MA/SA/BA

Differentiable Simulation for Efficient and Versatile Quadruped Locomotion

Background

Recent work by Song et al. [1] has demonstrated the potential of differentiable simulation to dramatically accelerate quadruped locomotion learning. This approach overcomes the high sample complexity inherent in traditional reinforcement learning (RL) methods, which rely on high-variance, zeroth-order gradient estimates. By contrast, differentiable simulation computes low-variance, first-order gradients by backpropagating through the system's dynamics, enabling a quadruped to learn stable walking policies in minutes without extensive parallelization. While this approach marks a significant advance, its reliance on pre-defined gait schedules and heuristics for foot placement limits its autonomy and adaptability.

This thesis aims to push the frontier of differentiable simulation for quadruped locomotion. Beyond improving efficiency, the project will investigate emergent gait discovery—where natural gaits arise without manual scheduling. Furthermore, this work will explore enhancing the generalization of learned policies across a wider range of terrains. The fundamental goal is to investigate how differentiable simulation can support robust and versatile locomotion in real-world environments.

Your Tasks

Through this project, you will gain hands-on expertise in state-of-the-art robot learning algorithms and frameworks.

Your tasks include:

- Perform an in-depth literature review on differentiable simulation-based robot learning.
- Implement a hybrid differentiable simulation pipeline.
- Explore free-gait (schedule-free) locomotion across diverse terrains.
- Compare performance against RL baselines (e.g., PPO).
- Investigate sim-to-real transfer to our **Unitree Go2** robot

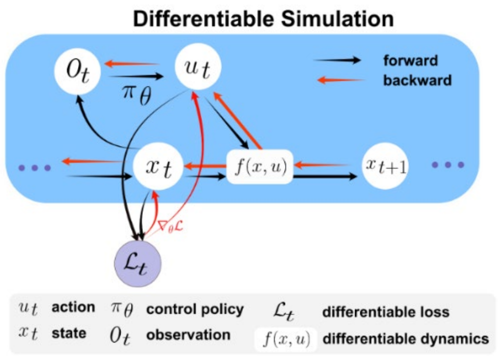

Graphical model for policy learning using differentiable simulation. [1]

Requirements

- High self-motivation and passion for research.
- Six months working time.
- Background in robotic dynamics, deep learning and reinforcement learning is an advantage.
- Programming skills (Python, PyTorch; CUDA helpful).

Supervisor: Prof. Alois Knoll.

Advisor: Yulong Xiao, yulong.xiao@tum.de

Lehrstuhl für Robotik, Künstliche Intelligenz und Echtzeitsysteme, TUM School of Computation, Information and Technology

[1] Y. Song, S. Kim, and D. Scaramuzza, "Learning Quadruped Locomotion Using Differentiable Simulation," 2024.

*Unitree Go2 picture from https://www.unitree.com