Unified Omnidirectional Jumping and Walking for a Quadruped in Reduced Gravity using Reinforcement Learning

Abstract: As planetary exploration targets scientifically valuable but inaccessible regions such as Martian lava tubes, steep crater rims, and boulder-filled terrain, the need for agile robotic mobility becomes paramount. While the Olympus jumping quadruped has demonstrated high-performance jumping and robust walking, current control strategies developed for the Olympus quadruped rely on separate, task-specific policies for vertical and horizontal maneuvers. This project aims to bridge this gap by developing a unified reinforcement learning policy within the olympus_lab framework. The main objective is to enable the robot to seamlessly transition between walking and jumping by utilizing a 3D takeoff velocity vector as a primary control input. This approach will allow for arbitrary jump trajectories, providing the flexibility required for complex exploration scenarios where fixed jump profiles are insufficient. We also want to incorporate yaw rotation targets. The work will leverage existing curriculum-based learning strategies and physics-informed reward densification to achieve precise landing and robust traversal in simulated Martian gravity. If progress is sufficient, the developed policy will be validated on the Olympus hardware to demonstrate successful Sim2Real transfer in earth gravity.

Tasks:

- Study and understand Reinforcement Learning fundamentals, specifically Proximal Policy Optimization (PPO) and curriculum-based reference state initialization.
- Study the existing olympus_lab framework and the current implementation of separate walking and jumping policies.
- Develop a unified policy architecture that accepts a 3D takeoff velocity vector to enable omnidirectional jumping and transitions.
- Design and implement a densified reward structure based on projectile motion equations to facilitate learning of arbitrary trajectories. Including yaw rotations
- Train and evaluate the control policy in a simulated Martian gravity environment using NVIDIA Isaac Sim.
- Implement and validate the control policy on the Olympus quadruped hardware to assess landing precision and stability.

Literature (indicative):

[1] J. A. Olsen, L. R. Pettersen, and K. Alexis. "Towards quadrupedal jumping and walking for dynamic locomotion using reinforcement learning." IEEE Robotics and Automation Letters (RA-L), 2026.

[2] J. A. Olsen, G. Malczyk, and K. Alexis. "Olympus: A jumping quadruped for planetary exploration utilizing reinforcement learning for in-flight attitude control." IEEE ICRA, 2025.

[3] J. A. Olsen and K. Alexis. "Towards low-gravity planetary exploration using reinforcement learning for walking, jumping, and in-flight attitude control." 2025. [4] N. Rudin, D. Hoeller, P. Reist, and M. Hutter. "Learning to walk in minutes using massively parallel deep reinforcement learning." Conference on Robot Learning (CoRL), 2022.

[5] N. Rudin, H. Kolvenbach, V. Tsounis, and M. Hutter. "Cat-Like Jumping and Landing of Legged Robots in Low Gravity Using Deep Reinforcement Learning." IEEE Transactions on Robotics, 2021.

[6] Relevant project framework: https://github.com/ntnu-arl/olympus_lab

[7] N. Rudin, D. Hoeller, P. Reist, and M. Hutter. "Learning to walk in minutes using massively parallel deep reinforcement learning." Conference on Robot Learning (CoRL), 2022.

[9] Olympus page: https://ntnu-arl.github.io/olympus/

Main supervisor: Kostas Alexis, Professor, NTNU | **Co-supervisor:** Jørgen Anker Olsen, PhD