Learning Surfing-like Balance without Water Simulation

Hauk Nam*
Hanyang University
Seoul, Republic of Korea
namhoyog@naver.com

Changho Lee*
Hanyang University
Seoul, Republic of Korea
changholee@hanyang.ac.kr

Yoonsang Lee†
Hanyang University
Seoul, Republic of Korea
yoonsanglee@hanyang.ac.kr

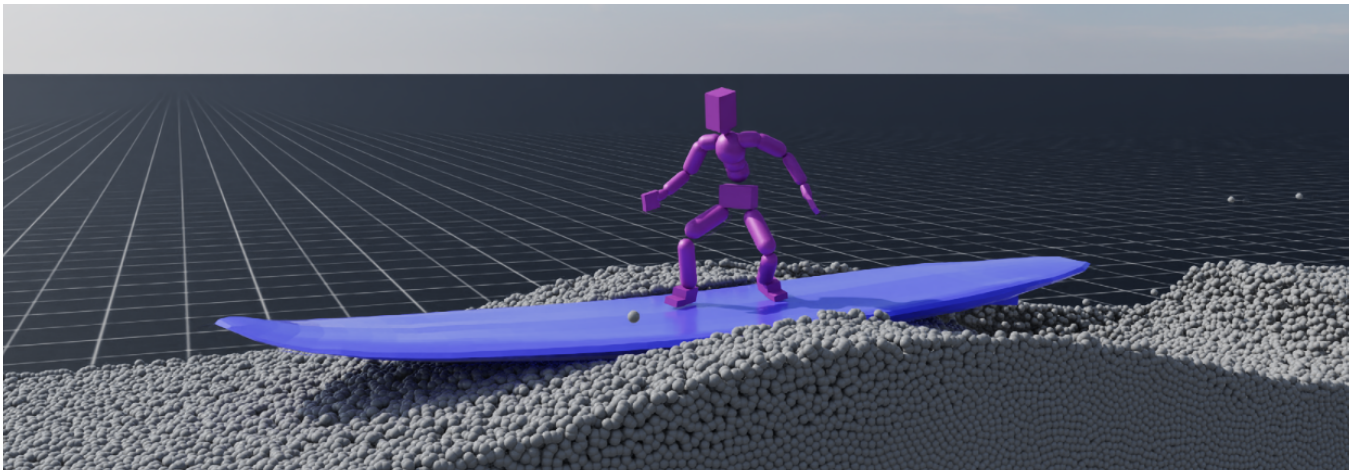

Figure 1: Our method learns surfing-like balance control without water simulation. The learned policy maintains stable balance on a moving board under wave conditions at runtime.

Abstract

Controlling characters in fluid environments such as water remains challenging due to complex dynamics and high simulation cost. In particular, surfing requires maintaining balance on a continuously moving support, making stable control difficult without accurate physical modeling. In this work, we present a method for learning surfing-like balance control without relying on water simulation. Our approach combines a hierarchical control framework with a stage-based training scheme that progressively introduces dynamic board behavior. The low-level policy generates full-body motion from partial joint targets, while the high-level policy adapts these targets based on environmental conditions. Despite being trained entirely in a non-fluid setting, the learned policy maintains stable balance on a moving board under wave conditions at runtime. These results suggest that surfing-like balance can be achieved without explicitly modeling fluid dynamics through appropriate control structures and training strategies.

*These authors contributed equally.
†Corresponding author.

SIGGRAPH Posters '26, Los Angeles, CA, USA

ACM ISBN 979-8-4007-2548-7/26/07
https://doi.org/10.1145/3799825.3818707

CCS Concepts

• **Computing methodologies → Physical simulation; Reinforcement learning**.

ACM Reference Format:
Hauk Nam, Changho Lee, and Yoonsang Lee. 2026. Learning Surfing-like Balance without Water Simulation. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH Posters '26), July 19–23, 2026, Los Angeles, CA, USA*. ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3799825.3818707

1 Introduction

Controlling characters in complex dynamic environments, such as water, remains challenging. In particular, surfing involves maintaining balance on a moving board while continuously reacting to external disturbances, making it a demanding control problem. This tight interaction with external dynamics makes stable control difficult without accurate physical modeling. However, fluid environments introduce highly nonlinear dynamics and significant computational cost, limiting the efficiency of reinforcement learning.

Prior work on physics-based character control has demonstrated strong performance in generating diverse motions [Kim et al. 2025; Peng et al. 2018]. However, incorporating fluid environments significantly increases simulation cost and instability. Recent approaches attempt to approximate character–fluid interactions, but still rely on modeling fluid dynamics or incorporating additional simulation components [Dou et al. 2025]. These limitations motivate learning control policies for dynamic environments without explicitly simulating fluids.

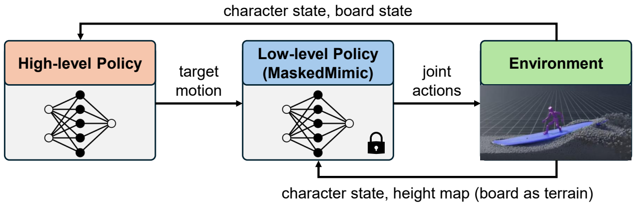

Figure 2: Overview of our hierarchical control framework. The high-level policy predicts target motions based on the character and board states, while the low-level policy produces joint actions that track these targets and is pretrained and kept fixed during training.

In this work, we present a method for learning surfing-like balance control without relying on water simulation. Instead of modeling fluid dynamics, we expose the policy to progressively changing board dynamics through a stage-based training scheme. We employ a hierarchical control framework, where a low-level policy drives the character using partial joint targets, and a high-level policy adapts these targets based on environmental conditions.

Despite being trained entirely without water, the resulting policy demonstrates stable balance behavior on boards under wave conditions at runtime. These results suggest that surfing-like balance can be achieved without explicitly modeling fluid dynamics through appropriate control structures and training strategies.

2 Related Work

Reinforcement learning for whole-body control has largely been driven by imitation learning-based approaches [Peng et al. 2021; Zhang et al. 2025]. Methods such as MaskedMimic [Tessler et al. 2024] enable full-body control from partial motion constraints, providing flexibility under sparse or multi-modal inputs. Balance control on dynamic support surfaces has also been studied [Xi et al. 2019; Yoon et al. 2024], but these approaches typically consider controlled or simplified platform dynamics. Learning in fluid settings remains challenging due to high computational cost and complex interactions. Recent work such as CFC [Dou et al. 2025] explores approximations of character–fluid coupling using a two-level world model to enable policy learning in such environments. In contrast, we avoid modeling fluid interactions altogether and train policies without water, while still achieving surfing-like balance under wave conditions.

3 Method

Hierarchical Control Structure. We employ a hierarchical control framework composed of a low-level policy and a high-level policy (Figure 2). The low-level policy controls the character by tracking target motions defined on a subset of joints, while the high-level policy adjusts these targets based on the current environmental state. This separation allows the low-level policy to focus on stable motion control, while the high-level policy adapts the behavior to external dynamics such as board movement, enabling effective control in dynamic environments.

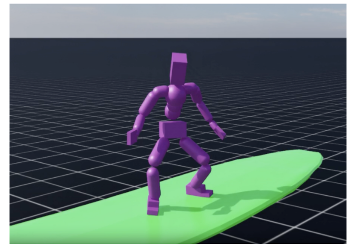

Figure 3: Predefined surfing pose.

Low-level Control with Target Motion. We adopt a pretrained MaskedMimic [Tessler et al. 2024] model as our low-level policy, as it enables full-body control from partial joint targets without requiring explicit reference motions. This is particularly suitable for tasks such as surfing, where constructing task-specific motion datasets is challenging. By specifying a small set of key joint targets, we define desired poses in a structured manner, while the low-level policy tracks these targets to produce consistent full-body motion. Furthermore, as the policy is pretrained on diverse motion data, it generalizes to previously unseen states without requiring task-specific motion datasets.

The low-level policy takes as input target motions defined on a subset of key joints and outputs joint actions that track these targets. We consider six key joints: the root, head, both hands, and both feet. Target motions are constructed by combining a predefined surfing pose (Figure 3) with offsets predicted by the high-level policy. Position targets are defined for these joints, while orientation targets are applied only to the feet. Velocity targets are specified for the root as the sum of the board velocity and high-level offsets. The remaining joints are implicitly controlled by the policy.

The policy additionally takes the current character state following MaskedMimic and heightmap observations that encode the board as terrain.

High-level policy. The high-level policy takes as input the character state and board state (orientation and motion), and predicts residual offsets that define the target motion by adjusting a predefined surfing pose. The policy is trained with the following reward function:

$$r = w_1 r_{\text{board}} + w_2 r_{\text{contact}} + w_3 r_{\text{stability}} + w_4 r_{\text{foot}} + w_5 r_{\text{head}} + w_6 r_{\text{smooth}}. \tag{1}$$

Here, r_{board} encourages stable board orientation. Contact-related terms ($r_{\text{contact}}, r_{\text{stability}}$) promote consistent foot contact and reduce slipping. Foot-related terms (r_{foot}) guide foot placement and height relative to the board. r_{head} encourages maintaining a forward-facing head orientation. Finally, r_{smooth} penalizes abrupt changes in target motion. Each term is implemented using nonlinear penalty functions (e.g., exponential or tanh) that penalize deviations from desired states.

Stage-based Training. We employ a stage-based training scheme that progressively introduces dynamic board behavior without using fluid simulation. In Stage 1, the board is suspended in the air with a restoring torque applied to its pitch, while no stabilization is applied to roll. This setup allows the policy to learn balance control by compensating for lateral instability, while the board remains level in the forward direction. In Stage 2, we build upon Stage 1 by introducing forward motion and a gradually increasing target pitch. Specifically, the board moves forward while its target pitch shifts from horizontal to a forward-tilted configuration over time. This results in coupled translational and rotational dynamics that resemble surfing conditions. By progressively augmenting the board dynamics in this manner, the policy learns to handle motion and tilt interactions without requiring explicit fluid simulation.

4 Results

We evaluate the learned policy in wave environments at runtime, despite being trained entirely without water simulation (Figure 1). The policy maintains balance on a moving board under various wave conditions, including changes in wave height and speed. While performance degrades in more extreme conditions (e.g., large waves), the policy often recovers from disturbances and regains balance. These results demonstrate that surfing-like balance can emerge without explicitly modeling fluid dynamics.

We further evaluate the effect of stage-based training by comparing policies trained with Stage 1 only and with both Stage 1 and Stage 2. Under the same evaluation setting, the success rate improves from 15% to 60% in the presence of forward motion and time-varying board pitch, indicating that Stage 2 substantially improves the policy's ability to maintain balance under dynamic board conditions. The average survival duration also increases from 198.4 to 283.3 steps (1 step = 1/30 s), although this metric is capped by the episode limit of 350 steps.

5 Conclusion and Future Work

We presented a method for learning surfing-like balance control without relying on water simulation. By combining a hierarchical control framework with stage-based training, the policy learns to maintain balance on dynamically moving boards and exhibits stable behavior under wave conditions at runtime. This result suggests that surfing-like balance can be achieved without explicitly modeling fluid dynamics, through appropriate control structures and training strategies. Future work includes extending to more realistic wave conditions and exploring fine-tuning in environments with learned approximations of fluid interactions.

Acknowledgments

This work was supported by Culture, Sports and Tourism R&D Program through the Korea Creative Content Agency grant funded by the Ministry of Culture, Sports and Tourism in 2024 (RS-2024-00399136).

References

Zhiyang Dou, Chen Peng, Xinyu Lu, Xiaohan Ye, Lixing Fang, Yuan Liu, Wenping Wang, Chuang Gan, Lingjie Liu, and Taku Komura. 2025. CFC: Simulating Character-Fluid Coupling using a Two-Level World Model. *ACM Trans. Graph.* 44, 6, Article 199 (Dec. 2025), 17 pages. doi:10.1145/3763318

Minsu Kim, Eunho Jung, and Yoonsang Lee. 2025. PhysicsFC: Learning User-Controlled Skills for a Physics-Based Football Player Controller. *ACM Transactions on Graphics* 44, 4 (July 2025), 1–21. doi:10.1145/3731425

Xue Bin Peng, Pieter Abbeel, Sergey Levine, and Michiel van de Panne. 2018. DeepMimic: example-guided deep reinforcement learning of physics-based character skills. *ACM Trans. Graph.* 37, 4, Article 143 (July 2018), 14 pages. doi:10.1145/3197517.3201311

Xue Bin Peng, Ze Ma, Pieter Abbeel, Sergey Levine, and Angjoo Kanazawa. 2021. AMP: adversarial motion priors for stylized physics-based character control. *ACM Trans. Graph.* 40, 4, Article 144 (July 2021), 20 pages. doi:10.1145/3450626.3459670

Chen Tessler, Yunrong Guo, Ofir Nabati, Gal Chechik, and Xue Bin Peng. 2024. MaskedMimic: Unified Physics-Based Character Control Through Masked Motion Inpainting. *ACM Transactions on Graphics (TOG)* (2024).

Ao Xi, Thushal Wijekoon Mudiyanselage, Dacheng Tao, and Chao Chen. 2019. Balance control of a biped robot on a rotating platform based on efficient reinforcement learning. *IEEE/CAA Journal of Automatica Sinica* 6, 4 (2019), 938–951. doi:10.1109/JAS.2019.1911567

Minsung Yoon, Heechan Shin, Jeil Jeong, and Sung-Eui Yoon. 2024. Learning-based Adaptive Control of Quadruped Robots for Active Stabilization on Moving Platforms. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 701–708. doi:10.1109/IROS58592.2024.10802150

Ziyu Zhang, Sergey Bashkirov, Dun Yang, Yi Shi, Michael Taylor, and Xue Bin Peng. 2025. Physics-Based Motion Imitation with Adversarial Differential Discriminators. In *Proceedings of the SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25)*. ACM, 1–12. doi:10.1145/3757377.3763819