

Comparative Analysis of Reinforcement Learning Algorithms for Autonomous Driving

Authors : Migena Mana, Ahmed Khalid Syed, Abdul Malik, Nikhil Cherian

Abstract : In recent years, advancements in deep learning enabled researchers to tackle the problem of self-driving cars. Car companies use huge datasets to train their deep learning models to make autonomous cars a reality. However, this approach has certain drawbacks in that the state space of possible actions for a car is so huge that there cannot be a dataset for every possible road scenario. To overcome this problem, the concept of reinforcement learning (RL) is being investigated in this research. Since the problem of autonomous driving can be modeled in a simulation, it lends itself naturally to the domain of reinforcement learning. The advantage of this approach is that we can model different and complex road scenarios in a simulation without having to deploy in the real world. The autonomous agent can learn to drive by finding the optimal policy. This learned model can then be easily deployed in a real-world setting. In this project, we focus on three RL algorithms: Q-learning, Deep Deterministic Policy Gradient (DDPG), and Proximal Policy Optimization (PPO). To model the environment, we have used TORCS (The Open Racing Car Simulator), which provides us with a strong foundation to test our model. The inputs to the algorithms are the sensor data provided by the simulator such as velocity, distance from side pavement, etc. The outcome of this research project is a comparative analysis of these algorithms. Based on the comparison, the PPO algorithm gives the best results. When using PPO algorithm, the reward is greater, and the acceleration, steering angle and braking are more stable compared to the other algorithms, which means that the agent learns to drive in a better and more efficient way in this case. Additionally, we have come up with a dataset taken from the training of the agent with DDPG and PPO algorithms. It contains all the steps of the agent during one full training in the form: (all input values, acceleration, steering angle, break, loss, reward). This study can serve as a base for further complex road scenarios. Furthermore, it can be enlarged in the field of computer vision, using the images to find the best policy.

Keywords : autonomous driving, DDPG (deep deterministic policy gradient), PPO (proximal policy optimization), reinforcement learning

Conference Title : ICCV 2019 : International Conference on Connected Vehicles

Conference Location : Paris, France

Conference Dates : September 19-20, 2019