



HAL
open science

Interaction-Aware Driver Maneuver Inference in Highways Using Realistic Driver Models

David Sierra González, Víctor Romero-Cano, Jilles Steeve Dibangoye,
Christian Laugier

► **To cite this version:**

David Sierra González, Víctor Romero-Cano, Jilles Steeve Dibangoye, Christian Laugier. Interaction-Aware Driver Maneuver Inference in Highways Using Realistic Driver Models. Proceedings of the 2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC 2017), Oct 2017, Yokohama, Japan. hal-01589493

HAL Id: hal-01589493

<https://inria.hal.science/hal-01589493>

Submitted on 18 Sep 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Interaction-Aware Driver Maneuver Inference in Highways Using Realistic Driver Models

David Sierra González¹, Víctor Romero-Cano², Jilles Steeve Dibangoye¹, and Christian Laugier¹

Abstract—In this work, we address the problem of lane change maneuver prediction in highway scenarios using information from sensors and perception systems widely used in automated driving. Our prediction approach is two-fold. First, a driver model learned from demonstrations via Inverse Reinforcement Learning is used to equip a host vehicle with the anticipatory behavior reasoning capability of common drivers. Second, inference on an interaction-aware augmented Switching State-Space Model allows the approach to account for the dynamic evidence observed. The use of a driver model that correctly balances the driving and risk-averse preferences of a driver allows the computation of a planning-based maneuver prediction. Integrating this anticipatory prediction into the maneuver inference engine brings a degree of scene understanding into the estimate and leads to faster lane change detections compared to those obtained by relying on dynamics alone. The performance of the presented framework is evaluated using highway data collected with an instrumented vehicle. The combination of model-based maneuver prediction and filtering-based state and maneuver tracking is shown to outperform an Interacting Multiple Model filter in the detection of highway lane change maneuvers regarding accuracy, detection latency—by an average of 0.4 seconds—and false-positive rates.

I. INTRODUCTION

Interacting with human agents is one of the major challenges faced by Advanced Driver Assistance Systems (ADAS) and autonomous driving vehicles. The difficulty stems from the fact that human behavior is determined by a complex set of interdependent factors, which are very hard to model (e.g., intentions, perception, emotions). As a consequence, most existing systems build upon simple short-term motion models assuming constant velocity or acceleration [1]. In spite of the complexity of the problem, human drivers are extremely good at predicting the intentions of surrounding vehicles. That is due to their innate ability to: 1) interpret driver motion cues, and 2) put themselves in the place of other drivers to reason over their most likely risk-averse behavior.

In this paper, we present a probabilistic framework to estimate and predict the development of traffic situations. Our approach integrates the two key points mentioned above to produce accurate human-like predictions. In the first place, we make inference on an interaction-aware augmented Switching State-Space Model (aSSSM) [2], [3], which allows

us to take into account the motion cues of drivers. Additionally, a driver model learned from demonstrated highway driving data using Inverse Reinforcement Learning (IRL) is used to identify the most likely, risk-averse, anticipatory maneuver for each driver in the scene. By integrating both sources of information, the main drawbacks of both model- and dynamics-based vanilla approaches are overcome. Model-based scene prediction approaches typically assume a risk-minimizing behavior (encoded in the model) for all agents, whether it is selfish [4] or altruist [5], [6]. Hence, these approaches fail to predict dangerous maneuvers that the model cannot explain (e.g., a dangerous lane change maneuver in the highway). In contrast, dynamics-based prediction approaches proceed by identifying the motion model that best fits the current dynamics of a maneuvering target without any reasoning about the traffic scene (i.e., they disregard the interactions between vehicles) [7], [8], [9]. In consequence, they are limited to predicting only the most immediate future and are prone to false positives under real-world conditions (in the context of highway lane change prediction, consider the response of one such system to the deviations from the lane center characteristic of human lane keeping). Our approach brings the best of both worlds, providing accurate short- and long-term predictions by combining dynamic evidence with the scene understanding brought in by the model.

In this work, we specifically apply the proposed framework to the estimation of lane change maneuvers in 2-lane highways. Particular attention is given to studying the influence of the model-based prediction in the final maneuver estimate. The performance is evaluated on a dataset gathered using an instrumented vehicle in French highways and compared qualitatively and quantitatively against an Interacting Multiple Model filter [9].

The rest of the document is organized as follows: Section II surveys the existing literature in the field of motion prediction for Intelligent Vehicles; Section III describes in detail the proposed approach; Section IV presents the different test scenarios and shows the qualitative and quantitative experimental results. Lastly, Section V concludes.

II. RELATED WORK

Research in the field of behavior inference and motion prediction for Intelligent Vehicles has intensified significantly in the past decade. Approaches are typically classified into three main groups: Physics-based, Maneuver-based and Interaction-aware models [10].

This work was supported by Toyota Motor Europe.

¹The authors are with Inria Grenoble Rhône-Alpes, 38330 Montbonnot-Saint-Martin, France - {david.sierra-gonzalez, jilles.dibangoye, christian.laugier}@inria.fr

² Víctor Romero-Cano is with Universidad Autónoma de Occidente, Cali, Colombia varomero@uao.edu.co. He contributed to this work while he was part of Inria Grenoble Rhône-Alpes.

In Physics-based approaches, the future motion is predicted using dynamic or kinematic models, limiting their applicability to the immediate-term prediction. Approaches based on the Interacting Multiple-Model filter (IMM) are included in this group [7], [9]. They rely on a bank of Kalman Filters to track the state of the target, yielding as a byproduct a distribution over the current dynamics of the target. The switching between dynamics is Markovian, causing the performance to depend heavily on a proper tuning of the regime transition matrix. Furthermore, they fail to consider the interactions between the traffic participants in the switching process.

In Maneuver-based approaches, the future trajectory of the target is predicted by identifying the maneuver in execution from a finite set of maneuvers contained in a database. Methods to identify the maneuver include Hidden Markov Models (HMM) [11] and Gaussian Processes (GP) [12]. Typically, these approaches also fail to consider the interactions between vehicles.

In practice, the behavior of the traffic participants is highly interdependent. Dynamic Bayesian Networks (DBN) can explicitly formalize the interactions between traffic participants in the long-term prediction of traffic situations [13]. This approach infers the behavior of each vehicle by identifying its current situation (e.g., close to the vehicle in front) as a function of its local situational context (e.g., distance to the vehicle in front). This method makes much sense conceptually but relies strongly on particle filtering, which may limit its applicability in complex real-world scenarios.

Agamennoni et al. presented another approach following the same line of thought [3]. In this approach, high-level discrete contexts determine the evolution of low-level dynamics. In contrast to the IMM approaches discussed above, the switching process is only conditionally Markov, also depending on the continuous state at the preceding time step through a linear feature-based function. The major difference with [13] lies in the inference engine. Instead of relying on particle filtering, approximate inference is performed using a variation of the forward pass for augmented Switching Linear Dynamical Systems (aSLDS) presented in [2]. The approach is evaluated on its ability to track interacting trucks in an open-cast mine. The model describing the interactions is fairly straightforward, essentially building upon the time-to-collision and the right-of-way. Finally, this method learns a different model for each high-level context.

More sophisticated, full-fledged, interaction-aware driver models have been used in [14], [4]. In [14], a cost-function that models the preferences and the risk-averse behavior of drivers is used to estimate the interaction-aware motion intention of each vehicle in the scene. The model-based prediction is then complemented with a learning-based prediction to account for physical evidence of lane changes. In this approach, most of the cost-function terms correspond to those presented in [15]. No information is available on how to tune the parameters or balance the weights of the cost function. In [4], in contrast, the balance between the different driver model terms is learned automatically from demon-

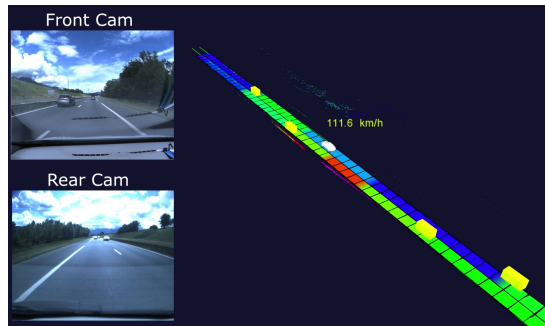


Fig. 1: Slice of the cost function that encodes the risk-averse behavior of drivers superimposed on the ground on an exemplary highway situation. Red: high cost; green: intermediate; blue: low.

strations using IRL. A dynamic model learned from driving demonstrations is shown to produce human-like, interaction-aware, long-term predictions in highway scenarios. However, this approach is unable to predict any behavior not explained by the risk-averse model (e.g., keeping the lane motivated by a long-term goal or a dangerous lane change).

The proposed approach builds upon [4] and [3], adopting the model-based behavior inference of the former and the DBN structure and inference engine of the latter. By planning over a complex driver model that can handle lane changes, we successfully infer the most likely, interaction-aware behavior of highway drivers. This anticipatory prediction is then used in the switching process of the state-space model, leading to faster detections of new maneuvers and an increased robustness in the estimations.

III. MANEUVER ESTIMATION FRAMEWORK

A. Driver modeling

To learn the feature-based cost function describing the preferences of a driver from driving demonstrations, we use a three-fold method. First, we model the dynamics of the driver as a Markov Decision Process (MDP). Then, we select a number of relevant features using background knowledge. And finally, we use the Maximum Entropy IRL algorithm [16] to learn the balance between the different terms of the cost function according to the behavior demonstrated. More details can be found in [4]. The cost function that results from this method is linear on the features:

$$\mathcal{C}(s) = \boldsymbol{\theta} \cdot \mathbf{f}(s) \quad (1)$$

where $\boldsymbol{\theta} = (\theta_1, \dots, \theta_K)$ is the weight vector and $\mathbf{f}(s) = (f_1(s), \dots, f_K(s))$ is the feature vector that parameterizes state s .

In order for the model to capture the preferences and the risk-averse behavior of a highway driver, we require the following features: lane; speed deviation; and time-headway.

a) Lane: aims to capture the preference of an agent to drive on a particular lane.

b) Speed deviation: encodes the penalty of deviating from the agent's desired speed, which is set to the maximum speed reached by the agent since the last change in the speed limit.

c) *Time-headway*: defines the time elapsed between the back of the lead vehicle passing a point and the front of the following vehicle passing the same point. It indicates potentially dangerous situations. We consider two separate features associated to the time-headway: the time-headway to the closest car in front in the current lane and the time-headway to the closest trailing car in the current lane.

For further details regarding the shape of the feature function components, the reader can refer to [4]. Fig. 1 shows an example of a cost function learned from demonstrations using these features. In this 2-lane highway scene, we illustrate the slice of the cost function associated to the current speed of the vehicle of interest. At the current speed, a lane change maneuver would be discouraged as it leads to high-cost states caused by the traffic in the neighboring lane.

B. Overview and Notation

We propose a state and maneuver estimation framework based on a combination of discrete model-based maneuver prediction and discrete-continuous Bayesian filtering. More generally, the probabilistic model proposed can be categorized as a Switching State Space Model (SSSM), in which a high-level layer reasons about the maneuvers being performed by the different interacting vehicles and determines the evolution of the low-level dynamics. The high-level reasoning is performed through the fusion of: 1) an anticipatory, interaction-aware, model-based prediction; and 2) probabilistic messages coming from the bottom nodes of our SSSM and representing evidence about low level dynamics. Fig. 2 shows the graphical model that specifies the conditional independence assumptions of our model. Bold arrows indicate multi-vehicle dependencies. Focusing on the slice of the graphical model for vehicle i , we can observe three layers of abstraction:

- The highest level corresponds to the maneuver m_t^i being executed by the vehicle. This is a discrete hidden random variable. The different maneuvers considered in this work and their corresponding dynamics are specified in subsection III-D.
- The second level describes the state of the vehicle in a curved road frame through the continuous state vector \mathbf{x}_t^i . The state is not directly observable.
- Finally, the shaded nodes in the graphical model are the observations. The observation vector \mathbf{z}_t^i is a noisy measurement of the state \mathbf{x}_t^i .

The factorization of the joint distribution given the model assumptions is the following:

$$P(\mathbf{x}_{1:T}^{1:N}, m_{1:T}^{1:N}, \mathbf{z}_{1:T}^{1:N}) = P(\mathbf{x}_1^{1:N}, m_1^{1:N}, \mathbf{z}_1^{1:N}) \prod_{t=2}^T \prod_{i=1}^N [P(\mathbf{x}_t^i | m_{t-1:t}^i, \mathbf{x}_{t-1}^{1:N}) P(m_t^i | m_{t-1}^{1:N}, \mathbf{x}_{t-1}^{1:N}) P(\mathbf{z}_t^i | \mathbf{x}_t^i)] \quad (2)$$

where the notation $\mathbf{x}_{1:T}^{1:N}$ is shorthand for the tuple $(\mathbf{x}_1^1, \dots, \mathbf{x}_T^1, \dots, \mathbf{x}_1^N, \dots, \mathbf{x}_T^N)$, T indicates the number of timesteps considered and N denotes the number of vehicles involved.

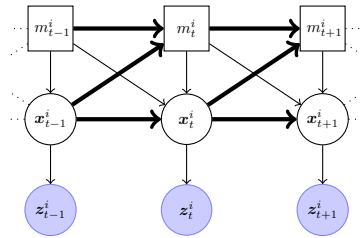


Fig. 2: Graphical model representation of the proposed switching state-space model. A slice of the Bayesian network corresponding to vehicle i is shown. Bold arrows indicate multi-vehicle dependencies.

The term $P(\mathbf{x}_t^i | m_{t-1:t}^i, \mathbf{x}_{t-1}^{1:N})$ describes the dynamic evolution of the state of a target given the distribution over maneuvers at the current and previous timesteps, and the previous states of all vehicles. Detailed descriptions of the motion model considered and of the dynamics associated to the different maneuvers are given in subsections III-C and III-D respectively.

The maneuvering behavior of drivers is described in the predictive term $P(m_t^i | m_{t-1}^{1:N}, \mathbf{x}_{t-1}^{1:N})$. The probability of a driver choosing a maneuver depends heavily on the state of the other traffic participants. We take advantage of the risk-averse IRL driver model to take into account these interactions and to forecast the most likely distribution over maneuvers at each timestep. This process is detailed in section III-E.

The term $P(\mathbf{z}_t^i | \mathbf{x}_t^i)$ is the measurement model and relates the hidden states of the vehicles in the scene with the observations through a linear measurement equation with added Gaussian noise $\mathbf{z}_t^i | \mathbf{x}_t^i \sim \mathcal{N}(\mathbf{C}\mathbf{x}_t^i, \mathbf{R})$, where $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes a multivariate Gaussian distribution with mean vector $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$; \mathbf{C} is the output matrix, and \mathbf{R} is the covariance of the Gaussian observation noise.

To complete the description of the framework we present the inference engine to estimate the filtered posterior probability distribution over states and maneuvers for all agents in subsection III-F.

C. Vehicle dynamics and maneuvers

We assume that vehicles move on a 2D environment and that the configuration of a vehicle i in a curved road frame is given by the state vector $\mathbf{x}_t^i = [x, y, \psi, v, \omega]^T \in \mathbb{R}^5$, where x and y are the target coordinates, v is the vehicle's absolute linear speed along its direction of travel ψ , and ω is the yaw rate. All variables in the state vector except for the yaw rate are observed. We assume that the dynamics satisfy the following differential equations:

$$\begin{aligned} \dot{x}(t) &= v(t) \cos \psi(t) \\ \dot{y}(t) &= v(t) \sin \psi(t) \\ \dot{\psi}(t) &= \omega(t) \\ \dot{v}(t) &= a_{idm} \\ \dot{\omega}(t) &= 0 \end{aligned} \quad (3)$$

where we have neglected the influence of the yaw rate ω on the position. The term a_{idm} is the longitudinal acceleration of the target, which is set using the Intelligent Driver Model (IDM) [17]. Hence, the acceleration is calculated at each timestep as a function of the states of all the vehicles in the scene. A maneuver-dependent predictive function g_λ integrates (3) over an interval of time Δt to obtain the next state of the vehicle. The motion is assumed to be perturbed by Gaussian noise to account for the maneuver-specific modeling errors:

$$\mathbf{x}_t^i | \mathbf{x}_{t-1}^{1:N}, m_t^i = M_\lambda \sim \mathcal{N}(g_\lambda(\mathbf{x}_{t-1}^{1:N}), \mathbf{Q}_\lambda) \quad (4)$$

where \mathbf{Q}_λ is the noise covariance matrix, which is dependent on the maneuver being performed.

Notice that it makes a lot of sense to have the motion being dependent not only on the maneuver being performed, but on the states of the other traffic participants. For instance, a lane keeping maneuver involves adjusting the acceleration of the vehicle depending on the state of the preceding traffic.

D. Maneuvers

In this work, we consider only two possible high-level maneuvers for the vehicles in the scene: lane keeping and lane changing. Their acceleration is set, as described above, using the IDM [17]:

$$a_{idm} = a_0 \left[1 - \left(\frac{v_{t-1}}{v_0} \right)^\delta - \left(\frac{s^*(v_{t-1}, \Delta v_{t-1})}{s_{t-1}} \right)^2 \right] \quad (5)$$

where the desired minimum gap s^* is given by:

$$s^*(v, \Delta v) = s_0 + vT + \frac{v\Delta v}{2\sqrt{a_0 b_0}} \quad (6)$$

In (5) and (6), s represents the gap with the leading vehicle, a_0 the maximum acceleration, b_0 the maximum comfortable deceleration, v_0 the desired speed, T the desired front time-headway, s_0 the traffic jam minimum distance, and $\Delta v = v - v_l$ the velocity difference with the leading vehicle. The parameter δ controls the change in acceleration as a function of velocity.

The IDM encodes different driving modes such as a free-road acceleration strategy when the gap to the leading vehicle is sufficiently large, or an ‘‘intelligent’’ braking strategy that keeps the deceleration below the comfort threshold b_0 unless an emergency situation takes place. More details can be found in [17].

The details regarding the dynamics of each of the two available maneuvers are detailed below:

1) *Lane change (LC)*: The target vehicle turns and moves towards the neighboring lane. The dynamics of this maneuver are perfectly specified with the differential equations (3) and its process noise covariance matrix \mathbf{Q}_{LC} .

2) *Lane keeping (LK)*: The vehicle of interest remains on its lane, aligned with the direction of the road, and driving at its desired speed unless it is slowed down by the leading vehicle. We do not adopt the constant yaw assumption of other approaches for this maneuver [9]. Due to the particularities of the inference engine used in this work,

the constant yaw assumption does not suffice to tell apart the LK and the LC dynamics. Instead, we parametrize our lane keeping motion model using (3) and the corresponding process noise covariance matrix \mathbf{Q}_{LK} , along with an artificial observation ω' for the yaw rate:

$$\omega' = -\omega_{max}(\psi(t)/\psi_{max}) \quad (7)$$

The intuition behind this artificial observation is simple: if a target vehicle that is performing a lane keeping maneuver has a high yaw (in road coordinates, i.e. it is not aligned with the road), a steering action is expected in order to re-align the vehicle with the road. The magnitude of the expected yaw rate will be proportional to the misalignment of the target with the road, and is parametrized by ψ_{max} and ω_{max} , which have been obtained experimentally.

E. Interaction-Aware, Model-Based Maneuver Forecasting

By means of the risk-averse driver model obtained with IRL, we can forecast the probability of each driver’s next maneuver in response to the states and maneuvers of the other traffic participants. This full-fledged single driver model contrasts with the approach presented in [3], where each maneuver has an associated simple feature-based model. The driver model used here balances the (navigational and risk) preferences of drivers and enables us to predict their anticipatory behavior. A driver will perform a maneuver at the current timestep if, given his prediction for the behavior of the surrounding drivers, this leads to a sequence of F future states that agree with his/her preferences (encoded in the driver model):

$$P(m_{t+1}^i = M | \mathbf{x}_t^{1:N}, m_t^{1:N}) \propto \quad (8)$$

$$\mathbb{E}_{(\mathbf{x}_t^{1:N}, m_t^{-i}) \sim P(\mathbf{x}_t^{1:N}, m_t^{1:N} | z_t^{1:N})} \left[1 - \frac{\sum_{k=0}^F \theta^T \mathbf{f}(\mathbf{x}_{t+k, M}^i)}{\sum_{M'} \sum_{k=0}^F \theta^T \mathbf{f}(\mathbf{x}_{t+k, M'}^i)} \right]$$

where we have overloaded the notation for the state to explicitly indicate the maneuver being used to propagate it between timesteps, and the notation m^{-i} indicates the maneuvers for all agents except agent i . The expectation is taken with respect to the posterior at the previous timestep, which factorizes across agents.

Fig. 3 illustrates the procedure encoded in (8). At any given timestep we sample a state and a maneuver from the posterior distribution of each traffic participant. Vehicles other than the vehicle of interest are propagated forward F timesteps according to the sampled maneuver. The lane change maneuver is propagated by setting a predefined angular speed ω_{LC} until the yaw reaches a given threshold. The lane keeping maneuver fixes the relative lateral position of the vehicle of interest and propagates it only longitudinally. When the maneuver being performed is a lane change and the vehicle reaches the centerline of the neighboring lane, the maneuver is switched automatically to lane keeping.

In contrast, the vehicle of interest is propagated forward F timesteps with all available maneuvers. For each of them, the cost associated to the resulting sequence of states is

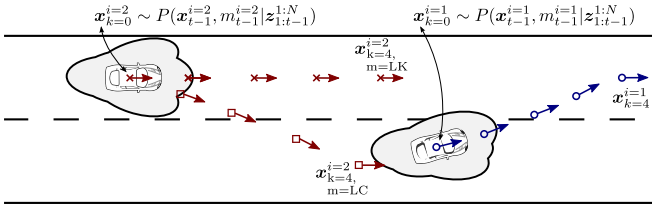


Fig. 3: Schematic showing the anticipatory, model-based maneuver forecasting approach.

calculated. Averaging over all samples, the maneuver producing the lowest sum of costs will be more likely than the others, capturing thus the planning-based behavior of drivers. As an example, we can consider a vehicle being slowed down behind a truck in the highway and considering whether to keep driving behind the truck or to overtake. Using the features presented in subsection III-A, we know that if the driver keeps driving behind the truck he will be penalized for deviating from his desired speed. In contrast, a lane change will mean a small penalty for not driving in the right-most lane but will enable the driver to accelerate, minimizing thus the cost due to speed. This example highlights the importance of: 1) The need to consider a sufficiently large planning-horizon when evaluating different maneuvers; and 2) Using a cost function that realistically balances the different features.

F. Approximate Inference

Exact inference of the posterior $P(\mathbf{x}_t^i, m_t^i | \mathbf{z}_{1:t}^i)$ is intractable in both the standard SLDS and in the aSLDS, scaling exponentially with time [2], [3]. The proposed approximate inference engine is similar to the filtering approach presented in [2], with an extension to account for non-linear dynamics. Such an engine was already used in [3], although with a non-Gaussian emission model. Approximate inference is performed individually per agent, which makes the algorithm highly parallelizable. The key idea is to approximate the intractable posterior with a simpler distribution (a Gaussian mixture), and to establish a recursion to track it over time. The posterior can be decomposed as:

$$P(\mathbf{x}_t^i, m_t^i | \mathbf{z}_{1:t}^i) = P(\mathbf{x}_t^i | m_t^i, \mathbf{z}_{1:t}^i) P(m_t^i | \mathbf{z}_{1:t}^i) \quad (9)$$

A recursion is established for each of the terms on the r.h.s. The first term is approximated with a Gaussian mixture distribution with C components:

$$P(\mathbf{x}_t^i | m_t^i, \mathbf{z}_{1:t}^i) \approx \sum_{c_t=1}^C P(\mathbf{x}_t^i | c_t, m_t^i, \mathbf{z}_{1:t}^i) P(c_t | m_t^i, \mathbf{z}_{1:t}^i)$$

where the term $P(\mathbf{x}_t^i | c_t, m_t^i, \mathbf{z}_{1:t}^i)$ is a Gaussian with mean $f(c_t, m_t^i)$ and covariance $F(c_t, m_t^i)$, and the term $P(c_t | m_t^i, \mathbf{z}_{1:t}^i)$ indicates the weight of the mixture component. The first recursion is then established as:

$$\begin{aligned} P(\mathbf{x}_{t+1}^i | m_{t+1}^i, \mathbf{z}_{1:t+1}^i) &= \sum_{m_t^i, c_t} P(\mathbf{x}_{t+1}^i | m_t^i, c_t, m_{t+1}^i, \mathbf{z}_{1:t+1}^i) \\ &= \sum_{m_t^i, c_t} P(\mathbf{x}_{t+1}^i | c_t, m_{t+1}^i, \mathbf{z}_{1:t+1}^i) P(m_t^i, c_t | m_{t+1}^i, \mathbf{z}_{1:t+1}^i) \end{aligned}$$

The term $P(\mathbf{x}_{t+1}^i | c_t, m_{t+1}^i, \mathbf{z}_{1:t+1}^i)$ is obtained by propagating forward with all the available dynamics m_{t+1}^i each component of the Gaussian mixture composing $P(\mathbf{x}_t^i | m_t^i, \mathbf{z}_{1:t}^i)$, and conditioning on the new observation \mathbf{z}_{t+1}^i . This leads indeed to an exponential increase in the number of Gaussian components that will need to be collapsed back to C components per maneuver. The prediction and update steps are performed using an Extended Kalman Filter (EKF). For clarity, we adopt the following shorthand notation for the indexes: $\alpha \equiv m_t^i$, $\beta \equiv m_{t+1}^i$. Hence, the non-linear predictive function \mathbf{g}_β is associated to the dynamics of the maneuver m_{t+1}^i . In this setting, the EKF equations are:

$$\begin{aligned} \bar{\boldsymbol{\mu}}_{t+1}^{(\alpha\beta c_t)} &= \mathbf{g}_\beta(f(c_t, \alpha)) \\ \bar{\boldsymbol{\Sigma}}_{t+1}^{(\alpha\beta c_t)} &= \nabla \mathbf{g}_\beta(f(c_t, \alpha))^T F(c_t, \alpha) \nabla \mathbf{g}_\beta(f(c_t, \alpha)) + \mathbf{Q}_\beta \\ \mathbf{S}_{t+1}^{(\alpha\beta c_t)} &= \mathbf{C} \bar{\boldsymbol{\Sigma}}_{t+1}^{(\alpha\beta c_t)} \mathbf{C}^T + \mathbf{R}_\beta \\ \mathbf{K}_{t+1}^{(\alpha\beta c_t)} &= \bar{\boldsymbol{\Sigma}}_{t+1}^{(\alpha\beta c_t)} \mathbf{C}^T (\mathbf{S}_{t+1}^{(\alpha\beta c_t)})^{-1} \\ \boldsymbol{\mu}_{t+1}^{(\alpha\beta c_t)} &= \bar{\boldsymbol{\mu}}_{t+1}^{(\alpha\beta c_t)} + \mathbf{K}_{t+1}^{(\alpha\beta c_t)} (\mathbf{z}_{t+1}^i - \mathbf{C} \bar{\boldsymbol{\mu}}_{t+1}^{(\alpha\beta c_t)}) \\ \boldsymbol{\Sigma}_{t+1}^{(\alpha\beta c_t)} &= (\mathbf{I} - \mathbf{K}_{t+1}^{(\alpha\beta c_t)} \mathbf{C}) \bar{\boldsymbol{\Sigma}}_{t+1}^{(\alpha\beta c_t)} \end{aligned} \quad (10)$$

where $\mathbf{K}_{t+1}^{(\alpha\beta c_t)}$ represents the Kalman gain and $\mathbf{S}_{t+1}^{(\alpha\beta c_t)}$ the innovation covariance. To obtain the weights of the new mixture components we consider:

$$\begin{aligned} P(m_t^i, c_t | m_{t+1}^i, \mathbf{z}_{1:t+1}^i) &\propto P(\mathbf{z}_{t+1}^i | c_t, m_{t+1}^i, \mathbf{z}_{1:t}^i) \\ P(m_{t+1}^i | c_t, m_t^i, \mathbf{z}_{1:t}^i) &P(c_t | m_t^i, \mathbf{z}_{1:t}^i) P(m_t^i | \mathbf{z}_{1:t}^i) \end{aligned} \quad (11)$$

The terms $P(c_t | m_t^i, \mathbf{z}_{1:t}^i)$ and $P(m_t^i | \mathbf{z}_{1:t}^i)$ are available from the previous step in the recursion; the term $P(\mathbf{z}_{t+1}^i | c_t, m_{t+1}^i, \mathbf{z}_{1:t}^i)$ is the likelihood of the observation \mathbf{z}_{t+1}^i under the Gaussian $\mathcal{N}(\mathbf{C} \bar{\boldsymbol{\mu}}_{t+1}^{(\alpha\beta c_t)}, \mathbf{S}_{t+1}^{(\alpha\beta c_t)})$, and the prior $P(m_{t+1}^i | c_t, m_t^i, \mathbf{z}_{1:t}^i)$ is calculated by means of the driver model as seen in subsection III-E. This is where the fusion between dynamics- and model-based estimates takes place.

Coming back to (9), the recursion for the second term on the r.h.s. is established as:

$$P(m_{t+1}^i | \mathbf{z}_{1:t+1}^i) \propto \sum_{c_t, m_t^i} P(c_t, m_t^i, m_{t+1}^i, \mathbf{z}_{1:t+1}^i) \quad (12)$$

which can be calculated with the same components from (11). Finally, at each step of the inference procedure, the number of mixture components have increased from $|m|C$ to $|m|C^2$, each with a weight $P(m_t^i, c_t | m_{t+1}^i, \mathbf{z}_{1:t+1}^i)$. The C^2 components associated to each maneuver m_{t+1}^i are collapsed back to C components using a simple procedure: we retain the $|C-1|$ components with the largest mixture weights and merge the remaining to a single Gaussian.

IV. EXPERIMENTAL EVALUATION

We have presented a filtering framework that keeps track of the posterior distribution over states and maneuvers of any target within sensor range of the ego-vehicle in highway scenarios. Low-level dynamic evidence is combined with a model-based, risk-averse prediction, as shown in (11). By

fusing the dynamic evidence with an anticipatory, model-based prediction some degree of scene understanding is brought into the estimate. In this section, we focus on analyzing: 1) how both sources of information contribute to the final maneuver estimate; and 2) the ability of the proposed framework to estimate lane change maneuvers in real highway scenes.

To perform the analysis, we have gathered data on a French 2-lane highway using an instrumented vehicle equipped with a Velodyne HDL-64E LiDAR sensor, an Xsens MTi-100 IMU, and forward- and rear-facing IDS cameras. By means of a grid-based target tracker [18] and a lane tracker, we are able to localize the targets with respect to the ego-vehicle (EV) and the road network. A total of 10 lane-change highway scenes have been annotated, accounting for a total of 152s of driving data. The selected scenes include a variety of situations with 6 cut-in and 4 cut-out lane changes, taking place both behind (3) and in front (7) of the EV.

This section is organized as follows: first, the values of all the framework parameters used in the evaluation are presented. Next, the maneuver filtering performance of our approach is analyzed qualitatively on two of the traffic scenes of the dataset. Finally, a quantitative analysis is presented in order to be able to compare the lane-change detection performance of our framework against alternative approaches.

A. Framework parameters

The values of the framework parameters used in our evaluation are discussed here in roughly the same order that they have been introduced in the document. In the first place, the driver model weight parameters have been obtained by running the MaxEnt IRL algorithm on a dataset of 2 minutes of highway driving demonstrations performed on a simulator, as described in [4].

The values for the rest of parameters are shown in table I. The parameters for the IDM are similar to those suggested by Treiber et al. [17] with the exception of the desired time-headway T , which has been reduced from 1.6s to 1.0s.

The noise covariance matrices have been set using the discrete white noise approach [8], [19]. The process noise standard deviations to account for the constant acceleration and constant yaw-rate assumptions have been adopted from [9] and are shown in table I. The parameters ω_{max} and ψ_{max} used to generate the synthetic observations that encode the lane keeping behavior have been chosen empirically, as well as the components of the measurement noise covariance matrix \mathbf{R} , which is set to be diagonal.

The time horizon T_h and timestep dt for the model-based maneuver prediction were set by considering the trade-off between foresight, accuracy, and computation time. Finally, the number of Gaussian mixture components used during filtering has been set to $C = 3$. With this configuration, each filtering step took an average of 53ms in a Python implementation of our framework executed on an Intel i7-6700HQ CPU running at 2.60GHz.

TABLE I: Framework parameters

IDM	$a_0 = 1.5 \text{ m/s}^2$ $s_0 = 2 \text{ m}$	$b_0 = 1.67 \text{ m/s}^2$ $\delta = 4$	$T = 1.0 \text{ s}$
Noise LK	$\sigma_a = 4.0 \text{ m/s}^3$	$\sigma_w = 0.0205 \text{ rad/s}^2$	
Noise LC	$\sigma_a = 4.0 \text{ m/s}^3$	$\sigma_w = 0.15 \text{ rad/s}^2$	
LK dynam.	$\psi_{max} = 0.04 \text{ rad}$	$\omega_{max} = 0.28 \text{ rad/s}$	
R	$\sigma_x = 0.2 \text{ m}$ $\sigma_v = 0.2 \text{ m/s}^2$	$\sigma_y = 0.2 \text{ m}$ $\sigma_{w'} = 0.06 \text{ rad/s}^2$	$\sigma_\psi = 0.01 \text{ rad}$
Model pred.	$T_h = 3 \text{ s}$	$dt = 0.1 \text{ s}$	
Filtering	$C = 3$		

B. Qualitative analysis

We analyze the performance of our framework on two of the annotated highway scenes and compare the results obtained against the IMM-based lane change prediction approach from [9] (Model-set A).


Fig. 4 shows the results obtained for the two scenes. In the first one (left column), the target vehicle drives behind the EV until the left lane is free of traffic and then performs a lane change to overtake. The first row in Fig. 4c shows the lateral position of the target in road coordinates (the center lane-marking corresponds to $y = 0$); the second row shows the relative longitudinal position of the target with respect to the EV. As can be seen in the third row, the model-based prediction initially concedes a low probability to the LC maneuver due to the presence of traffic in the neighboring lane. Once the left lane becomes free, the probability for a LC begins to grow as the distance with the preceding traffic increases. The increase in the LC maneuver probability is due to the cost penalty induced by driving behind the EV at a speed lower than the target's desired speed, which had been set earlier during the target's approach.

Our framework's dynamics-based maneuver filtering estimate is shown in the fourth row of Fig. 4c. This estimate is obtained by using an uninformative model prediction in (11), i.e. a model that assigns the same probability to all maneuvers. The lane change maneuver is detected only 0.51s after it begins and roughly 1.5s before the target crosses the lane marking. This is 0.2s faster than the IMM-based estimate, which is shown in the last row. Our framework's final maneuver estimate is shown in the fifth row. As it can be seen, the effect of the model-based prediction in the LC maneuver estimate is a slightly quicker detection both of the beginning and of the end of the maneuver.

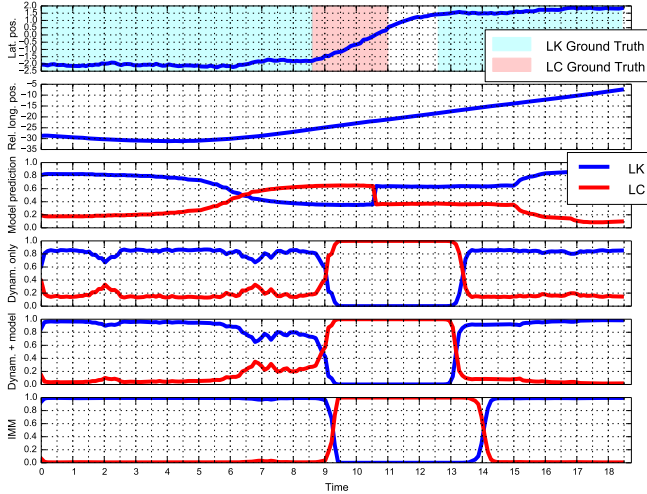
The right column in Fig. 4 shows the filtering results for a different scene, in which the target vehicle overtakes the EV and then merges back to the right lane. Focusing on Fig. 4f, we can observe again how integrating the model-based prediction into the filtering framework results in a quicker detection of the start and the end of the LC maneuver. An additional effect can be observed in this figure. The slight oscillations in the lateral position typical of a human-driven vehicle following a highway lane can be observed here around $t = 3.5s$. As a consequence of this lateral movement, both the dynamics-based and the IMM LC estimate produce a false positive. This does not occur in the final estimate of our approach, since the model-based prediction assigns a



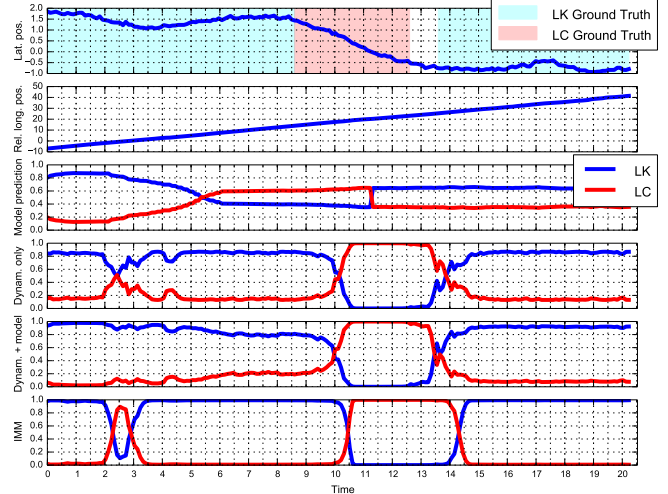
(a) Back view at $t = 2.0$ s. (b) Back view at $t = 10.6$ s.



(d) Front view at $t = 4.0$ s. (e) Front view at $t = 11.2$ s.



(c) Filtering results for scene 3.



(f) Filtering results for scene 7.

Fig. 4: Maneuver filtering results obtained for a rear cut-out lane change (left column) and a forward cut-in lane change (right column).

very low probability to a lane change maneuver around $t = 3$, moment in which the target drives right next the EV. This highlights the potential of our approach to reduce the number of false positives through the scene understanding brought in by the model-based prediction. However, does it imply that our framework cannot predict dangerous maneuvers? To answer this question we fixed the output of the model-based prediction to $P(m_t = \text{LK}) = 0.8$ and $P(m_t = \text{LC}) = 0.2$, and obtained again the filtered maneuver estimate for the two scenes discussed, showing the results in Fig. 5. It can be observed that the lane changes are still detected, with only a slight delay with respect to the estimate obtained with the true model prediction. The detections are still as fast as those of the IMM approach.

C. Quantitative analysis

In order to produce a quantitative analysis, we can interpret the filtered estimate as the output of a binary classifier in which the LC maneuver is the positive class, and each timestep corresponds to a classification sample. A prediction is counted as positive when $P(m_t = \text{LC}) > 0.5$. We have manually annotated the maneuver ground truth for the 10 test scenes, as shown in the first row of figures 4c and 4f. In the annotation, the lane change is assumed to start at the moment in which the target begins to move towards the neighboring lane. The end of the LC maneuver is assumed to occur when the center of mass of the target has reached $0.5m$ into the new lane. The LK maneuver is resumed once

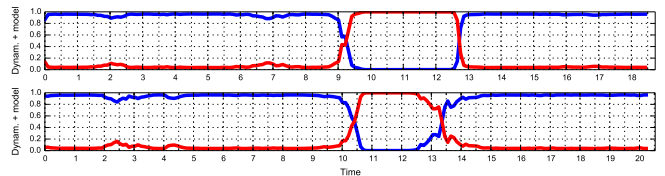


Fig. 5: Filtering results obtained with an incorrect model prediction for the scenes shown in Fig. 4c (top), and Fig. 4f (bottom). The model prediction is manually set to $P(m_t = \text{LK}) = 0.8$ and $P(m_t = \text{LC}) = 0.2$ throughout the whole scene. Despite this, the lane change maneuvers are still detected.

the target stabilizes its lateral position. The time between the end of the LC maneuver and the beginning of the LK maneuver is not considered in our analysis. It is unimportant whether the algorithm predicts the maneuver switch at the moment in which the target begins to steer or at the moment in which it has finally aligned with the road.

With this setup, we can calculate the following metrics to evaluate the robustness of the estimation: 1) Accuracy: proportion of samples correctly classified; 2) Precision: number of correct positive predictions divided by the total number of positive class values predicted; 3) Recall: number of positive predictions divided by the number of positive class values in the test data; and 4) False Positive Rate: number of false positives divided by the total number of negative class values in the test data. The results are presented in table II. The overall performance of our approach is satisfactory, beating

TABLE II: Performance metrics from a classification perspective.

	Accuracy	Precision	Recall	FPR
Model	0.7463	0.4455	0.7305	0.2493
Dynamics	0.8945	0.7557	0.7532	0.0668
IMM	0.8029	0.5332	0.6786	0.1630
Dynamics+Model	0.9203	0.8277	0.7955	0.0454

the IMM filter in every metric. The effect of the model-based prediction is also clear in the results, improving the performance with respect to the dynamics-based prediction in all four metrics. A result of particular relevance for the potential application of this system as a collision detection warning system is the extremely low false-positive rate.

Another relevant performance measure that needs to be quantified is the lane change detection delay, that is, the delay between the beginning of the lane change maneuver and the actual detection. The individual detection delays for each of the test scenes and for each approach are presented in table III. Again, the proposed framework beats the IMM filter by detecting the beginning of a lane change 0.39s earlier on average. Integrating the model-based prediction into the filtering framework leads to detecting a lane change 0.13s earlier on average in comparison with the dynamics-only estimate.

TABLE III: Lane change detection delay. We measure the time gap between the beginning of the lane change maneuver and its detection on each of the test scenes.

Detection delay: t_D [s]											
Scene	1	2	3	4	5	6	7	8	9	10	Mean
IMM	1.40	0.78	0.71	0.84	0.46	1.27	2.01	1.25	1.05	0.78	1.05
Dyn	1.09	0.48	0.51	0.74	0.15	1.07	1.71	0.95	0.75	0.48	0.79
Dyn+Mod	0.99	0.28	0.31	0.64	0.15	1.07	1.51	0.74	0.65	0.28	0.66

V. CONCLUSIONS

In this paper we have presented a probabilistic filtering framework to infer the maneuvers of drivers in highway scenarios. The proposed approach integrates the anticipatory, risk-averse maneuver prediction capabilities of a driver model learned with IRL into a filtering framework. Experimental results on real highway data gathered with an instrumented vehicle show that our approach can detect a lane change maneuver 0.39s earlier on average than a state-of-the-art IMM filter. Including the model-based prediction into the filtering framework has been shown to produce faster detections and increased robustness in the maneuver estimations in comparison with a filter relying only on dynamics. The overall high accuracy and low false positive rate in the detection of lane changes obtained in our experiments open the door to the application of this work as part of a collision warning system.

Future work will explore enhancing the maneuver forecasting step by including destination forecasting (is the target going to take the next exit?), increasing the complexity of the model (e.g., with time-to-collision as a feature), and using different driver models (aggressive, relaxed, heavy vehicle...).

ACKNOWLEDGMENT

We would like to thank Nicolas Vignard and Gabriel Othmezzouri for the helpful discussions and the feedback provided during the preparation of this work. We acknowledge as well Jean-Alix David, Jérôme Lussereau and Lukas Rummelhard for their assistance during the data collection.

REFERENCES

- [1] L. Rummelhard, A. Nègre, M. Perrollaz, and C. Laugier, "Probabilistic Grid-based Collision Risk Prediction for Driving Application," in *ISER*, Marrakech/Essaouira, Morocco, June 2014.
- [2] D. Barber, "Expectation correction for smoothed inference in switching linear dynamical systems," *Journal of Machine Learning Research*, vol. 7, pp. 2515–2540, 2006.
- [3] G. Agamennoni, J. I. Nieto, and E. M. Nebot, "Estimation of multivehicle dynamics by considering contextual information," *IEEE Transactions on Robotics*, vol. 28, no. 4, pp. 855–870, Aug 2012.
- [4] D. Sierra González, J. S. Dibangoye, and C. Laugier, "High-Speed Highway Scene Prediction Based on Driver Models Learned From Demonstrations," in *Proceedings of the 2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC 2016)*, Rio de Janeiro, Brazil, Nov. 2016.
- [5] A. Lawitzky, D. Althoff, C. F. Passenberg, G. Tanzmeister, D. Wollherr, and M. Buss, "Interactive scene prediction for automotive applications," in *2013 IEEE Intelligent Vehicles Symposium (IV)*, Gold Coast City, Australia, June 23–26, 2013, 2013, pp. 1028–1033.
- [6] W. Schwarting and P. Pascheka, "Recursive conflict resolution for cooperative motion planning in dynamic highway traffic," in *Intelligent Transportation Systems (ITSC), 2014 IEEE 17th International Conference on*, Oct 2014, pp. 1039–1044.
- [7] K. Weiss, N. Kaempchen, and A. Kirchner, "Multiple-model tracking for the detection of lane change maneuvers," in *IEEE Intelligent Vehicles Symposium, 2004*, June 2004, pp. 937–942.
- [8] N. Kaempchen, K. Weiss, M. Schaefer, and K. C. J. Dietmayer, "Imm object tracking for high dynamic driving maneuvers," in *IEEE Intelligent Vehicles Symposium, 2004*, June 2004, pp. 825–830.
- [9] R. Toledo-Moreo and M. A. Zamora-Izquierdo, "Imm-based lane-change prediction in highways with low-cost gps/ins," *IEEE Transactions on Intelligent Transportation Systems*, vol. 10, no. 1, pp. 180–185, March 2009.
- [10] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *ROBOMECH Journal*, vol. 1, no. 1, pp. 1–14, 2014.
- [11] C. Tay, "Analysis of Dynamic Scenes: Application to Driving Assistance," Theses, Institut National Polytechnique de Grenoble - INPG, Sept. 2009.
- [12] G. Aoude, B. Luders, J. Joseph, N. Roy, and J. How, "Probabilistically safe motion planning to avoid dynamic obstacles with uncertain motion patterns," *Autonomous Robots*, vol. 35, no. 1, pp. 51–76, 2013.
- [13] T. Gindele, S. Brechtel, and R. Dillmann, "A probabilistic model for estimating driver behaviors and vehicle trajectories in traffic environments," in *13th International IEEE Conference on Intelligent Transportation Systems*, Sept 2010, pp. 1625–1631.
- [14] M. Bahram, C. Hubmann, A. Lawitzky, M. Aeberhard, and D. Wollherr, "A combined model- and learning-based framework for interaction-aware maneuver prediction," *IEEE Transactions on Intelligent Transportation Systems*, vol. 17, no. 6, pp. 1538–1550, June 2016.
- [15] M. T. Wolf and J. W. Burdick, "Artificial potential functions for highway driving with collision avoidance," in *2008 IEEE International Conference on Robotics and Automation*, May 2008, pp. 3731–3736.
- [16] B. D. Ziebart, A. L. Maas, J. A. Bagnell, and A. K. Dey, "Maximum entropy inverse reinforcement learning," in *Proceedings of the Twenty-Third AAAI Conference on Artificial Intelligence, AAAI 2008, Chicago, Illinois, USA, July 13-17, 2008*, 2008, pp. 1433–1438.
- [17] M. Treiber, A. Hennecke, and D. Helbing, "Congested traffic states in empirical observations and microscopic simulations," *Phys. Rev. E*, vol. 62, pp. 1805–1824, Aug 2000.
- [18] L. Rummelhard, A. Nègre, and C. Laugier, "Conditional monte carlo dense occupancy tracker," in *2015 IEEE 18th International Conference on Intelligent Transportation Systems*, Sept 2015, pp. 2485–2490.
- [19] Y. Bar-Shalom and X.-R. Li, *Estimation with Applications to Tracking and Navigation*. New York, NY, USA: John Wiley & Sons, Inc., 2001.