Amazon DeepRacer

Ira A. Fulton Schools of Engineering
Arizona State University

Ryan Roberti, Saket Shanbhag, Faisal Alghamdi, Alex Beardsley, Abdulaziz Ben Salim
Mentor: Trevor Thornton

Acknowledgments:
Glen Uehara/General Dynamics,
Todd Bevins/AWS,
Skysong CIC

Problem Statement

Autonomous systems have been developed and used increasingly over time and have proved successful in various engineering solutions. The team chose to explore reinforcement machine learning as a method to develop an autonomous driving system. The objective of the Amazon DeepRacer project was to design, train, and evaluate a reinforcement learning model capable of successfully navigating the DeepRacer vehicle around the track in the shortest time possible.

Concept and Theory

Reinforcement learning is a type of machine learning that uses reward-based training to incentivize desired behaviors. The agent (vehicle) receives its state from the environment using the images from its camera. It then takes an action and is rewarded based on the definition of the reward function. The reward function is designed to reward actions that help the vehicle complete a lap in the shortest time possible while staying on track. The agent updates its policy by adding actions that maximize cumulative reward. A reinforcement learning model consists of three key components: state (input), policy, and action (output).

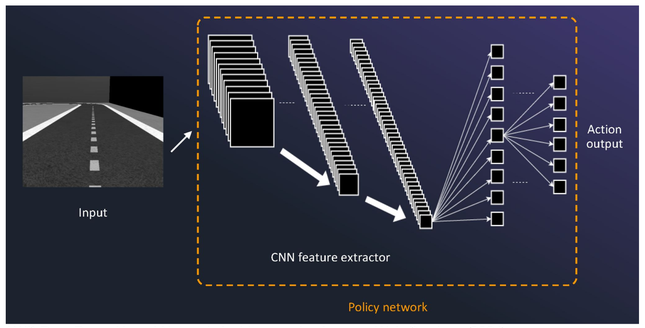

AWS DeepRacer Neural Network Architecture by Amazon Web Services (2024)

Simulation Results

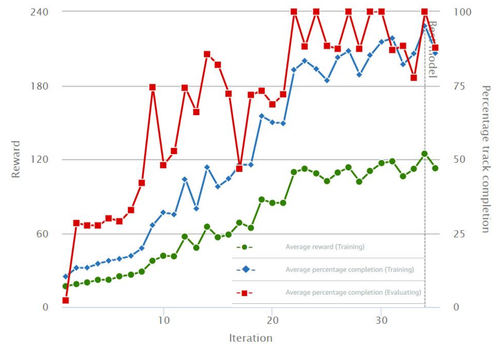

Reward Graph of Ryan-Model3

Evaluation Results of Ryan-Model3

Trial	Time (MM:SS.mmm)	Trial results (% track completed)	Status	Off-track	Off-track penalty
1	00:10.380	100%	Lap complete	0	-
2	00:11.884	100%	Lap complete	0	-
3	00:11.746	100%	Lap complete	0	-

The models were trained and evaluated through the DeepRacer console. Log analysis was used to analyze training and make improvements to the models. The best performing model achieved an average lap time of 11.34 seconds over three laps with zero off-tracks. The reward graph and evaluation results of the model are shown above and a heatmap of reward values obtained from log analysis is shown below.

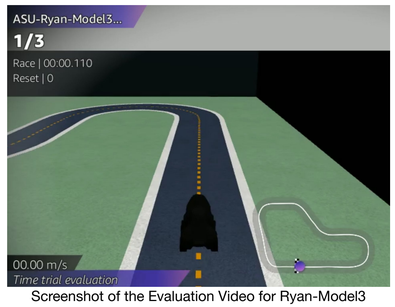

Screenshot of the Evaluation Video for Ryan-Model3

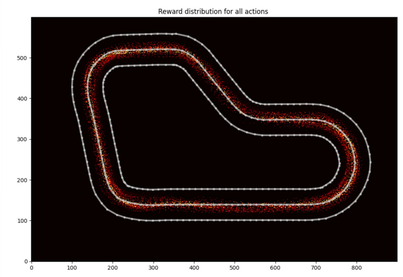

Heatmap of Rewards for Ryan-Model3 from Log Analysis

Testing Results

Testing Day Results of Ryan-Model3

Trial	Time (MM:SS.mmm)	Trial results (% track completed)	Status	Off-track	Off-track penalty
1	00:11.21	100%	Lap complete	0	-
2	00:09.96	100%	Lap complete	0	-
3	00:09.93	100%	Lap complete	0	-
4	00:10.93	100%	Lap complete	0	-
5	00:09.26	100%	Lap complete	0	-

The top-performing model achieved an average lap time of 10.26 seconds over five laps with zero off-tracks. The simulated and real-world results were closely correlated. The real-world vehicle achieved a slightly faster average lap time than the simulated vehicle.

DeepRacer on Track at Testing Day 2

Summary

The team successfully designed, trained, and evaluated a reinforcement learning model capable of keeping the DeepRacer on track while achieving fast lap times. Using the DeepRacer console and a physical vehicle and track, the team successfully created models capable of navigating simulated and real-world environments. The simulation and real-world results were heavily correlated, demonstrating the effectiveness of training in a simulated environment for real-world running.