Generalized Robot Learning from One Scenario

Nikitas Klapsis, ndklapsis@gmail.com
Crescenta Valley High School, Class of 2022
USC Viterbi Department of Computer Science, SHINE 2020

Introduction

A major bottleneck to effective robot learning is collecting and processing enough data to solve complex tasks. The robot must interact with example scenarios many times to solve each scenario, and must be presented with a diversity of scenarios in order to generalize. We asked the following research question:

Can a robot interacting with only **one** scenario still learn to solve complex tasks?

Methods

To solve the general task, the robot must extract as much information as possible from the single scenario. We evaluate our approach on the lunar lander domain from OpenAI Gym where the objective is to land a spacecraft safely. Each scenario differs in terrain and initial conditions of the lander. In order to generalize to all scenarios, we want our robot to explore **different** solutions to landing the lander such as impacting the ground at different velocities or landing at different positions. To explore different solutions we trained our robot with the quality diversity algorithm Covariance Matrix Adaptation MAP-Elites (CMA-ME).

Research and Results

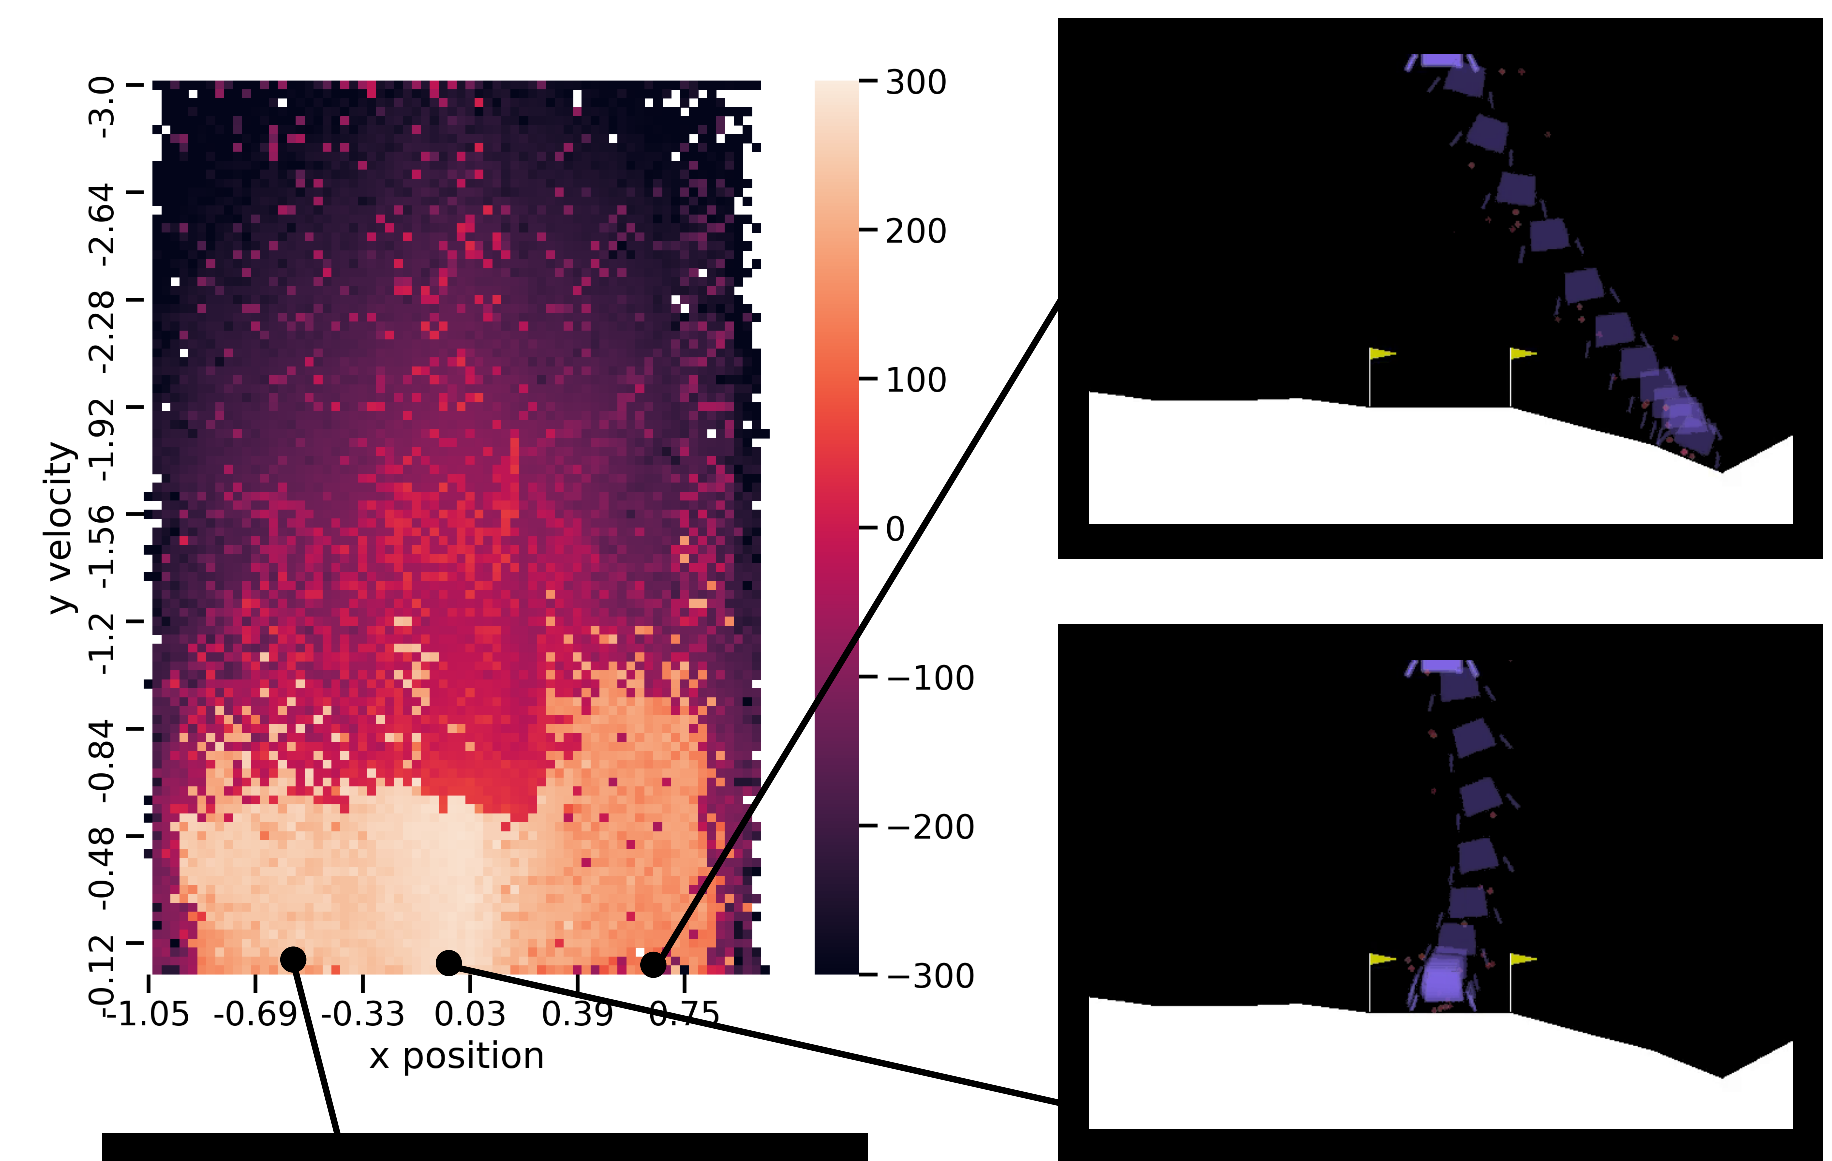

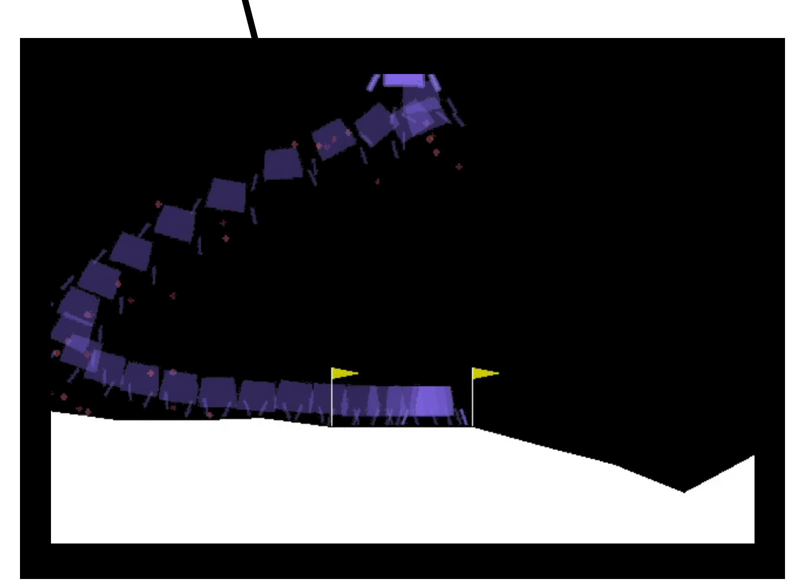

Figure 1: Each cell represents a solution generated from CMA-ME on the training scenario. Intensity marks solutions with high rewards. QR codes link to videos of the lander.

Figure 2: Each cell represents the average score from 200 different environments generated from the model made by CMA-ME for the first training environment.

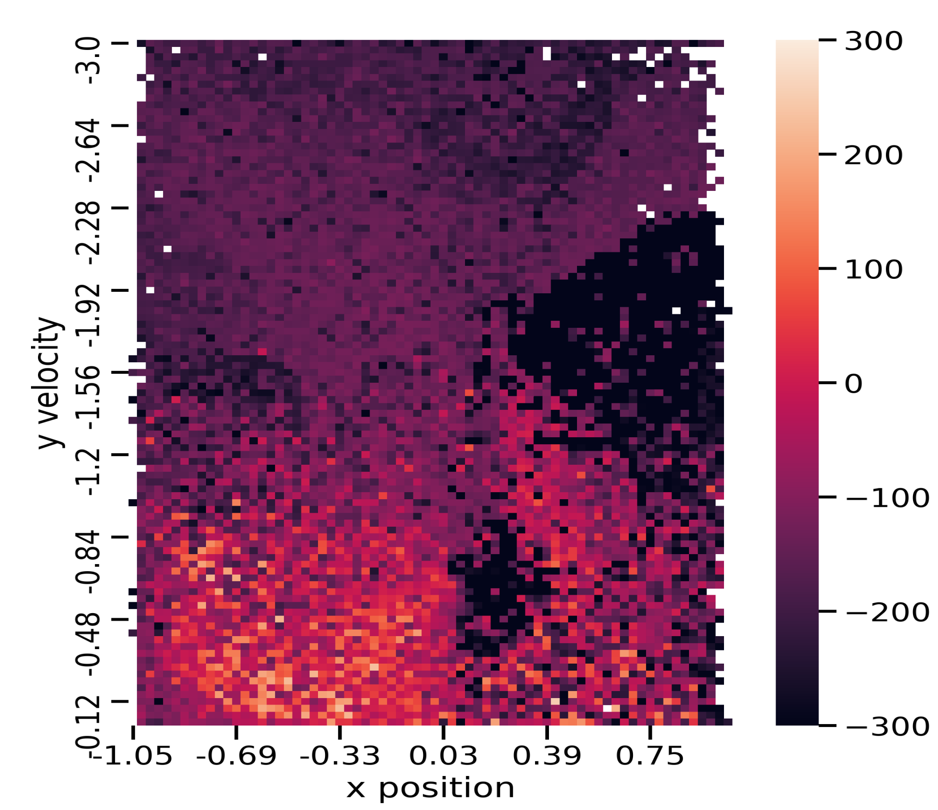

Next Steps and Broader Impacts

The next steps in our project would be seeing if this method also worked in a different OpenAI Gym environment like car racing and bipedal walking. Furthermore, by implementing a neural network instead of a linear model, our method could solve more complex tasks. After demonstrating the generality of our method, we could train physical robots with fewer example scenarios.

A possible application for this method of learning is to use it where the data from an environment is extremely limited, like rovers on different planets. Diverse solutions for both design and control of rovers could lead to greater chances of mission success, and more robust exploration of planets. The applications for CMA-ME are vast and continue to grow as machine learning becomes an essential part of research and daily life.

Acknowledgements

I would like to thank Professor Nikolaidis for giving me the opportunity to work in his lab this summer. I would also like to thank my Ph.D. mentor, Matthew Fontaine, for guiding me through every step of the process of learning about machine learning and its applications with quality diversity, my lab mate Ruth Berkun, and the entire SHINE team for working through the pandemic and persevering to still educate all of us.