Learning to Step: Deep Reinforcement Learning for Robust Foothold Planning in Complex Terrains

Description

Traditional quadruped locomotion relies heavily on "Manual Rules" and heuristics for foothold selection. While efficient on flat ground, these hard-coded rules often fail when facing the chaos of the real world—loose rubble, slippery slopes, or irregular gaps. If the engineer didn't code a rule for it, the robot falls.
We are ditching the manual rulebook in favor of Deep Reinforcement Learning (Deep RL). Instead of telling the robot how to step, we create a simulation environment where it learns to step through millions of trials and errors. We aim to train a neural network policy that autonomously perceives the terrain and selects the optimal foothold, just like a biological brain.

Background

The primary advantage of legged robots over wheeled ones is their potential to go anywhere: climbing stairs, stepping over gaps, and navigating disaster zones. However, achieving this potential requires the robot to interact intelligently with the ground. Currently, most miniature quadrupeds rely on "Blind Locomotion." They execute a fixed rhythm and react only after they hit an obstacle. While this works on flat floors, it is dangerous in the wild. If a robot steps blindly onto the edge of a loose rock or into a gap, it loses balance instantly. To survive, the robot must transition from Reactive (stumbling and recovering) to Proactive (perceiving and planning).
In the past, engineers tried to solve this by writing explicit rules (e.g., "If the surface is tilted > 30 degrees, don't step"). However, the real world is too chaotic to be captured by if-else statements. Hand-crafted heuristics are brittle; they break whenever the robot encounters a scenario the engineer didn't foresee.
This is where Data-Driven Control enters the picture. Nature doesn't program rules; nature learns. By using Deep Reinforcement Learning, we can empower the robot to discover its own strategies. Through massive parallel training, the agent learns to interpret terrain geometry and identify stable footholds implicitly, achieving a level of robustness and adaptability that traditional control theory cannot match.

Technische Universität München

TUM School of Computation, Information and Technology

Lehrstuhl für Robotik, Künstliche Intelligenz und Echtzeitsysteme

Supervisor:
Prof. Dr.-Ing. Alois Knoll

Advisor:
Qian Huang M.Sc.

Type:
MA,SA,BA

Research area:
Deep Reinforcement Learning, Robotics, Sim-to-Real, Locomotion

Programming language:

Python (PyTorch), C++ (Deployment)

Requirements:
High self-motivation and passion for AI-driven robots; At least six-month working time; (Optional) Experience with Python and Deep Learning frameworks ro Physics Simulators (NVIDIA Isaac Sim, MuJoCo)

Language:
English

For more information please contact us:

E-Mail: qian.huang@tum.de

Internet: www.ce.cit.tum.de/air

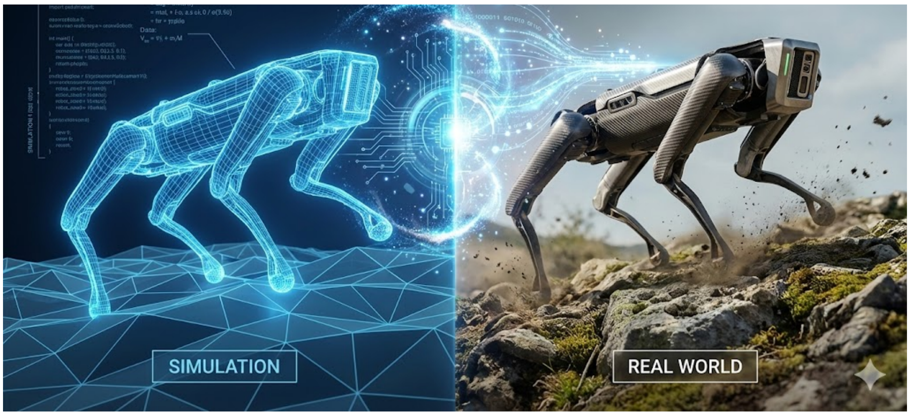

Technische Universität München

TUM School of Computation,
Information and Technology

Lehrstuhl für Robotik, Künstliche
Intelligenz und Echtzeitsysteme

Tasks

1. Foundation: Reproduce State-of-the-Art:
Task: Your first mission is to reproduce an open-source gait control framework (e.g., based on legged_gym or similar SOTA works).
Goal: Get a stable "Blind Walking" policy running in simulation. This creates the essential locomotion backbone upon which your foothold planner will be built.
2. Simulation Setup:
Build a challenging training environment in Isaac Sim or MuJoCo.
Procedurally generate complex terrains: stairs, slopes, stepping stones, and random debris fields to challenge your baseline.
3. Perceptive RL Training:
The Core Task: Extend the baseline by integrating exteroceptive perception.
Design the network to process Height Maps and output specific Foothold Selections.
Train the agent using PPO to prefer safe, flat surfaces over risky edges.
4. Sim-to-Real:
Implement Domain Randomization (randomizing friction, mass, sensor noise) to make the policy robust.
(Optional) Deploy your trained policy onto the real robot and validate that it can intelligently choose footsteps in the real world.

What You Will Gain

AI + Robotics Mastery: Get hands-on experience with the hottest tech stack in robotics: PyTorch + NVIDIA Isaac Gym/Mujoco.

Algorithm Design: Gain a deep, intuitive understanding of PPO, Actor-Critic architectures, and Value Function approximation.

Research Impact: This is a frontier research topic. Excellent results here have potential for publication in top-tier conferences.

Mentorship & Support

This topic, which focuses on designing for advanced robotic control, may sound challenging. However, you will not be starting from scratch.
Your mentor (Ph.D. Student) has extensive research experience in this specific domain and has already established a solid foundation for this project. We are fully prepared to provide comprehensive, step-by-step guidance to ensure you get up to speed quickly. You will receive dedicated, hands-on support throughout the entire research and implementation process.
This is a unique opportunity to tackle a high-impact problem with expert, full-time mentorship.

For further discussion on specific tasks, welcome to direct contact me via email.