Reinforcement Learning for Autonomous Intersection Management

Jack Lowe, Master of Arts in Engineering

University of Dublin, Trinity College, 2022

Supervisor: Melanie Bouroche

This project consists of the investigation of data needs for efficient Q-Learning for distributed Autonomous Intersection Management (AIM). Specifically, this entails the design and implementation of a lightweight Q-Learning framework that uses Connected Autonomous Vehicles (CAVs) as agents to decrease the delay in an unsignalised intersection. The literature review concludes that while Q-Learning is and has been an efficient tool in the design of scheduling policies for traffic light control, its potential for distributed and unsignaled AIM remains globally unresolved, specifically due to the neglect of current implementations to consider the materialistic need for Q-functions to be unencumbered by excessive data if they are to be stored locally inside CAVs. By minimising both agents and observations in our Q-Learning model, this project seeks to validate the hypothesis that minimal-data Q-functions are viable solutions to AIM scheduling policies considering the complexity of modern AIM modelling.

The results show that while the framework is capable of performing Q-learning, the fine line between an excessive number of states and sufficient agent information was not met. Various choices of observations and actions are nevertheless explored and the consequences they have on performance are analysed.