

HUSKY: Humanoid Skateboarding System via Physics-Aware Whole-Body Control

Jinrui Han^{1,2*} Dewei Wang^{1,3*} Chenyun Zhang¹ Xinzhe Liu^{1,4}

Ping Luo² Chenjia Bai^{1†} Xuelong Li^{1†}

¹Institute of Artificial Intelligence (TeleAI), China Telecom ²The University of Hong Kong

³University of Science and Technology of China ⁴ShanghaiTech University

*Equal contribution [†]Corresponding author

Webpage: husky-humanoid.github.io Code: [TeleHuman/humanoid_skateboarding](https://github.com/TeleHuman/humanoid_skateboarding)




Fig. 1: Overview. (a) Our proposed framework HUSKY enables the humanoid robot to perform complete real-world skateboarding, including pushing, steering, and phase transitions. (b) Generalization to diverse outdoor scenarios and skateboards with consistent stability and control. (c) Reliable indoor skateboarding performance. (d) Lean-to-steer behaviors achieved by exploiting robot body tilt. (e) Robustness against external disturbances.

Abstract—While current humanoid whole-body control frameworks predominantly rely on the static environment assumptions, addressing tasks characterized by high dynamism and complex interactions presents a formidable challenge. In this paper, we address humanoid skateboarding, a highly challenging task requiring stable dynamic maneuvering on an underactuated wheeled platform. This integrated system is governed by non-holonomic constraints and tightly coupled human-object interactions. Successfully executing this task requires simultaneous mastery of hybrid contact dynamics and robust balance control on a mechanically coupled, dynamically unstable skateboard. To overcome the aforementioned challenges, we propose HUSKY, a learning-based framework that integrates humanoid-skateboard system modeling and physics-aware whole-body control. We first model the coupling relationship between board tilt and truck steering angles, enabling a principled analysis of system

dynamics. Building upon this, HUSKY leverages Adversarial Motion Priors (AMP) to learn human-like pushing motions and employs a physics-guided, heading-oriented strategy for lean-to-steer behaviors. Moreover, a trajectory-guided mechanism ensures smooth and stable transitions between pushing and steering. Experimental results on the Unitree G1 humanoid platform demonstrate that our framework enables stable and agile maneuvering on skateboards in real-world scenarios. More videos and code are available on our [project page](#).

I. INTRODUCTION

Recent advancements in robotic hardware and control algorithms have enabled humanoid robots to perform a wide range of whole-body control tasks, including locomotion over complex terrains [28, 37, 20], dancing [12, 18, 41, 10],

manipulation [4, 30, 31], and fall recovery [14, 15]. However, learning skateboarding skills for humanoid robots remains a highly challenging and largely unexplored problem. While existing approaches primarily focus on contact-rich whole-body control tasks [48, 43] and humanoid–object interaction scenarios [38, 39], humanoid skateboarding involves even more intricate contact dynamics and hybrid motion transitions. Furthermore, the shift from interacting with static objects to an underactuated wheeled platform challenges existing control frameworks, rendering them inadequate for stable humanoid skateboarding system.

Traditional model-based approaches typically leverage simplified dynamic models for trajectory generation and Model Predictive Control (MPC) to synthesize skateboarding maneuvers for legged robots [32, 33, 42]. While these classical methodologies can generate controlled motions, the high computational cost of solving high-dimensional, non-convex optimization problems often precludes the real-time responsiveness required for dynamic skateboarding. Moreover, simplified models are insufficient to capture the complex non-holonomic and underactuated dynamics inherent to skateboarding, resulting in limited robustness to unmodeled physical effects and environmental variations. Although agile skateboarding skills can be achieved in simulation [22], they overlook physical hardware constraints and the sim-to-real gap, making direct real-world deployment highly challenging.

Deep Reinforcement Learning (DRL) has recently emerged as a powerful alternative, leveraging large-scale parallel training to synthesize complex behaviors with minimal modeling assumptions and robust performance. Recent researches have demonstrated the effectiveness of RL in diverse tasks, from force-controlled interaction [49] and end-effector control [17] to quadrupedal skateboarding [19, 3] and humanoid roller skating [6], highlighting RL’s capacity to manage complex contact dynamics and stabilize high-speed, underactuated systems.

However, humanoid skateboarding introduces substantially greater challenges due to the emergence of a tightly coupled, highly underactuated system. While the humanoid’s high-dimensional state and action spaces already complicate whole-body control, this difficulty is compounded by the dynamic nature of the skateboard. The board’s motion, governed by intricate interactions between the wheels, ground, and the robot’s feet, creates a scenario where the robot must indirectly control its own moving support base. The inherent instability of this task becomes even more evident when compared to quadrupedal platforms. Unlike quadrupeds [19], humanoids possess significantly higher Degrees of Freedom (DoF) and an elevated Center of Mass (CoM), paired with a much smaller and more precarious support polygon. These physical constraints, combined with the need for complex leg reorientation and side-on steering maneuvers, result in highly nonlinear dynamics and severe sim-to-real discrepancies. Consequently, existing learning-based controllers that may succeed on simpler platforms struggle to maintain robust coordination here. Table I further highlights these unique structural and task-specific difficulties.

TABLE I: Comparison of Humanoid and Quadrupedal Skateboarding Tasks.

Term	Humanoid (G1)	Quadruped (Go1) [19]
Degrees of Freedom	23	12
Contact Points	1–2	2–4
CoM Height	0.6–0.8 m	0.2–0.3 m
Support Polygon	0.05–0.10 m^2	0.10–0.15 m^2
Side-on Steering	✓	×
Leg Reorientation	✓	×

In this paper, we propose HUSKY, a physics-aware whole-body controller for **humanoid skateboarding system**. HUSKY first explicitly analyzes the coupled humanoid–skateboard system and derives an equality constraint between board tilt and truck steering angles, enabling physics-informed policy learning. The skateboarding task is formulated as a hybrid dynamical system with discrete contact phases, capturing the distinct pushing and steering behaviors inherent to skateboard locomotion. Based on this modeling, we employ DRL to train the humanoid to acquire the essential skateboarding skills, including propulsion through pushing and directional control through steering. To enhance the robustness and naturalness of pushing behaviors, HUSKY leverages Adversarial Motion Priors (AMP) [27] to guide the policy toward human-like propulsion motions. Steering is achieved through a physics-aware strategy that exploits the intrinsic coupling between humanoid body lean and skateboard truck steering. To facilitate exploration across distinct phases and ensure smooth transitions, HUSKY integrates a trajectory planning mechanism that stabilizes transitions between pushing and steering, resulting in continuous and coordinated skateboarding behaviors.

Taken together, HUSKY enables the humanoid to achieve agile, stable skateboarding maneuvers, as validated in both simulation and real-world experiments. Our main contributions are summarized as follows:

- We model the humanoid skateboarding system by explicitly incorporating the tilt–steering constraint and formulate the skateboarding task as a hybrid dynamical system characterized by distinct contact phases.
- We develop a learning-based whole-body control framework that integrates AMP-based pushing, physics-guided steering, and trajectory planning for phase transitions.
- We validate HUSKY in both simulation and real-world experiments, demonstrating agile, stable, and human-like humanoid skateboarding maneuvers.

II. RELATED WORK

A. Learning-based Humanoid Whole-Body Control

Recent advances in learning-based humanoid whole-body control have enabled robots to tackle complex control tasks, leveraging state-of-the-art algorithms and high-fidelity simulators [45, 23, 21]. Existing research primarily focuses on locomotion [28, 7], motion tracking [26, 12], and teleoperation [11, 46, 47]. These approaches have achieved robust blind locomotion over uneven terrains [8, 40], highly dynamic motion tracking [41, 10], agile parkour [50], and complex

manipulation tasks [5]. To enhance environmental perception and reactive capabilities, many studies have incorporated exteroceptive sensors such as LiDAR and depth cameras, enabling more sophisticated terrain traversal [37, 20, 1] and object interaction [38, 39, 49]. Our work further investigates the task of humanoid skateboarding, which demands significantly higher levels of control agility and whole-body coordination.

B. Legged Robots on Dynamic Platforms

The interaction between legged robots and dynamic platforms poses significant challenges. Model-based approaches have enabled quadrupedal robots to perform skateboarding through offline trajectory optimization [42]. Learning-based methods further allow quadrupeds to acquire more robust skateboarding skills, including mounting [3] and maneuvering the board [19]. By simulating dynamic platform motion in virtual environments, Huang et al. [13] demonstrated effective balance control across a variety of unstable surfaces.

Humanoid skateboarding, however, remains especially demanding for model-based frameworks [33, 32, 22], as it requires agile, real-time control and highly accurate physical modeling [36, 16]. In contrast, while RL has shown great potential for enabling robust and dexterous humanoid behaviors in other domains [2, 9], its application to skateboarding remains largely unexplored. Our framework leverages these capabilities to realize learning-based agile skateboarding skills on a humanoid robot in the real world.

III. HUMANOID-SKATEBOARD SYSTEM

A. System Modeling

We consider a humanoid robot interacting with a fully passive articulated skateboard via foot contacts, with all actuation provided by the humanoid. As illustrated in Fig. 2, the skateboard comprises three main components: the deck, trucks, and wheels. The deck supports the humanoid's weight and transmits forces to the trucks, which connect the deck to the wheels. In real skateboards, the trucks use a kingpin-based suspension with elastic bushings, which converts deck tilting into truck steering while providing force damping. While directly modeling the full suspension in simulation is challenging due to its structural complexity and compliance. Therefore, we adopt a simplified kinematic model that preserves the key lean-to-steer behavior while remaining tractable for control and simulation.

Under this formulation, wheel-ground contacts are modeled as non-holonomic rolling constraints that restrict lateral slip. Instead of explicitly resolving the constraint forces, we exploit the truck geometry to relate the board tilt angle γ to the rotation of the truck axes. Due to the mechanical structure of the trucks, tilting the board by γ induces a coordinated rotation of the truck axes, resulting in a kinematic coupling [36] described by:

$$\tan \sigma = \tan \lambda \sin \gamma, \quad (1)$$

where λ is the constant rake angle of the skateboard, and σ is the resulting truck steering angle. Intuitively, this implies




Fig. 2: Skateboard Model. We analyze the skateboard kinematic structure and derive the coupling relationships among the board tilt, truck steering, and rake angles, which form the basis of the lean-to-steer behavior.

that the truck steering angle is determined by the board tilt angle, with larger tilts producing greater steering deflections. Detailed derivation is provided in Appendix A.

This abstraction directly links board tilt to truck steering, capturing the essential equality constraint of humanoid-driven maneuvers while avoiding unnecessary mechanical redundancies. The resulting model is both computationally efficient and physically representative, providing a tractable foundation for subsequent control and learning of humanoid skateboarding.

B. Humanoid Skateboarding Task

We formulate humanoid skateboarding as a hybrid control task, comprising two periodically recurring phases, *pushing* and *steering*, which arise from discrete changes in the system's contact topology. Each phase is governed by distinct dynamics and control objectives. The framework is shown in Fig. 3.

During the pushing phase (Section IV-B), one foot of the humanoid maintains contact with the skateboard to stabilize the body, while the other intermittently contacts the ground to generate propulsion. Forward motion is produced by tangential ground reaction forces at the pushing foot, whereas the supporting foot ensures balance on the moving board.

In the steering phase (Section IV-C), both feet of the humanoid remain on the skateboard, producing a passive gliding motion. Steering is achieved by modulating the board tilt through body leaning, which induces truck rotation and allows the humanoid to follow a desired heading.

While transitions between pushing and steering are challenging due to the distinct dynamics and control objectives of each phase. In Section IV-D, we describe strategies to explicitly bridge these phases, enabling smooth and stable phase transitions during humanoid skateboarding.




Fig. 3: Framework of HUSKY. (a) We first analyze and model the humanoid-skateboard system, deriving a physics-inspired lean-to-steer coupling mechanism. Due to the distinct contact dynamics and control objectives across skateboarding phases, we adopt a phase-wise learning strategy. (b) The learning-based whole-body control framework integrates an AMP-based pushing style for active forward propulsion, a steering strategy guided by kinematics-aware tilt references, and a trajectory-guided transition mechanism to enable stable switching between pushing and steering phases.

IV. METHOD

A. Problem Formulation

In this work, we adopt the Unitree G1 humanoid robot [35] with 23 controllable degrees of freedom (DoF), excluding the three DoFs of each wrist. We formulate the humanoid skateboarding task as a RL problem, in which the agent interacts with the environment through a policy π to maximize cumulative reward. At each timestep t , the policy receives the state s_t and outputs an action $a_t \in \mathbb{R}^{23}$, which is subsequently mapped to motor torques via a PD controller. This defines the policy as $\pi(a_t|s_t)$. The environment dynamics $p(s_{t+1}|s_t, a_t)$ determine the next state, while a dense reward $r(s_t, a_t)$ evaluates skateboarding performance while providing regularization. The agent aims to maximize the expected return $J = \mathbb{E} \left[\sum_{t=0}^{T-1} \beta^t r_t \right]$, where β denotes the discount factor.

We utilize an asymmetric actor-critic framework for policy training, where the actor observes only proprioceptive information, while the critic has access to additional privileged information. At each timestep t , the policy state s_t consists of a short history of proprioceptive states $\mathbf{o}_{t-4:t}^{\text{prop}}$, where each proprioceptive state $\mathbf{o}_t^{\text{prop}} \in \mathbb{R}^{78}$ is defined as

$$\mathbf{o}_t^{\text{prop}} = [\mathbf{c}_t, \boldsymbol{\omega}_t, \mathbf{g}_t, \boldsymbol{\theta}_t, \dot{\boldsymbol{\theta}}_t, \mathbf{a}_{t-1}, \Phi], \quad (2)$$

where $\mathbf{c}_t = [v_{cmd}, \psi] \in \mathbb{R}^2$ contains the desired skateboard forward velocity v_{cmd} and heading target ψ , $\boldsymbol{\omega}_t \in \mathbb{R}^3$ is the base angular velocity, $\mathbf{g}_t \in \mathbb{R}^3$ is the base projected gravity, $\boldsymbol{\theta}_t \in \mathbb{R}^{23}$ and $\dot{\boldsymbol{\theta}}_t \in \mathbb{R}^{23}$ are the joint angles and velocities, and $\mathbf{a}_{t-1} \in \mathbb{R}^{23}$ is the previous action. The phase cycle is represented by a normalized variable $\Phi = (t \bmod H)/H \in [0, 1)$, where H is the cycle duration. This periodic variable governs transitions between pushing and steering phases, serving as a temporal reference for the policy.

To expose the skateboard state to the critic during training, we define the privileged observation $\mathbf{o}_t^{\text{priv}}$ as

$$\mathbf{o}_t^{\text{priv}} = [\mathbf{v}_t, \mathbf{p}_t^b, \mathbf{r}_t^b, \mathbf{v}_t^b, \boldsymbol{\omega}_t^b, \boldsymbol{\theta}_t^b, \mathbf{f}_t^g, \mathbf{f}_t^b], \quad (3)$$

where $\mathbf{v}_t \in \mathbb{R}^3$ denotes the robot base linear velocity, $\mathbf{p}_t^b \in \mathbb{R}^3$, $\mathbf{r}_t^b \in \mathbb{R}^6$, $\mathbf{v}_t^b \in \mathbb{R}^3$ and $\boldsymbol{\omega}_t^b \in \mathbb{R}^3$ represent the skateboard base position, orientation, linear velocity and angular velocity in the robot's frame, respectively. $\boldsymbol{\theta}_t^b \in \mathbb{R}^3$ is the skateboard's tilt and truck joint angles, while $\mathbf{f}_t^g \in \mathbb{R}^6$ and $\mathbf{f}_t^b \in \mathbb{R}^6$ represent the humanoid's feet contact forces with ground and skateboard, respectively. To guide the humanoid toward the desired behaviors in each phase, we detail phase-specific reward functions in Appendix C. The overall reward is formulated as:

$$r_t = \mathbb{I}^{\text{push}} \cdot r_t^{\text{push}} + \mathbb{I}^{\text{steer}} \cdot r_t^{\text{steer}} + \mathbb{I}^{\text{trans}} \cdot r_t^{\text{trans}} + r_t^{\text{reg}}, \quad (4)$$

where $\mathbb{I}^{\text{push}}$, $\mathbb{I}^{\text{steer}}$, and $\mathbb{I}^{\text{trans}}$ are binary indicators denoting the current contact phase, and r_t^{reg} is a global regularization term.

B. Adversarial Motion Prior for Pushing Style

During the pushing phase, the primary objective is to achieve a target skateboard velocity v_{cmd} by generating physically plausible propulsion through intermittent foot-ground contacts. Traditional reference-tracking methods enforce strict adherence to predefined joint trajectories [26], resulting in rigid behaviors that limit adaptation to varying environments and velocity demands. To encourage more natural and robust behaviors, we build on the AMP framework, which employs a discriminator D_ϕ to distinguish policy rollouts from reference human pushing motions.

The AMP observation is defined as the robot joint angles $\mathbf{o}_t^{\text{amp}} = \boldsymbol{\theta}_t \in \mathbb{R}^{23}$. To provide the discriminator with temporal context, we construct a motion window $\tau_t = \{\mathbf{o}_{t-3}^{\text{amp}}, \dots, \mathbf{o}_{t+1}^{\text{amp}}\}$ spanning five timesteps. Reference transitions are sampled from a human motion dataset $\mathcal{M}$, while generated transitions are collected from the current policy rollouts $\mathcal{P}$.

The discriminator is trained to distinguish between expert and agent motions using a least-squares loss, augmented with

a gradient penalty to promote training stability:

$$\arg \max_{\phi} \mathbb{E}_{\tau \sim \mathcal{M}}[(D_{\phi}(\tau) - 1)^2] + \mathbb{E}_{\tau \sim \mathcal{P}}[(D_{\phi}(\tau) + 1)^2] + \frac{\alpha^d}{2} \mathbb{E}_{\tau \sim \mathcal{M}}[\|\nabla_{\phi} D_{\phi}(\tau)\|_2]. \quad (5)$$

The discriminator output $d = D_{\phi}(\tau_t)$ is then mapped to a bounded style reward:

$$r_t^{\text{style}}(s_t) = \alpha \cdot \max\left(0, 1 - \frac{1}{4}(d - 1)^2\right), \quad (6)$$

where α is a scaling coefficient. The total reward for the pushing phase combines task-specific and style components:

$$r_t^{\text{push}} = r_t^{\text{task}} + r_t^{\text{style}}. \quad (7)$$

Here, r_t^{task} encourages accurate tracking of the commanded skateboard velocity, while r_t^{style} guides the humanoid to exhibit human-like pushing behaviors.

C. Kinematics-Guided Heading-Oriented Steering

In the steering phase, the humanoid receives a target heading command ψ and leans its body to induce a board tilt γ , which, through the truck geometry coupling in Eq. (1), produces the corresponding truck steering angle σ and achieves the desired turning behavior. To exploit this physical coupling, we provide the humanoid with a kinematics-guided reference specifying the desired board tilt angle for heading regulation.

Assuming a planar kinematic model of the skateboard, we adopt a bicycle-model approximation commonly used in vehicle dynamics [44], which captures the yaw motion of the skateboard while neglecting lateral slip and vertical dynamics. Under this formulation, the yaw rate $\dot{\psi}$ is given by:

$$\dot{\psi} = \frac{v}{L} \tan \sigma, \quad (8)$$

where v is the forward velocity of the skateboard in its base frame, and L is the wheelbase. Combining this with the lean-to-steer relationship in Eq. (1), we obtain

$$\dot{\psi} = \frac{v}{L} \tan \lambda \sin \gamma, \quad (9)$$

which links the board tilt angle γ to the resulting yaw rate.

Given a desired heading change $\Delta\psi = \psi - \psi_{\text{board}}$ over a steering horizon Δt , where ψ_{board} denotes the current heading of the skateboard, we assume a constant yaw rate $\dot{\psi} \approx \Delta\psi/\Delta t$ to achieve a smooth and gradual turn. Substituting this into Eq. (9) yields the kinematics-guided tilt reference:

$$\gamma_{\text{ref}} = \arcsin\left(\frac{L \Delta\psi}{v \Delta t \tan \lambda}\right), \quad (10)$$

where γ_{ref} represents the desired board tilt angle that minimizes the heading error. To ensure numerical stability at low speeds or large desired turns, we clip v to a minimum threshold and constrain γ_{ref} within feasible lean limits.

We define a tilt reward to guide the policy toward this physically consistent lean:

$$r_{\text{tilt}} = \exp\left(-\frac{\|\gamma - \gamma_{\text{ref}}\|^2}{\sigma_{\gamma}^2}\right), \quad (11)$$

where σ_{γ} sets the tolerance. This is combined with a heading-tracking reward $r_{\text{heading}} = \exp(-(\Delta\psi)^2/\sigma_{\psi}^2)$, ensuring that the policy aligns the skateboard with the target direction while maintaining a physically feasible lean angle.

To further promote balance and stable foot placement during steering, two virtual markers are defined above the skateboard trucks. These markers indicate preferred feet contact locations, and the policy is encouraged to place the feet close to them, providing auxiliary guidance for the humanoid's stance.

D. Trajectory Planning for Phase Transition

Direct transitions between the pushing and steering phases are challenging due to differences in body poses, contact dynamics, and learning objectives. Policies trained without explicit transition guidance often fail to explore new contact phases, which can lead to convergence to local optima [48].

To address this problem, we introduce a *trajectory-guided transition* mechanism that explicitly bridges the pushing and steering phases by generating intermediate reference states. Specifically, we select a set of key bodies, denoted by $\mathcal{K}$, which capture the humanoid's essential motion during phase transitions. The positions and orientations of these key bodies relative to the skateboard are then used to plan smooth and physically consistent transition trajectories.

At the end of the current phase, the humanoid's key body poses ($\mathbf{p}_{\text{end}}^{\mathcal{K}}, \mathbf{q}_{\text{end}}^{\mathcal{K}}$) are obtained online and serve as the initial condition for trajectory planning. Unlike predefined motions, these terminal poses vary across episodes and policies, resulting in a state-dependent transition problem.

For each phase, we also define a canonical stable pose ($\mathbf{p}_{\text{ref}}^{\mathcal{K}}, \mathbf{q}_{\text{ref}}^{\mathcal{K}}$) relative to the skateboard frame, representing the nominal start configuration of the subsequent phase. The transition is therefore formulated as a trajectory connecting the online terminal pose of the current phase to the fixed reference pose of the next phase. For example, the end pose of the pushing phase is planned toward the steering phase reference pose, and vice versa.

For body translations, the planned Cartesian position at timestep t is generated using an n -th order Bézier curve:

$$\mathbf{p}^{\mathcal{K}}(t) = \sum_{i=0}^n \binom{n}{i} (1-s)^{n-i} s^i \mathbf{p}_i^{\mathcal{K}}, \quad s = \frac{t-t_0}{t_f-t_0}, \quad (12)$$

where $\mathbf{p}_0^{\mathcal{K}} = \mathbf{p}_{\text{end}}^{\mathcal{K}}$ and $\mathbf{p}_n^{\mathcal{K}} = \mathbf{p}_{\text{ref}}^{\mathcal{K}}$ denote the initial and target positions, and $\{\mathbf{p}_i^{\mathcal{K}}\}_{i=1}^{n-1}$ are intermediate control points shaping the trajectory. Here, t_0 and t_f are the start and end times of the transition.

Body orientations are interpolated using spherical linear interpolation (slerp) between quaternions:

$$\mathbf{q}^{\mathcal{K}}(t) = \frac{\sin((1-s)\Omega)}{\sin \Omega} \mathbf{q}_{\text{end}}^{\mathcal{K}} + \frac{\sin(s\Omega)}{\sin \Omega} \mathbf{q}_{\text{ref}}^{\mathcal{K}}, \quad (13)$$

where $s = \frac{t-t_0}{t_f-t_0}$ and $\Omega = \arccos(\langle \mathbf{q}_{\text{end}}^{\mathcal{K}}, \mathbf{q}_{\text{ref}}^{\mathcal{K}} \rangle)$ is the angular distance between the quaternions. This ensures smooth and physically consistent orientation transitions while stepping onto or off the skateboard.

At each transitional timestep t , the policy receives a tracking reward that encourages it to follow the planned trajectories $(\mathbf{p}^{\mathcal{K}}(t), \mathbf{q}^{\mathcal{K}}(t))$. This trajectory-guided guidance shapes the exploration space, enabling reliable phase switching while retaining sufficient flexibility for the policy to adapt to dynamic interactions. By integrating trajectory-guided transitions with phase-specific learning objectives, the humanoid achieves stable and repeatable transitions between pushing and steering, facilitating consecutive skateboarding behaviors.

E. Sim-to-Real Transfer

1) *Skateboard Physical Identification*: To bridge the sim-to-real gap, we identify an equivalent PD model for the skateboard’s passive tilt dynamics (back view of Fig. 2) through a sequential analytical identification procedure based on its free-decay roll response [24]. Specifically, the skateboard is perturbed in roll and released, allowing the truck bushings and pivot friction to produce a naturally decaying oscillation. We first quantify energy dissipation by measuring two successive roll angle peaks, $\phi(t)$ and $\phi(t+T)$. The logarithmic decrement δ and the corresponding damping ratio ζ are computed as

$$\delta = \ln \frac{\phi(t)}{\phi(t+T)}, \quad \zeta = \frac{\delta}{\sqrt{4\pi^2 + \delta^2}}, \quad (14)$$

where T is the observed oscillation period. This damping ratio captures the intrinsic dissipative effects of the truck bushings and pivot interfaces. Given the identified ζ , we estimate the equivalent torsional stiffness of the truck assembly by matching the observed roll oscillation frequency. The undamped natural frequency is computed as:

$$\omega_n = \omega_d / \sqrt{1 - \zeta^2}, \quad (15)$$

where $\omega_d = 2\pi/T$ is the damped frequency. Approximating the skateboard’s roll inertia I using a rigid cuboid model dominated by the deck geometry, the torsional stiffness is identified as $k = I\omega_n^2$, which represents the restorative torque generated by the bushings under roll deformation. Finally, the velocity-dependent damping coefficient is obtained as $d = 2\zeta\sqrt{kI}$, modeling the dissipative torque arising from bushing hysteresis and pivot friction. Embedding the identified parameters (k, d) into simulation yields a physically grounded tilt spring-damper model that faithfully reproduces the passive truck mechanics observed on the real skateboard.

2) *Domain Randomization*: To further improve policy robustness and facilitate effective sim-to-real transfer, we employ domain randomization (DR) during training. The detailed configurations are provided in Table II.

V. EXPERIMENTS

A. Experiment Setup

In simulation, we use a standard $80\text{cm} \times 20\text{cm} \times 12\text{cm}$ skateboard model, the coupling mechanism in Eq. (1) is

TABLE II: Domain Randomization Parameters.

DR terms	Range	Unit
Robot Center of Mass	$\mathcal{U}(-2.5, 2.5)$	cm
Skateboard Center of Mass	$\mathcal{U}(-2.5, 2.5)$	cm
Default Root Position	$\mathcal{U}(-2.0, 2.0)$	cm
Default Joint Position	$\mathcal{U}(-0.01, 0.01)$	rad
Push Robot Base	$\mathcal{U}(-0.5, 0.5)$	m/s
Robot Body Friction	$\mathcal{U}(0.3, 1.6)$	-
Skateboard Deck Friction	$\mathcal{U}(0.8, 2.0)$	-

embedded into the simulator as an equivalent constraint. The details of skateboard model are provided in Appendix B. All trainings and evaluations are implemented in mjlab [45], which integrates MuJoCo physics [34] with the IsaacLab API [23]. Humanoid skateboarding policies are trained using the PPO [29] algorithm across 4,096 parallel environments, with each episode lasting 20 seconds. For each experiment, we evaluate 1,000 rollout episodes in simulation and report results over five random seeds. Additional training details are given in Appendix D. In real-world experiments, the policy is deployed on a Unitree G1 humanoid robot at 50 Hz, with joint positions tracked by a 500 Hz PD controller.

1) *Evaluation Metrics*: We adopt the following metrics to evaluate the humanoid skateboarding performance:

- **Task success rate** (E_{succ}): Defined as the percentage of episodes completed without termination.
- **Velocity tracking error** (E_{vel}): Measured as the mean absolute error between the commanded velocity v_{cmd} and the actual skateboard forward velocity v_{board} during the pushing phase: $E_{\text{vel}} = \mathbb{E}[|v_{\text{cmd}} - v_{\text{board}}|]$.
- **Heading tracking error** (E_{yaw}): Quantifies the accuracy of heading control during steering phase using the heading error: $E_{\text{yaw}} = \mathbb{E}[|\psi - \psi_{\text{board}}|]$.
- **Motion smoothness** (E_{smth}): Evaluated by aggregating joint angle variations across consecutive control steps, promoting temporally smooth and plausible behaviors.
- **Contact error** (E_{contact}): Defined as the per-step violation of the foot-board contact pattern, i.e., single-foot contact during pushing and double-foot contact during steering.

2) *Baselines*: To evaluate the effectiveness of HUSKY, we compare it with the following ablated baseline variants:

- **Pushing style**: We consider two baselines: (i) a **Tracking-based** variant [26] that directly tracks reference skateboarding motions from $\mathcal{M}$, and (ii) a **Gait-based** variant [25] in which the humanoid’s pushing follows a fixed gait schedule.
- **Steering strategy**: We compare our steering design with a **w/o Tilt Guidance** baseline, where the policy implicitly learn board tilt angles for steering without physics-guided tilt reference.
- **Transition mechanism**: We evaluate three variants: (i) **AMP Transition**, which uses phase-specific reference motions, including stepping onto and off the skateboard, and relies on style rewards to guide transitions, (ii) **Translation-only**, which guides only translational motion while ignoring orientation, underscoring the importance

TABLE III: Simulation Results.

Method	$E_{\text{succ}} \uparrow$	$E_{\text{vel}} \downarrow$	$E_{\text{yaw}} \downarrow$	$E_{\text{smth}} \downarrow$	$E_{\text{contact}} \downarrow$
Ablation on Pushing Style					
HUSKY-Tracking-Based	11.12 ± 3.86	0.435 ± 0.101	0.568 ± 0.092	0.044 ± 0.025	0.015 ± 0.010
HUSKY-Gait-Based	82.38 ± 7.25	0.102 ± 0.035	0.302 ± 0.041	0.043 ± 0.011	0.130 ± 0.072
HUSKY (ours)	100.00 ± 0.00	0.056 ± 0.013	0.208 ± 0.014	0.033 ± 0.005	0.001 ± 0.001
Ablation on Steering Strategy					
HUSKY-w/o-Tilt Guidance	96.72 ± 2.10	0.071 ± 0.010	0.233 ± 0.027	0.035 ± 0.017	0.002 ± 0.002
HUSKY (ours)	100.00 ± 0.00	0.056 ± 0.013	0.208 ± 0.014	0.033 ± 0.005	0.001 ± 0.001
Ablation on Transition Mechanism					
HUSKY-AMP Transition	85.12 ± 4.11	0.053 ± 0.025	0.265 ± 0.050	0.040 ± 0.007	0.394 ± 0.015
HUSKY-Translation-only	89.55 ± 2.30	0.064 ± 0.020	0.294 ± 0.075	0.039 ± 0.012	0.038 ± 0.012
HUSKY-Mixed Initialization	86.03 ± 1.70	0.067 ± 0.024	0.278 ± 0.092	0.037 ± 0.015	0.371 ± 0.025
HUSKY (ours)	100.00 ± 0.00	0.056 ± 0.013	0.208 ± 0.014	0.033 ± 0.005	0.001 ± 0.001

of jointly controlling both position and rotation during phase transitions, and (iii) *Mixed Initialization* [19] utilizes a balanced 50/50 mixture of pushing and steering phases during reset to facilitate comprehensive phase exploration.

B. Main Results

Table III summarizes the overall performance of HUSKY compared to the baselines across metrics. These results demonstrate that HUSKY enables robust humanoid skateboarding with stable, accurate, smooth motions and continuous, reliable phase transitions. The contributions of the key design choices are summarized as follows:

AMP enables flexible and natural pushing. For pushing, the *Tracking-based* baseline strictly follows reference motions, limiting its adaptability to varying contact timings and velocities and leading to the lowest success rate. The *Gait-based* baseline, lacking human motion priors, produces less smooth motions and higher contact errors, reflecting its difficulty in generating coordinated, human-like pushing behaviors.

Physics-guided steering enhances heading control. For steering, the *w/o Tilt Guidance* baseline learns board tilt angles implicitly, which reduces heading accuracy and highlights the importance of explicitly incorporating the lean-to-steer coupling for precise directional control.

Trajectory guidance is essential for phase transitions. For phase transitions, the *AMP Transition* baseline, which relies solely on full reference motions, achieves moderate success rates but incurs significant contact errors, as the robot fails to explore proper phase switching. The *Translation-only* baseline, which guides translation while omitting orientation alignment, exhibits poor heading tracking, causing the robot to align with the board direction rather than adopt a side-on posture and making steering control difficult.

C. More Analysis

Skateboard Modeling (Fig. 4). We validate the effect of skateboard modeling and visualize the resulting steering trajectories in Fig. 4 (a). Similar to prior simplified models [19], omitting the equality constraint in Eq. (1) prevents board tilting




Fig. 4: Steering Trajectories Visualizations. (a) Omitting lean-to-steer coupling prevents effective steering. (b) Incorporating physics-guided tilt guidance substantially increases the reachable heading range and precision.

from inducing truck steering, leaving the skateboard able only to glide straight forward with negligible turning capability.

Fig. 4 (b) compares policies trained without tilt guidance under identical steering commands by measuring the range of reachable headings. Without tilt guidance, the achievable heading range is narrow. In contrast, incorporating tilt guidance produces smooth turning trajectories and enables the humanoid to reach a substantially wider range of headings with higher precision. These results demonstrate that tilt guidance is essential for effective and precise turning control.

Phase Exploration (Fig. 5). To further analyze the necessity of transition guidance, we train policies without it under two settings: (i) all episodes are initialized to start in the pushing phase, matching our setup, and (ii) a 50/50 mixture of pushing and steering phases, as used in prior work [19] to promote phase exploration. We measure episode length and a steering contact reward, which encourages double-foot board contact and penalizes ground contact during steering.

In both settings, episode length increases rapidly during early training, yet the steering contact reward remains low, indicating persistent incorrect foot-board contact patterns. This suggests that policies trained without transition guidance fail to learn phase transitions, collapsing to a trivial pushing-only behavior and failing to properly execute the steering phase, even under mixed initialization. In contrast, HUSKY discovers correct contact patterns by mid-training, successfully learns foot-mounting transitions, and achieves higher rewards, demonstrating that trajectory-guided transitions are essential




Fig. 5: Training Performance Comparison. Episode length (left) and steering contact reward (right). Without transition guidance, the policy fails to maintain correct foot-board contacts. In contrast, the policy discovers correct contact patterns early in training and establishes stable phase transitions.




Fig. 6: Transition trajectories analysis. Representative trajectories during phase transitions. The humanoid maintains smooth, coordinated whole-body motions, with seamless transitions between pushing and steering.

for enabling phase switching and avoiding local optima.

Trajectories Analysis (Fig. 6). We further examine the humanoid’s skateboarding transitions by analyzing collected trajectory sequences. As shown in Fig. 6, the robot maintains smooth, coordinated whole-body motions with seamless transitions between pushing and steering. The trajectories display consistent foot placement and gradual body pose adjustments, reflecting strong temporal coherence and physical plausibility enabled by our trajectory-guided transitions.

D. Real-World Experiments

1) Overall Performance: Overall, Fig. 1 and the supplementary videos demonstrate that HUSKY achieves robust real-world skateboarding across diverse conditions. The system executes complete skateboarding behaviors including pushing, steering, and phase transitions, generalizes to different skateboard platforms, and performs reliably in both indoor and outdoor environments. Lean-to-steer control emerges naturally through body tilting, enabling heading regulation. The humanoid maintains smooth and coordinated foot placements during transitions, withstands external disturbances, and sustains multiple continuous skateboarding cycles.




Fig. 7: Details of feet motions during transitions. The humanoid pushes off the ground, mounts the skateboard, and adjusts orientation on the board.




Fig. 8: Effect of Skateboard Physical Identification. (a) Applying the compliant board parameters to the stiff board prevents mounting, as the real stiff board does not tilt under the robot’s step. (b) Applying the stiff board parameters to the compliant board causes over-leaning and loss of stability during steering.

2) Transition Details: We analyze the transitions from pushing to steering, focusing on detailed humanoid foot motion. In Fig. 7, the transition begins with the foot pushing against the ground to generate propulsion, followed by lifting and placing the foot onto the skateboard. Once on the board, the humanoid performs in-place adjustments, rotating the body to align the torso perpendicular to the skateboard deck, thereby facilitating stable steering. These motions closely match human behaviors and are consistent with our simulation data.

To account for the high friction of real skateboard decks, we apply DR during training, varying friction parameters to improve sim-to-real transfer of foot-mounting and body-alignment behaviors.

3) Role of System Identification: We evaluate the importance of skateboard system identification by focusing on the board’s tilt compliance, which determines how readily the board responds to humanoid leaning. Experiments are conducted using two skateboards with distinct tilt stiffness: one stiffer (black) and one more compliant (pink). As shown in Fig. 1 (a) and (b), when using the identified parameters corresponding to each physical board, the robot successfully performs push-and-steer maneuvers and achieves smooth board mounting and dismounting.

Cross-applying the identified parameters reveals the limitations of mismatched system models. In Fig. 8 (a), when using parameters identified from the compliant board on the stiffer board, the robot fails to mount. In simulation, the board

tilts under the robot’s stepping motion, allowing the policy to exploit this compliance for mounting, whereas the real stiff board remains nearly flat, breaking this assumption and preventing successful mounting. Conversely, in Fig. 8 (b), applying the stiff-board parameters to the compliant board causes excessive leaning and loss of stability during steering, since the policy is not trained for such compliant dynamics. These results highlight that accurate identification of board tilt compliance is critical for sim-to-real transfer of humanoid skateboarding policies.

VI. CONCLUSION AND FUTURE WORK

In this work, we present HUSKY, a physics-aware learning framework that enables humanoid robots to perform agile skateboarding with active propulsion and lean-to-steer control. By modeling the coupled humanoid–skateboard dynamics and explicitly capturing the relationship between board tilt and truck steering, we decompose the task into pushing, steering, and transition phases and develop phase-specific learning strategies. We further introduce a trajectory-guided transition mechanism to enable smooth and reliable phase switching, which is critical for long-horizon stability. Extensive simulation and real-world experiments demonstrate that HUSKY enables continuous and robust skateboarding behaviors on the Unitree G1 humanoid robot platform.

We emphasize that HUSKY is a principled paradigm with utility beyond skateboarding. Its core mechanisms—kinematic coupling for underactuated tools and transitions between contact topologies—extend to broader tasks. Despite strong simulation and real-world performance, several limitations remain and warrant future investigation.

Perception Integration. The limited camera field of view prevents reliable observation of the board and wheel–ground interactions, restricting perception-driven feedback in the control loop. Incorporating visual state estimation is important for enabling perception-aware skateboarding control.

Complex Terrains. Current experiments are conducted on relatively simple terrains, whereas human skateboarders routinely perform in complex environments such as skateparks while executing acrobatic maneuvers. Extending our framework to such scenarios will require richer motion priors and terrain-adaptive control strategies.

ACKNOWLEDGMENTS

Our training pipeline builds on mjlabs [45], we gratefully acknowledge the open-source communities and contributors behind the frameworks, and thank Unitree Robotics [35] for their support of the humanoid hardware used in this work.

REFERENCES

- [1] Arthur Allshire, Hongsuk Choi, Junyi Zhang, David McAllister, Anthony Zhang, Chung Min Kim, Trevor Darrell, Pieter Abbeel, Jitendra Malik, and Angjoo Kanazawa. Visual imitation enables contextual humanoid control. In *Proceedings of The Conference on Robot Learning*, Proceedings of Machine Learning Research, 2025.
- [2] Hongjun An, Wenhan Hu, Sida Huang, Siqi Huang, Ruanjun Li, Yuanzhi Liang, Jiawei Shao, Yiliang Song, Zihan Wang, Cheng Yuan, et al. Ai flow: Perspectives, scenarios, and approaches. *Vicinagearth*, 3(1):1, 2026.
- [3] Danil Belov, Artem Erkhov, Elizaveta Pestova, Ilya Osokin, Dzmitry Tsetserukou, and Pavel Osinenko. Quadrupedal robot skateboard mounting via reverse curriculum learning. *arXiv preprint arXiv:2505.06561*, 2025.
- [4] Qingwei Ben, Feiyu Jia, Jia Zeng, Junting Dong, Dahua Lin, and Jiangmiao Pang. Homie: Humanoid loco-manipulation with isomorphic exoskeleton cockpit. In *Robotics: Science and Systems*, 2025.
- [5] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. In *Conference on Robot Learning (CoRL)*, 2024.
- [6] Junchi Gu, Feiyang Yuan, Weize Shi, Tianchen Huang, Haopeng Zhang, Xiaohu Zhang, Yu Wang, Wei Gao, and Shiwu Zhang. Skater: Synthesized kinematics for advanced traversing efficiency on a humanoid robot via roller skate swizzles. *arXiv preprint arXiv:2601.04948*, 2026.
- [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] Xinyang Gu, Yen-Jen Wang, Xiang Zhu, Chengming Shi, Yanjiang Guo, Yichen Liu, and Jianyu Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. In *Robotics: Science and Systems*, 2024.
- [9] Zhaoyuan Gu, Junheng Li, Wenlan Shen, Wenhao Yu, Zhaoming Xie, Stephen McCrory, Xianyi Cheng, Abdulaziz Shamsah, Robert Griffin, C Karen Liu, et al. Humanoid locomotion and manipulation: Current progress and challenges in control, planning, and learning. *IEEE/ASME Transactions on Mechatronics*, 31(2):2300–2330, 2026.
- [10] 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. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2026.
- [11] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris M. Kitani, Changliu Liu, and Guanya Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. In *8th Annual Conference on Robot Learning*, 2024.
- [12] Tairan He, Jiawei Gao, Wenli Xiao, Yuanhang Zhang, Zi Wang, Jiashun Wang, Zhengyi Luo, Guanqi He, Nikhil Sobanbabu, Chaoyi Pan, Zeji Yi, Guannan Qu, Kris Kitani, Jessica Hodgins, Linxi ”Jim” Fan, Yuke Zhu, Changliu Liu, and Guanya Shi. Asap: Aligning simula-

- tion and real-world physics for learning agile humanoid whole-body skills. *Robotics: Science and Systems*, 2025.
- [13] Kaihui Huang, Heming Feng, Wei Meng, Tianqi Wei, and Tianjiang Hu. Learning-based quadruped robot framework for locomotion on dynamic rigid platforms. In *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 17746–17753. IEEE, 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] Heejin Jeong and Daniel D Lee. Efficient learning of stand-up motion for humanoid robots with bilateral symmetry. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1544–1549. IEEE, 2016.
 - [16] Alexander S Kuleshov. Mathematical model of the skateboard. In *ISBS-Conference Proceedings Archive*, 2006.
 - [17] Yitang Li, Yuanhang Zhang, Wenli Xiao, Chaoyi Pan, Haoyang Weng, Guanqi He, Tairan He, and Guanya Shi. Hold my beer: Learning gentle humanoid locomotion and end-effector stabilization control, 2025. URL <https://arxiv.org/abs/2505.24198>.
 - [18] Qiayuan Liao, Takara E Truong, Xiaoyu Huang, 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.
 - [19] Hang Liu, Sangli Teng, Ben Liu, Wei Zhang, and Maani Ghaffari. Discrete-time hybrid automata learning: Legged locomotion meets skateboarding. *arXiv preprint arXiv:2503.01842*, 2025.
 - [20] Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9997–10003. IEEE, 2025.
 - [21] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, et al. Isaac gym: High performance gpu-based physics simulation for robot learning. *arXiv preprint arXiv:2108.10470*, 2021.
 - [22] Elias Mikkola. Online planning and control of physics-based skateboarding animation. Master’s thesis, Aalto University, School of Science, 2023. URL <https://aaltodoc.aalto.fi/items/3d76eb38-cedf-473b-b19f-e6e9d604215b>.
 - [23] Mayank Mittal, Pascal Roth, James Tigue, Antoine Richard, Octi Zhang, Peter Du, Antonio Serrano-Muñoz, Xinjie Yao, René Zurbrügg, Nikita Rudin, et al. Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning. *arXiv preprint arXiv:2511.04831*, 2025.
 - [24] Katsuhiko Ogata. *Modern control engineering*. Prentice hall, 2010.
 - [25] Open-X-Humanoid. Tienkung-lab: An open-source research platform for humanoid robots. <https://github.com/Open-X-Humanoid/TienKung-Lab>, 2024.
 - [26] 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.
 - [27] 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.
 - [28] Ilija Radosavovic, Tete Xiao, Bike Zhang, Trevor Darrell, Jitendra Malik, and Koushil Sreenath. Real-world humanoid locomotion with reinforcement learning. *Science Robotics*, 9(89):eadi9579, 2024.
 - [29] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
 - [30] Jiyuan Shi, Xinzhe Liu, Dewei Wang, Sören Schwertfeger, Chi Zhang, Fuchun Sun, Chenjia Bai, Xuelong Li, et al. Adversarial locomotion and motion imitation for humanoid policy learning. *Advances in Neural Information Processing Systems*, 38:73918–73949, 2026.
 - [31] Haoming Song, Delin Qu, Yuanqi Yao, Qizhi Chen, Qi Lv, Yiwen Tang, Modi Shi, Guanghui Ren, Maoqing Yao, Bin Zhao, et al. Hume: Introducing system-2 thinking in visual-language-action model. *arXiv preprint arXiv:2505.21432*, 2025.
 - [32] Noriaki Takasugi, Kunio Kojima, Shu nichu Nozawa, Yohei Kakiuchi, Kei Okada, and Masayuki Inaba. Real-time skating motion control of humanoid robots for acceleration and balancing. In *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1356–1363. IEEE, 2016.
 - [33] Noriaki Takasugi, Kunio Kojima, Shu nichu Nozawa, Fumihito Sugai, Kakiuchi Yohei, Kei Okada, and Masayuki Inaba. Extended three-dimensional walking and skating motion generation for multiple noncoplanar contacts with anisotropic friction: Application to walk and skateboard and roller skate. *IEEE Robotics and Automation Letters*, 4(1):9–16, 2018.
 - [34] 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.
 - [35] Unitree Robotics. Humanoid robot G1_Humanoid Robot Functions_Humanoid Robot Price — Unitree Robotics, 2025. URL <https://www.unitree.com/g1/>. <https://www.unitree.com/g1/>.
 - [36] Balazs Varszegi, Denes Takacs, and Gabo r Stepan. Skateboard: a human controlled non-holonomic system.

- m. In *International Design Engineering Technical Conferences and Computers and Information in Engineering Conference*, volume 57168, page V006T10A066. American Society of Mechanical Engineers, 2015.
- [37] Dewei Wang, Xinmiao Wang, Xinzhe Liu, Jiyuan Shi, Yingnan Zhao, Chenjia Bai, and Xuelong Li. More: Mixture of residual experts for humanoid life-like gaits learning on complex terrains. *arXiv preprint arXiv:2506.08840*, 2025.
- [38] Huayi Wang, Wentao Zhang, Runyi Yu, Tao Huang, Junli Ren, Feiyu Jia, Zirui Wang, Xiaojie Niu, Xiao Chen, Jiahe Chen, et al. Physshi: Towards a real-world generalizable and natural humanoid-scene interaction system. *arXiv preprint arXiv:2510.11072*, 2025.
- [39] 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.
- [40] 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. *arXiv preprint arXiv:2502.17219*, 2025.
- [41] Weiji Xie, Jinrui Han, Jiakun Zheng, Huanyu Li, Xinzhe Liu, Jiyuan Shi, Weinan Zhang, Chenjia Bai, and Xuelong Li. Kungfubot: Physics-based humanoid whole-body control for learning highly-dynamic skills. *Advances in Neural Information Processing Systems*, 38: 62406–62433, 2026.
- [42] Zhe Xu, Mohamed Al-Khulaqui, Hanxin Ma, Jiajun Wang, Quanbin Xin, Yangwei You, Mingliang Zhou, Diyun Xiang, and Shiwu Zhang. Optimization based dynamic skateboarding of quadrupedal robot. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 8058–8064. IEEE, 2024.
- [43] 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.
- [44] Shaopu Yang, Yongjie Lu, and Shaohua Li. An overview on vehicle dynamics. *International Journal of Dynamics and Control*, 1(4):385–395, 2013.
- [45] Kevin Zakka, Qiayuan Liao, Brent Yi, Louis Le Lay, Koushil Sreenath, and Pieter Abbeel. mjlal: A lightweight framework for gpu-accelerated robot learning, 2026. URL <https://arxiv.org/abs/2601.22074>.
- [46] Yanjie Ze, Zixuan Chen, João Pedro Araújo, Zi an Cao, Xue Bin Peng, Jiajun Wu, and C. Karen Liu. Twist: Teleoperated whole-body imitation system. *arXiv preprint arXiv:2505.02833*, 2025.
- [47] 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, 2025. URL <https://arxiv.org/abs/2511.02832>.
- [48] Chong Zhang, Wenli Xiao, Tairan He, and Guanya Shi. Wococo: Learning whole-body humanoid control with sequential contacts, 2024. URL <https://arxiv.org/abs/2406.06005>.
- [49] Yuanhang Zhang, Yifu Yuan, Prajwal Gurunath, Ishita Gupta, Shayegan Omidshafiei, Ali-akbar Agha-mohammadi, Marcell Vazquez-Chanlatte, Liam Pedersen, Tairan He, and Guanya Shi. Falcon: Learning force-adaptive humanoid loco-manipulation. *arXiv preprint arXiv:2505.06776*, 2025.
- [50] Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. *arXiv preprint arXiv:2406.10759*, 2024.