Predicting Velocity Profiles of Road Users at Intersections Using Configurations

Matthias Platho, Horst-Michael Groß and Julian Eggert

Abstract— Intersections are among the most complex traffic situations that motorists encounter, which is reflected by the fact that in Europe more than 40 percent of accidents resulting in injury occur at intersections. In order to support the driver in crossing an intersection an advanced driver assistance system is required to predict the behavior of other drivers, like acceleration and braking maneuvers, as accurately as possible. Such a prediction is a challenging task when considering the complexity and variability of situations encountered at urban intersections. We propose to tackle this problem using a two-staged approach. In the first stage the situation is decomposed into small, more manageable sets of related road users to prevent a combinatorial explosion of possibilities. For each set the road user's driving situation is estimated. In the second stage the velocity profiles of all road users are predicted, taking advantage of the previously estimated driving situation by employing prediction models that are specific to the situation type. The proposed method is evaluated on a simulated intersection situation where the two-staged approach clearly outperforms prediction methods that work without assessing driving situations first. We also show qualitative results on real-world data that confirm the benefits of our approach.

I. INTRODUCTION

Crossing a busy intersection belongs to the most challenging tasks during everyday driving. In order to safely cross an intersection a driver needs to take numerous other road users into account, consider their interactions and predict their future behaviors. The complexity of this task provokes driver errors, which results in a significant amount of accidents occurring at intersections. Advanced driver assistance systems (ADAS) can help to prevent such type of accidents by supporting the driver in the assessment of situations and issuing warnings when upcoming conflicts are predicted. In order to arrive at a comprehensive prediction a high number of possibly relevant entities like vehicles, pedestrians and traffic lights has to be taken into account. But the more entities are considered the more complex a suitable prediction model will become - up to a point where the combinatorial explosion of possible future states renders the required computations infeasible. Instead of employing a single, all-encompassing prediction model an apparent alternative solution is to predict the behavior of each road user individually, taking only its relevant surrounding into account. While this solves the problem of combinatorial complexity it poses a new challenge: Determining what the *relevant surrounding entities* are.

M. Platho and H-M. Groß are with Department of Neuroinformatics, Technical University of Ilmenau, D-98684 Ilmenau, Germany {matthias.platho, hm.gross}@tu-ilmenau.de

J. Eggert is with Honda Research Institute Europe GmbH, D-63073 Offenbach am Main, Germany jeggert@honda-ri.de

In recent years an increasing amount of research has been carried out to tackle the problem of situation assessment and behavior prediction. Approaches have been proposed for both sub-problems. Related works can be coarsely divided by their primary goal, namely recognizing specific maneuvers [1], [2], inferring driver intent of a vehicle approaching an intersection [3], [4] or interpreting traffic situations as a whole [5], [6]. In [1] the recognition and prediction of turning maneuvers at an intersection is targeted and in [2] behaviors like overtake and sheer out are predicted. Estimating another driver's intent is the goal in [3]; the authors anticipate turning and stopping maneuvers of a preceding vehicle. In [4] the intent estimation is additionally used to derive a measure for the risk of the present situation. All of the aforementioned works show good results; however, they are tailored to a specific situation consisting of not more than two vehicles. Opposed to that, the authors of [5] consider all present entities and match encountered situations against a predefined set of situation prototypes. While this approach provides a more comprehensive situation assessment, its use of case-based reasoning raises issues concerning stability. The work presented in [6] employs description logic to reason about the relations between cars approaching an intersection. The authors report a successful inference of possible conflicts but also a high computational complexity for the reasoning process.

In this paper we argue that an accurate prediction of the evolution of an intersection situation requires a two-staged approach which *combines situation assessment with behavior prediction*. In the first stage each vehicle's driving situation, determined by its *relevant entities*, are estimated. Subsequently, in the second stage, based on the estimated driving situation, the vehicles behavior is predicted. In order to keep the behavior prediction for a vehicle computationally efficient we limit the prediction to the vehicle's future longitudinal velocity profile, as recent works have shown that it is an important cue for anticipating sudden stops and detecting hazardous maneuvers [7], [3]. In quantitative and qualitative evaluations we will show that our approach leads to simpler, more accurate prediction models, maintains computational feasibility and is superior to methods which do without a previous situation assessment.

The remainder of this paper is structured as follows. In the next sections the two-staged prediction approach is introduced, with Section II-A describing the first stage's method for comprehensive situation assessment and Section II-B detailing the velocity profile prediction. In Section III the evaluation methods for both a quantitative analysis on

simulated data and a qualitative analysis on real-world data are given. The results obtained are presented in Section IV and in Section V we give a brief outlook on future work.

II. PROPOSED APPROACH

In order to detect upcoming conflicts in time, driver assistance systems for intersection scenarios have to predict the evolution of the current situation with acceptable accuracy. As an all-encompassing prediction model incorporating all entities at once will grow overly complex, performing the prediction for road users individually appears to be a more viable approach. For obtaining correct predictions the individual road users must not be considered isolated from their context but with respect to their role in the current situation. Therefore, situation assessment is an essential part of the prediction process and we propose to employ a two-staged method.

The first stage aims at interpreting the present intersection situation comprehensively. Such an interpretation is a challenging problem as intersection scenarios are often highly variable and comprise a large number of relevant entities that are also interdependent. To overcome the problem's high complexity a decomposition of situations into smaller parts that are easier to handle has been suggested, e.g. in [8] and [9]. The intuition behind the decomposition is that each road user adapts its behavior according to a small set of entities that are currently 'relevant'. An entity can be relevant when for example its consideration is necessary to obey a traffic rule - like a red traffic light - or to avoid risks by keeping a safe distance to the car in front. A vehicle's driving situation can be estimated based on type and relationship of its relevant entities, as depicted in step 1 of Figure 1.

In a second stage the velocity profile of each vehicle is predicted. This is accomplished by employing situation-specific prediction models which are selected based on the driving situation estimation in stage one. The prediction model thus needs to take only these entities into account that were identified as relevant in the previous step. This lowers the complexity of the prediction models and allows them to focus on the important features for the prediction. The working principle of the second stage and the prediction results are illustrated in Figure 1 in steps 2 and 3, respectively.

In the following both stages are described in detail.

A. Driving Situation Assessment

As described above, situation assessment is accomplished by decomposing the situation into its constituent parts, using a method presented in [8]. In the cited work the parts that result from situation decomposition are termed *configurations*. Each configuration consists of two entities:

- **Reference entity**: An *affected* entity, in our cases a vehicle.
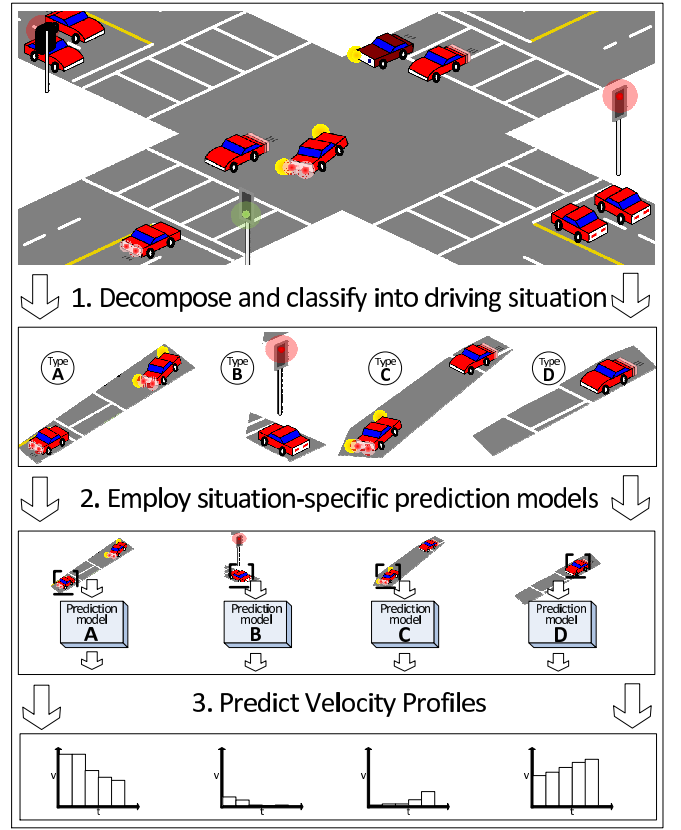- **Affecting entity**: An entity that is relevant for the behavior of the reference entity.

Fig. 1. Overview of the proposed system. At first for each vehicle its driving situation is estimated. Based on this estimation a situation-specific prediction model is selected which predicts the longitudinal velocity profile of the vehicle.

The presence of a configuration is thus expressing a behavioral dependency between two entities. A vehicle is said to "be in a configuration X" if it is the reference entity in configuration X. An exemplary decomposition of a traffic situation into its configurations is given in Figure 2.

For our analysis we considered three types of configurations:

- **Stopped by red traffic light (TL)**: The reference entity has to slow down or stop in order to obey a red traffic light
- **Stopped by leading car (LC)**: The reference entity has to slow down or stop in order to keep a safe distance to a leading car
- **Stopped by intersection (IS)**: The reference entity has to slow down or stop in order to yield to a car with right of way

Note that two cars in Figure 2 are in no configuration (NC) which means that there is currently no entity present that impacts their behavior. This type of driving situation resembles free driving, where the driver's behavior is mainly determined by his own goals.

All configurations have in common that they specify a decelerating maneuver. Such types of maneuvers have been chosen as they are critical for the avoidance of e.g. rear-end crashes and their causes can generally be inferred from the
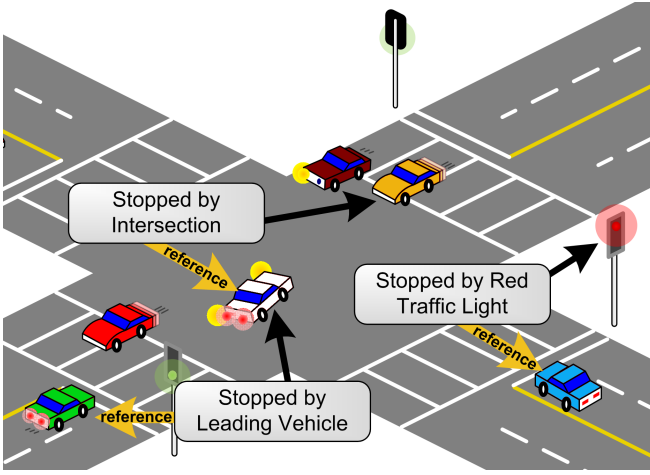
Fig. 2. An exemplary decomposition of an intersection situation into configurations. Each configuration contains an affecting and an affected (reference) entity.

environmental and behavioral context like present obstacles or applicable traffic rules. Additionally, obeying traffic lights, respecting right of way and adapting the velocity to the leading vehicle covers a broad share of everyday driving behaviors.

The structure of each configuration is specified by experts while its parameters are learned from data. The structure comprises a label and involved entities including their relations. Relations serve as features for detecting the configuration and can both be (dynamic) states like *Velocity* and *Traffic Light State* or relative measures like *Distance*. The structure of a configuration consists mainly of relational data which is one of the reasons why configurations are specified using graphs, like the one shown in Figure 3.
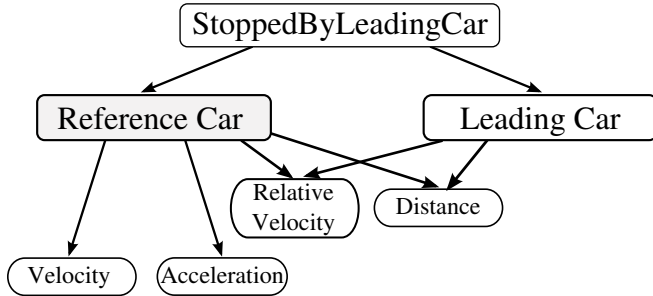

Fig. 3. A schematic representation of the configuration "Stopped by leading vehicle". It contains both unary relations (Velocity, Acceleration) as well as binary relations (Relative Velocity, Distance).

Another reason for choosing graphs is that they can be conveniently mapped to Bayesian Networks [10], which serve as classifiers for recognizing configurations. Bayesian Networks have many desirable properties like the ability to model sensor noise and to cope with missing features. Furthermore their probabilistic nature can be exploited to significantly decrease the complexity of situation assessment [11]. The parameters of the Bayesian Network, its *Conditional Probability Tables*, are learned from training data.

B. Velocity profile prediction

The second stage in the targeted situation prediction system is concerned with predicting the velocity profiles of vehicles with the aid of their previously estimated configuration.

In this paper a velocity profile comprises a prediction horizon of 3 seconds that is discretely sampled at 10 Hz, yielding 30 individual values. For each of the considered configurations plus the case of no configuration a separate prediction model is learned, denoted as $P_{TL}, P_{LC}, P_{IS}, P_{NC}$. The prediction models take the current dynamics of the vehicle and, with the exception of P_{NC}, additional configuration specific features as input, and return the predicted velocities for the next 3 seconds. In the following the implementation of the prediction models will be detailed.

All models employ Random Forest Regressors [12] as prediction method. Random Forest Regression is a state-of-the-art regressor that is known for its robustness against overfitting. In a first step, in order to verify the method's suitability for velocity profile prediction its performance was compared to models using Multiple Linear Regression (MLR). As Random Forest models excelled even complex MLR models significantly they were selected as prediction method.

Altogether, 7 features are used as independent variables for the prediction process which are described in the following. Note that 'target car' refers to the vehicle for which the velocity profile is predicted.

- *Velocity (VEL)*: Velocity of target car in m/s
- *Acceleration (ACC)*: Acceleration of target car in m/s^2
- *Traffic light distance (TLD)*: Distance to the stopping line of the next, relevant traffic light in m
- *Car ahead relative speed (CAS)*: Relative velocity between target car and its leading car in m/s
- *Car ahead distance (CAD)*: Distance between target car and its leading car in m
- *Intersection distance (ID)*: Distance to the entry point of the next intersection in m
- *Time (TIME)*: Time instance for which the velocity is predicted in s. Values are $0.1, 0.2, 0.3...3.0$

Each of the four prediction models uses a proper subset of these features. The reasons are twofold: Firstly, some of the features might not be present at all, for example if there is no traffic light ahead. Secondly, and more importantly, a major benefit of using configuration-specific models is that these models can be limited to the currently relevant features. This does not only lower the complexity of the prediction models but also aids the corresponding learning algorithm in focusing on the important attributes. Table I lists the four prediction models along with their incorporated features.

A close examination of the training data revealed that some features exhibit strong interactions, for example the relative speed of the car ahead (CAS) varies in its impact on the velocity profile depending on the distance to it (CAD). Other features again show for certain prediction models nonlinear, mainly quadratic relationships with the velocity profile. In

order to aid the Random Forest Regressors in capturing these characteristics the prediction models P_{TL} and P_{LC} use non-linearly transformed and multiplicatively combined features as input. For the other models transformations offered no significant gain in accuracy and so they use their features directly.

The most important parameters for adapting the Random Forest algorithm to a given problem are the number of trees learned and their maximal allowed depth. Based on the results of a grid search the number of trees is set to 400 and the maximal depth to 8.

III. EVALUATION

In the evaluation the gain from combining situation assessment and prediction is investigated. The investigation addresses two questions: Firstly, does the two-staged prediction process improve prediction accuracy significantly? Secondly, what is the gain of a preceding situation assessment by itself? The second question aims at determining whether driving situations with the same estimated configuration share a common quality that can be exploited. If this is the case then the performance of situation-specific prediction models will be superior to a general prediction model, even if they are limited to the same set of features.

Both questions were addressed in a quantitative evaluation where data was obtained from a microscopic traffic simulator. Afterwards we present a qualitative analysis that gives an outlook on the applicability of the methods to real-world situations.

In each of the evaluations the goal is to predict the velocity profiles of all vehicles currently present in a scene. A velocity profile consists of 30 velocity values representing a vehicle's velocity for the next 30 time steps, where each time step lasts 0.1 seconds. As error measure the sum of squared distances is used, so that the prediction error e between a predicted profile $\hat{V}$ and the actual profile V is:

$$e = \sum_{i=1}^{30} (\hat{V}_i - V_i)^2 \qquad (1)$$

A. QUANTITATIVE ANALYSIS ON SIMULATED DATA

Data for training and testing is generated using a microscopic traffic simulator presented in [8]. In this simulator each vehicle is controlled by an autonomous agent which actively perceives its environment to select behaviors according to its goals, i.e. arriving at a target position. The implementation gives access to the agent's decision level

TABLE I
FEATURES USED BY THE FOUR PREDICTION MODELS.

Prediction model	VEL	ACC	TLD	CAS	CAD	ID	TIME
$P_{TrafficLight}$	x	x	x				x
$P_{LeadingCar}$	x	x		x	x		x
$P_{InterSection}$	x	x				x	x
$P_{NoConfiguration}$	x	x					x

by which the information *"Which entity lead to the agent's behavior?"* can be obtained. This information serves as ground truth for the vehicle's current configuration.

The simulated scenario comprises a 4-way intersection consisting of a major road with two lanes in each direction that is crossed by a minor road with a single lane in each direction. The intersection is signalized. The simulation creates for each present vehicle at each time instance an individual data sample containing the vehicle's state, its configuration and all the features needed for configuration recognition and velocity profile prediction. As the simulation runs at 10 Hz and simulates on average about 20 vehicles at a time a large number of samples is generated in short time. A run of the simulation was recorded and split into training and testing set, yielding 15361 samples for training and 7084 samples for testing.
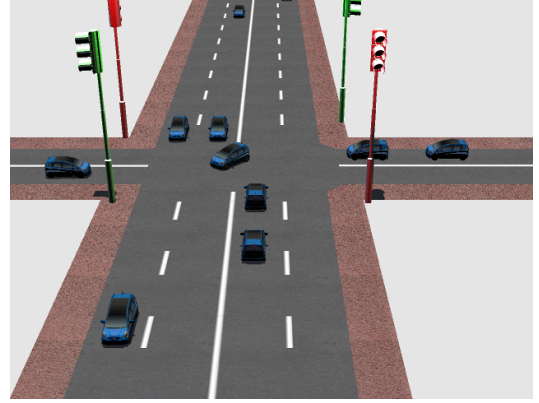

Fig. 4. The simulated intersection as visualized by the CarD microscopic traffic simulator. Traffic lights are not in scale.

In order to answer the first question and evaluate the gain of a two-staged approach its prediction accuracy is compared to two other methods which we term *KINEMATIC* and *PREDONLY*. This part of the evaluation estimates the benefit of the proposed model. *KINEMATIC* serves as a baseline model for judging the difficulty of the prediction task. It predicts the future velocity profile by a straightforward physical extrapolation from the current velocity and acceleration. The predicted velocity $\hat{V}_i$ for time step i given the current velocity v_0 and acceleration a_0 is thus:

$$\hat{V}_i = max(0, v_0 + i * a_0) \qquad (2)$$

Note that the max operator prevents the generation of negative velocities.

The *PREDONLY* model uses also velocity and acceleration as inputs, but in contrast to the *KINEMATIC* model it is learned from the data using a Random Forest, as described in section II-B. This enables the model to capture various characteristics present in the training data, like typical acceleration and deceleration behaviors or maximally allowed speed.

The two-staged prediction model proposed in this paper will be referred to as *TWO-STAGED*. For its configuration-specific prediction models the recorded ground truth of a

vehicle's configuration is only used for training. During testing the configuration is estimated based on recorded features and using the same Bayesian Network as in [8]. Based on the recognized configuration the corresponding prediction model is selected for predicting the velocity profile. The configuration estimation is thus prone to two sources of error for which the overall performance of the two-staged approach has to compensate. A wrong estimation leads to the selection of an inappropriate prediction model which will provide a different profile than the correct model. Additionally, vehicles might change their true configuration during the prediction horizon, for example if a formerly red traffic light turns green. This error, however, affects the other prediction methods as well.

If the *TWO-STAGED* method performs better than the other models it could be attributed to the fact that it uses more features. Though this possibility does not lower the methods usefulness it would be interesting to quantify the benefit of a preceding situation assessment itself. This is accomplished in the second part of the evaluation, where two methods with the same number of features are compared.

For this comparison the *PREDONLY* model is compared with a variant of the *TWO-STAGED* method, where all configuration-specific features are left aside. Thus the method uses solely the two features VEL and ACC as *PREDONLY* does. We refer to this stripped-down method as *TWO-STAGED BASIC*.

B. QUALITATIVE ANALYSIS ON REAL-WORLD DATA

In this analysis the approach is evaluated on exemplary situations taken from a real-world data set. The data set contains 29 minutes of recordings from a test vehicle equipped with video cameras and two laser scanners driving in an urban environment.

In the selected situations the ego vehicle is approaching a signalized intersection. There are four variations of this situation: The relevant traffic light can be red or green and there can be another vehicle driving ahead or not. Based on the data obtained from lidars, cameras and CAN bus the ego vehicle's velocity and acceleration and, if present, distance and relative speed of a leading vehicle are obtained. The traffic light distance is obtained by integrating the velocity backwards from the manually labeled point in time when the vehicle reaches the stop line. A more detailed description of the data and the feature extraction algorithms used can be found in [13].

For this part of the evaluation the prediction performance of *KINEMATIC*, *PREDONLY* and *TWO-STAGED* for three exemplary situations from the data set is compared. The configuration estimation for both training and testing data is obtained from the Bayesian Network that was learned on simulated data for the quantitative analysis. We are aware that in general simulated traffic data can not be used for learning real-world models but we expect that for this restricted approaching scenario it will be sufficient. Because the IS configuration can not be detected due to the limited viewing angle of the sensors used, only NC, LC and TL are modeled.

For a meaningful quantitative evaluation the configuration estimation would have to be learned from the real-word data sets. As this is subject to future work we here assess the performance of the proposed method qualitatively.

The Random Forests for Velocity Profile prediction were learned on 14 minutes of the data set and tested on the remaining 15 minutes. For the qualitative analysis three situations, one for each configuration, were selected. The selection was based on the goal to present both advantages and challenges of the proposed method. For the selected situations the velocity profiles predicted by the three methods are compared.

IV. RESULTS

The results of the quantitative evaluation are shown in Table II. In order to facilitate the comparison between the evaluated methods the *relative* overall velocity prediction error is given, which was obtained by dividing through the overall error of *PREDONLY*.

TABLE II
RELATIVE VELOCITY PREDICTION ERROR

Method	Relative error
TWO-STAGED	0.79
TWO-STAGED (BASIC)	0.93
PREDONLY	1
KINEMATIC	1.8

The relative prediction error of the baseline method *KINEMATIC* is nearly a factor 2 higher than for the learned *PREDONLY* method, which shows that the trivial extrapolation of vehicle dynamics as performed by *KINEMATIC* is not sufficient for an accurately predicted velocity profile. The proposed approach, which considers configuration information, is in turn 21% better than the *PREDONLY* method, which neglects this information. To provide an impression of the performance the mean prediction error over time, relative to the actual velocity, is given in Figure 5(a).
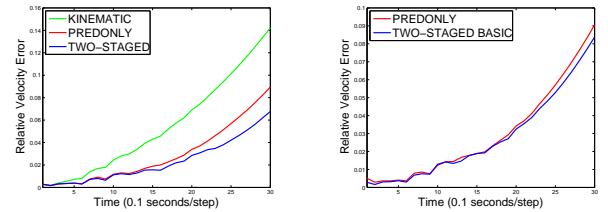
(a) With configuration-specific features

(b) Without configuration-specific features

Fig. 5. The relative prediction error over time is lowest when employing the proposed two-staged method (a), even if no configuration-specific features are used (b).

The plot shows that the *KINEMATIC* method looses track of the actual velocity already after one second, while the other two methods remain accurate for the first two seconds

Estimated Configuration: *No Configuration*

Estimated Configuration: *Stopped by leading car*

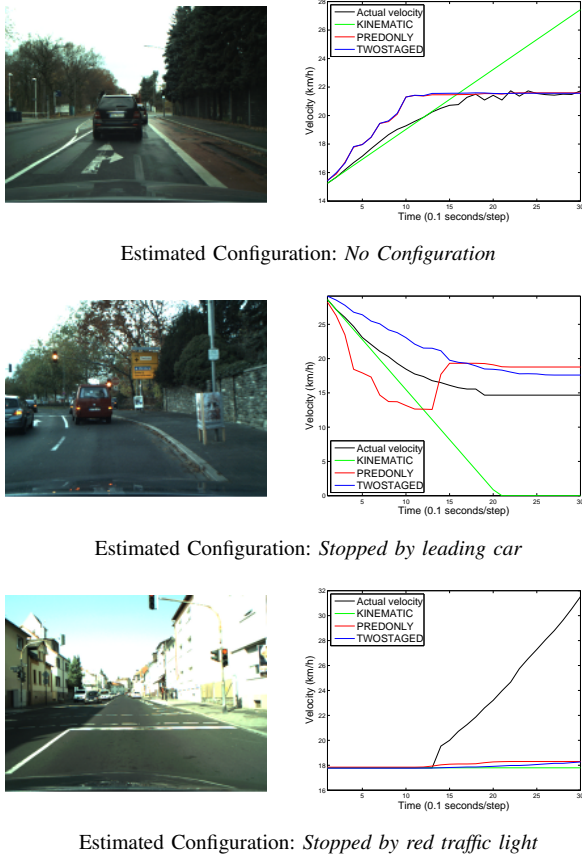Estimated Configuration: *Stopped by red traffic light*

Fig. 6. Camera images of three exemplary driving situations and the corresponding predicted and actual velocity profiles.

of the prediction horizon. While *PREDONLY* and *TWO-STAGED* are hardly discernible for the first 1.3 seconds the proposed method exhibits a significantly less steep rise in prediction error for longer horizons. While on the first look the relative prediction errors of the three methods seem to be too low to actually make a difference a closer look reveals that they actually do: Firstly, about 20% of the training cases contain a stopped car for which the velocity prediction is trivial and so the mean relative error is greatly reduced. Secondly, and more importantly, the proposed method performs far less severe mispredictions than the other methods: When integrating the predicted velocity profiles to obtain a future position it shows that *KINEMATIC* is out by more than an average car length (4 meters) in about 19% of cases and *PREDONLY* is out in 10% of cases while this happens for the *TWO-STAGED* method only in 5%.

In the second part of the quantitative evaluation the pure benefit of combining situation assessment and prediction was examined. The second row in Table II shows that using configuration specific prediction models (*TWO-STAGED BASIC* offers a 7% increase in prediction accuracy even if the specific prediction models are limited to the same features that *PREDONLY* uses. The relative error plot given in Figure 5(b) confirms that a two-staged method provides an increased benefit for longer prediction horizons.

An important finding of the quantitative evaluation is that the significant gain in performance from *TWO-STAGED BASIC* to *TWO-STAGED* was accomplished by using only one to two additional features. This means that the situation assessment in the first stage accomplished to identify the relevant determinants for future behavior which could then be used to improve prediction performance.

For the qualitative analysis three exemplary velocity profiles, one for each considered configuration, are given in Figure 6. Additionally the camera image at the time the prediction was made is given.

The predicted velocity profiles for the *No Configuration* situation match the actual velocity profile closely for all but the *KINEMATIC* method. This demonstrates that a straightforward extrapolation becomes very inaccurate for longer time horizons, which is also confirmed in the other two examples.

The *Stopped by leading car* situation is also quite accurately predicted, though the error is noticeably higher. Nevertheless the *TWO-STAGED* method is able to predict the course of the profile better than the *PREDONLY* method.

The predicted velocity profiles for the third example, a *Stopped by red traffic light* driving situation, show an important challenge of velocity profile prediction: After 1 second the traffic light turns green and all the prediction methods naturally fail to anticipate this state change. While traffic lights are an evident case of sudden, significant changes to the driving situation there are many possible scenarios where such a change can be triggered. A challenge for driving situation specific prediction models will be to capture also the possibility of a change in the driving situation during the prediction horizon.

V. CONCLUSIONS AND FUTURE WORK

In this paper an approach for predicting the evolution of complex traffic situations has been presented. The proposed method first assesses for each individual vehicle its driving situation and uses this information to select the corresponding, specific model for predicting its future velocity profile. Both an extensive quantitative evaluation on simulated data as well as a qualitative evaluation on real-world data confirms the benefits of the proposed method.

The working principle of the proposed two-staged system can be interpreted as the consecutive application of two models: one classifier for the configuration recognition and one regressor for the velocity profile prediction. It could be argued that this can be replaced by a single regressor that uses all of the features of both models as inputs. While such a single regressor is advantageous in case sufficient training data is available to cover the now larger state space adequately there are multiple objections to this approach. Besides increasing the complexity of the learning task substantially it also hinders the use of selective perception methods. Such methods, where in each time step only the currently most informative features are measured, are able

to save valuable computational resources of online systems [11].

In future work the applicability to a real-world system will be further investigated. In order to cope with configurations changing during the prediction horizon an additional prediction model on configuration level will be examined. Further research directions are the inclusion of more configuration types and obtaining the ability to learn the configuration estimation models from real-world data.

REFERENCES

[1] E. Käfer, C. Hermes, C. Wöhler, H. Ritter, and F. Kummert, "Recognition of situation classes at road intersections," in *Proc. IEEE Int Robotics and Automation (ICRA) Conference*, 2010, pp. 3960–3965.

[2] T. Gindele, S. Brechtel, and R. Dillmann, "A probabilistic model for estimating driver behaviors and vehicle trajectories in traffic environments," in *Proc. 13th Int Intelligent Transportation Systems (ITSC) IEEE Conf*, 2010, pp. 1625–1631.

[3] M. Liebner, M. Baumann, F. Klanner, and C. Stiller, "Driver intent inference at urban intersections using the intelligent driver model," in *Proc. IEEE Intelligent Vehicles Symp. (IV)*, 2012, pp. 1162 – 1167.

[4] S. Lefèvre, C. Laugier, and J. Ibanez-Guzmán, "Risk assessment at road intersections: Comparing intention and expectation," in *Proc. IEEE Intelligent Vehicles Symp. (IV)*, 2012, pp. 165 – 171.

[5] S. Vacek, T. Gindele, J. M. Zöllner, and R. Dillmann, "Situation classification for cognitive automobiles using case-based reasoning," in *Proc. IEEE Intelligent Vehicles Symp*, 2007, pp. 704–709.

[6] M. Hülsen, J. M. Zöllner, and C. Weiss, "Traffic intersection situation description ontology for advanced driver assistance," in *Proc. IEEE Intelligent Vehicles Symp. (IV)*, 2011, pp. 993–999.

[7] P. Angkititrakul, C. Miyajima, and K. Takeda, "Analysis and prediction of deceleration behavior during car following using stochastic driver-behavior model," in *Proc. IEEE Intelligent Transportation Systems (ITSC)*, 2012.

[8] M. Platho, H.-M. Groß, and J. Eggert, "Traffic situation assessment by recognizing interrelated road users," in *Proc. IEEE Intelligent Transportation Systems (ITSC)*, 2012.

[9] T. Schamm and J. M. Zöllner, "A model-based approach to probabilistic situation assessment for driver assistance systems," in *Proc. IEEE Intelligent Transportation Systems (ITSC)*, 2011.

[10] J. Pearl, "Bayesian networks: A model of self-activated memory for evidential reasoning," in *Proceedings of the 7th Conference of the Cognitive Science Society, Irvine*, Aug. 1985, pp. 329–334.

[11] M. Platho and J. Eggert, "Deciding what to inspect first: Incremental situation assessment based on information gain," in *Proc. IEEE Intelligent Transportation Systems (ITSC)*, 2012.

[12] L. Breiman, "Random forests," *Machine Learning*, vol. 45, pp. 5–32, 2001.

[13] M. Garcia Ortiz, J. Schmüdderich, F. Kummert, and A. Gepperth, "Situation-specific learning for ego-vehicle behavior prediction systems," in *Proc. IEEE Intelligent Transportation Systems (ITSC)*, 2011.