Rapid Quadrupedal Locomotion on Deformable Terrain

Mohammed Azharudeen Farook Deen, Omer Kemal Adak and Raul Fuentes

***Abstract*— Animals exhibit an extraordinary capability for agile locomotion across various natural terrains. However, replicating this adaptability by developing controllers capable of traversing on soft and deformable terrains, especially at high speed, remains a challenge. Our study proposes an end-to-end learned approach for a fast and robust controller that aims to guide quadruped robots through a wide range of deformable substrates such as soft soil, loose gravel, rough, rocky, wet, muddy, etc. We use a simulation environment that can be parameterized to represent diverse types of terrain that closely mimic the natural world and train a policy network via model-free reinforcement learning. Our proposed approach includes the key components (1) implementing a multiscale granular media model for real-time simulation of terrain dynamics that combines an active zone resolved as distinct particles on top of a resting compliant terrain, (2) using a terrain adaptive curriculum for command velocities, and (3) an online adaptation module strategy that can implicitly identify the terrain properties for easy sim-to-real transfer to physical robot leveraged from prior works.**

I. Introduction

Recent studies on legged robots have primarily focused on locomotion over rigid ground [1], [2], [3], [4]. As a result, many of the developed controllers exhibit compromised performance on soft and deformable terrains, especially as the robot speed increases [5], [6]. The primary reason is that reinforcement learning approaches, in general, are not effective beyond the data distribution as the agent cannot perform well in environments that it has not experienced.

We conducted preliminary real-world tests using a quadruped (Unitree A1) with its default onboard controller on sand, gravel, and mixed sand-gravel terrains. As demonstrated in Movies S1–S3 in the supplementary material, these trials resulted in frequent failures, emphasizing the challenges of navigating deformable substrates and motivating our approach to more advanced, terrain-adaptive control strategies. To address this gap, we are exploring a framework to train controllers for quadruped robots with non-rigid deformable terrains. At its core, we use a fast Position-based Dynamics (PBD) solver to simulate interaction with complex granular particles on top of compliant viscoelastic terrain. We use this as the main idea but have found that it is not sufficient to achieve robust locomotion on deformable terrain. As robots strive for high speeds [3], the influence of terrain variability on controller efficacy intensifies. Certain physical factors, such as contact force regulation and actuator limitations, only start to affect the robot's dynamics at elevated speeds. Training can typically be facilitated by initially assigning the agent straightforward tasks and progressively escalating their complexity through a structured curriculum. Hence, the second idea is to implement a terrain adaptive velocity curriculum that modifies the conventional velocity curriculum by introducing separate velocity command ranges for specific terrain types. This approach enables training quadrupeds to navigate all terrain types simultaneously, enhancing generalization. However, task difficulty is a function of both the terrain dynamics and the tracking commands. So, we also added a progressive terrain difficulty curriculum inspired by previous work [4], where training begins on flat terrain, progresses to compliant terrain, and culminates with granular terrain, which combines a compliant terrain with granular particles on top, creating a multiscale terrain representation [7].

The third key idea is to modify the existing adaptation module to explicitly incorporate deformable terrain parameters. The adaptation module is responsible for implicitly identifying terrain properties in real-time, allowing the policy to adjust its behavior dynamically. We aim to extend it to estimate parameters such as terrain stiffness/damping, friction, and granular particle parameters. This ensures improved sim-to-real transfer by enabling the controller to generalize better across a diverse range of deformable substrates [8].

Our ultimate goal is to construct a system that can traverse diverse terrains at a wide range of linear and angular velocities, potentially offering robust performance and effective zero-shot generalization. This extended abstract outlines our ongoing research directions to achieve these objectives.

II. Methods

An ideal training strategy for quadrupedal locomotion should enable learning a unified policy that generalizes across diverse terrains without requiring extensive fine-tuning. However, major constraints arise from the limited number of discrete particles that can be simulated in real-time, restricting both environment size and the number of parallel robot instances [4]. While larger granular domains would provide a more realistic terrain representation, they come at a significant computational cost. Additionally, our reliance on on-policy methods (such as Proximal Policy Optimization, PPO), which demands large amounts of robot data for stable convergence, further exacerbates these challenges, typically requiring up to 4096 parallel robots for efficient training.

The authors are with the Institute of Geomechanics and Underground Technology (GUT), RWTH Aachen University, Aachen, Germany. azharudeen.farook@rwth-aachen.de, (adak,raul.fuentes)@gut.rwth-aachen.de

Fig. 1. Terrain types used for training and testing in simulation: (a) Flat (b) Compliant, and (c) Granular Terrain

A. Simulation Pipeline

One of the primary reasons for limited progress in the field is the lack of an efficient simulation pipeline for soft and deformable terrain simulation. One promising direction for granular media simulation is the discrete element method (DEM), which models individual grains and their interactions. However, this approach requires prohibitively long computation times, making it unsuitable for reinforcement learning. Traditional DEM solver integrates Newton's equations of motion to determine particle trajectories, which are computationally impractical, especially for large numbers of particles. We use position-based dynamics (PBD) within NVIDIA's Isaac Sim framework, which, on the other hand, offers a more efficient alternative for real-time applications by directly manipulating particle positions to satisfy constraints, thereby avoiding the direct computation of forces.

As shown in Fig. 1, we use three different terrain types with variations within each type and resample parameters that induce variations of each terrain periodically. We employ a multi-scale terrain representation comprising a rigid flat ground, a range of compliant terrains (modeled through NVIDIA Isaac Sim's inbuilt compliant contact model using implicit springs) characterized by varying stiffness and damping, and granular terrains consisting of discrete particles atop a compliant substrate. The compliant category spans "soft-compliant," with minimal deformation, to "hard-compliant," which exhibits more pronounced yield under load. Granular terrains capture both vertical and horizontal reaction forces, accounting for lateral motion effects, and include low-cohesivity configurations akin to loose sand as well as highly cohesive setups resembling wet sand or mud (Movie S4 in supplementary demonstrates the robot traversing a granular terrain within our simulation environment). This multi-scale model is intended to reflect real-world conditions, where terrain can exhibit both compliant and granular characteristics simultaneously.

B. Terrain Adaptive Velocity Curriculum

When velocity commands are evenly sampled from a narrow range at the onset of training, the agent is capable of learning to track them. Nonetheless, when commanded velocities are consistently sampled from an extensive distribution, we observed that learning is unsuccessful. Failure arises from the difficulty of high-speed mobility; when most commands require high velocity, the agent is unable to accumulate sufficient rewards. This problem may be alleviated by initially exposing the agent to low-velocity commands and progressively augmenting the desired speed through a structured curriculum.

However, it is compounded by the varying difficulty levels across terrains: low-speed commands on soft terrain can still pose a challenge comparable to high-speed commands on more rigid terrain. Consequently, a single, global velocity command range can generate infeasible high-speed tasks for robots on certain difficult terrain that hinders training. To address this, we introduce a terrain-adaptive curriculum that adjusts the velocity command distribution for each terrain type based on its inherent difficulty.

C. System Identification

Achieving robust locomotion on deformable terrains demands real-time estimation of key properties like compliance, damping, and granular resistance—especially on substrates such as deep sand or mud where robot interactions significantly alter terrain dynamics. To facilitate sim-to-real transfer, we employ privileged learning: a teacher policy is first trained using detailed, simulation-only terrain data, after which a student policy learns to mimic the teacher using only onboard sensor inputs. Our enhanced adaptation module reconstructs these hidden states by integrating proprioceptive feedback with richer signals—including foot reaction forces, robot–particle contact states (weighted by parameters such as particle density and cohesion), terrain compliance properties, and joint contact information over a 50-timestep window—to form a latent representation of the terrain's extrinsic parameters. This approach, combined with dynamic domain randomization of critical physical parameters (tuned via tests such as the static angle-of-repose), enables robust adaptation across a diverse range of deformable terrains.

III. Conclusion and Future Work

While our approach remains under active development, we anticipate the proposed framework will serve as a scalable foundation for robust legged locomotion in unstructured environments. Furthermore, we plan to calibrate environment parameters to more accurately reflect real-world outdoor conditions.

In ongoing work, we are actively exploring more sample-efficient algorithms, such as Soft Actor-Critic (SAC). By leveraging SAC, we aim to potentially reduce the number of parallel simulation robots to fewer than 50 instances, thus allowing us to accommodate a greater number of granular grids in our training, without significantly increasing computational costs. Also, SAC's ability to incorporate off-policy data makes it a promising candidate for fine-tuning in real-world deployments.

Supplementary Material

Movies S1–S3: Testing cases of the A1 quadruped robot with its default onboard controller on sand, gravel, and sand-gravel mix.

Movie S4: Demonstration of our robot traversing granular terrain within our simulation environment.

Manuscript 235 submitted to 1st German Robotics Conference.
Received February 5, 2025.

REFERENCES

[1] J. Lee, J. Hwangbo, L. Wellhausen, V. Koltun, and M. Hutter, “Learning quadrupedal locomotion over challenging terrain,” *Sci. Robot.*, vol. 5, no. 47, p. eabc5986, Oct. 2020, doi: 10.1126/scirobotics.abc5986.

[2] Y. Ji, G. B. Margolis, and P. Agrawal, “DribbleBot: Dynamic Legged Manipulation in the Wild,” Apr. 03, 2023, arXiv: arXiv:2304.01159. Accessed: Apr. 05, 2023. [Online]. Available: http://arxiv.org/abs/2304.01159

[3] G. B. Margolis, G. Yang, K. Paigwar, T. Chen, and P. Agrawal, “Rapid Locomotion via Reinforcement Learning,” *in Robotics: Science and Systems XVIII, Robotics: Science and Systems Foundation*, Jun. 2022. doi: 10.15607/RSS.2022.XVIII.022.

[4] N. Rudin, D. Hoeller, P. Reist, and M. Hutter, “Learning to Walk in Minutes Using Massively Parallel Deep Reinforcement Learning,” Aug. 19, 2022, arXiv: arXiv:2109.11978. Accessed: Sep. 10, 2023. [Online]. Available: http://arxiv.org/abs/2109.11978.

[5] S. Choi et al., “Learning quadrupedal locomotion on deformable terrain,” *Sci. Robot.*, vol. 8, no. 74, p. eade2256, Jan. 2023, doi: 10.1126/scirobotics.ade2256.

[6] I. Radosavovic, S. Kamat, T. Darrell, and J. Malik, “‘Learning Humanoid Locomotion over Challenging Terrain,” Oct. 04, 2024, arXiv: arXiv:2410.03654. doi: 10.48550/arXiv.2410.03654.

[7] M. Servin, T. Berglund, and S. Nystedt, “A multiscale model of terrain dynamics for real-time earthmoving simulation,” *Adv. Model. Simul. Eng. Sci.*, vol. 8, no. 1, p. 11, May 2021, doi: 10.1186/s40323-021-00196-3.

[8] A. Kumar, Z. Fu, D. Pathak, and J. Malik, “RMA: Rapid Motor Adaptation for Legged Robots,” Jul. 08, 2021, arXiv: arXiv:2107.04034. Accessed: Sep. 25, 2023. [Online]. Available: http://arxiv.org/abs/2107.04034

Manuscript 235 submitted to 1st German Robotics Conference.
Received February 5, 2025.