

Interaction-Aware Probabilistic Behavior Prediction in Urban Environments

Jens Schulz¹, Constantin Hubmann¹, Julian Löchner¹, and Darius Burschka²

Abstract—Planning for autonomous driving in complex, urban scenarios requires accurate prediction of the trajectories of surrounding traffic participants. Their future behavior depends on their route intentions, the road-geometry, traffic rules and mutual interaction, resulting in interdependencies between their trajectories. We present a probabilistic prediction framework based on a dynamic Bayesian network, which represents the state of the complete scene including all agents and respects the aforementioned dependencies. We propose Markovian, context-dependent motion models to define the interaction-aware behavior of drivers. At first, the state of the dynamic Bayesian network is estimated over time by tracking the single agents via sequential Monte Carlo inference. Secondly, we perform a probabilistic forward simulation of the network’s estimated belief state to generate the different combinatorial scene developments. This provides the corresponding trajectories for the set of possible, future scenes. Our framework can handle various road layouts and number of traffic participants. We evaluate the approach in online simulations and real-world scenarios. It is shown that our interaction-aware prediction outperforms interaction-unaware physics- and map-based approaches.

I. INTRODUCTION

While autonomous driving has already been pioneered in the 1980s by universities such as Carnegie Mellon and the Bundeswehr University Munich, it is still considered a challenge to integrate autonomous vehicles into real traffic. A major difficulty is the interaction with human drivers. Autonomous vehicles need to estimate the intentions and anticipate the future behavior of humans in order to plan collision-free trajectories and drive in a foresighted, efficient and cooperative manner. As intentions cannot be measured directly and humans exhibit individual and complex behavior, predictions will always be afflicted with uncertainty.

Simple prediction approaches such as constant turn rate and velocity may be sufficient for short term predictions and non-interactive situations. However, they quickly come to a limit in complex urban scenarios. The mixture of crossing, merging and diverging lanes and corresponding traffic rules create a complex structure and a stronger need for interaction between traffic participants, as can be seen in Fig. 1: the behavior of a driver depends on his intentions, the interactions with surrounding traffic and the static context, such as the road geometry. Furthermore, the future trajectory of a vehicle does also depend on how the complete situation evolves over time and, therefore, on how other agents are going to act. This introduces the need for combinatorial and




Fig. 1. Interaction-aware probabilistic trajectory prediction in an urban intersection scenario: the three vehicles have multiple possible routes, overlapping lanes and have to interact with each other.

interaction-aware motion prediction, which still represents a great challenge today [1].

In this paper, a behavior prediction framework is presented, which explicitly considers the intentions of drivers and the interdependencies between their future trajectories. We model the development of a traffic situation as a stochastic process consisting of multiple interacting agents. The decision making process of an agent is divided into three hierarchical layers: which route it is going to follow (route intention), whether it is going to pass a conflict area at an intersection before or after another agent (maneuver intention), and what continuous action it is going to execute. First, the set of possible routes and maneuvers is queried online given a digital map and the agents’ poses. Each agent then acts according to context-dependent behavior models given their route and maneuver intentions and the current environment. Describing this process as a dynamic Bayesian network (DBN) allows to specify causal as well as temporal dependencies and consider uncertainty in measurements and human behavior. Sequential Monte Carlo inference, also known as particle filtering, enables the use of hybrid, non-linear system models and the representation of arbitrary probability distributions. Using observations of the agents’ poses and velocities, Bayesian statistics allow for an estimation of the intentions and, therefore, for a more accurate probabilistic trajectory prediction by forward simulation of the DBN.

In this work, we focus on unsignalized intersections due to the prevalence of interdependencies between vehicles.

II. RELATED WORK

In the area of autonomous vehicles, intention estimation and motion prediction of traffic participants has been widely studied. Although these problems are highly coupled, in the existing literature, they are often tackled separately.

¹Jens Schulz, Constantin Hubmann, and Julian Löchner are with BMW Group, Munich, Germany {jens.schulz | constantin.hubmann | julian.loechner}@bmw.de

²Darius Burschka is with the Department of Computer Science, Technical University of Munich, Germany burschka@tum.de © 2018 IEEE

A. Intention Estimation

Popular methods for estimating route and maneuver intentions are discriminative classifiers (e.g., support vector machines (SVMs) [2], random forests (RFs) [3], artificial neural networks (ANNs) [4]) and probabilistic graphical models (e.g., hidden Markov models (HMMs) [5], Bayesian networks (BNs) [6]). For this purpose, the set of possible intentions is typically predefined offline and the models are learned for these fixed number of classes. For highways, this set usually consists of *lane change left*, *lane change right*, and *keep lane* (e.g., [7], [8]). For intersections, the desired route is mostly represented by the turning directions *left*, *right*, and *straight* (e.g., [5], [6]). Besides the intention of a lane change or the desired route, more detailed intentions can be distinguished. A longitudinal classification whether to yield or stop before an intersection has already been investigated (e.g., [2], [9], [3]). In [10], the set of possible intentions is generate online: the possible route alternatives and the corresponding yield positions are determined online using a map and the intentions are estimated.

Interactions between traffic participants are often not considered (e.g., [7], [5], [2], [3]). When the motion of multiple vehicles is interdependent, however, this may result in inaccurate predictions, especially for longer prediction horizons (e.g., if a vehicle approaching an intersection has to decelerate because of a slow vehicle in front, without considering interactions, it might be misleadingly inferred that it intends to turn). Investigating the so-called *freezing robot problem*, [11] has shown that agents typically engage in *joint collision avoidance* and cooperatively make room to create feasible trajectories. Therefore, possible future interactions between agents should be taken into account.

Others works on intention estimation have already explicitly modeled interdependencies between vehicles: In [6] and [9], the dependency on the preceding vehicle is considered in order to improve the estimation at intersections. In [12], interdependencies between multiple vehicles are modeled using object oriented probabilistic relational models with learned probability tables. They automatically extract the possible routes from the map and distinguish different interaction types depending on the route relations (*merge*, *cross*, *diverge*, *follow*) of vehicles. Promising results are shown by [4] with a long short-term memory (LSTM) based route classification for intersections, considering the states of up to seven surrounding vehicles, therefore, respecting possible interactions implicitly.

All of these works focus on intention estimation with discrete classes, but do not predict continuous trajectories needed for many motion planning algorithms.

B. Trajectory Prediction

The most simple trajectory prediction methods are physics-based and assume models like constant velocity, not considering the situational context [1]. Especially at intersections and for long prediction horizons, these models tend to have low accuracy due to the high dependency of the drivers' actions on the road geometry, traffic rules and

interactions to surrounding vehicles. Trajectory prediction that incorporates contextual information is often based on regression methods (e.g., Gaussian processes (GPs) [13], [14], RFs [15], ANNs [16]) or planning-based methods (e.g., [17]). In [14], velocity profiles with heteroscedastic variance for stopping at an intersection are learned using GPs. They include knowledge about the upcoming intersection, but do not consider other vehicles. In [18], seven different regression methods for interaction-aware microscopic driver behavior are learned and compared to each other in highway scenarios. An ANN based mapping from Markovian scene state to a continuous action distribution of an agent is learned for highway scenarios by [16]. These models allow an interaction-aware forward simulation, but do not explicitly infer route or maneuver intentions.

C. Estimation and Prediction

Besides the work that is either concerned about intention estimation or trajectory prediction, there has been effort to solve these problems together: A two-staged approach is employed by [19], in which they first classify a traffic situation into one of multiple predefined driving situations and then predict the velocity profile using situation-specific models. As these profiles only depend on features of the current situation (e.g., states of preceding vehicles), but do not incorporate the prediction of the surrounding vehicles, future interdependencies are ignored. Another combined approach can be found in [8], where highway maneuvers are first estimated based on multi-agent simulations and then used as input for a continuous trajectory prediction. Thus, both works solve the two problems separately, but improve their trajectory prediction by their maneuver and route estimates.

In [15], learned context-dependent action models of traffic participants are embedded into a DBN in order to estimate the state of the current situation applying sequential Monte Carlo (SMC) inference and predict the future motion of drivers. They outperform a Bayesian filter with constant velocity and heading assumption in simulations in terms of position accuracy. Although the different route options are modeled within the DBN, driver intentions are not explicitly inferred and evaluated. GP regression is utilized by [13] to estimate the predefined route intention also using SMC. In our previous work [17], we address the interrelated problems of behavior generation of the ego vehicle and behavior prediction of the surrounding vehicles in a combined fashion. Multi-agent maneuvers based on the concept of homotopy and corresponding trajectories are planned and used for intention estimation and ego vehicle control.

In contrast to the work presented in this section, we aim to propose a model for combined intention estimation and state prediction that can handle

- automatic generation of route hypotheses and maneuver hypotheses given the map and agent poses
- a varying number of traffic participants and various intersection layouts
- uncertainty in both measurements and human behavior
- combinatoric interaction between traffic participants.

III. PROBLEM STATEMENT

A traffic scene S consists of a set of agents $\mathcal{V} = \{V^0, \dots, V^K\}$, with $K \in \mathbb{N}_0$, in a static environment (map) with discrete time, continuous state, and continuous action space. The map consists of a road network with topological, geometric and infrastructure (yield lines, traffic signs, etc.) information as well as the prevailing traffic rules. At time step t , the set of agents \mathcal{V} is represented by their kinematic states $X_t = [x_t^0, \dots, x_t^K]^\top$, route intentions $R_t = [r_t^0, \dots, r_t^K]^\top$, and maneuver intentions $M_t = [m_t^0, \dots, m_t^K]^\top$. The kinematic state $x_t^i = [x_{t,i}^i, y_{t,i}^i, \theta_{t,i}^i, v_{t,i}^i]^\top$ of agent V^i consists of the Cartesian position, heading, and absolute velocity. Its length and width are considered to be given deterministically by the most recent measurement and, for the sake of brevity, are not included within x^i . The route intention r_t^i defines a path through the road network the agent desires to follow, the maneuver intention m_t^i the desired order relative to other agents in cases of intersecting or merging routes (see Sec. IV-D and IV-E for detailed definitions). Other types of maneuvers such as lane changes or overtaking are not considered within this work. At each time step, each agent executes an action a_t^i that depends on its intentions, the map and the kinematic states of all agents, transforming the current kinematic state x_t^i to a new state x_{t+1}^i . The actions of all agents are denoted as $A = [a_t^0, \dots, a_t^K]^\top$. The complete dynamic part of a scene is thus described by $S_t = [X_t, R_t, M_t, A_t]^\top$. At each time step, a noisy measurement $Z_t = [z_t^0, \dots, z_t^K]^\top$ with $z_t^i = [z_{x,t}^i, z_{y,t}^i, z_{\theta,t}^i, z_{v,t}^i]^\top$ is observed according to the distribution $P(Z_t|X_t)$, that contains information about the kinematic states of all agents.

The objective of this work is twofold: one part is to estimate the route intentions R and maneuver intentions M of all agents at the current time. The other part is to predict the future kinematic states X up to a temporal horizon T .

IV. APPROACH

In this work, we model the development of a traffic scene as a Markov process in the form of a DBN, consisting of all agents in a scene. This allows to explicitly model relations between agents, include domain knowledge and handle the uncertainty of measurements and human behavior. Each agent follows its own decision making process, which is divided into three hierarchical layers: the route intention, the maneuver intention and the continuous action. The random variables of the presented DBN and their causal and temporal dependencies are depicted in Fig. 2 and are explained in detail later in this section. In order to account for changing situations, the network structure is adapted online (creating and deleting agents as well as route and maneuver hypotheses). Thus, it can be applied to varying situations with an arbitrary number of agents, intention hypotheses and different road layouts. As our DBN describes a hybrid, non-linear system with a multi-modal, non-Gaussian belief, sequential importance resampling is used for inference, allowing to represent arbitrary probability distributions.

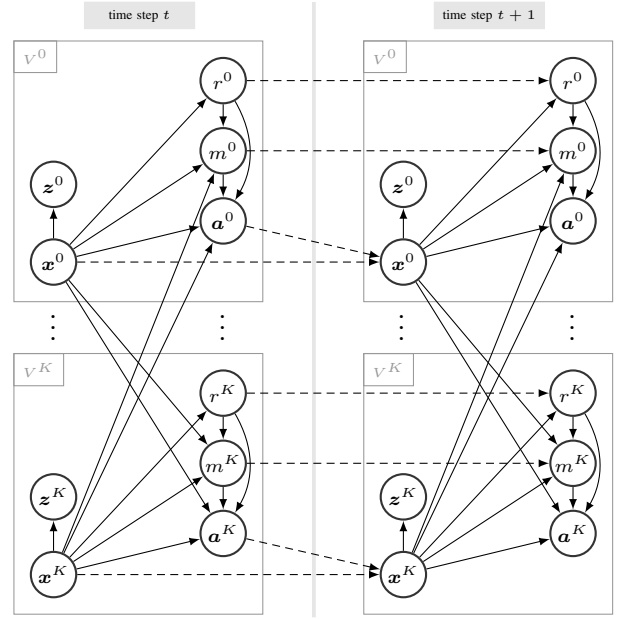


Fig. 2. DBN showing the interdependencies between agents. Random variables are drawn as circles, causal and temporal dependencies as solid and dashed arrows, respectively. r , m , and a also depend on the map.




Fig. 3. Exemplary initial sample generation (1-4), motion prediction (5), and particle weighting (6), shown for a single agent. Distributions are depicted simplified as being one dimensional. One particle represents the complete state space, i.e., kinematic state, route, maneuver, and action.

A. Estimation and Prediction

The goal of our framework is to estimate all drivers' intentions (route and maneuver) and to predict their future trajectories. The general procedure is exemplarily depicted in Fig. 3: Initially, a set of particles $S_0 = \{S_0^1, \dots, S_0^N\}$, with $S^i = [X^i, R^i, M^i, A^i]^\top$ representing the complete scene, is sampled (steps 1-4) according to the measurement and the map: $S_0^i \sim P(X_0, R_0, M_0, A_0 | Z_0, \text{map})$, with corresponding weights $\omega_0^i = 1/N$. Then, each particle is predicted to the subsequent time step (step 5) according to the transition probability: $S_{t+1}^i \sim P(S_{t+1}^i | S_t^i)$. As soon as a new measurement is available, the particle weights get updated according to the measurement likelihood (step 6): $\omega_t^i = P(Z_t | X_t^i) \omega_{t-1}^i$. The actual probability of a set of intentions (R_t, M_t) is given by

$$P(R_t, M_t) = \frac{\sum_{j \in \mathcal{J}} \omega_t^j}{\sum_{i=1}^N \omega_t^i}, \quad \mathcal{J} := \{j \mid R_t^j = R_t, M_t^j = M_t\}. \quad (1)$$

For the intention estimation process, the DBN is thus applied as a filter, comparing the different model hypotheses to the actual observations. The intention of a single agent can be derived through marginalization of the belief.




Fig. 4. Possible combinations of routes $[r_{\text{left}}^0, r_{\text{right}}^1]$ and $[r_{\text{right}}^0, r_{\text{right}}^1]$ at a roundabout: Although V^1 is currently not influenced by V^0 , it has to slow down in the future if V^0 stays inside the roundabout. This influence has to be taken into account for the trajectory prediction using forward simulation.

As DBNs are generative models, i.e., they can generate values of any of their random variables, it is possible to do a probabilistic forward simulation by iteratively predicting the current belief (including the estimated intentions) into subsequent time steps, applying the same models as for the filtering. As can be seen in Fig. 2, the action of an agent is modeled to be not *directly* dependent on the actions or intentions of others, but only on its own intentions and the current context (given by (map, X)). Thus, cyclic dependencies are avoided and one prediction step of an agent can be executed independently of the prediction steps of other agents. However, as the current context depends on other agents' past actions, an interdependency between their trajectories emerges *over time*, as shown in Fig. 4.

In order to reduce complexity and improve interpretability of the trajectory prediction, the forward simulation is not done for each particle, but for the mean kinematic state of all agents given their route and maneuver intentions. For each combination (R, M) within \mathcal{S} , one multi-agent trajectory is generated and weighted with the corresponding probability $P(R, M)$. Due to the interdependencies of multiple agents' future trajectories, this combinatorial aspect cannot be neglected within the prediction of the scene development.

The remainder of this section gives a detailed explanation of the single DBN nodes and their probability distributions.

B. Vehicle Kinematics

The action of each agent is defined as $\mathbf{a} = [a, \dot{\theta}]^\top$ with the longitudinal acceleration a and the yaw rate $\dot{\theta}$. It is the result of the decision making process, which is influenced by the current context and the agent's intentions, and is also estimated as a random variable of the DBN (see Sec. IV-F). The transition of the kinematic state is given by the probability distribution $P(\mathbf{x}'|\mathbf{x}, \mathbf{a}) = \mathcal{N}(\hat{\mathbf{x}}', \mathbf{Q})$, with

$$\hat{\mathbf{x}}' = \begin{pmatrix} \hat{x}' \\ \hat{y}' \\ \hat{\theta}' \\ \hat{v}' \end{pmatrix} = \begin{pmatrix} x + v\Delta T \cos(\theta') + \frac{1}{2}a\Delta T^2 \cos(\theta') \\ y + v\Delta T \sin(\theta') + \frac{1}{2}a\Delta T^2 \sin(\theta') \\ \theta + \dot{\theta}\Delta T \\ v + a\Delta T \end{pmatrix} \quad (2)$$

and $\mathbf{Q} = \text{diag}(\sigma_x^2, \sigma_y^2, \sigma_\theta^2, \sigma_v^2)$. Although this model is simplistic, we argue that it is sufficient for prediction purposes.




Fig. 5. Breadth-first search for possible routes of length l_H on lane graph.

C. Measurement

The proposed algorithm uses high-level cuboid objects as measurements, which can be derived by a magnitude of different sensors. Hence, low-level sensor specifics are abstracted. The data association, i.e., object detection and tracking, is handled by a separate algorithm and is considered to be given within this work. The kinematic state \mathbf{x} is measured with zero-mean Gaussian noise. The measurement $\mathbf{z} = [x_z, y_z, \theta_z, v_z]^\top$ is distributed according to $P(\mathbf{z}|\mathbf{x}) = \mathcal{N}(\hat{\mathbf{z}}, \mathbf{R})$, with $\hat{\mathbf{z}} = \mathbf{x}$ and $\mathbf{R} = \text{diag}(\sigma_{z_x}^2, \sigma_{z_y}^2, \sigma_{z_\theta}^2, \sigma_{z_v}^2)$.

D. Route Intention

The route $r \in \mathcal{R}$ forms the first layer of an agent's decision making process and serves as a path that guides its behavior. It is represented by a sequence of consecutive lanes. In every time step, the set of possible routes \mathcal{R} is determined given the agent's pose, the topological map, and a specified metric horizon l_H . We apply breadth-first search on the lane graph starting with the current lane matching (see Fig. 5).

The route of an agent mainly serves two purposes: Firstly, it allows to define relevant features along an agent's planned path such as the road curvature ahead or longitudinal distances to stop lines (see Sec. IV-F). Secondly, the routes of multiple agents allow to build relationships between agents on complex road layouts. Two routes are related by dividing them into parts that either *merge*, *diverge*, *cross*, are *identical*, or have no relevant relation at all. Different road junction types such as roundabouts, intersections or highway entrances can thus be broken down into these types of relations, allowing for a better generalization. Typical relations between agents consist of distances to merging or crossing areas of their routes and corresponding right of way rules (see Sec. IV-F). As each route has a different geometry and may imply different traffic rules and relations to other agents, the route directly influences a driver's actions.

Initially, the desired route r is sampled uniformly from the set of possible routes \mathcal{R} according to $P(r_i|\mathbf{x}, \text{map}) = |\mathcal{R}|^{-1}$, $\forall r_i \in \mathcal{R}$. Due to the fact that the route is only considered up to a specific horizon, a binary matching function $s_r(r', r) : \mathcal{R}' \times \mathcal{R} \rightarrow \{0, 1\}$ is used to determine which of the routes $r' \in \mathcal{R}'$ are possible successors of the current route r (i.e., imply the same decisions at each contained intersection) and which are not. If there are multiple candidates (in case of a route split), again, the route is sampled uniformly:

$$P(r'_j|r_i, \mathbf{x}, \text{map}) = \frac{s_r(r'_j, r_i)}{\sum_{r' \in \mathcal{R}'} s_r(r', r_i)}. \quad (3)$$




Fig. 6. (a): Possible conflict areas from V^0 's perspective for going straight, resulting from the three route hypotheses of V^1 . The actual route of V^1 is unknown to V^0 . (b): Four possible maneuvers for V^0 turning left, representing the sequence of agents passing the conflict areas.

E. Maneuver Intention

The maneuver $m \in \mathcal{M}$ forms the second layer of the decision making process and describes the desired sequence, in which agents are going to merge or cross at intersections. Therefore we introduce the notion of conflict areas: Given two agents on two routes, their conflict area is defined by the intersecting set of the areas of both routes, i.e., the area in which their lanes overlap. We assume an agent doesn't know which route other agents are going to follow, thus, all possible conflict areas are considered (see Fig. 6(a)).

In order to avoid collisions at conflict areas, agents have to schedule their passing sequence. A maneuver of agent V^i states for all pairs $\langle V^i, V^j \rangle$ that have a *potential* conflict (at least one route hypothesis of agent V^j has a conflict with V^i 's intended route), whether V^i will pass their conflict area first ($V^i < V^j$) or not ($V^i > V^j$). This definition follows our previous work [17], where maneuvers are based on the pseudo-homotopy of trajectories. Our data suggests, that vehicles that have right of way are typically not influenced by other vehicles approaching the intersection. Thus, different maneuvers are only considered for vehicles that do not have right of way. The set of possible maneuvers \mathcal{M} can be derived given the agent's route, the map, and the kinematic states of all agents. An example can be seen in Fig. 6(b). A more detailed description of this concept of maneuvers, also including lane changes, can be found in [17].

The desired maneuver m is initially sampled uniformly from the set of possible maneuvers \mathcal{M} according to $P(m_i|X, r, \text{map}) = |\mathcal{M}|^{-1}$, $\forall m_i \in \mathcal{M}$. As situations change over time, the set of possible maneuvers may change as well (e.g., a new agent arrives or an existing agent traverses a conflict area). Hence, for further time steps, a matching function $s_m(m', m) : \mathcal{M}' \times \mathcal{M} \rightarrow \{0, 1\}$ determines which of the new maneuvers $m' \in \mathcal{M}'$ are possible successors of the current maneuver m (i.e., there are no contradictory passing sequences). If there are multiple matching candidates, again, the maneuver is sampled uniformly.

F. Action Model

The action $\mathbf{a} = [a, \theta]$ of an agent depends on his route and maneuver intentions, the kinematic states of all agents, and the map. It forms the third layer of the decision making process. Within this section, a heuristics-based probabilistic

TABLE I
INFLUENCES, FEATURES, AND ACTION RANGES FOR AGENT V^i

Influence	Features	Action Range
vehicle dynamics	-	$[a_{vd}^{\min}, a_{vd}^{\max}]$
speed limit	$d_{v_{lim}}, v_{lim}, v^i$	$[-\infty, a_{IDM}^{\max}]$
preceding agent V^p	d^p, v^p, v^i	$[-\infty, a_{IDM}^{\max}]$
road curvature	d_ρ, ρ, v^i	$[-\infty, a_{curve}^{\max}]$
conflicting agent V^c	$\chi^{i,c}, d_{entry}^c, d_{exit}^c, v^c,$ $d_{yield}^i, d_{entry}^i, d_{exit}^i, v^i$	$[a_{conf}^{\min}, a_{conf}^{\max}]$

action model $P(\mathbf{a}|r, m, X, \text{map})$ is defined to show the potential of the general framework. In order to narrow down the large number of dependencies, we define a set of submodels, each handling one so-called *influence*. Each influence consists of a subset of the available features and constrains the acceleration to a range $[a^{\min}, a^{\max}]$ that is plausible (e.g., not leading to collisions or violations of traffic rules) given that specific influence. Tab. I shows the influences considered within this work, their corresponding features and action ranges. These influences represent the context on which an agent's actions are based on and are derived deterministically given the variables of the DBN.

The influence *vehicle dynamics* restricts the range of possible accelerations to the constant range $[a_{vd}^{\min}, a_{vd}^{\max}]$. *Speed limits* are defined by a set of pairs of speed limit v_{lim} and distance along the route $d_{v_{lim}}$ where it becomes effective. A *preceding agent* V^p is described by its relative distance d^p and its velocity v^p . For both of these influences, the so-called Intelligent Driver Model (IDM) presented by [20] is employed, dictating a maximum reasonable acceleration

$$a_{IDM}^{\max} = a_d \left(1 - \left(\frac{v^i}{v_{lim}} \right)^\delta - \left(\frac{d_d + v^i T_d + \frac{v^i(v^i - v^p)}{2\sqrt{|a_d b_d|}}}{d^p} \right)^2 \right). \quad (4)$$

The parameters minimum spacing d_d , desired time headway T_d , comfortable acceleration a_d , braking deceleration b_d , and acceleration exponent δ have to be specified. Although not part of the evaluation, the influences *red traffic light* and *stop sign* are also handled using (4) by setting $v^p = 0$ and d^p to the corresponding distance.

As the IDM was primarily designed for highway scenarios, the curvature of the road as well as merging or intersecting lanes are not considered. Thus, we define the following models allowing the prediction in urban scenarios: The model for the influence *curvature* is based on a desired maximum lateral acceleration a_{lat}^{\max} that implies a maximum velocity $v_\rho = \sqrt{\rho a_{lat}^{\max}}$ at a given curve radius ρ . The maximum acceleration of V^i for one time step ΔT to still be able to reach the velocity v_ρ at the corresponding distance d_ρ with the comfortable braking deceleration b_d is

$$a_{v_\rho, d_\rho}^{\max} = \tilde{a} = \frac{-2v + \Delta T b_d + \sqrt{4v \Delta T b_d + \Delta T^2 b_d^2 - 8b_d d_\rho + 4v_\rho^2}}{2\Delta T}, \quad (5)$$




Fig. 7. Example of possible upper and lower bounds of the action models of the single influences, used to define the action probability distribution.

which can be determined by the following equations:

$$\begin{aligned}
 v_1 &= v + \tilde{a}\Delta T \\
 v_\rho &= v_1 + b_d\Delta T_2 \\
 d_1 &= v\Delta T + \frac{1}{2}\tilde{a}\Delta T^2 \\
 d_\rho &= d_1 + v_1\Delta T_2 + \frac{1}{2}b_d\Delta T_2^2
 \end{aligned}$$

The smallest allowed acceleration of all curvature distance pairs along the route is used. This results in a foresighted curvature approach.

The *conflict* model is based on conflict areas at overlapping lanes where vehicles have to coordinate a specific sequence of passing. A conflict of agent V^i with another agent V^c is described by the right of way $\chi^{i,c}$, their velocities and distances to entering and exiting the conflict area and their distances to potential yield lines d_{entry} , d_{exit} , and d_{yield} , respectively. If agent V^i has right of way, results indicate that it is sufficient to assume that it is not influenced by the other agent ($[a_{\text{conf}}^{\min}, a_{\text{conf}}^{\max}] = [-\infty, \infty]$). If agent V^i has to yield, it acts according to its desired maneuver m . Each agent that is going to pass before V^i introduces an upper bound of acceleration (a_{conf}^{\max}), each agent that is going to pass after V^i introduces a lower bound (a_{conf}^{\min}). These accelerations are determined such that a minimum time gap between the two passing vehicles at the overlapping areas is ensured, assuming others drive with constant velocity.

The ranges of feasible accelerations of the single influences are combined as shown in Fig. 7 to the overall range

$$a_{\max} = \min\{a_{\text{vd}}^{\max}, a_{\text{curv}}^{\max}, a_{\text{IDM}}^{\max}, a_{\text{int}}^{\max}, a_{\text{conf}}^{\max}\}, \quad (6)$$

$$a_{\min} = \max\{a_{\text{vd}}^{\min}, a_{\text{conf}}^{\min}\}. \quad (7)$$

Our measurement data suggests that drivers tend to minimize driving time while not exceeding the plausible acceleration range. Thus, accelerations are sampled from the distribution $P(a|r, m, X, \text{map}) = \mathcal{N}(\mu_a, \sigma_a^2)$, with a mean close to the lowest maximum bound: $\mu_a = a_{\max} - \sigma_a$. The yaw rate is sampled from $P(\dot{\theta}|r, \mathbf{x}, a, \text{map}) = \mathcal{N}(\mu_{\dot{\theta}}, \sigma_{\dot{\theta}}^2)$, given a mean yaw rate $\mu_{\dot{\theta}}$ that keeps the agent close to the center of its lane, which is calculated based on simple heuristics.

V. EVALUATION

In order to assess the necessity of interaction-aware prediction, we compare our model to simpler models in simulated and real driving scenarios. Scenes with interactive behavior are recorded with a measuring vehicle on real roads and on a test-track and are generated with a proprietary traffic simulator. The measuring vehicle's pose and velocity is estimated using GPS/INS. Both lidar and radar sensors are used to detect and track objects nearby. The evaluation parameters can be seen in Tab. II. To avoid particle deprivation due to resampling, new particles are sampled from

TABLE II
EVALUATION PARAMETERS

ΔT	0.2 s	δ	4	$\sigma_{z_{x/y}}$	15 m
N	1000	a_{lat}^{\max}	2 m s^{-2}	σ_{z_θ}	3.14
l_H	30 m	$\sigma_{x/y}$	0.5 m	σ_{z_v}	15 m s^{-1}
d_d	2 m	σ_θ	0.05	$\sigma_{s_{x/y}}$	1 m
T_d	0.1 s	σ_v	1.5 m s^{-1}	σ_{s_θ}	0.03
a_d	0.7 m s^{-2}	σ_a	1.5 m s^{-2}	σ_{s_v}	1 m s^{-1}
b_d	-0.5 m s^{-2}	$\sigma_{\dot{\theta}}$	0.05 s^{-1}		

the current measurement distribution with probability 0.001. The computing time of one time step of a scene with three vehicles, each having three route options, is approximately 0.3 s on an Intel Core i7-5820K CPU @ 3.30GHz with non-optimized C++ code.

A. Intention Estimation

The imprecision of the route (and analogously maneuver) estimate is measured using the Kullback-Leibler divergence

$$D_{\text{KL}}(r_{\text{GT}}^i || r^i) = \sum_{j=1}^{|R|} r_{\text{GT},j}^i \log \frac{r_{\text{GT},j}^i}{P(r_j^i)} \quad (8)$$

from estimate $r^i = [P(r_1^i), \dots, P(r_{|R|}^i)]$ to ground truth $r_{\text{GT}}^i = [r_{\text{GT},1}^i, \dots, r_{\text{GT},|R|}^i]$, with

$$r_{\text{GT},j}^i = \begin{cases} 1 & \text{if } V^i \text{ follows } r_j^i \\ 0 & \text{else} \end{cases}. \quad (9)$$

We evaluate the intention estimation of the presented model, which we call *interactive* model, and a solely *map-based* model. The *map-based* model uses all of the features given by the map but ignores surrounding vehicles and, therefore, is interaction-unaware. Thus, agents are predicted as if there were no other vehicles around.

As our dataset mostly consists of scenes with little interaction and both models are identical for scenes without interaction, a statistical evaluation of the complete dataset produces similar results. In order to highlight their differences, we specifically determined situations in which multiple vehicles cross an intersection, and hence, containing interdependencies between vehicles. Though intersection crossings are statistically rare in our dataset, these situations tend to be most critical and therefore require explicit evaluation. The intention estimation is evaluated in detail for three scenes:

1) *Yielding vehicle*: In the simulated *scene 1* (first row of Fig. 8), V^1 has right of way and goes straight, V^0 has to yield and wants to turn left. To improve readability, at first it is assumed that V^0 is actually yielding and therefore only has one possible maneuver ($V^1 \prec V^0$), but multiple possible routes. While V^0 's routes for going straight and turning left demand yielding, the route for turning right is free. As V^0 waits for V^1 ($t=10-18$ s), it is inferred by the interactive model that turning right is unlikely (as waiting would not be necessary) and turning left and going straight is equally likely (as both routes are blocked). As soon as V^1 has left the conflict area, V^0 accelerates again and turns,




Fig. 8. Detailed evaluation of agent V^0 in scene 1 (first row, simulated data) and scene 2 (second row, real data): Comparison of route estimation and trajectory prediction for the different tracking methods (CTRV, map-based, interactive), also showing the scaled velocity profile $\tilde{v} = v/(50\text{km/h})$.

whereby the left route is inferred correctly. The map-based model, however, infers incorrectly that V^0 wants to turn right ($t=13\text{ s}$), as this route has the highest curvature, implying the lowest velocity. For $t>13\text{ s}$, as V^0 even becomes too slow for turning right, none of the map-based models can explain the actual behavior anymore. Thus, only the particles sampled newly from the measurement survive, resulting in a random oscillation and a momentary improvement of the D_{KL} .

1b) Maneuver distinction: The combined maneuver and route estimation is analyzed in Fig. 9, where scene 1 is modified, such that V^0 crosses first (scene 1b). The interactive model with maneuver distinction is compared to the interactive model without maneuver distinction (assuming V^0 will yield): At $t=0\text{ s}$, all routes are equally likely, but as V^0 does not decelerate strongly ($t=2-10\text{ s}$), the probability to yield decreases, whereas the probabilities to either turn right (no conflict) or merge / cross before V^1 increase. As V^0 slows down in order to respect the upcoming curvature ($t=9-11\text{ s}$), the straight route becomes unlikely. Finally ($t=12-20\text{ s}$), as the velocity is still too high for turning right, it is correctly inferred that V^0 will turn left and merge before V^1 . Without the distinction of the two possible maneuvers, assuming V^0 is going to yield, it is incorrectly inferred that V^0 wants to turn right (as this lane has no conflict), resulting in a higher estimation and trajectory prediction error.

2) Preceding vehicle: In the real driving scene 2 (second row of Fig. 8), V^0 follows V^1 approaching an intersection. As V^1 has to yield and therefore decelerates, V^0 decelerates as well in order to keep the desired headway distance. All three possible routes of V^0 are blocked by the preceding vehicle, hence, it is not possible to infer the route until the preceding agent has passed the intersection ($t=17\text{ s}$). A uniform distribution is the desired result, which is generated by the interactive method. The map-based method incorrectly infers that V^0 wants to turn right ($t=10\text{ s}$), as it is slowing down (actually caused by the preceding vehicle). For $t>10\text{ s}$, none of the map-based models can explain the observations anymore, also resulting in a random oscillation.

B. Trajectory Prediction

The accuracy of the trajectory prediction of all agents at time t for the future time step τ is quantified using the

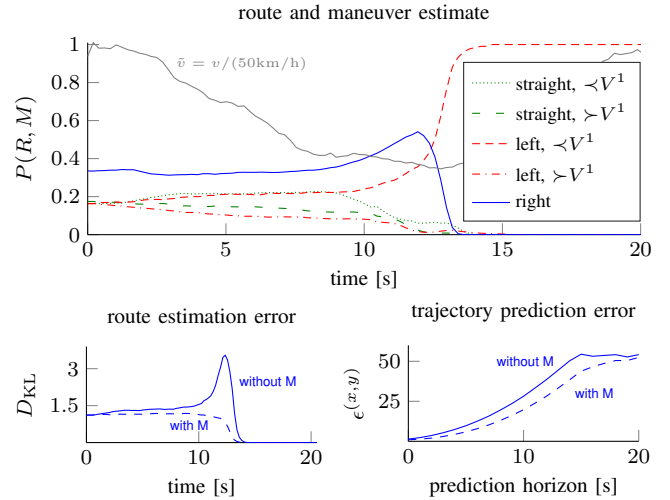


Fig. 9. Route and maneuver estimation of agent V^0 in scene 1b and comparison of interaction-aware model with and without maneuver distinction.

position components of the weighted root mean square error between prediction and measurement

$$\epsilon_{\tau|t}^{(x,y)} = \sqrt{\sum_{R_t, M_t} P(R_t, M_t) \left(\hat{X}_{\tau|t, R_t, M_t}^{(x,y)} - Z_{\tau}^{(x,y)} \right)^2}, \quad (10)$$

and the measurement likelihood

$$\mathcal{L}_{\tau|t}^{(x,y)} = \prod_{R_t, M_t} P(R_t, M_t) P(Z_{\tau}^{(x,y)} | \hat{X}_{\tau|t, R_t, M_t}^{(x,y)}). \quad (11)$$

The *interactive* model is compared to the *map-based* model and a *constant turn rate and velocity (CTRV)* model [21], which serves as a simple baseline algorithm. It is independent of both the map and surrounding vehicles. The error of the trajectory prediction of V^0 for scenes 1 and 2 are depicted in the most right column of Fig. 8. The CTRV model performs worse in scene 1, as V^0 changes its velocity and orientation more intensely. For the map-based model, the first scene is also more challenging, as V^0 stops for a long time, which cannot be explained by the model at all. Its high route estimation error negatively affects its prediction accuracy. The interactive model outperforms the other two approaches in both scenes.



Fig. 10. Camera view of measurement vehicle while yielding to oncoming traffic in order to turn left into a parking lot.




Fig. 11. Prediction error and likelihood in five different driving scenarios.

Furthermore, in order to compare the models in a quantitative manner, five different real driving scenes have been recorded on a test track and on real roads (Fig. 10). These scenes altogether consist of 15 vehicles, two four-way intersections, two T-junctions, and a roundabout. The statistical results showing the prediction error and measurement likelihood over all scenes and vehicles are depicted in Fig. 11. It can be seen that the interaction-aware model outperforms both CTRV and map-based models. Although the differences between the map-based and the interactive model might seem to be rather small, it has to be noted that the time steps in which traffic participants actually interact with each other do not predominate. As shown in Fig. 8, however, in scenes where the behaviors of drivers are highly interdependent, interaction-aware prediction becomes essential.

A video of the approach with exemplary scenes is included in the conference proceedings and can also be found at <https://mediatum.ub.tum.de/1449806>.

VI. CONCLUSIONS

In this work, we proposed an interaction-aware prediction framework that is able to estimate route and maneuver intentions of drivers and predict complete scene developments in a combined fashion. Possible routes and maneuvers are generated online given a map and the current belief state. The framework can handle a varying number of traffic participants and different road layouts without the need to predefine a discrete set of classes. It is capable of dealing with uncertainty in measurements and human behavior and interdependencies between drivers. Its particle filtering nature allows to represent the non-linear system dynamics and the multi-modal and hybrid belief state.

Due to the combinatorial aspect of long-term motion prediction, the complexity of inference grows exponentially with the number of considered agents and possible intentions. Nevertheless, we show that in cases with close interaction between traffic participants, their interdependencies cannot be neglected. Future work will focus on reducing this complexity and improving behavior model accuracy, e.g., by learning the action model from data, which in turn will

reduce the number of needed particles. Furthermore, different possible plans of the ego vehicle could be taken into account within the forward simulation, in order to evaluate them with respect to how surrounding drivers are likely going to react. Therefore, less conservative actions could be executed, respecting the influence of the ego vehicle on others.

REFERENCES

- [1] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *Robomech J.*, vol. 1, no. 1, p. 1, 2014.
- [2] G. S. Aoude, V. R. Desaraju, L. H. Stephens, and J. P. How, "Behavior classification algorithms at intersections and validation using naturalistic data," in *Intell. Veh. Symp. (IV)*, pp. 601–606, IEEE, 2011.
- [3] M. Barbier, C. Laugier, O. Simonin, and J. Ibañez-Guzmán, "Classification of Drivers Manoeuvre for Road Intersection Crossing with Synthetic and Real Data," in *Intell. Veh. Symp. (IV)*, p. 7, IEEE, 2017.
- [4] D. J. Phillips, T. A. Wheeler, and M. J. Kochenderfer, "Generalizable Intention Prediction of Human Drivers at Intersections," in *Intell. Veh. Symp. (IV)*, pp. 1665–1670, IEEE, 2017.
- [5] T. Streubel and K. H. Hoffmann, "Prediction of driver intended path at intersections," in *Intell. Veh. Symp. (IV)*, pp. 134–139, IEEE, 2014.
- [6] M. Liebner, M. Baumann, F. Klanner, and C. Stiller, "Driver intent inference at urban intersections using the intelligent driver model," in *Intell. Veh. Symp. (IV)*, pp. 1162–1167, IEEE, 2012.
- [7] P. Kumar, M. Perrollaz, S. Lefèvre, and C. Laugier, "Learning-based approach for online lane change intention prediction," in *Intell. Veh. Symp. (IV)*, pp. 797–802, IEEE, 2013.
- [8] M. Bahram, C. Hubmann, A. Lawitzky, M. Aeberhard, and D. Wollherr, "A Combined Model- and Learning-Based Framework for Interaction-Aware Maneuver Prediction," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, pp. 1538–1550, June 2016.
- [9] S. Klingelschmitt, M. Platho, H.-M. Groß, V. Willert, and J. Eggert, "Combining behavior and situation information for reliably estimating multiple intentions," in *Intell. Veh. Symp. (IV)*, pp. 388–393, IEEE, 2014.
- [10] S. Lefèvre, C. Laugier, and J. Ibañez-Guzmán, "Risk assessment at road intersections: Comparing intention and expectation," in *Intell. Veh. Symp. (IV)*, pp. 165–171, IEEE, 2012.
- [11] P. Trautman and A. Krause, "Unfreezing the robot: Navigation in dense, interacting crowds," in *Int. Conf. Intell. Robot. and Syst. (IROS)*, pp. 797–803, IEEE, 2010.
- [12] F. Kuhnt, J. Schulz, T. Schamm, and J. M. Zöllner, "Understanding interactions between traffic participants based on learned behaviors," in *Intell. Veh. Symp. (IV)*, pp. 1271–1278, IEEE, 2016.
- [13] Q. Tran and J. Firl, "Online maneuver recognition and multimodal trajectory prediction for intersection assistance using non-parametric regression," in *Intell. Veh. Symp. (IV)*, pp. 918–923, IEEE, 2014.
- [14] A. Armand, D. Filliat, and J. Ibañez-Guzmán, "Modelling stop intersection approaches using gaussian processes," in *Int. Conf. Intell. Transp. Syst. (ITSC)*, pp. 1650–1655, IEEE, 2013.
- [15] T. Gindele, S. Brechtel, and R. Dillmann, "Learning context sensitive behavior models from observations for predicting traffic situations," in *Int. Conf. Intell. Transp. Syst. (ITSC)*, pp. 1764–1771, IEEE, 2013.
- [16] D. Lenz, F. Diehl, M. T. Le, and A. Knoll, "Deep neural networks for Markovian interactive scene prediction in highway scenarios," in *Intell. Veh. Symp. (IV)*, 2017 IEEE, pp. 685–692, IEEE, 2017.
- [17] J. Schulz, K. Hirsenkorn, J. Löchner, M. Werling, and D. Burschka, "Estimation of collective maneuvers through cooperative multi-agent planning," in *Intell. Veh. Symp. (IV)*, pp. 624–631, IEEE, 2017.
- [18] T. A. Wheeler, P. Robbel, and M. J. Kochenderfer, "Analysis of microscopic behavior models for probabilistic modeling of driver behavior," in *Int. Conf. Intell. Transp. Syst. (ITSC)*, pp. 1604–1609, IEEE, 2016.
- [19] M. Platho, H.-M. Groß, and J. Eggert, "Predicting velocity profiles of road users at intersections using configurations," in *Intell. Veh. Symp. (IV)*, pp. 945–951, IEEE, 2013.
- [20] 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.
- [21] R. Schubert, E. Richter, and G. Wanielik, "Comparison and evaluation of advanced motion models for vehicle tracking," in *Int. Conf. Inform. Fusion*, pp. 1–6, IEEE, 2008.