Available online at www.sciencedirect.com

ScienceDirect

IFAC PapersOnLine 50-1 (2017) 290–296

Optimal Driving Policies for Autonomous Vehicles Based on Stochastic Drift Counteraction

Robert A. E. Zidek and Ilya V. Kolmanovsky

University of Michigan, Department of Aerospace Engineering, Ann Arbor, MI 48109 USA (e-mails: robzidek@umich.edu, ilya@umich.edu).

Abstract: This paper focuses on high-level control and decision-making for an autonomous car. We develop a hybrid probabilistic model that describes the motion of a car and its surrounding traffic on a two-lane highway/road, where the acceleration of the car and the lane changes serve as control variables. Using approximate dynamic programming (ADP) techniques and an enhanced version of the value iteration algorithm, a control policy is obtained that maximizes the expected time that the car maintains a prescribed minimum (safe) headway. Simulation results for different settings are provided to validate the approach.

Keywords: Autonomous vehicles, stochastic optimal control, stochastic hybrid systems.

1. INTRODUCTION

Autonomous (driverless) cars may improve road safety, provide greater convenience for humans, and lower cost of transportation in the future. While feasibility of autonomous driving has already been demonstrated, see, for example, Markoff (2013), advanced control algorithms need to be developed to achieve high levels of safety and ride comfort. We approach this objective by developing a model-based control strategy for autonomous driving that maximizes the time that a prescribed minimum headway relative to the in-front vehicle is maintained.

Our approach is based on a hierarchical control structure, where a high-level controller provides optimal decision-making and low-level controllers execute the decisions by regulating the longitudinal and lateral motion of the car. While the focus of this paper is on high-level control, relevant low-level controllers were discussed, for example, by Hatipoglu et al. (2003), Guo et al. (2014), and Hu et al. (2016). One of the earliest studies on modeling decision-making in driving can be found in Worrall et al. (1970) who used Markov chains calibrated from real traffic data. Gipps (1986) proposed a set of deterministic rules for lane-changing decisions for cars traveling at a constant velocity. More complex probabilistic approaches were developed, for example, by Wu et al. (2000), Toledo et al. (2007), and Schubert et al. (2010). Wang et al. (2015) proposed a game theoretic approach for lateral and longitudinal decision-making, where the control problem was decomposed into car following and lane-changing sub-problems in order to minimize a cost function.

The contribution of this paper is a novel approach based on *stochastic drift counteraction optimal control* (SDCOC), see Kolmanovsky et al. (2008). SDCOC was applied to the car following problem in Kolmanovsky and Filev (2009) and Zidek and Kolmanovsky (2016) with the objective of maintaining a prescribed headway for as long as possible. We extend this work by taking into account lateral motion as well, i.e., allowing lane changes. The application of SDCOC provides a systematic approach to generating driving policies that may enhance safety and ride comfort for autonomous cars. SDCOC is based on *dynamic programming* (DP) and the optimal control policy is characterized by the value function associated with the optimal control problem. However, previous applications of DP were impeded by the *curse of dimensionality* and limited to lower-dimensional problems.

In this paper, we formulate an *approximate/adaptive dynamic programming* (ADP) approach for SDCOC which allows the treatment of higher-dimensional problems. ADP obtains suboptimal solutions by approximating the value function, usually through the use of neural networks (NNs), see Werbos (2012), Heydari (2014), and Wei et al. (2016). In our approach, we use a feedforward NN, see Hertz et al. (1991), combined with an enhanced version of the value iteration algorithm developed by Zidek and Kolmanovsky (2015).

The paper is structured as follows. Section 2 describes the driving model. The SDCOC problem is stated in Section 3.1 and the ADP approach to compute its solution is developed in Section 3.2. Numerical results are shown in Section 4, where we also compare the ADP approach to conventional DP. A conclusion is given in Section 5.

2. DRIVING MODEL

In order to compute a SDCOC policy (see Section 3), a discrete-time stochastic hybrid model that approximately describes the motion of a car and its surrounding traffic is formulated in this section. While the developments in this paper are limited to roads with two lanes, which constitute the largest fraction of the multi-lane roads, the modeling framework may readily be extended to more than two lanes.

Peer review under responsibility of International Federation of Automatic Control.
10.1016/j.ifacol.2017.08.048

The proposed model considers three cars. Subscript "m" denotes the controlled car ("my car") and subscripts "c" and "o" denote the closest cars ahead of the controlled car in its current lane and in the other lane, respectively. The state vector at a time instant $t \in \mathbb{Z}_{\geq 0}$ is given by $x_t = [s_{\mathrm{c},t}, s_{\mathrm{o},t}, v_{\mathrm{m},t}]^{\mathrm{T}}$, where s_{c} and s_{o} are the respective headways relative to the closest cars ahead in each lane, and v_{m} is the velocity of the controlled car. In addition, $w_t = [v_{\mathrm{c},t}, v_{\mathrm{o},t}]^{\mathrm{T}}$ is a random disturbance, where v_{c} and v_{o} are the respective velocities of the two cars ahead. Note that the closest car ahead of the controlled car in the other lane is defined to be the closest car with a relative distance $s \geq -(d+\gamma)$, where d is the length of the controlled car and γ provides a margin of safety (in this paper $\gamma = 0$).

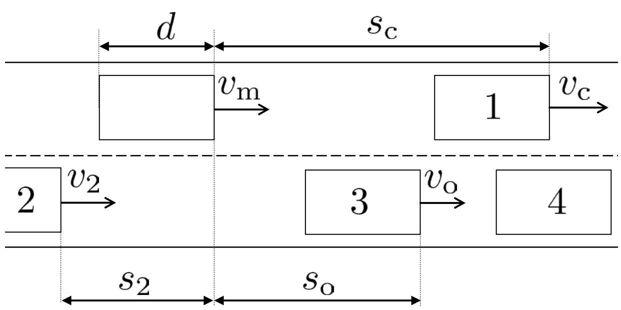

Fig. 1. Driving model: traffic example.

Figure 1 shows a traffic situation at which the controlled car is driving in the upper lane (current lane) and car 1 is the closest car ahead in its current lane. The closest car ahead in the other lane is car 3. If the velocity difference between car 2 and the controlled car, $v_2 - v_{\mathrm{m}}$, is positive, the distance between the two cars will eventually become $s_2 \geq -d$ and car 2 becomes the closest car ahead in the other lane, i.e., $s_{\mathrm{o}} = s_2$. Moreover, if car 3 cuts in between the controlled car and car 1, car 3 becomes the closest car ahead in the current lane and either car 2 (if $s_2 \geq -d$) or car 4 becomes the closest car ahead in the other lane. In addition to these scenarios, there are several other possible transitions, which are addressed by introducing a discrete-valued variable θ.

The variable θ models possible scenarios/transitions from time instant t to $t+1$. For the two-lane case, we can identify seven different transitions, $\theta \in \{1,2,3,4,5,6,7\}$. In the following, c_t and o_t denote the closest cars ahead in the current and in the other lane, respectively, at the time instant t. Similarly c_{t+1} and o_{t+1} denote the closest cars ahead in the respective lanes at the time instant $t+1$. The possible transitions are given by

- $\theta = 1$: c_t remains the closest car ahead in the current lane ($\mathrm{c}_t \to \mathrm{c}_{t+1}$) and o_t remains the closest car ahead in the other lane ($\mathrm{o}_t \to \mathrm{o}_{t+1}$).
- $\theta = 2$: c_t remains the closest car ahead in the current lane ($\mathrm{c}_t \to \mathrm{c}_{t+1}$) and a car other than o_t becomes the closest car ahead in the other lane (new car $\to \mathrm{o}_{t+1}$).
- $\theta = 3$: a car other than c_t or o_t becomes the closest car ahead in the current lane (new car $\to \mathrm{c}_{t+1}$) and a car other than c_t or o_t becomes the closest car ahead in the other lane (new car $\to \mathrm{o}_{t+1}$).
- $\theta = 4$: a car other than c_t becomes the closest car ahead in the current lane (new car $\to \mathrm{c}_{t+1}$) and o_t remains the closest car ahead in the other lane ($\mathrm{o}_t \to \mathrm{o}_{t+1}$).
- $\theta = 5$: o_t becomes the closest car ahead in the current lane ($\mathrm{o}_t \to \mathrm{c}_{t+1}$) and a car other than c_t becomes the closest car ahead in the other lane (new car $\to \mathrm{o}_{t+1}$).
- $\theta = 6$: o_t becomes the closest car ahead in the current lane ($\mathrm{o}_t \to \mathrm{c}_{t+1}$) and c_t becomes the closest car ahead in the other lane ($\mathrm{c}_t \to \mathrm{o}_{t+1}$).
- $\theta = 7$: a car other than o_t becomes the closest car ahead in the current lane (new car $\to \mathrm{c}_{t+1}$) and c_t becomes the closest car ahead in the other lane ($\mathrm{c}_t \to \mathrm{o}_{t+1}$).

The control input vector at a time instant t is given by $u_t = [a_{\mathrm{m},t}, l_{\mathrm{m},t}]^{\mathrm{T}} \in U = A \times \{0,1\}$, where $a_{\mathrm{m}} \in A$ denotes the acceleration of the controlled car and $l_{\mathrm{m}} \in \{0,1\}$ indicates whether to initiate a lane change ($l_{\mathrm{m}} = 1$) or not ($l_{\mathrm{m}} = 0$). In case $l_{\mathrm{m},t} = 1$, the current lane at the time instant t becomes the other lane at $t+1$, whereas the other lane at t becomes the current lane at $t+1$. Furthermore, we define the relative time gaps

$$T_{\mathrm{g,c}} = s_{\mathrm{c}}/v_{\mathrm{m}},\ T_{\mathrm{g,o}} = s_{\mathrm{o}}/v_{\mathrm{m}}. \tag{1}$$

We can now state the driving model, which is as follows

$$x_{t+1} = f(x_t, u_t, w_t, \theta_t) = \begin{bmatrix} s_{\mathrm{c},t+1} \\ s_{\mathrm{o},t+1} \\ v_{\mathrm{m},t} + \Delta t a_{\mathrm{m},t} \end{bmatrix}, \tag{2}$$

where Δt is the sampling time. Introducing $\Delta v_{\mathrm{c}} = v_{\mathrm{c}} - v_{\mathrm{m}}$ and $\Delta v_{\mathrm{o}} = v_{\mathrm{o}} - v_{\mathrm{m}}$, $s_{\mathrm{c},t+1}$ and $s_{\mathrm{o},t+1}$ in (2) are given by

$$s_{\mathrm{c},t+1} = \begin{cases} \min\{s_{\max}, s_{\mathrm{c},t} + \Delta t \Delta v_{\mathrm{c},t}\}, \text{ if } \theta_t \in \{1,2\} \\ \mathrm{init_c}(T_{\mathrm{g,c},t}, T_{\mathrm{g,o},t}, \theta_t), \text{ if } \theta_t \in \{3,4,7\} \\ \min\{s_{\max}, s_{\mathrm{o},t} + \Delta t \Delta v_{\mathrm{o},t}\}, \text{ if } \theta_t \in \{5,6\}, \end{cases} \tag{3}$$

$$s_{\mathrm{o},t+1} = \begin{cases} \min\{s_{\max}, s_{\mathrm{o},t} + \Delta t \Delta v_{\mathrm{o},t}\}, \text{ if } \theta_t \in \{1,4\} \\ \mathrm{init_o}(T_{\mathrm{g,c},t}, T_{\mathrm{g,o},t}, \theta_t), \text{ if } \theta_t \in \{2,3,5\} \\ \min\{s_{\max}, s_{\mathrm{c},t} + \Delta t \Delta v_{\mathrm{c},t}\}, \text{ if } \theta_t \in \{6,7\}, \end{cases} \tag{4}$$

where $s_{\max}$ is the maximum headway at which a car can be detected. If no car is ahead of the controlled vehicle in the respective lane, we set $s_{\mathrm{c}} = s_{\max}$ or $s_{\mathrm{o}} = s_{\max}$, respectively. The functions $\mathrm{init_c}$ and $\mathrm{init_o}$ in (3) and (4) set the value for s_{c} in case $\theta_t \in \{3,4,7\}$ and for s_{o} in case $\theta_t \in \{2,3,5\}$, respectively, depending on the current relative time gaps. Both $\mathrm{init_c}$ and $\mathrm{init_o}$ are defined below.

As in Kolmanovsky and Filev (2009), the velocities v_{c} and v_{o} are random variables that are modeled as Markov chains and take values in the discrete set $\mathcal{V} = \{v^j : j \in I_v\}$. The probability of transitioning from $w_t = [v^i, v^j]^{\mathrm{T}}$ to $w_{t+1} = [v^q, v^r]^{\mathrm{T}}$, given $\theta_t = p \in \{1,2,...,7\}$, is $P_w(v^q, v^r | v^i, v^j, p) \in [0,1]$ for all $i,j,q,r \in I_v$. Similarly, θ is a random variable and the probability that $\theta_t = p \in \{1,2,...,7\}$, given $T_{\mathrm{g,c},t} = T^i \in \mathcal{T}$, $T_{\mathrm{g,o},t} = T^j \in \mathcal{T}$, and $l_{\mathrm{m},t} = q \in \{0,1\}$, is $P_\theta(p | T^i, T^j, q) \in [0,1]$ for all $i,j \in I_T$, where $\mathcal{T} = \{T^j : j \in I_T\}$ is a discrete set.

Unlike v_{c} and v_{o}, $T_{\mathrm{g,c}}$ and $T_{\mathrm{g,o}}$ are continuous variables and nearest-neighbor interpolation is used to map $T_{\mathrm{g,c}} \notin \mathcal{T}$ and $T_{\mathrm{g,o}} \notin \mathcal{T}$ onto $\mathcal{T}$ when computing P_θ. The nearest-neighbor operator that maps a point $r \in \mathcal{R}$, where $\mathcal{R}$ is a