

# Graph Neural Networks for Modelling Traffic Participant Interaction

Frederik Diehl<sup>†</sup> and Thomas Brunner<sup>†</sup> and Michael Truong Le<sup>†</sup> and Alois Knoll<sup>‡</sup>

**Abstract**—By interpreting a traffic scene as a graph of interacting vehicles, we gain a flexible abstract representation which allows us to apply Graph Neural Network (GNN) models for traffic prediction. These naturally take interaction between traffic participants into account while being computationally efficient and providing large model capacity. We evaluate two state-of-the-art GNN architectures and introduce several adaptations for our specific scenario. We show that prediction error in scenarios with much interaction decreases by 30% compared to a model that does not take interactions into account. This suggests that interaction is important, and shows that we can model it using graphs. This makes GNNs a worthwhile addition to traffic prediction systems.

## I. INTRODUCTION

Short-term accurate behavior prediction of traffic participants is important for applications such as automated driving or infrastructure-assisted human driving[1]. A major open research question is how to model *interaction* between traffic participants. In the past, interactions have been modelled by either creating a representation of one or several traffic participants [2], [3] or by using a fixed environment representation such as a simulated lidar beam [4].

However, these methods impose certain disadvantages: A fixed environment representation poses a much harder problem to learn, since we cannot use data we might have extracted previously. Traffic participant representations, on the other hand, scale computationally with the amount of possible interactions, require a human to decide on a useful representation, and underspecify the problem one should learn.

By modelling each vehicle a node and possible interactions between vehicles as edges (see Fig. 1 for a visualization), we gain a sparse and high-level representation of a traffic scene as a graph.

At the same time, it has been shown [3]–[5] that machine learning models and particularly (deep) neural networks perform well on this problem. Yet most available deep learning models operate on data of a fixed size and with a fixed spatial organization such as single data points, time series, or images.

Only fairly recently [6], [7] have GNNs, i.e. neural networks operating on graph data, seen research interest and enjoyed successes. Later models [8], [9] only operate on a node’s local neighbourhood. This greatly improves scalability while improving performance.

<sup>†</sup>Frederik Diehl, Thomas Brunner, and Michael Truong Le are with fortiss GmbH, affiliated institute of Technische Universität München, Munich, Germany

<sup>‡</sup>Alois Knoll is with the Chair of Robotics, Artificial Intelligence and Real-time Systems, Technische Universität München, Munich, Germany




Fig. 1. Interaction graph of a traffic situation. Interactions are assumed to occur between the ego vehicle (one vehicle and connections from the vehicles it interacts with are highlighted in orange, but all vehicles are processed simultaneously) and its up to eight neighbours, assigned by current lane (black lines separate lanes). These interactions are modelled as directed edges leading from the neighbouring vehicles to the ego vehicle. Except for the example ego vehicle, edge directions have been omitted for clarity. As can be seen, this representation is sparse and models the whole traffic situation at once. Best viewed in color.

Marrying the representation of a traffic situation as a graph with the modelling capabilities of GNN models promises a clear method to take interactions between traffic participants into account, good predictive performance, and efficient computation.

To evaluate this, we conduct traffic participant prediction on two real-world datasets, evaluating their predictive performance and comparing them to three baseline models. We show that prediction error decreases by 30% compared to our baseline when interaction is plentiful and performs no worse when little interaction occurs. At the same time, computational complexity remains reasonable and scales with linearly in the number of interactions.

This suggests a graph interpretation of interacting traffic participants is a worthwhile addition to traffic prediction systems.

Our main contributions are:

- 1) We show that representing interactions as graphs leads to better performance.
- 2) We introduce several adaptations to two state-of-the-art GNN models.
- 3) We study both the results of different graph construction techniques and our introduced adaptations on two different datasets.

## II. RELATED WORK

Since traffic participant prediction is a key feature of autonomous driving and traffic simulations, it has been

a focus of extensive research for decades. This has led to a multitude of different algorithms useful for varying prediction timespans and computational resources.

### A. Traffic Prediction

Following the survey by Lefèvre *et al.* [10], we roughly categorize traffic participant prediction into three subgroups of ascending complexity: Physics-based, maneuver-based, and interaction-aware.

*Physics-based models* usually assume little vehicle action and instead use constant velocity or acceleration. The vehicle motion is then predicted from a physical model only. These models can be used for tracking [11] but often fail for predictions longer than a second or when vehicle interaction plays an important role.

*Maneuver-based models* use a set of maneuver prototypes and either match the past trajectory directly using cluster-based approaches [12] or from vehicle features using machine learning methods [13]–[15]. While these models are now able to include more complex maneuver, they also cannot take interaction into account.

*Interaction-aware models* aim to include interactions between vehicles in their predictions. These include an expansion of maneuver-based models which account for collision probabilities [16], coupled Hidden Markov Models (HMMs), which model pairwise entity dependencies [17], or machine learning-based models.

Machine learning-based models also vary in complexity and goal. Lenz *et al.* [3] use simple feed-forward neural networks to create a fast model for use in a Monte-Carlo Tree Search algorithm. Morton *et al.* [5] evaluate Recurrent Neural Networks (RNNs) for the same task, also trained in a supervised fashion. Conversely, Kuefler *et al.* [4] use Generative Adversarial Imitation Learning to imitate human driving behavior using reinforcement learning. For all of these, performances crucially depend on the representation of the environment.

### B. Environmental Representation

Environmental representation can be differentiated by their abstractness: One can represent the environment as data close to sensor input such as LIDAR beams [4], camera images, or a simple gridmap. Alternatively, one can represent the environment as a list of discrete objects.

1) *Sensor-like representation*: A sensor-like representation does not require expert knowledge to define the features to use and remains of constant size independent of factors such as traffic density. At the same time, it can receive information on many vehicles. However, the representation is inefficient (requiring many LIDAR beams or pixels per vehicle) and we are forced to learn not just driving behavior but also the extraction of vehicles from sensor data.

2) *Discrete object representation*: Alternatively, we can represent each vehicle as an object with certain attributes. Predictions are then created per car. Interaction is then a matter of choosing the correct environment representation, which may be as simple as the distance and approach speed

to the preceding vehicle — as in the Intelligent Driver Model (IDM) [2] — or might contain a multitude of preprocessed features [3]. However, these models by design have to be simplistic in their assumptions of interaction between traffic participants and are therefore limited.

Several of these shortcomings can be avoided by thinking about traffic participants and their interactions as nodes and edges in a graph. A behavior prediction model then operates on that graph, producing predictions for each node.

## III. TRAFFIC PARTICIPANT PREDICTION FROM A GRAPH

While there are several different GNN architectures, experimental results suggest the relatively simple Graph Convolutional Network (GCN) model still performs best over a wide variety of tasks [18]. We also evaluate the Graph Attention Network (GAT) model, since it allows us to easily include edge features into the model.

### A. Graph Convolutional Networks

GCNs [8] are an approach for node-based classification or prediction on a graph. Analogous to convolutions on images or time series, a GCN applies the same operation on all nodes. Like other neural networks, it is defined by a series of differently-parameterized layers which are applied successively.

1) *The Base Model*: Each layer of the GCN uses

$$H^{(l+1)} = \sigma \left( \tilde{D}^{-\frac{1}{2}} \tilde{A} \tilde{D}^{-\frac{1}{2}} H^{(l)} W^{(l)} \right) \quad (1)$$

as a transformation. Here,  $H^{(l)}$  is the  $l$ th layer’s activations,  $\tilde{A}$  is the adjacency matrix with added self-connections between nodes,  $\tilde{D}$  is the degree vector of  $\tilde{A}$ , and  $W^{(l)}$  is the  $l$ th layer’s learnable weight matrix.

This is equivalent to a first-order approximation of a localized spectral filter, but has two crucial advantages: The Graph Laplacian does not need to be inverted (which would incur computational cost of  $O(n^3)$ ) and the transformation specified by  $l$  layers takes exactly the  $l$ -hop neighborhood of a node into account. Accordingly, computational complexity scales linearly in the number of edges and can take vehicles into account which are not directly connected to the ego vehicle. This makes it more efficient than a naive encoding of the neighboring vehicles.

2) *Adaptations for the GCN*: We originally applied the GCN exactly as described by Kipf and Welling [8]. However, we found several changes to be crucial:

- *Residual Weights*: GCNs compute the next layer’s features for a node from a spectral decomposition of that node’s neighborhood (and, with added self-connections, the ego node itself). However, this means a GCN cannot treat the ego node’s own features differently from any of its neighbors. In the prediction task, this appears to be a significant obstacle to good performance. Accordingly, we remove the self-connections but introduce a second weight matrix defining a transformation on

the ego node’s features. Our transformation equation is therefore

$$H^{(l+1)} = \sigma \left( D^{-\frac{1}{2}} A D^{-\frac{1}{2}} H^{(l)} W^{(l)} + H^{(l)} W_s^{(l)} \right). \quad (2)$$

- *Weight by Distance:* Kipf and Welling [8] note that the adjacency matrix can be binary or weighted. We evaluate weighting edges by the inverse distance, with self-loops set to 1.
- *Feed-forward output:* In addition, we no longer use a full GCN but replace the output layer with a feed-forward layer operating on each node’s features independently. This allows a better decoupling of the feature extraction (occurring in the first few GCN layers) and the final prediction from the extracted features.

### B. Graph Attention Networks

Graph Attention Network (GAT) [9] layers compute each node’s next representation by an attention mechanism over all of its neighbors.

1) *The Base Model:* Specifically, they compute attention coefficients

$$\alpha_{ij}^l = \text{softmax}_j \left( a(W^l h_i^{(l-1)}, W^l h_j^{(l-1)}) \right) \quad (3)$$

for each connected node pair, with  $h_i^l$  being the  $i$ th node’s feature in the  $l$ th layer and  $W^l$  being the learnable weight matrix for the  $l$ th layer.  $a$  is the learnable attention computation, implemented by a neural network. The node feature vector is then computed as

$$h_i^l = \sigma \left( \sum_{j \in N_i} \alpha_{ij}^l W^l h_j^{(l-1)} \right), \quad (4)$$

where  $\sigma$  is a non-linearity, usually ReLU.

In practice, Veličković *et al.* [9] note that learning is stabilized by using multi-head attention, i.e. using  $k$  differently-parameterized attention mechanisms and concatenating - or averaging in the last layer - the result. This allows features to be created from different subsets of nodes depending on the needs of these features.

As with GCNs, a GAT layer operates on local neighborhood only and therefore also scales linearly in the number of edges.

2) *Adaptations for the GAT:* As before, we also apply some adaptations to the base GAT model.

- *Edge attributes:* In the GAT as introduced by Veličković *et al.* [9], attention depends only on the features of the two nodes. However, we do have additional data - like the relative positions - available to us in this scenario. Accordingly, we augment the attention computation from Eq. 3 by including edge features, such that

$$\alpha_{ij}^l = \text{softmax}_j \left( a(W^l h_i^{(l-1)}, W^l h_j^{(l-1)}, e_{ij}) \right). \quad (5)$$

We do not learn successive edge features but instead use the relative positions for each layer.

- *Residual Weights:* While the GAT should be able to learn by itself to concentrate one attention head onto

the ego node, we also evaluate explicitly adding a transformation of the ego node’s features.

- *Feed-forward output:* As with the GCN, our final output is produced by a feed-forward layer.

### C. Graph and Feature Construction

Formulating the prediction problem as a graph still leaves open the task of how we construct said graph and the node features. While there is an obvious strategy to construct node features - namely to use the corresponding car features like position or velocity - no such strategy is apparent to construct connections between the nodes. However, four basic strategies are immediately apparent:

- *Self connections:* This only adds self-loops to the graph. It ignores all interaction performance and should perform identically to a simple model operating on the vehicle data only.
- *All connections:* Connecting all vehicles ensures that no interactions are ignored. However, this ignores previous knowledge on spatial position and interaction and scales quadratically in the number of vehicles.
- *Preceding connection:* Arguably the most important interaction is with the vehicle immediately in front of us. We can therefore construct interactions only between the current vehicle and its predecessor.
- *Close vehicles:* Alternatively, we can argue that the main interactions are with the vehicles in an ego vehicle’s direct environment, which are at most eight vehicles located to the front, rear, and sides of the ego vehicle. This construction is similar to the approach by Lenz *et al.* [19] and Wheeler and Kochenderfer [20].

While we would prefer to learn these connecting strategies, this is a very difficult open problem and scales quadratically with the number of considered vehicles. We therefore only evaluate the fixed strategies. Connected to this, we also leave the interesting question of what neighborhood size is necessary to future work.

## IV. EXPERIMENTS

In order to evaluate the newly proposed models, we conduct a prediction experiment on real-world traffic data. We purposely keep baselines and models simple to demonstrate whether the graph interpretation is beneficial without introducing a multitude of confounding factors. We therefore do not include RNN architectures, simulation steps, or imitation learning.

From this, we aim to answer three main questions: (A) Which of our adaptations to GNNs are necessary? (B) How do we construct an interaction graph? (C) Does a graph model increase prediction quality?

### A. Datasets

We conduct our experiment on two different datasets: The NGSIM I-80 dataset [21] and the HighD dataset [22].

1) *NGSIM*: The NGSIM project’s I-80 dataset contains trajectory data for vehicles in a highway merge scenario for three 15-minute timespans. These are tracked using a fixed camera system. As Thiemann *et al.* [23] show, position, velocity, and acceleration data contain unrealistic values. We therefore smooth the positions using double-sided exponential smoothing with a span of  $0.5s$  and compute velocities from these.

We use two of the recordings as the training set and split the last one equally into validation and test set. We subsample the trajectory data to 1 FPS and extract trajectories consisting of a total of  $10s$  of length. The goal of the model is to predict the second half of the trajectory given the first five seconds.

2) *HighD*: Since the NGSIM dataset still contains many artifacts (errors in bounding boxes, undetected cars, complete non-overlap of bounding box and true vehicle), we additionally conduct experiments on the new HighD dataset [22], which is a series of drone recordings and extracted vehicle features from about 400 meters each from several locations on the German Autobahn. A total of 16.5 h of data is available, containing 110 000 vehicles with a total driving distance of 45 000 km. However, since the dataset consists mainly of roads without on- or off-ramps and without traffic jams, interaction seems limited: Only about 5% of the cars experience a lane change.

To avoid information leakage, we split the dataset by recording. The last 10% of the recordings are used as test set, the 10% before that as validation set. Trajectory construction is then identical to the NGSIM dataset.

## B. Baselines

We compare our approach to two different model-based static approaches, and one learned approach.

1) *Constant Velocity Model (CVM)*: This model considers each car to continue moving at the same velocity (both laterally and longitudinally) as the last frame it was observed.

2) *Intelligent Driver Model (IDM)*: The IDM [2] is a commonly-used driver model for microscopic traffic simulation since it is interpretable and collision-free. We use this to predict the changes in longitudinal velocity and keep the in-lane position constant.

The IDM’s acceleration is computed from both a *free road* and an *interaction* term. The free road acceleration is computed as

$$a_{free} = a_{max} \left( 1 - \left( \frac{v}{v_0} \right)^\delta \right),$$

with the maximum acceleration  $a_{max}$ , the acceleration exponent  $\delta$  and the desired velocity  $v_0$  being tunable parameters and the current velocity  $v$ . The interaction term is defined as

$$a_{int} = -a_{max} \left( \frac{s_0 + v * \tau}{s} + \frac{v \Delta v}{2s \sqrt{a_{max} b}} \right),$$

where the minimum distance to the front vehicle  $s_0$ , the time gap  $\tau$ , and the maximum deceleration  $b$  are tunable parameters.  $v$  is the vehicle’s speed and  $\Delta v$  the closing speed to its predecessor. The total acceleration is the sum of both

TABLE I  
OPTIMIZED PARAMETERS OF THE IDM.

Parameter		HighD	NGSIM [5]	
Desired velocity	$v_0$	$\left[ \frac{m}{s} \right]$	58.87	17.8
Maximum acceleration	$a_{max}$	$\left[ \frac{m}{s^2} \right]$	0.14	0.76
Time gap	$\tau$	$[s]$	0.12	0.92
Comfortable deceleration	$b$	$\left[ \frac{m}{s^2} \right]$	12.17	3.81
Minimum distance	$s$	$[m]$	14.46	5.249

the free road and the interaction acceleration. Since the IDM only outputs a longitudinal acceleration, we assume no lateral motion when using the IDM

We take the IDM parameters for the NGSIM dataset from Morton *et al.* [5]. For the HighD dataset, we tune the IDM’s parameters using guided random search with a total of 20 000 samples. Both values are listed in Table I.

3) *Independent Feed-Forward Model*: In addition to the models taking interaction into account, we also add a simple feed-forward neural network predicting the trajectory from only the ego vehicle’s past data. We use this baseline model to measure the improvement we gain from including interaction into our models.

## C. Model Configuration

Each model uses a similar configuration: Two layers producing a 256-dimensional feature representation followed by a feed-forward layer producing the final output: The displacement in x and y direction. All models use the ReLU nonlinearity. The GAT employs four attention heads (and 64-dimensional feature representations each).

Since the GNN models use two layers, their effective receptive field is the two-hop neighborhood from the ego vehicle.

All models receive inputs and produce outputs in fixed-length timesteps without recurrence. They are trained to predict displacement relative to the last position and receive position and velocity for each past timestep. They train to minimize the mean squared error over all outputs. All models are implemented in pytorch [24] using and expanding upon the pytorch-geometric library [25].

## D. Performance Measure

We report performances of the model by measuring the error in position between ground truth and prediction. We both report mean displacement over five seconds, weighting each timestep identically, and final displacement after five seconds.

## E. Experimental Procedure

Our choice of experiments is guided by the three main questions (sections V-A–C). To ensure meaningful results, we repeat each evaluation a total of ten times using different, randomly-chosen random seeds. In tables, we report all results as mean  $\pm$  standard deviation. Figures are violin

TABLE II  
RESULTS FOR OUR GNN ABLATION STUDY ON THE NGSIM I-80  
DATASET.  
(\*) USES 3 INSTEAD OF 10 EVALUATIONS.

	Mean Displ. [m]	Displ. @5s [m]
GCN Adaptations		
Default	2.50 ± 0.07	4.68 ± 0.15
no ff output	2.52 ± 0.06	4.66 ± 0.11
with weighted edges	2.60 ± 0.08	4.91 ± 0.15
no residuals & weighted edges	2.87 ± 0.04	5.19 ± 0.08
no residuals	3.74 ± 0.11	6.42 ± 0.10
GAT Adaptations		
Default	1.92 ± 0.04	3.45 ± 0.08
no ff output	1.93 ± 0.07	3.38 ± 0.16
no residuals	2.32 ± 0.02	3.96 ± 0.04
no edge features	2.40 ± 0.05	4.48 ± 0.09
Connection Strategy (GAT)		
Self-Connections	2.68 ± 0.05	5.08 ± 0.08
Preceding Connection	2.70 ± 0.04	5.11 ± 0.07
Neighbour Connection	1.93 ± 0.08	3.47 ± 0.13
All Connections (*)	2.41 ± 0.02	4.42 ± 0.03

plots, showing both individual results and the total result distribution.

We optimize both network adaptations and graph construction strategies on the NGSIM I-80 dataset since it is both smaller and contains more interactions. We then use these insights to pick the best-performing models and evaluate them on both the NGSIM I-80 and the HighD dataset.

## V. DISCUSSION

We structure our evaluation according to three research questions which answer (A) whether our proposed architectural adaptations are worthwhile, (B) which of the graph construction strategies should be preferred, and (C) whether the inclusion of interaction graph information improves performance.

### A. Which of our adaptations to GNNs are necessary?

In Section III-C, we proposed several changes to the GCN and GAT architectures. To answer which of these changes are beneficial, we conducted an ablation study whose results are listed in Table II. We evaluated this using the Neighbour Connection graph construction strategy.

For both models, removing the residual connections decreases the prediction error by at least 20%. This is likely because there is a clear difference between a neighbouring and the ego node in this task. We also found that using a feed-forward layer as last layer does produce a small increase in performance but also stabilizes training.

Introducing relative positions as edge features into the GAT seems to be a clear success, reducing the final displacement by about a meter. Contrary to that, edge weights for the GCN slightly decrease performance, especially when omitting residual weights. We believe that the main contribution of edge weights in our scenario is to discern between the ego

and surrounding vehicles, which is already more effectively modelled through residual weights.

We therefore evaluate the graph construction using the GAT model.

### B. How do we construct an interaction graph?

In Section III-C, we proposed four construction strategies for the interaction graph. We evaluate the quality of predictions with each of these strategies using the GAT models, since these seemed to perform best. We note that in practical scenarios, a tradeoff might be necessary between prediction quality and computational complexity. Table II shows results.

As expected, the Self-Connection strategy performs identically to the FF baseline model, and the Neighbour-Connection graph construction method performs best. Somewhat surprisingly, the Preceding-Connection strategy performs no better than the baseline.

We especially note that using the All-Connection strategy imposes significant computational disadvantages with quadratic instead of linear runtime and, in our experiments, a slowdown of about 50x.

We therefore use the Neighbour Connection graph construction strategy for our evaluation.

### C. Does a graph model increase prediction quality?

The motivation of our work is to evaluate whether it is beneficial to model interaction between traffic participants and whether this can be modelled in a graph construction. To answer this question, we compare models with interaction to a model without (FF). We also include a comparison with two classical models (CVM and IDM).

We chose the GAT model as best-performing GNN. We also included a GAT model without edge features (called GAT NEF in our figures and tables) for a fair comparison with the GCN model.

1) *NGSIM*: Fig. 2 and Table III show the performance of both our learning baseline and three of our GNN models.

As can be clearly seen, every GNN model performs better than the baseline. At the same time, there are clear performance differences between them: Both GCN and GAT NEF perform worse, which we assume is because these models cannot take relative positions directly into account and instead only act on the existence or non-existence of edges.

At the same time, the introduction of fixed edge features to the GAT model clearly shows its performance advantage, reducing the prediction error by a 30% compared to the FF baseline.

We note that the comparatively bad performance of the IDM in shorter timescales is consistent with previous work [3] and it is likely to achieve better performances in a closed- or open-loop simulation.

2) *HighD*: On the HighD dataset, our results are different: As Fig. 3 and Table IV show, there is no significant performance difference between either of the learned models, and no significant performance difference between the IDM and Constant Velocity Model (CVM) models. We believe




Fig. 2. Performance on the NGSIM Dataset. As can be seen, all three GNN models perform better than the baseline which does not take interaction into account.

TABLE III  
PERFORMANCE COMPARISON ON THE NGSIM DATASET.

	Mean Displ. [m]	Displ. @5s [m]
GAT	$1.89 \pm 0.02$	$3.40 \pm 0.04$
GAT NEF	$2.39 \pm 0.03$	$4.46 \pm 0.06$
GCN	$2.51 \pm 0.07$	$4.69 \pm 0.12$
FF	$2.78 \pm 0.08$	$5.22 \pm 0.14$
CVM	2.58	5.00
IDM	3.10	6.60

this to be a consequence of little interaction between the cars, which makes all learned models degenerate to the non-interaction case and makes the interaction term of the CVM model irrelevant. This shows that, even with no interaction, including interaction representations into our models does not cause performance degradation.

In summary, we show that (A) several of our changes result in better performance, (B) as does a good interaction graph construction strategy. (C) In total, our model retains performance on a dataset with little interaction and greatly improves it on a dataset with plentiful interaction.

## VI. CONCLUSION

We have proposed modelling a traffic scene as a graph of interacting vehicles. Through this interpretation, we gain a flexible and abstract model for interactions. To predict future traffic participant actions, we use Graph Neural Networks (GNNs), neural networks operating on graph data. These naturally take the graph model and therefore interaction into account. We evaluated two computationally efficient GNNs and proposed several adaptations for our scenario.

In a traffic dataset with plentiful interaction, including interactions decreased prediction error by over 30% compared to the best baseline model. At the same time, we saw no increase in prediction error on a dataset with little interaction.




Fig. 3. Performance on the HighD Dataset. As can be seen, the algorithms perform similar. We believe this to be due little interaction occurring in the dataset.

TABLE IV  
PERFORMANCE COMPARISON ON THE HIGHD DATASET.

	Mean Displ. [m]	Displ. @5s [m]
GAT	$0.47 \pm 0.04$	$1.04 \pm 0.07$
GAT NEF	$0.49 \pm 0.07$	$1.06 \pm 0.08$
GCN	$0.47 \pm 0.08$	$1.10 \pm 0.14$
FF	$0.45 \pm 0.06$	$1.09 \pm 0.09$
CVM	1.09	2.66
IDM	1.12	2.66

While we have improved prediction quality, much work remains to be done: This work is only a proof-of-concept that modelling interactions as a graph is worthwhile and should thus be seen as only one technique for one aspect of traffic prediction. Integrating this model into existing state-of-the-art methodology, particularly RNNs, remains an open task. At the same time, we would like to explore other graph construction strategies, particularly automatically finding relevant interactions.

## REFERENCES

- [1] G. Hinz, M. Buechel, F. Diehl, *et al.*, “Designing a far-reaching view for highway traffic scenarios with 5g-based intelligent infrastructure,” presented at the 8. Tagung Fahrerassistenz, 2017, p. 8.
- [2] M. Treiber, “Congested traffic states in empirical observations and microscopic simulations,” *Physical Review E*, vol. 62, no. 2, pp. 1805–1824, 2000.
- [3] D. Lenz, F. Diehl, M. T. Le, and A. Knoll, “Deep neural networks for markovian interactive scene prediction in highway scenarios,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 685–692.
- [4] A. Kuefler, J. Morton, T. Wheeler, and M. Kochenderfer, “Imitating driver behavior with generative adversarial networks,” in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 204–211.
- [5] J. Morton, T. A. Wheeler, and M. J. Kochenderfer, “Analysis of Recurrent Neural Networks for Probabilistic Modeling of Driver Behavior,” *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–10, 2016.

- [6] M. Gori, G. Monfardini, and F. Scarselli, "A new model for learning in graph domains," en, in *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, vol. 2, Montreal, Que., Canada: IEEE, 2005, pp. 729–734.
- [7] F. Scarselli, M. Gori, Ah Chung Tsoi, M. Hagenbuchner, and G. Monfardini, "The Graph Neural Network Model," en, *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 61–80, Jan. 2009.
- [8] T. N. Kipf and M. Welling, "Semi-supervised classification with graph convolutional networks," *arXiv:1609.02907 [cs, stat]*, Sep. 9, 2016. arXiv: 1609.02907.
- [9] P. Veličković, G. Cucurull, A. Casanova, *et al.*, "Graph Attention Networks," en, Oct. 2017.
- [10] S. Lefèvre, D. Vasquez, and C. Laugier, "A survey on motion prediction and risk assessment for intelligent vehicles," *Robomech Journal*, 2014.
- [11] R. Schubert, E. Richter, and G. Wanielik, "Comparison and evaluation of advanced motion models for vehicle tracking," *11th International Conference on Information Fusion*, pp. 1–6, 2008.
- [12] D. Vasquez and T. Fraichard, "Motion prediction for moving objects: A statistical approach," *Proc. IEEE International Conference on Robotics and Automation*, vol. 4, pp. 3931–3936, 2004.
- [13] M. Garcia Ortiz, J. Fritsch, F. Kummert, and A. Geppert, "Behavior prediction at multiple time-scales in inner-city scenarios," *Proc. IEEE Intelligent Vehicles Symposium*, pp. 1068–1073, 2011.
- [14] B. Morris, A. Doshi, and M. M. Trivedi, "Lane change intent prediction for driver assistance: Design and on-road evaluation," *Proc. IEEE Intelligent Vehicles Symposium*, pp. 895–901, 2011.
- [15] P. Kumar, M. Perrollaz, C. Laugier, *et al.*, "Learning-based approach for online lane change intention prediction," *Proc. IEEE Intelligent Vehicles Symposium*, pp. 797–802, 2013.
- [16] A. Lawitzky, D. Althoff, C. Passenberg, *et al.*, "Interactive scene prediction for automotive applications," in *IEEE Intelligent Vehicles Symposium*, 2013.
- [17] M. Brand, N. Oliver, and A. Pentland, "Coupled hiddenMarkov models for complex action recognition," *International Conference on Computer Vision and Pattern Recognition*, pp. 1–6, 1997.
- [18] O. Shchur, M. Mumme, A. Bojchevski, and S. Günnemann, "Pitfalls of Graph Neural Network Evaluation," en, *arXiv:1811.05868 [cs, stat]*, Nov. 2018, arXiv: 1811.05868.
- [19] D. Lenz, F. Diehl, M. T. Le, and A. Knoll, "Deep neural networks for Markovian interactive scene prediction in highway scenarios," in *2017 IEEE Intelligent Vehicles Symposium (IV)*, Jun. 2017, pp. 685–692.
- [20] T. A. Wheeler and M. J. Kochenderfer, "Factor graph scene distributions for automotive safety analysis," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, Nov. 2016, pp. 1035–1040.
- [21] Federal Highway Administration (FHWA), "US highway 80 dataset," FHWA-HRT-06-137, Tech. Rep., 2005.
- [22] R. Krajewski, J. Bock, L. Kloecker, and L. Eckstein, "The highd dataset: A drone dataset of naturalistic vehicle trajectories on german highways for validation of highly automated driving systems," in *2018 IEEE 21st International Conference on Intelligent Transportation Systems (ITSC)*, 2018.
- [23] C. Thiemann, M. Treiber, and A. Kesting, "Estimating Acceleration and Lane-Changing Dynamics Based on NGSIM Trajectory Data," *Transportation Research Record: Journal of the Transportation Research Board*, vol. 2088, pp. 90–101, 2008.
- [24] A. Paszke, S. Gross, S. Chintala, *et al.*, "Automatic differentiation in pytorch," in *NIPS-W*, 2017.
- [25] M. Fey, J. E. Lenssen, F. Weichert, and H. Müller, "SplineCNN: Fast geometric deep learning with continuous B-spline kernels," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.