Soft Contact Model for Robust Locomotion of Legged Robots

Yong-Hoon Lee[1], Keuntae Kim[2], Jaehyun Park[1], Chung Hyuk Park[2], and Hae-Won Park[1†]

Abstract— **We present a soft contact model to simulate diverse soft terrains, enabling robust legged locomotion through reinforcement learning. The model extends a standard spring-damper formulation with Stribeck-Coulomb friction and introduces randomized parameters, such as stiffness, damping, and friction coefficients, to capture the variability of real-world soft surfaces, including soil and mattresses. By replacing the default contact model in the simulator with our formulation, we train a locomotion policy using an existing learning framework. The resulting policy demonstrates stable walking on both flat and inclined soft terrains with the *Unitree Go1* robot in simulation. Notably, it generalizes to rigid ground without explicit training, highlighting improved robustness across contact conditions. This work offers a lightweight and flexible alternative to high-fidelity contact modeling for scalable locomotion training.**

I. INTRODUCTION

Legged robots are increasingly envisioned to operate in human-centric environments such as homes, where they must traverse not only rigid floors but also soft and deformable surfaces, including couches and cushions. However, most reinforcement learning-based locomotion controllers to date have been developed and tested on hard, rigid ground. As a result, policies trained on standard rigid terrains often suffer degraded performance when faced with soft or deformable substrates [1]. Bridging this gap is crucial, since a substantial portion of real-world surfaces—from sand and soil outdoors to carpeting indoors—are non-rigid and yield underfoot, causing foot sinking and altered frictional forces that can destabilize controllers tuned for rigid ground [2]. Robust locomotion over soft terrain is thus an essential capability for legged robots in many applications, yet it remains an open challenge.

The difficulty stems from the complex and unpredictable dynamics of deformable contact. Soft ground can significantly deform, dissipate energy, and slip, in ways difficult to capture with conventional rigid-body models. Terrain properties such as compliance, damping, and friction may vary widely and can change with environmental conditions. Even advanced perception systems cannot reliably infer these hidden properties of the terrain from vision alone [3], meaning a robot often will not know how soft a surface is until it steps on it. While high-fidelity physics models for soft materials exist, they are computationally intensive. Accurate simulation of a robot walking on soft surfaces might require modeling millions or even billions of material particles, which is impractical for fast simulation [4]. This makes brute-force simulation of deformable terrain infeasible in standard reinforcement learning training loops. As a result, prior research has often resorted to simplifying assumptions or avoided soft ground altogether, yielding controllers that work well on rigid terrain but generalize poorly to deformable surfaces.

To address the challenges of locomotion on soft and varying terrains, two primary strategies have been explored: model-based and learning-based approaches. Model-based methods incorporate explicit knowledge of terrain properties or estimate ground reaction forces using physical models [5]–[7]. While effective in controlled settings, these methods often require prior knowledge and manual tuning, limiting their applicability in unstructured environments.

Learning-based methods, particularly reinforcement learning, have shown promise in training robust locomotion policies in simulation [8]–[13]. Notably, Lee et al. demonstrated that a reinforcement learning policy trained only on rough rigid terrain was able to zero-shot generalize to deformable terrains like mud and snow, retaining robustness even without direct exposure during training [9]. A key enabler of such generalization is domain randomization, wherein physical parameters, such as friction coefficients and ground perturbations, are randomized during training to expose the policy to a wide range of dynamics [14]. However, its effectiveness diminishes when real-world conditions fall outside the randomized training domain, especially on highly deformable surfaces with qualitatively different contact dynamics. To bridge this gap, recent work has augmented simulators with deformable terrain models. For instance, Choi et al. [3] introduced a granular media simulation adjustable from soft sand to firm ground, coupled with an adaptive policy using foot sensor feedback. This approach enabled a quadruped robot to run at 3.03 m/s on sand and trot across an air mattress.

Building on these insights, this article proposes a lightweight simulation to achieve robust quadruped locomotion over diverse soft terrains. Instead of relying on expensive high-dimensional models of soil or foam, we extend a standard spring-damper contact model with Stribeck-Coulomb friction to approximate the behavior of compliant, dissipative surfaces. This soft contact model is implemented in the RaiSim simulator [15], which supports fast rigid-body simulation. By adjusting a small number of contact parameters, such as the contact stiffness, damping, or friction coefficients, we can emulate a spectrum of ground hardness

[1]Yong-Hoon Lee, Jaehyun Park, and Hae-Won Park are with the Department of Mechanical Engineering, Korea Advanced Institute of Science and Technology, Yuseong-gu, Daejeon 34141, Republic of Korea. haewonpark@kaist.ac.kr

[2]Keuntae Kim and Chung Hyuk Park are with the Department of Biomedical Engineering, The George Washington University, Washington, DC 20052, USA.

[†]Corresponding author

from soft, asking terrain to effectively rigid flooring. We leverage randomization over these contact parameters during reinforcement learning training to expose the policy to many virtual terrains. Through this approach, the robot learns to dynamically adjust its locomotion strategy as if feeling out each step, without any explicit terrain estimator.

In our experiments, a quadruped robot trained with our proposed method walks stably across soft terrains in RaiSim, while retaining high performance on rigid ground. These contributions take a step toward natural, reliable legged locomotion in human environments, enabling quadruped robots to traverse the same soft terrains humans do.

II. METHODS

A. Soft Contact Model

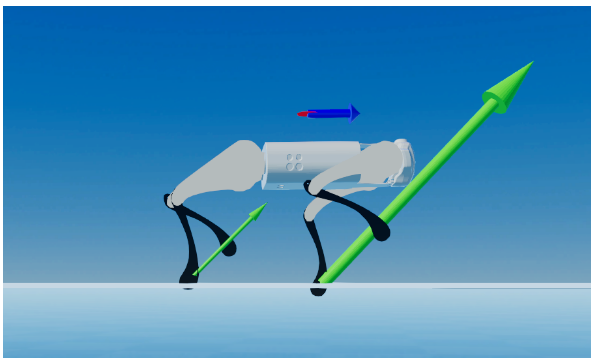

Fig. 1. Visualization of the contact forces generated by the proposed soft contact model. Green arrows indicate the computed contact forces applied to each foot, while the blue arrow represents the commanded body velocity.

To simulate diverse soft terrains with varying mechanical properties, we developed a soft contact model that captures the behavior of deformable surfaces such as foam and mattresses. The model is designed to support reinforcement learning in simulation by providing physically consistent and realistic contact interactions. We compute the contact force for each foot based on its penetration into the terrain along the surface normal, with contact forces comprising two main components: a nonlinear spring-damper model in the normal direction and a Stribeck-Coulomb friction model in the tangential direction. An overview of the resulting contact forces during locomotion is illustrated in Fig. 1, where force vectors vary in direction and magnitude depending on the phase of foot contact and terrain response. Unlike the previous model tailored for granular media [3], which assumes no ground reaction force during the upward movement of the foot due to the yielding behavior of loosely packed particles, our model continuously applies contact forces throughout all phases of penetration. This reflects the physical characteristics of compliant continuous surfaces, which generate reaction forces even as the foot moves upward, allowing the policy to learn behaviors consistent with real-world soft terrain.

The normal contact force is defined as:

$$\mathbf{F}_n = \left(k_{\text{soft}}(d)\, d - c_{\text{soft}}(v_n)\, v_n\right)\mathbf{n} \tag{1}$$

where d is the penetration depth, $\mathbf{n}$ is the unit normal vector of the contact surface, $v_n = \|\mathbf{v}_n\|$ is the magnitude of the velocity along the normal direction, k_{soft} is the depth-dependent stiffness, and c_{soft} is a velocity-dependent damping coefficient. The stiffness increases nonlinearly with penetration depth, mimicking the progressive resistance of soft materials:

$$k_{\text{soft}}(d) = k_{\min} + (k_{\max} - k_{\min})\left(\frac{\min(d, l_t)}{l_t}\right)^{1.2} \tag{2}$$

Here, $k_{\min}$ and $k_{\max}$ are sampled per environment, with $k_{\max}$ ranging from 1.5 to 2.5 times $k_{\min}$, to reflect the variability of soft terrains. l_t is the transition length that limits the growth of stiffness at large penetrations. This formulation is inspired by nonlinear contact mechanics models in which contact stiffness increases with deformation [16], [17], and has been effectively applied to robotic systems [18], [19].

To reflect asymmetric damping behavior during foot contact, c_{soft} is defined as:

$$c_{\text{soft}}(v_n) = \begin{cases} c_{\text{down}}, & \text{if } v_n < 0 \text{ (descending)} \\ c_{\text{up}}, & \text{if } v_n \geq 0 \text{ (ascending)} \end{cases} \tag{3}$$

This asymmetric damping models the observation that compressive motion into soft terrain (e.g., foot landing) typically induces greater damping due to material compaction, internal friction, and energy dissipation, whereas the recovery phase (e.g., foot lifting) involves less resistance as the material passively returns to its original shape. Accordingly, c_{down} is set to be larger than c_{up}, with c_{up} randomized between 0.2 and 0.8 times c_{down} across environments.

Tangential force is modeled using a combination of Stribeck and Coulomb friction, which together capture the transition from static to kinetic friction. The magnitude of the tangential friction force is determined as:

$$\begin{aligned} \|\mathbf{F}_{\text{fric}}\| = {} & \mu_k \|\mathbf{F}_{\text{n}}\| \tanh\left(\tfrac{v_t}{v_{\text{Coul}}}\right) \\ & + \sqrt{2\mathrm{e}}(\mu_s - \mu_k)\|\mathbf{F}_{\text{n}}\| \exp\left(-\left(\tfrac{v_t}{v_{\text{St}}}\right)^2\right)\tfrac{v_t}{v_{\text{St}}} \end{aligned} \tag{4}$$

where μ_k and μ_s are the kinetic and static friction coefficients, $v_t = \|\mathbf{v}_t\|$ is the magnitude of the tangential velocity, and v_{Coul} and v_{St} are velocity thresholds for the Coulomb and Stribeck transitions, respectively.

The resulting tangential friction force vector is applied in the direction opposite to the tangential velocity:

$$\mathbf{F}_{\text{fric}} = -\|\mathbf{F}_{\text{fric}}\|\frac{\mathbf{v}_t}{\|\mathbf{v}_t\|} \tag{5}$$

By stochastically sampling terrain stiffness, damping, and friction parameters, this model enables training in a wide range of soft contact conditions, promoting the emergence of robust locomotion policies.

B. Learning based Controller

Building on the soft terrain simulation environment described in Section II-A, we train a learning-based controller that enables the Unitree Go1 robot to achieve robust and adaptive locomotion on soft surfaces. The soft contact model, characterized by nonlinear stiffness, asymmetric damping, and friction, provides the physical substrate on which the controller is developed. Rather than using the simulator's default contact model, we apply the contact forces computed from our soft contact model directly as external forces to the robot's foot links, ensuring that the learned policy adapts to the dynamics of the soft terrain.

For policy training, we adopt the learning framework proposed by Kim et al. [20], which leverages a relaxed logarithmic barrier reward to softly enforce motion styles such as foot clearance, body height, joint posture, and preferred gait. The barrier reward is shaped by user-defined lower and upper bounds, along with δ values that control gradient steepness in the constraint-violation region. These terms are critical in guiding the policy toward desirable behaviors without imposing hard restrictions that might hinder learning in soft and variable terrain.

TABLE I
BARRIER REWARD FUNCTION.

Constraint Variable		d^{lower}	d^{upper}	δ
Gait	f_i	-0.6	-[†]	0.05
Foot clearance	$l_i\,[m]$	-0.05	-[†]	0.01
Joint position	$q_{\text{roll},i} - q^{\text{nom}}_{\text{roll},i}\,[rad]$	$-\pi/6$	$\pi/6$	0.02
	$q_{\text{hip},i} - q^{\text{nom}}_{\text{hip},i}\,[rad]$	$-\pi/4$	$\pi/4$	
	$q_{\text{knee},i} - q^{\text{nom}}_{\text{knee},i}\,[rad]$	$-2\pi/5$	$\pi/8$	
Body height	${}_bh_F\,[m]$	0.35	0.43	0.04
	${}_bh_H\,[m]$			
Target velocity	$v^{\text{cmd}}_x - v_x\,[m/s]$	-0.4	0.4	0.2
	$v^{\text{cmd}}_y - v_y\,[m/s]$			
	$\omega^{\text{cmd}}_z - \omega_z\,[rad/s]$			
Base motion	$\omega_x\,[rad/s]$	-0.3	0.3	0.3
	$\omega_y\,[rad/s]$			
	$v_z\,[m/s]$			0.2
Joint velocity	$\dot{q_j}\,[rad/s]$	-8	8	2.0

$(\cdot)^{nom}$ and $(\cdot)^{cmd}$ denote nominal and commanded values, respectively. $(\cdot)_i$ refers to the i-th leg, and $(\cdot)^j$ to the j-th joint. x, y, and z are defined in the body frame. f, l, q, $\dot{q}$, v, and ω represent gait constraint variable, foot clearance, joint position, joint velocity, body linear velocity, and body rotational velocity, respectively. ${}_bh_F$ and ${}_bh_H$ denote the heights of the front and hind roll joints from the ground, respectively, in the body frame. [†]For f_i and l_i, the upper bounds are unnecessary, thus set to non-reachable values (2.0 and 1.0, respectively).

The original framework, implemented on the KAIST HOUND platform [21], is adapted here for the Unitree Go1. To accommodate the differences in robot morphology and to enable stable walking on soft terrains, we modify the constraint boundaries and δ values of the barrier reward. Table I summarizes the constraint variables, their operational ranges, and δ values. By integrating soft contact model into the learning loop and tailoring the reward function accordingly, we achieve a locomotion controller that demonstrates robust performance on soft terrains.

C. *Training in Simulation*

The policy is trained entirely in simulation with the Unitree Go1 quadruped robot. Given that our ultimate goal is to deploy the policy on real hardware, we explicitly consider the Go1 robot's control latency. In real-world experiments, we observed a 10–20 ms latency between action computation and execution. To account for this in simulation, we inject a randomized delay, sampled uniformly from this range, before applying actions. This latency modeling ensures that the learned policy is robust to timing discrepancies that may arise during deployment.

In addition, to mitigate the dynamics mismatch introduced by such delays, we augment the observation with short-term histories of critical features. Specifically, we include 10 ms, 20 ms, 30 ms, and 40 ms past values of joint position errors, joint velocities, and relative foot positions in the body frame. These are added on top of the original long-term history inputs at 60 ms, 120 ms, and 180 ms used in the baseline learning framework [20].

Terrain variation is introduced via a randomized slope curriculum, where the robot is trained on inclines sampled uniformly from 0 % to 51 %. Our soft contact model is compatible with these inclined surfaces, as the surface normal and penetration computations are adapted to local slope angles. To facilitate early learning, particularly during gait discovery, we begin training with the simulator's default hard contact model. From 4000 to 7000 training iterations, we gradually increase the ratio of soft contact environments, smoothly transitioning from hard to soft terrain. This progressive switch allows the policy to stably adapt to the compliant dynamics of soft terrain while retaining the locomotion skills acquired during early training.

By integrating latency modeling, observation enhancements, and a soft-terrain-aware curriculum, our training process yields a locomotion policy that is not only robust in simulation but also designed with real-world deployment in mind.

III. RESULTS

Training and Deployment Setup We utilized the RaiSim simulator for training [15]. The training was conducted on an AMD Ryzen Threadripper PRO 5995WX and a single NVIDIA GeForce RTX 3080 Ti for 20 hours, over 20,000 iterations. During the learning process, 400 environments collected data at 100 Hz in 4-second episodes, resulting in a batch size of 160,000. To ensure stability in contact dynamics with the soft terrain model, the simulation time step was set to 5000 Hz. The control policy was deployed on the quadruped robot Unitree Go1.

A. *Performance on Variable Soft Terrains*

To quantitatively assess the robustness of the learned locomotion policy, we evaluate its performance across a range of soft terrain configurations characterized by different stiffness (k_{min}) and damping (c_{down}) parameters. Fig. 2 shows a 3D surface plot of the velocity tracking error norm over the evaluated parameter space.

The results indicate that the learned policy achieves consistently low velocity tracking errors across a wide range of soft terrain conditions. This demonstrates the policy's ability to adapt to varying ground properties without explicit knowledge of the terrain type. Performance degradation is primarily observed when stiffness is low, but damping is excessively high. In such cases, the foot-ground interaction becomes

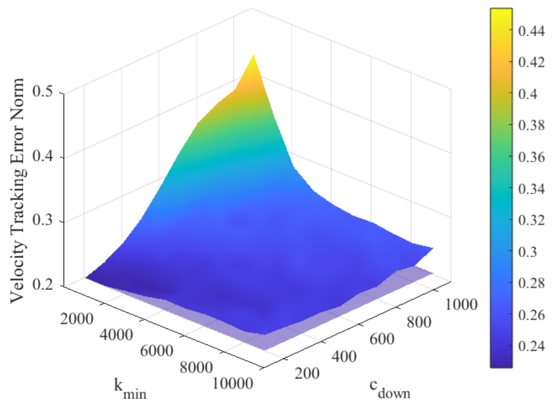

Fig. 2. The 3D surface represents the average tracking error as a function of terrain stiffness (k_{min}) and damping (c_{down}). The errors are computed from the norm of v_x, v_y, and ω_z over a 5-second period across 1000 environments where command velocities are uniformly sampled within the ranges of ± 1 m/s for v_x and v_y, and ± 1 rad/s for ω_z. A horizontal semi-transparent plane indicates the error norm (0.2259) achieved on rigid terrain, serving as a reference for comparison. The proximity of the surface to this baseline across a wide range of terrain parameters highlights the policy's robustness and generalization capability under diverse soft contact conditions.

either unstable or overdamped, resulting in increased tracking error. Nevertheless, even under these challenging configurations, the error remains bounded, indicating a smooth degradation rather than abrupt failure. These observations support the conclusion that the proposed training scheme yields a robust and generalizable locomotion strategy.

IV. CONCLUSION

We proposed a soft contact model for robust quadruped locomotion over diverse soft terrains by training a policy with randomized terrain parameters. The model captures key characteristics of deformable surfaces through nonlinear stiffness, asymmetric damping, and friction, enabling efficient and physically realistic simulation. The learned policy achieved stable and accurate velocity tracking, approaching that of rigid-ground locomotion across a wide range of terrain conditions, with tracking errors remaining bounded even in challenging configurations. Future work will focus on applying the method to real-world deployment and extending its applicability to structured/unstructured soft environments such as compliant blocks and articulated surfaces.

ACKNOWLEDGMENT

This work was partially funded by the NSF Grant # 2229885 Institute for Trustworthy AI in Law and Society (TRAILS) as a seed project entitled "Trustworthy Human-Animal Robot Interaction (THARI) for Autistic Children" and partially supported by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MSIT) (No. 2022R1A2C2011927).

REFERENCES

[1] C. Li, P. B. Umbanhowar, H. Komsuoglu, D. E. Koditschek, and D. I. Goldman, "Sensitive dependence of the motion of a legged robot on granular media," *Proceedings of the National Academy of Sciences*, vol. 106, no. 9, pp. 3029–3034, 2009.

[2] H. Kolvenbach, P. Arm, E. Hampp, A. Dietsche, V. Bickel, B. Sun, C. Meyer, and M. Hutter, "Traversing steep and granular martian analog slopes with a dynamic quadrupedal robot," *Field Robotics*, vol. 2, pp. 910–939, 2022.

[3] S. Choi, G. Ji, J. Park, H. Kim, J. Mun, J. H. Lee, and J. Hwangbo, "Learning quadrupedal locomotion on deformable terrain," *Science Robotics*, vol. 8, no. 74, p. eade2256, 2023.

[4] C. Lin, J. Fan, Y. Wang, Z. Yang, Z. Chen, L. Fang, T.-H. Wang, Z. Xian, and C. Gan, "Ubsoft: A simulation platform for robotic skill learning in unbounded soft environments," 2024.

[5] V. Vasilopoulos, I. S. Paraskevas, and E. G. Papadopoulos, "Compliant terrain legged locomotion using a viscoplastic approach," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 4849–4854.

[6] C. M. Hubicki, J. J. Aguilar, D. I. Goldman, and A. D. Ames, "Tractable terrain-aware motion planning on granular media: An impulsive jumping study," in *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2016, pp. 3887–3892.

[7] D. J. Lynch, K. M. Lynch, and P. B. Umbanhowar, "The soft-landing problem: Minimizing energy loss by a legged robot impacting yielding terrain," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 3658–3665, 2020.

[8] J. Hwangbo, J. Lee, A. Dosovitskiy, D. Bellicoso, V. Tsounis, V. Koltun, and M. Hutter, "Learning agile and dynamic motor skills for legged robots," *Science Robotics*, vol. 4, no. 26, p. eaau5872, 2019.

[9] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, "Learning quadrupedal locomotion over challenging terrain," *Science robotics*, vol. 5, no. 47, p. eabc5986, 2020.

[10] W. Yu, J. Tan, Y. Bai, E. Coumans, and S. Ha, "Learning fast adaptation with meta strategy optimization," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 2950–2957, 2020.

[11] G. Ji, J. Mun, H. Kim, and J. Hwangbo, "Concurrent training of a control policy and a state estimator for dynamic and robust legged locomotion," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 4630–4637, 2022.

[12] Z. Zhuang, Z. Fu, J. Wang, C. Atkeson, S. Schwertfeger, C. Finn, and H. Zhao, "Robot parkour learning," *arXiv preprint arXiv:2309.05665*, 2023.

[13] D. Hoeller, N. Rudin, D. Sako, and M. Hutter, "Anymal parkour: Learning agile navigation for quadrupedal robots," *Science Robotics*, vol. 9, no. 88, p. eadi7566, 2024.

[14] X. B. Peng, M. Andrychowicz, W. Zaremba, and P. Abbeel, "Sim-to-real transfer of robotic control with dynamics randomization," in *2018 IEEE international conference on robotics and automation (ICRA)*. IEEE, 2018, pp. 3803–3810.

[15] J. Hwangbo, J. Lee, and M. Hutter, "Per-contact iteration method for solving contact dynamics," *IEEE Robotics and Automation Letters*, vol. 3, no. 2, pp. 895–902, 2018.

[16] K. H. Hunt and F. R. E. Crossley, "Coefficient of restitution interpreted as damping in vibroimpact," 1975.

[17] J. W. Hutchinson, "Fundamentals of the phenomenological theory of nonlinear fracture mechanics," *Journal of Applied Mechanics*, vol. 50, no. 4b, pp. 1042–1051, 1983.

[18] N. Xydas and I. Kao, "Modeling of contact mechanics and friction limit surfaces for soft fingers in robotics, with experimental results," *The International Journal of Robotics Research*, vol. 18, no. 9, pp. 941–950, 1999.

[19] N. Diolaiti, C. Melchiorri, and S. Stramigioli, "Contact impedance estimation for robotic systems," *IEEE Transactions on Robotics*, vol. 21, no. 5, pp. 925–935, 2005.

[20] G. Kim, Y.-H. Lee, and H.-W. Park, "A learning framework for diverse legged robot locomotion using barrier-based style rewards," *arXiv preprint arXiv:2409.15780*, 2024.

[21] Y.-H. Shin, S. Hong, S. Woo, J. Choe, H. Son, G. Kim, J.-H. Kim, K. Lee, J. Hwangbo, and H.-W. Park, "Design of kaist hound, a quadruped robot platform for fast and efficient locomotion with mixed-integer nonlinear optimization of a gear train," in *2022 International Conference on Robotics and Automation (ICRA)*. IEEE, 2022, pp. 6614–6620.