

# Terrain Trekking: Achieving Campus Autonomy with Robot Dogs

Jing Guo,Xin Duan

June 26, 2024

## Abstract

Achieving campus autonomy with robot dogs is a significant challenge that requires the robots to navigate various terrains and environments. Existing methods either focus on specialized skills or lack vision-based capabilities. In this work, we plan to propose a system for learning vision-based autonomy skills without relying on reference animal data or complex reward systems. We will utilize a reinforcement learning approach inspired by direct collocation to enable the robot dogs to autonomously navigate campus terrains, including climbing over obstacles, traversing uneven surfaces, crossing gaps, and adapting to different terrain types. Our system utilizes the robot dogs' onboard sensors, such as depth cameras, to perceive and react to the environment. Through experiments and evaluations, we demonstrate the effectiveness of our approach in enabling robot dogs to autonomously navigate and achieve campus autonomy over various terrains.

## 1 Introduction

Humans and animals possess amazing athletic intelligence. Parkour is an exemplar of athletic intelligence of many biological beings capable of moving swiftly and overcoming various obstacles in complex environments by running, climbing, and jumping. Such agile and dynamic movements require real-time visual perception and memorization of surrounding environments, tight coupling of perception and action, and powerful limbs to negotiate barriers. One of the grand challenges of robot locomotion is building autonomous parkour systems. Boston Dynamics Atlas robots have demonstrated stunning parkour skills. However, the massive engineering efforts needed for modeling the robot and its surrounding environments for predictive control and the high hardware cost prevent people from reproducing parkour behaviors given a reasonable budget. Recently, learning-based methods have shown robust performance on walking, climbing stairs, mimicking animals and legged mobile manipulation by learning a policy in simulation and transferring it to the real world while avoiding much costly engineering and design needed for robot-specific modeling.

We introduces a robot parkour learning system for low-cost quadrupedal robots that can perform various parkour skills, such as climbing over high obstacles, leaping over large gaps, crawling beneath low barriers, squeezing through thin slits, and running. Our reinforcement learning method is inspired by direct collocation and consists of two simulated training stages: RL pre-training with soft dynamics constraints and RL fine-tuning with hard dynamics constraints. In the RL pre-training stage, we allow robots to penetrate obstacles using an automatic curriculum that enforces soft dynamics constraints. This encourages robots to gradually learn to overcome these obstacles while minimizing penetrations. In the RL fine-tuning stage, we enforce all dynamics constraints and fine-tune the behaviors learned in the pre-training stage with realistic dynamics. In both stages, we only use a simple reward function that motivates robots to move forward while conserving mechanical energy. After each individual parkour skill is learned, we use DAggerto distill them into a single vision-based parkour policy that can be deployed to a legged robot using only onboard perception and computation power.

## 2 State of the Art

### 2.1 Robot Locomotion

Dynamic robot locomotion needs real-time visual perception of surroundings.

[1] presents a learning-based control approach for quadrupedal robots that achieves high-level locomotion skills, including accurate and energy-efficient motions, through training in simulation. The approach enables the robot to adapt to different tasks and environments, including high-speed locomotion and dynamic recovery from falls, without extensive tuning on the physical robot. The results highlight the potential of learning-based methods in enhancing the performance and robustness of legged robots, offering a promising direction for the development of autonomous and agile robotic systems . [2] introduces a novel system for enabling low-cost quadrupedal robots to perform various parkour skills such as climbing over high obstacles, leaping over large gaps, crawling beneath low barriers, and squeezing through thin slits. The system is based on a two-stage reinforcement learning (RL) method that overcomes difficult exploration problems by using a pre-training stage with soft dynamics constraints and a fine-tuning stage with hard dynamics constraints. This approach allows robots to gradually learn to overcome obstacles while minimizing penetrations and conserving mechanical energy. The learned skills are then distilled into a single vision-based parkour policy that can be deployed on a robot using only onboard perception and computation power . [3] addresses the limitations of learned locomotion policies by introducing a structured family of locomotion strategies, known as Multiplicity of Behavior (MoB), which allows for rapid adaptation to new tasks and environments. The proposed approach eliminates the need for time-consuming retraining by selecting different strategies in real-time based on task requirements. The authors also present Walk These Ways, an open-source MoB locomotion controller capable of executing diverse gaits and enabling a wide range of downstream tasks, including crouching, hopping, high-speed running, stair traversal, bracing against shoves, and rhythmic dance. [4] This paper addresses the limitations of the standard modular control pipeline for legged manipulators with attached arms, which require extensive engineering and can result in non-smooth motions. The authors propose a unified policy for whole-body control using reinforcement learning, leveraging Regularized Online Adaptation and Advantage Mixing techniques. The proposed approach demonstrates dynamic and agile behaviors across various task setups, showcasing the potential of the low-cost legged manipulator design. [5] explores robot parkour using a small, low-cost robot with imprecise actuation and a single front-facing depth camera. A neural network policy trained through reinforcement learning enables the robot to perform precise parkour movements despite the limitations, including high jumps, long jumps, handstands, and running on tilted ramps.

### 2.2 Robot Navigation

[6] presents a fully learned approach for agile navigation of quadrupedal robots in parkour-like scenarios, utilizing advanced locomotion skills and a high-level navigation policy. The approach demonstrates successful real-world transfer of skills trained in simulation, enabling the robot to navigate and cross challenging obstacles at speeds of up to 2 meters per second. This method provides a foundation for future research in legged robot capabilities on challenging terrains without the need for expert demonstration, offline computation, or a priori knowledge of the environment . [7] introduces a learning framework that combines imitation learning, reinforcement learning, and curriculum learning to develop agile and dynamic motor skills for legged robots. The framework successfully trains a quadruped robot to perform various dynamic tasks, such as trotting, pacing, and bounding gaits, as well as jumping and turning maneuvers. This approach demonstrates the potential of machine learning techniques in developing versatile and robust motor skills for legged robots, paving the way for their application in complex and unstructured environments .




Figure 1: Train the network with Issac Gym

## 3 System Design

### 3.1 System Description

We first pre-train the model with sfot dynamics constraints in simulation. We design a training method to train on Unitree-go2 based on the legged gym framework. Then fine-tuning with hard dynamics constraints on the real go-2 robot. Sim2real gap is always a great challenge in robot learning. So we will design the reward function based on the attributes of go-2, and use PPO to train the policy network. Then we distill the parkour policies to a single parkour policy and apply the model to the real go-2 robot. Finally, we will deploy our policy onto real robot and give a walk demo. After successfully deploying high dynamic robot skills including running, climbing and jumping, we will use high level planner with osmAG map to do navigation in campus.

The MoveIt package for ROS is a powerful and widely used software framework for motion planning, control, 3D measurement, kinematics, control and navigation of robots. It is mainly used in robotic arms and mobile robots in industry and research. MoveIt is developed based on ROS and provides a series of tools and functions to enable robots to perform complex movements and tasks. It provides helpful functions for our project. Gazebo is a popular physics simulator for robotics development, which can simulate robots in a 3D environment. The gzserver and the gzclient are two programmes that guarantee the simulation. The gzserver does all the calculations for the simulation, generates sensor data and is responsible for things that related to the processing part. The gzclient generates the user interface, visualizing the simulation. Gazebo uses environment variables that are valid throughout the system to describe gazebo model path, resource path, plugin path and etc. Gazebo also provides nice Graphical User Interface(GUI) for users to visualize and control the simulation. We could spawn our robot models in gazebo using .urdf, .sdf and .world files.

### 3.2 System Evaluation

Successfully run on complex terrains where the parkour policy can work properly. Give a video demo to show the capability of the robot.

## 4 Experiment

We first trained our models in simulation using Isaac Gym. Then we deploy our terrain-over policy on the Unitree Go2 robot dog with 12 leg joints. We use the Intel RealSense D435i to get the depth image in real time at frequency of 30Hz. Our robot now can easily and stably climb onto a 35cm high box and even a 45cm higher one sometimes. This is much higher than each stair we may encounter in the real campus environment, so our robot can easily go over such terrains, going up and down stairs.

### 4.1 Data during training steps

We evaluate the robot’s performance through various training phases, highlighting differences between privileged training and knowledge distillation processes. These evaluations are crucial for assessing the adaptability and efficiency of the implemented training protocols. The accompanying figures underline the progression and outcomes throughout the training, demonstrating how the robot refines and applies learned behaviors effectively. Figure 2 provides a visual representation of these data metrics, enhancing the textual analysis with empirical evidence.

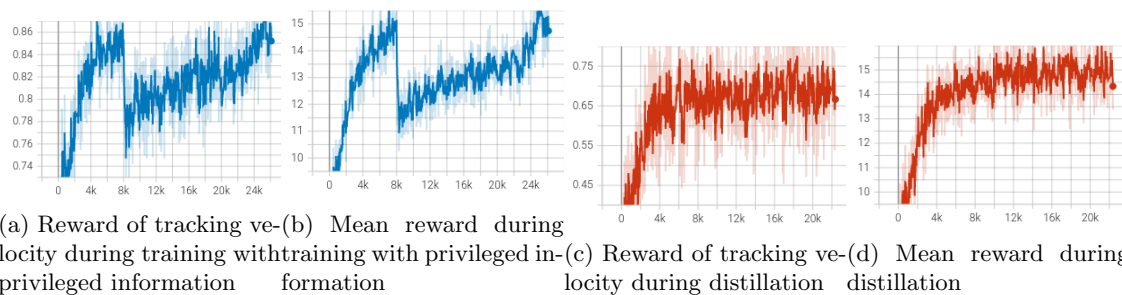


Figure 2: Key data during training steps

### 4.2 Checking onboard data

We utilize Foxglove Studio to facilitate real-time diagnostics and interaction with ROS2 messages, essential for monitoring the robot’s internal state and operational efficiency. The graphical interface displayed provides a comprehensive view of the system’s real-time functionality and communication flows, crucial for ongoing assessments and adjustments. Figure 3 captures this interface, offering insight into the diagnostic processes utilized during operational assessments.

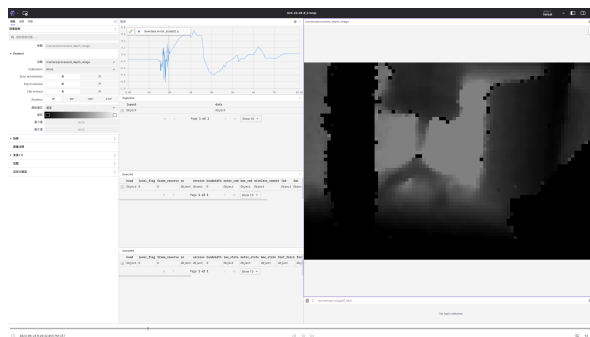


Figure 3: Checking ROS2 messages using Foxglove

### 4.3 Deploy on Unitree go2

We record the practical deployment of the robot in complex environments, showcasing its capabilities in navigating and interacting with various physical challenges. The focus is on demonstrating the robot’s maneuverability and adaptive responses during physically demanding tasks. The series of images in

