

Walking on Ice: Adaptive Gait Control of Humanoid Robots Based on Visual Prediction and Proprioceptive Estimation for Variable Friction Environments

Yixuan Shen^{1,2}, Rongqiang Zhao^{1,2}, Ruonan Li^{1,2,*} and Jie Liu^{1,2,*}

¹Faculty of Computing, Harbin Institute of Technology

²State Key Laboratory of Smart Farm Technologies and Systems, China

25B903128@stu.hit.edu.cn, zhaorq@hit.edu.cn, 20B951005@stu.hit.edu.cn, jieliu@hit.edu.cn

Abstract

Humanoid robots have been increasingly deployed in real-world environments such as intelligent manufacturing, healthcare, and agriculture, from indoor spaces to outdoor terrains where friction conditions vary significantly. However, existing approaches such as domain randomization or employing estimation struggle to handle low-friction gait control and fail to address variable friction conditions, making it difficult to ensure control robustness. To address these challenges, we propose Walking on Ice (WoI), a friction-aware locomotion framework that combines proactive prediction with reactive adaptation for varying surface friction. We introduce a VLM-based friction predictor that enables the policy to take anticipatory actions before friction changing, thereby addressing variable-friction scenarios. Additionally, we design a friction estimator combined with a joint history encoder that perceives proprioceptive states to handle low-friction surfaces. Experimental results demonstrate that WoI improves success rate by 8.5% and reduces slip rate by 5.15% over the state-of-the-art method Denoising World Model Learning (DWL). Particularly, in extreme cases, WoI achieves 98.3% success rate with 51.0% higher survival rate, 3.20% lower slip rate, and 46.5% lower stuck rate than DWL, demonstrating stable and efficient gait adaptation. Our work provides a robust and efficient gait control framework for stable locomotion in real-world scenarios with dynamic friction conditions.

have enabled impressive humanoid walking behaviors, including traversing stairs [Long *et al.*, 2025; Duan *et al.*, 2024], rough terrains [Cui *et al.*, 2024; Yu *et al.*, 2024], and sparse footholds [Radosavovic *et al.*, 2024b; Gu *et al.*, 2024; Wang *et al.*, 2025]. However, these methods predominantly assume that the ground provides sufficient friction for stable foot-ground contact, or handle friction variations implicitly through domain randomization [Tan *et al.*, 2018; Tobin *et al.*, 2017] to stabilize gait control. Therefore, perceiving ground features such as friction variations is crucial for the gait control of humanoid robots.

Existing locomotion methods can be broadly categorized by their treatment of friction. Methods relying solely on domain randomization do not explicitly address friction, instead focusing on other aspects such as taking adaptive gaits by terrain geometry [Duan *et al.*, 2024; Hoeller *et al.*, 2024; Loquercio *et al.*, 2023; Agarwal *et al.*, 2023; Miki *et al.*, 2022; He *et al.*, 2025]. Yet, without explicit awareness of surface friction, these policies typically employ conservative gaits, or fail catastrophically when encountering slippery surfaces. While some recent works attempt to estimate terrain properties from proprioceptive perceptions [Kumar *et al.*, 2021; Ji *et al.*, 2022; Gu *et al.*, 2024], or by leveraging visual information [Brandao *et al.*, 2016; Peng *et al.*, 2025]. Nevertheless, policies estimating friction from proprioception must first step onto the slippery surface before recognizing danger. Visual approaches also have limitations: CNN-based methods [Brandao *et al.*, 2016] lack generalization to novel surfaces and are sensitive to lighting variations, while using Vision-Language Models (VLMs) with Retrieval-Augmented Generation (RAG) introduces additional retrieval latency that limits real-time applicability [Arslan *et al.*, 2024].

In summary, enabling humanoid robots to walk robustly on surfaces with varying friction coefficients faces three key challenges: (1) *Anticipating surface friction before contact*. Unlike geometric features, friction is not directly observable from conventional visual representations and requires semantic understanding of material properties. (2) *Adapting gait in real-time under imperfect friction prediction*. Latency and uncertainty require the policy to handle prediction errors gracefully. (3) *Learning a unified policy that generalizes across a wide friction spectrum*. The optimal walking strategy differs between high and low-friction regimes.

To address these challenges, we propose Walking on Ice

1 Introduction

Humanoid robots hold great promise for operating in human-centric environments, from household assistance to industrial inspection [Tang *et al.*, 2025; Chen *et al.*, 2025; Kojima *et al.*, 2025]. A fundamental capability for such deployment is robust locomotion control across diverse scenarios. The context of advanced control technologies such as reinforcement learning [Li *et al.*, 2024a; Li *et al.*, 2024b]

*Corresponding authors.




Figure 1: Overview of WoI. (a) We build a terrain friction dataset containing 10 types of surfaces with friction ranging from high to low. (b) We fine-tune Qwen-2.5-VL-7B-Instruct by LoRA to predict friction coefficients from images. (c) The friction-aware locomotion policy combines VLM-based friction prediction with proprioceptive states for adaptive humanoid locomotion on varying surfaces.

(WoI), a friction-aware locomotion framework for humanoid robots. Our approach consists of three main components: (1) a VLM friction predictor which is fine-tuned by Low-Rank Adaption (LoRA)[Hu *et al.*, 2022] that anticipates upcoming terrain friction from egocentric images, (2) a friction estimator that provides real-time friction perception, and (3) a friction-aware locomotion policy that combines proactive visual prediction with reactive proprioceptive estimation. Our main contributions are summarized as follows:

- We propose WoI, the first humanoid locomotion framework that jointly leverages both current and anticipated friction information, enabling proactive gait adaptation on slippery surfaces.
- We design a friction-awareness mechanism that combines proactive VLM-based prediction with reactive proprioceptive estimation, allowing the robot to adapt gait before contact and in real-time, thereby avoiding the overly conservative gaits caused by domain randomization or catastrophic failures from delayed reaction.
- We introduce a VLM friction predictor, which leverages the world knowledge of VLMs to achieve robust generalization across novel surfaces and lighting variations.

2 Related Work

2.1 Handling Friction by Domain Randomization

Recent years have witnessed remarkable progress in legged locomotion using reinforcement learning. However, most existing methods focus on traversing diverse terrain geometries and achieving agile skills such as running, jumping, and parkour and overlook terrain friction properties, instead relying on domain randomization[Tobin *et al.*, 2017] to achieve robustness, both in humanoid robots [Radosavovic *et al.*, 2024b; Radosavovic *et al.*, 2024a; Singh *et al.*, 2024; Cui *et al.*, 2024; Yu *et al.*, 2024; Long *et al.*, 2025; Duan *et al.*, 2024; Bang *et al.*, 2024; Wang *et al.*, 2025; Siekmann *et al.*, 2021] and quadruped robots[Rudin *et al.*, 2022; Miki *et al.*, 2022; Tan *et al.*, 2018; Choi *et al.*, 2023; Kumar *et al.*, 2021; Hoeller *et al.*, 2024; Agarwal *et al.*, 2023; Jenelten *et al.*, 2024; He *et al.*, 2025].

While domain randomization partly enables policies to handle friction variations and enhance robustness, it leads to inherently conservative gaits that sacrifice efficiency for stability, as the policy cannot distinguish between high and low friction conditions and must adopt a one-size-fits-all strategy. [He *et al.*, 2025] achieves impressive performance on sparse foothold scenarios by introducing attention modules to process height maps, but focuses primarily on terrain geometry and assumes non-slippery surfaces. Similarly, [Radosavovic *et al.*, 2024b] demonstrates real-world humanoid locomotion with causal transformer architecture, yet the training and deployment environments implicitly assume sufficient ground friction. [Long *et al.*, 2025] proposes a perceptive internal model for humanoid locomotion over challenging terrain, but focusing on terrain elevation rather than friction properties.

2.2 Friction-Aware Locomotion

To address the limitations of domain randomization, some approaches estimate current friction from contact interaction[Lee *et al.*, 2020; Gu *et al.*, 2024; Ji *et al.*, 2022; Jenelten *et al.*, 2019], while other approaches tried vision-based methods to predict friction[Brandao *et al.*, 2016; Chen *et al.*, 2024; Peng *et al.*, 2025].

Proprioceptive-based methods estimate friction through force and motion analysis during ground contact. [Gu *et al.*, 2024] proposes a denoising world model that learns to predict terrain properties from proprioceptive history, enabling humanoid walking on snow and ice. [Ji *et al.*, 2022] concurrently trains a state estimator with the control policy to infer ground properties. However, these methods can only estimate friction *after* contact occurs, which may result in falls when encountering sudden friction transitions.

Vision-based methods attempt to predict friction before contact by using visual cues. [Brandao *et al.*, 2016] employs CNNs for material classification followed by friction database lookup, but this approach is constrained to predefined material categories and cannot generalize to novel surfaces. [Chen *et al.*, 2024] proposes a self-supervised framework to predict terrain physical parameters from vision, yet requires extensive real-world interaction data for training. Recently, [Peng *et al.*, 2025] explores VLMs with RAG for

friction estimation in wheeled robots, but the RAG-based approach introduces retrieval latency [Arslan *et al.*, 2024] that may limit real-time applicability in dynamic locomotion.

3 Method

As illustrated in Fig. 1, we present WoI, a friction-aware locomotion framework for humanoid robots. In contrast to approaches mentioned in Sec. 2, our key insight is to decompose friction awareness into two complementary components: anticipated visual prediction and current proprioceptive estimation, enabling proactive gait adaptation while maintaining real-time reactivity. (1) For proactive prediction, we fine-tune Qwen-2.5-VL-7B-Instruct by LoRA to directly regress friction coefficients from egocentric images. VLMs possess rich world knowledge about material properties, making our predictor inherently robust to lighting variations and capable of zero-shot generalization to novel surfaces without requiring extensive data collection across diverse environmental conditions. (2) For reactive estimation, we design a proprioception-based friction estimator within the policy network that infers current contact friction from joint states and forces, providing real-time correction when predictions are inaccurate. (3) We represent joint state history as a 2D “past-state-picture” and employ convolutional networks for temporal feature extraction, enabling the policy to capture motion patterns indicative of surface interactions. The locomotion policy takes the predicted and estimated friction as inputs along with proprioceptive observations. We first formulate the problem as a Markov Decision Process, then describe each component in detail.

3.1 Problem Formulation

We formulate friction-aware humanoid locomotion as a Partially Observable Markov Decision Process (POMDP) defined by the tuple $(\mathcal{S}, \mathcal{O}, \mathcal{A}, \mathcal{P}, r, \gamma)$, where $\mathcal{S}$ is the state space, $\mathcal{O}$ is the observation space, $\mathcal{A}$ is the action space, $\mathcal{P} : \mathcal{S} \times \mathcal{A} \rightarrow \mathcal{S}$ is the transition dynamics, $r : \mathcal{S} \times \mathcal{A} \rightarrow \mathbb{R}$ is the reward function, and $\gamma \in (0, 1)$ is the discount factor. The goal is to learn a policy $\pi_\theta(a_t|o_t)$ that maximizes the expected cumulative reward:

$$J(\theta) = \mathbb{E}_{\pi_\theta} \left[\sum_{t=0}^T \gamma^t r(s_t, a_t) \right] \quad (1)$$

We train an actor-critic policy using Proximal Policy Optimization (PPO) [Schulman *et al.*, 2017] with an asymmetric architecture [Pinto *et al.*, 2018], where the actor receives only observations available at deployment, while the critic has access to privileged information during training.

The actor observation at time t is defined as:

$$o_t = \left(\underbrace{(c_t, \omega_t, g_t, q_t, \dot{q}_t, a_{t-1}, \tau_t, \phi_t)}_{\text{Proprioception}}, \underbrace{H_t}_{\text{History}}, \underbrace{\mu_t^{\text{pred}}}_{\text{VLM Prediction}} \right) \quad (2)$$

where $c_t \in \mathbb{R}^3$ denotes velocity commands, $\omega_t \in \mathbb{R}^3$ is base angular velocity, $g_t \in \mathbb{R}^3$ is projected gravity, $q_t, \dot{q}_t \in \mathbb{R}^{N_j}$ are joint positions and velocities, a_{t-1} is the previous action, τ_t represents joint torques, $\phi_t = [\sin(\varphi_t), \cos(\varphi_t)]$ encodes gait phase, $H_t \in \mathbb{R}^{C \times N_j \times L}$ is the joint state history, and $\mu_t^{\text{pred}} \in [0, 1]$ is the VLM-predicted friction.




Figure 2: Policy network structure. Joint state history is processed as a 2D “past-state-picture” through a CNN. The friction estimator uses LSTM to infer current contact friction from proprioception. Both predicted and estimated friction are concatenated with other observations and fed to the action MLP.

The critic additionally observes privileged information:

$$o_t^{\text{critic}} = (o_t, v_t, \dot{p}_t^{\text{feet}}, F_t^{\text{contact}}, h_t^{\text{feet}}, \underbrace{\mu_t^{\text{contact}}}_{\text{Ground-truth Friction}}) \quad (3)$$

Privileged States

where v_t is base linear velocity, $\dot{p}_t^{\text{feet}}$ and F_t^{contact} are feet velocities and contact forces, h_t^{feet} is feet height, and μ_t^{contact} is the ground-truth friction at current contact. Notably, the actor sees predicted future friction μ_t^{pred} while the critic observes actual contact friction μ_t^{contact} , enabling the policy to learn anticipatory gait adjustment.

The action $a_t \in \mathbb{R}^{N_j}$ specifies target joint positions, controlled by a PD controller. Reward function combines velocity tracking, style regularization, and friction-adaptive rewards (details provided in supplementary material).

3.2 VLM-based Friction Predictor

We leverage the world knowledge embedded in Vision-Language Models, which have learned associations between visual appearances and physical properties (e.g., “rubber provides grip”) from large-scale image-text pretraining.

Dataset Construction. We construct a terrain friction dataset containing 33261 images of 10 surface types by splitting online videos and label with the corresponding friction coefficient, including dry and rainy street, wooden and ceramic floors, icy, grassy, sandy, snowy, unpaved and marshy surfaces (details see supplementary material).

Model Architecture. We adopt Qwen2.5-VL-7B-Instruct [Bai *et al.*, 2025] as our base VLM and fine-tune it using Low-Rank Adaptation (LoRA) [Hu *et al.*, 2022]. Given an egocentric RGB image I_t capturing the terrain ahead of the robot, the model predicts the friction coefficient μ_t^{pred} :

$$\mu_t^{\text{pred}} = f_{\text{VLM}}(I_t; \theta_{\text{base}}, \Delta\theta_{\text{LoRA}}) \quad (4)$$

where θ_{base} represents the frozen pretrained weights and $\Delta\theta_{\text{LoRA}}$ denotes the trainable low-rank adaptation matrices. LoRA decomposes the weight update as $\Delta W = BA$, where $B \in \mathbb{R}^{d \times r}$ and $A \in \mathbb{R}^{r \times k}$ with rank $r \ll \min(d, k)$, significantly reducing the number of trainable parameters while preserving the model’s pretrained knowledge.

3.3 Friction-Aware Locomotion Policy

As illustrated Fig. 2, the locomotion policy receives the friction coefficient as an external input, allowing the policy to adapt its gait based on anticipated terrain properties.

Joint History as Past-State-Picture. To capture temporal patterns indicative of terrain interactions and slips, we maintain a history of joint states over the past L timesteps. We reshape it into a 2D image-like tensor:

$$H_t = \text{reshape}([q_{t-L+1:t}, \dot{q}_{t-L+1:t}, \tau_{t-L+1:t}]) \in \mathbb{R}^{C \times N_j \times L} \quad (5)$$

where the channels $C = 3$ correspond to joint positions, velocities, and torques. This “past-state-picture” representation treats the degrees of freedom as the height dimension and time as the width dimension, enabling a 2D CNN to simultaneously capture temporal patterns hidden in history and cross-joint correlations through spatial convolutions.

Latent Feature Extraction. The past-state-picture is processed by a 2D CNN encoder to extract a compact latent representation:

$$z_t^H = f_{\text{CNN}}(H_t) \quad (6)$$

The CNN consists of 2 convolutional layers with kernel size $[3 \times 3]$, followed by adaptive average pooling and ELU activations. The latent vector z_t^H encodes motion patterns that reflect terrain properties through foot-ground interactions.

Friction Estimator. In addition to the proactive VLM-based prediction μ_t^{pred} , we incorporate a reactive friction estimator that infers the current contact friction from proprioceptive observations. The estimator takes proprioceptive inputs excluding velocity commands, as friction estimation should be independent of the commanded motion. The estimator uses an LSTM to capture temporal dynamics, followed by an MLP to regress the friction coefficient. During training, the estimator is supervised by the ground-truth contact friction μ_t^{contact} using MSE loss.

$$o_t^{\text{est}} = (\omega_t, g_t, q_t, \dot{q}_t, a_{t-1}, \tau_t^{\text{dof}}, \tau_t, \phi_t) \quad (7)$$

$$\mu_t^{\text{est}} = f_{\text{MLP}}(f_{\text{LSTM}}(o_{t-L:t}^{\text{est}})) \quad (8)$$

$$\mathcal{L}_{\text{est}} = (\mu_t^{\text{est}} - \mu_t^{\text{contact}})^2 \quad (9)$$

where o_t^{est} is the estimator input, μ_t^{est} is the estimated friction coefficient, and $\mathcal{L}_{\text{est}}$ is the supervision loss.

To prevent gradient interference between the estimator and the policy, we detach μ_t^{est} before feeding it to the action head.

4 Experiments

4.1 Experimental Setup

We conduct both simulation experiments and real-world deployment both with the 29-DoF Unitree G1 humanoid robot. Simulation is performed in Genesis [Authors, 2024] and locomotion policy training is completed on NVIDIA A100, with 40GB of memory. During deployment, the policy is zero-shot transferred on NVIDIA A5000.

Evaluation Metrics

We evaluate the performance using the following metrics, where f_i indicates whether the robot fell during trial i :

- **Survival Rate.** The percentage of trials where the robot does not fall, regardless of whether it reaches the target distance in time (20s):

$$E_{\text{surv}} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[\neg f_i] \quad (10)$$

- **Success Rate.** The percentage of trials where the robot traverses a target distance d_{target} in time, where N is the total number of trials and d_i is the distance traveled in trial i :

$$E_{\text{succ}} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[d_i \geq d_{\text{target}}] \quad (11)$$

- **Stuck Rate.** The percentage of trials where the robot fails to reach the target distance but does not fall:

$$E_{\text{stuck}} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}[d_i < d_{\text{target}} \wedge \neg f_i] \quad (12)$$

- **Slip Rate.** The ratio of slip events to total contact steps during locomotion, where $\mathcal{F}$ is the set of feet, $\mathbf{v}_{j,t}^{\text{contact}}$ is the tangential velocity of foot j at time t , v_{slip} is the slip velocity threshold, and $c_{j,t}$ indicates foot-ground contact:

$$E_{\text{slip}} = \frac{\sum_{t=1}^T \sum_{j \in \mathcal{F}} \mathbf{1}[\|\mathbf{v}_{j,t}^{\text{contact}}\|_2 > v_{\text{slip}} \wedge c_{j,t} = 1]}{\sum_{t=1}^T \sum_{j \in \mathcal{F}} \mathbf{1}[c_{j,t} = 1]} \quad (13)$$

Baselines

To assess the effectiveness of WoI, we compare WoI against (1) **PPE**[Chen *et al.*, 2024]; (2) **FSL**[Peng *et al.*, 2025]; (3) **DLSG**[Jenelten *et al.*, 2019]; (4) **MRCNN**[Brandao *et al.*, 2016]; (5) **LeggedGym** [Rudin *et al.*, 2022]; (6) **DWL** [Gu *et al.*, 2024].

Ablations

To validate the contribution of each component, we evaluate the following ablation variants: (1) **w/o VLM**: removing the VLM-based friction predictor μ_t^{pred} , relying solely on proprioceptive friction estimation; (2) **w/o Estimator**: removing the proprioceptive friction estimator μ_t^{est} , using only VLM-predicted friction; (3) **w/o Both**: removing both friction inputs, equivalent to a friction-unaware baseline; (4) **WoI_{1D CNN}**: replacing the 2D CNN history encoder with a 1D temporal convolution; (5) **WoI_{LSTM}**: replacing the 2D CNN history encoder with an LSTM network.

4.2 Simulation Results

Comparison with Baselines

Methods in Different Environments. Tab. 1 presents a comparison of methods in different environments. Unlike approaches that rely on proprioceptive estimation only (DLSG, DWL), vision-based prediction only (PPE, FSL, MRCNN), or domain randomization without explicit friction awareness (LeggedGym), WoI uniquely combines both reactive proprioceptive estimation and proactive VLM-based prediction,

Metrics	Algorithm						
	PPE [Chen <i>et al.</i> , 2024]	FSL [Peng <i>et al.</i> , 2025]	DLSG [Jenelten <i>et al.</i> , 2019]	MRCNN [Brandao <i>et al.</i> , 2016]	LeggedGym [Rudin <i>et al.</i> , 2022]	DWL [Gu <i>et al.</i> , 2024]	WoI (Ours)
Applications	Friction Prediction	Locomotion	Locomotion	Locomotion	Locomotion	Locomotion	Locomotion
Robot Type	Quadruped	Wheeled-Legged	Quadruped	Humanoid	Quadruped & Biped	Humanoid	Humanoid
$R_{friction}$	No DR	Not Provided	No DR	No DR	[0.50, 1.25]	[0.20, 2.00]	[0.10, 1.00]
$Awa_{friction}$	Predictive	Predictive	Reactive	Predictive	None	Reactive	Predictive & Reactive
$M_{friction}$	Vision (DINOv2-Based)	VLM+RAG	Proprioception	Vision (CNN-Based)	None	Proprioception	VLM+LoRA & Proprioception
Success Rate (%)	-	100.0	-	-	100.0	100.0	100.0
T_{trial} (s)	-	20.0	69.8	60.0	20.0	10.0	6.1

Table 1: Comparison of different methods. $R_{friction}$ stands for the domain randomization range of friction during training phase. No DR means that the method does not apply domain randomization in training phase. $Awa_{friction}$ stands for the awareness of friction. $M_{friction}$ stands for the perception method of friction. T_{trial} stands for the time used to finish a trial.

Method	$\mu = 0.1$				$\mu = 0.2$				$\mu = 0.3$				$\mu = 0.5$			
	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}
LeggedGym [Rudin <i>et al.</i> , 2022]	100.0	0.0	100.0	0.24	100.0	55.1	44.9	0.11	100.0	71.3	28.7	0.11	100.0	88.1	11.9	0.17
DWL [Gu <i>et al.</i> , 2024]	70.6	0.0	70.6	2.68	93.3	0.0	93.3	1.89	100.0	92.4	7.6	5.55	100.0	98.8	1.2	6.77
WoI (Ours)	100.0	100.0	0.0	1.44	100.0	100.0	0.0	0.74	100.0	100.0	0.0	0.64	100.0	100.0	0.0	1.28

Method	$\mu = 0.8$				$\mu = 1.0$				$1.0 \rightarrow 0.1$				$1.0 \rightarrow 0.3$			
	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}
LeggedGym [Rudin <i>et al.</i> , 2022]	100.0	89.4	10.6	0.19	100.0	75.9	24.1	0.20	100.0	76.1	23.9	0.20	100.0	76.7	23.3	0.20
DWL [Gu <i>et al.</i> , 2024]	100.0	99.4	0.6	6.87	100.0	98.3	1.7	6.72	47.3	0.8	46.5	4.64	100.0	91.5	8.5	5.78
WoI (Ours)	100.0	100.0	0.0	1.08	100.0	100.0	0.0	1.37	98.3	98.3	0.0	1.44	100.0	100.0	0.0	0.63

Table 2: Comparison with different methods on uniform friction and friction transition scenarios. The robot is tasked to walk 5 meters within 1000 simulation steps at 50Hz (20s) with a ahead command of 1.0m/s. For friction transition scenarios, the robot starts on a surface with $\mu = 1.0$ and encounters a friction change after traveling 0.5m. All results are averaged over 1000 episodes.




Figure 3: Foot Velocity Comparison on fixed friction scenarios. The robot is tasked to walk with a ahead command of 1.0m/s, and data is collected over 1000 episodes with per episode of 1000 steps. The X axis stands for step count, Y axis stands for the **tangential** velocity of contact foot, which suggest the slip level.

achieving both anticipatory gait adjustment and real-time correction. WoI covers a more extensive low-friction range (0.1) compared to LeggedGym (0.5) and DWL (0.2). Moreover, our method completes trials in 6.1 seconds on average, reducing trial time by approximately 90% compared to DLSG (69.8s) and MRCNN (60.0s).

Uniform Friction Surfaces. Tab. 2 presents that WoI achieves 100.0% survival rate and 100.0% success rate across all friction conditions. On extremely low friction ($\mu = 0.1$), WoI improves survival rate by 29.4% and success rate by 100% compared to DWL, while reducing slip rate by 1.24%, demonstrating strong effectiveness and robustness on low-friction surfaces. As shown in Fig. 3, WoI shows lower contact tangential velocity than DWL, indicating better slip control in extreme scenarios. At $\mu = 0.2$, DWL achieves 93.3% survival rate, but has a slip rate of 1.89%, compared to 0.74% for WoI, indicating that our method avoids overly conserva-

tive behavior while maintaining robustness. At moderate-low friction ($\mu = 0.3$), all methods achieve 100.0% survival rate, but WoI improves success rate by 7.6% over DWL’s 92.4% and reduces slip rate by 4.91% from DWL’s 5.55%. On moderate to high friction surfaces ($\mu = 0.5, 0.8, 1.0$), DWL achieves 100.0% survival rate, but WoI further improves success rate by 1.2%, 0.6%, and 1.7% respectively, while reducing slip rate by 5.47%, 5.79%, and 5.35%.

For LeggedGym, which relies solely on domain randomization without explicit friction awareness, it achieves 100% survival rate (see Fig. 4(a)) and the lowest slip rate across all surfaces. However, even on high friction surfaces ($\mu = 0.8, 1.0$), it exhibits 10.6% and 24.1% stuck rates respectively, while WoI has 0.0% stuck rate. On extremely low friction ($\mu = 0.1$), LeggedGym reaches 100.0% stuck rate, while WoI still has a lower max tangential velocity of contact foot and similar min tangential velocity of contact foot,




Figure 4: Comparison on E_{surv} , E_{succ} , E_{stuck} with different methods on uniform friction and friction transition scenarios. (a) Uniform friction scenarios. The X axis is the friction coefficient and the Y axis is the rate. (b) Friction transition scenarios. The X axis stands for the evaluation metrics and the Y axis stands for the rate.




Figure 5: Foot Velocity Comparison on friction transition scenarios. The robot is tasked to walk with a ahead command of 1.0m/s with a initial area of 0.5m long with friction coefficient of 1.0, and data is collected over 1000 episodes with per episode of 1000 steps. The X axis stands for step count, Y axis stands for the tangential velocity of contact foot, which suggest the slip level.

indicating better control of slipping and efficiency.

These results demonstrate that WoI effectively handles the full spectrum of friction conditions, while baseline methods either adopt overly conservative gaits (LeggedGym) or fail to adequately handle low-friction surfaces (DWL).

Friction Transition Scenarios. Tab. 2 presents results on friction transition scenarios, including high-to-moderate ($1.0 \rightarrow 0.3$) and high-to-extreme ($1.0 \rightarrow 0.1$) transitions. WoI achieves the highest success rate across all methods and scenarios. In the extreme transition ($1.0 \rightarrow 0.1$), although LeggedGym achieves 100.0% survival rate (See Fig. 4(b)), it exhibits a stuck rate of 23.9% and show very low tangential velocity of contact foot (See Fig. 5), indicating conservative gaits. DWL only achieves 47.3% survival rate, and obviously high tangential contact foot velocity, demonstrating that methods relying solely on reactive friction estimation have limited adaptability to friction transitions. In contrast, WoI improves success rate by 22.2% over LeggedGym and 97.5% over DWL, and shows low but suitable tangential velocity of contact foot, demonstrating strong robustness and traversabil-

ity under varying friction conditions. In the $1.0 \rightarrow 0.3$ transition, WoI achieves 100.0% survival rate and success rate, improving success rate by 23.3% over LeggedGym and 8.5% over DWL, while reducing slip rate by 5.15% compared to DWL, demonstrating that WoI maintains excellent robustness, efficiency, and gait stability when transitioning from high to low friction surfaces. These results demonstrate that LeggedGym survives through conservative gaits but fails to make forward progress, DWL falls due to the lack of proactive friction anticipation, while only WoI achieves both high survival and traversability.

Ablation Studies

Friction Awareness Components. Tab. 3(a) presents the ablation on friction awareness components. All methods achieve 100.0% survival rate and success rate on uniform friction surfaces ($\mu = 0.3$ and $\mu = 0.8$) and moderate friction transitions ($1.0 \rightarrow 0.3$), with w/o VLM and w/o Both achieving the lowest slip rates at $\mu = 0.3$ and $\mu = 0.8$. However, in extreme friction transitions ($1.0 \rightarrow 0.1$), WoI achieves the highest success rate of 98.3% and the lowest slip rate of 1.44%. Compared to w/o VLM, WoI improves success rate by 4.8%, demonstrating the effectiveness and necessity of proactive friction prediction for friction transition scenarios. Compared to w/o Estimator, WoI improves success rate by 0.6% and reduces slip rate by 0.43%, confirming that the proprioceptive estimator contributes to improved performance in friction transition scenarios. The w/o Both variant achieves the highest survival rate of 99.9%, but also exhibits the highest stuck rate of 4.0%, resulting in 2.4% lower success rate than WoI, indicating relatively conservative gaits due to the complete lack of friction awareness. In summary, the VLM predictor enables proactive adaptation to friction changes, the proprioceptive estimator provides real-time refinement, and their combination achieves optimal performance across all scenarios.

History Encoder Architecture. Tab. 3(b) compares different history encoder architectures. All methods achieve 100.0% survival rate and success rate on uniform friction surfaces ($\mu = 0.3$ and $\mu = 0.8$) and moderate friction transitions ($1.0 \rightarrow 0.3$). Since all variants retain the dual friction awareness mechanism, they exhibit 0.0% stuck rate across all scenarios. Additionally, WoI with LSTM encoder and WoI with

Method	$\mu = 0.3$				$\mu = 0.8$				$1.0 \rightarrow 0.1$				$1.0 \rightarrow 0.3$			
	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}	E_{surv}	E_{succ}	E_{stuck}	E_{slip}
<i>(a) Ablation on Friction Awareness</i>																
WoI(Ours)	100.0	100.0	0.0	0.64	100.0	100.0	0.0	1.08	98.3	98.3	0.0	1.44	100.0	100.0	0.0	0.63
w/o VLM	100.0	100.0	0.0	0.62	100.0	100.0	0.0	0.72	93.5	93.5	0.0	1.74	100.0	100.0	0.0	0.66
w/o Estimator	100.0	100.0	0.0	1.14	100.0	100.0	0.0	1.06	97.7	97.7	0.0	1.87	100.0	100.0	0.0	1.08
w/o Both	100.0	100.0	0.0	0.86	100.0	100.0	0.0	0.56	99.9	95.9	4.0	1.64	100.0	100.0	0.0	0.64
<i>(b) Ablation on History Encoder</i>																
WoI _{2D} CNN (Ours)	100.0	100.0	0.0	0.64	100.0	100.0	0.0	1.08	98.3	98.3	0.0	1.44	100.0	100.0	0.0	0.63
WoI _{1D} CNN	100.0	100.0	0.0	1.43	100.0	100.0	0.0	0.82	94.6	94.6	0.0	1.58	100.0	100.0	0.0	1.23
WoI _{LSTM}	100.0	100.0	0.0	0.24	100.0	100.0	0.0	1.52	95.4	95.4	0.0	2.45	100.0	100.0	0.0	0.21

Table 3: Ablation study on friction awareness components and history encoder architectures. Experimental setup follows Tab. 2. All results are averaged over 1000 episodes.




Figure 6: Real-world experiments. We conduct real-world experiments on (a) Wooden floor with the left half dry and the right half watered. ($\mu \approx 0.4\text{--}0.6 \rightarrow \mu \approx 0.2\text{--}0.3$). (b) Normal tile floor ($\mu \approx 0.3\text{--}0.5$). (c) Frozen outdoor surface covered with ice and snow ($\mu \approx 0.05\text{--}0.15$).

1D CNN encoder achieve the lowest slip rates at $\mu = 0.3$ and $\mu = 0.8$, with LSTM achieving the lowest slip rate in the $1.0 \rightarrow 0.3$ transition. However, in extreme friction transitions ($1.0 \rightarrow 0.1$), WoI with 2D CNN encoder achieves the highest survival rate and success rate of 98.3% and the lowest slip rate of 1.44%, improving success rate by 3.7% over 1D CNN and 2.9% over LSTM. These results demonstrate that treating joint history as a 2D “past-state-picture” better captures the motion dynamics and contact information embedded in cross-joint correlations. In summary, while all encoder architectures perform comparably on moderate friction conditions, the 2D CNN encoder provides superior performance in challenging friction transitions by exploiting spatial correlations across joints.

4.3 Real-World Experiments

We deploy WoI on a real Unitree G1 humanoid robot without any fine-tuning. The VLM runs at 1 Hz, and the control loop runs at 50 Hz, consistent with simulation.

Results. As Fig. 6 illustrated, WoI successfully performed on surfaces varying from indoor to outdoor and with friction transition, demonstrating adaptive control across various scenarios. In the indoor wooden floor test (Fig. 6(a)), the right half of wood is watered to decrease friction and to create friction transition situation. WoI successfully traverse both the dry area and the watered area, validating the ability to han-

dle varying friction surfaces in real-world. In normal indoor tile floor test (Fig. 6(b)), WoI walks with a natural, efficient gait. On icy surfaces, as Fig. 6(c) shown, robot encounters a slippage at initialization and recovers immediately and successfully navigates surfaces covered with snow and ice.

5 Conclusion

We presented WoI, a friction-aware locomotion framework that enables humanoid robots to walk robustly on surfaces with varying friction coefficients. Our approach combines three key components: a VLM-based friction predictor that leverages world knowledge to anticipate surface friction from egocentric images, a proprioceptive friction estimator that provides reactive correction during contact, and a 2D CNN-based history encoder that processes joint state history as a “past-state-picture” for effective temporal feature extraction. Experiments in simulation and real-world deployment demonstrate that WoI achieves higher success rate both on fixed low-friction surfaces and friction transition scenarios compared to different methods, while maintaining robust performance and effectiveness across diverse terrain conditions. We will explore policy work both in environments with changes both in physical properties and geometric shape in the future.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (Grant No. 62350710797).

References

- [Agarwal *et al.*, 2023] Ananye Agarwal, Ashish Kumar, Jitendra Malik, and Deepak Pathak. Legged locomotion in challenging terrains using egocentric vision. In *Conference on robot learning*, pages 403–415. PMLR, 2023.
- [Arslan *et al.*, 2024] Muhammad Arslan, Hussam Ghanem, Saba Munawar, and Christophe Cruz. A survey on rag with llms. *Procedia computer science*, 246:3781–3790, 2024.
- [Authors, 2024] Genesis Authors. Genesis: A generative and universal physics engine for robotics and beyond, December 2024.
- [Bai *et al.*, 2025] Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- [Bang *et al.*, 2024] Seung Hyeon Bang, Carlos Arribalzaga Jové, and Luis Sentis. RL-augmented mpc framework for agile and robust bipedal footstep locomotion planning and control. In *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*, pages 607–614. IEEE, 2024.
- [Brandao *et al.*, 2016] Martim Brandao, Yukitoshi Minami Shiguematsu, Kenji Hashimoto, and Atsuo Takanishi. Material recognition cnns and hierarchical planning for biped robot locomotion on slippery terrain. In *2016 IEEE-RAS 16th International Conference on Humanoid Robots (Humanoids)*, pages 81–88. IEEE, 2016.
- [Chen *et al.*, 2024] Jiaqi Chen, Jonas Frey, Ruyi Zhou, Takahiro Miki, Georg Martius, and Marco Hutter. Identifying terrain physical parameters from vision-towards physical-parameter-aware locomotion and navigation. *IEEE Robotics and Automation Letters*, 2024.
- [Chen *et al.*, 2025] Xinglin Chen, Yishuai Cai, Minglong Li, Yunxin Mao, Zhou Yang, Wenjing Yang, Weixia Xu, and Ji Wang. Btpp: A platform and benchmark for behavior tree planning in everyday service robots. In James Kwok, editor, *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 8715–8723. International Joint Conferences on Artificial Intelligence Organization, 8 2025. Main Track.
- [Choi *et al.*, 2023] Suyoung Choi, Gwanghyeon Ji, Jeongsoo Park, Hyeongjun Kim, Juhyeok Mun, Jeong Hyun Lee, and Jemin Hwangbo. Learning quadrupedal locomotion on deformable terrain. *Science Robotics*, 8(74):eade2256, 2023.
- [Cui *et al.*, 2024] Wenhao Cui, Shengtao Li, Huaxing Huang, Bangyu Qin, Tianchu Zhang, Liang Zheng, Ziyang Tang, Chenxu Hu, NING Yan, Jiahao Chen, et al. Adapting humanoid locomotion over challenging terrain via two-phase training. In *8th Annual Conference on Robot Learning*, 2024.
- [Duan *et al.*, 2024] Helei Duan, Bikram Pandit, Mohitvishnu S Gadde, Bart Van Marum, Jeremy Dao, Chanh Kim, and Alan Fern. Learning vision-based bipedal locomotion for challenging terrain. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 56–62. IEEE, 2024.
- [Gu *et al.*, 2024] 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. *arXiv preprint arXiv:2408.14472*, 2024.
- [He *et al.*, 2025] Junzhe He, Chong Zhang, Fabian Jenelten, Ruben Grandia, Moritz Bächer, and Marco Hutter. Attention-based map encoding for learning generalized legged locomotion. *Science Robotics*, 10(105):eadv3604, 2025.
- [Hoeller *et al.*, 2024] David Hoeller, Nikita Rudin, Dhionis Sako, and Marco Hutter. Anymal parkour: Learning agile navigation for quadrupedal robots. *Science Robotics*, 9(88):ead7566, 2024.
- [Hu *et al.*, 2022] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2):3, 2022.
- [Jenelten *et al.*, 2019] Fabian Jenelten, Jemin Hwangbo, Fabian Tresoldi, C Dario Bellicoso, and Marco Hutter. Dynamic locomotion on slippery ground. *IEEE Robotics and Automation Letters*, 4(4):4170–4176, 2019.
- [Jenelten *et al.*, 2024] Fabian Jenelten, Junzhe He, Farbod Farshidian, and Marco Hutter. Dtc: Deep tracking control. *Science Robotics*, 9(86):eadh5401, 2024.
- [Ji *et al.*, 2022] Gwanghyeon Ji, Juhyeok Mun, Hyeongjun Kim, and Jemin Hwangbo. Concurrent training of a control policy and a state estimator for dynamic and robust legged locomotion. *IEEE Robotics and Automation Letters*, 7(2):4630–4637, 2022.
- [Kojima *et al.*, 2025] Takeshi Kojima, Yaonan Zhu, Yusuke Iwasawa, Toshinori Kitamura, Gang Yan, Shu Morikuni, Ryosuke Takanami, Alfredo Solano, Tatsuya Matsushima, Akiko Murakami, and Yutaka Matsuo. A comprehensive survey on physical risk control in the era of foundation model-enabled robotics. In James Kwok, editor, *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 10517–10527. International Joint Conferences on Artificial Intelligence Organization, 8 2025. Survey Track.
- [Kumar *et al.*, 2021] Ashish Kumar, Zipeng Fu, Deepak Pathak, and Jitendra Malik. Rma: Rapid motor adaptation for legged robots. *Robotics: Science and Systems XVII*, 2021.
- [Lee *et al.*, 2020] Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning quadrupedal locomotion over challenging terrain. *Science robotics*, 5(47):eabc5986, 2020.

- [Li *et al.*, 2024a] Jinlong Li, Ruonan Li, Lunhui Xu, and Jie Liu. Self-supervised generative adversarial learning with conditional cyclical constraints towards missing traffic data imputation. *Knowledge-Based Systems*, 284:111233, 2024.
- [Li *et al.*, 2024b] Ruonan Li, Yang Qin, Jie Liu, Lu Zang, and Jinlong Li. Path planning strategy based on principal component federation for multi-agent in connected vehicles. *IEEE Transactions on Automation Science and Engineering*, 2024.
- [Long *et al.*, 2025] 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.
- [Loquercio *et al.*, 2023] Antonio Loquercio, Ashish Kumar, and Jitendra Malik. Learning visual locomotion with cross-modal supervision. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 7295–7302. IEEE, 2023.
- [Miki *et al.*, 2022] Takahiro Miki, Joonho Lee, Jemin Hwangbo, Lorenz Wellhausen, Vladlen Koltun, and Marco Hutter. Learning robust perceptive locomotion for quadrupedal robots in the wild. *Science robotics*, 7(62):eabk2822, 2022.
- [Peng *et al.*, 2025] Bo Peng, Donghoon Baek, Qijie Wang, and Joao Ramos. Friction-aware safety locomotion for wheeled-legged robots using vision language models and reinforcement learning. In *2025 IEEE-RAS 24th International Conference on Humanoid Robots (Humanoids)*, pages 1–1. IEEE, 2025.
- [Pinto *et al.*, 2018] Lerrel Pinto, Marcin Andrychowicz, Peter Welinder, Wojciech Zaremba, and Pieter Abbeel. Asymmetric actor critic for image-based robot learning. *Robotics: Science and Systems XIV*, 2018.
- [Radosavovic *et al.*, 2024a] Ilija Radosavovic, Sarthak Kamat, Trevor Darrell, and Jitendra Malik. Learning humanoid locomotion over challenging terrain. *arXiv preprint arXiv:2410.03654*, 2024.
- [Radosavovic *et al.*, 2024b] 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.
- [Rudin *et al.*, 2022] Nikita Rudin, David Hoeller, Philipp Reist, and Marco Hutter. Learning to walk in minutes using massively parallel deep reinforcement learning. In *Conference on robot learning*, pages 91–100. PMLR, 2022.
- [Schulman *et al.*, 2017] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [Siekman *et al.*, 2021] Jonah Siekman, Kevin Green, John Warila, Alan Fern, and Jonathan Hurst. Blind bipedal stair traversal via sim-to-real reinforcement learning. *Robotics: Science and Systems XVII*, 2021.
- [Singh *et al.*, 2024] Rohan P Singh, Mitsuharu Morisawa, Mehdi Benallegue, Zhaoming Xie, and Fumio Kanehiro. Robust humanoid walking on compliant and uneven terrain with deep reinforcement learning. In *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*, pages 497–504. IEEE, 2024.
- [Tan *et al.*, 2018] Jie Tan, Tingnan Zhang, Erwin Coumans, Atil Iscen, Yunfei Bai, Danijar Hafner, Steven Bohez, and Vincent Vanhoucke. Sim-to-real: Learning agile locomotion for quadruped robots. *Robotics: Science and Systems XIV*, 2018.
- [Tang *et al.*, 2025] Shixiang Tang, Yizhou Wang, Lu Chen, Yuan Wang, Sida Peng, Dan Xu, and Wanli Ouyang. Human-centric foundation models: Perception, generation and agentic modeling. In James Kwok, editor, *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence, IJCAI-25*, pages 10678–10686. International Joint Conferences on Artificial Intelligence Organization, 8 2025. Survey Track.
- [Tobin *et al.*, 2017] Josh Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and Pieter Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. In *2017 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 23–30. IEEE, 2017.
- [Wang *et al.*, 2025] Huayi Wang, Zirui Wang, Junli Ren, Qingwei Ben, Tao Huang, Weinan Zhang, and Jiangmiao Pang. Beamdojo: Learning agile humanoid locomotion on sparse footholds. In *Robotics: Science and Systems (RSS)*, 2025.
- [Yu *et al.*, 2024] Shangqun Yu, Nisal Perera, Daniel Marew, and Donghyun Kim. Learning generic and dynamic locomotion of humanoids across discrete terrains. In *2024 IEEE-RAS 23rd International Conference on Humanoid Robots (Humanoids)*, pages 1048–1055. IEEE, 2024.