Discriminative Learning of Relevant Percepts for a Bayesian Autonomous Driver Model

Mark Eilers
Human Centered Design
OFFIS Institute for Information Technology
Oldenburg, Germany
e-mail: mark.eilers@offis.de

Claus Möbus
Learning and Cognitive Systems
Carl-von-Ossietzky University
Oldenburg, Germany
e-mail: claus.moebus@uni-oldenburg.de

***Abstract*—Models of the human driving behavior are essential for the rapid prototyping of assistance systems. Based on psychological studies, various percepts and measures have been proposed for the lateral and longitudinal control in driver models without demonstrating the generalizability of results to natural settings. In this paper, we present the learning of a probabilistic driver model. It represents and mimics the lateral and longitudinal human driving behavior on virtual highways by performing situation-adequate lane-following, car-following, and lane changing behavior. Because there is considerable uncertainty about the relevant percepts in natural driving behavior, we select hypothetically relevant percepts from the variety of possibilities based on their statistical relevance. This is a new approach to generate hypothesis about the relevant percepts and situation-awareness of drivers in dynamic traffic scenes. The percepts are revealed in a structure-learning procedure using a discriminative scoring criterion based on the Bayesian Information Criterion. Discriminative learning maximizes the conditional likelihood of probabilistic models, whereas the traditional generative learning maximizes the unconditional likelihood. This way, it attempts to find the structure with the best performance for the intended use, which in our application is the best prediction of driving actions given the available percepts.**

Keywords–Probabilistic Driver Models; Bayesian Autonomous Driver Models; Machine-Learning; Structure-Learning; Discriminative Learning.

I. INTRODUCTION

The Human Centered Design of intelligent transport systems requires computational models of human behavior and cognition. Particularly models of the human driving behavior (i.e., driver models) are essential for the rapid prototyping of error-compensating assistance systems [1]. Various authors proposed control-theoretic models (e.g., [2]), closely related perception-action models (e.g., [3][4]), and production-system models implemented in cognitive architectures (e.g., [5]). Due to the variability of human cognition and behavior, the irreducible lack of knowledge about underlying cognitive mechanisms, and the irreducible incompleteness of knowledge about the environment [6], we conceptualize, estimate, and implement models of human drivers as probabilistic models: Bayesian Autonomous Driver (BAD) models.

In earlier research [7], we developed a BAD model with Dynamic Bayesian Networks (DBNs), based on the assumption that a single DBN representing a single skill is sufficient for lateral and longitudinal control. Later, we realized that for modeling the complex competence of human drivers, a skill hierarchy (e.g., Figure 1) is necessary. We developed a hierarchical modular probabilistic architecture that allows to construct driver models by decomposing complex behaviors into pure behaviors and vice versa: Bayesian Autonomous Driver Mixture-of-Behaviors (BAD MoB) models [8][9][10].

Based on psychological studies (e.g., [11][12][13][14][15]), various percepts and measures have been recommended for the lateral and longitudinal control in driver models. These proposals are partly contradictory and often depend on special experimental settings, like straight roads, winding roads, low speed, and/or the absence of other traffic participants. Other and more natural scenarios may provide and require different perceptual cues that are not yet fully understood or formalized. A general computational vision theory of driving behavior is still pending.

Because there remains considerable uncertainty about the relevant percepts in natural driving behavior, we propose the use of structure-learning procedures to select hypothetically relevant percepts from the variety of possibilities based on their statistical relevance. We used the proposed procedure to learn the relevant percepts for a BAD MoB model that represents and mimics the lateral and longitudinal human driving behavior on virtual two-lane highways according to the skill hierachy shown in Figure 1. We assume that the overall complex driving behavior can be decomposed into the simpler behaviors lane-following, car-following, performing lane changes to the left and lane changes to the right.

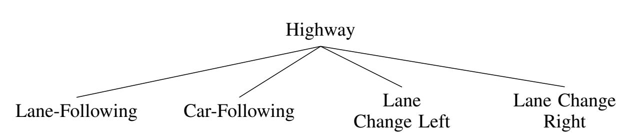

Figure 1. Skill hierarchy representing the human driving behavior on virtual two-lane highways.

 ISBN: 978-1-61208-340-7

The paper is organized as follows. In the following section we give a brief overview of percepts and measures that have been recommended in the literature for modeling the lateral and longitudinal human control behavior. Section III introduces the fundamentals of BAD MoB models. In Section IV, we describe a structure-learning procedure for selecting the pertinent percepts from a universe of hypothetically available percepts, using a discriminative version of the Bayesian Information Criterion. In Section V, we present a resulting BAD MoB model that mimics the human driving behavior on virtual highways and discuss the meaningfulness of the learned percepts with respect to the literature. Finally, we conclude with Section VI.

II. A Universe of Percepts

For the most part, the literature considers three types of percepts as important for lateral control: *bearing angles* [13][14][3][16], *splay angles* [13][14][17], and the *optic flow* [13][14][17] (Figure 2).

Bearing angles are defined as the angles between the driver's heading (the driver's body axis, assuming that he is belted in) and the direction to specific reference points in the driver's field of view (Figure 2a) [14]. If available, obvious choices for such reference points are the lane edges. When aligned with the course of the road, a conceivable strategy for lane-keeping on a straight path is to keep the bearing angle to a reference point on the left lane edge and the bearing angle to a reference point on the right lane edge constant [14].

Bearing angles to single reference points that drivers tend to *visually target* are also known as *visual direction angles* [13]. Notable proposed examples for such targeted reference points are points on the future path of the driver, the centerline, lead-cars, and for curved roadways, the tangent point [3][18][19].

Splay angles are defined as the optical projections of lane edges or the centerline around a reference point on the driver's retina relative to a vertical line in the driver's field of view, e.g., the heading [14] (Figure 2b). Similar to bearing angles, when aligned with the course of the road, a valid strategy for lane-keeping would be to keep the splay angle to the left lane edge and the splay angle to the right lane edge constant.

The optic flow denotes the global image motion of the environment projected on the retina when one moves in the world [14][16]. Similar to bearing and splay angles, we can define flow angles as the angles between the driver's heading and the direction of the optic flow of reference points in the driver's field of view (Figure 2c). A simple strategy for lane-keeping using the optic flow would be to align the focus of expansion (specified by the intersection of two flow angles) with the intended target direction.

For longitudinal control, the literature mainly discusses the time-to-contact/collision (TTC) [13][20][21] and the time-headway (THW) [21][22]. The TTC of a vehicle A with a speed v_A, following a vehicle B with a speed v_B, in a distance d, is defined as TTC $= d/(v_A - v_B)$ and denotes the remaining time until A reaches B. As a special case of the TTC, the THW denotes the remaining time until A will reach the current position of B and is defined as THW $= d/v_A$.

III. Bayesian Autonomous Driver Mixture-of-Behavior Models

Throughout this paper, we will be concerned with probability distributions over sets of discrete random variables. Variables and set of variables will be denoted by capital letters, while specific values taken by those (sets of) variables will be denoted by lowercase letters. For time series, we assume that the timeline is discretized into time-slices with a constant granularity of 50ms. We will index these time-slices by non-negative integers and will use X_i^t to represent the instantiation of a variable X_i at time t. A sequence $X_i^j, X_i^{j+1}, \ldots, X_i^k$ will be denoted by $X_i^{j:k}$ and we will use the notation $x_i^{j:k}$ for an assignment of values to such sequences.

A Bayesian Network (BN) is an annotated directed acyclic graph (DAG) that encodes a joint probability over a set of variables $X = \{X_1, \ldots, X_n\}$ [23]. Formally, a Bayesian Network B is defined as a pair $B = \{G, \theta\}$. The component G is a DAG, whose vertices correspond to the random variables $X_1, \ldots, X_n$, and whose arcs define the (in)dependencies between these variables, in that each variable X_i is independent of its non-descendants given its (possibly empty) set of parents $Pa\left(X_i\right)$ in G. The component θ represents a set of parameters that quantifies the probabilities of the network. We assume that θ contains a parameter $\theta_{x|pa(X)} = P\left(x|pa\left(X\right)\right)$ for each possible value $x \in X$ and $pa\left(X\right) \in Pa\left(X\right)$. Given G and θ, a Bayesian network B defines a unique joint probability distribution (JPD) over X as:

$$P_B\left(X\right) = \prod_{i=1}^{n} P\left(X_i|Pa\left(X_i\right)\right). \tag{1}$$

DBNs extend BNs to model the stochastic evolution of a set of variables $X = \{X_1, \ldots, X_n\}$ over time [24]. A DBN D is defined as a pair $D = \left(B^1, B^{\rightarrow}\right)$, where $B^1 = \left(G^1, \theta^1\right)$ is a BN that defines the probability distribution $P\left(X^1\right)$ and, under the assumption of first-order Markov and stationary processes, $B^{\rightarrow} = \left(G^{\rightarrow}, \theta^{\rightarrow}\right)$ is a two-slice Bayesian network (2TBN) that defines the CPD $P\left(X^t|X^{t-1}\right)$ for all t. The nodes in the first slice of the 2TBN do not have any parameters associated with them, but each node in the second slice of the 2TBN has an associated CPD which defines $P\left(X_i^t|Pa\left(X_i^t\right)\right)$, where a parent $X_j \in Pa\left(X_i^t\right)$ can either be in time-slice t or $t-1$. The JPD over any number of T time-slices is then given by:

$$P_D(X^{1:T}) = \prod_{t=1}^{T}\prod_{i=1}^{n} P\left(X_i^t|Pa\left(X_i^t\right)\right). \tag{2}$$

A. Definition of BAD MoB models

In essence, a BAD MoB model is a combination of several DBNs, whose functional interaction allow the generation of context dependent driving behavior by sequencing and mixing simpler behaviors according to a skill hierarchy [9][10] (e.g., Figure 1).

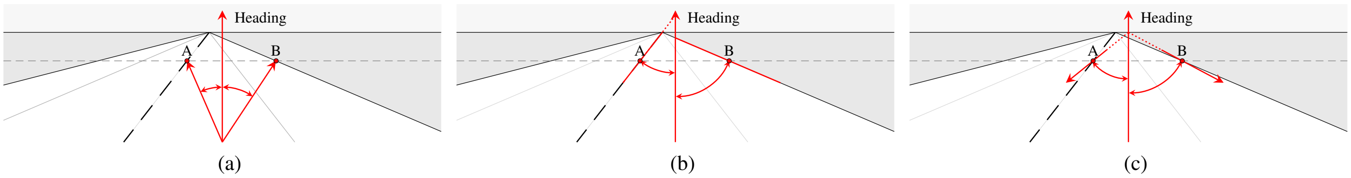

Figure 2. Illustration of bearing angles (a), splay angles (b) and flow angles (c).

Let A denote a set of discrete random variables representing the lateral and longitudinal control actions of a driver. In this paper, we assume the use of steering wheel angles for lateral control and the position of a combined acceleration and braking pedal for longitudinal control. $P = \{P_1, \dots, P_m\}$ denotes a set of discrete random variables representing hypothetical percepts that could be available to the driver due to foveal and ambient vision [25]. Given a skill hierarchy, decomposing a complex behavior in a number of n simple behaviors, B denotes a discrete random variable with n values for the n simple behaviors.

We assume, that the sensor-motor schema of each of the n simple behaviors in the skill hierarchy can be modeled by a distinct DBN π_i that defines a JPD $P_{\pi_i}\left(A^{1:T}, P^{1:T}\right)$. Due to their purpose, we refer to these models as *action-models*. In addition, we assume that the appropriateness of the simple behaviors in a given situation can be modeled by a DBN π_B that defines a JPD $P_{\pi_B}\left(B^{1:T}, P^{1:T}\right)$, which we refer to as a *behavior-classification-model*.

A BAD MoB model π is then defined as a DBN that combines both action- and behavior-classification-models, using a technique called *behavior-combination* [26]. The combination is achieved by specifying the CPDs of π as queries to be inferred by the action- and behavior-classification-models. Under the assumption of first-order Markov and stationary processes, the JPD of π for any number of T time slices is defined as:

$$\begin{aligned}&P_\pi\left(A^{1:T}, B^{1:T}, P^{1:T}\right)\\ &= \prod_{t=1}^{T} P\left(P^t\right) P\left(B^t|B^{t-1}, P^t\right) P\left(A^t|A^{t-1}, P^t, B^t\right). \quad (3)\end{aligned}$$

The CPD $P\left(B^t|B^{t-1}, P^t\right)$ represents the appropriateness of each simple behavior in the given situation and is defined as a query $P_{\pi_B}\left(B^t|B^{t-1}, P^t\right)$ to be inferred by the behavior-classification-model π_B. Each CPD $P\left(A^t|A^{t-1}, P^t, B^t = i\right)$ represents the motor-output of a specific behavior in a given situation and is defined as a query $P_{\pi_i}\left(A^t|A^{t-1}, P^t\right)$ to be inferred by the action-model π_i that realizes the sensor-motor-schema of the corresponding behavior. We would like to emphasize that the structure of π itself is predefined and fixed. In contrast, the structures of action- and the behavior-classification-models will be obtained via structure-learning procedures.

B. Definition of component-models

Action- and behavior-classification-models are defined in the same manner, and we will therefore simply refer to them as *component-models*. Each component-model π_c is a distinct DBN that defines a joint distribution $P_{\pi_c}\left(X^{1:T}, P^{1:T}\right)$ and will be used to infer the query $P_{\pi_c}\left(X^t|X^{t-1}, P^t\right)$ (where $X = A$ in the case of action-models and $X = B$ in the case of behavior-classification-models).

We model component-models in the fashion of state-observation models [23], in that they consist of a *transition model* $P\left(X^t|X^{t-1}\right)$ and an *observation model* $P\left(P^t|X^t\right)$. We rely on the assumption that not all of the available percepts P are necessarily relevant for the realization or classification of driving behaviors. Accordingly, we can separate P into two mutually exclusive sets $P_R \subseteq P$ and $P_I \subseteq P$, where P_R consists of the relevant percepts and P_I consists of the irrelevant percepts.

This allows us to decompose $P\left(P^t|X^t\right)$ to $P\left(P_R^t|X^t\right) P\left(P_I^t\right)$. For $P\left(P_R^t|X^t\right)$, we assume that the percepts are conditionally independent given X^t: $P\left(P_R^t|X^t\right) = \prod_{P_i \in P_R} P\left(P_i^t|X^t\right)$. As the irrelevant percepts P_I have no influence on the estimation of X, we can replace the CPD $P\left(P_I^t\right)$ by $\prod_{P_j \in P_I} P\left(P_j^t\right)$.

A schematic graph-structure of component-models is shown in Figure 3. The JPD over all variables for any number of T time-slices is then given by:

$$\begin{aligned}&P_\pi(X^{1:T}, P^{1:T})\\ &= \prod_{t=1}^{T}\left[P\left(X^t|X^{t-1}\right) \prod_{P_i \in P_R} P\left(P_i^t|X^t\right) \prod_{P_j \in P_I} P\left(P_j^t\right)\right]. \quad (4)\end{aligned}$$

The query $P_{\pi_c}(X^t|X^{t-1}, P^t)$ needed for the specification of the BAD MoB model can be inferred by:

$$\begin{aligned}&P_{\pi_c}(X^t|X^{t-1}, P^t)\\ &= \frac{P\left(X^t|X^{t-1}\right) \prod_{P_i \in P_R} P\left(P_i^t|X^t\right)}{\sum_{x^t \in X^t} P\left(x^t|X^{t-1}\right) \prod_{P_i \in P_R} P\left(P_i^t|x^t\right)}. \quad (5)\end{aligned}$$

As each component-model is a distinct DBN, each may use a different set of relevant percepts. The problem is to decide, which of the available percepts $P_i \in P$ are relevant and which are irrelevant. Due to the considerable uncertainty about the relevant percepts for realization and classification of natural driving behaviors, we use machine-learning methods to learn the graph structure of component-models and obtain the statistically relevant percepts in natural driving behaviors from the variety of proposed percepts.

 ISBN: 978-1-61208-340-7

IV. Learning BAD MoB models

We derive the structures of component-models by a machine-learning method, based on the score and search paradigm, where a search in the space of possible graph structures is guided by a scoring function that evaluates the degree of fitness between the model and a set of experimental data. From a Bayesian perspective such a scoring criterion can be defined as $P(G_\pi|\delta)$, the probability of a graph structure G_π of a model π given a dataset δ [23]. Using the Bayes' rule, $P(G_\pi|\delta)$ is given by:

$$P(G_\pi|\delta) = P(\delta|G_\pi)\, P(G_\pi) / P(\delta), \tag{6}$$

where $P(\delta|G_\pi)$ is the likelihood of the data given the graph structure, $P(G_\pi)$ is a prior over possible graph structures, and $P(\delta)$ is a constant that does not depend on the actual graph structure and can therefore be neglected [23]. The likelihood of the data given a graph structure can be computed by integrating over all possible parameters θ_π of π [23][24]:

$$P(\delta|G_\pi) = \int P(\delta|G_\pi, \theta_\pi)\, P(\theta_\pi|G_\pi)\, d\theta_\pi, \tag{7}$$

where $P(\delta|G_\pi, \theta_\pi)$ is the likelihood of the data given π and $P(\theta_\pi|G_\pi)$ is a prior distribution over the possible parameter values for a graph structure G_π. A common approach for evaluating the integral, is to use an approximation derived from the asymptotic behavior of (7) for infinite datasets, which results in a scoring criterion known as the *Bayesian Information Criterion* (BIC) [23][24][27][28].

Although their structure is fixed, we derive our structure-learning approach for component-models from the hypothetically learning of BAD MoB models. Let δ denote a *complete* database consisting of n samples $\delta^i = (a^i, b^i, p^i)$, $\hat{\theta}_\pi$ denote the maximum likelihood estimator, and $\mathrm{Dim}[\pi]$ denote the number of independent parameters, the BIC score for a BAD MoB model π is defined as:

$$\mathrm{BIC}(G_\pi : \delta) = \log P\left(\delta|G_\pi, \hat{\theta}_\pi\right) - \frac{\mathrm{Dim}[\pi]}{2} \log n. \tag{8}$$

For a BAD MoB model π, the log-likelihood $\log P(\delta|G_\pi, \hat{\theta}_\pi)$ is given by $\log P_\pi(a^{1:n}, b^{1:n}, p^{1:n} : \hat{\theta}_\pi)$ (in the following we will drop $\hat{\theta}_\pi$ for clarity). Using (3) this can be rewritten as a sum of terms:

$$\log P_\pi\left(a^{1:n}, b^{1:n}, p^{1:n}\right) = \sum_{i=1}^{n} \log P\left(b^i|b^{i-1}, p^i\right) + \sum_{j=1}^{n} \log P\left(a^j|a^{j-1}, b^j, p^j\right) + \sum_{k=1}^{n} \log P\left(k^i\right). \tag{9}$$

As the structure of a BAD MoB model is predefined, its score only depends on the structure of its component-models and their consequential ability to infer their corresponding queries. This translates the task of learning a BAD MoB model into the task of learning the graph structure for each component-model individually. Consequently, using (9), we can decompose (8) in order to define a scoring criterion for component-models.

A. Discriminative learning of component-models

Let δ_c denote a subset of δ consisting of only the n_c samples $\delta^i = (x^i, x^{i-1}, p^i)$ related with a component-model π_c, the portion of the BIC score for π_c would be given by:

$$\sum_{i=1}^{n_c} \log P_{\pi_c}\left(x^i|x^{i-1}, p^i\right) - \frac{\mathrm{Dim}[\pi_c]}{2} \cdot \log n. \tag{10}$$

In contrast to maximizing the unconditional log-likelihood in (8), we now aim to maximize a conditional log-likelihood $\mathrm{CL}\left(\hat{\theta}_{\pi_c} : \delta_c\right) = \sum_{i=1}^{n_c} \log P_{\pi_c}\left(x^i|x^{i-1}, p^i\right)$. Learning in order to maximize a conditional (log-)likelihood is known in the literature as *discriminative* learning [23][28][29][30][31]. Accordingly we will refer to the scoring criterion for component-models as a *Discriminative BIC* (DBIC).

For discriminative learning based on the BIC, it has been recognized that its penalty term tends to have a too high impact on the score, resulting in too simple model structures [29][30]. As a consequence, [30] propose to adjust the penalty by multiplying it by a parameter $\beta < 1$, they proposed as $\beta = 1/10$. Following this, the DBIC for a component-model is then defined as:

$$\mathrm{DBIC}(G_{\pi_c} : \delta_c) = \mathrm{CL}\left(\hat{\theta}_{\pi_c} : \delta_c\right) - \beta \frac{\mathrm{Dim}[\pi_c]}{2} \cdot \log n. \tag{11}$$

B. Learning Procedure

For each component-model, the goal of the learning procedure is to find the graph structure from the space of possible graph structures that maximizes the DBIC. Even given the severe structural constraints of component-models (cf. Section III-B), there exist 2^m possible graph structures for a number of m available percepts. As it is not feasible to evaluate all these possibilities, we rely on heuristic methods to find a good but not necessarily optimal solution.

By now, we use a common *greedy hill-climbing* search procedure [23]. We start with an initial "blind" model that does not utilize any percepts (hence with the observation model $P(P^t|X^t) = \prod_{i=1}^{m} P(P_i^t)$) and compute its DBIC. For each available percept $P_j \in P$, we then construct a model in which we utilize P_j by adding an edge in

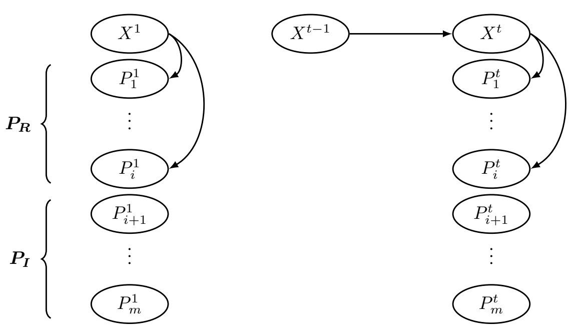

Figure 3. Schematic structure of action- and behavior-classification-models, defined by a BN and a 2TBN.

the graph from X^t to P_j^t (therefore replacing $P\left(P_j^t\right)$ in the observation model with $P\left(P_j^t|X^t\right)$) and compute its DBIC score. Intuitively, a percept P_k^t leading to the highest improvement of the DBIC when utilized can be seen as the most pertinent percept of the given possibilities and will be permanently added to the initial model. The process is then repeated until no further added percept improves the DBIC. Eventually, this learning procedure selects a minimal set of relevant percepts.

V. A BAD MoB model representing the human driving behavior on virtual highways

We used the described method to learn the relevant percepts for a BAD MoB model based on the skill hierarchy shown in Figure 1, representing the lateral and longitudinal human driving behavior on virtual highways. For this, we selected approx. 800 percepts that hypothetically could be relevant for the human lateral and longitudinal driving behavior.

Primarily, we selected bearing, splay, and flow angles (c.f. Figure 2), utilizing a variety of possible reference points. Reference points were placed in different fixed $(25m, 50m, \ldots, 250m)$ and time-dependent $(i \cdot speed,\ i = \{1s, 2s, \ldots, 10s\})$ distances, both relative to the driver's current lane (on the lane edge or centerline left to the driver, on the middle of the driver's lane, and on the lane edge or centerline right to the driver) and absolute (on the left lane edge, the middle of the fast lane, the centerline, the middle of the slow lane, and the right lane edge). As additional reference points only applicable for bearing angles, we selected the far point, as proposed by [3] (placed on the tangent point if available and on the vanishing point otherwise), and traffic participants in the vicinity of the driver (the nearest cars in front and behind the driver on the two lanes of the highway).

To enable the possible use of strategies that utilize two angles resp. reference points simultaneously (cf. Section II), we considered percepts that represent the differences between bearing, splay, and flow angles of the reference points on the left and right lane edges or one of these lane edges and the centerline.

From the percepts obtainable from traffic participants in the vicinity of the driver, we selected distances, speed differences to the driver, TTCs, and THWs of the nearest cars in front and behind the driver on the two lanes of the highway.

As further percepts that obviously could have an effect, esp. on longitudinal control, we included the driver's speed v_{ego}, the prescribed speed limit v_{limit}, and the combination of both, which we will refer to as the speed potential, representing the allowed speed gain (resp. prescribed speed reduction) defined as $v_{\text{pot}} = v_{\text{limit}} - v_{\text{ego}}$.

A. Experimental Data

The database needed for the learning procedure was obtained in a simulator study using a fixed based driving simulator that comprises a mockup of the driver's cab of a real car, positioned amidst three projection surfaces for a simulated 3D-environment, providing a realistic field of view of 170° (Figure 4).

The study was conducted with eight participants (four male, four female) between the age of 24 and 30 and with normal or corrected-to-normal vision. The scenario comprised approximately 37 km of a four-lane highway based on a section of the German highway A1, with two lanes in each direction and moderate traffic, generated by a number of non-controlled automated vehicles traveling at varying desired speeds and able to pass other vehicles (including the driver's car).

The experimental procedure consisted of three phases. In the first phase, each participant was introduced to the simulator and performed a training session. In the second phase, the participants drove one trial without other traffic. They were instructed to obey the shown speed-limits (100-130 km/h) and perform lane changes when being asked by the instructor. In the third phase, the participants drove two trials with other traffic. They were instructed to obey the German traffic rules. For the second trial with other traffic, left-hand curvatures were inverted to right-hand curvatures and vice versa. A single participant (participant 1, male) attended two times and in each case performed an additional third trial with other traffic.

During the trials, with a frequency of 60 Hz, we recorded the values of all defined percepts, the steering wheel angle, and the position of a combined acceleration-braking pedal. This led to an experimental database of approx. 1900000 data samples $\delta^i = (a^i, p^i)$ comprising a total time of approx. 525 minutes. In order to define the missing behavior values, we manually completed each sample δ^i with the shown behavior b^i according to the skill hierarchy (Figure 1).

B. Results and Discussion

By now, we used the experimental data of participant 1 (male, approx. 550000 data samples) to learn a BAD MoB model representing an individual driver. In total, twelve percepts were selected during the learning procedure (Table I). In the following, we will attempt to discuss their meaningfulness with respect to the literature.

1) Lane-Following: A single percept was learned as relevant for lateral control, representing the bearing angle

Figure 4. Fixed based driving simulator.

TABLE I. PERTINENT PERCEPTS FOR ACTION-MODELS REPRESENTING THE BEHAVIORS FOR LANE-FOLLOWING, CAR-FOLLOWING, LANE CHANGES TO THE LEFT, AND LANE CHANGES TO THE RIGHT.

Behavior	Relevant Percepts for Lateral Control	Relevant Percepts for Longitudinal Control
Lane-following	1. Bearing angle between the heading and a reference point on the middle of the driver's lane in a distance of $v_{\text{ego}} \cdot 5s$	1. Speed potential
Car-following	1. Flow angle between heading and the optic flow direction of a reference point on the centerline in a distance of $v_{\text{ego}} \cdot 3s$	1. TTC to the lead-car (bumper-to-bumper distance, allowing positive and negative values)
Lane changes to the left	1. Bearing angle between the heading and a reference point on the middle of the right lane in a distance of 100m	1. TTC to the lead-car on the fast lane (bumper-to-bumper distance, only positive values)
	2. Bearing angle between the heading and a reference point on the left lane edge in a distance of $v_{\text{ego}} \cdot 2s$	
Lane changes to the right	1. Flow angle between the heading and the optic flow direction of a reference point on the middle of the right lane in a distance of 75m	1. TTC to the lead-car on the fast lane (bumper-to-bumper, only positive values)
	2. Bearing angle between the heading and a reference point on the right lane edge in a distance of $v_{\text{ego}} \cdot 3s$	
	3. TTC to the lead-car on the fast lane (bumper-to-bumper distance, allowing positive and negative values)	
	4. TTC from the car behind on the slow lane (Euclidean distance, only positive values)	

between the driver's heading and a reference point on the middle of the driver's lane in a distance of $v_{\text{ego}} \cdot 5s$. Under the assumption that the middle of the lane can be seen as an approximation of the future path of the driver, this percept is consistent with the proposals of [13] for roads with gentle curvatures, and findings of [19] and [32], who report that drivers often fixate the center of the road resp. the future path. In contrast, the far point, as proposed by [3] for lane-following, was only the 114^{th} highest-rated percept.

The single relevant percept selected for the longitudinal control represents the speed potential. This seems reasonable, as it should be sufficient for a driver to keep an intended target speed as implied by the current speed limit. However, as German traffic rules prohibits to overtake on the right, we expected a second percept, associated with the lead-car on the lane left to the driver. Indeed, the percept that would have been selected as the second relevant percept (but was rejected due to the increasing penalty) represents the distance to the lead-car on the fast lane.

2) Car-Following: The single percept learned for the lateral control in car-following represents the angle between the driver's heading and the optic flow direction of a reference point on the centerline in a distance of $v_{\text{ego}} \cdot 3s$. This is in contrast to [33], who found that drivers tend to primarily fixate the lead-car during car-following. Based on these findings, [3] concluded that in the presence of a lead-car, the lead-car would act as the primary reference point and consequently proposed the angle between the heading and the lead-car as the most relevant percept for lateral control during car-following. Maybe surprisingly, this angle was only the 353^{th} best-rated percept.

However, the single relevant percept for longitudinal control in car-following represents the TTC to the lead-car. This is consistent with the proposals of [20] and [21] and would imply that the driver indeed primarily focuses the lead-car. Our model would therefore imply that during car-following, the driver primarily depends on ambient vision for lateral control, while using the foveal vision for longitudinal control.

3) Lane Changes: For the lateral control during lane changes to the left, two percept were selected. They represent the bearing angle to the middle of the right lane in a distance of 100m and the bearing angle to the left lane edge in a distance of $v_{\text{ego}} \cdot 2s$. These percepts are consistent with findings of [34], who report that during lane changes to the left lane, drivers direct their gazes primarily and almost equally to the left and the right lane.

In contrast, during lane changes to the right, drivers direct the majority of their gazes to the right lane, while dividing the rest of their gazes equally between the left lane and the mirror [34]. As shown in Table I, the selected percepts are indeed consistent with these findings. The two most relevant percepts represent the angle between the heading and the optic flow direction of a reference point on the middle of the right lane and the bearing angle between the heading and a reference point on the right lane edge. The third percept represents the TTC to the lead-car on the fast lane, which implies a certain attention to the left lane. The last percept represents the TTC to the car behind on the slow lane, which implies a certain attention to the mirror.

Concerning the longitudinal control, for both lane changes to the left and to the right, the single selected percept represents the TTC to the lead-car on the fast lane. This may be explained by a rare and unpredictable tendency of non-controlled traffic participants to surprisingly and recklessly change lanes, which enforced the participant to perform all-out brakings during lane changes. However, this also gives us a hint to explore the use of separated skill hierarchies for lateral and longitudinal control in our future research.

VI. Conclusion

We presented the learning of a hierarchical and modular probabilistic driver model that represent and mimics the lateral and longitudinal human driving behavior on virtual highways. Its relevant percepts were selected in a discriminative structure-learning procedure from a set of hypothetical percepts proposed in literature. The performance of the learned BAD MoB model is very promising (videos available at http://www.lks.uni-oldenburg.de/46350.html). The selected percepts are sufficient for the modeling and simulation of the lateral and longitudinal human driving behavior on virtual highways, including situation-adequate lane-following, car-following, and lane changing behavior. The selected percepts seem reasonable and for the most part consistent with findings reported in psychological studies. This indicates that the proposed method can be used to generate hypothesis about the relevant percepts and situation-awareness of drivers in dynamic traffic scenes to be validated by experiments with human drivers.

In our future work, we will expand our selection of hypothetical percepts and will explore the use of different and separated skill hierarchies for lateral and longitudinal control.

ACKNOWLEDGMENT

This work has been supported by the projects IMoST2 (Integrated Modeling for Safe Transportation II), sponsored by the Government of Lower Saxony, Germany under contracts ZN2245, ZN2253, ZN2366, and the project HoliDes (Holistic Human Factors and System Design of Adaptive Cooperative Human-Machine Systems), funded by the ARTEMIS Joint Undertaking Grant agreement no.: 332933.

REFERENCES

[1] C. Cacciabue, Ed., Modelling Driver Behaviour in Automotive Environments: Critical Issues in Driver Interactions with Intelligent Transport Systems. Springer, 2007.

[2] T. Jürgensohn, Modelling Driver Behaviour in Automotive Environments. Springer, 2007, ch. Control Theroy Models of the Driver, pp. 277–292.

[3] D. Salvucci and R. Gray, “A two-point visual control model of steering,” Perception, vol. 33, 2004, pp. 1233–1248.

[4] R. Wilkie and J. Wann, “Controlling steering and judging heading: Retinal flow, visual direction and extra-retinal information,” Journal of Experimental Psychology: Human Perception and Performance, vol. 29, 2003, pp. 363–378.

[5] D. Salvucci, Integrated Models of Cognitive Systems. New York: Oxford University Press, 2007, ch. Integrated Models of Driving Behavior, pp. 356–367.

[6] P. Bessière, C. Laugier, and R. Siegwart, Eds., Probabilistic Reasoning and Decision Making in Sensory-Motor Systems. Berlin: Springer, 2008.

[7] C. Möbus and M. Eilers, “Further steps towards driver modelling according to the Bayesian programming approach,” in Digital Human Modeling, ser. LNCS (LNAI). San Diego: Springer, 2009, pp. 413–422.

[8] C. Möbus and M. Eilers, “Mixture of behaviors and levels-of-expertise in a Bayesian autonomous driver model,” in Advances in Applied Digital Human Modeling, V. Duffy, Ed. Boca Raton: CRC Press, Taylor & Francis Group, 2010, pp. 425–435.

[9] M. Eilers and C. Möbus, “Learning of a Bayesian autonomous driver mixture-of-behaviors (bad-mob) model,” in Advances in Applied Digital Human Modeling, V. Duffy, Ed. Boca Raton: CRC Press, Taylor & Francis Group, 2010, pp. 436–445.

[10] M. Eilers and C. Möbus, “Learning the relevant percepts o modular hierarchical Bayesian driver models using a Bayesian information criterion,” in Digital Human Modelling, ser. LCNS 6777, V. Duffy, Ed. Heidelberg: Springer, 2011, pp. 463–472.

[11] M. Chattington, M. Wilson, D. Ashford, and D. Marple-Horvat, “Eye-steerig coordination in natural driving,” Experimental Brain Research, vol. 180, 2007, pp. 1–14.

[12] M. Land, Vision and Action. Cambridge: Cambridge University Press, 1998, ch. The Visual Control of Steering, pp. 163–180.

[13] M. Land and B. Tatler, Looking and Acting: Vision and eye movements in natural behavior. Oxford: Oxford University Press, 2009.

[14] L. Li and J. Chen, “Relative contribution of optic flow, bearing, and splay angle information to lane keeping,” Journal of Vision, vol. 10, 2010, pp. 1–14.

[15] R. Wilkie, D. Poulter, and J. Wann, “Where you look when you learn to steer,” Journal of Vision, vol. 4, no. 8, 2004, doi:10.1167/4.8.1.

[16] J. Wann and R. Wilkie, Optical flow and beyond. Norwell, MA, USA: Kluwer Academic Publishers, 2004, ch. How do we control high speed steering?, pp. 401–419.

[17] A. Chatziastros, G. Wallis, and H. Bülthoff, “The use of optical flow and splay angle in steering a central path,” Max Planck Institute for Biological Cybernetics, Tübingen, Germany, Tech. Rep. 71, October 1999.

[18] M. Land and D. Lee, “Where we look when we steer,” Nature, vol. 369, 1994, pp. 742–744.

[19] O. Lappi, E. Lehtonen, J. Pekkanen, and T. Itkonen, “Beyond the tangent point: Gaze targets in naturalistic driving,” Journal of Vision, vol. 13, 2013, pp. 1–18.

[20] D. Lee, “A theory of visual control of braking based on information about time to collision,” Perception, vol. 5, 1976, pp. 437–459.

[21] W. van Winsum, “The human element in car following models,” Tranportation Research Part F, vol. 2, 1999, pp. 207–211.

[22] M. Gouy, C. Diels, N. Reed, A. Stevens, and G. Burnett, “Preferred or adopted time headway? A driving simulator study,” in Proceedings of the International Conference on Ergonomnics & Human Factors, M. Anderson, Ed., 2013, pp. 153–159.

[23] D. Koller and N. Friedman, Probabilistic Graphical Models: Principles and Techniques. MIT Press, 2009.

[24] N. Friedman, K. Murphy, and S. Russell, “Learning the structure of dynamic probabilistic networks,” in Proceedings of the 14th conference on Uncertainty in artificial intelligence, 1998, pp. 139–147.

[25] W. Horrey, C. Wickens, and K. Consalus, “Modeling driver’s visual attention allocation while interacting with in-vehicle technologies,” Journal of Experimental Psychology, vol. 12, no. 2, 2006, pp. 67–78.

[26] P. Bessière et al., “Survey: Probabilistic methodolgy and techniques for artefact conception and development,” INRIA - Institut National de Recherce en Informatique et en Automatique, Tech. Rep., 2003.

[27] G. Schwarz, “Estimating the dimension of a model,” Ann. Stat., vol. 6, 1978, pp. 461–464.

[28] K. Murphy, Machine Learning: A Probabilistic Perspective. The MIT Press, 2012.

[29] Y. Guo and R. Greiner, “Discriminative model selection for belief net structures,” in Proceedings of the 20th National Conference on Artificial Intelligence, 2005, pp. 770–776.

[30] S. Natarajan, W.-K. Wong, and P. Tadepalli, “Structure refinement in first order conditional influence language,” in Proceedings of the Workshop on Open Problems in Statistical Relational Learning (SRL), 2006, 8 pages.

[31] G. Santafe, J. Lozano, and P. Larranaga, “Discriminative vs. generative learning of Bayesian network classifiers,” in ECSQARU 2007, ser. LNAI 4724, K. Mellouli, Ed. Berlin, Heidelberg: Springer, 2007, pp. 453–464.

[32] R. Wilkie and J. Wann, “Eye-movements aid the control of locomotion,” Journal of Vision, vol. 3, no. 11, 2003, pp. 677–684.

[33] D. Salvucci, E. Boer, and A. Liu, “Toward an integrated model of driver behavior in a cognitive architecture,” Transportation Research Record, vol. 1779, 2001, pp. 9–16.

[34] D. Salvucci and A. Liu, “The time course of a lane change: Driver control and eye-movement behavior,” Tranportation Research Part F, vol. 5, 2002, pp. 123–132.