Learning Soccer Skills for Humanoid Robots: A Progressive Perception-Action Framework

Jipeng Kong[1,2*]   Xinzhe Liu[1,2*]   Yuhang Lin[1,3]   Jinrui Han[1,4]
Sören Schwertfeger[2]   Chenjia Bai[1†]   Xuelong Li[1]
[1]Institute of Artificial Intelligence (TeleAI), China Telecom   [2]ShanghaiTech University
[3]Zhejiang University   [4]Shanghai Jiao Tong University
*Equal contribution   †Corresponding author

Fig. 1: Humanoid learning soccer skills. (a) The robot performs goal-directed kicks referring to different motions conditioned on the ball position. (b) The robot achieves accurate shots while imitating styles of different professional players. (c) Stable and precise kicking on a grass field. (d) The robot can successfully kick a moving ball.

Abstract—Soccer presents a significant challenge for humanoid robots, demanding tightly integrated perception-action capabilities for tasks like perception-guided kicking and whole-body balance control. Existing approaches suffer from inter-module instability in modular pipelines or conflicting training objectives in end-to-end frameworks. We propose Perception-Action integrated Decision-making (PAiD), a progressive architecture that decomposes soccer skill acquisition into three stages: motion-skill acquisition via human motion tracking, lightweight perception-action integration for positional generalization, and physics-aware sim-to-real transfer. This staged decomposition establishes stable foundational skills, avoids reward conflicts during perception integration, and minimizes sim-to-real gaps. Experiments on the Unitree G1 demonstrate high-fidelity human-like kicking with robust performance under diverse conditions—including static or rolling balls, various positions, and disturbances—while maintaining consistent execution across indoor and outdoor scenarios. Our divide-and-conquer strategy advances robust humanoid soccer capabilities and offers a scalable framework for complex embodied skill acquisition. The project page is available on project page.

I. INTRODUCTION

Soccer represents a globally recognized benchmark for evaluating the integrated perception-action capabilities of humanoid robots [19]. Unlike the miniature Robotis OP3 platform (height: 0.5 m, mass: 3.5 kg) [8], today's humanoid robots stretch well beyond 1.2 m in height and 30 kg in mass. This scale escalation intensifies challenges in whole-body balancing and human-like kicking motions due to gravitational and inertial effects [22, 34, 44]. Meanwhile, unlike

ordinary humanoid locomotion tasks [7, 50], learning soccer demands real-time adaptation to unstructured environments: robots must simultaneously process dynamic scene changes, regulate whole-body locomotion, and execute precise ball-kicking maneuvers while maintaining balance. This holistic fusion of perception and action integrates multiple core competencies, including dynamic visual perception, whole-body control, and spatiotemporally accurate striking.

Existing approaches predominantly adopt two paradigms: (i) modular hierarchical architectures that segregate perception, motion planning, and low-level control into decoupled modules [1, 16]; and (ii) end-to-end reinforcement learning (RL) frameworks pursuing joint optimization of visual inputs to motor control [36, 49]. The former paradigm often incurs inter-module representation gaps, while the latter exhibits training instability stemming from reward conflicts between locomotion, ball kicking, and post-impact recovery. Critically, although recent vision-driven successes [41] have exploited adversarial motion priors (AMP) [31], they still depend on manually engineered rewards, severely limiting scalability and real-world robustness.

To address these challenges, we propose *Perception-Action integrated Decision-making (PAiD)*, a progressive architecture for learning humanoid soccer skills. Our core insight rethinks skill acquisition: complex behaviors should not be treated as monolithic optimization problems but rather acquired through a curriculum that first masters foundational motions, then integrates perceptual feedback, and finally solves real-world deployment. Specifically: (i) At the *motion-skill acquisition* stage, we adopt whole-body motion tracking to learn various kicking skills from human soccer players in different styles. The robot learns natural, high-fidelity kicking behaviors for various goal locations. Crucially, this stage decouples how to kick from where to kick by isolating low-level policy learning from perceptual noise, establishing stable perception-free skills. (ii) Then, we perform *perception-action integration* by introducing lightweight environmental perception for reaching and kicking balls in different positions. In this stage, the robot learns to chase and kick positionally generalizable targets and even rolling balls. A minimal prior reward guides the policy to adjust whole-body gaits and orientation toward the ball. Due to the distinct stage separation with lightweight perception modules, our method avoids reward conflicts inherent in single-stage methods. (iii) For *sim-to-real transfer*, we identify that minor discrepancies in physics properties (e.g., restitution and friction of the ball) between simulation and real world drastically degrade performance. We thus propose an iterative physics alignment strategy to identify key parameters of ball in simulation to ensure efficient sim-to-real alignment. For real-world perception, we combine visual and radar-based localization to obtain perception input.

We validate PAiD on a Unitree G1 humanoid platform. The experiments demonstrate high-fidelity replication of human kicking biomechanics, with kinematic trajectories closely matching human player motions. The robot exhibits robust kicking capabilities, achieving 91.3% kick success rate under diverse ball positions, lighting variations, and physical disturbances. Our method achieves superior sim-to-real transfer, where the physics alignment strategy yields better sample efficiency than baselines, with consistent performance across indoor and outdoor surfaces.

Our contributions are threefold: (i) A novel learning framework that decomposes soccer skills into progressive stages of motion acquisition, perception-action integration, and sim-to-real transfer; (ii) A versatile motion tracking system for learning diverse human kicking motions and a lightweight perception integration mechanism ensuring accurate ball interaction; (iii) An alignment strategy that resolves physical and visual gaps between simulation and real-world deployment. Our work advances humanoid soccer capabilities and provides a divide-and-conquer strategy for perception-based human-skill acquisition in embodied agents.

II. RELATED WORKS

A. Humanoid Whole-Body Control

Humanoid whole-body control advances along several major directions. (i) Proprioception-driven methods focus on gait control [50], standing up [13, 14, 57], jumping [21], running [1], and even parkour [24, 58] via large-scale parallel simulation, advanced algorithms and RL-based policy optimization [3, 28, 32]. Other methods also employ AMP-based reward to learn human-style gaits [25, 31, 40]. However, each specific task and human-like behavior requires meticulous engineering of human-designed or independently learned rewards, limiting their generalizability to diverse tasks. (ii) Motion-based whole-body control extracts human behaviors through motion capture [26, 27, 54], video extraction [45], or motion generation [18] with motion retargeting [4]. Then, the RL-based motion tracking techniques [11, 30], such as DeepMimic [30], ASAP [12], ExBody [6], KungfuBot [45], and BeyondMimic [23] design motion-tracking rewards to learn whole-body motion imitation policy. Recent works such as TWIST2 [53], KungfuBot2 [9], GMT [5], and Any2Track [55] extended this paradigm to general motion tracking for whole-body teleoperation. However, existing motion-tracking frameworks neglect external environmental perception and fail to address scenarios under dynamic perception changes—such as soccer playing where robots must adapt to moving goals. (iii) Perception-integrated works mainly focus on terrain awareness and contact-based object interaction (HOI) [43, 46, 56] via human-object retargeting [51]. For example, VideoMimic [2] represents a significant advancement by leveraging real-to-sim 4D reconstruction for perception and whole-body control. HITTER [37] combines high-level model-based planner for trajectory prediction and low-level controller for movement. However, these methods still operate in relatively controlled environments or lower perception requirements.

B. Robot Soccer

Learning soccer policies requires not only precise motion tracking but also accurate ball perception, real-time decision-making, and precise kicking execution, leading to substantially
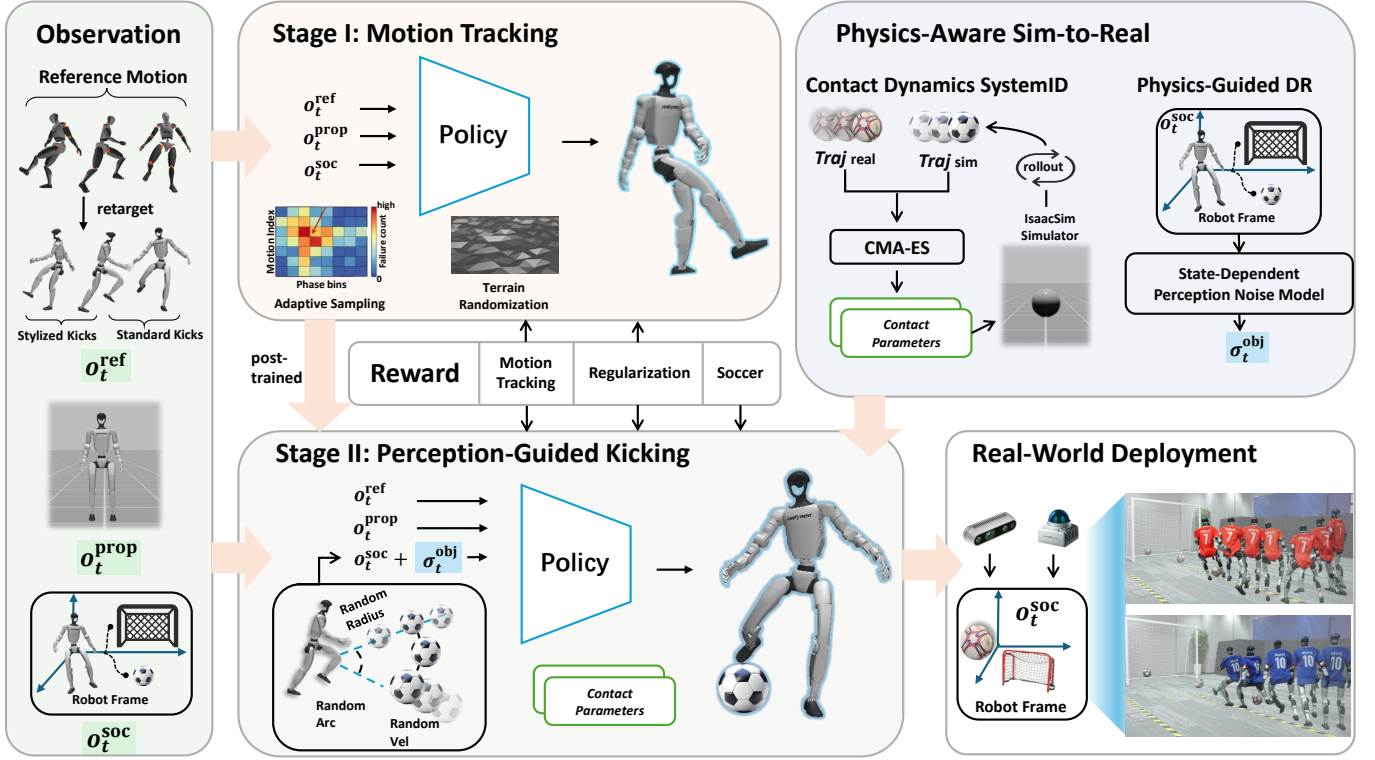
Fig. 2: Overview of the Perception-Action integrated Decision-making (PAiD) framework. Our pipeline progressively acquires robust soccer skills through three stages: (1) **Motion Tracking**: We retarget diverse human kicking motions (Standard & Stylized) to the humanoid and train a unified tracking policy using adaptive sampling to master fundamental skills without perceptual noise. (2) **Perception-Guided Kicking**: We equip the policy with egocentric perception and task-specific rewards to generalize kicking skills to randomized static and rolling ball targets. (3) **Physics-Aware Sim-to-Real Transfer**: We bridge the reality gap by aligning simulation contact dynamics with real-world measurements (ball drop & rolling tests) and incorporating physics-guided observation noise. (4) **Real-World Deployment** We successfully deploy PAiD on the Unitree G1.

higher requirements for the tight integration of perception and control [19]. Early research primarily focused on quadruped platforms, employing Bézier curves as motion priors or trajectory planners combined with end-to-end RL to learn shooting and goalkeeping strategies [15, 17, 35]. However, quadruped robots exhibit significantly lower control complexity and cannot replicate authentic human kicking motion patterns. Recent advances in humanoid soccer have explored coupled reward functions for staged chasing and kicking strategies [49], integrated locomotion and dribbling rewards [8], active perception fusion for dribbling skills [42], and environment reconstruction for soccer skill acquisition [38]; nevertheless, these approaches require complex manual reward engineering and coupled optimization processes that introduce significant challenges. While recent work [41] attempts to combine AMP with motion priors and perception integration for kicking skills, they exhibit limited capability in replicating human kicking postures and optimize perception and control through single-stage processes. In contrast, we propose a perception-action integrated framework that achieves robust and human-like soccer skills through progressive skill decomposition without reward conflicts.

III. METHOD

Our goal is to enable a humanoid robot to execute precise, human-like kicks toward a target from arbitrary ball positions. To mitigate reward conflicts between objectives, we adopt a two-stage training strategy. In Stage I, the robot learns diverse human-like soccer shooting skills via a unified motion tracking framework. In Stage II, we generalize the policy through randomized ball placements along with few task rewards, thus enabling the robot to kick the ball from arbitrary locations to the target, while maintaining human motion styles. Notably, our framework achieves this using only a compact dataset of reference motions. Finally, we perform physics-aware system identification and domain randomization to facilitate sim-to-real transfer.

A. Problem Formulation

We define the soccer shooting task as a scenario where a humanoid robot first approaches the ball and then kicks it toward the goal. The problem is formulated as a finite-horizon Markov Decision Process (MDP) defined by the tuple $\mathcal{M} = \langle \mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma \rangle$. The state space $\mathcal{S}$ represents the physical state of the robot and the environment. The action space $\mathcal{A}$ consists of target positions for the robot's joint PD

controllers. The transition dynamics $\mathcal{T} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}'$ are governed by the physical dynamics. The policy $\pi_\theta(\mathbf{a}_t|\mathbf{o}_t)$ is trained using Proximal Policy Optimization (PPO) [33], with the objective to maximize the expected cumulative reward $J(\theta) = \mathbb{E}[\sum_{t=0}^{T-1} \gamma^t r_t]$, where γ is the discount factor and r_t is the reward function at time t. To enable a single policy to master diverse soccer shooting skills and adapt to dynamic environments, we construct the observation space $\mathbf{o}_t$ as:

$$\mathbf{o}_t = \left[\, \mathbf{o}_t^{\text{prop}},\ \mathbf{o}_t^{\text{ref}},\ \mathbf{o}_t^{\text{soc}} \,\right], \tag{1}$$

where $\mathbf{o}_t^{\text{prop}}$ is the robot's proprioception, comprising projected gravity, base angular velocity, joint positions, joint velocities, and the previous action; $\mathbf{o}_t^{\text{ref}}$ is the motion tracking targets, including reference joint positions, joint velocities and root angular velocity; $\mathbf{o}_t^{\text{soc}}$ consists of the ball position $\mathbf{g}_{ball}$ and goal position $\mathbf{g}_{goal}$ relative to the robot's root frame. The ball position is fixed according to the reference motion in Stage I, and is randomized to facilitate generalization in Stage II. The goal position is sampled uniformly from a $1\,\text{m} \times 0.5\,\text{m}$ rectangular region centered 5 meters in front of the robot's initial position.

B. Stage I: Soccer-skill Acquisition via Motion Tracking

The primary objective of this stage is to acquire high-fidelity, physically feasible kicking primitives that are robust across diverse terrains. By learning these skills in a perception-free setting, the policy can focus on reproducing structured kicking behaviors, which later serves as a stable motion prior for perception-guided generalization in Stage II.

a) Data Preparation: To ensure the policy masters a diverse repertoire of kicking skills and covers a broad soccer shooting zone, we curated a specialized motion dataset comprising 13 distinct human soccer kicking motions using a motion capture system. The dataset is structured into two categories: (i) *standard kicks*, which are recorded with uniformly distributed ball placements (approximately 30 cm intervals) to cover diverse distances and shooting angles; and (ii) *stylized kicks*, which mimic signature moves of famous soccer players. For each motion, the striking leg was naturally selected based on the ball's relative location to ensure biomechanical feasibility. These motions are then retargeted to the robot via GMR [4] and annotated with the *striking leg* (left or right). This information is vital for the subsequent stage, serving as a gating signal for the contact-based rewards, ensuring the robot learns to strike the ball with the intended foot. Detailed information about the motion dataset is provided in Appendix A.

b) Unified Motion Tracking: We follow Beyond-Mimic [23] to compute a yaw-only alignment $\Delta\mathbf{R}_t$ from the relative anchor rotation and use it to rotate the reference motion into the robot's current anchor yaw before evaluating tracking errors. This alignment facilitates the robot in adjusting its heading to strike the ball towards various angles. However, relying on a single reference motion causes overfitting and restricts positional generalization, and the large variation in task difficulty further hinders unified training. In particular,

motions like stylized kicks, which involve highly dynamic movements, are substantially harder to learn than standard kicks; and even among standard kicks, differences in ball placement introduce varying requirements for precision and balance.

To tackle this challenge, we introduce an adaptive sampling strategy that learns a *single*, unified policy from a collection of reference kicking motions. Concretely, we extend the adaptive sampling approach by incorporating a motion index as an additional dimension in the sampling space. We maintain a failure histogram $F \in \mathbb{R}^{M \times B}$ over the motion indices $\mathcal{M} = \{1, \ldots, M\}$ and discretized phase bins $\mathcal{B} = \{1, \ldots, B\}$. This histogram records the failure count $F_{m,b}$ for each phase segment of each kick. At the start of every episode, we sample a motion–phase pair (m, ϕ) from a smoothed distribution that is proportional to F. As a result, training is automatically steered toward the most difficult segments across the entire motion library, promoting uniform proficiency across all skills despite their varying difficulties. Meanwhile, we incorporate terrain randomization—including irregular and uneven surfaces—directly into Stage I training. This not only enables the policy to adapt to a wide range of terrains but also removes the need to explicitly handle terrain-related difficulty when designing rewards in Stage II.

Regarding reward design for the perception-free Stage I and the subsequent perception-driven stage, we employ motion-tracking rewards along with regularization in Stage I, and then incorporate ball-centric task shaping (covering ball proximity as well as striking and outcome components) in Stage II. Table I provides an overview of all reward terms and indicates in which stage they are used. We set the episodes terminate either upon reaching a time limit or when tracking violates safety constraints, e.g., due to large attitude errors inferred from projected gravity, substantial deviations in anchor height, or pronounced kinematic inconsistencies. As a result, Stage I yields stable, human-like kicking primitives that can later be specialized to different ball setups with only minimal additional objectives.

C. Stage II: Perception-Guided Positional Generalization

The primary objective of this stage is to learn perceptually efficient kicking behaviors that generalize across varying ball positions and reliably produce goal-directed shots. To this end, we retain motion tracking as a structural prior and introduce lightweight perception integration mechanism, while avoiding interference with the learned kicking dynamics.

a) Ball Position Sampling: For each selected motion, we first place the ball at a motion-consistent nominal location determined by its terminal position, and then sample randomized ball spawns around this nominal location by (i) applying an angular perturbation around the nominal direction and (ii) sampling a radius offset. This yields a curved arc of feasible ball locations around the intended kick direction, promoting positional generalization without disrupting the motion timing. A detailed description and schematic illustration of the positional generalization setup are provided in Appendix A.

TABLE I: Reward terms used in Stage I/II.

Type	Term	Description	Weight (I)	Weight (II)
Motion-Tracking Reward	anchor-pos	Anchor position tracking: $\exp(-\|\mathbf{p}_a^{\text{robot}} - \tilde{\mathbf{p}}_a^{\text{ref}}\|_2^2/\sigma^2)$.	1.0	-
	anchor-ori	Anchor orientation tracking: $\exp(-\mathrm{d}(\mathbf{R}_a^{\text{robot}}, \tilde{\mathbf{R}}_a^{\text{ref}})^2/\sigma^2)$.	1.0	0.5
	body-pos	Body position tracking (relative): mean over tracked bodies i of $\exp(-\|\mathbf{p}_i^{\text{robot}} - \tilde{\mathbf{p}}_i^{\text{ref}}\|_2^2/\sigma^2)$.	1.0	0.8
	body-ori	Body orientation tracking (relative): mean over tracked bodies i of $\exp(-\mathrm{d}(\mathbf{R}_i^{\text{robot}}, \tilde{\mathbf{R}}_i^{\text{ref}})^2/\sigma^2)$.	1.0	0.8
	lin-vel	Body linear-velocity tracking: $\exp(-\|\mathbf{v}^{\text{robot}} - \mathbf{v}^{\text{ref}}\|_2^2/\sigma^2)$.	1.0	0.8
	ang-vel	Body angular-velocity tracking: $\exp(-\|\boldsymbol{\omega}^{\text{robot}} - \boldsymbol{\omega}^{\text{ref}}\|_2^2/\sigma^2)$.	1.0	0.8
	foot-pos	Foot position tracking (relative): $\exp(-\|\mathbf{p}_{\text{foot}}^{\text{robot}} - \tilde{\mathbf{p}}_{\text{foot}}^{\text{ref}}\|_2^2/\sigma^2)$.	-	1.0
Soccer Reward	ball-prox	Ball proximity: $\exp(-d_{xy}^2/\sigma^2)$, where d_{xy} is the horizontal base-to-ball distance. Frozen after the first valid contact.	-	1.0
	contact	Correct-foot first-contact reward (gated by the kicking-leg label).	-	50.0
	side-kick	Sideways-kick prior (leg-conditioned lateral swing).	-	50.0
	vel-align	Ball velocity direction alignment after contact.	-	30.0
	speed	Planar ball speed shaping after contact.	-	10.0
	z-speed	Ball vertical-speed penalty after contact.	-	-0.2
Regularization Reward	action-rate	Action-rate regularization (smoothness): $-\|\mathbf{a}_t - \mathbf{a}_{t-1}\|_2^2$.	-0.1	-0.1
	joint-limit	Joint-limit penalty.	-10.0	-10.0
	undesired-contact	Undesired-contact penalty (non-foot contacts).	-0.1	-0.1
	foot-sep	Foot separation reward to reduce crossing (distance-based).	-	0.2
	waist-rate	Waist action-rate regularization.	-	-0.25
	upright	Pelvis uprightness penalty (roll/pitch).	-	-1.0

Notably, we sample an initial linear velocity for the ball within a small range, and we employ an LSTM-based policy architecture capable of temporal aggregation. Consequently, the policy can implicitly predict the short-horizon trajectory of rolling balls, ensuring consistent striking performance despite ball motion.

b) Perception Observation: The policy always receives the ball and goal positions expressed in the robot pelvis frame:

$$\mathbf{g}_{ball} = \mathbf{R}_{\text{pelvis}}^{-1}\left(\mathbf{p}_{ball}^w - \mathbf{p}_{\text{pelvis}}^w\right), \qquad (2)$$

$$\mathbf{g}_{goal} = \mathbf{R}_{\text{pelvis}}^{-1}\left(\mathbf{p}_{goal}^w - \mathbf{p}_{\text{pelvis}}^w\right). \qquad (3)$$

In real-world deployment, the ball position is obtained via fused visual and radar-based localization, so $\mathbf{g}_{ball}$ is continuously available.

c) Lightweight Perception-Action Rewards: We retain the tracking rewards but relax global position constraints to allow locomotion and stance adjustment. In particular, we set the anchor position tracking weight to zero in Stage II. Furthermore, we add only a small set of task rewards: (i) *Ball proximity:* an exponential reward on the horizontal distance between the robot's pelvis and the ball. After the first valid kick contact, we freeze this term to avoid post-contact pursuit behaviors from dominating the optimization. (ii) *Contact correctness:* a one-time reward at the first ball contact, granted only if the contacting foot is consistent with the motion's left/right kicking-leg label. (iii) *Strike direction prior:* a directional prior encouraging the foot swing to align with the expected lateral striking direction conditioned on the kicking leg. (iv) *Post-kick outcome shaping:* after a correct-foot contact, we activate a short reward window to encourage the ball velocity direction to align with the desired direction from the initial ball position

to the target goal, and to encourage sufficient planar ball speed while penalizing excessive vertical speed. These terms are gated by a minimum ball-speed threshold to avoid applying outcome shaping before a physical strike.

d) Stabilization Terms: To maintain balance under perception-driven adjustments, we include stabilization terms such as waist action-rate regularization, pelvis uprightness (penalizing roll/pitch deviation via projected gravity), and a foot-separation reward to reduce foot crossing. These additions are deliberately lightweight, helping preserve the learned human-like kick kinematics while avoiding the unstable coupled optimization commonly observed in perception-to-control training.

Overall, the synergy of these mechanisms and rewards enables the robot to achieve robust and precise human-like kicking capabilities while ensuring effective positional generalization with only a compact set of reference motions.

D. Physics-Aware Sim-to-Real Transfer

The ball's physical properties in simulation are crucial for the policy to transfer to the real world, as even minor discrepancies in friction or restitution can cause significant trajectory deviations. Moreover, during real-world deployment, ball and goal observations are inherently noisy throughout the robot's movement. To this end, we adopt a physics-aware sim-to-real transfer strategy that combines contact dynamics identification with structured domain randomization (DR). Rather than relying on unstructured parameter perturbations, our approach explicitly incorporates physical priors at both the contact dynamics and observation levels.

a) Contact Dynamics Identification: To reduce discrepancies in the ball's behavior in policy transfer, we calibrate
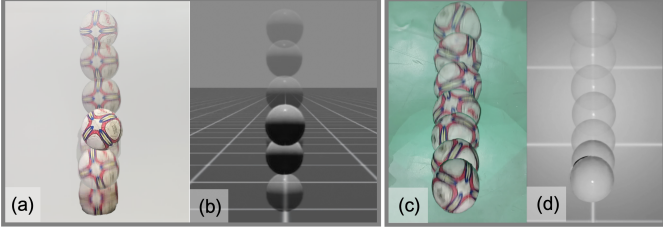
Fig. 3: Comparison of the soccer ball's physical behavior in the real world and in simulation after parameter identification. (a)–(b) compare ball drop experiments, while (c)–(d) compare rolling experiments.

the ball–ground contact dynamics through simple real-world experiments. In particular, we perform a ball drop experiment to characterize normal impact behavior and a rolling experiment to capture tangential frictional energy loss. We then replicate both experiments in IsaacSim [29], which is used for policy training, under identical initial conditions. We identify a compact set of contact parameters employed by IsaacSim, including the ball's static and dynamic friction, restitution, and linear and angular damping coefficients. Ball dynamics are represented by time-series state observations sampled at a fixed interval, yielding real-world and simulated trajectories $\mathbf{h} = \{h_i\}_{i=0}^{N_d}$, $\mathbf{d} = \{d_i\}_{i=0}^{N_r}$ and $\mathbf{h}' = \{h'_i\}_{i=0}^{N_d}$, $\mathbf{d}' = \{d'_i\}_{i=0}^{N_r}$, which denote the ball height and horizontal displacement, respectively. Optimization is performed by minimizing the trajectory matching loss

$$\mathcal{L}_{\text{sysid}} = \lambda_1 \sum_{i=0}^{N_d} (h_i - h'_i)^2 + \lambda_2 \sum_{i=0}^{N_r} (d_i - d'_i)^2, \quad (4)$$

where λ_1 and λ_2 balance the contributions of impact and rolling dynamics. Optimization is performed using a derivative-free evolutionary strategy based on CMA-ES [10].

We perform system identification separately on two surface types: a rigid hard ground and a soccer field surface, yielding two sets of nominal contact parameters. This design allows the contact behavior of the soccer ball to be accurately captured across ground conditions with distinct contact properties. As illustrated in Fig. 3, a soccer ball simulated with the identified parameters exhibits physical behaviors that closely match real-world observations, demonstrating the fidelity of the system identification process. During policy training, environments are evenly split between these two parameter sets. To account for residual modeling errors, each parameter is further perturbed by Gaussian noise sampled as $\theta \sim \mathcal{N}(\theta_{\text{nominal}}, \mathbf{I})$, where θ_{nominal} denotes the identified parameter vector for the corresponding surface. The more detailed procedure for parameter identification is provided in Appendix B.

b) Physics-Guided Domain Randomization: To account for perception uncertainty, we incorporate DR into the observation space alongside dynamics randomization. Beyond standard DR commonly used in motion tracking approaches [9, 23] summarized in Table II, we introduce a physics-guided noise model tailored to soccer ball and goal perception

TABLE II: Domain randomization terms and ranges.

Term	Range
Robot static friction	$\mathcal{U}(0.3, 1.6)$
Robot dynamic friction	$\mathcal{U}(0.3, 1.2)$
Robot restitution	$\mathcal{U}(0.0, 0.5)$
Joint default pos	$\mathcal{U}(-0.01, 0.01)$
Base CoM	$x \sim \mathcal{U}(-0.025, 0.025), y\ z \sim (-0.05, 0.05)$
Push robot	$\mathcal{U}(-0.5, 0.5)$

[20, 48].

As mentioned in Sec. III-A, the policy observes the relative positions of both the ball and the goal in the robot's local coordinate frame. Due to robot motion and partial observability, both quantities are dynamic in the robot frame and subject to state-dependent estimation errors. Empirically, these errors tend to increase with the speed of the observed object and its distance from the robot.

To capture this effect, we inject zero-mean Gaussian noise into object-related observations, including both the ball and the goal. For each object $obj \in \{\text{ball}, \text{goal}\}$, the noise magnitude is defined as

$$\sigma_t^{obj} = c_{\min} + \frac{\|\mathbf{v}_t^{obj}\|}{c_{\text{vel}}} + \frac{\|\mathbf{p}_t^{obj} - \mathbf{p}_t^{\text{robot}}\|}{c_{\text{dist}}}, \quad (5)$$

where $\mathbf{p}_t^{obj}$ and $\mathbf{v}_t^{obj}$ denote the position and velocity of object obj expressed in the robot frame. The noisy observation is then obtained by adding $\sigma_t^{obj} \cdot \mathcal{N}(0, I)$.

By grounding both contact dynamics and observation randomization in physical intuition, our strategy exposes the policy to realistic variations in ball motion and perception, leading to improved robustness in real-world deployment.

IV. EXPERIMENTS

In this section, we systematically evaluate the proposed PAiD framework from three perspectives: motion tracking quality, soccer shooting proficiency, and real-world deployment performance. The experiments are designed to answer the following research questions. **Q1**: How effectively can the robot reproduce human-like kicking motions? **Q2**: How accurate and robust is the robot's soccer shooting capability? **Q3**: How well does the learned policy transfer to the physical world? For each evaluation, we specify the experimental setup and quantitative metrics in advance.

A. Motion Tracking Quality

To answer **Q1**, we evaluate the tracking quality in Mujoco[39] using four standard metrics: *Global Mean Per-Joint Position Error (G-MPJPE)* for global pose alignment, *Mean Per-Joint Position Error (MPJPE)* for local pose precision (ignoring global drift), *Velocity Error (Vel Err)*, defined as the mean Euclidean global root velocities error for dynamic consistency, and the *AUJ Score* for physical smoothness, measured as the absolute deviation of the time-averaged peak link jerk from the reference dataset average. For a comprehensive evaluation, we divide our motion dataset into two difficulty categories: (i) *Standard Kicks* (10 motions), representing fundamental kicking mechanics with moderate complexity and

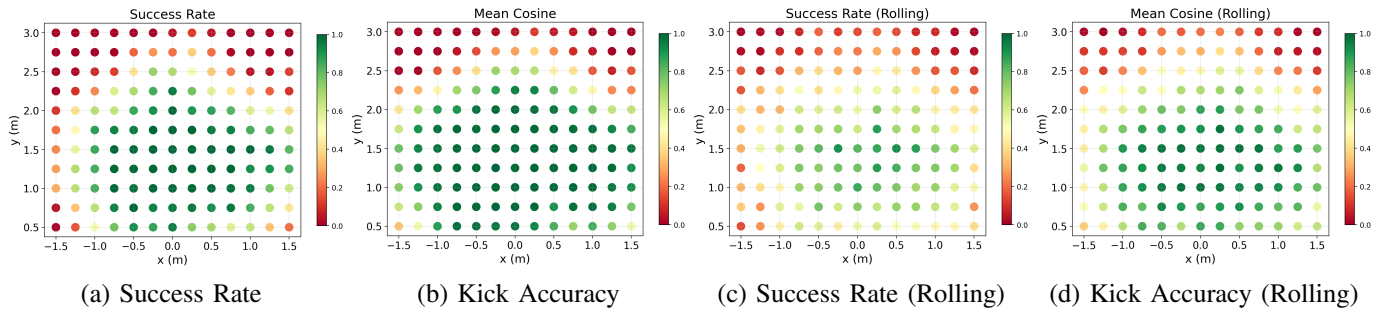(a) Success Rate   (b) Kick Accuracy   (c) Success Rate (Rolling)   (d) Kick Accuracy (Rolling)

Fig. 4: Quantitative analysis of soccer shooting proficiency across the workspace. The heatmaps visualize the spatial distribution of success rates and kicking accuracy for both static ball scenarios (a, b) and dynamic rolling ball interception (c, d).

TABLE III: Comparison of motion tracking performance

Method	MPJPE ↓	G-MPJPE ↓	Vel Err ↓	AUJ Score ↓
	Standard Kicks			
GMT	59.77	195.79	3.99	0.0217
Any2Track	104.20	468.56	8.18	0.1043
TWIST2	30.06	161.96	4.40	0.0158
BeyondMimic	**25.26**	**104.41**	**3.06**	0.0283
Ours (PAiD)	**27.75**	213.38	**3.75**	**0.0129**
	Stylized Kicks			
GMT	64.22	376.45	6.30	**0.0107**
Any2Track	172.23	675.63	11.58	0.5420
TWIST2	41.09	318.93	6.98	0.0216
BeyondMimic	**39.13**	**162.39**	**4.26**	0.0161
Ours (PAiD)	62.01	**230.36**	**5.04**	0.0444

generally benign dynamics; and (ii) *Stylized Kicks* (3 motions), comprising high-dynamic, professional-level maneuvers that impose substantial biomechanical demands on the controller.

We benchmark our framework against state-of-the-art motion tracking approaches, including GMT [5], Any2Track [55], TWIST2 [53], and BeyondMimic [23]. As shown in Table III, our method achieves competitive or leading results across most metrics. While BeyondMimic achieves the best scores on some metrics, our approach provides a balanced trade-off between precision and versatility, supporting robust tracking of a diverse repertoire of soccer skills.

B. Soccer Shooting Proficiency

To answer **Q2**, we systematically evaluate soccer shooting proficiency in Mujoco through two distinct experimental protocols designed to test spatial generalization and dynamic adaptation: (1) *Grid-Based Static Evaluation*: To ensure comprehensive coverage, we discretize the evaluation area into a 11×13 grid, uniformly distributing initial ball positions across the $[0.5, 3.0] \times [-1.5, 1.5]$ m region with 0.25 m spacing. At each grid point, we perform 50 independent trials with randomized robot initializations sampled uniformly within a 0.5 m radius circle to assess robustness. (2) *Rolling Ball Interception*: Based on the first protocol, we introduce dynamic complexity by initializing the ball with a velocity sampled from $[0.1, 0.3]$ m/s. We employ two metrics: *Success Rate*, calculated as the number of goals divided by the total number of kicks, and *Kick Accuracy*, computed as the cosine similarity

TABLE IV: Comparison of soccer shooting proficiency

Method	*Static Evaluation*		*Rolling Interception*	
	Success Rate ↑	Accuracy ↑	Success Rate ↑	Accuracy ↑
Pure RL	33.0%	0.5235	23.5%	0.3592
AMP-based	46.8%	0.7146	35.5%	0.5314
Single-Stage	78.1%	0.9218	54.3%	0.7977
Ours (PAiD)	**91.3%**	**0.9689**	**71.9%**	**0.8892**

$\cos(\theta)$, where θ denotes the angle between the ball's outgoing velocity vector and the target direction vector connecting the ball's position at impact to the center of the goal.

We further provide a spatial performance analysis in Fig. 4. The heatmaps demonstrate that our policy maintains strong performance within the training range and demonstrates some generalization beyond it, though proficiency decreases at the boundaries due to kinematic limits.

In addition, we compare the PAiD framework against three baselines: (i) an *AMP-based* approach trained with adversarial motion priors; (ii) a *Pure RL* baseline trained from scratch using manually designed rewards without reference motions; and (iii) a *Single-Stage* baseline that jointly optimizes tracking and task rewards. Quantitative results are reported in Table IV, evaluated within a ball placement range of $[0.5, 2.0] \times [-1.0, 1.0]$ m, which corresponds to the effective workspace for our generalized motions. Within this region, PAiD achieves a high static success rate of 91.3% and a rolling interception rate of 71.9%, both significantly outperforming the baselines. In contrast, the AMP baseline often produces unnatural kicks, the Pure RL baseline struggles with reward engineering, and the Single-Stage baseline suffers from reward conflicts between locomotion and kicking objectives.

C. Real-World Performance

To answer **Q3**, we deploy our policy on the Unitree G1 platform to evaluate the real-world performance.

a) Deployment: We implement a perception-and-decision framework for real robot deployment. In the perception module, a depth camera D435i is employed to continuously obtain the ball's relative position in the robot frame using YOLOv8 to obtain the ball's mask [52] when the ball is within the camera's field of view. To address out-of-view scenarios, Fast-LIO [47] with a LiDAR sensor is utilized to estimate the robot's global pose, enabling transformation

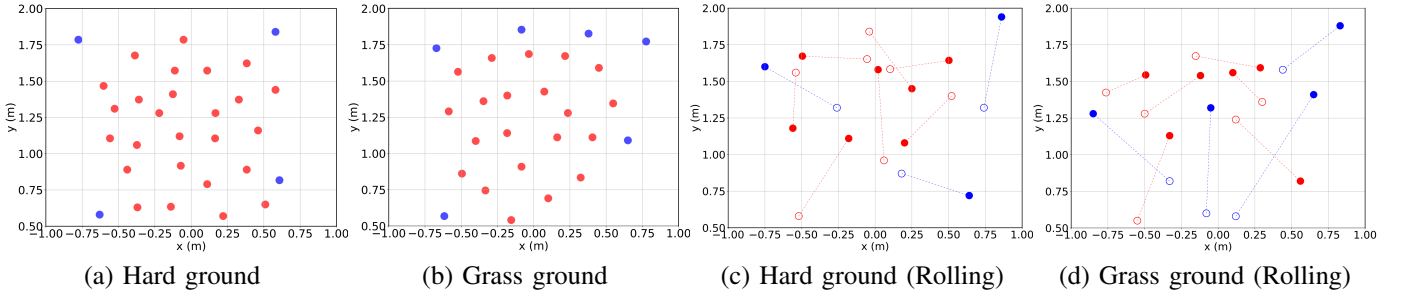(a) Hard ground   (b) Grass ground   (c) Hard ground (Rolling)   (d) Grass ground (Rolling)

Fig. 5: Real-world shooting tests with randomly placed soccer balls. (a)–(b) show results on hard ground and grass, with 30 trials per surface. (c)–(d) show tests on rolling balls, where hollow and solid markers denote the rolling start and end positions, respectively, with 10 trials per surface. Red and blue dots indicate success and failure.
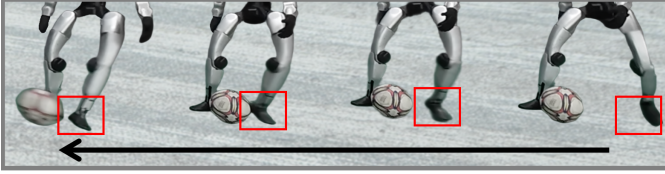

Fig. 6: Detailed foot motions during ball kicking.

of the ball's relative position to the world frame. When the ball leaves the camera view, its previous world-frame position is used to recover the relative ball position. This integrated RGB-D and LiDAR pipeline provides robust ball localization throughout the shooting process. In the decision module, we adopt a simple yet effective motion selection mechanism: Before shooting starts, a set of candidate motions is preloaded. The motion whose final-frame robot root position is closest to the estimated ball position is selected and provided as input to a single policy, which is then executed.

b) Soccer Shooting Proficiency and Terrain Robustness Test: In Fig. 1 (a), (c), and (d), we set up a goal with a width of 2 m and a height of 1.5 m in the real world. During the experiments, the robot is positioned 5 m in front of the goal, and a soccer ball is randomly placed within an area in front of the robot that roughly matches the depth camera's field of view, across multiple trials. Fig. 5 shows that our method achieves a high success rate for most ball locations on both hard ground and grass surfaces, except for a few corner cases where the success rate is lower, as these positions require the robot to perform large body rotations or travel a longer distance. During testing, the robot successfully kicked soccer balls placed at random positions for up to 11 consecutive trials, demonstrating the robustness of the proposed policy. As shown in Fig. 6, at ball–foot impact, the foot achieves a well-formed strike through the foot arch, which underlies the high shooting success rate. In particular, Fig. 5 (c) and (d) demonstrate that our method enables the real robot to successfully kick rolling soccer balls, which is a critical capability for practical soccer gameplay.

We also conduct ablation studies on two core components of our physics-aware sim-to-real transfer strategy on the real robot, as summarized in Table V. Specifically, we evaluate the impact of contact dynamics identification and observation noise by selectively removing each component during training.

TABLE V: Ablation experiment of physics-aware sim-to-real transfer strategy

Method	Hard Ground			Grass Ground		
	(-0.6,0.8)[1]	(0.0,0.8)	(0.6,0.8)	(-0.6,0.8)	(0.0,0.8)	(0.6,0.8)
Ours	**5/5**[2]	**5/5**	**5/5**	**4/5**	**5/5**	**5/5**
w/o SystemID	4/5	5/5	4/5	2/5	3/5	1/5
w/o Obs Noise	2/5	3/5	3/5	3/5	3/5	2/5

[1] The values indicate the soccer ball placement position, measured in meters.
[2] For each placement position, we conducted 5 trials and recorded the successful number.

The results indicate that incorporating either component leads to noticeable performance gains in real-world deployment, while combining both yields the best overall results. These findings demonstrate that accurate modeling of contact dynamics, together with realistic observation noise, plays a crucial role in bridging the sim-to-real gap, enabling the learned policy to robustly adapt to varying ground conditions encountered in practice.

c) Motion Human-Likeness Test: We evaluate three professional-player–style shooting motions on the real robot, with the soccer ball placed in the vicinity of the original ball positions defined by each motion. The qualitative results are reported in Fig. 1 (b). Despite the inherent complexity and coordination required by these highly human-like motions, the robot consistently executes them in a stable and controlled manner, achieving accurate and goal-directed shots. These results demonstrate the high fidelity, expressiveness, and scalability of our method in transferring diverse motions to the real robot.

V. CONCLUSION

In this work, we presented our Perception-Action integrated Decision-making (PAiD) framework for learning robust humanoid soccer skills. By decoupling motion-skill acquisition, perception-action integration, and sim-to-real transfer, our approach overcomes the limitations of single-stage RL and imitation learning. Extensive experiments demonstrate that PAiD achieves superior tracking accuracy, shooting proficiency, and real-world robustness compared to state-of-the-art baselines.

While PAiD shows strong generalization and transferability, its performance at extreme ball positions and under highly dynamic conditions can be further improved. Future work will explore multi-modal perception and adaptive policy architectures to enhance robustness in more complex environments.

REFERENCES

[1] M. Abreu, L. P. Reis, and N. Lau. Learning to run faster in a humanoid robot soccer environment through reinforcement learning. In *Robot World Cup (RoboCup)*, 2019.

[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. *ArXiv*, abs/2505.03729, 2025.

[3] Hongjun An, Wenhan Hu, Sida Huang, Siqi Huang, Ruanjun Li, Yuanzhi Liang, Jiawei Shao, Yiliang Song, Zihan Wang, Cheng Yuan, Chi Zhang, Hongyuan Zhang, Wenhao Zhuang, and Xuelong Li. Ai flow: Perspectives, scenarios, and approaches, 2025. URL https://arxiv.org/abs/2506.12479.

[4] Joao Pedro Araujo, Yanjie Ze, Pei Xu, Jiajun Wu, and C. Karen Liu. Retargeting matters: General motion retargeting for humanoid motion tracking. *ArXiv*, 2025.

[5] 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.

[6] 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.

[7] Xinyang Gu, Yen-Jen Wang, and Jianyu Chen. Humanoid-gym: Reinforcement learning for humanoid robot with zero-shot sim2real transfer. *arXiv preprint arXiv:2404.05695*, 2024.

[8] Tuomas Haarnoja, Ben Moran, Guy Lever, Sandy H Huang, Dhruva Tirumala, Jan Humplik, Markus Wulfmeier, Saran Tunyasuvunakool, Noah Y Siegel, Roland Hafner, et al. Learning agile soccer skills for a bipedal robot with deep reinforcement learning. *Science Robotics*, 9(89):eadi8022, 2024.

[9] 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.

[10] Nikolaus Hansen and Andreas Ostermeier. Completely derandomized self-adaptation in evolution strategies. In *Evolutionary Computation*, 2001.

[11] 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 *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 8944–8951. IEEE, 2024.

[12] Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbab, Chaoyi Pan, et al. Asap: Aligning simulation and real-world physics for learning agile humanoid whole-body skills. *arXiv preprint arXiv:2502.01143*, 2025.

[13] Xialin He, Runpei Dong, Zixuan Chen, and Saurabh Gupta. Learning getting-up policies for real-world humanoid robots. *arXiv preprint arXiv:2502.12152*, 2025.

[14] 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. *arXiv preprint arXiv:2502.08378*, 2025.

[15] Xiaoyu Huang, Zhongyu Li, Yan-Ling Xiang, Yiming Ni, Yufeng Chi, Yunhao Li, Lizhi Yang, Xue Bin Peng, and Koushil Sreenath. Creating a dynamic quadrupedal robotic goalkeeper with reinforcement learning. *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2715–2722, 2022.

[16] Y Ji, G. B. Margolis, and P. Agrawal. Dribblebot: Dynamic legged manipulation in the wild. *arXiv preprint arXiv:2304.01159*, 2023.

[17] Yandong Ji, Zhongyu Li, Yinan Sun, Xue Bin Peng, Sergey Levine, Glen Berseth, and Koushil Sreenath. Hierarchical reinforcement learning for precise soccer shooting skills using a quadrupedal robot. *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1479–1486, 2022.

[18] Nan Jiang, Zimo He, Wanhe Yu, Lexi Pang, Yunhao Li, Hongjie Li, Jieming Cui, Yuhan Li, Yizhou Wang, Yixin Zhu, et al. Uniact: Unified motion generation and action streaming for humanoid robots. *arXiv preprint arXiv:2512.24321*, 2025.

[19] Hiroaki Kitano, Minoru Asada, Yasuo Kuniyoshi, Itsuki Noda, and Eiichi Osawa. Robocup: The robot world cup initiative. In *Proceedings of the first international conference on Autonomous agents*, pages 340–347, 1997.

[20] 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. *arXiv preprint arXiv:2506.08931*, 2025.

[21] Zhongyu Li, Xue Bin Peng, Pieter Abbeel, Sergey Levine, Glen Berseth, and Koushil Sreenath. Robust and versatile bipedal jumping control through reinforcement learning. *arXiv preprint arXiv:2302.09450*, 2023.

[22] Qiayuan Liao, Bike Zhang, Xuanyu Huang, Xiaoyu Huang, Zhongyu Li, and Koushil Sreenath. Berkeley humanoid: A research platform for learning-based control. *arXiv preprint arXiv:2407.21781*, 2024.

[23] 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.

[24] Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. *arXiv preprint arXiv:2411.14386*, 2024.

[25] Chenhao Lu, Xuxin Cheng, Jialong Li, Shiqi Yang, Mazeyu Ji, Chengjing Yuan, Ge Yang, Sha Yi, and Xiaolong Wang. Mobile-television: Predictive motion

priors for humanoid whole-body control. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 5364–5371. IEEE, 2025.

[26] Jiaxin Lu, Chun-Hao Paul Huang, Uttaran Bhattacharya, Qixing Huang, and Yi Zhou. Humoto: A 4d dataset of mocap human object interactions. *ArXiv*, abs/2504.10414, 2025.

[27] Naureen Mahmood, Nima Ghorbani, Nikolaus F. Troje, Gerard Pons-Moll, and Michael J. Black. Amass: Archive of motion capture as surface shapes. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 5441–5450, 2019.

[28] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science robotics*, 7(62):eabk2822, 2022.

[29] NVIDIA. Isaac Sim. URL https://github.com/isaac-sim/IsaacSim.

[30] 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.

[31] Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. Amp: Adversarial motion priors for stylized physics-based character control. *ACM Transactions on Graphics (ToG)*, 40(4):1–20, 2021.

[32] Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Conference on robot learning*, pages 91–100. PMLR, 2022.

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

[34] Jiyuan Shi, Xinzhe Liu, dewei wang, ouyang lu, Sören Schwertfeger, Chi Zhang, Fuchun Sun, Chenjia Bai, and Xuelong Li. Adversarial locomotion and motion imitation for humanoid policy learning. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=hVoIz6xD9Q.

[35] Zhi Su, Yuman Gao, Emily Lukas, Yunfei Li, Jiaze Cai, Faris Tulbah, Fei Gao, Chao Yu, Zhongyu Li, Yi Wu, and Koushil Sreenath. Toward real-world cooperative and competitive soccer with quadrupedal robot teams. *ArXiv*, abs/2505.13834, 2025.

[36] Zhi Su, Yuman Gao, Emily Lukas, Yunfei Li, Jiaze Cai, Faris Tulbah, Fei Gao, Chao Yu, Zhongyu Li, Yi Wu, et al. Toward real-world cooperative and competitive soccer with quadrupedal robot teams. *arXiv preprint arXiv:2505.13834*, 2025.

[37] 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*, abs/2508.21043, 2025.

[38] Dhruva Tirumala, Markus Wulfmeier, Ben Moran, Sandy H. Huang, Jan Humplik, Guy Lever, Tuomas Haarnoja, Leonard Hasenclever, Arunkumar Byravan, Nathan Batchelor, Neil Sreendra, Kushal Patel, Marlon Gwira, Francesco Nori, Martin A. Riedmiller, and Nicolas Manfred Otto Heess. Learning robot soccer from egocentric vision with deep reinforcement learning. In *Conference on Robot Learning*, 2024.

[39] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 5026–5033. IEEE, 2012. doi: 10.1109/IROS.2012.6386109.

[40] Dewei Wang, Xinmiao Wang, Xinzhe Liu, Jiyuan Shi, Yingnan Zhao, Chenjia Bai, and Xuelong Li. More: Mixture of residual experts for humanoid lifelike gaits learning on complex terrains, 2025. URL https://arxiv.org/abs/2506.08840.

[41] Yushi Wang, Changsheng Luo, Penghui Chen, Jianran Liu, Weijian Sun, Tong Guo, Kechang Yang, Biao Hu, Yangang Zhang, and Mingguo Zhao. Learning vision-driven reactive soccer skills for humanoid robots. *arXiv preprint arXiv:2511.03996*, 2025.

[42] Zhuoheng Wang, Jinyin Zhou, and Qi Wu. Dribble master: Learning agile humanoid dribbling through legged locomotion. *ArXiv*, abs/2505.12679, 2025.

[43] Haoyang Weng, Yitang Li, Nikhil Sobanbabu, Zihan Wang, Zhengyi Luo, Tairan He, Deva Ramanan, and Guanya Shi. Hdmi: Learning interactive humanoid whole-body control from human videos. *arXiv preprint arXiv:2509.16757*, 2025.

[44] Weiji Xie, Chenjia Bai, Jiyuan Shi, Junkai Yang, Yunfei Ge, Weinan Zhang, and Xuelong Li. Humanoid whole-body locomotion on narrow terrain via dynamic balance and reinforcement learning. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2025.

[45] 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. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025.

[46] Sirui Xu, Hung Yu Ling, Yu-Xiong Wang, and Liang-Yan Gui. Intermimic: Towards universal whole-body control for physics-based human-object interactions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12266–12277, June 2025.

[47] Wei Xu and Fu Zhang. Fast-lio: A fast, robust lidar-inertial odometry package by tightly-coupled iterated kalman filter. *IEEE Robotics and Automation Letters*, 6(2):3317–3324, 2021. doi: 10.1109/LRA.2021.3064227.

[48] Zifan Xu, Myoungkyu Seo, Dongmyeong Lee, Hao Fu, Jiaheng Hu, Jiaxun Cui, Yuqian Jiang, Zhihan Wang, Anastasiia Brund, Joydeep Biswas, and Peter Stone.

Learning agile striker skills for humanoid soccer robots from noisy sensory input, 2025. URL https://arxiv.org/abs/2512.06571.

[49] 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.

[50] Yufei Xue, Wentao Dong, Minghuan Liu, Weinan Zhang, and Jiangmiao Pang. A unified and general humanoid whole-body controller for versatile locomotion. *arXiv preprint arXiv:2502.03206*, 2025.

[51] 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.

[52] Muhammad Yaseen. What is yolov8: An in-depth exploration of the internal features of the next-generation object detector, 2024. URL https://arxiv.org/abs/2408.15857.

[53] Yanjie Ze, Siheng Zhao, Weizhuo Wang, Angjoo Kanazawa, Rocky Duan, Pieter Abbeel, Guanya Shi, Jiajun Wu, and C Karen Liu. Twist2: Scalable, portable, and holistic humanoid data collection system. *arXiv preprint arXiv:2511.02832*, 2025.

[54] He Zhang, Shenghao Ren, Haolei Yuan, Jianhui Zhao, Fan Li, Shuangpeng Sun, Zhenghao Liang, Tao Yu, Qiu Shen, and Xun Cao. Mmvp: A multimodal mocap dataset with vision and pressure sensors. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21842–21852, 2024.

[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] Siheng Zhao, Yanjie Ze, Yue Wang, C. Karen Liu, Pieter Abbeel, Guanya Shi, and Rocky Duan. Resmimic: From general motion tracking to humanoid whole-body loco-manipulation via residual learning. *ArXiv*, abs/2510.05070, 2025.

[57] Yingnan Zhao, Xinmiao Wang, Dewei Wang, Xinzhe Liu, Dan Lu, Qilong Han, Peng Liu, and Chenjia Bai. Towards adaptive humanoid control via multi-behavior distillation and reinforced fine-tuning. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*, 2026.

[58] Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. *arXiv preprint arXiv:2406.10759*, 2024.

A. Data Preparation for Policy Learning

a) Motion Dataset Collection: The motion dataset consists of a curated set of human demonstration trajectories, designed to provide representative spatial and stylistic coverage for policy learning. In line with the main text, the dataset is structured into two categories: (i) *standard kicks*, which are collected with systematically varied ball placements to span a range of distances and approach angles; and (ii) *stylized kicks*, which capture high-difficulty, professional-player–inspired maneuvers. For each motion, the striking leg is selected based on the ball's relative position to ensure biomechanical plausibility.

- **Distribution and Coverage:** The standard kicks are sampled to achieve broad, but not exhaustive, coverage of typical ball positions and approach angles encountered in soccer. The resulting dataset spans a substantial portion of the robot's effective workspace, as visualized by the motion trajectory spread in Fig. 7.
- **Stylized Motions:** The stylized kicks are selected to evaluate the policy's ability to reproduce challenging, high-velocity, and stylistically distinct actions. Their spatial patterns are shown in Fig. 8.
- **Trajectory Characteristics:** Each trajectory records the full 3D pose of the robot at every frame, including root and limb kinematics, enabling precise motion retargeting and imitation. The temporal color gradient in the figures highlights the dynamic evolution of each kick.

These visualizations and statistics provide further insight into the spatial and stylistic diversity of the dataset, as well as the challenges posed to the learning algorithm.
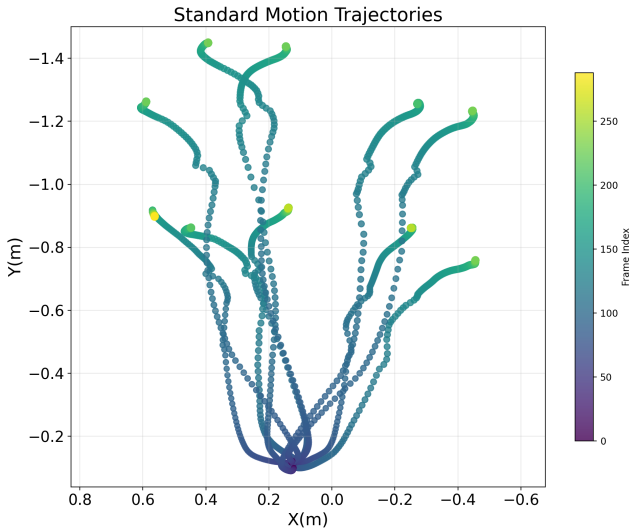

Fig. 7: Root trajectories of all standard kicks in the dataset.

b) Ball Placement and Generalization: The ball placement strategy is critical for enabling the policy to generalize beyond the original motion dataset. Consistent with the methodology described in the main text, the initial ball
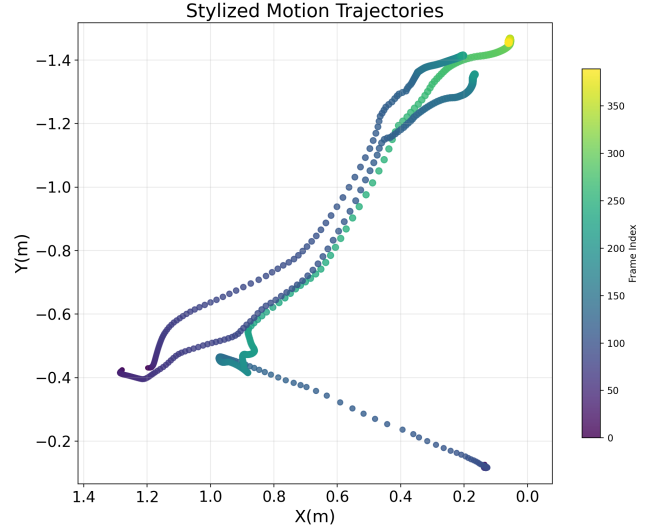

Fig. 8: Root trajectories of all stylized kicks in the dataset.

positions are anchored to the terminal foot contact locations of the reference motions, ensuring biomechanical feasibility and natural transitions. To further enhance spatial generalization, we randomize the ball position within a set of angular and radial sectors centered in front of the robot, as illustrated in Fig. 9.
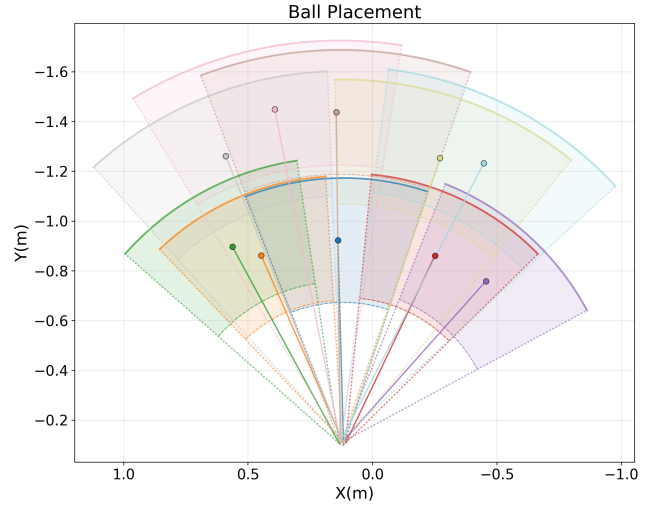

Fig. 9: Schematic of the ball placement strategy. Colored sectors represent the feasible regions for each reference motion, with solid dots indicating nominal motion positions. The randomization process samples ball positions within these regions, enabling the policy to generalize across a broad workspace.

Each colored sector in the figure corresponds to the feasible region associated with a specific reference motion, while the solid dots indicate the nominal ball positions used in the original motions. During training, ball positions are sampled within these sectors, exposing the policy to a diverse set of approach

TABLE VI: Ball Physical Parameters for CMA-ES Optimizer.

Parameter	Initial Value	Search Range
Static friction	0.5	$[0, 1.0]$
Dynamic friction	0.5	$[0, 1.0]$
Restitution	0.5	$[0, 1.0]$
Linear damping	1.0	$[0, 5.0]$
Angular damping	1.0	$[0, 5.0]$

TABLE VII: Hyperparameter Settings of the CMA-ES Optimizer.

Hyperparameter	Value
Search scale	0.2
Population size	4

angles and distances. This design ensures that the learned policy is not restricted to a discrete set of placements, but can robustly handle balls distributed throughout the effective workspace defined by the motion dataset.

By combining motion-conditioned placement with controlled randomization, our approach achieves a balance between physical plausibility and spatial diversity, facilitating both high-fidelity imitation and strong generalization.

B. Details of Contact Dynamic Identification

In this section we provide the details of how to identify the contact dynamic parameters.

To ensure that the simulated soccer ball exhibits physical behaviors consistent with real-world dynamics, we optimize five physical parameters of the ball asset in the simulator, including *static friction*, *dynamic friction*, *restitution*, *linear damping*, and *angular damping*. Notably, the ground-related parameters are kept at their default values, as the contact interaction between the ball and the ground can be equivalently adjusted by tuning the corresponding coefficients of the ball alone.

In the real-world experiments, a high-frame-rate camera is employed to record the ball's motion trajectories.

- **Ball drop experiment:** We record the initial height of the ball and its height at intervals of $0.1\,\text{s}$ after release, collecting measurements over a duration of $2\,\text{s}$.
- **Ball rolling experiment:** The initial velocity of the ball is estimated from the first few frames at the onset of rolling. Subsequently, the traveled distance is recorded at intervals of $0.1\,\text{s}$ for a total duration of $2\,\text{s}$.

To mitigate measurement noise and experimental variability, each experiment is repeated five times. The resulting trajectories are treated as ground-truth data for subsequent parameter optimization.

In the simulator, we adopt the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) for parameter optimization, which is a population-based, gradient-free optimization algorithm well suited for non-convex and black-box problems. The hyperparameter settings of CMA-ES are summarized in Table VI and VII, and the overall optimization procedure is illustrated in Algorithm 0.

Algorithm 1 Parameter Identification via CMA-ES

1: **Input:** Real-world trajectories $\mathbf{h} = \{h_i\}_{i=0}^{N_d}$, $\mathbf{d} = \{d_i\}_{i=0}^{N_r}$
2: set initial parameter x_0, initial search scale σ_0
3: set parameter bounds and population size P
4: set tolerance ϵ (loss change)
5: Initialize CMA-ES with $(x_0, \sigma_0, \text{bounds}, P)$
6: $L_{\text{prev}} \leftarrow +\infty$
7: **for** $t = 1 \rightarrow T$ **do**
8: $\quad \{x^{(j)}\}_{j=1}^{P} \leftarrow \text{ask}()$
9: $\quad \text{losses} \leftarrow [\,]$
10: $\quad$ **for** $j = 1 \rightarrow P$ **do**
11: $\quad\quad$ Set simulation parameters $\leftarrow x^{(j)}$
12: $\quad\quad$ Reset environment and ball
13: $\quad\quad$ Perform drop and rolling rollouts
14: $\quad\quad$ Obtain simulated trajectories
$$\mathbf{h}' = \{h_i'\}_{i=0}^{N_d}, \quad \mathbf{d}' = \{d_i'\}_{i=0}^{N_r}$$
15: $\quad\quad \ell \leftarrow \mathcal{L}_{\text{sysid}}(\mathbf{h}', \mathbf{d}', \mathbf{h}, \mathbf{d}) \qquad \triangleright \text{see Eq. (4)}$
16: $\quad\quad$ append ℓ to losses
17: $\quad$ **end for**
18: $\quad \text{tell}(\{x^{(j)}\}, \text{losses})$
19: $\quad L_{\text{best}} \leftarrow \min(\text{losses})$
20: $\quad$ **if** $|L_{\text{prev}} - L_{\text{best}}| < \epsilon$ **then**
21: $\quad\quad$ **break**
22: $\quad$ **end if**
23: $\quad L_{\text{prev}} \leftarrow L_{\text{best}}$
24: **end for**
25: **return** optimized parameters $x^\star$ (best CMA-ES solution)

We conduct data collection and system identification separately on rigid hard ground and soccer grass surfaces. The identified parameter sets for these two surface types are reported in Table VIII.

TABLE VIII: Optimized soccer ball physical parameters on different surface types.

Parameter	Hard Ground	Grass Surface
Static Friction	0.77	0.98
Dynamic Friction	0.07	0.15
Restitution	0.75	0.71
Linear Damping	0.01	0.01
Angular Damping	4.28	4.95