

A Biomechanics-Inspired Approach to Soccer Kicking for Humanoid Robots

Daniel Marew¹, Nisal Perera¹, Shangqun Yu¹, Sarah Roelker², and Donghyun Kim¹

Abstract—Soccer kicking is a complex whole-body motion that requires intricate coordination of various motor actions. To accomplish such dynamic motion in a humanoid robot, the robot needs to simultaneously: 1) transfer high kinetic energy to the kicking leg, 2) maintain balance and stability of the entire body, and 3) manage the impact disturbance from the ball during the kicking moment. Prior studies on robotic soccer kicking often prioritized stability, leading to overly conservative quasi-static motions. In this work, we present a biomechanics-inspired control framework that leverages trajectory optimization and imitation learning to facilitate highly dynamic soccer kicks in humanoid robots. We conducted an in-depth analysis of human soccer kick biomechanics to identify key motion constraints. Based on this understanding, we designed kinodynamically feasible trajectories that are then used as a reference in imitation learning to develop a robust feedback control policy. We demonstrate the effectiveness of our approach through a simulation of an anthropomorphic 25 DoF bipedal humanoid robot, named PresToe, which is equipped with 7 DoF legs, including a unique actuated toe. Using our framework, PresToe can execute dynamic instep kicks, propelling the ball at speeds exceeding 11 m/s in full dynamics simulation.

I. INTRODUCTION

The RoboCup Federation has established an ambitious goal for 2050: to develop a humanoid soccer team capable of defeating the human FIFA World Cup champions of that year [1], [2]. For this dream to become a reality and to effectively compete with humans, the robots must not only walk, run, and leap, but also master the art of kicking a ball as swiftly and as accurately as humans. Achieving a powerful and accurate kick in bipedal robots is a complex endeavor. It requires the substantial transfer of kinetic energy to the kicking leg, all while maintaining balance with only one foot grounded, and ensuring effective recovery after ball impact. However, existing studies on robotic soccer kicks have not addressed the full spectrum of the challenges associated with such dynamic motion. For example, [3]–[5] anchored the robot to the ground, thereby circumventing the challenges of balance and post-kick recovery. While this allows researchers to concentrate on the dynamic interaction between the foot and ball during impact, the results are hard to extend to biped robots with significant under-actuation. In contrast, [6]–[10] explore soccer kick motions using fully mobile humanoid robots. However, these approaches tend to favor conservative stationary kick movements, resulting in less powerful kicks to guarantee stability both before and after ball contact.

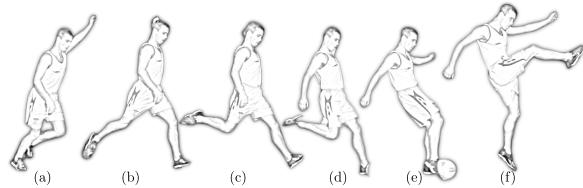


Fig. 1. Phases of an in-step soccer kick: (a) Approach/Run-up, (b) Planting/Support, (c) Wind-up/Backswing, (d) Cocking, (e) Swing/Acceleration, and (f) Follow-through. Adapted from [15].

One common missing perspective in both approaches is the lack of focus on momentum-building preparatory steps before a kick. In human soccer, players use such steps to accumulate momentum, enabling more powerful strikes. Many robotic approaches overlook this critical aspect, often keeping the robot's supporting leg stationary, which limits the kinetic energy transferred to the ball at impact. There exists only a handful of research utilizing stepping before kicking [11]–[13]. However, [11], [12] used the walking step mainly to reposition the robot rather than to build momentum for a stronger kick. A video demonstration by Honda Research featuring the Asimo robot [13] shows more powerful kicks achieved in part through preparatory steps. However, the robot's motion lacks key human-like traits of an instep kick: it omits arm and torso movements crucial for angular momentum regulation and balance, and follows a straight-line trajectory instead of a human's typical curved path that is associated with more powerful kicks [14]. Consequently, the kick resembles a walk-and-kick motion rather than a dynamic human instep kick.

Often, Biomechanics provides interesting perspectives and useful information for robot control. The previous studies of soccer kick biomechanics identified several key attributes that determine the power and quality of a kick [14], [16]. Among these, the foremost determinant of kick power is the kinodynamic whip action [14], [16]. This action results from the initial loading of the trunk, hip, and knee joints, often referred to as the tension arc [14], [16]. Potential energy stored in the hip flexors, quadriceps, hamstrings, and calf muscles is then quickly released in a coordinated manner, moving from proximal to distal parts. This sequential action, resembling a whip, transfers large kinetic energy to the kicking foot and subsequently to the ball.

Another key attribute is the approach angle towards the ball, which for professional players typically ranges between 24° and 43° [14], [16], [17]. This approach angle enhances

¹ College of Information and Computer Sciences, University of Massachusetts Amherst, MA, U.S. (donghyunkim@cs.umass.edu)

² School of Public Health and Health Sciences, University of Massachusetts Amherst, MA, U.S.

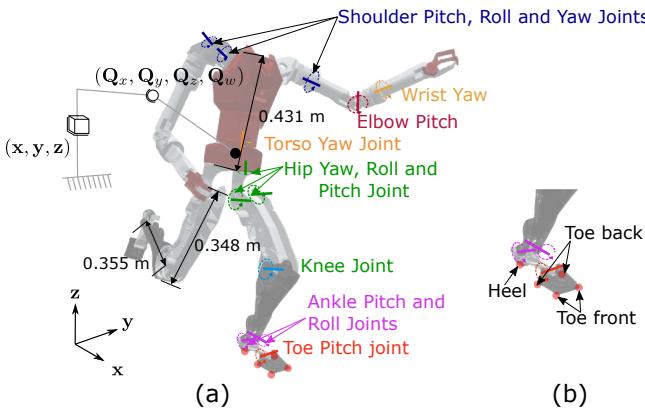


Fig. 2. PresToe: (a) Robot degrees of freedom (b) Foot contact points

kicking power by permitting a greater range of motion for trunk rotation, thus leading to a more substantial loading of the trunk prior to the whip action’s onset [16]. Furthermore, an equally pivotal attribute is the placement of the supporting foot. In this regard, expert advice emphasizes positioning the foot next to and slightly behind the ball, parallel to the desired kick direction. Moreover, the length of the step taken immediately before planting the foot strongly correlates with kick power, with longer strides resulting in more powerful kicks [14], [16]. Lastly, the ankle locks before impact to ensure a firmer surface contact, which, in turn, correlates with greater power transfer [18].

In this work, we focus on the instep soccer kick which is, a powerful and widely studied skill, comprises six stages as illustrated in Fig. 1: (a) the approach or run-up, where the player advances towards the ball; (b) the planting or support phase, involving placement of the non-kicking foot; (c) the wind-up or backswing, where the kicking leg is retracted; (d) the cocking phase, characterized by forward thigh motion with a flexed knee; (e) the swing or acceleration phase, involving rapid hip flexion and knee extension; and (f) the follow-through, continuing the leg’s forward motion after ball impact [14], [16], [18]–[22].

Our strategy to implement the principles learned from biomechanics into robot control is developing a novel motion planning and control framework that combines motion retargeting, kinodynamic trajectory optimization, and imitation learning. Our framework is divided into two parts: 1) motion planning based on human motion capture data, and 2) imitation learning to find robust control policy using the found motions as reference. In the motion planning phase, we use kinodynamic trajectory optimization [23] to properly account for the system’s physical limits, which are often ignored in human-to-robot motion retargeting studies [24]–[27]. Considering the robot’s dynamics is crucial not only because the resultant motions are within the joint torque/velocity/position limits, but also due to the balance stability. [28] demonstrated effective retargeting of human dance motion by using a Linear Inverted Pendulum Model (LIPM) of the robot and enforcing Zero Moment Point (ZMP) constraints. Furthermore, [29] showed that incorpo-

rating dynamics leads to more realistic-looking retargeting of human motion to animated characters. Their approach uses a two-stage process: first, obtaining a trajectory based on a Single Rigid Body (SRB) model, then performing inverse kinematics to recover a kinematic trajectory consistent with the Center of Mass (CoM) and foot contact locations.

Our approach is comprehensive, using kinodynamic trajectory optimization that leverages the centroidal dynamics of the robot, which are more expressive than SRB or LIPM approximations. Additionally, our resulting kinematic trajectory aligns not only with the CoM and foot contact positions but also with linear and angular momenta, yielding realistic and expressive motion retargeting. Beyond modeling and optimization, we incorporate biomechanical observations from the six phases of the instep kick as constraints and cost terms in our trajectory optimization formulation, complementing the implicit priors obtained from the MoCap data.

After obtaining optimal kicking trajectories that include essential preparatory momentum-building steps, we train a reinforcement learning control policy with an imitation objective. This policy aims to track the trajectories obtained in the motion planning phase while maintaining the robot’s balance in Isaac Gym [30] – accounting for factors such as the dynamic interaction between the robot and ball.

The main contributions of this paper are as follows: 1) We present a new motion planning and imitation learning frameworks for performing human-like soccer kicks by identifying key constraints and insights from the biomechanics of soccer kicks. 2) We showcase instep kicks that propel the ball at speeds exceeding 11 m/s. This is almost double the 6.6 m/s instep kick achieved by the ground-anchored single-leg robot [4]. Notably, the demonstrated ball speed is about 40% of the speed an average human player can achieve [31], which is an impressive result considering the robot’s size and limited torque capability. The robot weighs 30kgs and has a maximum hip pitch torque of only 48 N.m, compared to the average human player’s hip joint flexion/extension capability of over 250 Nm [32]. 3) We highlight the importance of proper consideration of dynamics in the offline motion planning stage for sample-efficient imitation learning. This is demonstrated by comparing purely kinematics-based reference motions with kinodynamically consistent ones.

II. METHODOLOGY

A. Offline Kicking Trajectory Generation

Our objective is to emulate the powerful kicks demonstrated by humans using a robotic platform. While earlier works attempted to manually craft kicking motions from scratch, we believe in harnessing the natural expertise of humans [6]–[10]. To this end, we utilize MoCap data derived from humans executing soccer kicks as a reference for generating robot trajectories that are consistent with the robot’s kinematic and torque capability. Our methodology comprises four stages. Initially, we preprocess the MoCap data to ensure it aligns with the robot’s morphology. Subsequently, using inverse kinematics, we derive purely kinematic joint

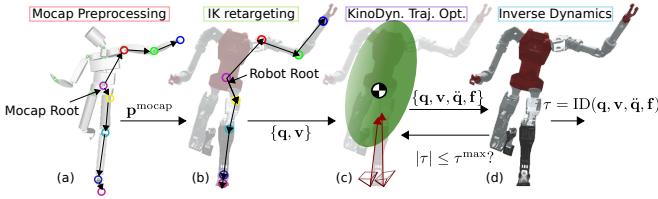


Fig. 3. Offline Motion Planning Steps: (a) MoCap Link Rescaling: Corresponding joints are highlighted in the same color (on the MoCap skeleton in (a) and robot in (b)). For clarity, only the left-side joint correspondence between the MoCap data and the robot is shown. (b) Inverse kinematics-based retargeting, ignoring dynamics. (c) Kinodynamic trajectory optimization. (d) Full-body dynamics-based torque limit violation verification.

position and velocity trajectories that closely mirror the link position trajectory present in the MoCap data. These kinematic trajectories serve both as a reference and an initial guess to inform the more complex, non-linear kinodynamic trajectory optimization problem. Here, we generate trajectories that are consistent with the robot's centroidal dynamics and kinematics. Finally, we use inverse dynamics to verify that the robot's torque limits are not violated, and we make necessary adjustments to the trajectory if any violations are detected. A detailed depiction of our framework is illustrated in Fig. 3.

1) *MoCap Data Preprocessing*: To derive a robot trajectory compatible with the robot's morphology, we preprocess the original MoCap data by adjusting link lengths to align with the robot's structure. We initiate this process by establishing a correspondence between human and robot joints, as shown in Fig. 3 (a). Starting from the root joint, we recursively scale the link lengths in the MoCap data to achieve a match. Moreover, during this stage, we establish a ground foot contact schedule and other important timing parameters. Specifically, we identify the beginning of the swing phase (T^{swing}), the ankle locking time (T^{lock}), and the end of the ball impact phase (T^{impact}). We derive the ground contact schedule by monitoring when the feet initiate or break contact with the ground, using heuristics based on foot height and velocity thresholds [29]. This information is vital for subsequent stages as it enables the imposition of constraints that prevent unwanted artifacts such as foot skating.

2) *Kinodynamic Trajectory Optimization*: We use kinodynamic trajectory optimization, first proposed by [23], to generate physically consistent human-like soccer kicking trajectories, using the solution from the prior step as a reference. This approach, based on the robot's centroidal dynamics and full-body kinematics, strikes a balance between full-body dynamics and reduced-order models like Linear Inverted Pendulum (LIP) or Single Rigid Body (SRB) based approaches. It has fewer non-linearities compared to full-body dynamics-based trajectory optimization, hence can be solved more efficiently while being more realistic and expressive than LIP or SRB models, which rely on restrictive assumptions [23], [33]. In kinodynamic trajectory optimization, the centroidal dynamics of the robot, given by Eqs. (1)

and (2), relates the ground reaction forces to the robot's linear and angular momentum, while the kinematics model is used to determine a corresponding kinematic trajectory consistent with the center of mass, contact location, and momenta. Thus, the optimization encompasses centroidal variables such as the position $\mathbf{r} \in \mathbb{R}^3$, velocity $\dot{\mathbf{r}} \in \mathbb{R}^3$, and acceleration $\ddot{\mathbf{r}} \in \mathbb{R}^3$ of the robot's center of mass (CoM), the robot's centroidal angular momentum (CAM) $\mathbf{h} \in \mathbb{R}^3$ and its rate $\dot{\mathbf{h}} \in \mathbb{R}^3$, the location of contact points $\mathbf{c}_i \in \mathbb{R}^3$, as well as the external ground reaction forces acting at the contact points $\mathbf{f}_i \in \mathbb{R}^3$. The optimization also involves kinematic variables: joint position $\mathbf{q} \in \mathbb{R}^{n+7}$ and joint velocity $\mathbf{v} \in \mathbb{R}^{n+6}$, where n is the number of joints of the robot. The additional seven dimensions in the position vector refer to the body position $\in \mathbb{R}^3$ and orientation represented by a unit quaternion $\in \mathbb{R}^4$, while the velocity vector includes six dimensions for the body's linear and angular velocities.

$$m\ddot{\mathbf{r}} = \sum_{i=1}^{n_c} \mathbf{f}_i + mg \quad (1)$$

$$\dot{\mathbf{h}} = \sum_{i=1}^{n_c} (\mathbf{c}_i - \mathbf{r}) \times \mathbf{f}_i \quad (2)$$

Here, n_c represents the number of contact points. For the PresToe robot (see Fig. 2) in this study, we utilize 10 contact points: 5 for each foot. Specifically, there are two contact points on the front part of the toe, two on the back part, and one for the heel, as depicted in Fig. 2(b). Following [23], the continuous dynamics of the robot is sampled at regular intervals of Δt , in order to transcribe the optimization problem into a nonlinear programming problem. Equations 1-2 describe the discretized linear and angular centroidal dynamics of the robot, respectively at timestep k . The optimization objective comprises several components. Firstly, reference tracking objective represented by

$$J_k^{\text{ref}} = \|\phi_{FK}(\mathbf{q}_k) - \mathbf{p}_k^{\text{mocap}}\|_{Q_m}^2 \quad (3)$$

which ensures the resulting trajectory is close to the one obtained in the previous step. Additionally, we included a cost term to encourage kicking foot acceleration during the swing phase of the motion. This term is applied from the start of the swing phase to the end of the impact phase, corresponding to phases (d) and (e) in Fig. 1.

$$J_k^{\text{imp}} = -\zeta_k \|\mathbf{J}_{foot,k}(\mathbf{q}_k) \mathbf{v}_k\|_{Q_i}^2 \quad (4)$$

where ζ_k is a weighting function given by

$$\zeta_k = \begin{cases} 1 & \text{if } k\Delta t \in [T^{\text{swing}}, T^{\text{impact}}], \\ 0 & \text{otherwise.} \end{cases} \quad (5)$$

Here $\mathbf{J}_{foot,k}$ is the Jacobian of the kicking leg's foot link in the world frame. Lastly, a regularization cost is included as

$$J_k^{\text{reg}} = (1 - \zeta_k) \|\mathbf{v}_k\|_{Q_v}^2 + \|\mathbf{h}_k\|_{Q_h}^2 + \sum_i^{n_c} \|\mathbf{f}_{i,k}\|_{Q_f}^2 \quad (6)$$

Where Q_m , Q_v , Q_i , Q_h and Q_f positive definite diagonal weighting matrices. In addition to the dynamic constraints

given by Eqs. (1)-(2), we enforce consistency between the centroidal and kinematic variables using Eqs. (7) to (9).

$$\mathbf{h}_k = \mathbf{A}_{CAM}(\mathbf{q}_k) \mathbf{v}_k \quad (7)$$

$$\mathbf{r}_k = \phi_{CoM}(\mathbf{q}_k) \quad (8)$$

$$\mathbf{c}_{i,k} = \phi_{c_i}(\mathbf{q}_k) \quad (9)$$

where \mathbf{A}_{CAM} represents the centroidal angular momentum matrix [23]. In this context, ϕ_{CoM} and ϕ_{c_i} refer to the forward kinematic maps for the center of mass and the i^{th} contact point, respectively. To ensure kinodynamic feasibility of the resulting trajectories, we incorporate additional contact constraints. Utilizing the contact schedule obtained from previous steps, we impose the following constraints when contact points are active:

$$\mathbf{f}_{i,k} \in \mathcal{F} \quad (10)$$

$$c_{i,k} = c_{i,k-1} \quad (11)$$

$$c_{i,z} = \gamma_{\text{gnd}}(c_{i,x}, c_{i,y}) \quad (12)$$

Equation (10) represents the friction cone constraint, where \mathcal{F} is the linear approximation of the friction cone. Equation (11) enforces a no-slip constraint when a contact point is active in consecutive time steps, whereas Eq. (12) ensures contact with the ground. Here, γ_{gnd} represents a height map function that returns the ground height at a specified (x, y) position. Additionally, we enforce a constraint at every timestep to prevent ground penetration by the feet, given by:

$$c_{i,z} \geq \gamma_{\text{gnd}}(c_{i,x}, c_{i,y}). \quad (13)$$

To address the numerous instances of self-collisions present in the MoCap data, we use convex approximations of selected links that are prone to collision and impose the following distance constraint:

$$\mathcal{C}_{ij}(\mathbf{q}_k) \geq r_i + r_j \quad \forall (i, j) \in P \quad (14)$$

Where \mathcal{C}_{ij} is a capsule distance function representing the minimum distance between two capsules representing the i^{th} and j^{th} links. P represents the set of relevant link pairs, and r_i and r_j are the radii of the capsules for links i and j respectively. Moreover, we impose joint velocity and torque limits in Eqs.(16)-(18). Here we approximate the necessary torque as equivalent to that required for producing ground reaction forces and establish limits accordingly [33], [34], which is given by

$$\boldsymbol{\tau}_{j,k} \approx -\mathbf{S}_j^\top \left(\sum_{i=1}^{n_c} \mathbf{J}(\mathbf{q}_k)_i^\top \mathbf{f}_{i,k} \right), \quad (15)$$

where $\boldsymbol{\tau}_{j,k} \in \mathbb{R}^n$ is the required joint torque at j -th joint, $\mathbf{S}_j \in \mathbb{R}^{n \times n+6}$ is a selection matrix, and $\mathbf{J}_i \in \mathbb{R}^{3 \times n+6}$ is the Jacobian of the i -th contact point.

$$\mathbf{q}_{\min} \leq \mathbf{q}_k \leq \mathbf{q}_{\max} \quad (16)$$

$$\mathbf{v}_{\min} \leq \mathbf{v}_k \leq \mathbf{v}_{\max} \quad (17)$$

$$\boldsymbol{\tau}_{\min} \leq \boldsymbol{\tau}_k \leq \boldsymbol{\tau}_{\max} \quad (18)$$

Furthermore, we incorporate the following explicit biomechanics-inspired constraints: Ankle locking and Approach Angle constraints given by Eqs. (19) and (20) respectively.

$$\mathbf{v}_k^{\text{ankle}} = 0, \quad \text{for } k\Delta t \in [T^{\text{lock}}, T^{\text{impact}}] \quad (19)$$

$$\theta_{\min} < \cos^{-1} \left(\frac{(\mathbf{p}^{\text{ball}}_{xy} - \mathbf{p}^{\text{body}}_{xy}) \cdot \hat{\mathbf{x}}}{\|\mathbf{p}^{\text{ball}}_{xy} - \mathbf{p}^{\text{body}}_{xy}\|} \right) < \theta_{\max} \quad (20)$$

Equation. (19) inspired by [18] ensures that the kicking foot ankle joint velocity $\mathbf{v}_k^{\text{ankle}}$ is zero for before the impact moment, starting from $T_{\text{start}}^{\text{lock}}$ up to the end of the impact period T^{impact} . Equation (20) constrains the approach angle of the robot relative to the ball. It ensures that the angle between the vector from the initial body position $\mathbf{p}^{\text{body}}_{xy}$ to the ball position $\mathbf{p}^{\text{ball}}_{xy}$ (both projected onto the xy-plane) and the x-axis falls within the range $[\theta_{\min}, \theta_{\max}]$. In this work, we set $\theta_{\min} = 24^\circ$ and $\theta_{\max} = 43^\circ$, based on biomechanics literature [14], [16], [17]. Lastly, we add the dynamics integration constraints.

$$\mathbf{q}_{k+1} = \mathbf{q}_k \oplus \mathbf{v}_k \Delta t, \quad (21)$$

$$\mathbf{r}_{k+1} = \mathbf{r}_k + \dot{\mathbf{r}}_k \Delta t \quad (22)$$

$$\dot{\mathbf{r}}_{k+1} = \dot{\mathbf{r}}_k + \ddot{\mathbf{r}}_k \Delta t \quad (23)$$

$$\mathbf{h}_{k+1} = \mathbf{h}_k + \dot{\mathbf{h}}_k \Delta t \quad (24)$$

The final kinodynamic optimization problem is then given by Eqs. (25) to (32).

$$\min_{\substack{\mathbf{r}_k, \dot{\mathbf{r}}_k, \mathbf{h}_k, \mathbf{q}_k, \\ \dot{\mathbf{q}}_k, \mathbf{f}_k, \mathbf{c}_k}} \sum_{k=1}^N J_k^{\text{ref}} + J_k^{\text{reg}} + J_k^{\text{imp}} \quad (25)$$

s.t.

Integration constraints (Eqs. (22) - (24)) (26)

Centroidal dyn. constraints (Eqs. (1)-(2)) (27)

Consistency constraints (Eqs. (7)-(9)) (28)

Contact constraints (Eqs. (10)-(13)) (29)

Biomechanics constraints (Eqs. (19)-(20)) (30)

Collision constraint (Eqs. (14)) (31)

Joint Limit constraints (Eqs. (16)-(18)) (32)

This optimization problem involves considerable non-linear constraints, although not as much as fullbody dynamics based trajectory optimization, so can be difficult to solve efficiently. We address this challenge using a two-stage approach: First, we perform a kinematics-only optimization stage where we ignore the centroidal dynamics, and consistency constraints and obtain a purely kinematic trajectory. This stage is equivalent to the optimization-based inverse kinematics approach described in [24] and illustrated in Fig. 3(b). We then use the result from the first stage as a warm start for the full optimization problem, which includes both kinematics and centroidal dynamics. Finally, we use the joint position \mathbf{q} , velocity \mathbf{v} , and ground reaction force \mathbf{f} from the previous step to compute the required torque and verify that torque limits are not violated when full-body

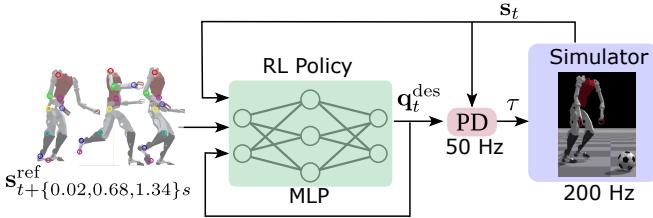


Fig. 4. Imitation learning framework

dynamics is considered. We do this by first obtaining the joint acceleration $\ddot{\mathbf{q}}$ trajectory through finite differencing of the joint velocity \mathbf{v} , and then applying inverse dynamics:

$$\tau^{ID} = \text{ID}(\mathbf{q}, \mathbf{v}, \ddot{\mathbf{q}}, \mathbf{f}) \quad (33)$$

If necessary, we iterate between the kinodynamics optimization and torque verification steps until all constraints are satisfied.

B. Imitation Learning

Once the physically consistent trajectory is obtained, we use the Proximal Policy Optimization (PPO) algorithm [35] to train a Reinforcement Learning (RL) control policy in a simulated environment in Isaac Gym [30] that includes a standard sized soccer ball. This policy has an imitation objective to robustly track the reference trajectory, while making necessary adjustments to account for the mismatch between the centroidal dynamics used in trajectory optimization and the full-body dynamics of the simulation. The RL policy is trained to maximize the following tracking rewards:

$$r_k^{\text{imitation}} = r_k^{\text{keypoint}} + r_k^{\text{joint}} + r_k^{\text{CoM}}, \quad (34)$$

where

$$r_k^{\text{keypoint}} = w_k e^{k_k \sum_{m=1}^M \|\phi_{FK}^m(\mathbf{q}_k^{\text{ref}}) - \phi_{FK}^m(\mathbf{q}_k)\|_2}, \quad (35)$$

$$r_k^{\text{joint}} = w_q e^{k_q \|\mathbf{q}_k^{\text{ref}} - \mathbf{q}_k\|_2} + w_v e^{k_v \|\mathbf{v}_k^{\text{ref}} - \mathbf{v}_k\|_2}, \quad (36)$$

$$r_k^{\text{CoM}} = w_c e^{k_c \|\mathbf{p}_{comk}^{\text{ref}} - \mathbf{p}_{comk}\|_2}. \quad (37)$$

Here, w_k , w_q , w_v , and w_c are weight parameters for each reward component k_k , k_q , k_v , and k_c are scaling factors. $\phi_{FK}^m(\cdot)$ is the forward kinematics map for the m -th keypoint. In addition to the imitation objective, we incorporate the reward to maximize the ball's velocity toward the desired direction:

$$r_k^{\text{ball}} = w_{\text{ball}} * (e^{\max(0, \mathbf{v}^{\text{ball}} \cdot \mathbf{n}^{\text{target}})} - 1) \quad (38)$$

where \mathbf{v}^{ball} is the velocity vector of the ball and $\mathbf{n}^{\text{target}}$ is the unit vector representing the desired direction of ball travel.

C. Sample Efficient RL Training Through Reference Guided Early Termination Strategy

Previous works in the graphics community have demonstrated the importance of early termination for sample-efficient training of motion imitation [36], [37]. Early termination prevents unnecessary local optima behaviors by assigning zero reward after the agent enters undesirable

states, thereby avoiding unnecessary exploration and improving sample efficiency. Luo et al. [37] used a Reference-based Early Termination (RET) strategy, where the agent is terminated if the average link position of the robot deviates from the reference motion. However, this kind of termination strategy based on motion references without consideration of the robot's dynamics faces challenges in yielding the desired behavior, especially for highly dynamic motions like a soccer kick. This is because faithfully following the kinematics reference alone does not guarantee that physical limits, such as torque limits, are not violated. Therefore, the agent might need to deviate from the reference significantly to accommodate the limits not considered in the motion planning stage.

This challenge is a primary reason why we spent significant efforts to create a kinodynamic trajectory optimization framework. The obtained optimal trajectories are guaranteed the system to operate within the physical limits; thereby the trajectories can serve as a trustful prior for constraining the search space during RL training through early termination. In particular, we use the following distance threshold-based RET strategy to achieve sample-efficient exploration:

$$\text{Terminate if } \max_{i \in \mathcal{L}} \|\phi_{FK}^i(\mathbf{q}) - \phi_{FK}^i(\mathbf{q}^{\text{ref}})\|_2 > \bar{d} \quad (39)$$

That is, terminate the episode if the position of any link of the robot in the link set \mathcal{L} deviates from the reference by a distance greater than the threshold \bar{d} .

We consider two kinds of thresholds for our termination strategy. The first is a fixed threshold, where a single distance threshold is used throughout the training process. The second one a time-varying threshold, where the threshold is updated as the training progresses, given by Eq. (40). Specifically, we continuously reduce the distance threshold over time following a staircase-like function, as depicted in Fig. 5, as the policy becomes better at tracking the reference. This approach allows for finer and more precise tracking, enabling the policy to maintain biomechanical characteristics captured in the human MoCap data while allowing for necessary adaptations to the robot's specific dynamics.

$$\bar{d}(t) = \begin{cases} d_{\max} & \text{if } t < t_{\text{start}} \\ d_{\min} & \text{if } t \geq t_{\text{end}} \\ d_{\max} - \sum_{i=1}^n \Delta h \cdot \sigma\left(\gamma \cdot \frac{t - (t_{\text{start}} + i\Delta t)}{\Delta t}\right) & \text{otherwise} \end{cases} \quad (40)$$

where σ is the sigmoid function given by $\sigma(x) = \frac{1}{1+e^{-x}}$, γ is a smoothness parameter, $\Delta d = \frac{d_{\max} - d_{\min}}{n}$ and $\Delta t = \frac{t_{\text{end}} - t_{\text{start}}}{n}$.

III. EXPERIMENTAL SETUP

A. Experimental Platform

In this study, we use PresToe, illustrated in Fig. 2, for all our simulation experiments. PresToe is a 25 DoF humanoid robot, comprised of two 7 DoF legs, each with a distinct 1 DoF toe actuation, two 5 DoF arms, and a 1 DoF torso. The robot stands at a height of 1.3 meters in its nominal configuration and weighs 29.5 kg. The joint torque and

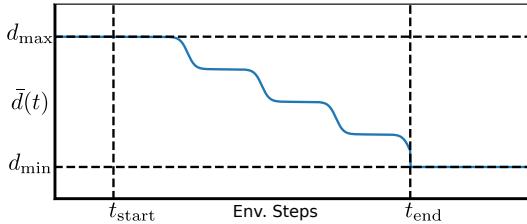


Fig. 5. Stair Case Distance Threshold Function

TABLE I
PRESTOE JOINT TORQUE AND VELOCITY LIMITS

Group	Joint	Max Torque	Max Speed
Leg	Hip	48 N·m	20 rad/s
	Knee	200 N·m	10 rad/s
	Ankle	100 N·m	10 rad/s
	Toe	10 N·m	20 rad/s
Arm	Shoulder		
	Elbow	18 N·m	40 rad/s
Torso	Torso	18 N·m	40 rad/s

velocity limits are listed in Table I. PresToe is an ideal candidate for replicating human-like movements, given its comparable size and anthropomorphic design. The leg design is based on our earlier work on Staccatoe [34] and has a unique co-actuation coupling mechanism that allows the robot to generate large knee and ankle torques of up to 200 N·m for the knee and 100 N·m for the ankle, depending on the robot's joint configuration.

B. Reference Motion Generation and RL Policy Training

We generate reference motion by kinodynamically retargeting MoCap data of a human performing a soccer kick, obtained from the CMU Mocap Dataset [38]. The original motion is 3.7 seconds long, recorded at 120 Hz. We subsample it to 30 Hz, resulting in a trajectory optimization horizon of 112 time steps. We follow the four steps described in Section II-A.2 to perform the retargeting.

The overall Imitation Learning framework is depicted in Fig. 4. As the figure shows, the observation space of the policy is given by

$$\mathbf{o}_t = \left[\mathbf{s}_{t+\{0.02, 0.68, 1.34\}s}^{\text{ref}}, \mathbf{s}_t, \mathbf{q}_{t-0.02s}^{\text{des}} \right], \quad (41)$$

where $\mathbf{s}_{t+\{0.02, 0.68, 1.34\}s}^{\text{ref}}$ represents the future desired reference states at 0.02s, 0.68s, and 1.34s ahead of the current time, \mathbf{s}_t is the current state represented by \mathbf{q}, \mathbf{v} , and $\mathbf{q}_{t-0.02}^{\text{des}}$ is the action at the previous time step which is a desired joint position command. We convert this position command to a torque command using the following PD control law:

$$\boldsymbol{\tau} = \mathbf{k}_p(\mathbf{q}^{\text{des}} - \mathbf{q}) - \mathbf{k}_d\mathbf{v}, \quad (42)$$

where \mathbf{k}_p and \mathbf{k}_d are PD gains. In all experiments, the episode begins with the robot initialized with a random state sampled from the reference motion similar to [36]. For both the policy and value networks, we use a multi-layer perceptron neural network consisting of two hidden

layers with sizes 1024 and 512, respectively, alongside ReLU activation.

C. Sample Efficiency Experiment

We assess the effectiveness of introducing dynamics in the offline motion planning stage by comparing motion tracking policies trained with solely imitation objectives using three different reference types. The first is a kinodynamic reference trajectory that takes into consideration the centroidal dynamics of the robot and imposes appropriate torque limits. The second is a purely kinematics-based motion reference similar to [24], which retains artifacts such as foot skating present in the reference MoCap data. The third is another purely kinematic reference but that does not involve artifacts like foot sliding. This reference is obtained by leveraging the same contact schedule used to obtain the kinodynamic trajectory.

D. Ball Kicking Experiment

In this experiment, we first train a motion tracking policy solely with the motion imitation objective using a time-varying distance threshold approach in Eq. (40). The threshold changes from 30 cm to 10 cm over 1 billion training steps. The goal here is to track the reference as closely as possible and retain biomechanical insights captured in the human MoCap data. We then adapt this policy for a high-impact soccer ball kicking task. To achieve this, we introduce a standard FIFA size 5 ball (weighing 430g with a diameter of 22 cm) into the simulation. We further train the policy to maximize ball velocity in the forward direction using the impact reward in Eq. (38), in addition to the imitation objective. Note that in this work, we assume a fixed ball position and we do not provide it as an observation to the policy. Instead, it is inferred indirectly through the ball impact reward.

IV. RESULTS AND DISCUSSION

Fig. 7 illustrates the imitation reward return as training progresses under two scenarios. Fig. 7 (a) shows the performance without reference-based early termination (RET), while Fig 7 (b) demonstrates the results with RET using a distance threshold of 30 cm. Both Figs 7 (a) and (b) display the training performance when three different reference types are used: kinematic-only with motion artifact like foot sliding, kinematic-only but without motion artifacts, and kinodynamic references. As evident from the figure, the policy trained with kinodynamically consistent trajectories learns faster and achieves higher returns in both scenarios. When RET is deployed, there is a noticeable advantage, especially when motion artifacts like foot slippage are removed. The policy tends to learn faster in this case, although not as effectively as with references that consider dynamics.

Table II presents the reference tracking performance of the trained policies, measured using the average link position error between the policy and the reference at each time step. The table clearly demonstrates the importance of RET. In fact, as can be observed in the attached video, without RET,

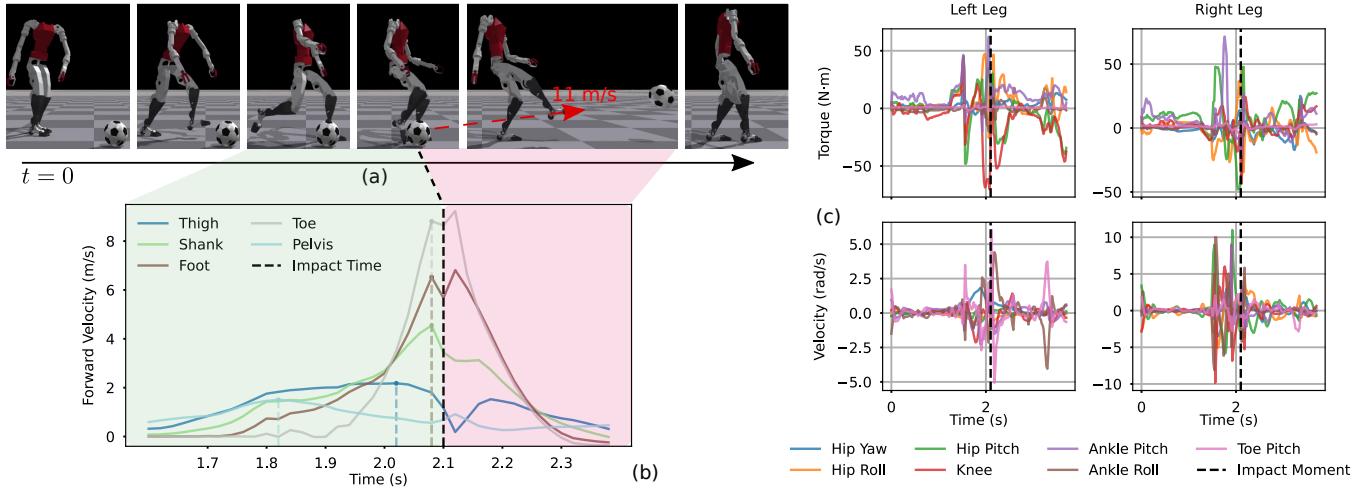


Fig. 6. Human-like instep soccer kicking with Prestoe: (a) Motion timelapse from left to right: the robot executes the classic instep kick phases, i.e., momentum-building step, windup, kicking leg acceleration, impact, and follow-through. (b) Forward velocity of the kicking leg links moments before impact and during follow-through phase; vertical dashed lines indicate the moment each link reaches its maximum speed. (c) Torque and velocity profiles throughout the motion for both kicking and non-kicking legs.

the policies fail to learn to emulate the motion with just reward signals alone in one billion steps for all three reference types. However, with when RET is used, all three policies successfully learned to imitate the motion, with the policy trained on kinodynamically consistent trajectories learning significantly faster and more closely resembling the reference. The kinodynamic reference's superior performance can be attributed to its consideration of both kinematics and dynamics of the robot as well as its physical limits, providing a more realistic and achievable trajectory for the policy to imitate.

Fig. 6(a) presents a snapshot sequence of the robot performing an instep kick using the kicking policy. The kick follows the six steps of an instep kick: approach, support phase, windup, leg swing, ball contact, and follow-through phases, achieving ball speeds exceeding 11 m/s. This speed is approximately double the 6.6 m/s instep kick achieved by the ground-anchored single-leg robot detailed in [4], and about 40% of the speed an average human player can achieve [31]. This performance is noteworthy considering the robot's size and limited torque capability. The robot has a maximum hip pitch torque of 48N.m, compared to the average human player's hip joint flexion/extension capability of over 250N.m [32]. Fig. 6(b) showcases the velocities of the kicking leg's links in the forward direction before and after ball impact. We observe that the pelvis reaches its maximum speed and starts to decelerate before the thigh reaches its maximum speed and decelerates. Moreover, the shank, foot, and toe links reach their maximum speeds at the moment of impact. This pattern exhibits the classic proximal-to-distal kinetic energy transfer present in optimal soccer kicks, as discussed in biomechanics literature [14], [16], [19], [22]. Finally, Fig. 6(c) shows the torque and velocity profiles of the kicking (Right leg) and non-kicking leg joints throughout the entire motion and demonstrates that they remain within the physical limits of the robot.

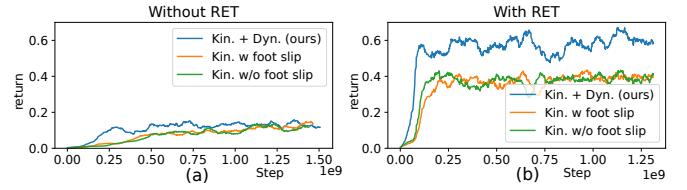


Fig. 7. Training progress measured by normalized imitation reward return vs. environment steps for three reference types: Kinematic-only with foot slippage artifacts (Kin w/ foot slip), Kinematic-only without foot slippage (Kin w/o foot slip), and Kinodynamic reference. (a) Without reference-based early termination (RET). (b) With reference-based early termination.

TABLE II
AVERAGE LINK POSITION TRACKING ERROR

Ref. Type	Without RET	With RET
Kin. w/ foot slip	0.66 +/- 0.47 m	0.14 +/- 0.03 m
Kin. w/o foot slip	0.61 +/- 0.51 m	0.114 +/- 0.06 m
Kin. + Dyn. (ours)	0.60 +/- 0.50 m	0.09 +/- 0.02 m

V. CONCLUSION AND FUTURE WORK

In this study, we developed a motion planning and imitation learning-based control framework informed by human biomechanics to facilitate dynamic and powerful soccer kicks with a humanoid robot. This approach significantly surpasses the limited kick-power outputs observed in previous works which used quasi-static and walk-kick type kicking techniques. Additionally, we demonstrated the effectiveness of kinodynamic motion retargeting, which produces physically consistent trajectories for efficiently learning control policies, as opposed to kinematics-only approaches that don't take into account the dynamics of the robot. In future work, we aim to extend and validate our approach with other dynamic motions beyond soccer kicks. Additionally, as all experiments in this study were conducted in a simulated environment, we plan to validate these trajectories on a physical robot.

ACKNOWLEDGMENT

This material is based upon work supported by the National Science Foundation under Grant No. 2220924.

REFERENCES

- [1] H. Kitano and M. Asada, "The RoboCup humanoid challenge as the millennium challenge for advanced robotics," *Advanced Robotics*, vol. 13, no. 8, pp. 723–736, Jan. 1998. [Online]. Available: <https://doi.org/10.1163/156855300x00061>
- [2] R. Gerndt, D. Seifert, J. H. Baltes, S. Sadeghnejad, and S. Behnke, "Humanoid robots in soccer: Robots versus humans in RoboCup 2050," *IEEE Robotics Automation Magazine*, vol. 22, no. 3, pp. 147–154, Sep. 2015. [Online]. Available: <https://doi.org/10.1109/mra.2015.2448811>
- [3] H. Schempf, C. Kraeuter, and M. Blackwell, "Roboleg: a robotic soccer-ball kicking leg," in *Proceedings of 1995 IEEE International Conference on Robotics and Automation*, vol. 2, 1995, pp. 1314–1318 vol.2.
- [4] S. Haddadin, T. Laue, U. Frese, S. Wolf, A. Albu-Schäffer, and G. Hirzinger, "Kick it with elasticity: Safety and performance in human–robot soccer," *Robotics and Autonomous Systems*, vol. 57, no. 8, pp. 761–775, Jul. 2009. [Online]. Available: <https://doi.org/10.1016/j.robot.2009.03.004>
- [5] M. Vahidi and S. A. A. Moosavian, "Dynamics of a 9-DoF robotic leg for a football simulator," in *2015 3rd RSI International Conference on Robotics and Mechatronics (ICROM)*. IEEE, Oct. 2015. [Online]. Available: <https://doi.org/10.1109/icrom.2015.7367803>
- [6] J. Müller, T. Laue, and T. Röfer, "Kicking a ball – modeling complex dynamic motions for humanoid robots," 01 2010, pp. 109–120.
- [7] S. Barrett, K. Genter, T. Hester, M. Quinlan, and P. Stone, "Controlled kicking under uncertainty," in *The Fifth Workshop on Humanoid Soccer Robots at Humanoids*, 2010.
- [8] T. Hester, M. Quinlan, and P. Stone, "Generalized model learning for reinforcement learning on a humanoid robot," in *2010 IEEE International Conference on Robotics and Automation*. IEEE, May 2010. [Online]. Available: <https://doi.org/10.1109/robot.2010.5509181>
- [9] L. Leottau, C. Celegin, and J. Ruiz-del Solar, "Ball dribbling for humanoid biped robots: A reinforcement learning and fuzzy control approach," in *RoboCup 2014: Robot World Cup XVIII*, R. A. C. Bianchi, H. L. Akin, S. Ramamoorthy, and K. Sugiura, Eds. Cham: Springer International Publishing, 2015, pp. 549–561.
- [10] D. Closson, Ed., *Tectonics*. InTech, Feb. 2011. [Online]. Available: <https://doi.org/10.5772/567>
- [11] H. Teixeira, T. Silva, M. Abreu, and L. P. Reis, "Humanoid robot kick in motion ability for playing robotic soccer," in *2020 IEEE International Conference on Autonomous Robot Systems and Competitions (ICARSC)*. IEEE, Apr. 2020. [Online]. Available: <https://doi.org/10.1109/icarsc49921.2020.9096073>
- [12] A. Rezaeipanah, P. Amiri, and S. Jafari, "Performing the kick during walking for RoboCup 3d soccer simulation league using reinforcement learning algorithm," *International Journal of Social Robotics*, vol. 13, no. 6, pp. 1235–1252, Nov. 2020. [Online]. Available: <https://doi.org/10.1007/s12369-020-00712-2>
- [13] A. Express, "Honda's asimo: the penalty-taking, bar-tending robot," 2014, YouTube video. [Online]. Available: <https://www.youtube.com/watch?v=QdQL11uWWCI>
- [14] G. Shan and X. Zhang, "From 2d leg kinematics to 3d full-body biomechanics—the past, present and future of scientific analysis of maximal instep kick in soccer," *Sports Medicine, Arthroscopy, Rehabilitation, Therapy Technology*, vol. 3, no. 1, Oct. 2011. [Online]. Available: <https://doi.org/10.1186/1758-2555-3-23>
- [15] BiomechanicsMMU, "Football free kick - slow motion video," <https://www.youtube.com/watch?v=lbMA2wWuqh8>, 2009, accessed: Wed Oct 25 22:32:03 EDT 2023.
- [16] G. Shan, "?—a state-of-the-art review," *Applied Sciences*, vol. 12, no. 21, p. 10886, Oct. 2022. [Online]. Available: <https://doi.org/10.3390/app122110886>
- [17] C. D. Egan, M. H. G. Verheul, and G. J. P. Savelsbergh, "Effects of experience on the coordination of internally and externally timed soccer kicks," *Journal of Motor Behavior*, vol. 39, no. 5, pp. 423–432, Sep. 2007. [Online]. Available: <https://doi.org/10.3200/jmbr.39.5.423-432>
- [18] W. R. Barfield, "The biomechanics of kicking in soccer," *Clinics in sports medicine*, vol. 17, no. 4, pp. 711–728, 1998.
- [19] A. Lees and L. Nolan, "The biomechanics of soccer: A review," *Journal of Sports Sciences*, vol. 16, no. 3, pp. 211–234, Jan. 1998. [Online]. Available: <https://doi.org/10.1080/026404198366740>
- [20] E. Kellis and A. Katis, "Biomechanical characteristics and determinants of instep soccer kick," *J Sports Sci Med*, vol. 6, no. 2, pp. 154–165, 2007.
- [21] R. H. Brophy, S. I. Backus, B. S. Pansy, S. Lyman, and R. J. Williams, "Lower extremity muscle activation and alignment during the soccer instep and side-foot kicks," *Journal of Orthopaedic Sports Physical Therapy*, vol. 37, no. 5, pp. 260–268, May 2007. [Online]. Available: <https://doi.org/10.2519/jospt.2007.2255>
- [22] A. Lees, T. Asai, T. B. Andersen, H. Nunome, and T. Sterzing, "The biomechanics of kicking in soccer: A review," *Journal of Sports Sciences*, vol. 28, no. 8, pp. 805–817, Jun. 2010. [Online]. Available: <https://doi.org/10.1080/02640414.2010.481305>
- [23] H. Dai, A. Valenzuela, and R. Tedrake, "Whole-body motion planning with centroidal dynamics and full kinematics," in *2014 IEEE-RAS International Conference on Humanoid Robots*, 2014, pp. 295–302.
- [24] A. Tang, T. Hiraoka, N. Hiraoka, F. Shi, K. Kawaharazuka, K. Kojima, K. Okada, and M. Inaba, "Humanmimic: Learning natural locomotion and transitions for humanoid robot via wasserstein adversarial imitation," *arXiv preprint arXiv:2309.14225*, 2023.
- [25] Y. Yan, E. V. Mascaro, and D. Lee, "Imitationnet: Unsupervised human-to-robot motion retargeting via shared latent space," in *2023 IEEE-RAS 22nd International Conference on Humanoid Robots (Humanoids)*. IEEE, 2023, pp. 1–8.
- [26] X. Cheng, Y. Ji, J. Chen, R. Yang, G. Yang, and X. Wang, "Expressive whole-body control for humanoid robots," *arXiv preprint arXiv:2402.16796*, 2024.
- [27] T. He, Z. Luo, W. Xiao, C. Zhang, K. Kitani, C. Liu, and G. Shi, "Learning human-to-humanoid real-time whole-body teleoperation," *arXiv preprint arXiv:2403.04436*, 2024.
- [28] S. Nakaoka, A. Nakazawa, F. Kanehiro, K. Kaneko, M. Morisawa, and K. Ikeuchi, "Task model of lower body motion for a biped humanoid robot to imitate human dances," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2005, pp. 3157–3162.
- [29] D. Rempe, L. J. Guibas, A. Hertzmann, B. Russell, R. Villegas, and J. Yang, "Contact and human dynamics from monocular video," in *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16*. Springer, 2020, pp. 71–87.
- [30] V. Makoviychuk, L. Wawrzyniak, Y. Guo, M. Lu, K. Storey, M. Macklin, D. Hoeller, N. Rudin, A. Allshire, A. Handa *et al.*, "Isaac gym: High performance gpu-based physics simulation for robot learning," *arXiv preprint arXiv:2108.10470*, 2021.
- [31] A. Radja, L. P. Kuvavcic, Goran and De Giorgio, Andrea and Sellami, Maha and Ardigo, N. L. Bragazzi, and J. Padulo, "The ball kicking speed: A new, efficient performance indicator in youth soccer," *Plos one*, vol. 14, no. 5, p. e0217101, 2019.
- [32] K. Sakamoto, N. Numazu, S. Hong, and T. Asai, "Kinetic analysis of instep and side-foot kick in female and male soccer players," *Procedia engineering*, vol. 147, pp. 214–219, 2016.
- [33] M. Chignoli, D. Kim, E. Stanger-Jones, and S. Kim, "The MIT humanoid robot: Design, motion planning, and control for acrobatic behaviors," *CorR*, vol. abs/2104.09025, 2021. [Online]. Available: <https://arxiv.org/abs/2104.09025>
- [34] N. Perera, S. Yu, D. Marew, M. Tang, K. Suzuki, A. McCormack, S. Zhu, Y.-J. Kim, and D. Kim, "Staccatoe: A single-leg robot that mimics the human leg and toe," *arXiv preprint arXiv:2404.05039*, 2024.
- [35] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, "Proximal policy optimization algorithms," *arXiv preprint arXiv:1707.06347*, 2017.
- [36] X. B. Peng, P. Abbeel, S. Levine, and M. Van de Panne, "Deepmimic: Example-guided deep reinforcement learning of physics-based character skills," *ACM Transactions On Graphics (TOG)*, vol. 37, no. 4, pp. 1–14, 2018.
- [37] Z. Luo, J. Cao, K. Kitani, W. Xu *et al.*, "Perpetual humanoid control for real-time simulated avatars," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 10895–10904.
- [38] F. De la Torre, J. Hodgins, A. Bargteil, X. Martin, J. Macey, A. Collado, and P. Beltran, "Guide to the carnegie mellon university multimodal activity (cmu-mmact) database," 2009.