Transportation Research Part F 21 (2013) 295–314

Contents lists available at ScienceDirect

Transportation Research Part F

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

The learning of longitudinal human driving behavior and driver assistance strategies

Mark Eilers [a,*], Claus Möbus [b], Fabio Tango [c], Olivier Pietquin [d]

[a] OFFIS – Institute for Information Technology, Escherweg 2, 26121 Oldenburg, Germany
[b] Carl von Ossietzky University, Ammerländer Heerstraße 114-118, 26129 Oldenburg, Germany
[c] CRF – Centre Ricerche Fiat, Strada Torino, 50, 10043 Orbassano, Italy
[d] SequeL project team (University Lille 1, LIFL CNRS UMR 8022, INRIA Lille - Nord Europe), Parc Scientifique de la Haute-Borne, 40 avenue Halley, 59650 Villeneuve d'Ascq, FRANCE

ARTICLE INFO

Keywords:
Smart assistance
Probabilistic driver models
Markov Decision Processes
Dynamic Bayesian Networks
Structure- and Parameter Learning of probabilistic graphical models
Discriminative learning

ABSTRACT

Models of the human driving behavior are essential for the rapid prototyping of error-compensating assistance systems. Various authors proposed control-theoretic and production-system models. Here we present machine-learning alternatives to train assistance systems and estimate probabilistic driver models from human behavior traces. We present a partially autonomous driver assistance system based on Markov Decision Processes. Its assistance strategies are trained from human behavior traces using the Least Square Policy Iteration algorithm. The resulting system is able to reduce the number of collisions encountered when following a lead-vehicle. Furthermore, we present a Bayesian Autonomous Driver Mixture-of-Behaviors model for the longitudinal control of human drivers based on the modular and hierarchical composition of Dynamic Bayesian Networks. Their parameters and structures are estimated from human behavior traces using a discriminative scoring criterion based on the Bayesian Information Criterion. This allows the selection of pertinent percepts from the variety of percepts proposed for driver models according to their statistical relevance. The resulting driver model is able to reproduce the longitudinal control behavior of human drivers while driving unassisted or assisted by the presented assistance system.

1. Introduction

The Human or Cognitive Centered Design of intelligent transport systems requires computational models of human behavior and cognition. Particularly models of human *driving* behavior are essential for the rapid prototyping of error-compensating assistance systems (Cacciabue, 2007). By now, such models are mainly developed and used as driver models in traffic scenario simulations to provide safety assertions and support risk-based design. However, due to the need for smarter and more intelligent assistance and considering that the problem of transferring human skills into the envisioned technical systems is becoming more and more apparent (Xu & Lee, 2005), driver models that are able to predict the driver's behavior and maneuver intentions could be utilized to improve the future generation of assistance systems (e.g. Aoude, Desaraju, Stephens, & How, 2012; Kasper et al., 2011; Mabuchi & Yamada, 2011; Möbus, Eilers, Garbe, & Zilinski, 2009; Möbus & Eilers, 2011a; Morris, Doshi, & Trivedi, 2011; Ortiz, Fritsch, Kummert, & Gepperth, 2011).

* Corresponding author. Tel.: +49 4417982381.
E-mail addresses: mark.eilers@offis.de (M. Eilers), claus.moebus@uni-oldenburg.de (C. Möbus), fabio.tango@crf.it (F. Tango), olivier.pietquin@univ-lille1.fr (O. Pietquin).

1369-8478/$ - see front matter
http://dx.doi.org/10.1016/j.trf.2013.09.021

In this paper, we present machine-learning approaches for developing driver assistance systems and driver models. In particular, a partially autonomous driver assistance system (PADAS) called LOSS (Longitudinal Support System, developed by the third and fourth author), is implemented by a Markov Decision Process (MDP). The LOSS enhances common forward collision warning and adaptive cruise control systems, adding assisted and emergency braking functionalities. In order to decide the correct strategy (e.g. warning or intervention) for a given situation, the system was trained using human behavior traces obtained in a driver simulator study. Furthermore, we present a model of the human driver, realized by a Bayesian Autonomous Driver Mixture-of-Behaviors (BAD MoB) model based on the modular and hierarchical composition of Dynamic Bayesian Networks (DBNs) (developed by the first and second author). The parameters and structures of the DBNs were estimated from multivariate time-series of human behavior traces in a separate driving simulator study applying the LOSS. The resulting BAD MoB model is able to reproduce the longitudinal control behavior of human drivers while driving unassisted and being assisted with the LOSS. We believe that the proposed LOSS and BAD MoB model can be combined into an improved PADAS. In this PADAS, the MDP will be used to determine what the driver *should* do to obtain safe driving behavior while the BAD MoB model will be used to predict and anticipate what the driver most probably *will* do in the near future. Here, we attempt to lay the foundation to this approach.

This work was primarily conducted in the European project "Integrated Human Modelling and Simulation to support Human Error Risk Analysis of Partially Autonomous Driver Assistance Systems" (ISi-PADAS). The objective of ISi-PADAS was to provide a methodology for the risk assessment of a PADAS based on simulations with driver models in order to replace potential expensive and intense testing of new assistance systems in experimental studies with human drivers. As a demonstration of this methodology, both LOSS and the BAD MoB model were simultaneously utilized in a project-spanning use-case.

This paper is organized as follows. Section 2 presents the development of a new PADAS called LOSS based on Markov Decision Processes (MDPs). Section 2.1 elaborates the general architecture of artificial co-drivers in general and the LOSS in particular, the building blocks and the technologies that are needed to bring them to life, pointing out which parts have been already researched and which gaps still remain. We clarify what "understanding the driver" actually means and how a joint system can be obtained. Section 2.2 addresses the problem of finding an optimal strategy for the PADAS, which is constituted by a set of decision-rules that determines, as a function of the vehicle situation, the sequence of signals sent to the driver (warning) and how the driver's vehicle is decelerated (intervention). Section 2.3 presents driving simulator studies conducted to obtain human behavior traces needed to learn the optimal strategy. Section 2.4 then presents early results for the PADAS. Section 3 presents the development of a model of the human driver realized by a Bayesian Autonomous Driver Mixture-of-Behaviors (BAD MoB) model, learnt from human behavior traces. Section 3.1 gives a firm introduction to the structure and utilization of BAD MoB models. Section 3.2 addresses the machine-learning approach to obtain BAD MoB models from human behavior traces. Section 3.3 presents a driving simulator study conducted to obtain human behavior traces and presents the resulting learnt BAD MoB model. A validation and comparison of the BAD MoB model behavior and experimental data is presented in Section 3.4. The paper closes with a summary and conclusion given in Section 4.

2. Learning optimal strategies for driver assistance systems

The goal of a PADAS is to aid human users to drive safely. This can be achieved by developing a co-driver, which is able to "understand" drivers and to form symbiotic systems with them, providing the pertinent information thanks to a natural interface. Such a system gives a continuous support, able to assist the driver in many different traffic situations and scenarios, intervening continuously from warning up to automatic braking in the whole longitudinal control of the vehicle (Tango, Minin, Aras, & Pietquin, 2011). The basic metaphor idea defines the system as a sentient co-driver, which is usually silent but can advise the driver or even intervene when a dangerous situation happens. Communication between driver and co-driver occurs at different levels through multiple HMI channels with extensive use of haptic feedbacks and interaction on driver commands.

2.1. Modeling approach

As Da Lio, Biral, Galvani, & Saroldi (2012) observed, mankind has used animals, and especially horses, as transportation systems for thousand years. Last century they were eventually replaced by motor vehicles because of range, speed, capacity, and costs. However something was lost: the intelligence of the animals was traded with power. Norman (2007) recalls the interaction between a rider and a horse as one example of how future intelligent things should work. In this context, the "co-driver" is an artificial system smart enough to "understand" the driver and to form with the driver a symbiotic system (Möbus & Eilers, 2009, 2011a) (thus implementing the aforementioned vision, for example adopting the H-metaphor, Flemish et al., 2003). In other words, the co–driver is a sort of a "virtual driver" in the car (Fig. 1), able to detect the road scenario, including obstacles, road geometry, vehicle position and motion, as well as all the other information necessary to drive. This "virtual-driver" "optimizes" the situation, selecting the most appropriate actions, based on the current state, in order to move to the next (safer) state. Then it interacts accordingly with the driver, by means of a dedicated HMI.

The term "co-driver" is adapted from the co-pilot in the aviation domain, where – strictly speaking – it indicates the second pilot of an aircraft. In a broader sense it may also indicate an automatism that can take care of some guidance tasks. For a complete review of the State of the Art, the interest reader can see Da Lio et al. (2012), Saroldi, Tango, & Da Lio (2012), and the projects HAVEit and DIPLECS. In our case, the co-driver is a model, which produces sensory-motor strategies for many