

Multi Agent Reinforcement Learning

Our Sponsors

Dr. Matt Brown
Dr. Daniele Reda

Wayve Technologies Ltd.


Our Team

Sasan Ghasaei | Felipe Garavelli
William Gibbs | Mahdi Shakouri | Itai Boss

Project # 2561



WAYVE



Motivation

Background


Wayve trains autonomous driving agents in simulation to develop their core capabilities. Achieving reliable sim-to-real transfer is essential to safely scale these models and reduce reliance on extensive real-world driving.

Problem


However, full-scale real-world testing remains costly, time consuming, and heavily constrained by safety regulations. This creates a need for an efficient, miniature testing platform to validate these simulation-trained policies.

Our Solution

We built an end-to-end, miniature testing platform using RC cars and an overhead vision system to seamlessly validate policies trained in Simulation.




System Overview



Simulation

- Isaac Lab as simulation framework for RL training
- Training in many parallel environments at once thanks to GPU utilization
- Collecting many training experiences in a short time due to high parallelization
- Picture alongside showing the training of line follower agents on multiple tracks at once!



Why training in simulation is useful?




Training in simulation to build basic skills

Training in reality to learn real physics and adapt to real world


Time

Hardware Platform

1:64 RC Car with Ackermann Steering




Aruco + Battery Upgrade




Training Tasks


Line Following: Make the car follow line




Object Avoidance: Line follow but be aware of obstacles




Car Resetter: Given a goal position, move the car to the goal and stay at the goal



Computer Vision



Model Architecture



Real-World Training

