ominus$ measures the geodesic distance on $SO(3)$ between the current orientation θ_k and the reference $\hat{\theta}_k^{src}$.

Interaction Objective (E_{inter}). To enforce the relative interaction geometry, we treat the inter-agent edges as extrinsic constraints driven by the Unified Reference $\mathcal{M}_{uni}$. We formulate this as a variable-stiffness spring potential:

$$E_{inter}(q) = \sum_{(i,j)\in\mathcal{E}_{inter}} \omega_{ij}(d_{ij}) \cdot \|(p_i - p_j) - (\hat{p}_i^{uni} - \hat{p}_j^{uni})\|^2,$$

Here, p_i and p_j denote the current vertex positions of different agents, while $\hat{p}_i^{uni}$ and $\hat{p}_j^{uni}$ represent the corresponding target coordinates derived from the reference $\mathcal{M}_{uni}$. To naturally prioritize close-range geometry (e.g., contact) over distant relations, we define the stiffness ω_{ij} as a continuous exponential decay function of the source distance d_{ij}:

$$\omega_{ij}(d_{ij}) = \omega_{max} \cdot e^{-\gamma d_{ij}},$$

where ω_{max} denotes peak stiffness and γ the decay rate. This effectively models the interaction as a non-linear spring system that stiffens during close contact to prevent penetration, while becoming compliant at a distance to allow free motion.

Topological Interaction Priors. Beyond kinematic trajectories, IAMR explicitly extracts inter-agent topological structures to serve as interaction priors for downstream policy learning, as illustrated in the left part of Fig. 2. We generate two cross-agent graph representations: (1) An **Interaction Graph** (derived from $\mathcal{E}_{inter}$), which encodes the binary connectivity *bridging* the keypoints of the two agents [54]; (2) A **Contact Graph** (constructed via collision detection), which records the binary physical contact states *between* the links of the

two distinct robots [42]. These graphs provide the essential interaction topology required for the graph-based rewards in the subsequent training phase.

B. Interaction-Guided Reinforcement Learning (IGRL)

To master the coupled dynamics of dual-humanoid interaction, we propose **IGRL**, a multi-agent reinforcement learning module. Unlike standard motion tracking policies that treat agents as isolated entities, IGRL explicitly models interaction geometry through graph-based rewards and incorporates specific training strategies to bridge the reality gap.

1) Multi-Agent Policy Design: As shown in the middle part of Fig. 2, we formulate the problem as a Multi-Agent Markov Decision Process (MA-MDP) and adopt the *Centralized Training with Decentralized Execution* (CTDE) paradigm using MAPPO [10, 27].

Interaction-Centric Observation. Effective collaboration requires a comprehensive awareness of both self and partner states. The policy operates on a compact observation space $o_t = \{o_{prop}, o_{peer}, o_{ref}\}$:

- Proprioception (o_{prop}): Defines the agent's internal state, including joint positions, velocities, base angular velocities, and the previous action.
- Peer Perception (o_{peer}): Encodes the peer's state relative to the ego agent. It includes the peer's joint positions and the relative root transform expressed in the ego-centric frame: the relative position $P_{rel} = R_{ego}^T(P_{peer} - P_{ego})$ and orientation $R_{rel} = R_{ego}^T R_{peer}$. This explicit spatial formation enables the agent to anticipate and handle coupled dynamics.
- Reference Motion (o_{ref}): Contains the future reference trajectories and the reference relative state, serving as the correct interaction topology.

Network Architecture. We design a hierarchical policy architecture where 1D-CNN temporal encoders process historical observations and future references, feeding latent features along with the current observation into an MLP action decoder. More details are provided in Appendix B.

Robust Training Strategy. To ensure transferability and handle complex interaction phases, we implement two strategies, with mathematical formulations provided in Appendix B:

- **Curriculum-based Adaptive Sampling:** Standard RSI [31] relies on sparse failure counts, neglecting non-terminal interaction violations. We propose an error-aware sampling strategy based on a multi-objective landscape composed of integrating failure, tracking, and interaction metrics. The curriculum dynamically evolves from prioritizing stability in early stages to focusing on tracking and interaction precision as the policy matures.
- **Dual-Agent Domain Randomization:** To bridge the sim-to-real gap, we simulate wireless latency via noisy, delayed peer observations and apply initial state perturbations to enforce recovery from physical misalignment.

2) Graph-based Rewards: Standard tracking rewards typically treat agents as isolated entities, failing to enforce topological consistency or capture coupled dynamics. To bridge this gap, we introduce graph-based rewards to **guide** the learning of dynamic control policies by translating the *Topological Interaction Priors* established in IAMR, as shown in Fig. 2.

Interaction Graph Reward (r_{inter}). To ensure precise relative geometry during interaction, we penalize deviations of the *interaction graph*. The reward is formulated as:

$$r_{inter} = \exp\left(-\frac{1}{\sigma_{inter}} \sum_{(i,j)\in\mathcal{E}_{inter}} \omega_{ij} \|d_{ij}^{sim} - \hat{d}_{ij}^{ref}\|^2\right),$$

where σ_{inter} is a sensitivity scaling factor, and d_{ij} denotes the relative position vector connecting joints i and j in the simulation (sim) and reference (ref) states. By inheriting the distance-aware dynamic weights ω_{ij} from IAMR, the policy inherently learns to prioritize the same geometric constraints that were optimized during retargeting.

Contact Graph Reward ($r_{contact}$). Since kinematic references lack force information, we utilize *contact graph* to regularize physical interaction. This reward is designed with two goals: (1) *Contact Consistency*, which penalizes mismatches between the simulated and reference contact states; and (2) *Force Regularization*, which constrains contact forces to realistic ranges when active, and penalizes non-zero forces during non-contact phases. This explicitly discourages interpenetration and encourages compliant physical interaction.

C. Real-World Deployment

Deploying **Rhythm** on physical hardware faces the *Sim-to-Real Gap*, specifically the lack of global observability and the issue of asynchronous execution.

1) State Estimation and Relative Localization: Constructing the observation space for dual-humanoid interaction requires precise global and relative state information. We employ a robust localization system, utilizing *POINT-LIO* [14] for high-frequency local odometry and *GICP* [35] to register real-time point clouds against a pre-built map for drift-free global positioning. A Kalman Filter fuses these estimates to ensure robustness under highly dynamic motions.

Robots broadcast global poses $\{P, R\}$ via *LCM* [21]. The ego-agent transforms received data into its local frame to derive $\{P_{rel}, R_{rel}\}$, enabling real-time reconstruction of o_{peer} consistent with the simulation.

2) Inter-Agent Synchronization: We use the *motion phase* ϕ [31] as the continuous variable representing the temporal execution progress of the interaction policy. In distributed settings, inherent hardware clock drift inevitably causes the phases of the two agents to diverge over time.

To maintain temporal alignment, we implement a soft synchronization mechanism based on proportional feedback. Agents exchange their current phase ϕ via the wireless bridge. Upon receiving the peer's phase ϕ_{peer}, the ego agent dynamically modulates its phase progression rate $\dot{\phi}_{ego}$. Instead of a

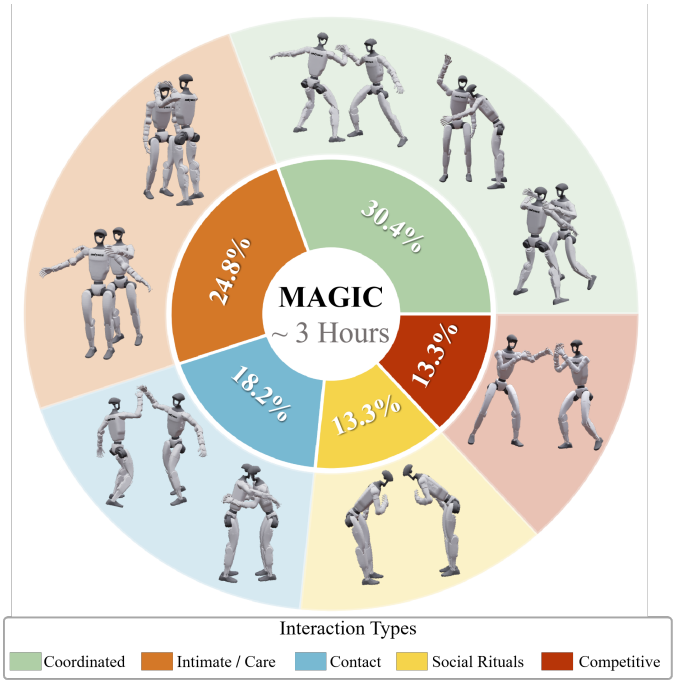

Fig. 3: Overview of MAGIC. MAGIC contains ∼3 hours of high-fidelity interaction data balanced across five semantic categories (inner chart). Representative snapshots (outer ring) illustrate the diversity ranging from loose spatiotemporal coordination to intensive contact.

fixed increment ($\dot{\phi}_{base} = 1.0$), we apply a correction:

$$\dot{\phi}_{ego} = 1.0 + k(\phi_{peer} - \phi_{ego}),$$

where k is the synchronization gain. This compensates for temporal drift by modulating execution rate, ensuring smooth alignment without the discontinuities of hard phase resets.

IV. Experiments

We design our experiments to systematically validate the proposed framework by answering three core questions:

- **Q1 (Retargeting Quality):** Can IAMR synthesize kinematically feasible and topologically consistent trajectories that preserve the interaction geometry of the source data?
- **Q2 (Policy Efficacy):** Can IGRL utilize the interaction guidance to capture the coupled dynamics, overcoming the limitations of treating agents as isolated entities?
- **Q3 (Real-World Robustness):** Can the learned policies be successfully deployed to physical dual-humanoid systems, overcoming the limitations of noisy global observability and asynchronous communication?

A. The MAGIC Dataset

High-quality motion data is the cornerstone of learning interactions. Addressing the lack of clean human-human interaction data, we introduce **MAGIC**, a high-fidelity motion capture dataset comprising ∼3 hours of valid motion sequences.

1) Acquisition and Diversity: Distinct from existing datasets [11, 25, 45, 51], we ensured physical plausibility for robot transfer through *Anthropometric Consistency* (matched actor heights) and *Temporal Continuity* (long-horizon sequences $>$ 10s). As illustrated in Fig. 3, MAGIC covers a diverse semantic spectrum: *Coordinated* actions (30.4%), *Intimate/Care* behaviors (24.8%), *Contact* (18.2%), *Social Rituals* (13.3%), and *Competitive* interactions (13.3%). All sequences are captured via high-fidelity optical motion capture in BVH format. Interaction graphs are extracted from the inter-agent edges $\mathcal{E}_{inter}$ of the retargeted motion, while contact labels are derived post-retargeting via collision detection on the retargeted humanoid kinematics.

2) Data Release: To facilitate future research, we will publicly release both the raw motion capture data and the retargeted humanoid reference trajectories generated by IAMR.

B. Experimental Setup

Datasets. We use the proposed MAGIC dataset for both training and evaluation. To systematically assess performance across diverse physical dynamics, we categorize the tasks into three physics-based groups: **Collaborate** (non-contact synchronization), **Light Contact** (transient interaction), and **Intensive Contact** (continuous force transmission). Additionally, we employ the external *Inter-X* dataset [45] to assess robustness under significant anthropometric mismatches.

Baselines. We benchmark **Rhythm** against representative methods under two evaluation aspects, with full implementation details in Appendix C:

- **Retargeting Quality (Q1):** We compare against **GMR** [3], which performs Cartesian optimization, **OR** [49], a single-agent interaction mesh method, and **DOR**, our constructed dual-agent extension of OR.
- **Policy Efficacy (Q2):** To validate the effectiveness of IGRL, we compare against a **Single-Agent** (Status Quo) baseline that performs isolated tracking, as well as key ablated variants of our method, including **w/o Peer Obs**, **w/o Contact Rew**, and **w/o Interaction Rew**.

Evaluation Metrics. We adopt task-specific evaluation metrics for **Retargeting Quality (Q1)** and **Policy Efficacy (Q2)**. For conciseness, detailed metric definitions and mathematical formulations are deferred to Appendix C.

C. Retargeting Quality (Q1)

We evaluate IAMR against baselines across three dimensions: **Safety**, measured by Inter-Penetration Rate (IPR) and Max Penetration Depth (MPD); **Fidelity**, quantified by Interaction Edge Error (IEE) and Contact F1 Score [20, 29]; and **Utility**, assessed via Downstream Success Rate (DSR).

Quantitative Analysis. Table I presents the performance comparison across four interaction categories. We observe three key trends. First, isolated baselines (GMR, OR) fundamentally fail to ensure physical feasibility. Notably, in Intensive Contact, OR exhibits a prohibitive Inter-Penetration Rate (IPR) of 47.3%, rendering the motions physically infeasible and unsuitable for policy learning. Second, while the coupled baseline (DOR) guarantees safety (IPR=0), its rigid formulation compromises interaction fidelity under anthropometric mismatch. This is evident on the Inter-X dataset, where DOR's F1-Strict drops to 0.589 due to its inability to reconcile conflicting kinematic constraints. In contrast, IAMR achieves

TABLE I: Quantitative Results of Retargeting. Comparison across four interaction categories. Metrics include Safety (IPR, MPD), Fidelity (IEE, F1), and Utility (DSR). IAMR achieves the best balance, strictly eliminating penetration (IPR=0) while maximizing contact F1 scores.

	Safety		Fidelity			Utility
	IPR(%)↓	MPD(cm)↓	IEE(%)↓	F1-S↑	F1-L↑	DSR(%)↑
MAGIC: Collaborate						
GMR	0.14	1.2	4.3	0.602	0.804	85.5
OR	0.20	1.4	4.1	0.747	0.902	87.4
DOR	**0.00**	**0.0**	3.9	0.711	0.899	**89.5**
IAMR	**0.00**	**0.0**	**3.7**	**0.785**	**0.936**	89.0
MAGIC: Light Contact						
GMR	2.18	3.3	4.6	0.738	0.893	48.5
OR	7.62	5.9	3.6	0.844	0.912	63.1
DOR	**0.00**	**0.0**	3.6	0.810	0.918	69.4
IAMR	**0.00**	**0.0**	**3.1**	**0.905**	**0.935**	**75.3**
MAGIC: Intensive Contact						
GMR	35.2	3.8	9.6	0.864	0.928	45.5
OR	47.3	5.3	8.0	0.884	0.929	56.5
DOR	**0.00**	**0.0**	7.8	0.883	0.925	63.3
IAMR	**0.00**	**0.0**	**6.6**	**0.932**	**0.941**	**78.3**
Inter-X						
GMR	11.7	1.7	8.0	0.598	0.752	31.7
OR	18.4	2.6	6.8	0.587	0.791	46.3
DOR	**0.00**	**0.0**	6.7	0.589	0.795	52.9
IAMR	**0.00**	**0.0**	**4.9**	**0.843**	**0.860**	**69.9**

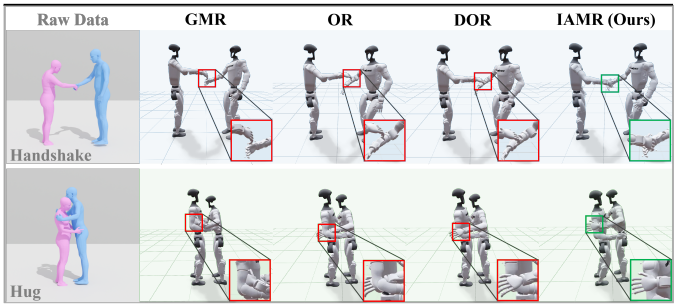

Fig. 4: Qualitative Visualization of Retargeting on Inter-X. Top: Baselines suffer from contact loss ("air handshakes"), whereas IAMR preserves precise interaction geometry. **Bottom:** OR leads to severe penetration while DOR forces unnatural stiff postures; IAMR maintains close-proximity topology without collisions.

the optimal safety-fidelity trade-off. By adaptively decoupling self-motion from interaction objectives, our method eliminates penetration while outperforming DOR by 43% in F1-Strict on Inter-X. Crucially, this strict preservation of interaction topology translates to superior downstream utility, yielding 78.3% DSR in contact-rich tasks.

Qualitative Visualization (Inter-X Cases). Fig. 4 validates robustness under significant mismatches. In transient "Handshakes", IAMR utilizes dynamic weighting to prioritize critical interaction geometry, effectively resolving the contact loss ("air handshakes") observed in baselines. Conversely, in continuous "Hugs", our decoupled optimization reconciles kinematic conflicts, preventing the penetration of OR and the stiffness of DOR to maintain valid close-proximity topology.

TABLE II: Quantitative Results of Policy. We evaluate the contribution of each component. Our full method achieves the most robust balance, effectively integrating coarse-grained geometric alignment (low IEE) with fine-grained physical contact fidelity (high CSR).

	Interaction		Contact	
	ISR(%)↑	IEE(%)↓	CSR(%)↑	CER↓
MAGIC: Collaborate				
Single Agent	18.7	38.9	100.0	0.000
w/o Peer Obs	19.5	47.0	100.0	0.000
w/o Contact Rew	**93.4**	**4.7**	100.0	0.000
w/o Interact Rew	58.1	15.1	100.0	0.000
Ours (Full)	92.9	4.8	100.0	0.000
MAGIC: Light Contact				
Single Agent	34.3	19.9	24.1	0.283
w/o Peer Obs	48.9	13.9	18.6	0.268
w/o Contact Rew	85.9	5.4	52.1	0.203
w/o Interact Rew	48.7	19.3	28.1	0.243
Ours (Full)	**90.0**	**4.2**	**78.0**	**0.120**
MAGIC: Intensive Contact				
Single Agent	24.0	29.9	37.5	0.312
w/o Peer Obs	34.1	21.4	43.7	0.280
w/o Contact Rew	**77.3**	**7.7**	70.6	0.174
w/o Interact Rew	51.3	17.0	56.8	0.211
Ours (Full)	75.2	7.9	**78.8**	**0.159**
Inter-X				
Single Agent	25.7	26.9	57.4	0.256
w/o Peer Obs	73.2	7.6	75.3	0.143
w/o Contact Rew	**95.1**	**3.4**	68.3	0.208
w/o Interact Rew	63.2	9.7	**78.6**	**0.110**
Ours (Full)	92.8	3.5	77.4	0.125

D. Policy Efficacy (Q2)

We assess the robustness and fidelity of the learned control policy across two dimensions: **Interaction Performance**, measured by Interaction Edge Error (IEE) and Interaction Success Rate (ISR); and **Contact Performance**, quantified by Contact Success Rate (CSR) and Contact Error Rate (CER).

Quantitative Analysis. Table II presents the quantitative ablation results, revealing two critical insights. First, Interaction Awareness is non-negotiable. The Single Agent (Vanilla) baseline, limiting agents to isolated tracking, fails to coordinate effectively, achieving only 18.7% ISR in Collaborate scenarios. Similarly, removing peer observations (w/o Peer Obs) severs the closed-loop synchronization, causing physical coupling to collapse (CSR drops to 18.6% in Light Contact). Second, there exists a functional hierarchy between coarse-grained geometric guidance and fine-grained contact regulation. The ablation results reveal that the interaction reward serves as a necessary foundation: removing it (w/o Interaction Rew) causes performance to collapse across all metrics as the policy fails to guide agents into the interaction envelope where contact can be established. For example, ISR drops to 58.1% in Collaborate and CSR drops to 28.1% in Light Contact. Once this spatial proximity is achieved, the contact reward becomes critical for physical realism. Notably, the w/o Contact Rew variant exhibits a "ghosting" phenomenon: it

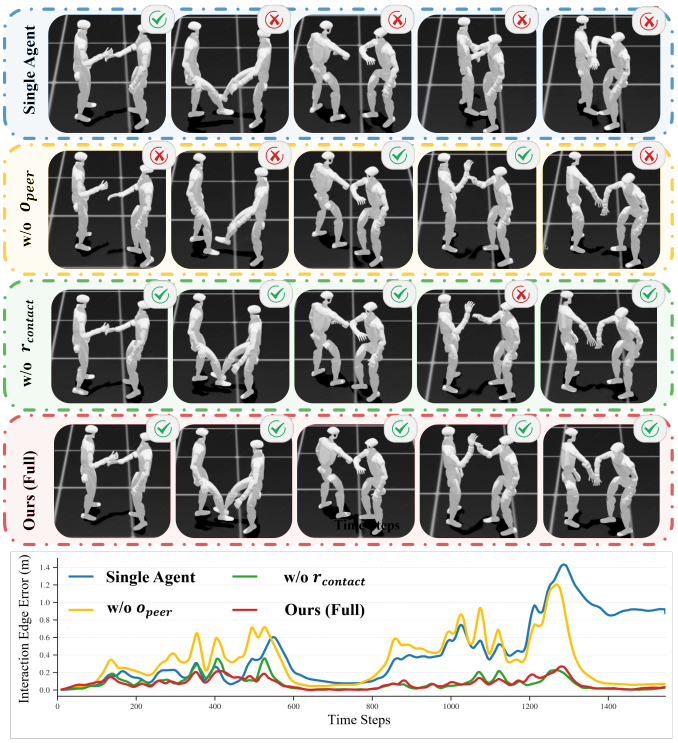

Fig. 5: Qualitative Visualization of Policy. Single Agent (blue) drifts into collisions. w/o Contact Rew (green) achieves low error but exhibits physical "ghosting". In contrast, Ours enforces valid physical contact.

attains high geometric precision (93.4% ISR in Collaborate) by disregarding physical collision constraints, yet fails to maintain valid physical contact (only 52.1% CSR in Light Contact). This is because the policy over-optimizes r_{inter}, exploiting physical interpenetration as a shortcut to minimize geometric error rather than establishing true physical coupling. Ours (Full) effectively integrates these components, leveraging geometric guidance to establish the spatial foundation while using contact regulation to enforce valid physical coupling (above 75% ISR and 77% CSR across all scenarios).

Qualitative Visualization. Fig. 5 highlights the behavioral divergence across methods in the Greeting task. The Single Agent, treating the peer merely as a dynamic obstacle, may initiate interaction but inevitably drifts into collision. While its low-level robustness prevents immediate termination, the interaction topology is destroyed. Similarly, without explicit peer monitoring, the w/o Peer Obs baseline suffers from severe desynchronization (large error spikes). Notably, the w/o Contact Rew variant successfully maintains coarse-grained geometric alignment, yielding an IEE curve comparable to the full method (Green versus Red lines). However, it lacks fine-grained contact fidelity, allowing hands to "ghost" through each other. Ours bridges this gap, achieving the precision as w/o Contact Rew while enforcing valid physical coupling at the contact interface. These visual observations directly align with the quantitative hierarchy in Table II, particularly explaining the discrepancy between geometric (ISR) and physical (CSR) metrics in the baselines.

TABLE III: Main Results for Real Robot Experiments. We conducted 10 trials for each task and evaluated success based on contact establishment at specific keyframes (K frames per trial).

Task	Method	Success / Total	Rate (%)
Hug	Single Agent	8 / 30	26.7%
	Ours	**26 / 30**	**86.7%**
Shoulder	Single Agent	6 / 30	20.0%
	Ours	**24 / 30**	**80.0%**
Greeting	Single Agent	11 / 90	12.2%
	Ours	**74 / 90**	**82.2%**

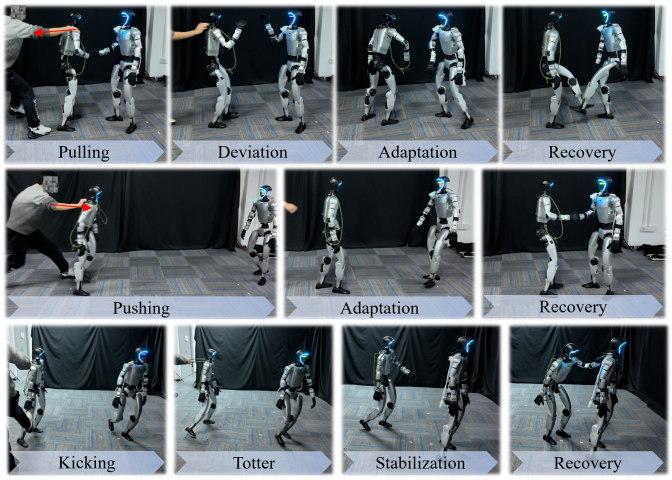

Fig. 6: Robustness to disturbances. Our policy demonstrates strong resilience against aggressive external perturbations (pulling, pushing, and kicking), successfully recovering balance and synchronization.

E. Real-World Robustness (Q3)

We validate the deployment of **Rhythm** on Unitree G1 humanoids to assess its generality, robustness, and success rate in physical environments.

Framework Generality. Fig. 1 demonstrates the system's versatility across diverse modalities. By strictly preserving fine-grained contact geometry in physical coupling tasks (Fig. 1a-c) and maintaining spatiotemporal coherence in long-horizon coordination (Fig. 1d), **Rhythm** effectively ensures interaction integrity across the spectrum using a unified formulation.

Quantitative Success Rate. We evaluate success based on valid contact establishment at task-specific keyframes: *Hug* uses 3 keyframes (shoulder pat, hand clasp, full hug); *Shoulder* uses 3 keyframes (start, mid, end of walking); and *Greeting* uses 9 keyframes (spanning hand, leg, shoulder, and elbow contacts). Table III reports success rates (measured over 10 trials). **Rhythm** consistently outperforms the Single Agent baseline by more than 60%, with the most pronounced gap in *Greeting* (12.2% versus 82.2%). While the baseline fails due to accumulated drift in this long-horizon task, our relative state estimation continuously compensates for misalignment, ensuring robust physical coupling.

Robustness to Disturbances. To evaluate system resilience, we subjected the robots to significant external perturbations during execution. As illustrated in Fig. 6, these disturbances

included aggressive pushing, pulling, and kicking forces. Despite the severity of these physical interferences, the agents successfully recovered their balance and actively adjusted their relative states to restore the interaction topology. This empirical evidence confirms that our policy has learned robust closed-loop synchronization capabilities, allowing for real-time recovery strategies rather than open-loop motion replay.

V. Conclusion

In this work, we present **Rhythm**, a unified framework that, to our knowledge, achieves the first robust transfer of physically coupled interactive behaviors to dual-humanoid hardware. By integrating Interaction-Aware Motion Retargeting (IAMR) to resolve kinematic conflicts and Interaction-Guided Reinforcement Learning (IGRL) to master coupled dynamics, our approach effectively bridges the Sim-to-Real gap. Furthermore, we release the MAGIC dataset to facilitate future research in multi-agent embodied intelligence.

While this work utilizes dual-humanoid setups to rigorously validate the modeling of coupled dynamics, our graph-based formulation is theoretically generic and supports extension to multi-agent systems. A current limitation is the reliance on pre-built maps for state estimation. Future work will focus on eliminating this dependency by shifting towards fully ego-centric perception for map-free collaboration, and scaling the framework to orchestrate complex multi-humanoid interactions in open-world scenarios.

Acknowledgments

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant 62403142, and in part by the Science and Technology Commission of Shanghai Municipality under Grant 24511103100. This work was led by Hongjin Chen during the internship at TARS.

References

[1] Marc Alexa. Differential coordinates for local mesh morphing and deformation. *The Visual Computer*, 19 (2):105–114, 2003.

[2] Arthur Allshire, Hongsuk Choi, Junyi Zhang, David McAllister, Anthony Zhang, Chung Min Kim, Trevor Darrell, Pieter Abbeel, Jitendra Malik, and Angjoo Kanazawa. Visual imitation enables contextual humanoid control. In *Conference on Robot Learning*, 2025.

[3] 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.

[4] Qingwei Ben, Feiyu Jia, Jia Zeng, Junting Dong, Dahua Lin, and Jiangmiao Pang. HOMIE: Humanoid loco-manipulation with isomorphic exoskeleton cockpit. In *Robotics: Science and Systems*, 2025.

[5] Longbing Cao. Humanoid robots and humanoid ai: Review, perspectives and directions. *ACM Computing Surveys*, 58(4):1–37, 2025.

[6] Haoran Chen, Yiteng Xu, Yiming Ren, Yaoqin Ye, Xinran Li, Ning Ding, Yuxuan Wu, Yaoze Liu, Peishan Cong, Ziyi Wang, et al. Symbridge: A human-in-the-loop cyber-physical interactive system for adaptive human-robot symbiosis. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*, pages 1–12, 2025.

[7] 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*, 2025.

[8] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. In *Robotics: Science and Systems*, 2024.

[9] Yushi Du, Yixuan Li, Baoxiong Jia, Yutang Lin, Pei Zhou, Wei Liang, Yanchao Yang, and Siyuan Huang. Learning human-humanoid coordination for collaborative object carrying. *arXiv preprint arXiv:2510.14293*, 2025.

[10] Jiawei Gao, Ziqin Wang, Zeqi Xiao, Jingbo Wang, Tai Wang, Jinkun Cao, Xiaolin Hu, Si Liu, Jifeng Dai, and Jiangmiao Pang. CooHOI: Learning cooperative human-object interaction with manipulated object dynamics. *Advances in Neural Information Processing Systems*, 37: 79741–79763, 2024.

[11] Anindita Ghosh, Rishabh Dabral, Vladislav Golyanik, Christian Theobalt, and Philipp Slusallek. ReMoS: 3D motion-conditioned reaction synthesis for two-person interactions. In *European Conference on Computer Vision (ECCV)*, 2024.

[12] Xinyang Gu, Yen-Jen Wang, Xiang Zhu, Chengming Shi, Yanjiang Guo, Yichen Liu, and Jianyu Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. In *Robotics: Science and Systems*, 2024.

[13] 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.

[14] Dongjiao He, Wei Xu, Nan Chen, Fanze Kong, Chongjian Yuan, and Fu Zhang. Point-LIO: robust high-bandwidth light detection and ranging inertial odometry. *Advanced Intelligent Systems*, 5(7):2200459, 2023.

[15] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. OmniH2O: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. In *Conference on Robot Learning*, 2024.

[16] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2024.

[17] Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbabu, Chaoyi Pan, Zeji Yi, Guannan Qu, Kris Kitani, Jessica K. Hodgins, Linxi Fan, Yuke Zhu, Changliu

Liu, and Guanya Shi. ASAP: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. In *Robotics: Science and Systems*, 2025.

[18] Xialin He, Runpei Dong, Zixuan Chen, and Saurabh Gupta. Learning getting-up policies for real-world humanoid robots. In *Robotics: Science and Systems*, 2025.

[19] Tao Huang, Junli Ren, Huayi Wang, Zirui Wang, Qingwei Ben, Muning Wen, Xiao Chen, Jianan Li, and Jiangmiao Pang. Learning humanoid standing-up control across diverse postures. In *Robotics: Science and Systems*, 2025.

[20] Wei-Jin Huang, Yue-Yi Zhang, Yi-Lin Wei, Zhi-Wei Xia, Juantao Tan, Yuan-Ming Li, Zhilin Zhao, and Wei-Shi Zheng. Learning whole-body human-humanoid interaction from human-human demonstrations. *arXiv preprint arXiv:2601.09518*, 2026.

[21] Yih Huang, Eric Fleury, and Philip K McKinley. LCM: A multicast core management protocol for link-state routing networks. In *ICC'98. 1998 IEEE International Conference on Communications. Conference Record. Affiliated with SUPERCOMM'98 (Cat. No. 98CH36220)*, volume 2, pages 1197–1201. IEEE, 1998.

[22] Jialong Li, Xuxin Cheng, Tianshu Huang, Shiqi Yang, Ri-Zhao Qiu, and Xiaolong Wang. AMO: Adaptive motion optimization for hyper-dexterous humanoid whole-body control. In *Robotics: Science and Systems*, 2025.

[23] Junheng Li, Ziwei Duan, Junchao Ma, and Quan Nguyen. Gait-Net-augmented implicit kino-dynamic MPC for dynamic variable-frequency humanoid locomotion over discrete terrains. In *Robotics: Science and Systems*, 2025.

[24] Yixuan Li, Yutang Lin, Jieming Cui, Tengyu Liu, Wei Liang, Yixin Zhu, and Siyuan Huang. CLONE: Closed-loop whole-body humanoid teleoperation for long-horizon tasks. In *Conference on Robot Learning*, 2025.

[25] Han Liang, Wenqian Zhang, Wenxuan Li, Jingyi Yu, and Lan Xu. InterGen: Diffusion-based multi-human motion generation under complex interactions. *International Journal of Computer Vision (IJCV)*, 2024.

[26] 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. *arXiv preprint arXiv:2508.08241*, 2025.

[27] Michael L. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the Eleventh International Conference on Machine Learning (ICML)*, volume 157, pages 157–163. Morgan Kaufmann, 1994.

[28] Chenhao Liu, Leyun Jiang, Yibo Wang, Kairan Yao, Jinchen Fu, and Xiaoyu Ren. Humanoid whole-body badminton via multi-stage reinforcement learning. *arXiv preprint arXiv:2511.11218*, 2025.

[29] Zuhong Liu, Junhao Ge, Minhao Xiong, Jiahao Gu, Bowei Tang, Wei Jing, and Siheng Chen. It Takes Two: Learning interactive whole-body control between humanoid robots. *arXiv preprint arXiv:2510.10206*, 2025.

[30] Zhengyi Luo, Jinkun Cao, Alexander Winkler, Kris Kitani, and Weipeng Xu. Perpetual humanoid control for real-time simulated avatars. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10895–10904, 2023.

[31] Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel Van de Panne. DeepMimic: Example-guided deep reinforcement learning of physics-based character skills. *ACM Transactions On Graphics (TOG)*, 37(4):1–14, 2018.

[32] Xue Bin Peng, Yunrong Guo, Lina Halper, Sergey Levine, and Sanja Fidler. ASE: Large-scale reusable adversarial skill embeddings for physically simulated characters. *ACM Transactions on Graphics (TOG)*, 41 (4):1–17, 2022.

[33] Skand Peri, Akhil Perincherry, Bikram Pandit, and Stefan Lee. Non-conflicting energy minimization in reinforcement learning based robot control. In *Conference on Robot Learning*, 2025.

[34] Junli Ren, Junfeng Long, Tao Huang, Huayi Wang, Zirui Wang, Feiyu Jia, Wentao Zhang, Jingbo Wang, Ping Luo, and Jiangmiao Pang. Humanoid Goalkeeper: Learning from position conditioned task-motion constraints. *arXiv preprint arXiv:2510.18002*, 2025.

[35] Aleksandr Segal, Dirk Haehnel, and Sebastian Thrun. Generalized-ICP. In *Robotics: Science and Systems*, volume 2, page 435. Seattle, WA, 2009.

[36] Yiyang Shao, Bike Zhang, Qiayuan Liao, Xiaoyu Huang, Yuman Gao, Yufeng Chi, Zhongyu Li, Sophia Shao, and Koushil Sreenath. LangWBC: Language-directed humanoid whole-body control via end-to-end learning. In *Robotics: Science and Systems*, 2025.

[37] Qincheng Sheng, Zhongxiang Zhou, Jinhao Li, Xiangyu Mi, Pingyu Xiang, Zhenghan Chen, Haocheng Xu, Shenhan Jia, Xiyang Wu, Yuxiang Cui, et al. A comprehensive review of humanoid robots. *SmartBot*, 1(1):e12008, 2025.

[38] Sebastian Starke, Yiwei Zhao, Taku Komura, and Kazi Zaman. Local motion phases for learning multi-contact character movements. *ACM Transactions on Graphics*, 39(4), 2020.

[39] Zhi Su, Bike Zhang, Nima Rahmanian, Yuman Gao, Qiayuan Liao, Caitlin Regan, Koushil Sreenath, and S Shankar Sastry. HITTER: A humanoid table tennis robot via hierarchical planning and learning. *arXiv preprint arXiv:2508.21043*, 2025.

[40] Huayi Wang, Zirui Wang, Junli Ren, Qingwei Ben, Tao Huang, Weinan Zhang, and Jiangmiao Pang. BeamDojo: Learning agile humanoid locomotion on sparse footholds. In *Robotics: Science and Systems*, 2025.

[41] Lipeng Wang, Hongxing Fan, Haohua Chen, Zehuan Huang, and Lu Sheng. InterMoE: Individual-specific 3d human interaction generation via dynamic temporal-selective moe. In *Proceedings of the AAAI Conference*

on Artificial Intelligence, 2026.

[42] Yinhuai Wang, Qihan Zhao, Runyi Yu, Hok Wai Tsui, Ailing Zeng, Jing Lin, Zhengyi Luo, Jiwen Yu, Xiu Li, Qifeng Chen, Jian Zhang, Lei Zhang, and Ping Tan. Skillmimic: Learning basketball interaction skills from demonstrations. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 17540–17549, June 2025.

[43] Jungdam Won, Deepak Gopinath, and Jessica Hodgins. Control strategies for physically simulated characters performing two-player competitive sports. *ACM Transactions on Graphics (TOG)*, 40(4):1–11, 2021.

[44] Weiji Xie, Jinrui Han, Jiakun Zheng, Huanyu Li, Xinzhe Liu, Jiyuan Shi, Weinan Zhang, Chenjia Bai, and Xuelong Li. KungfuBot: Physics-based humanoid whole-body control for learning highly-dynamic skills. *Advances in Neural Information Processing Systems*, 2025.

[45] Liang Xu, Xintao Lv, Yichao Yan, Xin Jin, Shuwen Wu, Congsheng Xu, Yifan Liu, Yizhou Zhou, Fengyun Rao, Xingdong Sheng, Yunhui Liu, Wenjun Zeng, and Xiaokang Yang. Inter-X: Towards versatile human-human interaction analysis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

[46] Liang Xu, Chengqun Yang, Zili Lin, Fei Xu, Yifan Liu, Congsheng Xu, Yiyi Zhang, Jie Qin, Xingdong Sheng, Yunhui Liu, et al. Perceiving and acting in first-person: A dataset and benchmark for egocentric human-object-human interactions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12535–12548, 2025.

[47] Zifan Xu, Myoungkyu Seo, Dongmyeong Lee, Hao Fu, Jiaheng Hu, Jiaxun Cui, Yuqian Jiang, Zhihan Wang, Anastasiia Brund, Joydeep Biswas, et al. Learning agile striker skills for humanoid soccer robots from noisy sensory input. *arXiv preprint arXiv:2512.06571*, 2025.

[48] Yufei Xue, Wentao Dong, Minghuan Liu, Weinan Zhang, and Jiangmiao Pang. A unified and general humanoid whole-body controller for fine-grained locomotion. In *Robotics: Science and Systems*, 2025.

[49] 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.

[50] Wei Yao, Yunlian Sun, Chang Liu, Hongwen Zhang, and Jinhui Tang. PhysiInter: Integrating physical mapping for high-fidelity human interaction generation. *arXiv preprint arXiv:2506.07456*, 2025.

[51] Yifei Yin, Chen Guo, Manuel Kaufmann, Juan Jose Zarate, Jie Song, and Otmar Hilliges. Hi4D: 4D instance segmentation of close human interaction. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023.

[52] Yanjie Ze, Zixuan Chen, Joao Pedro Araujo, Zi-ang Cao, Xue Bin Peng, Jiajun Wu, and Karen Liu. TWIST: Teleoperated whole-body imitation system. In *Conference on Robot Learning*, 2025.

[53] Tong Zhang, Boyuan Zheng, Ruiqian Nai, Yingdong Hu, Yen-Jen Wang, Geng Chen, Fanqi Lin, Jiongye Li, Chuye Hong, Koushil Sreenath, and Yang Gao. HuB: Learning extreme humanoid balance. In *Conference on Robot Learning*, 2025.

[54] Yunbo Zhang, Deepak Gopinath, Yuting Ye, Jessica Hodgins, Greg Turk, and Jungdam Won. Simulation and retargeting of complex multi-character interactions. In *ACM SIGGRAPH 2023 Conference Proceedings*, pages 1–11, 2023.

[55] Zhikai Zhang, Jun Guo, Chao Chen, Jilong Wang, Chenghuai Lin, Yunrui Lian, Han Xue, Zhenrong Wang, Maoqi Liu, Jiangran Lyu, et al. Track any motions under any disturbances. *arXiv preprint arXiv:2509.13833*, 2025.

[56] Kun Zhou, Jin Huang, John Snyder, Xinguo Liu, Hujun Bao, Baining Guo, and Heung-Yeung Shum. Large mesh deformation using the volumetric graph laplacian. *ACM Transactions on Graphics (TOG)*, 24(3):496–503, 2005.