

Interaction-Aware Tracking and Lane Change Detection in Highway Scenarios Using Realistic Driver Models

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

► **To cite this version:**

David Sierra González, Víctor Romero-Cano, Jilles Dibangoye, Christian Laugier. Interaction-Aware Tracking and Lane Change Detection in Highway Scenarios Using Realistic Driver Models. ICRA 2017 Workshop on Robotics and Vehicular Technologies for Self-driving cars, Jun 2017, Singapore, Singapore. 2017, <http://cim.mcgill.ca/malika/ICRA17_WS/>. <hal-01534094>

HAL Id: hal-01534094

<https://hal.inria.fr/hal-01534094>

Submitted on 7 Jun 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.

Interaction-Aware Tracking and Lane Change Detection in Highway Scenarios Using Realistic Driver Models (Extended Abstract)

David Sierra González, Víctor Romero-Cano, Jilles S. Dibangoye, and Christian Laugier

Abstract—We address the problem of multi-vehicle tracking and motion prediction in highway scenarios using information from sensors and perception systems widely used in automated driving. In particular, we focus on the detection of lane change maneuvers. Dangerous lane changing constitutes the main cause of highway accidents and a reliable detection system is still lacking on modern cars. 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 behaviors that deviate from those learned from demonstrations. In this paper, we show how to combine model-based behavior prediction and filtering-based state and maneuver tracking in order to detect lane changes in highway scenarios, and present the results obtained on real data gathered with an instrumented vehicle.

I. INTRODUCTION

Interacting with human agents is one of the major challenges faced by Advanced Driver Assistance Systems (ADAS) and autonomous driving vehicles. In spite of the complexity of the problem, human drivers are extremely good at predicting the intentions of surrounding vehicles. This is due to their innate ability to interpret the motion cues of other drivers and to reason over their most likely risk-averse behavior.

In previous work, our efforts were focused on modeling the risk-averse behavior of drivers from driving demonstrations using Inverse Reinforcement Learning (IRL), and on using the resulting models to predict the development of highway traffic scenes [1]. The main drawback of this approach is that it fails to consider all the dynamic evidence that might give away potentially dangerous maneuvers that deviate from the expected risk-averse behavior of highway drivers.

In order to exploit the predictive potential of IRL driver models without disregarding any physical evidence, we introduce a probabilistic framework that predicts the lane change intentions of highway drivers by merging model and tracking-based maneuver inference.

II. RELATED WORK

Research in the fields of state estimation and scene prediction for Intelligent Vehicles has increased significantly in the past decade. In the first place, we can identify tracking-based

maneuver inference approaches, which are based on identifying the motion model that fits best the current dynamics of a maneuvering target from among a discrete set of models [2], [3]. These approaches are typically based on an Interacting Multiple-Model (IMM) filter, in which 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.

Dynamic Bayesian Networks (DBN) have been used to explicitly consider the interactions between traffic participants in the long-term prediction of traffic situations [4]. The behavior for each vehicle is inferred 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 approach makes conceptually a lot of sense but relies strongly on particle filtering, which may limit its applicability in real-world complex scenarios.

Another approach along the same lines is presented in [5]. 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, depending also on the continuous state at the preceding time step through a linear feature-based function. The major difference with [4] 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 Dynamical Systems (aSLDS) presented in [6]. The approach is evaluated on its ability to track interacting trucks in an opencast mine. The model describing the interactions is relatively simple, considering only as factors the time-to-collision and the right-of-way.

The proposed approach builds upon [1] and [5]. The model-based prediction from [1] is integrated into a filtering framework in order to: 1) reason about the most likely high-level, interaction-aware behavior of drivers by considering its near-future consequences; and 2) use the expected model-based behavior to reinforce and accelerate the detection of changes in the dynamics of the targets.

III. FRAMEWORK

A. Driver modeling

In order to learn the feature-based cost function describing the preferences of a driver from driving demonstrations, we first model the dynamics of the driver as a Markov Decision Process (MDP), then select a number of relevant features using background knowledge, and finally use an IRL algorithm to learn the balance between the different terms of the

This work has been supported by Toyota Motor Europe. The authors are with INRIA Grenoble Rhone-Alpes, 38330 Montbonnot-Saint-Martin, France - {david.sierra-gonzalez, victor.romero-cano, jilles.dibangoye, christian.laugier}@inria.fr

cost function according to the behavior demonstrated. More details can be found in [1]. The resulting cost function is linear on the features $\mathcal{C}(s) = \boldsymbol{\theta} \cdot \mathbf{f}(s)$, 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 .

B. Framework overview

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. Fig. 1 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. In this work we consider two possible maneuvers: lane keeping (LK) and lane change (LC).
- The second level describes the state of the vehicle in a curved road frame through the continuous 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.
- Finally, the shaded nodes in the graphical model are the observations.

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

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

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.

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. The simplified kinematic bicycle model is used, with constant acceleration (CA) dynamics for the LK maneuver, and constant turn-rate and acceleration dynamics (CTRA) for the lane change maneuver.

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 subsection III-C.

Finally, 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.




Fig. 1: Graphical model representation of the proposed switching state-space model.

C. 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. 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 other surrounding drivers, this leads to a sequence of F future states that agree with his/her preferences (encoded in the driver model):

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

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.

D. Approximate Inference

Exact inference of the posterior $P(\mathbf{x}_t^i, m_t^i | \mathbf{z}_{1:t}^i)$ is intractable in the aSLDS, scaling exponentially with time [5], [6]. The proposed approximate inference engine is similar to the filtering approach presented in [6], with an extension to account for non-linear dynamics. Inference is performed individually per agent. 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:

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

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 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:

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

The term $P(\mathbf{x}_{t+1}^i | c_t, m_t^i, m_{t+1}^i, \mathbf{z}_{1:t+1}^i)$ is obtained by propagating forward with all the available dynamics m_{t+1}^i each component of the Gaussian mixture and by conditioning on the new observation \mathbf{z}_{t+1}^i . This leads to an exponential increase in the number of Gaussian components that is collapsed back to C components per maneuver at the end of each inference step. The prediction and update steps are performed using an Extended Kalman Filter (EKF). To obtain the weights of the new mixture components we consider:

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

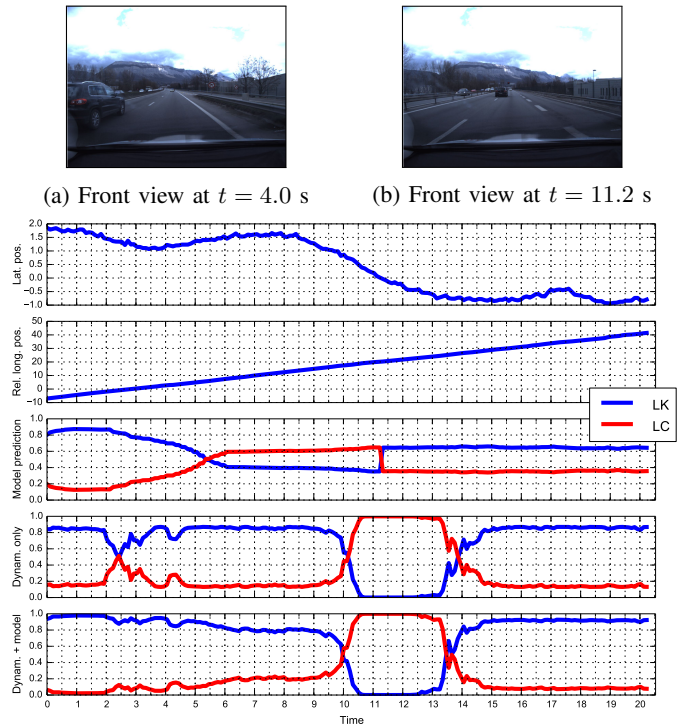
The terms $P(c_t | m_t^i, \mathbf{z}_{1:t}^i)$ and $P(m_t^i | \mathbf{z}_{1:t}^i)$ are available from the previous step in the recursion; the term $P(\mathbf{z}_{t+1}^i | c_t, m_t^i, m_{t+1}^i, \mathbf{z}_{1:t}^i)$ is the likelihood of the observation \mathbf{z}_{t+1}^i under the respective Kalman prediction, 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-C. This is where the fusion between model and dynamics takes place. For further details, including the second recursion not detailed here, see [6]. This inference engine leads to a computational complexity that grows linearly in the number of vehicles.

IV. EXPERIMENTAL EVALUATION

In this section, we focus on the validation of our framework’s ability to discriminate between lane change and lane keeping maneuvers in highway scenarios. Furthermore, we study how the model and dynamics-based predictions interact to determine the maneuver being executed by the target.

To validate our approach, we have gathered data on a French highway using an instrumented vehicle. By using a grid-based target tracker and a lane tracker, we are able to localize the targets with respect to the ego-vehicle (EV) and the road network. Fig. 2 shows one of the situations encountered in which a fast-driving vehicle overtakes the EV and merges in front of it. The lateral position of the target in road coordinates (the separating lane marking corresponds to $y = 0$) is shown in the first row of Fig. 2c. The longitudinal position of the target with respect to the EV is shown in the second row. The third row shows the model-based prediction considering a prediction horizon of 3s and sampling interval of 0.1s. It can be seen that the predicted probability of a LC maneuver is low while the target is in parallel with the EV, and begins to increase as the target gets farther away in front since drivers typically drive on the right-most lane.

The fourth row shows the dynamics-based estimate, obtained using an uninformative model (a model that assigns the same probability to all maneuvers). The last row shows the estimate obtained by fusing dynamic evidence with the model-based prediction, as explained in the previous section. The lane change maneuver is detected at around $t = 10.0s$, approximately 1.5s before the target crosses the lane marking and 0.2s earlier than with the dynamics-only estimate. The dynamics-only estimate is prone to false-positive detections (see $t = 2.5s$) caused by the oscillations of drivers around the lane center. The scene understanding brought into the estimate by the model-based prediction leads to a slower



(c) Data from the target and maneuver probability estimates.

Fig. 2: Typical highway scenario: a driver overtakes the ego-vehicle and merges in front of it.

response to unlikely maneuvers, potentially reducing the number of false positives. We can thus conclude that including the model-based prediction into the filtering framework results in faster detections and an increased robustness in the maneuver estimations. Finally, in cases where the dynamics-only estimate is confident enough and the model disagrees (see $t = 6.0s$ to $t = 9.5s$), the final estimate is dominated by the dynamics. This makes conceptually a lot of sense and constitutes a significant improvement over the behavior of pure model-based prediction approaches [1].

This example highlights the potential of the presented framework to identify lane change maneuvers in highways by combining dynamics and model-based inference.

REFERENCES

- [1] D. Sierra González, J. S. Dibangoye, and C. Laugier, “High-Speed Highway Scene Prediction Based on Driver Models Learned From Demonstrations,” in *IEEE 19th International Conference on Intelligent Transportation Systems*, Rio de Janeiro, Brazil, Nov. 2016.
- [2] 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.
- [3] R. Toledo-Moreo and M. 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, 2009.
- [4] 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.
- [5] 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.
- [6] D. Barber, “Expectation correction for smoothed inference in switching linear dynamical systems,” *Journal of Machine Learning Research*, vol. 7, pp. 2515–2540, 2006.