

The Blind Skateboarder: Unlocking Dexterous Robotic Skateboarding via Interactive Perception

Aditya Bhatt^{1,2}, Ansh Prakash^{1,2}, Joao Carvalho^{1,2}, Puze Liu³, and Jan Peters^{1,2}




Fig. 1: Proprioception-only humanoid skateboarding. Two examples from an eight-skill repertoire executed by the Unitree G1 using onboard proprioception without direct board sensing. *Funnel*-trained policies acquire maneuver-specific footholds through contact before performing dynamic tricks. External links to watch the skills in action: One uncut trick video, and one single-tricks compilation video.

Abstract—Skateboarding combines dynamic whole-body control with manipulation of an underactuated object. We present, to our knowledge, the first humanoid system to execute skateboard tricks, with a diverse skill library comprising mounting, pushing, dismounting, manuals, ollies, shove-its, and kickflips, and walk-the-dog. Crucially, these skills are *blind*: they rely on the robot’s onboard proprioception alone, without any access to board pose, vision, tracking markers, tactile sensors, or force/torque feedback. To make narrow reference-imitation policies usable from diverse initial states, we introduce a *funneling* curriculum that teaches each policy to seek trick-specific footholds through contact before execution. Simulation experiments quantify foothold acquisition and skill execution for four skills, while we deploy the complete eight-skill repertoire on the real robot. State-estimation analysis and controlled interventions provide evidence that contact-driven interactive perception enables blind manipulation of the moving board.

I. INTRODUCTION

Humanoid robots can now execute increasingly dynamic whole-body motions, yet most examples of this either interact with static supports (e.g. walls, stairs, or raised platforms) or, during manipulation, assume access to external object state. Skateboarding pushes the limits of this paradigm into

the regime of *agile whole-body manipulation*: the robot must balance on and deliberately manipulate an underactuated board that moves in response to every contact. Existing robotic skateboarding has largely focused on balance, propulsion, and steering. [1]–[4] However, the practice of skateboarding is defined by its *tricks*, which to the best of our knowledge have not been attempted in robotics. Tricks are especially demanding because each maneuver requires a distinct foot configuration and a precisely timed contact sequence, and visual tracking is often unreliable due to fast motion and self-occlusion. In this preliminary report of ongoing work on athletic interactive perception, we train a Unitree G1 to execute a unified library of eight behaviors comprising access, riding, dismounting, balance, aerial and flip tricks.

At deployment, the robot receives only histories of on-board proprioceptive measurements and previous actions; it never observes the skateboard pose. This creates a structural challenge: before executing a maneuver, the feet must reach narrow, trick-specific regions of a board whose position changes under contact and is frequently occluded from view. We treat this preparatory behavior as *interactive perception* [5], [6]. By probing and moving across the deck, the robot receives contact-dependent changes in its joint and inertial measurements, allowing it to locate mechanically

¹DFKI GmbH, Darmstadt, Germany.

²Intelligent Autonomous Systems Group, TU Darmstadt, Darmstadt, Germany.

³Tongji University, Shanghai Research Institute for Autonomous Intelligent Systems, Shanghai, China.

meaningful regions such as the trucks, tail, and board edges. Each policy can therefore acquire its required footholds through interaction before initiating the commanded maneuver.

We train all eight behaviors using a DeepMimic-style reference-motion imitation framework [7] that builds upon the reward terms and PD gains used in BeyondMimic [8]. Motion imitation alone, however, produces policies that succeed only near their reference initial states, whereas deployed skills must start from the diverse configurations left by preceding behaviors. Our funnel curriculum adds a preparatory idle phase and progressively widens the reset distribution away from the reference motion. Each policy thereby learns to transform varied initial footholds into the narrow preconditions required by its maneuver, execute the trick, and return to a controllable on-board state. Our experiments quantify success rates for the tricks, foothold convergence across randomized trials, and use board-state estimation and controlled interventions to show that successful acquisition depends on contact-driven interactive perception.

In summary, our contributions are: (1) a proprioception-only humanoid system spanning eight skateboard behaviors; (2) a unified funnel curriculum that turns reference-imitation skills into policies capable of acquiring trick-specific footholds from diverse states; and (3) empirical evidence that this acquisition emerges from contact-driven interactive perception.

II. APPROACH

To illustrate the fact that proprioception can be sufficient for advanced underactuated dexterous manipulation, we choose the robot’s onboard IMU and joint angle encoders as the only sensors used for the policy’s observations.

a) Proprioception-only control: Each behavior is a partially observed control problem. At runtime, the policy receives a history [9] of joint positions and velocities, inertial measurements, and previous actions, but no skateboard pose, vision, tracking markers, tactile sensing, or force/torque feedback. It outputs desired joint positions at 50 Hz, tracked by low-level PD controllers. External contact creates differences between commanded and realized motion, allowing the observation history to encode information about the board’s position, geometry, and mechanical response. We refer to the board-relative foot configuration required to initiate a maneuver as its *foothold*.

b) Reference-imitation skills: We train each of the eight behaviors as a separate DeepMimic-style policy [7] using PPO [10] in Isaac Lab [11]. A reference motion supplies phase-conditioned whole-body targets, while the actor operates from a 5-step proprioceptive history and consumes the skill’s reference motion streamed during execution. During training only, an asymmetric critic receives privileged simulator state, including the true board state. We randomize motor gains, body and board dynamics, friction, and sensor biases to support transfer to hardware; none of this privileged information is available to the deployed policy.

c) Funnel curriculum: Standard reference imitation teaches a policy to execute a motion when initialized near its trajectory, but not to enter that trajectory from realistic on-board states. We therefore augment every skill with an idle phase before and after execution. Training begins from reference-motion resets and progressively blends in randomized states with varied robot, board, and foot configurations. During the pre-execution idle phase, the policy must remain balanced and prepare for the commanded maneuver; when triggered, the reference phase advances, and the trick is executed. This curriculum turns each narrow imitation policy into a *funnel*: through contact, it drives a broad set of initial footholds into the small region *required* by the maneuver, executes the reference motion, and returns to a stable state from which another skill can begin.

d) Proprioceptive state probe: To measure what information about the skateboard is decodable from the observation history, we train an auxiliary estimator to predict the board pose from the same five-step proprioceptive history available to the actor. The estimator is supervised using ground-truth simulator state during training, but its predictions are never provided to the policy and do not affect control. It therefore serves only as a diagnostic probe: by aligning its prediction error with contact events and maneuver phases, we can quantify when interaction makes the otherwise unobserved board state identifiable from proprioception.


III. EXPERIMENTS

We evaluate three questions: whether the learned policies execute the intended skateboard behaviors, whether diverse initial foot placements converge to trick-specific footholds, and whether this convergence depends on information acquired through contact. We deploy all eight skills on the real robot, while our quantitative analysis focuses on the four skills for which completed randomized evaluations are available: mounting, manual, ollie, and pop shove-it. All policies use only proprioceptive observations during evaluation; privileged board state is used solely to compute metrics and train the diagnostic state estimator. Policies are trained in Isaac Lab and evaluated in MuJoCo without additional fine-tuning.


a) Skill execution: We evaluate each of the four quantitative skills from 100 randomized initial conditions. Mounting, manual, ollie, and pop shove-it achieve success rates of 75%, 89%, 91%, and 73%, respectively. A trial succeeds when the maneuver completes without either foot touching the ground, the pelvis tilting more than 45° from the upright orientation, or the pelvis falling below 0.5 m. We additionally deploy the complete eight-skill repertoire on hardware; quantitative hardware evaluation is left to future work.

b) Controlled intervention: We test whether foothold acquisition relies on interaction rather than memorized board geometry. When placed near the nose of a board one meter longer than its training board, the pop shove-it policy traverses to the opposite end and acquires its tail foothold in 3.3 s during a 6 s idle phase.

A. Pop Shove-it policy rollout in MuJoCo



B. 100 independent Pop Shove-It trials



C. Board estimation error and foot contact forces




Fig. 2: Pop Shove-it execution study. In A, an episode begins with the robot standing at a random position on the board. In the first 1s idle phase, policy moves the robot’s feet to the trick’s required foothold. In the next 1.7s, it executes the trick by jumping off the board while simultaneously *scooping* the tail backwards to spin the board 180°, while popping the tail to lift the board off the ground. The core airborne phase lasts only approx. 0.3s. Just like a skilled skateboarder, it catches the airborne board with its left foot and then lands it securely with both feet. After landing, the policy immediately brings the feet back to the required foothold. **B.** We record 100 independent trials; despite diverse on-board starts, we observe strong convergence to the required foothold, graduating to 73 stable landings. **C.** We analyze the error of the proprioceptive board pose estimator in those 73 successful trials and find that the policy is able to localize the board sharply in the foothold acquisition phase, driving an average positioning error of 11cm to under 2cm. The high-velocity airborne flip phase—when the robot completely loses contact with the board—accompanies an *increase* in the uncertainty, which is again sharply tamped down upon landing. Shaded regions indicate one standard deviation.




Fig. 3: Generalization to altered board geometry. We train the policy with a standard skateboard size (top), but test it additionally with a board extended by 1 m (bottom). Starting near the nose of the board, the blind robot policy actively traverses the unfamiliar deck—without accidentally stepping off the board—and acquires the pop-shove-it foothold at the tail end, all without directly sensing any board pose.

IV. CONCLUSION

We presented a unified framework for eight proprioception-only humanoid skateboard skills. A funnel curriculum expands reference-imitation policies beyond narrow trajectory initializations, enabling them to acquire maneuver-specific footholds through contact before

execution. Quantitative evaluation of mounting, manuals, ollies, and pop shove-its shows robust execution and convergence from randomized states, while the diagnostic board-state estimator and controlled interventions provide evidence for contact-driven interactive perception. Our current quantitative study is limited to four skills in simulation, with the complete repertoire deployed on




Fig. 4: Foothold funnel trajectories. Board-relative trajectories of the *left* and *right* foot across 100 independent trials from diverse initial conditions, for each skill, from the preparatory idle phase, through trick execution, until idling again. Despite being blind, our purely proprioceptive policies *funnel* these highly varied feet positions into a narrow set of foothold outcomes before and after trick execution: corresponding to the maneuver-specific regions learned from the reference motions. Convergence is tighter for manuals and ollies (narrow required footholds) than for pop shove-its

hardware but evaluated qualitatively. Broader real-world evaluation remains an important direction for future work.

REFERENCES

- [1] H. Liu, S. Teng, B. Liu, W. Zhang, and M. Ghaffari, “Discrete-time hybrid automata learning: Legged locomotion meets skateboarding,” *arXiv preprint arXiv:2503.01842*, 2025.
- [2] M. Yoon, J. Jeong, and S.-E. Yoon, “Phase-aware policy learning for skateboard riding of quadruped robots via feature-wise linear modulation,” *arXiv preprint arXiv:2602.09370*, 2026. [Online]. Available: <https://arxiv.org/abs/2602.09370>
- [3] W. Thibault, V. Rajendran, W. Melek, and K. Mombaur, “Learning skateboarding for humanoid robots through massively parallel reinforcement learning,” *arXiv preprint arXiv:2409.07846*, 2024.
- [4] J. Han, D. Wang, C. Zhang, X. Liu, P. Luo, C. Bai, and X. Li, “Husky: Humanoid skateboarding system via physics-aware whole-body control,” *arXiv preprint arXiv:2602.03205*, 2026, accepted to Robotics: Science and Systems (RSS).
- [5] J. Sinapov, T. Bergquist, C. Schenck, U. Ohiri, S. Griffith, and A. Stoytchev, “Interactive object recognition using proprioceptive and auditory feedback,” *The International Journal of Robotics Research*, vol. 30, no. 10, pp. 1250–1262, 2011.
- [6] W. Fu, H. Li, I. X. He, S. Tellex, and S. Sridhar, “Unitac: Whole-robot touch sensing without tactile sensors,” *arXiv preprint arXiv:2507.07980*, 2025.
- [7] 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*, vol. 37, no. 4, pp. 143:1–143:14, jul 2018. [Online]. Available: <https://doi.org/10.1145/3197517.3201311>
- [8] Q. Liao, T. E. Truong, X. Huang, Y. Gao, G. Tevet, K. Sreenath, and C. K. Liu, “BeyondMimic: From motion tracking to versatile humanoid control via guided diffusion,” *arXiv preprint arXiv:2508.08241*, 2025. [Online]. Available: <https://arxiv.org/abs/2508.08241>
- [9] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “RMA: Rapid motor adaptation for legged robots,” in *Robotics: Science and Systems*, 2021.
- [10] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017. [Online]. Available: <https://arxiv.org/abs/1707.06347>
- [11] M. Mittal *et al.*, “Isaac lab: A gpu-accelerated simulation framework for multi-modal robot learning,” *arXiv preprint arXiv:2511.04831*, 2025. [Online]. Available: <https://arxiv.org/abs/2511.04831>