Karlsruhe Institute of Technology

Autonome Lernende Roboter (ALR)

Prof. Gerhard Neumann

Project Type

- ☑ Master Thesis
- ☑ Bachelor Thesis
- ☑ Research Project

Supervisors

Florian Seligmann

florian.seligmann@kit.edu

Onur Celik

onur.celik@kit.edu

Difficulty

Algorithmic

Math

Application

Predicting PD Constants for Dynamic Humanoid Locomotion

Description

In modern sim-to-real reinforcement learning (RL) for legged locomotion, control policies typically predict target joint positions, which are then tracked by a low-level Proportional-Derivative (PD) controller. However, the PD gains for stiffness (K_p) damping (K_d) typically remain fixed. This fundamentally limit the robot's agility and compliance: A robot with high stiffness can move fast but suffers from rigid, loud, high-impact interactions with the ground, whereas low stiffness leads to sluggish tracking.

In this project, we want to overcome this limitation by allowing the RL agent to predict PD constants alongside joint angles. This is highly inspired by biological locomotion, where, during running, muscle stiffness dynamically adapts to absorb impacts upon heel strike and to provide explosive force during push-off. By porting this concept to humanoid robots, we aim to achieve significantly more silent and compliant walking, while improving the energy efficiency of the locomotion policy. For quadruped robots, it has already been shown that variable stiffness improves the policy's energy efficiency [1].

Specifically, applied to the Unitree G1, the robot should learn to instantaneously reduce its stiffness the moment its foot touches the floor to absorb the shock, minimizing acoustic noise and mechanical wear. Simultaneously, it can adaptively increase stiffness in the stance leg to enable fast, highly dynamic movements like running.

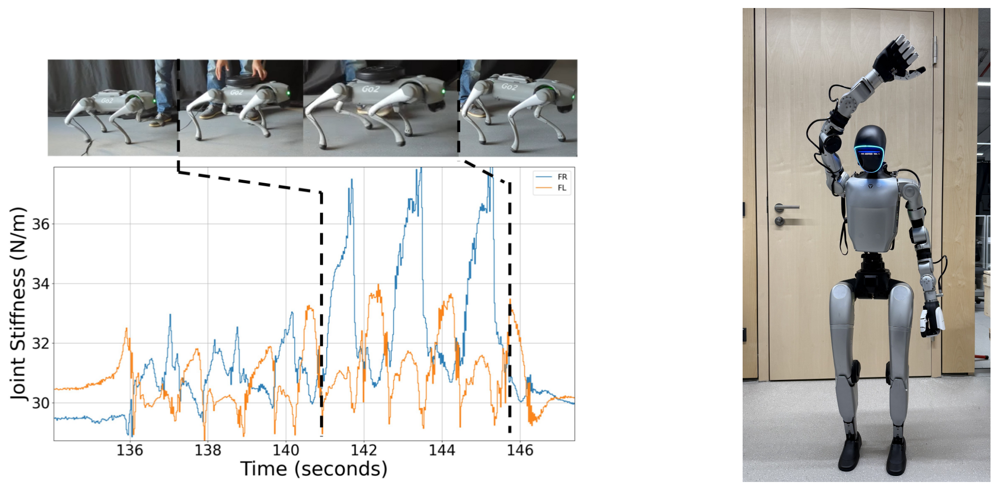

Figure 1: Left: Variable Stiffness for a Quadruped Robot (taken from [1]). Right: Our Unitree G1 humanoid that you will deploy your policy on.

Tasks

Depending on the type of thesis (Bachelor/Master), the tasks in this project will involve:

- Literature Research: Research existing approaches for dynamic PD gains and get familiar with our existing sim-to-real locomotion pipeline.
- Action Space Extension: Modify the RL architecture so the policy outputs the target joint angles alongside the K_p and K_d constants per timestep.
- Reward Engineering: If necessary, design reward functions that specifically penalize high impact forces and joint torques.
- Sim-to-Real Deployment: Transfer the learned policy to the physical Unitree G1 robot and evaluate various metrics such as acoustic footprint, running speed, and energy efficiency.

References

[1] Dario Spoljaric, Yashuai Yan, and Dongheui Lee. Variable stiffness for robust locomotion through reinforcement learning, 2025.