

# Teaching Vehicles to Anticipate: A Systematic Study on Probabilistic Behavior Prediction Using Large Data Sets

Florian Wirthmüller<sup>1</sup>, Julian Schleichtriemen<sup>1</sup>, Jochen Hipp<sup>1</sup>, and Manfred Reichert<sup>2</sup>

**Abstract**—By observing their environment as well as other traffic participants, humans are enabled to drive road vehicles safely. Vehicle passengers, however, perceive a notable difference between non-experienced and experienced drivers. In particular, they may get the impression that the latter ones anticipate what will happen in the next few moments and consider these foresights in their driving behavior. To make the driving style of automated vehicles comparable to the one of human drivers with respect to comfort and perceived safety, the aforementioned anticipation skills need to become a built-in feature of self-driving vehicles. This article provides a systematic comparison of methods and strategies to generate this intention for self-driving cars using machine learning techniques. To implement and test these algorithms we use a large data set collected over more than 30 000 km of highway driving and containing approximately 40 000 real-world driving situations. We further show that it is possible to classify driving maneuvers upcoming within the next 5 s with an Area Under the ROC Curve (AUC) above 0.92 for all defined maneuver classes. This enables us to predict the lateral position with a prediction horizon of 5 s with a median lateral error of less than 0.21 m.

**Index Terms**—Automated driving, advanced driver assistance systems, maneuver classification, trajectory prediction, vehicle position prediction, Gaussian mixture regression, mixture of experts.

## I. INTRODUCTION

**A**UTOMATED driving has the potential to radically change our mobility habits as well as the way goods are transported. To enable driving automation, several processing steps have to be executed. Fig. 1 illustrates this thought: In the first step, the current traffic scene has to be sensed and a proper representation of the environment needs to be

Manuscript received October 2, 2019; revised April 19, 2020; accepted May 21, 2020. The Associate Editor for this article was M. Mesbah. (Corresponding author: Florian Wirthmüller.)

Florian Wirthmüller is with Mercedes-Benz AG Research and Development, 71034 Böblingen, Germany, and also with the Institute of Databases and Information Systems (DBIS), Ulm University, 89081 Ulm, Germany (e-mail: florian.wirthmueller@daimler.com).

Julian Schleichtriemen is with Mercedes-Benz AG Research and Development, 71034 Böblingen, Germany, and also with the Institute of Realtime Learning Systems, University of Siegen, 57076 Siegen, Germany.

Jochen Hipp is with Mercedes-Benz AG Research and Development, 71034 Böblingen, Germany.

Manfred Reichert is with the Institute of Databases and Information Systems (DBIS), Ulm University, 89081 Ulm, Germany.

This article has supplementary downloadable material available at <http://ieeexplore.ieee.org>, provided by the authors.

Digital Object Identifier 10.1109/TITS.2020.3002070




Fig. 1. Long-term driving behavior predictions in the context of trajectory planning for automated driving (equal symbols denote simultaneity).

generated. Using this information, the given traffic situation needs to be interpreted and the behavior of others has to be anticipated. Subsequently, a plan, i.e. a trajectory, is derived based on this knowledge. Finally, this plan is executed in the last step of this process. How long the trajectory stays viable, before it has to be re-planned, is strongly influenced by the capability of the prediction component.

As opposed to other research works dealing with techniques to interconnect vehicles through a so called car-to-car communication, we aim to solve this anticipation task locally. On one hand, it is not foreseeable when an adequate market penetration of vehicles with such techniques will be reached. On the other, a local prediction component always becomes necessary, as there are several traffic participants without communication abilities such as bicyclists. In addition, local predictions might become necessary to bypass transmission times in certain cases as emphasized by [1]. Moreover, it is reasonable to approach the topic from the perspective of highway driving, as this use case is easier to realize than others due to its clear constraints (e.g. structured setting, absence of pedestrians). However, for the prediction task this implies the challenge to create precise long-term predictions (2 to 5 s) rather than short forecasts (up to 2 s), as in highway scenarios higher velocities can be expected than in urban or rural areas.

### A. Problem Statement

We tackle the challenge of anticipating the behavior of other traffic participants in highway scenarios. In particular, we aim to generate information that can be processed by trajectory planning algorithms to implement an anticipatory driving style.

In this context, our objective is to model future vehicle positions within a time  $t$  in longitudinal  $x_t$  and lateral  $y_t$  direction as spatial distributions  $x_t \sim p_x$ ,  $y_t \sim p_y$  rather than estimating single shot predictions  $\hat{x}_t$  and  $\hat{y}_t$  respectively. Note that these distributions are more useful for down-streamed criticality assessments as they enable us to represent several alternative hypotheses at a time with their particular frequencies. Despite the focus on highway driving, the presented methods shall be general enough to be appropriate in other environments as well.

### B. Problem Resolution Strategy

This article presents a systematic workflow for the design and evaluation of a lightweight maneuver-based model [2], which uses standard sensor inputs to perform long-term driving behavior predictions. Methodically, we build on [3] and use a two-step Mixture of Experts (*MOE*) approach. This includes a maneuver classification and a down-streamed behavior prediction. The maneuver probabilities  $\{P_m\}_{\forall m \in M}$  determined by the classifier are used in the Mixture of Experts approach as gating nodes. Specifically, the probabilities control the weighting  $w_m$  of the respective expert distributions  $p_{y,m}$ , while calculating the overall distribution of future vehicle positions  $p_y$ . Eq. 1 summarizes this procedure for the lateral direction (equivalent for  $x$ ):

$$\begin{aligned} y_t &\sim p_y(\Theta_y, I, t) \\ &= \sum_{m \in M} p_{y,m}(\theta_{y,m}|I, t) \cdot w_m(I) \end{aligned} \quad (1)$$

The set of maneuvers  $M$  is defined as follows:

$$M = \{LCL, FLW, LCR\} \quad (2)$$

Different weighting approaches based on the maneuver probabilities are presented in Sec. VII. The expert distributions  $p_{y,m}$  are modeled as Gaussian Mixture Models (*GMMs*) in the combined input and output space with  $K$  components according to Eq. 3, and are used in a Gaussian Mixture Regression manner. Hence, they are conditioned by the input features  $I$  and the prediction time  $t$  (cf. Eq. 1).

$$p_{y,m}(\theta_{y,m}) = \sum_{i=1}^K \phi_{y,m,i} \cdot \mathcal{N}(\mu_{y,m,i}, \Sigma_{y,m,i}) \quad (3)$$

The parameters of the *GMMs* are subsumed in  $\Theta_y$ :

$$\Theta_y = \{\theta_{y,m}\}_{\forall m \in M} = \{\phi_{y,m}, \mu_{y,m}, \Sigma_{y,m}\}_{\forall m \in M} \quad (4)$$

In addition, we introduce an alternative methodology to the Mixture of Experts approach, integrating the outputs of the gating nodes into one single model. This simplifies Eq. 1 as follows:

$$y_t \sim p_y(\theta_{y,IGMM}|I, t, P_{LCL}(I), P_{LCR}(I)) \quad (5)$$

For implementing the models, we use out-of-the-box modules from the widely used frameworks Apache Spark MLlib [4] (classifiers) and Scikit-learn [5] (*GMMs*).

Altogether, we contribute a systematic workflow for designing and evaluating the prediction models as well as methodical

extensions to known approaches. Moreover, we assess the performance of the developed modules for the two tasks of predicting (1) driving maneuvers and (2) probability distributions of future positions both separately and in combination. To evaluate the modules, we utilize a large data set comprising real-world measurements. As will be shown, our prediction models outperform established state-of-the-art approaches.

The remainder of this article is organized as follows: Sec. II discusses related work on object motion prediction, emphasizing the value added by our approach. Sec. III introduces the data set and describes the preprocessing steps applied to it. Sec. IV outlines the training of the considered maneuver classifiers, whereas Sec. V deals with the experimental evaluation and the performance of the classifiers. Based on these findings, Sec. VI develops different approaches for estimating probability distributions of future vehicle positions, which are then assessed in Sec. VII. Finally, Sec. VIII summarizes the article and gives an outlook on future work.

## II. RELATED WORK

Regarding the understanding and prediction of the behavior of other traffic participants in highway scenarios, various aspects were investigated in literature. Accordingly, this section is sub-divided into three parts: Sec. II-A presents approaches inferring the kind of maneuver that will be executed by a vehicle. Note that applications like collision checkers or trajectory planning algorithms cannot directly process such kind of information. Instead, probabilities of future vehicle positions or trajectories need to be predicted. Related research on this topic is presented in Sec. II-B. Bringing together the aspects of maneuver classification and position prediction, Sec. II-C gives an overview of hybrid prediction approaches. Finally, Sec. II-D closes the section with a brief literature discussion, leading to the contributions of this article in Sec. II-E.

### A. Classification Approaches

Classification approaches for maneuver recognition are described in [1], [6]–[8]. In [1], a system is introduced, which is capable of detecting lane changes with high accuracies (>99%), approximately 1s before their occurrence. For this purpose, dynamic Bayesian networks are used. Another approach, which is capable of detecting lane changes approximately 1.5s before their occurrence, is presented in [6]. To achieve this, the lane change probability is decomposed into a situation- and a movement-based component, resulting in an  $F_1$ -score better than 98%. The approach presented in [7], in turn, shows that it is possible to detect lane changes up to time horizons of 2s when using feature selection for scene understanding, with an Area Under the Curve (*AUC*) better than 0.96. Moreover, [8] combines interaction-aware heuristic models with an interaction-unaware learned model. The interaction-aware component relies on a multi agent simulation based on game theory, in which each agent simultaneously tries to minimize different cost functions. These cost functions are designed using expert knowledge and consider traffic rules. In a second step, the output of the interaction model is used to

condition an interaction-unaware classifier based on Bayesian networks. The approach is able to detect lane changes on average 1.8 s in advance, with an *AUC* better than 0.93.

### B. Trajectory and Position Prediction Approaches

Approaches dealing with the prediction of trajectories and positions are presented in [9]–[13]: [9] uses a fully-connected Deep Neural Network to learn the parameters of a two-dimensional *GMM*. For each situation, an adapted Gaussian Mixture distribution models the probability density in the output dimensions  $a_x$  and  $v_y$  (cf. Tab. XII). This distribution is then sampled to estimate trajectories. The authors evaluate their approach with the widely used NGSIM data set [14] and show that a root weighted square error (comparable to *RMSE*) of approximately 0.5 m in lateral direction at a prediction horizon of 5 s can be achieved.

Another approach, also evaluated with the NGSIM data set, is presented in [10]. The authors propose the use of a Long Short Term Memory network for predicting trajectories. In particular, the approach is able to compute single shot predictions with an *RMSE* of approximately 0.42 m at a prediction horizon of 5 s. Reference [11] deals with the prediction of spatial probability density functions, especially at road intersections. More precisely, a conditional probability density function, which models the relationship between past and future motions, is inferred from training data. Finally, standard *GMMs* and variational approaches are compared. In [12], this approach is extended by a hierarchical Mixture of Experts that allows to incorporate categorical information. The latter includes, for example, the topology of a road intersection.

In [13], a Gaussian Mixture Regression approach for predicting future longitudinal positions as well as a procedure for estimating the prediction confidence are introduced.

### C. Hybrid Approaches

Approaches that combine strategies for both maneuver detection and trajectory or position prediction, similar to the approach presented in this article, are described in [15]–[20]. In the following, we denote such approaches as hybrid.

Reference [15] presents a two-staged approach: In the first step, a Multilayer Perceptron (*MLP*) is used to estimate the future lane of a vehicle. In a second step, a concrete trajectory realization is estimated with an additional *MLP*. As a result, the lane estimation module is able to detect lane changes 2 s in advance with an *AUC* better than 0.90. The evaluation of the trajectory prediction module shows a median lateral error of approximately 0.23 m at a prediction horizon of 5 s.

Reference [16] proposes another hybrid approach that uses the prediction of future trajectories to forecast lane change maneuvers. Moreover, the intention of drivers is modeled using a Support Vector Machine. Subsequently, the resulting action is checked for collisions. This enables the approach to model interrupted lane changes. During the evaluation, an  $F_1$ -score of 98.1 % with a detection time up to 1.74 s is achieved.

In turn, [17] does not follow such a hybrid approach, but contains an intermediate step before predicting trajectories. Instead of learning maneuver probabilities, the authors present

a regression technique for estimating the time span to the next lane change relying on Random Forests. In [18], this approach is extended and combined with findings from [6]. The estimated time up to the next lane changes to the left and to the right are used as input for a cubic polynomial which is intended to predict future trajectories. Finally, the approach is evaluated with the mentioned NGSIM data set, showing a median lateral error of approximately 0.5 m at a prediction horizon of 3 s for lane changing scenarios, assuming a perfect maneuver classification.

Reference [19] proposes the use of a maneuver recognition based on a Hidden Markov Model, distinguishing between ten maneuver classes. Based on this model, a position prediction module, which combines several maneuver specific variational *GMMs* (according to [11]) and an Interacting Multiple Model, which weights different physical models against each other, are implemented. As the approach uses ten maneuver classes and as the errors are only measured in terms of Euclidean distance, the results are difficult to compare with the ones of other approaches. Additionally, the approach is evaluated on a rather small data set. Finally, in [20] these findings are pursued by the use of a Long Short Term Memory network. The authors demonstrate certain improvements compared to their previous work, while using the NGSIM data set for evaluation purposes.

Reference [3] presents an approach predicting future lateral vehicle positions based on Gaussian Mixture Regression and a Mixture of Experts with a Random Forest as gating network. The approach is evaluated based on a small data set, leading to noisy results, especially in case of lane changes. The evaluation shows that the approach is able to perform maneuver classifications with an *AUC* better than 0.84 and lateral position predictions with a median error of less than 0.2 m at a prediction horizon of 5 s.

### D. Discussion

The findings of our literature survey can be summarized as follows: Many works provide meaningful algorithmic contributions. However, in numerous cases we miss structure regarding the problem resolution strategy. Often, it does not become clear how the approaches compare to any baseline (e. g. [19]). Moreover, parameters (e. g. [16]) and feature sets (e. g. [10]) are selected manually, and are thus difficult to retrace. In addition, most approaches focus on short or medium prediction horizons (e. g. [1]), or lack a good prediction performance for larger time-horizons (e. g. [18]). When analyzing the approaches that aim to resolve the long-term prediction problem, it becomes clear that the latter is challenging as the prediction models become significantly more complex as, e. g., pointed out by [7], [8] and [21].

Moreover, many approaches (e. g. [10]) aim to predict single trajectories or single shot predictions rather than probabilistic distributions of future vehicle positions. Therefore, the objective to be optimized is mostly the root-mean-square error (*RMSE*). As opposed to these works, we consider the objective of the learning problem as generating an estimator that models a probability distribution of positions reflecting the frequencies of all observed positions, e. g., for different drivers



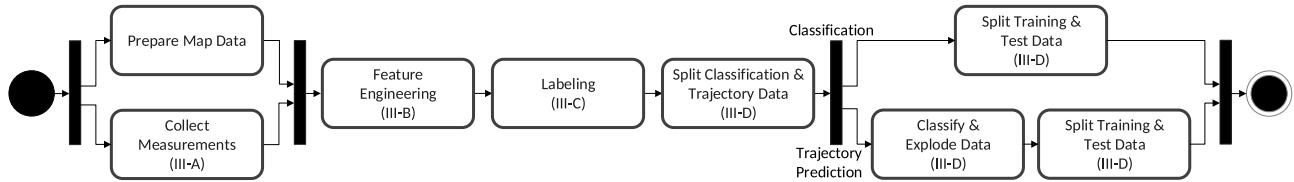


Fig. 2. Preprocessing steps used in the proposed workflow (respective sections are referred in the boxes).

in the same situation. Thus, we aim to maximize the likelihood of truly occupied positions given the model. As reasoning behind this design choice, such distributions contain significantly more information than single shot predictions. Thus, they are more useful for applications that need to consider risks, like, for example, maneuver planning approaches as presented in [11], [22], [23].

### E. Contributions

The contribution of this article is threefold:

- 1) We apply a heuristic-free machine learning workflow to generate a model capable of predicting maneuvers and precise distributions of future vehicle positions for time horizons up to 5 s (reasonable in terms of comparability). This is achieved with a machine learning workflow that omits any human tuned (hyper-) parameters when constructing the classifiers. Note that this includes all aspects involving feature engineering, labeling, feature selection, and hyperparameter optimization for different classification algorithms. Regarding feature engineering and selection, this means that we construct a data set with a large superset of all features, which are potentially relevant for the problem solution beforehand. Afterwards we select a more or less small feature set that still ensures maximum predictive power through an automated feature selection process.
- 2) We evaluate the modules for maneuver classification and position prediction, where both parts are not only evaluated separately, as in other works (e.g. [18]), but as a combined prediction system as well. This concerns the lateral as well as the longitudinal behavior. In this context, we show that directly feeding the results of the classifier into the regression problem produces results comparable to an Mixture of Experts approach. Additionally, we show that relying on the Markov assumption and not modeling the interactions between the traffic participants explicitly, allows producing superior results compared to existing approaches. As opposed to these works, we integrate the different aspects of behavior prediction, which comprise the prediction of driving maneuvers and positions both in lateral and longitudinal direction. In addition, we introduce new methodologies and conduct a large-scale evaluation.
- 3) We demonstrate that the presented methods not only have the potential to outperform state-of-the-art approaches when feeding them with a sufficient number of data. Additionally, we show that our approach is able to provide a meaningful estimate of the prediction

uncertainty to the consumer of the information, which is beneficial for collision risk calculation and trajectory planning (e.g. [22]).

## III. DATA PREPARATION & EXPERIMENTAL SETUP

Sec. III-A introduces the considered data set and the experimental setup. Sec. III-B then gives a detailed overview of the features used to train our models. Afterwards, Sec. III-C introduces the labeling process. Finally, Sec. III-D deals with the data set split for training, validating and testing the constructed models as well as further preprocessing steps. Fig. 2 summarizes the overall preprocessing workflow.

### A. Data Collection

For modelling and evaluating our modules, we use measurement data from a fleet of testing vehicles [24] equipped with common series sensors. The sensor setup includes a front-facing camera detecting lane markings as well as two radars observing the traffic situation in the back. In addition, the vehicles have a front-facing automotive radar to sense the distances and velocities of surrounding vehicles. The data has been collected with different vehicles and drivers at varying times of the day during all seasons. The data collection campaign spanned over more than a year and was mainly restricted to the area around Stuttgart in Germany. Through the wide variance, we are expecting our models to achieve good generalization characteristics.

Unlike other contributions (e.g. [3]), we are not using the actual object-vehicles as prediction target  $o$  in this work, but rather the ego- (or measurement-) vehicle itself. However, as our work of course focuses on the prediction of surrounding vehicles, we solely use features that are observable from an external point of view, as postulated in other works (e.g. [1] or [16]). Note that this constraint excludes features like driver status or steering wheel angle. Thus, the models remain applicable to actual object-vehicles, assuming a good sensing of their surrounding. Working with the ego-vehicle data offers several advantages concerning the modeling of situations: First, each situation can be described in a similar way, as situations in which relevant neighboring vehicles to the target-vehicle are hidden for the measurement-vehicle can not occur. In addition, all measurements span longer time periods as the target-vehicle can never disappear from the field of view. This way of data handling is widespread in literature (e.g. [6]). In addition, one can expect that future sensor setups will minimize measurement uncertainty for perceived objects and will get closer to the data quality that is nowadays available for the ego-vehicle.




Fig. 3. Environment model used for our investigations.

Basically, our investigations rely on a similar environment model than the one presented in [7], modeling the surrounding with a fixed grid of eight relation partners. But opposed to [7], we use the ego-vehicle as prediction target. For this purpose, we slightly adapt the environment model: As the sensors facing the rear traffic in the testing vehicles are less capable than the ones facing the front, our environment model (cf. Fig. 3) distinguishes between relation partners behind (index  $rb$ ) and in front of (index  $rf$ ) the prediction target  $o$ . Thus, the relation vectors of the rear objects  $R_{rb}$  are shortened compared to the ones of the front objects  $R_{rf}$ . The relation vectors describe the relation between the respective object and the prediction target. Object-vehicles on the same lane as  $o$  and driving behind  $o$  are left out, as the current sensor setup is not able to sense them. Consequently, a traffic situation can be described by the feature vector  $F_{sit}$ , which contains the relations of  $o$  and its seven relation partners, its own status  $F_o$ , and the infrastructure description  $F_{infra}$  (cf. Eq. 6):

$$\begin{aligned}
 F_{sit} = [ & R_{rf}(r = fl), R_{rf}(r = f), R_{rf}(r = fr), \\
 & R_{rf}(r = l), R_{rf}(r = r), \\
 & R_{rb}(r = rl), R_{rb}(r = rr), \\
 & F_o, F_{infra}]^T \quad (6)
 \end{aligned}$$

A detailed listing of the particular elements of the relation vectors  $R_{rf}$  and  $R_{rb}$  as well as  $F_o$  and  $F_{infra}$  can be found in Tab. XII.

### B. Feature Engineering

To test and develop our system and to fill the described environment model, we use fused data originating from three different sources:

- 1) The basis for our investigations are measurement data produced by the testing fleet (cf. Sec. III-A).
- 2) As we identified additional features being of interest as inputs beforehand, we fuse the data with information from a navigation map (e.g. bridges, tunnels, and distances to highway approaches).

- 3) Besides, we calculate some higher order features out of the measurements, as e.g. a conversion to a curvilinear coordinate-system along the road [25].

### C. Labeling

Like previous works [3], we divide all samples into the three maneuver classes *LCL* (lane change left), *FLW* (lane following), and *LCR* (lane change right) and apply a labeling process that works as follows: First, for each measurement, the times up to the next lane change to the left neighboring lane (*TTLCL*) and to the right one (*TTLCR*) respectively are calculated. This is accomplished by a forecast in time with the distances to the lane markings. As the moment of the lane change, we define the point in time when the vehicle center has just crossed the lane marking. Subsequently, we determine the maneuver labels of each sample based on a defined prediction horizon  $T_h$  according to Eq. 7:

$$L = \begin{cases} LCL, & \text{if } (TTLCL \leq T_h) \wedge \\ & (TTLCL < TTLCR) \\ LCR, & \text{if } (TTLCR \leq T_h) \wedge \\ & (TTLCR < TTLCL) \\ FLW, & \text{otherwise} \end{cases} \quad (7)$$

We decided to use a horizon of 5s, as the duration of lane change maneuvers usually ranges from 3s to 5s (see [16]). Consequently, it is reasonable to label samples only to an upper boundary of 5s as potential lane change samples. Additionally, this value is widely used in literature as longest prediction time (e.g. [8], [15] or [16]) and, therefore, it allows for comparability. However, note that this style of labeling might result in decreased performance values, as detections being slightly more than 5s ahead of a lane change count as false positives in the evaluation.

### D. Data Set Split

As shown in Fig. 2, we split our data into several parts after executing the mentioned preprocessing steps. The first split divides our data into one part for the maneuver classification  $D^{Ma}$  and another one for the position prediction  $D^{Po}$ . This allows us to produce models based on independent data sets. An overview of the splits as well as the respective data set sizes and identifiers is given in Tab. I.

The first part  $D^{Ma}$  is then used as follows: To prepare the training, parametrization and evaluation of the developed classifiers as well as to stay methodically straight, we split data set  $D^{Ma}$  once more into six folds.<sup>1</sup> Thereof we use five folds  $D_{TV}^{Ma}$  in Sec. IV for the design and parametrization. The remaining fold  $D_6^{Ma} = D_{Te}^{Ma}$  is only used for the performance examinations presented in Sec. V. The split is performed based on entire situations as described in [3]. This means that the measurements of each situation solely occur in one of the folds. Note that this ensures the absence of unrealistic results, which might occur due to similar samples from the same

<sup>1</sup>As shown in the following sections, the amount of folds is a trade-off between computability and correctness

TABLE I  
DATA SET IDENTIFIERS AND SIZES

Maneuver Data: $D^{Ma}$						Position Data: $D^{Po}$		
Training & Validation: $D_{TV}^{Ma}$					Test: $D_{Te}^{Ma}$	Training: $D_T^{Po}$ 130 623 Trajectories (7 s; variable sampling time)		Test: $D_{Te}^{Po}$ 20 000 Trajectories (5 s; 10 Hz)
$D_1^{Ma}$	$D_2^{Ma}$	$D_3^{Ma}$	$D_4^{Ma}$	$D_5^{Ma}$	$D_6^{Ma}$	$D_{T,LCR}^{Po}$	$D_{T,FLW}^{Po}$	$D_{T,LCR}^{Po}$
Samples per Maneuver Class:						Selected <sup>3</sup> Trajectories:		
90 759	87 499	89 048	90 458	92 669	87 308	3 685	6 037	5 071

time series in the evaluation and trainings data otherwise. To achieve an even proportion of the three maneuver classes, we balance the number of samples within each fold by a random undersampling strategy. As the prediction problem is extremely unbalanced, as outlined in [10], classifiers would focus on the most frequent maneuver class *FLW* otherwise. In our case approximately 94 % of the data points belong to that class.

In addition, we only take situations into account that were collected continuously up to the prediction horizon of 5 s. This ensures that the folds are also balanced over time, which constitutes a prerequisite for performing fair evaluations. This is necessary, as the prediction task is obviously much more demanding when predicting a lane change 4 s in advance instead of 1 s in advance. Due to this strategy, the numbers of samples in the six folds are slightly different, but we consider this as uncritical. Overall,  $D^{Ma}$  contains approximately 8 hours of highway driving of which  $\frac{2}{3}$  are collected right during lane changes.

The second data set  $D^{Po}$ , which serves for the training and evaluation of the position prediction, is processed as follows: Initially, we add the lane change probabilities as estimated by the different classifiers to each sample. Furthermore, we only consider measurements that were collected when the vehicle was manually driven. Note that this restriction is essential as all vehicles of our testing fleet are equipped with an Adaptive Cruise Control (ACC) system. Thus, driving in a semi-automated mode is over-represented in our data set compared to reality.<sup>2</sup>

We further split data set  $D^{Po}$  into the subsets  $D_T^{Po}$  for training and  $D_{Te}^{Po}$  for evaluating the position predictions (cf. Sec. VI and Sec. VII). Afterwards, we expand each data point in  $D_T^{Po}$  with the desired prediction outputs, i.e., the true positions in  $x$  and  $y$  direction for all times  $t \in T_T = \{-1.0\text{s}, -0.9\text{s}, \dots, 6.0\text{s}\}$ . Note that the samples with negative times and the ones with times  $>5\text{s}$  are needed to train the distributions correctly. Strictly limiting the times to a certain range would generate areas in the data space, which are difficult to represent with *GMMs* due to discontinuities similar to the ones in the probability dimension (cf. Sec. VI-B). To overcome these problems, we integrated a mechanism performing a subsampling between -1 s and 0 s as well as between 5 s and 6 s

<sup>2</sup>We do not explicitly filter out ACC driving in the data set for maneuver classification, as we can assume that ACC is always deactivated during lane changes.



Fig. 4. Process of training and evaluating maneuver classifiers.

according to a Gaussian distribution (percentiles:  $P_{50} = 0.0\text{s}$ ;  $P_{-3\sigma} = -1.0\text{s}$ ; equivalent between 5 and 6 s).

Another mechanism performing a time interpolation ensures that the training data points are distributed continuously along the time dimension. Accordingly, we also have access to prediction times in between our sampling times during the training process. Moreover, the data points in the position test data set  $D_{Te}^{Po}$  are expanded with  $x$  and  $y$  positions as well as corresponding times  $t \in T_{Te} = \{0.0\text{s}, 0.1\text{s}, \dots, 5.0\text{s}\}$ .

Finally, we 'coil' the two data sets  $D_T^{Po}$  &  $D_{Te}^{Po}$  such that each of the newly constructed data points contains the features at the start point of the prediction, one corresponding prediction time, and the actual  $x$  and  $y$  positions at that point in time (in Fig. 2 this step is called 'Explode Data'). Hence, our data sets are multiplied by a factor of  $|T_T| = 71$  respectively  $|T_{Te}| = 51$  and are structured as described in Sec. VII-A. Note that  $D_T^{Po}$  is re-split along the maneuver labels and undersampled in Sec. VI-A, to train maneuver specific position prediction experts.

#### IV. MANEUVER CLASSIFIER TRAINING

This section gives an overview of the different techniques used for feature selection (cf. Sec. IV-A), classification algorithms (cf. Sec. IV-B), and techniques to tune the respective hyperparameters (cf. Sec. IV-C) for the maneuver classification. The corresponding activities are illustrated by Fig. 4.

##### A. Feature Selection

This section deals with the task of selecting a meaningful subset of features from the available superset. Such selection makes sense for two reasons: First, it can improve the prediction performance of the maneuver classifiers. Second, it can help to reduce calculation efforts, enabling predictions on devices with limited computational power as well. Our main goal here is to improve the overall prediction performance. Note that this slightly contrasts with an overall ranking of the available features, as some of them are highly redundant. Consequently, the most predictive variables shall be selected,

<sup>3</sup>for details see Sec. VI-A

TABLE II  
SUMMARY OF EXAMINED FEATURE SELECTION TECHNIQUES

Variant	Description
<i>A</i>	Superset as Baseline
<i>B</i>	Correlation Threshold
<i>C</i>	<i>CFS</i>
<i>D</i>	Wrapper Technique

while excluding redundant ones. In literature, one can find numerous works dealing with feature selection in machine learning applications. In our implementation, we rely on the findings from [26]. As we claim to solve the underlying classification problem through a systematic machine learning workflow, we start with simple techniques and move towards more sophisticated and computationally expensive ones. To demonstrate the performance of the used techniques, additionally, we test the classification with the entire superset as a baseline. The superset that contains all features is denoted as *A* in the following.

The first investigated technique is a simple correlation-based feature selection technique, which evaluates the correlation of all features and then applies a threshold (set to 0.15) to remove features showing a very low correlation with the maneuver class from the superset. More precisely, we compute Spearman's Correlation (see [27, p. 133 ff]) between each feature and the time up to the next lane change (*TTLC*). We selected this quantity instead of the maneuver label, as it enables a smooth fade-out. The resulting feature set is denoted as *B* in the following. Tab. II summarizes the examined variants and their abbreviations. Finally, the elements of the resulting feature sets can be found in Tab. XII.

The second technique uses the Correlation-based Feature Selection (*CFS*; cf. [28]) and is referred to as *C* in the following. For this technique, the correlation of entire feature sets instead of single features is calculated. More precisely, for all feature sets *S*, the 'merit'  $M_S$ , as a measure of the predictive performance, is computed according to Eq. 8:

$$M_S = \frac{n \overline{\rho_{cf}}}{\sqrt{n + n(n-1) \overline{\rho_{ff}}}} \quad (8)$$

$n$  describes the number of features and  $\overline{\rho_{cf}}$  corresponds to the mean correlation of all features with the class label or, in our case, *TTLC*. Variable  $\overline{\rho_{ff}}$ , in turn, describes the mean feature-feature inter-correlation of all features within *S*. As can be seen from Eq. 8, strongly correlated features in a feature set *S* minimize  $M_S$ , whereas a stronger correlation with the class label  $\overline{\rho_{cf}}$  maximizes the value of  $M_S$ . All these computations rely on the assumption that no strong feature inter-correlations are present in the data set, but that instead every relevant feature itself is at least weakly correlated with the class label (see also [28]). To meet the conditions of our data set and to be consistent with variant *B*, we use Spearman's correlation coefficient. As the computation of  $M_S$  is not feasible for all possible feature combinations, we use a backward selection strategy that, according to Guyon and Elisseeff [26], typically provides superior results compared to forward selection. When applying it in our research, we try to minimize the possible shortcomings of the *CFS* by applying cross-validation with the

five data folds for training and validation ( $D_{TV}^{Ma}$ ), as described in Sec. III-D.

The feature selection techniques described so far are limited in two aspects: Firstly, a proper incorporation of the properties of the used classification algorithm is missing. Secondly, features only being meaningful in combination with others are not considered in feature sets *B* and *C*. Therefore, when generating feature set *D*, we apply a wrapper feature selection technique as described in [29]. As the training of Random Forests already includes an implicit feature selection, we solely focus on wrapper techniques including the other classifiers presented in Sec. IV-B. The main idea of wrapper techniques is to incorporate the classifier itself as black box into the feature selection process. Within this process the prediction performance on a validation data set is used to determine the best feature set for the respective classifier. We build our investigations on a hyperparameter set that was optimized as described in Sec. IV-C, with the feature set of variant *C* being used for optimization. According to the process for deriving *C*, we perform the search for the most descriptive feature set with backward elimination. As for each of the approximately 5000 possible subsets, a classifier needs to be trained and evaluated, the wrapper technique becomes computationally expensive. To accelerate the computation, we are not performing the validation using cross-validation. Instead, we use one of the data folds constructed in Sec. III-D for training ( $D_1^{Ma}$ ) and one for validation ( $D_2^{Ma}$ ).

### B. Examined Classification Algorithms

For the task of maneuver classification, we consider three different algorithms for evaluation purposes, which have been successfully applied in reference works:

- 1) The first algorithm is based on a Gaussian Naïve Bayes (*GNB*) approach using *GMMs* instead of only using one Gaussian kernel per class and was presented in [7].
- 2) The second algorithm is based on a Random Forest (*RF*) and was presented in [3].
- 3) The third algorithm is based on a Multilayer Perceptron (*MLP*) approach and was presented similarly in [15]. As opposed to *GNB* and *RF*, this approach uses scaled features, as suggested by [30, p. 398 ff]. In contrast to [15], we use a modified labeling and a partly automated strategy to identify an optimal model structure, where we restrict the model to one hidden layer in order to keep the parameter optimization solvable in finite time.

### C. Hyperparameter Optimization

To achieve the best possible performance and to enable a fair comparison of the examined classifiers, we optimize their respective hyperparameters. For the *GNB*, this means to find the optimal number of Gaussian kernels *K* used for each feature and class. A Variational Bayesian Gaussian Mixture Model (*VBGMM*; see [31]) is used in this context. This technique was already successfully applied in [11]. The principle behind *VBGMMs* is to fit a distribution of the possible Gaussian Mixture distributions using a Dirichlet process.



TABLE III  
OPTIMIZED HYPERPARAMETERS PER CLASSIFIER

Classifier	Parameter	Description
MLP	$\alpha$	<i>Step size</i> : Controls how fast the weights of the network are adapted towards the direction of the gradient
	$n_{hidd}$	<i>Hidden neurons</i> : Describes the structure of the network as we are only working with one hidden layer
	$n_{iter}$	<i>Iterations</i> : Maximum number of training cycles
RF	$n_{tree}$	<i>Trees</i> : Number of parallel trees in the forest
	$n_{splt}$	<i>Splits</i> : Maximum number of splits in each tree
	$n_{smpl}$	<i>Samples</i> : Minimum number of samples necessary for a split

Hence, this technique ensures that the optimal value for  $k$  is determined automatically.

Regarding *RF* and *MLP* approaches, the parameter optimization is executed for each feature set using a grid-search. This means, that we vary the parameters and calculate for each parameter set a performance value. For the latter, we calculate the average balanced accuracy (see Sec. V-A) in a leave one out cross-validation manner. Thereby, we use the data of the five data folds for training and validation ( $D_{TV}^{Ma}$ ). The parameters to be optimized are summarized in Tab. III.

So far, we constructed different feature sets (cf. Sec. IV-A) and optimized the hyperparameters for the different classification algorithms (cf. Sec. IV-B & Sec. IV-C). Subsequently, we now execute a second training step with a larger amount of data for all algorithms, using the optimized feature sets and hyperparameters. The enlargement of the data set is achieved using all five folds that we previously used in the cross-validation  $D_{TV}^{Ma}$ . Note that through this step we derive the final models for the classifier evaluation (cf. Sec. V).

## V. MANEUVER CLASSIFIER EVALUATION

This section presents the experimental results obtained with the trained classification models (cf. Sec. IV). Sec. V-A introduces the used performance measures, whereas Sec. V-B presents and discusses the results measured with the constructed test data set (cf. Sec. III-A).

### A. Performance Measures

To be able to assess the performance of the developed classifiers, several metrics are needed, as we are simultaneously focusing on different objectives. Particularly, we are interested in predicting lane changes not only with high accuracies, but also as early as possible in advance of their execution.

To reflect that, we use the balanced accuracy (*BACC*), which enables us to perform an even weighting of the classification performance for the three maneuver classes. Basically, we use the definition presented in [32], but in a generalized form for multiclass problems (cf. Eq. 9):

$$BACC = \frac{1}{|M|} \cdot \sum_{m \in M} \frac{TP_m}{P_m} \quad (9)$$

TABLE IV  
DEFINITION OF THE DETECTION TIME METRICS

Metric	Definition
$\tau_f$	Time between the vehicle center crosses the centerline and the first detection of the correct maneuver class (as presented in [6])
$\tau_c$	Time between the vehicle center crosses the centerline and the moment at which the classifier becomes certain about its decision for a specific maneuver class and does not change it till the end of the situation. Note that this is a by far stricter definition than the one of $\tau_f$

$M$  is defined according to Eq. 2. Moreover,  $TP_m$  corresponds to the number of true positives for class  $m$  and  $P_m$  to the number of samples truly belonging to class  $m$  (positives). Thereby, the classifiers assign each sample to the class with the highest probability value.

Additionally, we use the Receiver Operator Characteristic (*ROC*) and Area Under the *ROC* Curve (*AUC*), which both are widely used metrics in this domain (e.g. [33, p. 180 ff]). As opposed to the *BACC*, the *ROC* curve is originally intended to assess binary classifiers. Accordingly, we transform our three-class problem into three binary classification problems. In contrast to the *BACC*, the *ROC* curves constructed this way enable us to show off the classification performance at different working points (WP). For example, this property allows us to assess the performance for the maneuver classes *LCL* and *LCR* with more conservative classifier parametrizations and, thus, less false positives. Additionally, the *AUC* helps to analyze the performance at all possible working points at once.

Besides, metrics which enable us to analyze the technically possible prediction time horizon are needed. As the point in time being referenced in this context is essential and most sources (e.g. [1], [15] and [16]) are not very exact in this respect, we introduce the two metrics  $\tau_f$  and  $\tau_c$  (cf. Tab. IV).

As opposed to the *BACC* evaluation, for which an unambiguous class assignment becomes necessary, the class assignment is at this point conducted in a way that matches the binary evaluation in the *ROC* curve: For the classes *LCL* and *LCR*, respectively, we select a binary decision threshold that keeps the false positive rate below 1%. The resulting working points are presented later on in Fig. 5 along with the *ROC* curves. The detection times calculated this way reflect an evaluation with a limited false positive rate and, hence, at a similar working point for the different classifiers. Note that this ensures a fair evaluation. We decide here for a very low false positive rate as the system should not produce too many lane change detections. Remember that in practice, lane changes occur very rarely compared to lane following.

### B. Results & Discussion

Tab. V shows the results (*BACC*, *AUC*,  $\tau$ ) for the different classifiers and feature sets measured based on the maneuver test data set  $D_{Te}^{Ma}$ . Probably, due to the large number of samples, a favorable classifier parametrization and selection seem to have a significantly higher impact on the classification performance than a clever feature selection has. Note that






Fig. 5. ROC curves for the developed maneuver classifiers with their respective best parameter sets and hyperparameters.

TABLE V  
SUMMARY OF EXAMINED CLASSIFIERS WITH  
PREFERRED HYPERPARAMETERS

Classifier	Feature Set	Performance on Test Data per Class ( $AUC$ ; $\bar{\tau}_f$ ; $\bar{\tau}_c$ )			
		BACC	LCL	FLW	LCR
GNB	A	0.704	0.924 2.86±1.46 s 1.86±1.40 s	0.815 - -	0.905 2.92±1.42 s 2.16±1.25 s
	B	0.692	0.910 2.82±1.38 s 1.91±1.24 s	0.801 - -	0.895 2.82±1.32 s 2.06±1.09 s
	C	0.651	0.874 2.57±1.31 s 1.85±1.21 s	0.770 - -	0.884 2.73±1.31 s 1.97±1.10 s
	$D_{GNB}$	<b>0.772</b>	<b>0.943</b> <b>3.26±1.28 s</b> <b>2.41±1.29 s</b>	<b>0.864</b> - -	<b>0.929</b> <b>3.11±1.14 s</b> <b>2.61±1.03 s</b>
MLP	A	0.822	0.973 3.67±1.26 s 2.58±1.51 s	0.909 - -	<b>0.961</b> 3.35±1.19 s 2.81±0.99 s
	B	<b>0.831</b>	0.974 3.73±1.07 s 2.72±1.40 s	0.912 - -	0.958 <b>3.60±1.15 s</b> <b>2.86±1.10 s</b>
	C	0.798	0.966 3.46±1.07 s 2.69±1.10 s	0.892 - -	0.953 3.47±1.11 s 2.81±0.89 s
	$D_{MLP}$	<b>0.831</b>	<b>0.976</b> <b>3.79±1.16 s</b> <b>2.72±1.45 s</b>	<b>0.915</b> - -	0.960 3.33±1.18 s 2.68±0.98 s
RF	A	<b>0.838</b>	<b>0.978</b> 3.81±1.14 s 3.11±1.35 s	<b>0.925</b> - -	<b>0.968</b> 3.60±1.19 s <b>3.06±1.17 s</b>
	B	0.834	0.976 3.73±1.13 s 3.11±1.29 s	0.918 - -	0.959 <b>3.61±1.17 s</b> <b>3.06±1.03 s</b>
	C	0.799	0.964 3.45±1.07 s 2.73±1.14 s	0.893 - -	0.953 3.49±1.10 s 2.93±0.92 s

this can be concluded, as the classifiers working with feature sets B and C only perform slightly worse regarding BACC and AUC than the other classifiers. However, applying a feature selection still remains reasonable as it ensures shorter computation times. In addition, the results indicate that the feature selection contributes to an increase of the prediction times in most cases. Note that this does not apply to the RF as this classifier performs an implicit feature selection.

Fig. 5 additionally shows the ROC curves for the respective best combination of classifier and feature set regarding

TABLE VI  
CONTEXTUAL FEATURES SOLELY IMPACTING SPECIAL SITUATIONS

Features	Providing Insights on
fog lamps, wiper, ...	weather conditions
tunnel, bridge, ...	structural characteristics
lane marking color, ...	road works
country, distance to next highway exit/approach, ...	geographic specialties

BACC and AUC for each of the three classifiers. As another result of our investigations, the classification performance for the lane following maneuver (FLW), which is neglected by most researchers in literature, is notably worse than for the lane changing maneuvers for all considered algorithms. This can be explained with the fact that nearly each sample, which can not be certainly assigned to one of the lane change maneuvers, is classified as lane following. This is caused, as confusions between a lane change to the right and one to the left are very rare. Thus, a significantly larger number of false positives arises for maneuver class FLW. In addition, we could reproduce the findings of [8], which showed that lane changes to the left are easier to predict than the ones to the right. One may explain this phenomenon with the observation that lane changes to the right are often motivated by the intention to leave the highway. The latter can be hardly predicted compared to lane changes to the left, which are often performed to overtake slower leading vehicles. Besides, it can be observed that the classification problem remains resolvable even with a significantly decreased number of features, as shown by the MLP classifier with feature set  $D_{MLP}$ , which only includes 24 features. This illustrates that a decreased number of features sometimes leads to an improved performance due to a lower dimension of the input space. This can be explained with the fact that numerous features, which we expected to provide insights into specific lane changing situations, seem to have nearly no effect concerning the general behavior in highway situations. Exemplary features showing this behavior are summarized in Tab. VI.

An explanation of this behavior is that situations, which are affected by these features, occur even rarer than lane changes. However, as automated driving is extremely demanding




Fig. 6. Histogram of detection times  $\tau_f$  (a) and  $\tau_c$  (b) for *RF* for maneuver class *LCL* with feature set *A*.

TABLE VII  
AUC VALUES IN COMPARISON TO REFERENCE WORKS

Approach	AUC			Prediction Horizon
	LCL	FLW	LCR	
[3]	0.863	0.661	0.836	5.0 s
[7]	0.970	-	0.991	2.0 s
[8]	0.947	-	0.942	2.5 s
[18]	0.934	-	0.993	2.0 s
<i>MLP</i>	0.976	0.915	0.960	5.0 s
<i>RF</i>	0.978	0.925	0.968	5.0 s




Fig. 7. Steps to train and evaluate the position predictors.

exactly in these situations, additional investigations are needed in these cases (cf. Sec. VIII).

It is noteworthy that the detection times  $\tau_f$  and  $\tau_c$  are limited to a maximum of 5 s due to our evaluation methodology. Therefore, the average values  $\bar{\tau}_f$  and  $\bar{\tau}_c$  presented in Tab. V will even be exceeded in practice. To substantiate this assumption, Fig. 6 shows a histogram of the detection times for the *RF*. The distribution shows numerous situations, that are detected 5 or more seconds in advance.

Altogether, our investigations show that a systematic machine learning workflow, combined with a large amount of data, is able to outperform current state-of-the-art approaches significantly. This becomes obvious when looking at the *AUC* in comparison to other approaches. Tab. VII shows that our approach outperforms the others, although we are working with a significantly larger prediction horizon, which makes the classification problem more demanding as aforementioned. Finally, note that the mentioned state-of-the-art approaches were designed and evaluated on considerably smaller data sets.

Our investigations show that the *GNB* classifier performs significantly worse than the two other approaches (i.e. *MLP* and *RF*). Thus, we only use these two classifiers in our further studies. Additionally, we are restricting ourselves to those feature sets and hyperparameter sets showing the best performance (cf. Tab. VIII).

## VI. POSITION PREDICTOR TRAINING

This section deals with the training of the models for position prediction. In particular, we show how to determine

TABLE VIII  
SELECTED FEATURE SETS AND HYPERPARAMETERS PER CLASSIFIER

Classifier	Parameter	Value
<i>MLP</i>	Feature Set	$D_{MLP}$
	$\alpha$	0.02
	$n_{hidd}$	27
	$n_{iter}$	800
<i>RF</i>	Feature Set	<i>A</i>
	$n_{tree}$	128
	$n_{split}$	16
	$n_{simpl}$	100




Fig. 8. Illustration of the mixture of experts (*MOE*) approach.

the *GMM* parameters  $\Theta$ . Sec. VI-A relies on the Mixture of Experts (*MOE*) approach, which was introduced in [3] for lateral predictions and which uses Gaussian Mixture Regression (cf. Eq. 1). An alternative approach is presented in Sec. VI-B. As opposed to the *MOE* approach, it solves the problem in one processing step (cf. Eq. 5). The entire procedure, including the evaluation process (cf. Sec. VII), is depicted in Fig. 7.

### A. Mixture of Experts Approach

To train the experts for the three maneuver classes, we divide the data set (cf. Sec. III-D) along the maneuver labels (cf. Fig. 7). Subsequently, we perform a random under-sampling of the data points for the *FLW* maneuver class to obtain approximately the same number of samples as for the other two classes. The basic idea behind this step is that the regression problem for the *FLW* class is less complex than for the two other classes. Thus, it should be solvable with the same amount of data. Amongst others, this data reduction helps to speed up training. As a consequence, the number of *FLW* samples is approximately decreased by 95% and the data sets  $D_{T,LCL}^{Po}$ ,  $D_{T,FLW}^{Po}$ , and  $D_{T,LCR}^{Po}$  are constructed (cf. Tab. I). Afterwards, we train an expert *GMM* with each of these data sets. These experts are later used in the *MOE* approach (cf. Fig. 8). We choose a maximum number of  $K = 50$  mixture components as well as full covariance matrices,<sup>4</sup> and fit the *GMM* in a variational manner again. Besides, we use the following input-feature set  $F_y^I$  and the true position  $y$  at a defined prediction time  $t$  to train the experts in lateral direction (cf. Eq. 10):

$$F_y^I = \{v_y, d_y^{cl}\} \quad (10)$$

Regarding the prediction in longitudinal direction, we need to distinguish whether or not a preceding vehicle is present.

<sup>4</sup>Preliminary investigations showed that *GMMs* with diagonal covariance matrices are faster to fit, but are by far less accurate.




Fig. 9. Illustration of the integrated approach.

If no vehicle is in sensor range, both the relative speed and distance for that vehicle are set to default values. As involving the latter in the training of the models would lead to bad fits, the input feature sets  $F_{x,Obj}^I$  and  $F_{x,Obj}^I$  are defined as follows (cf. Eq. 11 & Eq. 12):

$$F_{x,Obj}^I = \{v_x, a_x, d_v^{rel,f}, v_v^{rel,f}\} \quad (11)$$

$$F_{x,Obj}^I = \{v_x, a_x\} \quad (12)$$


As shown in [13], the prediction performance for the longitudinal direction can be significantly increased by learning the deviation from the constant velocity prediction  $\hat{x}_{CV}$  instead of the true target position  $x$ . Consequently, we use the output dimensions  $F_x^O$  (cf. Eq. 13):

$$F_x^O = \{x - \hat{x}_{CV}, t\} \quad (13)$$

### B. Integrated Approach

As alternative to the *MOE* approach, this section presents an integrated approach, which uses the unsplitted data set  $D_{Te}^{Po}$  (cf. Tab. I) and expands the feature sets ( $F_{x,Obj}^I, F_{x,Obj}^I, F_y^I$ ) with the maneuver probabilities  $P_{LCL}$  and  $P_{LCR}$  (cf. Fig. 9).  $P_{FLW}$  is left out here as this information would be redundant to the one provided by  $P_{LCL}$  and  $P_{LCR}$ , and we want to keep the models' dimension as low as possible. Consequently, the task of considering the maneuver probabilities is directly integrated in the model. The resulting one-block solution is both easier to implement and to use. In this context, we discovered that *GMMs* are not well suited to fit probabilities bounded to values between 0 and 1. Especially, this is the case if most of the probabilities tend against the extreme values (cf. Fig. 10 (a)). Hence, we expand our data set with a duplicate of each data point containing probability values, which are mirrored at 0 for original probabilities being lower than 0.5 and at 1 for all other original probabilities. This way, we are able to generate the density shown in Fig. 10 (b), which we identified as easier to fit with *GMMs*. Note that before our adjustment, the density contained an abrupt jump, especially at  $P_{LCL} = 0$ . As such discontinuities are only representable by numerous Gaussian components, which are symmetrical and smooth per definition, many components needed in other areas of the data space would be wasted for this purpose.

The actual training of the integrated *GMM* is performed similarly to the experts training in a variational fashion, with  $K = 50$  components and full covariance matrices, but with the entire training data set. Thus, no undersampling procedures are applied and the unbalanced nature of the maneuver classes and their actual frequencies are preserved.

Fig. 10. Density of  $P_{LCL}$  before (a) and after (b) adjustment.

## VII. POSITION ESTIMATION EVALUATION

In order to evaluate the position predictions, first of all, one has to decide which of the considered classifiers fits best as gating network in the Mixture of Experts (*MOE*) and in the integrated approach respectively. Hence, we calculate the average log-likelihoods  $\bar{\mathcal{L}}$  on the entire position test data set  $D_{Te}^{Po}$  (cf. Sec. III-D). Note that this data set is not balanced according to the maneuver labels as also suggested in [20]. In particular, the unbalanced nature of the data allows us to draw general conclusions about the performance, independent of the respective driving maneuver. In this context, the use of the average log-likelihood as quality criterion for comparing different approaches is beneficial, as it rates the quality of the predicted probability density distribution instead of assessing only the ability to predict one single position with maximized accuracy. Moreover, the log-likelihood is exactly the value to be maximized in the process of fitting a *GMM*. However, as  $\bar{\mathcal{L}}$  can not be interpreted as physical quantity, it is solely useful for comparison purposes. As we are also interested in assessing the performance concerning the spatial error and to achieve comparability, we additionally investigate this quantity for the approach working best in the following subsections.

Tab. IX shows the per sample log-likelihood of different approaches for the longitudinal ( $\bar{\mathcal{L}}_x$ ) as well as the lateral ( $\bar{\mathcal{L}}_y$ ) direction. In this context, we use the already introduced classifiers *RF* and *MLP* in combination with four different strategies to combine the experts' position estimates, as introduced in Eq. 1, as weighting function  $w_m(I)$ :

- 1) Raw probabilities (Raw): This strategy directly uses the raw probabilities as issued by the classifiers  $P_m^{clf}(I)$  as gating probabilities. This means that we concatenate the three *GMMs* and multiply the mixture weights with the probabilities issued by the respective classifier:  $w_m^{Raw}(I) = P_m^{clf}(I)$ .
- 2) Winner Takes it All (WTA): This strategy uses the outputs of the *GMM* for the maneuver class with the largest probability according to the respective classifier (cf. Eq. 14).

$$w_m^{WTA}(I) = \begin{cases} 1, & \text{if } P_m^{clf}(I) = \max_{\{q \in M\}} P_q^{clf}(I) \\ 0, & \text{else} \end{cases} \quad (14)$$

- 3) Prior Weighted Raw probabilities (PW-Raw): This strategy considers that the classifiers were trained on a balanced data set. Thus, it multiplies the raw probabilities with the prior probabilities for each maneuver class:  $w_m^{PWRaw}(I) = \text{norm}(P_m^{clf}(I) \cdot \pi_m)$ .

TABLE IX  
PER SAMPLE LOG-LIKELIHOODS WITH DIFFERENT  
CLASSIFIERS AND *MOE* STRATEGIES

Classifier	Strategy	$\bar{\mathcal{L}}_x$ (normalized [%])	$\bar{\mathcal{L}}_y$ (normalized [%])
Labels	-	-14.066 (100)	-7.547 (100)
Priors	-	-13.273 (106.0)	-7.769 (97.1)
NOCLF	-	<b>-13.171 (106.8)</b>	-7.762 (97.2)
<i>MLP</i>	Raw	-13.667 (102.9)	-7.900 (95.5)
	WTA	-16.279 (86.4)	-8.793 (85.8)
	PW-Raw	<b>-13.329 (105.5)</b>	<b>-7.608 (99.2)</b>
	I-GMM	-13.354 (105.3)	-7.691 (98.1)
<i>RF</i>	Raw	-13.568 (103.7)	-7.781 (95.9)
	WTA	-15.685 (89.7)	-8.369 (90.2)
	PW-Raw	-13.280 (105.9)	-7.626 (99.0)
	I-GMM	<b>-13.207 (106.5)</b>	<b>-7.611 (99.2)</b>

- 4) Integrated *GMM* (I-GMM): This strategy directly uses the integrated approach presented in Sec. VI-B to predict the probability distributions and follows Eq. 5.

To demonstrate the benefits of our approach, which combines maneuver classification and position prediction, we additionally analyze its performance compared to reference strategies. First, we use the labels as a perfect classifier according to Eq. 15:

$$w_m^{\text{Labels}} = \begin{cases} 1, & \text{if } m = L \\ 0, & \text{else} \end{cases} \quad (15)$$

Moreover, we use the pure prior probabilities ( $\pi_{LCL} = \pi_{LCR} = 0.03$ ;  $\pi_{FLW} = 0.94$ ) as most naive classifier ( $w_m^{\text{Priors}} = \pi_m$ ) and a strategy without a classifier, referred to as NOCLF in the following.

For the longitudinal direction, Tab. IX shows that the reference solution without any previous maneuver classification (*NOCLF*) is able to produce slightly better results than the other combinations. Although it seems to be trivial that lane changes have not to be taken into account when predicting the longitudinal behavior, this is noteworthy, as our expectations beforehand was that lane changes to the left mostly go along with an acceleration, whereas braking actions are extremely rare.

By contrast, the benefits of the Mixture of Experts (*MOE*) approach come into effect for the lateral direction. As shown in Tab. IX, the combination of prior weighting and *MLP* probabilities performs best. Furthermore, all combinations involving the integrated approach perform only slightly worse or even better (*RF*) than the combinations using prior weighted probabilities. As benefit, these models are easier to use and are more robust against poor or uncalibrated maneuver probabilities without needing an additional calibration step. This can be explained with the fact that these models perform an implicit probability calibration during the training of the *GMM*.

Moreover, we learned that the WTA strategy has no practical relevance, as it does not necessarily produce continuous position predictions over consecutive time steps as accomplished by the other strategies per definition. Besides, in case of a misclassification, the WTA strategy solely asks one specific expert model, which might not be applicable in that area of the data space, what clearly decreases the overall performance.

In the following, we investigate the spatial errors of the best combinations (lateral: *MLP* classifier with PW-Raw strategy; longitudinal: *NOCLF*), as previously introduced. For this purpose, we present the applied performance measures in Sec. VII-A and then show the obtained results in Sec. VII-B.

#### A. Performance Measures

To measure the spatial performance of our predictions, we rely on the unbalanced position evaluation data set  $D_{Te}^{Po}$ . The latter contains the needed inputs for the maneuver classifiers and position predictors ( $I$ ) as well as the true trajectories  $TR$  according to Eq. 16.

$$D_{Te}^{Po} = [I \ TR] \quad (16)$$

$TR$  contains  $N = 20\,000$  5 s-trajectories sampled with 10 Hz (hence 1 000 000 samples) according to Eq. 17:

$$TR = [tr^0 \ tr^1 \ \dots \ tr^N] \quad (17)$$

Each trajectory  $tr^i$  consists of 51 corresponding  $x$  and  $y$  positions, according to Eq. 18:

$$tr^i = \begin{bmatrix} x_{0,0}^i & y_{0,0}^i \\ x_{0,1}^i & y_{0,1}^i \\ \vdots & \vdots \\ x_{5,0}^i & y_{5,0}^i \end{bmatrix} \quad (18)$$

The predicted trajectories  $\hat{TR}$  are then calculated with the described classifiers and position predictors in the same format as  $TR$ . However, as the Gaussian Mixture Regression originally produces probability densities instead of point estimates, these have to be calculated first. This is accomplished by calculating the center of gravity of the density as described in [3]. Accordingly, the prediction error  $e_t^i$  of a specific prediction time  $t$  for one of the  $i$  trajectories is calculated separately for the two dimensions  $x$  and  $y$  as follows (Eq. 19):

$$e_t^i = [e_{x,t}^i \ e_{y,t}^i] = [|x_t^i - \hat{x}_t^i| \ |y_t^i - \hat{y}_t^i|] \quad (19)$$

Variables  $\hat{x}$  and  $\hat{y}$  describe the estimated positions, whereas  $x$  and  $y$  correspond to the actual ones. The individual errors  $e_t^i$  of all trajectories  $i$  are concatenated to  $E_t$  (cf. Eq. 20):

$$E_t = [E_{x,t} \ E_{y,t}] = [e_{x,t}^i \ e_{y,t}^i]_{\forall i} \quad (20)$$

At this point, we want to re-emphasize, that although this way of evaluating the performance produces easy to interpret results, it disregards that our original outputs (i.e. spatial probability densities) contain much more information than a single point estimation.

#### B. Results & Discussion

Fig. 11 shows the performance of the selected combinations of classifiers and mixing strategies (highlighted in Tab. IX) at a prediction horizon of 5 s for the longitudinal ( $E_{x,5}$ ) and the lateral ( $E_{y,5}$ ) direction on the left side. In comparison, a constant velocity (*CV*) prediction and a Mixture of Experts (*MOE*) with labels<sup>5</sup> are shown. The right-hand side of Fig. 11

<sup>5</sup>Using the *MOE* with the labels as input corresponds to the assumption of a perfect classifier.






Fig. 11. Visualization of the error distribution (left) in longitudinal and lateral direction and the median lateral error as function of the prediction time (right).

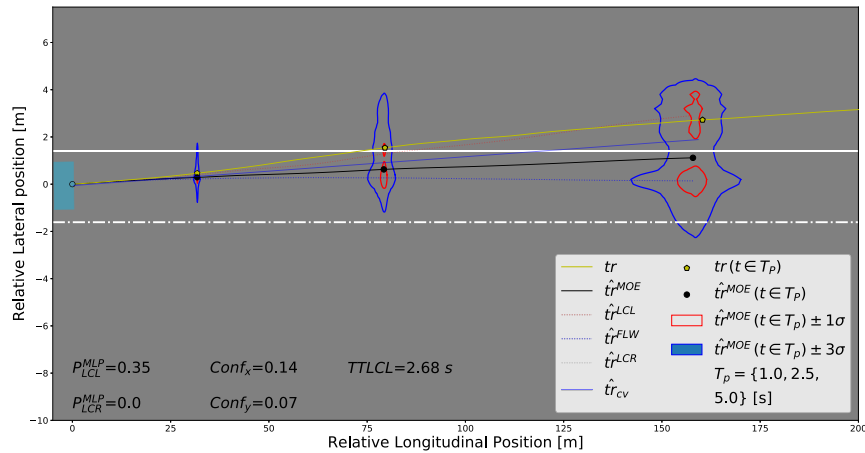


Fig. 12. Predicted probability distribution of future vehicle positions for an illustrative situation.

 TABLE X  
 COMPARING LATERAL PREDICTION PERFORMANCE  
 WITH RELATED WORKS

Approach	$\tilde{E}_{y,t}$ [m]	Prediction Horizon $t$ [s]
[3]	$\approx 0.18$	5.0
[15]	0.23	5.0
[18]	0.50	3.0
<i>MLP</i> (PW-Raw)	0.204	5.0

shows the development of the median lateral error  $\tilde{E}_{y,t}$  as function of the prediction time  $t$ .

As the plots indicate, our position prediction system is able to produce results comparable to the ones with a perfect maneuver classification, in both lateral and longitudinal direction. Additionally, the plots show that we are able to clearly outperform simple models as *CV* and reach a very small median lateral prediction error of less than 0.21m at a prediction horizon of 5s. As shown in Tab. X, this is remarkable compared to other approaches. Note that we did not include studies in this compilation, which report the root-mean-square error (*RMSE*), which we quantify with a value of 0.64m. On one hand, we follow [34], which points out that *RMSE* measures do not allow for a comparison over different data sets, as the values depend on the size of the data set. On the other, the challenge tackled by us (cf. Sec. I-A) is to predict the probability distribution of future vehicle positions

 TABLE XI  
 PREDICTION ERRORS PER CLASS AND DIRECTION

$d$	$\tilde{E}_{d,5}^{LCL}$ [m]	$\tilde{E}_{d,5}^{FLW}$ [m]	$\tilde{E}_{d,5}^{LCR}$ [m]
$x$	3.22	1.67	2.20
$y$	1.25	0.19	1.80




Fig. 13. Prediction confidence against lateral prediction errors.

rather than single shot estimates. Consequently, we did not optimize the predictions to minimize *RMSE*. Therefore, it is not surprising that other works which explicitly minimize this value, but ignore distribution estimations, perform better with respect to *RMSE*.

As shown in [3], these results are dominated by the most frequent maneuver class (*FLW*). Hence, Tab. XI complementarily shows the errors for 20000 maneuvers of each type.

TABLE XII  
DESCRIPTION OF THE EVALUATED FEATURES  $f$  OF AN OBSERVED VEHICLE  $o$  AND USAGE  
OF THE FEATURES IN THE CONSTRUCTED FEATURE SETS ( $A - D$ )

$R$	$f$	Description	Unit (Continuous) Range of Values (Nominal)	Element of				
				$B$ $ \cdot  = 40$	$C$ $ \cdot  = 29$	$D_{MLP}$ $ \cdot  = 24$	$D_{GNB}$ $ \cdot  = 48$	
$R_{rf}$	general information describing the related vehicle $r$ (cf. Eq. 6 & Fig. 3)							
	$actv^r$	activity status	{0: inactive, 1: active}	{f, l} <sup>6</sup>	{f, fl, fr, l}	{fr, r}	{fl, fr, l, r}	
	$mov^r$	movement status	{0: standing, 1: moving}	{f, l}	{f, fl, fr, l}	{r}	{fl, fr, r}	
	$class^r$	object class	{0: bicycle, 1: motorbike, 2: car, ..., 14: no class}	{f, l}	{f, fl, l}		{fl, fr, r}	
	$cutinlvl^r$	cut-in level	{0: $P \leq 0.5$ , 1: $P > 0.5$ , 2: $P > 0.66$ , 3: $P > 0.9$ }	{l}			{r}	
	relation between $o$ and related vehicle $r$ in $o$ 's coordinate-system							
	$d_x^{rel,r}$	longitudinal distance	$m$	{f, l}	{f, l}	{f, r}	{fr, r}	
	$d_y^{rel,r}$	lateral distance		{f, l}	{f, fr}	{r}	{fr, r}	
	$v_x^{rel,r}$	relative longitudinal speed	$m/s$	{f, r}		{r}	{f, fr, r}	
	$v_y^{rel,r}$	relative lateral speed		{f, fl, l, r}	{f}	{f, fr, r}	{f, fr, r}	
	$a_x^{rel,r}$	relative longitudinal acceleration	$m/s^2$			{fr}	{f, fl, fr, r}	
	relation between $o$ and related vehicle $r$ in curvilinear coordinates							
	$d_w^{rel,r}$	longitudinal distance	$m$	{f, l}	{l}	{fr}	{fl, fr, r}	
	$d_u^{rel,r}$	lateral distance		{f, l}	{fr}	{r}	{fr, r}	
	$v_w^{rel,r}$	relative longitudinal speed	$m/s$	{f, r}		{f}	{f, fl, fr, r}	
	$v_u^{rel,r}$	relative lateral speed		{f, fl, l, r}		{l}	{fr, r}	
	$R_{rb}$	$mov^r$	movement status	{0: standing, 1: moving}	{rr}	{rr}	{rr}	{rr}
		$d_y^{rel,r}$	lateral distance	$m$	{rl}		{rl}	{rl, rr}
$F_o$	$fog^f$	status of the front fog lamp	{0: off, 1: on}					
	$fog^r$	status of the rear fog lamp						
	$fog^{rl}$	status of the rear left fog lamp						
	$fog^{rr}$	status of the rear right fog lamp						
	$wpr$	wiper level	{0, ..., 15}					
	$d_y^{ml}$	distance between the center of $o$ and the left marking	$m$	✓	✓	✓	✓	
	$d_y^{mr}$	distance between the center of $o$ and the right marking		✓	✓			
	$d_y^{cl}$	distance between the center of $o$ and the centerline of the assigned lane		✓	✓			
	$v_x$	longitudinal speed of the observed vehicle	$m/s$	✓	✓		✓	
	$v_y$	lateral speed of the observed vehicle		✓	✓		✓	
	$a_x$	longitudinal acceleration of the observed vehicle	$m/s^2$	✓	✓	✓	✓	
	$a_y$	lateral acceleration of the observed vehicle		✓	✓	✓	✓	
	$\psi$	angle of the observed vehicle relative to the direction of the lane	$^\circ$	✓	✓	✓	✓	
$F_{infra}$	$t^{ml}$	type of the left marking	{0: no marking, 1: continuous, 2: broken}	✓	✓	✓	✓	
	$t^{mr}$	type of the right marking		✓	✓	✓	✓	
	$c^{ml}$	color of the left marking	{0: no marking, 1: white, 2: yellow}				✓	
	$c^{mr}$	color of the right marking					✓	
	$nlanes_{cam}$	number of parallel lanes observed via the camera	{0: 0, ..., 3: 3+}				✓	
	$nlanes_{map}$	number of lanes stored in the map	{0, ..., 5}					
	$cntr$	country	{0: GER, 1: US, ...}					
	$tnl$	indicator if the situation takes place in a tunnel	{0: False, 1: True}				✓	
	$brd$	indicator if the situation takes place on a bridge						
	$v^{lim}$	speedlimit of the current highway section	{1: $> 130 \frac{km}{h}$ , ..., 8: $< 11 \frac{km}{h}$ }					
	$t^a$	type of next approach to the highway	{0: unknown, 1: on ramp, 2: highway merge}					
	$t^e$	type of next exit of the highway	{0: unknown, 1: ramp, 2: highway divider}					
	$w^{ml}$	width of the left marking	$m$	✓	✓			
	$w^{mr}$	width of the right marking		✓	✓			
	$w^{lane}$	width of the lane				✓		
	$d_x^a$	distance to the next approach to the highway						
	$d_x^e$	distance to the next exit of the highway						
	$c_0$	curvature of the road	$1/m$					
	$c_1$	derrivation of the curvature	$1/m^2$					

<sup>6</sup> This means for example, that feature set  $B$  (introduced in Sec. IV-A), contains in total 40 elements, including the activity status of its surrounding vehicles in front (f) and to its left (l). In contrast to that, the activity states of its other front relation partners (r, fl, fr) are not included in  $B$ .

As can be seen, the errors for the lane change maneuvers are considerably larger than the ones for lane-following. On one hand, this can be explained with the more complex regression task. On the other, the predictions are subjected to higher uncertainties in case of a lane change, as shown by the predicted distributions (cf. Fig. 12). As opposed to that, the uncertainty is ignored in the single point estimates. Note that the increased uncertainties are caused by the lack of knowledge on the exact point in time at which the maneuver will be completed. This even holds true, if the classifier made the position prediction to know about an upcoming lane change.

Complementary to these quantitative evaluations, we performed qualitative testing and visualized single situations along with our predictions. To illustrate this, we attached a short video and present a single frame in Fig. 12. More precisely, Fig. 12 shows the predictions during an upcoming lane change, along with the described uncertainties. In addition, we show the confidence of our predictions ( $Conf_x$ ,  $Conf_y$ ), which provides an important hint concerning the reliability of the predictions to the consumer of the information. This value is calculated similarly to [13] through additional  $GMMs$  fitted in the input dimensions. To demonstrate its general usability, we visualized the confidence value divided by the standard

deviation against the lateral prediction errors at  $T_h = 5$  s in Fig. 13. As can be seen, and as expected the prediction errors decrease with increasing confidence values.

## VIII. SUMMARY AND OUTLOOK

This work introduces a machine learning workflow that enables calculations of long-term behavior predictions for surrounding vehicles in highway scenarios. For the first time, a combined compilation of prediction techniques for driving maneuvers and positions as well as lateral and longitudinal behavior is presented. The developed modules are evaluated in detail based on a large amount of real-world data, challenging established state-of-the-art approaches.

To further improve the quality of the presented behavior predictions, especially in complex situations, we are working on various enhancements and conducting additional studies. Currently, we migrate the prediction strategies to an experimental vehicle to enable detailed investigations regarding run time as well as resource usage. Meanwhile, we are about to apply our models to predict movements of surrounding vehicles in contrast to ego-vehicle movements. Besides, we plan to apply our predictor to a publicly available data set as highD [35] or NGSIM to improve comparability. In addition, we want to investigate up to which maximum prediction horizon (beyond 5 s), the maneuver detection produces useful insights.

Moreover, we see high potential in identifying demanding scenarios and explicitly integrating contextual knowledge (e. g. weather, traffic, time of day or local specialties) into our models. First experiments towards this direction have proven, that contextual properties can have a considerable impact on driving behavior.

## ACKNOWLEDGMENT

The authors would like to thank Mercedes-Benz AG Research and Development for providing real-world measurement data, which enabled us to perform our experiments. Furthermore, they would like to thank the Institute of Databases and Information Systems at Ulm University as well as Prof. Dr. Klaus-Dieter Kuhnert from the Institute of Realtime Learning Systems at the University of Siegen for supporting our studies.

## REFERENCES

- [1] G. Weidl, A. L. Madsen, S. Wang, D. Kasper, and M. Karlsen, "Early and accurate recognition of highway traffic maneuvers considering real world application: A novel framework using Bayesian networks," *IEEE Intell. Transp. Syst. Mag.*, vol. 10, no. 3, pp. 146–158, Jun. 2018.
- [2] 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.
- [3] J. Schlechtriemen, F. Wirthmueller, A. Wedel, G. Breuel, and K.-D. Kuhnert, "When will it change the lane? A probabilistic regression approach for rarely occurring events," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2015, pp. 1373–1379.
- [4] X. Meng *et al.*, "MLlib: Machine learning in apache spark," *J. Mach. Learn. Res.*, vol. 17, no. 1, pp. 1235–1241, 2016.
- [5] F. Pedregosa *et al.*, "Scikit-learn: Machine learning in python," *J. Mach. Learn. Res.*, vol. 12, pp. 2825–2830, Oct. 2011.
- [6] C. Wissing, T. Nattermann, K.-H. Glander, C. Hass, and T. Bertram, "Lane change prediction by combining movement and situation based probabilities," *IFAC-PapersOnLine*, vol. 50, no. 1, pp. 3554–3559, Jul. 2017.
- [7] J. Schlechtriemen, A. Wedel, J. Hillenbrand, G. Breuel, and K.-D. Kuhnert, "A lane change detection approach using feature ranking with maximized predictive power," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2014, pp. 108–114.
- [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, no. 6, pp. 1538–1550, Jun. 2016.
- [9] D. Lenz, F. Diehl, M. T. Le, and A. Knoll, "Deep neural networks for Markovian interactive scene prediction in highway scenarios," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2017, pp. 685–692.
- [10] F. Althe and A. de La Fortelle, "An LSTM network for highway trajectory prediction," in *Proc. IEEE 20th Int. Conf. Intell. Transp. Syst. (ITSC)*, Oct. 2017, pp. 353–359.
- [11] J. Wiest, M. Höffken, U. Kresel, and K. Dietmayer, "Probabilistic trajectory prediction with Gaussian mixture models," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2012, pp. 141–146.
- [12] J. Wiest, F. Kunz, U. Kreßel, and K. Dietmayer, "Incorporating categorical information for enhanced probabilistic trajectory prediction," in *Proc. 12th Int. Conf. Mach. Learn. Appl.*, vol. 1, Dec. 2013, pp. 402–407.
- [13] J. Schlechtriemen, A. Wedel, G. Breuel, and K.-D. Kuhnert, "A probabilistic long term prediction approach for highway scenarios," in *Proc. 17th Int. IEEE Conf. Intell. Transp. Syst. (ITSC)*, Oct. 2014, pp. 732–738.
- [14] J. Colyar and J. Halkias, "US highway 101 dataset," Federal Highway Admin., Washington, DC, USA, Tech. Rep. FHWA-HRT-07-030, 2007.
- [15] S. Yoon and D. Kum, "The multilayer perceptron approach to lateral motion prediction of surrounding vehicles for autonomous vehicles," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2016, pp. 1307–1312.
- [16] H. Woo *et al.*, "Lane-change detection based on vehicle-trajectory prediction," *IEEE Robot. Autom. Lett.*, vol. 2, no. 2, pp. 1109–1116, Apr. 2017.
- [17] C. Wissing, T. Nattermann, K.-H. Glander, and T. Bertram, "Probabilistic time-to-lane-change prediction on highways," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2017, pp. 1452–1457.
- [18] C. Wissing, T. Nattermann, K.-H. Glander, and T. Bertram, "Trajectory prediction for safety critical maneuvers in automated highway driving," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 131–136.
- [19] N. Deo, A. Rangesh, and M. M. Trivedi, "How would surround vehicles move? A unified framework for maneuver classification and motion prediction," *IEEE Trans. Intell. Vehicles*, vol. 3, no. 2, pp. 129–140, Jun. 2018.
- [20] N. Deo and M. M. Trivedi, "Convolutional social pooling for vehicle trajectory prediction," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. Workshops (CVPRW)*, Jun. 2018, pp. 1468–1476.
- [21] S. Klingelschmitt, M. Platho, H.-M. Groß, V. Willert, and J. Eggert, "Combining behavior and situation information for reliably estimating multiple intentions," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2014, pp. 388–393.
- [22] J. Schlechtriemen, K. P. Wabersich, and K.-D. Kuhnert, "Wiggling through complex traffic: Planning trajectories constrained by predictions," in *Proc. IEEE Intell. Vehicles Symp. (IV)*, Jun. 2016, pp. 1293–1300.
- [23] D. Sadigh, S. Sastry, S. A. Seshia, and A. D. Dragan, "Planning for autonomous cars that leverage effects on human actions," in *Proc. Robot., Sci. Syst.*, vol. 2. Ann Arbor, MI, USA: RSS Foundation, 2016.
- [24] S. Tattersall, U. Petersen, and J. Breuer, "Ein Messdatenmanagementsystem für die Feldabsicherung von neuen Fahrerassistenzsystemen," in *Proc. VDI-Berichte*, no. 2166, 2012, pp. 203–214.
- [25] A. Thorvaldsson and V. Bandi, "Reference path estimation for lateral vehicle control," M.S. thesis, Dept. Signals Syst., Chalmers Univ. Technol., Gothenburg, Sweden, 2015.
- [26] I. Guyon and A. Elisseeff, "An introduction to variable and feature selection," *J. Mach. Learn. Res.*, vol. 3, pp. 1157–1182, Mar. 2003.
- [27] L. Fahrmeir, C. Heumann, R. Künstler, I. Pigeot, and G. Tutz, *Statistik: Der Weg zur Datenanalyse*. Berlin, Germany: Springer, 2016.
- [28] M. A. Hall, "Correlation-based feature selection for discrete and numeric class machine learning," in *Proc. Int. Conf. Mach. Learn. (ICML)*. San Mateo, CA, USA: Morgan Kaufmann, 2000, pp. 359–366.
- [29] R. Kohavi and G. H. John, "Wrappers for feature subset selection," *Artif. Intell.*, vol. 97, nos. 1–2, pp. 273–324, Dec. 1997.
- [30] T. Hastie, J. Friedman, and R. Tibshirani, *The Elements of Statistical Learning*, vol. 1. New York, NY, USA: Springer, 2001.

- [31] A. Corduneanu and C. M. Bishop, "Variational Bayesian model selection for mixture distributions," in *Artificial Intelligence and Statistics*. Waltham, MA, USA: Morgan Kaufmann, 2001, pp. 27–34.
- [32] K. H. Brodersen, C. S. Ong, K. E. Stephan, and J. M. Buhmann, "The balanced accuracy and its posterior distribution," in *Proc. 20th Int. Conf. Pattern Recognit.*, Aug. 2010, pp. 3121–3124.
- [33] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. Cambridge, MA, USA: MIT Press, 2012.
- [34] C. Willmott and K. Matsuura, "Advantages of the mean absolute error (MAE) over the root mean square error (RMSE) in assessing average model performance," *Climate Res.*, vol. 30, pp. 79–82, 2005.
- [35] R. Krajewski, J. Bock, L. Kloeker, and L. Eckstein, "The highD dataset: A drone dataset of naturalistic vehicle trajectories on German highways for validation of highly automated driving systems," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 2118–2125.



**Florian Wirthmüller** received the B.Sc. and M.Sc. degrees in computer engineering with a study focus on cognitive technical systems from the Ilmenau University of Technology in 2015 and 2017, respectively. He is currently pursuing the Ph.D. degree with the Institute of Databases and Information Systems (DBIS), Ulm University, in cooperation with Mercedes-Benz AG Research and Development. His research interests include automated driving, big data analytics, machine learning, and backend architectures supporting manually driven as well as automated vehicles.



**Julian Schleichriemen** received the Diploma degree in applied computer science with electrical engineering as the main subject from the University of Siegen in 2012, where he is currently pursuing the Ph.D. degree with the Institute of Realtime Learning Systems, in cooperation with Mercedes-Benz AG Research and Development. His research interests include vehicle and driver prediction using machine learning techniques and the incorporation of this information in behavior and trajectory planning.



**Jochen Hipp** received the Diploma degree in computer science and economics and the Ph.D. degree from the University of Tübingen. Since then, deriving knowledge from massive data sets is part of his daily work at Mercedes-Benz AG Research and Development. Over the years, he has been active in different fields such as root cause analysis and early warning based on aftersales data, target-oriented endurance testing, customer profiles, advanced driver assistance systems, autonomous driving with a focus on high definition maps, vehicle localization, and currently working on the analysis of field data to improve current and future driver assistance system generations.



**Manfred Reichert** is a Full Professor with Ulm University, where he is the Director of the Institute of Databases and Information Systems (DBIS). His research interests include business process management, intelligent information systems, process and data mining, and mobile services. Furthermore, he served as the General Chair of the BPM 2009 and EDOC 2014 conferences as well as the BPM 2015 workshops. He was the PC Co-Chair of the BPM 2008, CoopIS 2011, and EDOC 2013 conferences. He has coauthored a Springer book on process flexibility and obtained the BPM Test of Time Award from the BPM Conference in 2013.