Low-cost Terrain Navigation for Education in End-to-End Robotics

Laura Wang, Caleb Hsu, Byron Boots, and Tyler Han

An open-source framework enabling sim-to-real transfer for depth- based navigation on a budget- friendly wheeled platform

Robots dancing to Michael Jackson, playing Minecraft, or driving around the city—these agents are all trained using reinforcement learning and simulation through a method called sim-to-real. Yet, sim-to-real remains inaccessible to the broader community due to a dependence on expensive or proprietary technology.

WheeledLab [1] has recently gained visibility [2] as fully open-source pipeline for sim-to-real reinforcement learning integrating low-cost wheeled robots.

Examples of drifting (left), elevation traversal (mid), and visual navigation (right) policies trained through WheeledLab

To further demonstrate applications of such pipelines, this project builds on WheeledLab to achieve a policy for **depth-based navigation**

Motivation

Prior Requirements	New Requirements
Perfect height map	Depth camera
External tracking	Onboard sensing
Preprogrammed object dimensions	No manual scene specification
~$50k* infrastructure	<$500 sensor**

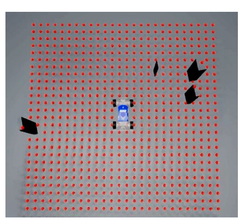

Sim visualization of sampled terrain heights

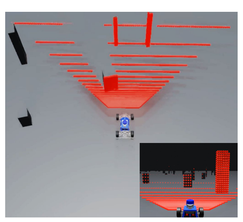

Sim visualization of sampled depths

*approx. $48,000 USD worth of equipment when using an Optitrack motion capture system
**approx. $420 USD Intel RealSense D455 depth camera

Methods

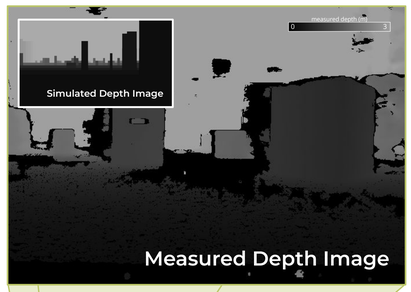

DOMAIN RANDOMIZATION

- noise
- pixel dropout

Physics randomization
- Wheel friction
- Robot mass
- Start and goal locations

HARDWARE

We use the MuSHR platform, estimated to cost 930 USD [3]. Among other features, the platform incorporates a
- 1/10 scale rally car chassis
- Realsense D455 depth camera
- NVIDIA Jetson Nano

We conducted training on a single NVIDIA RTX 3090 GPU.

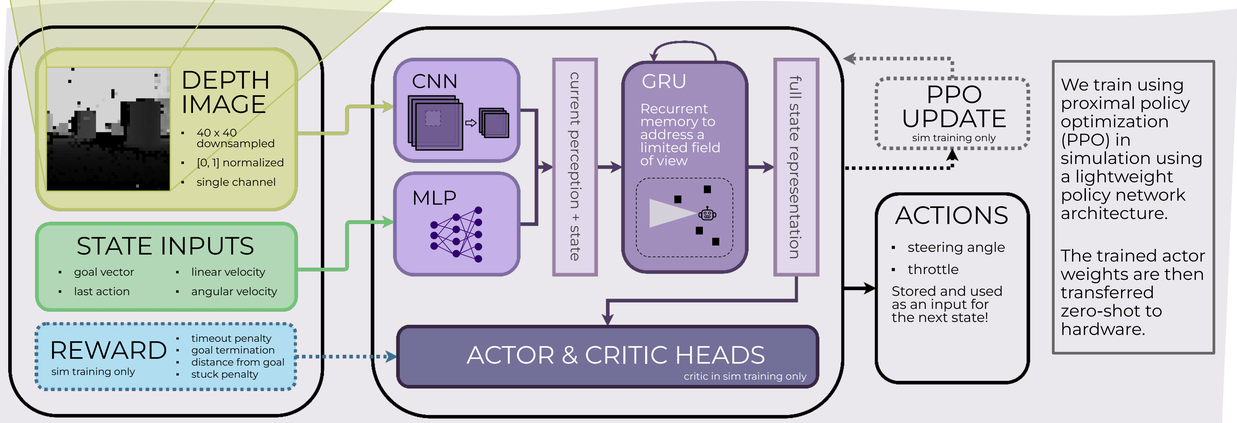

Results

Simulated Deployment

Policy Type	Success Rate†
Baseline (height map)	48.5%
MLP-only	20.0%
CNN-only	54.0%
CNN + GRU	84%
CNN + GRU, with simulated visual noise	79.4%

† relative frequency of navigating to a goal zone in under 10 seconds, without getting stuck or rolling over

Real world deployment:

- Baseline policy failed to reach the goal
- CNN-only policy bumped into unseen obstacles
- CNN+GRU policy's steering oscillated in response to noise
- ***Despite decreased sim performance***, the final policy smoothly evaded obstacles in real deployment

Real-world deployment of the final policy with snapshots of the depth camera view

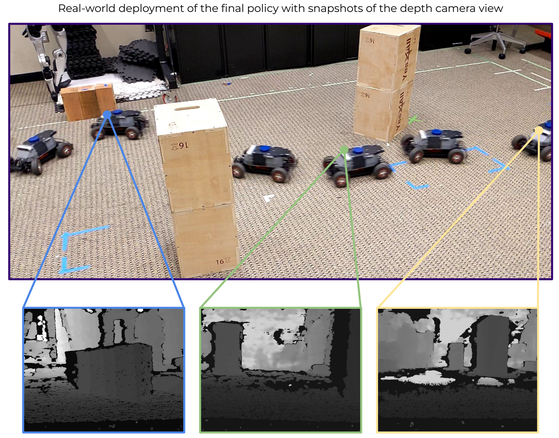

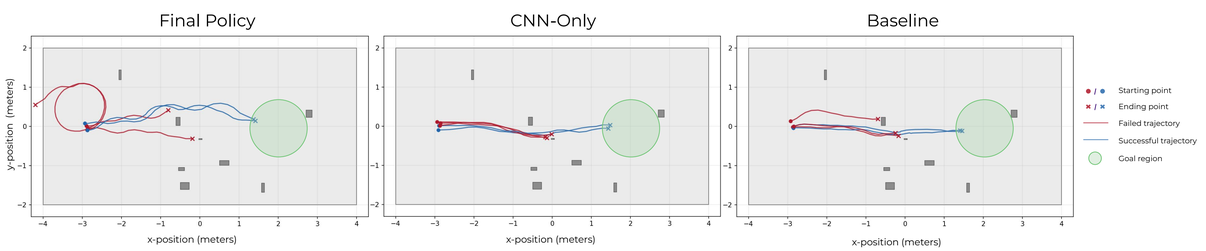

Implementation

REWARDS

- **Goal reached** — rewards and terminates upon reaching the goal
- **Goal distance** — linearly penalizes distance from the goal
- **Stuck** — penalizes and terminates upon being unable to move
- **Time out** — penalizes and terminates after 10 seconds

NOTES

- To avoid gradient starvation, we balanced vision and state latents as input to the GRU

- Some curriculum learning stabilized training but often degraded policy

- Recurrent policy training generally benefited from sparser rewards

Future Work

We have demonstrated the feasibility of an end-to-end terrain-aware navigation policy on a fully low-cost pipeline. However, we highlight some potential improvements:

Issue Faced	Potential Change
Slow training speed	Asymmetric actor-critic
Inconsistent policy behavior in the real	Alternative visual encoders
Instability with recurrent policies	Further iteration + other recurrent networks

Acknowledgements

Special thanks to Tyler Han, the Robot Learning Lab, and the CSE 494 staff for their mentorship and support!

[1] T. Han, P. Shah, S. Rajagopal, Y. Bao, S. Jung, S. Talia, G. Guo, B. Xu, B. Mehta, E. Romig, R. Scalise, and B. Boots. Wheeled lab: Modern sim2real for low-cost, open-source wheeled robotics. In Proceedings of the Conference on Robot Learning (CoRL), August 2025. URL http://arxiv.org/abs/2502.07380.

[2] LycheeAI and J. Zhou. Full Isaac Lab RL Project with WheeledLab, 2026. URL https://www.youtube.com/watch?v=caVh-2v4xz4. YouTube video. Accessed: 2026-05-27.

[3] S. S. Srinivasa, P. Lancaster, J. Michalove, M. Schmittle, C. Summers, M. Rockett, R. Scalise, J. R. Smith, S. Choudhury, C. Mavrogiannis, and F. Sadeghi. MuSHR: A low-cost, open-source robotic racecar for education and research. IEEE Robotics and Automation Letters (RA-L), 8(12):7917–24, December 2023. URL http://arxiv.org/abs/1908.08031.