

Payload-Adaptive Walking, Jumping, and In-Flight Attitude Control for a Jumping Quadruped in Reduced Gravity

Abstract: For future robotic missions to planetary bodies, legged platforms must be capable of carrying diverse scientific instruments such as cameras, spectrometers, and LiDAR systems. However, the addition of a significant payload, which can reach up to 2.5 kg, fundamentally alters the robot's mass distribution, center of gravity, and inertia tensor. These changes can lead to instability during walking or significant landing inaccuracies during jumping maneuvers if the control policies are not payload-aware. This project thesis aims to develop payload-adaptive reinforcement learning policies using the `olympus_lab` framework. The research will focus on creating a robust controller capable of maintaining high performance in walking, precise jumping, and in-flight attitude reorientation despite varying payload configurations. By utilizing techniques such as domain randomization and history-dependent policy architectures, the robot should learn to implicitly infer its physical properties from motion feedback. The results will be validated in a simulated Martian environment to ensure that the Olympus platform can safely transport scientific equipment across challenging terrain.



Tasks:

- Study the Reinforcement Learning problem formulation and the `olympus_lab` framework for walking, jumping, and attitude control.
- Analyze how variations in payload mass and mounting position affect the dynamics of the 5-bar leg mechanism and body reorientation during flight.
- Investigate and implement policy architectures, such as Recurrent Neural Networks (RNN) or Transformers, that allow the robot to adapt to changing dynamics through temporal context.
- Train a comprehensive policy (or separate policies) that maintains performance in walking, jumping, and in-flight attitude control under various payload scenarios.
- Evaluate the performance of the adaptive policy compared to baseline controllers in simulated Martian gravity to assess 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] G. B. Margolis and P. Agrawal. "Walk these ways: Tuning robot control for generalization with multiplicity of behavior" International Conference on Robotics and Automation (ICRA), 2023.
- [7] Rudin, N., Hoeller, D., Reist, P., & Hutter, M. (2022, January). Learning to walk in minutes using massively parallel deep reinforcement learning. In Conference on Robot Learning (pp. 91-100). PMLR.
- [8] Relevant project framework: https://github.com/ntnu-arl/olympus_lab
- [9] Olympus page: <https://ntnu-arl.github.io/olympus/>

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