Video results : FPV of quadrotor flight

Shut up and show me the code

Autonomous Quadrotor Flight in Simulation using RL

Ratnesh Madaan, Dhruv Mauria Saxena, Rogério Bonatti, Shohin Mukherjee

The Robotics Institute

Carnegie Mellon University

Motivation

- Despite advancements in sensing technologies, it is difficult to develop robust systems by separating perception and control.
- Learning to fly in the real world is impractical since it is time consuming and expensive.
- Learning to fly in simulation opens the possibility of transferring learned policies to the real world.

Environment

- Environment consists of randomly places cylindrical obstacles, simulated and rendered in Gazebo.
- Position of cylinders changes for each episode.
- Quadrotor is equipped with a planar laser rangefinder and a front-facing RGB-D camera.

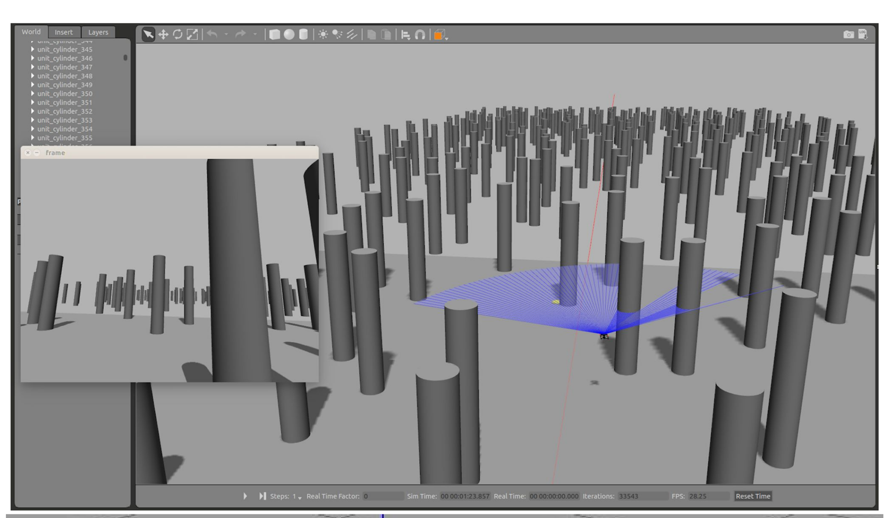

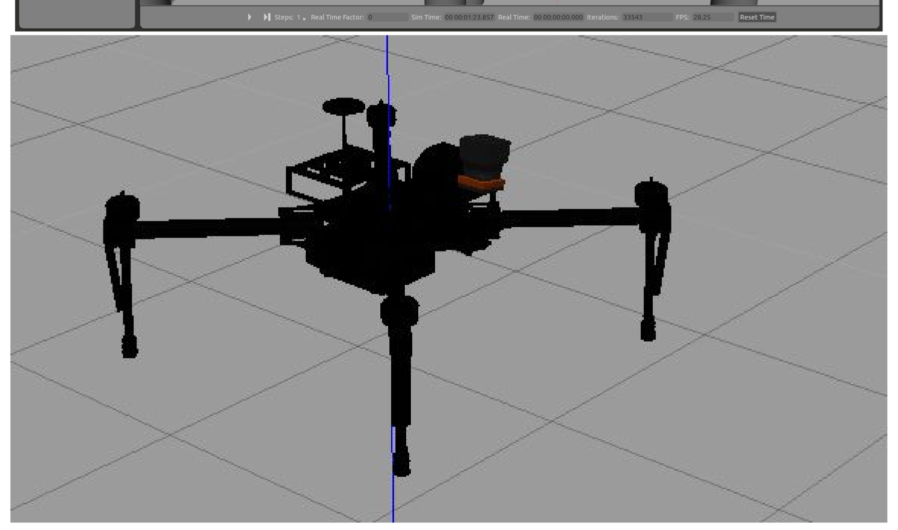

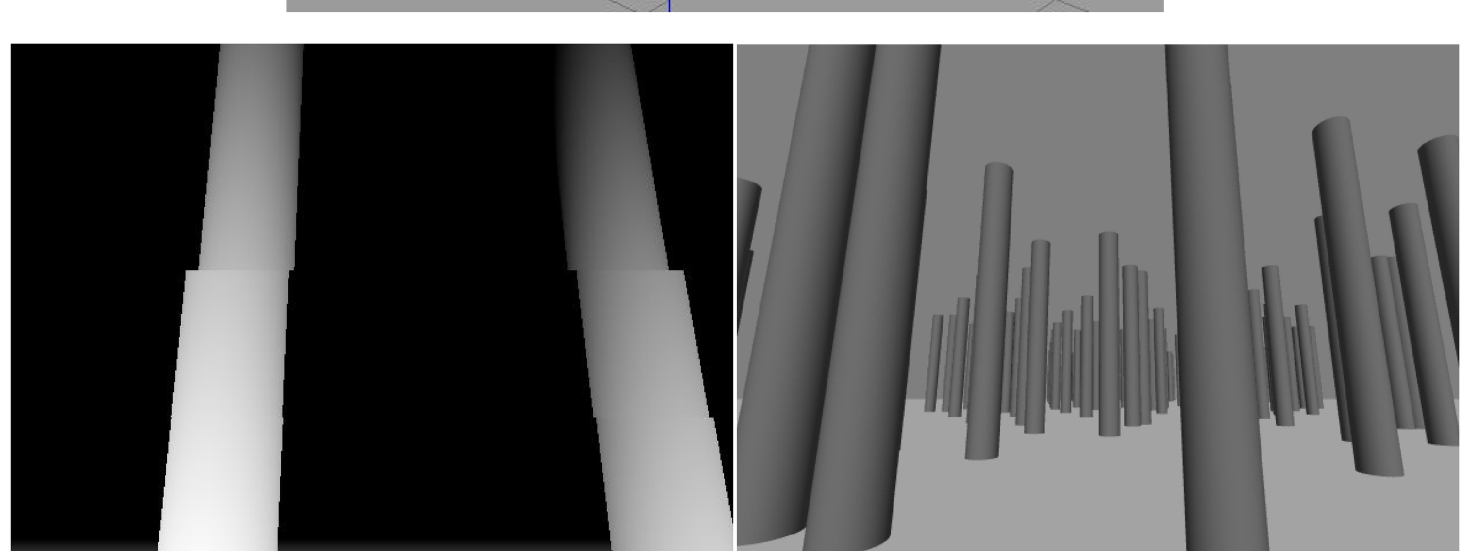

Objectives

- Develop an open-source Gazebo environment integrated with Gym which can be used by the community for reinforcement learning research.
- Train a deep Q-network capable of flying a drone autonomously in the Gazebo environment.

Learning On Images

Network architecture:

- Input: 84 X 84 X 4 images
- Conv layer 1: 32, 8 X 8 filters, stride 4
- Conv layer 2: 64, 4 X 4 filters, stride 2
- Conv layer 3: 64, 3 X 3 filters, stride 1
- Fully-connected layer: 512 units
- Output: 9 units (correspond to yaw angles)

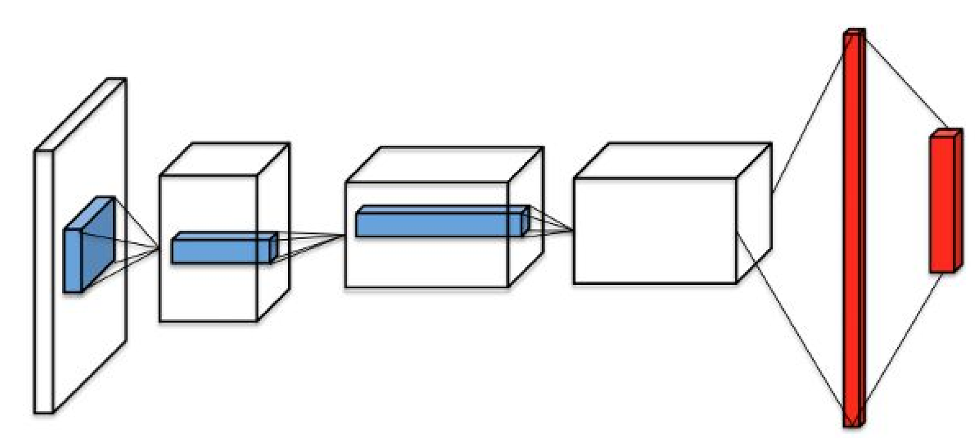

Learning On Laser Data

Network architecture:

- Input: 70 X 4 array
- Fully-connected layer: 512 units
- Fully-connected layer: 512 units
- Output: 9 units (number of actions)

Learning Algorithm - DQN

- Q_{w^-} : Target Network
- Q_w : Online Network

$$w := w + \alpha\left(r + \gamma \max_{a' \in A} Q_{w^-}(s', a') - Q_w(s, a)\right) \nabla_w Q_w(s, a)$$

Partial Results

Graphs for grayscale, monocular camera images

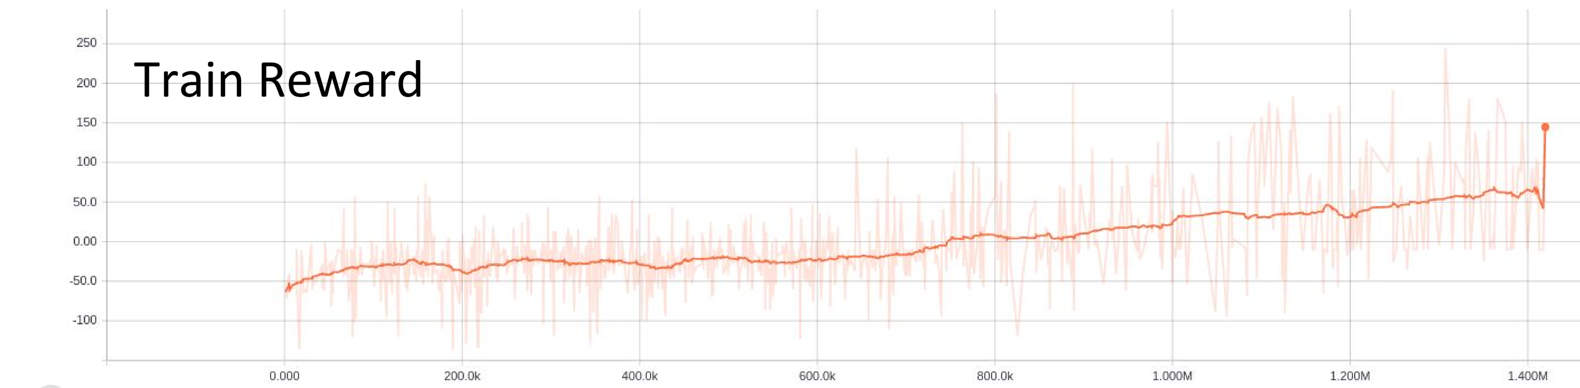

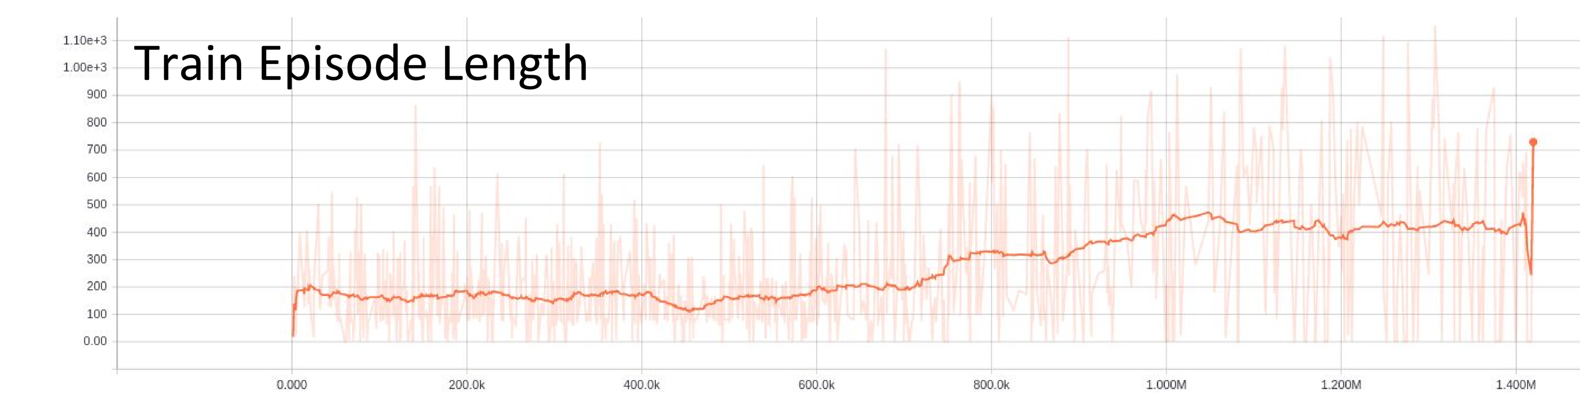

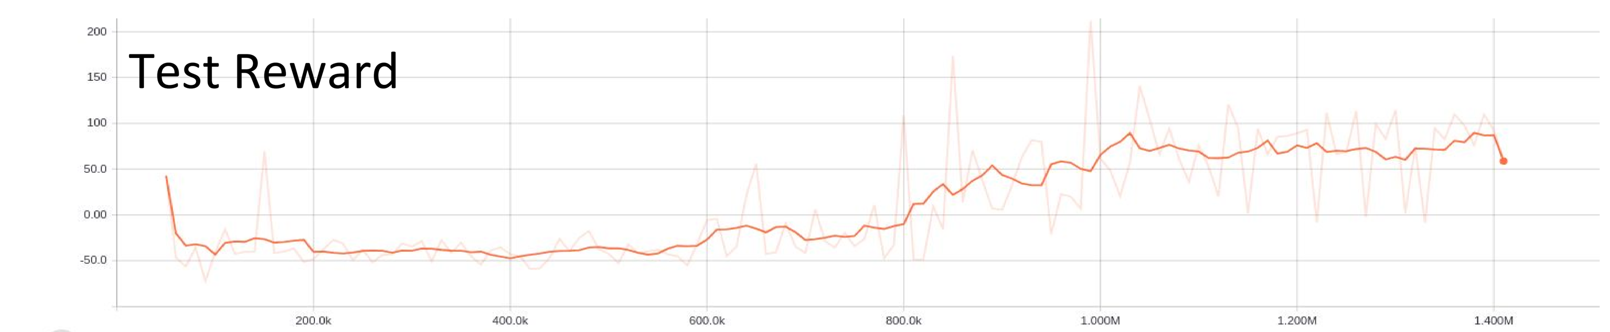

Conclusion And Future Work

- With less than 1M iterations, we still could not observe significant learning in the environments with laser and depth images as inputs.
- Learning on depth images and laser data can be more easily transferrable to real-life applications.
- During the summer we plan to test real quadcopters flying with the policies learned in simulation.

References

- [1] Mnih, Volodymyr, et al. "Playing atari with deep reinforcement learning." arXiv preprint arXiv:1312.5602 (2013).
- [2] F. Sadeghi and S. Levine, "(cad) 2 rl: Real single-image flight without a single real image," arXiv preprint arXiv:1611.04201, 2016