Teaching machines to direct traffic through deep reinforcement learning

August 17 2016

Rush hour—the dreaded time of day when traffic conditions seem bent on making you late. As your car slowly creeps in line behind countless others stuck at a stop light, you think to yourself, "Why aren't these lights changing faster?" Traffic control scientists have long tried to solve this signaling problem. Unfortunately, the complexity of traffic situations makes the job extremely hard. A recent study suggests that machines can learn how to plan traffic signals just right to reduce wait times and make traffic queues shorter.

Automating traffic control is notoriously tricky because it involves two challenging tasks: modeling traffic flow and then optimizing it. While traditional artificial intelligence and simulation techniques can learn and reproduce the dynamics of certain traffic situations, they cannot easily determine the best course of action. Despite this shortcoming, one emerging AI method that has shown promise is reinforcement learning.

Much like the reward system in our brain, reinforcement learning algorithms operate by determining what set of actions are most beneficial to a system in a given state. In video games, a popular proving ground for this type of algorithm, these steps are the strategic moves a player must make to earn the highest possible score. In traffic signaling, it's the pattern of traffic light switches that minimizes the time drivers spend waiting in a queue and/or the length of that queue. The problem with this algorithm is that the dataset it must comb through becomes exponentially larger as traffic conditions and ways of controlling traffic become increasingly complex. This is where another type of AI method called deep learning comes into play.

Inspired by how the human brain works, deep learning uses special algorithms called neural networks to search for hidden patterns in sets of data. When combined with reinforcement learning, this technique effectively reduces the space a machine must search through as it seeks the best solution. This greatly reduces the required computational time.

When deployed on a virtual four-way traffic intersection with four lanes heading east-west and four heading north-south, the deep reinforcement learning algorithm outperformed a conventional reinforcement learning algorithm. There were shorter lines. There was a better balance of traffic in both directions. And over the course of a full day, more than 1000 fewer vehicles came to a full stop. In the end, the average delay was cut by 14 percent, with vehicles spending 13 seconds less in traffic during peak morning hours.

Deep reinforcement learning algorithms aren't quite ready to hit the streets. But they are helping make real-world traffic scenarios a lot less complicated for researchers, which could pave the way toward a new understanding of how traffic flows work.

More information: DOI: 10.1109/JAS.2016.7508798

Provided by Chinese Association of Automation

provided for information purposes only.