Deep Reinforcement Learning Traffic Signal Controller

Osvaldo Castellanos
osvaldo.castellanos01@utrgv.edu

2/08/2019

Summary of the Proposal

In this project, we aim to apply multi-agent deep reinforcement learning techniques to traffic signal controllers. The project is set up to investigate whether optimal policies exist for traffic signal controllers in adaptive environments. Additionally, we are interested analyzing the effects that different reward designs have on the performances of multiple learning algorithms.

Background

We've developed a traffic simulator that allows us to test a traffic signal control agent in multiple environments. In addition to evaluating our algorithm to the state of practice simulators, we will investigate how multi-agent reinforcement learning algorithms can train a traffic signal agent to adapt to different events that cause traffic. For example, events such as cars stalling, accidents, and natural disasters like flooding. Our simulator gives us control over parameters that can simulate these conditions. We can specify environments with multiple roads, intersections, traffic lights, and specify if zones aren't accessible to drivers. Thus, given an environment where an event has taken place that causes traffic, the traffic signal controller can coordinate routes to circumvent traffic-heavy areas.

Current research in traffic signal control that also applies multi-agent reinforcement learning train agents in specific, or even single, environments. Our goal is to train a system that can be practical in multiple settings, under multiple conditions. Overall, in this project, our goal is to apply the system to simulations of large-scale areas such as the Rio Grande Valley in South Texas.

A summary of the main points that will distinguish this project from previous work:

- State-of-the-art approaches to this problem have investigated optimal policies under simplified environments, such as a single intersection. We will start with a similar environment and expand to more realistic environments.
- Recent progress has been extended to include environments with multiple intersections. However, to our knowledge, there has been no successful multi-agent traffic signal controller.

Goal and Objectives

The goal of this research is to analyze the optimal policies a traffic signal controller has during multiple environments. The goal will be met by achieving the following objectives:

- Formally define the traffic signal optimization problem as a Markov Decision Process.
- Run baseline learning algorithms such as DQN, DDPG, A2C, and more, on a simplified environment. Once performances have been recorded, the next step will be to expand the environment to include 2 to 4 intersections.
- The final goal will be to run multi-agent learning algorithms on the multiple inter-sectional environment and analyze the performances in order to identify areas of improvements.

Data and Methods

We will generate data through a traffic simulator. The simulation will provide a model of cars driving through an intersection that is being controlled by a traffic signal.

- The simulator can provide the position, speed, and velocity of the vehicles. Each feature of the data will be provided in a matrix.
- The current traffic light will be provided as an integer that encodes a specific light sequence.
- We will also record the waiting time of each car at a red light. This will help us define the average wait time, which is one of the reward signals we will be using.

References

- Juntao Gao, Yulong Shen, "Adaptive Traffic Signal Control: Deep Reinforcement Learning Algorithm with Experience Replay and Target Network,"
- Marco Wiering, "Multi-Agent Reinforcement Learning for Traffic Light Control,"
- Seyed Sajad Mousavi, Michael Schukat, Enda Howley, "Traffic light control using deep policy-gradient and value-function-based reinforcement learning,"
- Xiaoyuan Liang, Xusheng Du, "Deep Reinforcement Learning for Traffic Light Control in Vehicular Networks,"