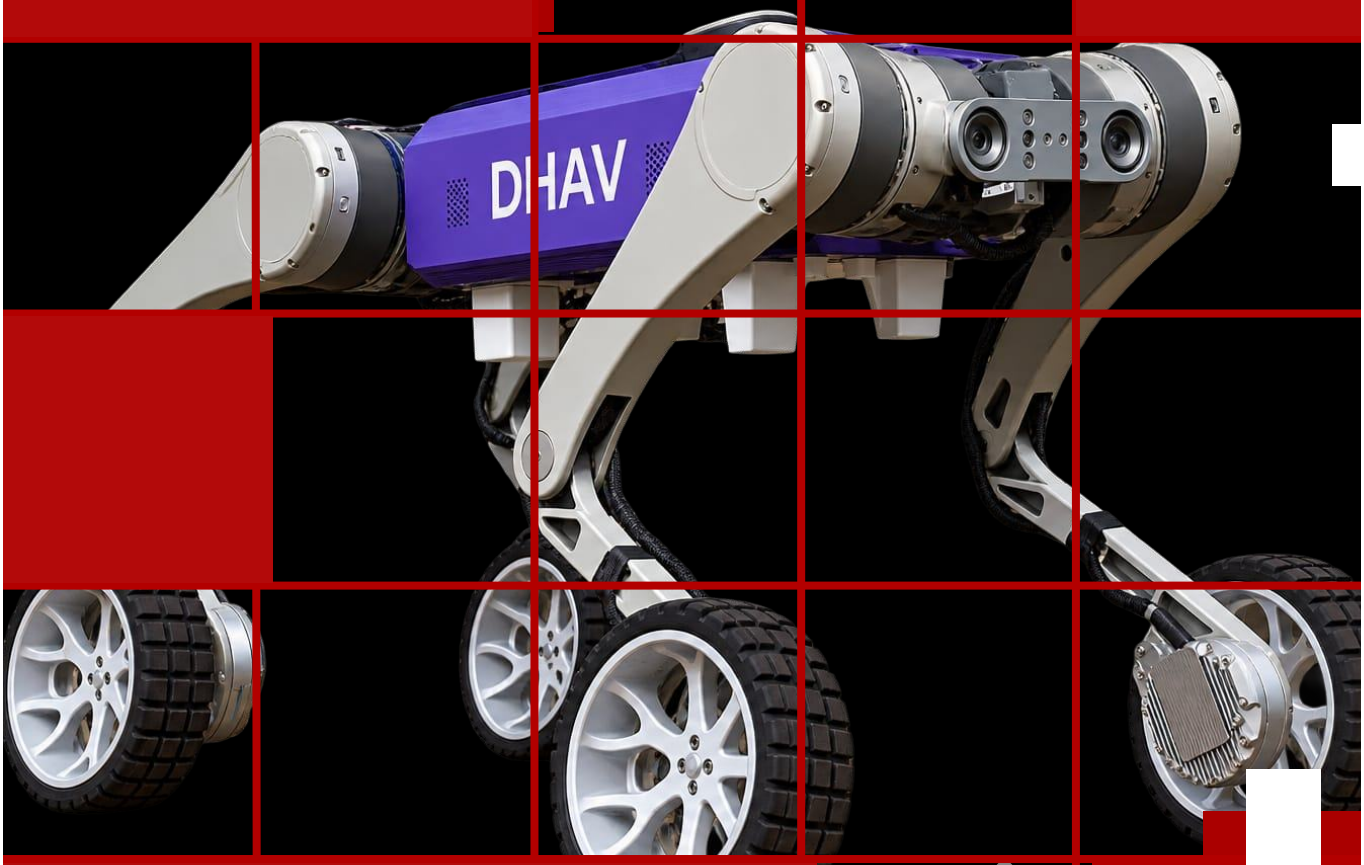
INTER IIT
TECH MEET 15.0



MID PREP

High speed locomotion policy for wheeled quadruped

Introduction

Wheeled-legged quadruped robots combine the speed and energy efficiency of wheeled mobility with the terrain adaptability of legged locomotion. However, achieving stable, high-speed movement in these hybrid systems remains a significant control challenge due to complex nonlinear dynamics, rapid changes in contact forces, sharp turns, and transitions between rolling and stepping.
This problem statement focuses on developing a robust, high-speed locomotion policy using reinforcement learning, trained in simulation and deployed on physical hardware. The objective is to enable DHAV-Hy, xTerra Robotics' indigenously developed wheeled-legged quadruped, to navigate dynamic environments with improved stability, agility, and mobility while leveraging the complementary strengths of wheels and legs.

About us

xTerra Robotics is an Indian deep-tech startup founded in 2023 , focused on developing autonomous legged robots and advancing embodied AI. We design, manufacture, and assemble our robotic systems in India, building indigenous capabilities in advanced robotics. Our vision is to become a global leader in autonomous robotics by creating intelligent, adaptable robots capable of operating across diverse environments. Through reliable, high-performance robotic solutions, our mission is to enable people to work more safely and efficiently while advancing India's technological self-reliance and strengthening its capabilities in robotics and embodied AI.

DHAV-Hy is India's first indigenously designed, manufactured, and assembled wheeled-legged quadruped robot for commercial applications.

https://youtu.be/AfRuGAZmWPY?si=jiHjrCK1qe3j0VBB

Problem Statement Description

Train a Reinforcement Learning policy for high speed blind locomotion of Dhav-Hy. It should be able to maneuver even on uneven terrain, without any perceptive sensors, joint states and imu data will be available. The goal is to make the robot go as fast as possible while being stable (flat + uneven terrain, upto ~10 cm obstacles). While chasing the goal you cannot forget the other parts of locomotion, which brings us to the solution requirements.

Policy Expectations:

1. The robot should achieve a very fast forward velocity (main criteria you will be judged on)
2. A low **cost of transport (COT)**
3. The velocity tracking for other low speed motions (lateral and yaw motion, you can surprise us with high velocities in this too)
4. The policy must not fall during operation and should handle minor disturbances
5. No self collisions
6. Controllability, the robot must be able to yaw comfortably even at its high speed forward velocity, etc.
7. Able to carry a payload of at least 10 kgs
8. Measures for some sim to real transferability

These are the criteria upon which your policy will be judged on

We will be providing the contingents with a basic code base so that they don't have to write training from scratch. The main innovation will be in architecture and methodology of training for high speed robust locomotion (read papers)

Technical criteria

Policy Interface and Control

1. Policy will run at 50Hz, outputting joint positions(16 dim). The input will be joint states and imu data(or a history of).
2. The leg joints will be position controlled, while the wheels will be velocity controlled.
3. The robot uses a normal radio transmitter for velocity inputs (like Flysky or Radiomaster).
4. The driver cannot toggle between policies, a single policy is expected, the policy doesn't need to have fall recovery
5. The teams will be provided with the URDF and a preliminary code base

Sensing

1. The locomotion must be blind, no perceptive (LiDAR, camera) inputs, the only available data are joint states and imu data (the ordering of the input and dimensions will be clear once the repo is shared).

Onboard Compute

1. The onboard compute will be equivalent to a Raspberry Pi. Will provide an upper cap to policy size later, should not be an issue regardless.

Operating Conditions and Robustness

1. Surface friction cannot be assumed, it's part of the simulation to real gap.
2. The payload is mounted on top of the robot, the robot's COM could easily vary with at least ±7 cm in x & y and +14 cm in z direction.
3. The lateral push will be of 250N for 0.1 seconds.
4. The policy is not recommended to have fall recovery, only push recovery (small disturbances).

5. The number of physical Dhav-Hy will depend on the availability of number of robots. The minor difference in mechanical performance of robots or between runs is what the sim to real transferability is for. (I don't recommend focusing on this part)

Safety and Failure Conditions

1. The hardware safety will trigger when the robot hits a 70 degree pitch or roll angle, indicating a fall and ending your run.
2. If the robot falls during normal operation because of policy behaviour, the team will not be allowed a re run.

Evaluation Setup

1. The final evaluation will be done in an open ground (like a football field).
2. There won't be a yaw test in the 100m dash, but controllability during high speed motion will be taken into account.
3. The robot will begin the 100m dash from a standing pose.
4. There will be two leaderboards, a sim evaluation and a real HW evaluation. Expect the sim evaluation to be more rigorous.
5. The sim evaluation script will be hidden from the teams.
6. Throughout the challenge, there will be a public scoreboard on the main repo showing the velocity each team is achieving, along with whether stability is met. It will be updated every 4 days. The main repository from which the teams will fork their repos will update the stability eval and a speed leaderboard.
7. Cost of Transport will be calculated via either of the two formulas below. E is Energy, P is power, m is mass, g is gravity, d is distance and v is velocity.

$$\text{CoT} = \frac{E}{m \cdot g \cdot d} = \frac{P}{m \cdot g \cdot v}$$

Deliverables

The contingent is expected to maintain a private fork of our repo, which can be accessed by us. The results should be reproducible via the repo, and it must always contain your latest policy as final.pt and the architecture of the policy being used in the runner files.

- Any files that are required for end to end reproduction of the performance are expected to be pushed/submitted to the repo.

Judging Criteria

The final score is split across three components:

Component	Weightage
Hidden sim evaluation	25%
Real test	60%
Short presentation on novelty of your solution	5%
Technical Report	10%

Simulation Safety Test (Pre-deployment Gate)

At the finals, before your policy is deployed on the robot, it must pass our simulation test to ensure safety on the robot. We will be using the final.pt in your repo to keep you updated on this throughout the duration the statement is active.

Hidden Sim Evaluation (25%)
A hidden sim evaluation will be done on your policy. This script is not shared with the teams.

Novelty Presentation (5%)
A short presentation on the novelty of your solution, given just before the real test.

Real Run/Test (60%)
The robot will have to complete a 100m dash, driven by a designated driver from the contingent, this will be timed. The quantities described in Policy Expectations will be measured from this run.

Technical Report (10%)
A Technical report on findings , experimentations , ablations , training setup , timeline of progress , etc

Hardware run procedure:

1. The xTerra team will do the zero state calibration of the robot. The teams will be directly given the control to the robot with the policy enabled.
2. The teams will get two physical attempts, the best will be counted. The flashing time is not a concern, that will be the xTerra team's responsibility.

Resources

1. We will be providing access to a private github repo to the contingents.
2. The contingents are encouraged to look at latest research papers regarding high speed locomotion (blind).