# Maneuver-based Trajectory Prediction for Self-driving Cars Using Spatio-temporal Convolutional Networks

Benedikt Mersch<sup>1,\*</sup>

Thomas Höllen<sup>1,\*</sup>

Kun Zhao<sup>2</sup>

Cyrill Stachniss<sup>1</sup>

Ribana Roscher<sup>1</sup>

Abstract-The ability to predict the future movements of other vehicles is a subconscious and effortless skill for humans and key to safe autonomous driving. Therefore, trajectory prediction for autonomous cars has gained a lot of attention in recent years. It is, however, still a hard task to achieve humanlevel performance. Interdependencies between vehicle behaviors and the multimodal nature of future intentions in a dynamic and complex driving environment render trajectory prediction a challenging problem. In this work, we propose a new, datadriven approach for predicting the motion of vehicles in a road environment. The model allows for inferring future intentions from the past interaction among vehicles in highway driving scenarios. Using our neighborhood-based data representation, the proposed system jointly exploits correlations in the spatial and temporal domain using convolutional neural networks. Our system considers multiple possible maneuver intentions and their corresponding motion and predicts the trajectory for five seconds into the future. We implemented our approach and evaluated it on two highway datasets taken in different countries and are able to achieve a competitive prediction performance.

#### I. INTRODUCTION

Being able to predict the future movement of other traffic participants is important when driving on roads. Experienced human drivers subconsciously anticipate possible maneuvers of other road users and react accordingly in advance. Especially with varying numbers and types of traffic participants in dynamic and complex road environments, predicting the intentions of other drivers is a challenging task and humans refine this skill over time while typically outperforming technical systems.

Intelligent vehicles require a reliable perception of the driving environment over time and space to make mediumor even long-term predictions. This is not only the case for autonomous driving, but also for advanced driver assistance systems. For example, adaptive cruise control or traffic jam assistants need to predict the behaviors of other road users in advance to help the system react on time. The longer the prediction horizon, the more time the system has to provide a safe and at the same time convenient user experience.

In contrast to pedestrians and bicycles, cars are more restricted in their motion due to larger inertia, traffic rules, and road geometries. This makes vehicle movements better predictable, especially in structured driving scenarios

<sup>2</sup> Aptiv Services Deutschland GmbH, Wuppertal, Germany



Fig. 1: The goal is to predict the future maneuver of the black target vehicle (T). The different future motions are estimated depending on the neighbors' states like position, velocity, and acceleration. Colors indicate possible maneuvers with corresponding trajectories. Gray arrows denote velocity vectors, their lengths indicate speed.

like highway driving. Despite these amendable properties, challenges arise due to the dynamic interactions among vehicles. Particularly lane change maneuvers require special attention for surrounding objects. Given a so-called target vehicle, for which one wants to predict the future trajectory, neighboring vehicles influence but also restrict the possible maneuvers, such as accelerating or changing a lane. The situation depicted in Fig. 1 illustrates such a scenario. In case the black car (T) is faster than the truck in front (F), it will need to slow down or make a lane change. Possible scenarios are shown as colorized trajectories. Their likelihood strongly depends on the positions, velocities, and accelerations of the neighboring cars F, L, R, RL, and FL. The example illustrates the inherent interdependencies among vehicle behaviors for a time horizon of several seconds. This paper proposes a method to tackle these challenges by paying attention to the other road users and their dynamics and by explicitly estimating the driver's maneuver intention in advance.

The main contributions of this paper towards vehicle motion prediction are two-fold. First, we present a novel semantic neighborhood representation of the scene around a target vehicle for joint aggregation of higher-level features in prediction tasks. The memory-efficient and dense 3D tensor encodes the time, neighbor positions, and past vehicle states as the dynamic context. Second, we propose the use of two 2D convolutional neural networks (CNNs) for joint spatio-temporal feature extraction from the proposed input representation. Our approach explicitly uses convolutions across time and the space of neighboring vehicles. To this end, we classify the future lane change intention of a target vehicle with respect to a lateral motion and then predict a trajectory based on the classified lane change intention. Using a spatio-temporal CNN for sequence prediction leads to a simpler and more compact architecture compared to recurrent network approaches resulting in fewer parameters to train but competitive prediction performance.

<sup>\*</sup> These authors contributed equally.

<sup>&</sup>lt;sup>1</sup> University of Bonn, Germany

This work has partially been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 - 390732324 - PhenoRob and by the European Union's Horizon 2020 research and innovation programme under grant agreement No 101017008 (Harmony).

This yields an approach that is able to (i) successfully classify lane change maneuvers and predict a corresponding trajectory for real-world scenarios by (ii) performing joint spatio-temporal feature aggregation with 2D CNNs, (iii) and outperform state-of-the-art methods. These three main claims are backed up by the paper, our experimental evaluation, and an ablation study.

## II. RELATED WORK

Brown et al. [5] break down the task of estimating and predicting human driver behavior into the tasks of state estimation, intention estimation, trait estimation, and motion prediction. In this work, we focus on predicting the future motion of a vehicle based on the estimated driver intention.

Inferring the future motion of a vehicle from past data has been approached from different perspectives. Closed-loop approaches roll out a control policy in a forward simulation for all target vehicles up to the prediction time, which results in interaction-aware trajectories [34]. Interactionawareness between different traffic participants is deepened with game-theoretic approaches [36], which condition an agent's future motion on the predicted motion of others. Since this dependency increases the computational complexity and can become intractable, other approaches model the future behaviors of vehicles to be independent of each other. Among these, physics-based models solely use kinematic and dynamic properties and apply filter- or sampling-based methods [26]. More advanced independent prediction models are data-based approaches, which are further addressed in this work.

Mozaffari et al. [30] differentiate data-based prediction approaches based on input representation, output type, and prediction method. If only the target's trajectory history is used for prediction, interdependencies between surrounding vehicles can not be considered [38], [32]. In contrast, adding information about neighboring agents [11], [9], stacking different sources of spatial information in Bird's Eye Views (BEV) [8], [18] or using raw sensor data of the target's surroundings makes interaction-awareness between past states possible [6], [27]. Deo and Trivedi [10] use a combination of these approaches and call it a social tensor. This tensor consists of a BEV image of the scene and is augmented by pre-processed temporal features. In this paper, we extend this idea by defining a spatio-temporal representation in advance and then jointly aggregating spatial and temporal features for prediction instead of processing them separately [10].

A common goal of motion prediction is the estimation of maneuver intentions. These methods provide a highlevel understanding of future behavior, usually defined for specific scenarios like intersections [41] or highway lane changes [13]. To predict low-level future motion, an occupancy grid map containing the probabilities of occupancy at future time steps can be used [18]. This output representation also allows for predicting multiple future modes but lacks accuracy for large grid cells resulting in less consistent trajectories. Another possibility is to predict trajectories directly, either in a uni- or multimodal fashion. Unimodal predictions come with a lower computational cost but tend to converge to a mean of different behavior modes [9], [28]. Maneuverbased predictions fix this problem by outputting multiple motion hypotheses incorporating different maneuvers. These models can be formulated in a probabilistic framework, meaning that they model or sample from a multimodal distribution conditioned on the input data [40], [15], [37]. Other approaches estimate intention modes in advance and use them to predict a trajectory [35], [38] based on the maneuver. Casas et al. [6] perform a multi-class classification with eight intention classes using CNNs on a voxelized LiDAR scan and a dynamic map containing road structures and traffic lights. The generated intention scores are then further processed to condition the trajectory estimation. In contrast to these approaches, we define an intention space with three lateral motions and classify a maneuver at each step. Combining these motions at each step results in a larger variety of maneuvers during the prediction.

Different methods have evolved for processing the spatial and temporal dependencies to predict future behavior. Recurrent neural networks (RNNs) can predict time series by maintaining a hidden state while processing temporal information. However, they can be difficult to train for larger time series due to vanishing or exploding gradients caused by their recurrent structure. Therefore, more advanced recurrent models like gated recurrent units (GRUs) [7] or long shortterm memory networks (LSTMs) [17] have been proposed and applied for trajectory prediction [1], [15], [22], [2]. Since only extracting temporal information does not account for dependencies between traffic participants, convolutional neural networks (CNNs) have been widely implemented for spatial information aggregation [18], [6], [19], [8]. Other methods are fully connected [20] or graph neural networks [12]. Generative adversarial methods (GANs) proposed by Goodfellow et al. [14] have also been considered for trajectory prediction [25], [40]. The already mentioned approach by Deo and Trivedi [10] uses a combination of LSTMs for temporal feature extraction for each vehicle trajectory and CNNs for processing the resulting grid representation. The final trajectory is generated by an LSTM decoder for each maneuver. Other combinations of the presented approaches can be found for trajectory prediction [33], [40], [4]. The main difference to our approach is that we solely use a convolutional architecture for jointly encoding temporal and spatial information as well as decoding the resulting features for prediction without the use of recurrent structures.

The use of CNNs instead of RNNs for sequence modeling has been proposed by Bai et al. [3]. They argue that such temporal convolutional networks (TCN) are capable of outperforming recurrent models for specific tasks while being simpler and easier to train. In addition to that, TCNs can process the data in parallel, which makes them faster at inference. Nikhil and Morris [31] propose to use CNNs for pedestrian trajectory prediction and show superior performance and speed compared to LSTM-based models. However, neighboring agents are not considered in their method since they only use 1D convolutions along the time dimension. The approach developed in this work is inspired by temporal convolutions and provides, in contrast to Nikhil and Morris [31], a joint aggregation of spatial and temporal features from the proposed novel tensor input representation.

## III. OUR APPROACH

The key idea of our approach is to first classify the lane change maneuver of a target vehicle for each step into the future and then predict the corresponding trajectory with the classification result as an additional input.

At the current time step t=0, the goal is to predict the sequence of P future positions  $\{(x_t, y_t)\}_{t=1}^P$  from H past states  $S = \{s_{-t}\}_{t=0}^{H-1}$ . Each past state  $s_{-t}$  consists of multiple channels like position, acceleration, and velocity of the target vehicle and its neighbors, which are transformed into a 3D tensor. This representation is addressed in Sec. III-A. The lane change maneuvers are pre-defined and cover sequences of straight driving, left, and right lane changes as explained in Sec. III-B. The classification and regression steps rely on the exploitation of past state sequences with spatio-temporal 2D convolutions, which are presented in Sec. III-C. Based on the CNN features, the classification module predicts a discrete maneuver (straight driving, left lane change, or right lane change) for each prediction step 1, ..., P into the future, see Sec. III-D. The classified sequence is then passed to the regression module as an additional input to predict a trajectory based on the maneuver as outlined in Sec. III-E. Both modules share the same convolution-based architecture. When training on common, publicly available highway datasets, lane changes are usually underrepresented since the vehicles drive straight most of the time. We found that training a single network with two heads and a joint multitask loss for classification and regression is hard to tune due to the imbalance of lane change labels. Therefore, we propose to use partitioned models for each task, which makes it easier to train them separately without shared parameters.

#### A. Neighborhood-based Input Representation

We take the direct neighborhood of the target vehicle into account to predict a trajectory from past data. Close neighbors influence the possible actions a vehicle can take. For example, a leading car with a lower velocity requires deceleration or a lane change of the target vehicle, but the corresponding neighboring lane needs to be free for a safe lane change maneuver. Whereas the short-time prediction of for example the next second is mainly restricted by the inertia of the car, we assume that for longer prediction horizons, including the neighborhood leads to better predictions. This claim is backed up by our experiment in Sec. IV-E. We define a semantic neighborhood consisting of up to seven vehicles around the target inspired by the work of Hu et al. [20]. As mentioned, the vehicles in front (F) and on the left (L) and right (R) play an important role. The remaining neighborhood is further discretized in front left (FL) and right (FR) as well as rear left (RL) and right (RR) vehicles. In total, this results in a maximum number of eight considered vehicles including the target. We illustrate the semantic



Fig. 2: Left: At each time step, we discretize the neighborhood around the target vehicle into seven neighbor positions. Right: For each neighbor, we extract the input channels x and y positions, velocity v, and acceleration a for each time step resulting in 2D tensors and concatenate them to a 3D input tensor.

neighborhood in Fig. 2. For each neighbor vehicle, we acquire an image-like 2D tensor containing the different input channels over time and concatenate them to a 3D tensor. In case a neighbor is not present, we set the entries to zero, which is a state that is not naturally present in the input data due to the choice of the reference location. The concatenation order of neighbors can be chosen randomly but needs to be consistent for training and evaluation. We experienced that it is advantageous to keep neighbor positions close to each other in the data representation. Based on our experiments in Sec. IV-E, we hypothesize that the convolutional network can learn a notion of absolute position and discriminates between neighbors. This is further discussed in Sec. III-C. In contrast to other approaches that encode the whole scene into a grid structure [10], our neighborhood representation is more dense and memory-efficient, since not all vehicles are taken into account.

The proposed multi-channel design makes it easy to add new information channels that can be useful for prediction. We use the information given by the publicly available highway datasets, which are x and y position, velocity, and acceleration. Note that velocity and acceleration can also be extracted from the positional information over time in case they are not directly provided. We normalize the data beforehand to guarantee a consistent input scaling.

We store the neighbors and channels per frame along the third, temporal dimension, see Fig. 2. In order to compare our results with existing methods, we use three seconds of input data and predict the maneuvers and positions for the following five seconds. In this setup with data sampled at 10 Hz, this results in an input tensor of dimension  $4 \times 8 \times 30$  with 4 channels, 8 neighbors, and 30 frames of history.

#### B. Output Parameterization

The number of predicted future positions P depends on the output size of the last dense layer in the classification and regression heads in Fig. 3 and is only limited by the resolution of the dataset used for training. We choose to predict and evaluate the future for the next 5 s represented via 5 positions being 1 s apart in time. This sampling rate is



Fig. 3: Two spatio-temporal convolutional networks aggregate features in parallel from the neighborhood-based 3D input tensor. The tensor from Fig. 2 contains the x and y positions, velocity v and acceleration a for all eight vehicles at the past time steps. Note that the transposed input tensor is depicted for better visualization. The upper module classifies a maneuver for each prediction step and passes the result to the lower regression module. Finally, the offsets between the future trajectory and the initial position are predicted. For each layer we provide the output dimensions.

commonly used for evaluating the performance of trajectory prediction methods [10], [40], [35], and therefore, we use it here. Note that the input data is still processed at 10 Hz, which is unaffected by the output resolution. Furthermore, a higher output-rate can be easily added. A ground truth maneuver output with 3 classes (straight, left lane change, right lane change) is of size  $5\times3$  and is computed from the lane IDs. The lane ID is provided by the datasets and not part of the input data. If the lane ID of a vehicle changes between subsequent frames, we assume a lane change with a duration of 2 s before and after the corresponding frames and infer the maneuver direction from the lane IDs.

The regression layer outputs trajectory offsets that describe the difference between the predicted position and the initial position at t=0 of the target vehicle. This improves the training since the identity mapping from input to output is achieved by predicting an offset of zero. At inference time, the model's output is de-normalized and added to the target's current position.

#### C. Spatio-temporal Convolutional Neural Network

Given the data representation proposed in Sec. III-A, we use four convolutional layers for feature aggregation. As carried out by Bai et al. [3], we apply convolutions along the time dimension to detect local correlations between states in the history sequence. In addition to that, we extend the temporal convolution along the spatial neighbor dimension resulting in 2D convolutions. The hierarchical stacking of layers ensures that the low-level features from the input data can be combined in the subsequent layers for higher-level representations. We conclude that higher layers learn to interpret the convolution along the neighborhood dimension and therefore account for the absolute position of the neighbor. This assumption is backed up by Kayhan et al. [21], who claim that CNNs can encode absolute spatial locations from boundary effects. Since we use zero padding and our spatial dimension is only eight, we can expect to make use of the boundary effects during training.

It is worth pointing out that the learning process of the temporal and spatial relations among all vehicles happens jointly. This makes the prediction fast since, in contrast to recurrent models, no previous hidden states need to be computed and taken into account. Analogously to visionbased 2D convolutions with multiple image channels, the filter depth is matched with the number of feature channels.

Stacking multiple convolutional layers increases the receptive field of the CNN. For temporal convolutional networks, the receptive field determines the number of past time steps that influence a single output. Especially for longer input sequences, this is an important factor for prediction performance. Besides deepening the architecture by adding layers, it is also possible to increase the size of the convolution kernels. However, since all techniques increase the number of training parameters and therefore the model complexity, we use dilated convolutions. A dilated convolution expands the receptive field without reducing the resolution or coverage as pointed out by Yu and Koltun [39]. We implement dilations along the temporal dimension by adding zeros between kernel entries depending on the dilation rate. We provide an experiment on the improvement by using dilated convolutions in Sec. IV-E.

The complete network architecture is illustrated in Fig. 3. Each convolutional layer uses a leaky rectified linear unit (ReLU) activation function [29]. At the first layer, the input tensor of size  $4 \times 8 \times 30$  is convolved with 24 filters with a kernel size of  $5 \times 10$  and dilation of one. Each output in the resulting feature map of size  $24 \times 4 \times 12$  has a receptive field of  $4 \times 5 \times 19$ . Furthermore, a dilated convolution with 40 filters of size  $3 \times 3$  is carried out resulting in a size of  $40 \times 2 \times 8$ . The third layer applies 56 filters of size  $2 \times 3$ with dilation, which leads to a map with  $56 \times 1 \times 4$  features. Finally, the last convolutional layer reduces the channel dimension with 24  $1 \times 1$  kernels to 24 feature channels and a final receptive field of  $4 \times 8 \times 27$ .

#### D. Classification

When only using a regression module to predict the trajectory, the model needs to capture multiple possible maneuvers as motivated in Fig. 1. Mozaffari et al. [30] point out that this can lead to predictions that are averaged over all possible modes. We decide to first classify the future maneuver of the target vehicle and to predict a trajectory that depends on the estimated maneuver.

To classify the future maneuver of the vehicle with our spatio-temporal CNN defined in Sec. III-C, we flatten the output of the last convolutional layer with a size of 24.4=96and pass it to a fully connected layer with a leaky ReLU activation function and a hidden dimension of 40. A final fully connected layer processes the hidden feature vector of length 40. The output layer predicts the class logits, which can be normalized to a probability for each of the three maneuvers "straight", "left", and "right" at each prediction step resulting in an output vector of size  $5 \times 3$ . Selecting the indices of the maneuvers with the highest predicted probability results in an estimated maneuver intention vector of size  $5 \times 1$  containing values between 0 and 2. Based on the three possible maneuvers at each of the P=5 prediction steps, there are in total  $3^5=243$  maneuver sequences the classification network can model. During training, we optimize the network parameters by minimizing the sum of negative log-likelihoods at each prediction step reading

$$\mathcal{L}_{\text{class}} = \sum_{t=1}^{P} -\log p(c_t), \tag{1}$$

with  $p(c_t)$  denoting the predicted probability for the ground truth class  $c_t$  at time step t. We analyze the improvements on the prediction performance achieved by our classification module in Sec. IV-E.

#### E. Regression

To make maneuver-based predictions, we feed the classified maneuver sequence directly to the regression module. As discussed above, there are 243 possible combinations of maneuver predictions for a given input tensor. We concatenate the resulting vector of size  $5 \times 1$  containing the indices of the five most likely maneuvers at each time step with the flattened CNN output vector resulting in 101 features. We pass this vector to a fully connected layer with the same architecture as in the classification module. The corresponding hidden vector of size 40 is the input for the final output layer with a linear activation function, which then predicts a vector of size  $5 \times 2$  containing the x and y positions at each of the 5 prediction steps. During training, we use the ground truth maneuvers as input for the regression module, whereas during evaluation, we use the maneuvers predicted by the classification module. The training aims at minimizing the batch-wise root mean squared error (RMSE) reading

$$\mathcal{L}_{\text{reg}} = \sqrt{\frac{1}{P} \sum_{t=1}^{P} \|\hat{\boldsymbol{x}}_t - \boldsymbol{x}_t\|^2}, \qquad (2)$$

with predicted locations  $\hat{x}_t$  and ground truth locations  $x_t$  for each time step t, where each location consists of the x and y positions.

With the proposed architecture, the temporal and spatial information of the highway scene around a target vehicle is jointly encoded by applying 2D convolutions. As pointed out by Bai et al. [3], the temporal convolution avoids a long backpropagation path resulting in a more compact model, which is easier to train compared to RNNs. Also, the prediction can be carried out in parallel due to the temporal convolutional structure and is therefore fast. The total amount of trainable parameters for the design depicted in Fig. 3 is 65,721. For comparison, the convolutional social pooling approach [10] has 194,954 and the multiple futures prediction approach [37] has 1,073,644 trainable parameters.

## IV. EXPERIMENTAL EVALUATION

The main focus of this work is to predict the lane change maneuver and the corresponding future trajectory of a vehicle based on past information about states including the neighborhood. We present our experiments to show the capabilities of our method and to support our key claims made about our work, which are: (i) Successfully classifying lane change maneuvers and predicting a corresponding trajectory for real-world scenarios by (ii) performing joint spatio-temporal feature aggregation with 2D CNNs, (iii) and outperforming state-of-the-art methods.

## A. Implementation Details

We train the classification and regression modules with the Adam optimizer [23] and a learning rate of  $7 \cdot 10^{-5}$ . Additionally, we augment the training set by overlapping the input data by 20 frames resulting in 506,180 training examples. This leads to better coverage of different lane change maneuvers in the training data. The total training time for each module is 17.5 h with 300 epochs on an Intel Xeon W-2133 CPU and an Nvidia Quadro RTX 5000 GPU. It takes around 0.3 ms on average to classify the future maneuver sequence and to predict the corresponding trajectory.

#### B. Experimental Setup

We consider two different datasets to train and evaluate our proposed approach on real-world scenarios. The highD dataset [24] contains 110,000 vehicle trajectories from 60 highway recordings. The highways are located in Germany and have 2 to 3 lanes. To show that our approach generalizes well on different driving behaviors and highway structures, we also use the NGSIM dataset [16] with 9,206 vehicles driving on two different US highways with 5 to 6 lanes. Both datasets provide access to position, velocity, and acceleration. For each dataset, we use 70% of all vehicle trajectories for training, 10% for validation, and 20% for testing. For a better comparison, we use the same vehicle trajectories for testing as Song et al. [35].

## C. Qualitative Analysis

The qualitative experiments evaluate the prediction results on real-world data and support the claim that our modular approach of classification and regression results in reasonable predicted maneuver-based trajectories. We show that the prediction relies on the spatio-temporal encoding of the target vehicle's surroundings by modifying the neighborhood leading to different results. The original scenes are taken from recording 56 of the highD dataset.



Fig. 4: Two scenarios with predictions for the black target vehicle (T). *Left*: The original scene with a dashed red line as ground truth. The solid line represents the first seconds of the predicted future trajectory. The velocity of neighboring vehicles is depicted by the length of the gray velocity vectors. Our approach successfully classifies a left lane change (blue) since the car in front (F) is slower. *Right*: In case we modify the neighborhood by removing the car in front, the prediction changes to a straight driving maneuver (black).



Fig. 5: Three scenarios with modified environments. *Top*: Original highD scene, the target can not change the lane and needs to decelerate. *Middle*: Removed leading vehicle, the target does not need to decelerate resulting in a longer trajectory. *Bottom*: Removed left rear neighbor, a left lane change (blue) is now possible and therefore predicted. Note that the velocity of neighboring vehicles is depicted by the length of the gray velocity vectors.

In Fig. 4, we consider two possible scenarios for the same highway scene. The original neighborhood is depicted on the left. A slower car (F) is in front of the black target vehicle (T) and forces it to change the lane or decelerate. In this case, the target will let the left center (L) vehicle pass and perform a lane change. The ground truth trajectory is represented by a red dashed line. Our approach successfully classifies a left lane change maneuver indicated by the blue color and predicts a trajectory that matches the ground truth. To demonstrate the influence of the semantic neighborhood on the prediction, we remove the leading car (F) by replacing the corresponding entries in the input data with zeros. Based on the new neighborhood, our model predicts a straight driving maneuver since there is no longer a need to change the lane.

The second example in Fig. 5 demonstrates a more advanced highway scene. The black target vehicle (T) is again facing a slower vehicle in front (F). Additionally, two cars at the rear and center of the left lane (RL and L) hinder the target from changing the lane. Our approach predicts that the car will stay on the lane and slow down. The ground truth is shown in dashed red and confirms our prediction. In the middle part of Fig. 5, we removed the truck in front (F) from the spatio-temporal input representation. One can see that the predicted trajectory is still straight, but now longer compared to the previous case. This can be reasoned by the fact that there is no longer a slower leading car that forces the target vehicle to decelerate. In a third scenario, we keep the truck in front (F) and remove the car on the rear left (RL). The classification module now outputs a lane change maneuver to the left (blue) and the corresponding trajectory moves to

the left lane. All three examples show that our approach can predict reasonable maneuver intentions and trajectories from the neighborhood of the target vehicle.

#### D. Quantitative Analysis

We present a second experiment to support the claim that our approach outperforms existing state-of-the-art methods for trajectory prediction. For a fair comparison, we evaluate our approach on the same test sets for each dataset as done by Song et al. [35]. We report the root mean squared error (RMSE) for each prediction step, which is a common metric for trajectory prediction. We compare the results with the baseline methods S-LSTM [1], CS-LSTM [10], S-GAN [15], MATF [40] and PiP-noPlan [35]. In case of S-GAN and MATF which are stochastic models, we take the best root mean squared error among 3 sampled trajectories. We refer to the PiP-noPlan implementation since the planning coupled module in PiP uses ground truth future trajectory of the target vehicle to predict the future trajectory of neighboring cars which would result in an unfair comparison.

The final results are shown in Tab. I. In general, our method shows better results for both datasets at the prediction steps. Especially for the highD dataset, the proposed approach outperforms other state-of-the-art methods at larger prediction horizons. All methods show larger RMSE values for the NGSIM dataset, which can be explained by larger noise in the data as pointed out by Krajewski et al. [24]. This noise leads to smaller performance improvements between the baselines and also explains why our method achieves a smaller margin compared to the evaluation on the highD

| Dataset | Time | S-LSTM [1] | CS-LSTM [10] | S-GAN [15] | MATF [40] | PiP-noPlan [35] | Our Approach |
|---------|------|------------|--------------|------------|-----------|-----------------|--------------|
| NGSIM   | 1 s  | 0.60       | 0.58         | 0.57       | 0.66      | 0.55            | 0.53         |
|         | 2 s  | 1.28       | 1.26         | 1.32       | 1.34      | 1.20            | 1.17         |
|         | 3 s  | 2.09       | 2.07         | 2.22       | 2.08      | 2.00            | 1.93         |
|         | 4 s  | 3.10       | 3.09         | 3.26       | 2.97      | 3.01            | 2.88         |
|         | 5 s  | 4.37       | 4.37         | 4.40       | 4.13      | 4.27            | 4.05         |
| highD   | 1 s  | 0.19       | 0.19         | 0.30       | -         | 0.18            | 0.10         |
|         | 2 s  | 0.57       | 0.57         | 0.78       | -         | 0.53            | 0.21         |
|         | 3 s  | 1.18       | 1.16         | 1.46       | -         | 1.09            | 0.41         |
|         | 4 s  | 2.00       | 1.96         | 2.34       | -         | 1.86            | 0.78         |
|         | 5 s  | 3.02       | 2.96         | 3.41       | -         | 2.81            | 1.34         |

TABLE I: Comparison of the root mean squared error at each prediction step evaluated on the NGSIM [16] and highD [24] datasets. Bold numbers indicate the best result. The baseline results are reported by Song et al. [35]. For the stochastic models S-GAN and MATF, the best root mean squared error among 3 sampled trajectories is reported. Note that for PiP, we refer to the PiP-noPlan implementation for a fair comparison, since the planning coupled module in PiP partially uses ground truth information.

| Time       | Only Input $x$ and $y$ | Without<br>Neighborhood | Shuffled<br>Neighborhood | Without<br>Dilation | Without<br>Maneuver | Sampled<br>Maneuver | Full<br>Approach |
|------------|------------------------|-------------------------|--------------------------|---------------------|---------------------|---------------------|------------------|
| 1 s        | 0.24                   | 0.08                    | 0.15                     | 0.15                | 0.14                | 0.14                | 0.10             |
| 2 s<br>3 s | 0.64                   | 0.18                    | 0.25                     | 0.30                | 0.26                | 0.33                | 0.21             |
| 4 s        | 2.02                   | 0.83                    | 0.82                     | 0.90                | 0.87                | 1.04                | 0.78             |
| 5 s        | 3.00                   | 1.52                    | 1.40                     | 1.45                | 1.48                | 1.60                | 1.34             |

TABLE II: Ablation study on the highD dataset [24].

dataset. We conclude that noise in for example the vehicle speed heavily influences the prediction due to the high average speed on highways.

# E. Ablation Study

Finally, we conduct an ablation study to evaluate how much each proposed component of our method contributes to the performance reported in Tab. I. We retrain and evaluate our approach with several modifications on the highD [24] dataset since the data is more accurate as previously described. The results are shown in Tab. II.

In a first ablation experiment, we retrain our method only using x- and y-position as input channels. The result shown in Tab. II indicates that adding velocity and acceleration as inputs leads to a major performance improvement, which justifies the multi-channel design of the proposed 3D input tensor.

If the local neighborhood defined in Sec. III-A is ignored for prediction as done by Nikhil and Morris [31], the prediction performance is slightly better for short-term predictions, but worse for larger prediction horizons. This indicates that the behavior of other traffic participants mainly influences the future trajectory at larger horizons and should not be ignored. Furthermore, we retrain our approach with a shuffled neighborhood order and the competitive result for prediction steps larger than 1 s indicates that the CNN is able to learn the absolute position across the vehicle dimension independent of the ordering.

We conducted an additional ablation study to investigate the effect of dilated convolutions used to increase the receptive field and to reduce the number of parameters. If no dilated convolutions are used, the size of the dense layers in Fig. 3 needs to be increased to account for the resulting larger feature maps at the output stage of the CNN. This results in a model with 90,681 parameters. Our experiment shows that the use of dilated convolutions leads to a better performance while having fewer parameters to train.

The last ablation study investigates the effect of our maneuver-based prediction. First, if no maneuvers are used, the performance degrades. We hypothesize that the regression module alone is not capable of capturing all maneuvers and therefore a prior classification improves distinguishing between these modes. Finally, we show that our classification module can predict different maneuvers based on the previous input. We test this by comparing to a model which randomly samples a maneuver sequence from the training distribution and passes it to the regression module. It can be seen that the sampled maneuvers lead to a worse performance compared to our trained classification module.

## V. CONCLUSION

In this paper, we presented a novel approach to classify the future lane change maneuver of a target vehicle and predict the corresponding trajectory. Our model operates on a 3D spatio-temporal input representation encoding the neighborhood around the target. We exploit the local correlations in neighboring state sequences with spatio-temporal 2D convolutions resulting in a simple, memory-efficient architecture with fast and parallelized inference. This allows us to successfully account for different possible driving maneuvers which are used for a more informed prediction. We implemented and evaluated our approach on two datasets taken at different locations and provided comparisons to other existing techniques and supported all claims made in this work. The experiments suggest that our approach can successfully predict maneuver-based trajectories and outperforms state-of-the-art methods.

#### ACKNOWLEDGMENTS

We thank J. Behley, X. Chen, and L. Wiesmann for fruitful discussions and Song et al. [35] for providing their data preprocessing code and corresponding test split.

#### References

- A. Alahi, K. Goel, V. Ramanathan, A. Robicquet, L. Fei-Fei, and S. Savarese. Social LSTM: Human Trajectory Prediction in Crowded Spaces. In Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2016.
- [2] J. Amirian, J.B. Hayet, and J. Pettre. Social Ways: Learning Multi-Modal Distributions of Pedestrian Trajectories With GANs. In Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition Workshops, 2019.
- [3] S. Bai, J.Z. Kolter, and V. Koltun. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv* preprint, abs/1803.01271, 2018.
- [4] M. Bansal, A. Krizhevsky, and A. Ogale. ChauffeurNet: Learning to Drive by Imitating the Best and Synthesizing the Worst. In Proc. of Robotics: Science and Systems (RSS), 2019.
- [5] K. Brown, K. Driggs-Campbell, and M.J. Kochenderfer. Modeling and Prediction of Human Driver Behavior: A Survey. *arXiv preprint*, abs/2006.08832, 2020.
- [6] S. Casas, W. Luo, and R. Urtasun. IntentNet: Learning to Predict Intention from Raw Sensor Data. In Proc. of the Conf. on Robot Learning (CoRL), 2018.
- [7] J. Chung, C. Gulcehre, K. Cho, and Y. Bengio. Empirical Evaluation of Gated Recurrent Neural Networks on Sequence Modeling. In Proc. of the Advances in Neural Information Processing Systems Workshops, 2014.
- [8] H. Cui, V. Radosavljevic, F.C. Chou, T.H. Lin, T. Nguyen, T.K. Huang, J. Schneider, and N. Djuric. Multimodal Trajectory Predictions for Autonomous Driving using Deep Convolutional Networks. In Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA), 2019.
- [9] S. Dai, L. Li, and Z. Li. Modeling Vehicle Interactions via Modified LSTM Models for Trajectory Prediction. *IEEE Access*, 7:38287– 38296, 2019.
- [10] N. Deo and M.M. Trivedi. Convolutional Social Pooling for Vehicle Trajectory Prediction. Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition Workshops, 2018.
- [11] N. Deo and M.M. Trivedi. Multi-Modal Trajectory Prediction of Surrounding Vehicles with Maneuver based LSTMs. In Proc. of the IEEE Vehicles Symposium (IV), 2018.
- [12] F. Diehl, T. Brunner, M. Truong-Le, and A. Knoll. Graph Neural Networks for Modelling Traffic Participant Interaction. In Proc. of the IEEE Vehicles Symposium (IV), 2019.
- [13] W. Ding, J. Chen, and S. Shen. Predicting Vehicle Behaviors Over An Extended Horizon Using Behavior Interaction Network. In Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA), 2019.
- [14] I.J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio. Generative Adversarial Networks. In Proc. of the Advances in Neural Information Processing Systems (NIPS), 2014.
- [15] A. Gupta, J. Johnson, L. Fei-Fei, S. Savarese, and A. Alahi. Social GAN: Socially Acceptable Trajectories with Generative Adversarial Networks. In Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2018.
- [16] J. Halkias and J. Colyar. Next Generation SIMulation Fact Sheet. Technical report, Federal Highway Administration, 2006.
- [17] S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9(8):1735–1780, 1997.
- [18] S. Hoermann, M. Bach, and K. Dietmayer. Dynamic Occupancy Grid Prediction for Urban Autonomous Driving: A Deep Learning Approach with Fully Automatic Labeling. In Proc. of the IEEE Intl. Conf. on Robotics & Automation (ICRA), 2018.
- [19] J. Hong, B. Sapp, and J. Philbin. Rules of the Road: Predicting Driving Behavior with a Convolutional Model of Semantic Interactions. In Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR), 2019.
- [20] Y. Hu, W. Zhan, and M. Tomizuka. Probabilistic Prediction of Vehicle Semantic Intention and Motion. In Proc. of the IEEE Vehicles Symposium (IV), 2018.

- [21] O.S. Kayhan and J.C. van Gemert. On Translation Invariance in CNNs: Convolutional Layers can Exploit Absolute Spatial Location. In Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR), 2020.
- [22] B. Kim, C. M. Kan, J. Kim, S. H. Lee, C. C. Chung, and J. W. Choi. Probabilistic Vehicle Trajectory Prediction over Occupancy Grid Map via Recurrent Neural Network. In Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC), 2017.
- [23] D.P. Kingma and J. Ba. Adam: A Method for Stochastic Optimization. In Proc. of the Int. Conf. on Learning Representations (ICLR), 2015.
- [24] 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. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC), 2018.
- [25] A. Kuefler, J. Morton, T. A. Wheeler, and M.J. Kochenderfer. Imitating Driver Behavior with Generative Adversarial Networks. In Proc. of the IEEE Vehicles Symposium (IV), 2017.
- [26] S. Lefèvre, D. Vasquez, and C. Laugier. A survey on motion prediction and risk assessment for intelligent vehicles. *Journal of Robotics and Mechanical Engineering Research*, 1:1–14, 2014.
- [27] L.L. Li, B. Yang, M. Liang, W. Zeng, and M. Ren. End-to-end Contextual Perception and Prediction with Interaction Transformer. *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems* (IROS), 2020.
- [28] W. Luo, B. Yang, and R. Urtasun. Fast and Furious: Real Time Endto-End 3D Detection, Tracking and Motion Forecasting with a Single Convolutional Net. In Proc. of the IEEE Conf. on Computer Vision and Pattern Recognition (CVPR), 2018.
- [29] A. Maas, A.Y. Hannun, and A.Y. Ng. Rectifier Nonlinearities Improve Neural Network Acoustic Models. In Proc. of the Int. Conf. on Machine Learning (ICML), 2013.
- [30] S. Mozaffari, O.Y. Al-Jarrah, M. Dianati, P. Jennings, and A. Mouzakitis. Deep Learning-Based Vehicle Behavior Prediction for Autonomous Driving Applications: A Review. arXiv preprint, abs/1912.11676, 2019.
- [31] N. Nikhil and B.T. Morris. Convolutional Neural Network for Trajectory Prediction. In Proc. of the Europ. Conf. on Computer Vision Workshops, 2018.
- [32] S.H. Park, B. Kim, C.M. Kang, C.C. Chung, and J.W. Choi. Sequenceto-Sequence Prediction of Vehicle Trajectory via LSTM Encoder-Decoder Architecture. In *Proc. of the IEEE Vehicles Symposium (IV)*, 2018.
- [33] A. Sadeghian, V. Kosaraju, A. Sadeghian, N. Hirose, S. H. Rezatofighi, and S. Savarese. SoPhie: An Attentive GAN for Predicting Paths Compliant to Social and Physical Constraints. In Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR), 2019.
- [34] J. Schulz, C. Hubmann, J. Lochner, and D. Burschka. Interaction-Aware Probabilistic Behavior Prediction in Urban Environments. In Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems (IROS), 2018.
- [35] H. Song, W. Ding, Y. Chen, S. Shen, M. Wang, and Q. Chen. PiP: Planning-informed Trajectory Prediction for Autonomous Driving. *Proc. of the Europ. Conf. on Computer Vision (ECCV)*, 2020.
- [36] L. Sun, W. Zhan, C.Y. Chan, and M. Tomizuka. Behavior Planning of Autonomous Cars with Social Perception. In Proc. of the IEEE Vehicles Symposium (IV), 2019.
- [37] Y. Tang and R. Salakhutdinov. Multiple Futures Prediction. Proc. of the Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [38] L. Xin, P. Wang, C.Y. Chan, J. Chen, S.E. Li, and B. Cheng. Intentionaware Long Horizon Trajectory Prediction of Surrounding Vehicles using Dual LSTM Networks. In Proc. of the IEEE Intl. Conf. on Intelligent Transportation Systems (ITSC), 2018.
- [39] F. Yu and V. Koltun. Multi-Scale Context Aggregation by Dilated Convolutions. In Proc. of the Int. Conf. on Learning Representations (ICLR), 2016.
- [40] T. Zhao, Y. Xu, M. Monfort, W. Choi, C. Baker, Y. Zhao, Y. Wang, and Y.N. Wu. Multi-Agent Tensor Fusion for Contextual Trajectory Prediction. In Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition (CVPR), 2019.
- [41] A. Zyner, S. Worrall, and E. Nebot. Naturalistic Driver Intention and Path Prediction Using Recurrent Neural Networks. *IEEE Trans. on Intelligent Transportation Systems (ITS)*, 21(4):1584–1594, 2018.