Engineering Applications of Artificial Intelligence 52 (2016) 203–212

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

Traffic flow optimization: A reinforcement learning approach

Erwin Walraven [a,*], Matthijs T.J. Spaan [a], Bram Bakker [b]

[a] Delft University of Technology, Mekelweg 4, 2628 CD Delft, The Netherlands
[b] Cygnify BV, Bargelaan 200, 2333 CW Leiden, The Netherlands

ARTICLE INFO

Article history:
Received 19 June 2015
Received in revised form
2 December 2015
Accepted 4 January 2016

Keywords:
Traffic flow optimization
Traffic congestion
Variable speed limits
Reinforcement learning
Neural networks

ABSTRACT

Traffic congestion causes important problems such as delays, increased fuel consumption and additional pollution. In this paper we propose a new method to optimize traffic flow, based on reinforcement learning. We show that a traffic flow optimization problem can be formulated as a Markov Decision Process. We use *Q*-learning to learn policies dictating the maximum driving speed that is allowed on a highway, such that traffic congestion is reduced. An important difference between our work and existing approaches is that we take traffic predictions into account. A series of simulation experiments shows that the resulting policies significantly reduce traffic congestion under high traffic demand, and that inclusion of traffic predictions improves the quality of the resulting policies. Additionally, the policies are sufficiently robust to deal with inaccurate speed and density measurements.

1. Introduction

Traffic congestion is a problem many people are faced with almost every day, causing not only delays, but also pollution and increased fuel consumption. In the United States, travel delay increased from 1.1 billion hours in 1982 to 5.5 billion hours in 2011 (Schrank et al., 2012). At the same time, the amount of wasted fuel increased from 1.9 billion liters to 11 billion liters. The total congestion costs were 121 billion dollars in 2011. Expanding the road network to increase its capacity would be a straightforward solution, but this is not always feasible in practice because of space and budget limitations.

Rather than increasing the capacity of the road network to reduce congestion, variable message signs can be installed, which communicate speed limits to car drivers. These speed limits can be adjusted depending on the current traffic conditions, and such speed limits have a positive influence on traffic flow (Papageorgiou et al., 2008). Although several control algorithms for variable message signs have been developed, most of these approaches are reactive in the sense that speed limits are assigned when congestion is actually detected. In this paper we show that artificial intelligence techniques can play an important role in realizing proactive control for assigning speed limits.

The application of artificial intelligence in the area of traffic and transportation can be motivated by observing that several systems in these areas start to rely more on autonomous and intelligent decision making. For example, autonomously driving vehicles (Bai et al., 2015) need to be able to determine an appropriate driving speed, taking into account the distance to other objects and potential congestion in case the vehicle density on the road increases. Intelligent lane keeping methods need to be able to reason about vehicles in the vicinity, such that appropriate control actions are taken (Lu et al., 2007). Such applications need to decide autonomously how problems are solved, need to respond to a dynamic and changing environment and should be adaptive in the sense that systems should continuously reflect the preferences of their users. All these characteristics motivate the introduction of intelligent software systems in this domain, also known as intelligent agents (Jennings and Wooldridge, 1998). Proactively assigning speed limits to highways also requires an adaptive control system which dynamically responds to sensor data and traffic predictions. In this paper we take a step in this direction by presenting a reinforcement learning algorithm which automatically learns when speed limits should be assigned to highways to reduce congestion, based on the characteristics of the highway as well as demand volumes occupying the highway and predictions regarding future traffic conditions. The resulting algorithm is able to learn proactive control rules in a highly complex domain, and shows that reinforcement learning has the potential to address traffic congestion problems.

The main contributions of our paper can be summarized as follows. First, we formulate a traffic flow optimization problem as a Markov Decision Process (Puterman, 1994), and we show that *Q*-learning (Watkins, 1989) can be applied to find policies dictating how speed limits should be assigned to highway sections to

* Corresponding author. Tel.: +31 15 27 86206.
E-mail address: e.m.p.walraven@tudelft.nl (E. Walraven).

http://dx.doi.org/10.1016/j.engappai.2016.01.001
0952-1976/

reduce traffic congestion. Second, we show how traffic predictions can be included in our method. Third, we discuss how artificial neural networks (Haykin, 1999) can be used to approximate policies defining the speed limits. Using simulations we show that our methods are able to reduce travel time and congestion under high traffic demand in small road networks, and the policies are sufficiently robust to deal with inaccurate speed and density measurements.

The structure of the paper is as follows. Section 2 discusses related work in the area of artificial intelligence, focusing on problems in traffic and transportation domains. Section 3 introduces the traffic flow optimization problem under consideration. Section 4 provides background information regarding traffic flow modeling and reinforcement learning. The traffic flow optimization problem is formulated as a Markov Decision Process in Section 5 and Section 6 introduces the reinforcement learning algorithm which can be used to obtain policies. Section 7 describes a series of experiments and Section 8 provides a discussion of our work. Section 9 summarizes our results and conclusions.

2. Related work

In this section we give an overview of related work that applies artificial intelligence methods to solve problems in traffic and transportation domains.

2.1. Traffic flow control using artificial intelligence methods

Artificial intelligence has been applied to regulate the number of vehicles entering a highway. For example, ramp metering devices can be controlled using reinforcement learning algorithms (Fares and Gomaa, 2014, 2015). Similar to our work, optimization of traffic flow is formulated as a sequential decision making problem, but these methods explicitly aim to keep the density of the vehicles close to the critical density, such that flow is optimized. Another method to control ramp metering devices with *Q*-learning can be found in work by Rezaee et al. (2012), where the number of vehicles passing a loop detector near an on-ramp is optimized. Similar work by Davarynejad et al. (2011) takes queue lengths of on-ramps into account during learning, which is not considered in our work. In contrast to ramp metering approaches, we control the speed of vehicles rather than controlling the number of vehicles entering a highway.

Neural networks have been applied in existing work to create controllers for traffic lights (Spall and Chin, 1994; Wei and Zhang, 2002). The road networks used for learning were relatively simple and demand volumes were assumed to be fixed. Neural networks have also been used to predict the exit demand of highways (Kwon and Stephanedes, 1994). Although the application of neural networks is related to our work, the authors did not combine neural networks with reinforcement learning.

Optimization of traffic lights in the urban area is also important to reduce congestion. Kuyer et al. (2008) discuss coordination of traffic lights in the urban area using reinforcement learning, where traffic lights are intelligent agents that coordinate their behavior. A multiagent reinforcement learning controller for traffic lights with multiple objectives is discussed in work by Khamis and Gomaa (2012, 2014), where the method minimizes travel time, increases safety and controls speed in such a way that less fuel is consumed. When modeling control of traffic lights as multiagent system, the system may become complex when scaling up to a large number of intersections with cooperative traffic lights. Abdoos et al. (2013) discuss a hierarchical control method to address this problem. A survey by Zhao et al. (2012) discusses intelligent solutions for control of traffic lights, such as neural networks, and acknowledges that additional research is needed for traffic control using intelligent systems.

2.2. Reinforcement learning for other transportation problems

Besides controlling vehicles from an infrastructural point of view (e.g., traffic lights and speed limits), reinforcement learning has been applied to make routing decisions to guide vehicles through a city (Zolfpour-Arokhlo et al., 2014). Similar to our work, significant travel time reductions can be obtained using reinforcement learning algorithms. Another traffic-related application of reinforcement learning can be found in the area of air traffic management (Tumer and Agogino, 2007), where major delay reductions can be realized. A more general application of reinforcement learning in this domain can be found in work by Cruciol et al. (2013), where different reward functions are investigated for decision-making in air traffic flow management with several stakeholders. In transportation and logistics, reinforcement learning has been applied to control ship unloading (Scardua et al., 2002). Similar to our work, control policies are learned using Markov Decision Processes, *Q*-learning and neural networks.

3. Traffic flow optimization problem

The problem of traffic congestion on highways occurs if the demand volume exceeds the highway capacity (Papageorgiou et al., 2003). Consequently, the density of the vehicles exceeds the critical density of the highway, the distance between vehicles decreases, and a lower speed is necessary to preserve safety. Furthermore, a congested highway leads to vehicular queueing, an increased travel time, increased cost, more fuel consumption and additional pollution in the environment. If an highway is not affected by congestion, then the flow on the highway is called free flow.

The delay incurred by vehicles as a consequence of congestion can be measured by computing the number of vehicle hours (Zhang et al., 2006). This is a metric that directly relates congestion to the travel time of vehicles. One vehicle hour represents one vehicle driving on a highway for 1 h, but can also be interpreted as 60 vehicles driving on a highway for 1 min. By aggregating over all vehicles, the total number of vehicle hours can be obtained. A similar metric is the total vehicle delay time, which is the total additional travel time compared to the travel time in case of free flow. For instance, if the travel time of a vehicle is 60 min in case of congestion, instead of 45 min in free flow, then its delay time is 15 min.

As mentioned in the introduction, expansion of the road network is not always feasible, and therefore other solutions are required. An example solution is assigning speed limits, which has shown to be able to reduce congestion on highways (Zhang et al., 2006). The problem to assign speed limits to highways is visualized in Fig. 1, where the gray area represents a congested area near the on-ramp and arrows indicate the direction of the vehicle flow. If the traffic demand volume of the on-ramp is high, speed limits can be assigned to upstream sections to reduce congestion. Assigning speed limits is not straightforward, however, because it is difficult to decide when speed limits should be issued and in which parts of the highway speed should be reduced. Additionally, speed limits should be increased and decreased smoothly to preserve safety and alternating sequences of speed limits should be prevented.