

# Exploring Navigation Maps for Learning-Based Motion Prediction

Julian Schmidt<sup>1,2</sup>, Julian Jordan<sup>1</sup>, Franz Gritschneider<sup>1</sup>, Thomas Monninger<sup>3</sup> and Klaus Dietmayer<sup>2</sup>

**Abstract**—The prediction of surrounding agents’ motion is a key for safe autonomous driving. In this paper, we explore navigation maps as an alternative to the predominant High Definition (HD) maps for learning-based motion prediction. Navigation maps provide topological and geometrical information on road-level, HD maps additionally have centimeter-accurate lane-level information. As a result, HD maps are costly and time-consuming to obtain, while navigation maps with near-global coverage are freely available. We describe an approach to integrate navigation maps into learning-based motion prediction models. To exploit locally available HD maps during training, we additionally propose a model-agnostic method for knowledge distillation. In experiments on the publicly available Argoverse dataset with navigation maps obtained from OpenStreetMap, our approach shows a significant improvement over not using a map at all. Combined with our method for knowledge distillation, we achieve results that are close to the original HD map-reliant models. Our publicly available navigation map API for Argoverse enables researchers to develop and evaluate their own approaches using navigation maps<sup>4</sup>.

## I. INTRODUCTION

Safe autonomous driving in complex traffic scenarios requires to predict the future motion of surrounding traffic participants. These predictions are crucial for tactical decision-making and subsequent motion planning.

Current state-of-the-art approaches for learning-based motion prediction (e.g., [1], [2], [3], [4]) are developed for and evaluated on datasets that are accompanied by High Definition (HD) maps. The core feature of HD maps is their centimeter-accurate information of lane geometry. Obtaining and maintaining HD maps is only possible in a partially automated manner, making HD map creation costly and time-consuming [5], [6]. Fully automated creation fails to work in urban scenarios with complex road topologies [7], [8], scenarios with occlusion [9] and in scenarios where important features, such as the lane markings, are hardly visible [6], [10]. In these cases, manual creation is required. Therefore, it is questionable (i) whether global coverage can ever be achieved with such maps and (ii) whether maintaining HD maps in fast-changing environments is even possible. Especially in urban environments, the routing of lanes can change on a daily basis, e.g., due to accidents, construction sites or road works.

The research leading to these results is funded by the BMWK within the project “KI Delta Learning” (Förderkennzeichen 19A19013A).

<sup>1</sup>Mercedes-Benz AG, Research & Development, Stuttgart, Germany  
 julian.s.j.schmidt@mercedes-benz.com

<sup>2</sup>Ulm University, Institute of Measurement, Control and Microtechnology, Ulm, Germany

<sup>3</sup>Mercedes-Benz Research & Development North America, Sunnyvale, CA, USA

<sup>4</sup><https://github.com/schmidt-ju/argoverse-navmap>

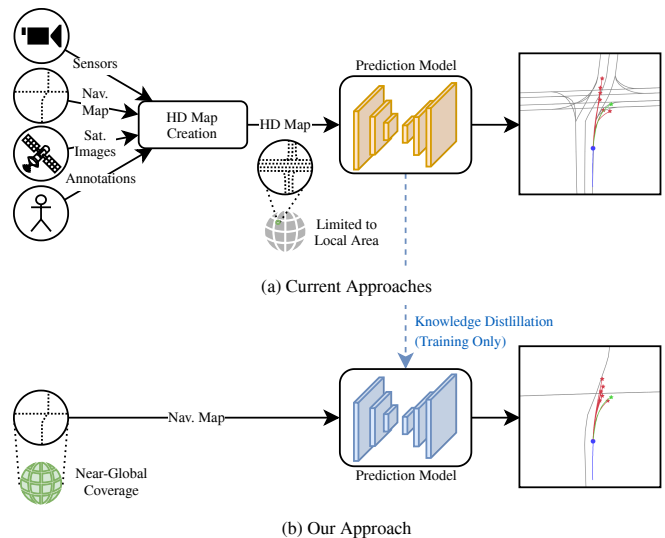


Fig. 1. Comparison of (a) current approaches for motion prediction and (b) our approach. We limit the map input of motion prediction models to navigation maps with near-global coverage. Optionally, if there is an HD map during training available, we propose a teacher-student-based method for knowledge distillation. This method encourages the navigation map-aware model to extract similar embeddings from the navigation map as the teacher from the HD map.

With this paper, we want to give an incentive for a new research direction: Motion prediction with navigation maps instead of HD maps. Navigation maps are an abstraction of HD maps. They do not contain precise geometric information on a lane-level and are therefore more robust with regard to fast-changing environments. Their geometric information is limited to a road-level, meaning multiple lanes are possibly summarized as one road, with varying quality. However, they have the advantage of being widely available around the whole globe. An example of a free navigation map with near-global coverage is the crowdsourced OpenStreetMap (OSM) [11]. To the best of our knowledge, it has never been investigated whether these navigation maps offer a viable alternative to HD maps for motion prediction.

In this work, we therefore investigate the use of navigation maps for the task of motion prediction. We focus on road-level geometry and omit geometric information of lanes, which form the core of every HD map and serve as the main source of information in the field of prediction. Our approach is shown in Fig. 1. On the basis of two state-of-the-art trajectory prediction models, we show that using navigation maps significantly improves prediction performance over not using a map. In the event that an HD map is locally available

for the training dataset, we propose a model-agnostic teacher-student method that uses knowledge distillation to guide the training process, helping to extract relevant information from the navigation map.

In summary, our main contributions are:

- We propose an approach to integrate free navigation maps into existing learning-based motion prediction models.
- We propose a model-agnostic teacher-student method that uses knowledge distillation to exploit locally available HD maps during training. This further improves prediction performance with navigation maps.
- We extensively evaluate the resulting models and our method for knowledge distillation on the publicly available Argoverse Motion Forecasting Dataset [12] and prove a performance close to their HD map-reliant counterparts.

## II. RELATED WORK

This section reviews related work regarding motion prediction with HD and navigation maps. Additionally, existing work related to knowledge distillation in the field of motion prediction is described.

### A. Motion Prediction with High Definition Maps

Motion prediction with HD maps is a constantly evolving topic related to autonomous driving. There are two main groups of approaches for extracting lane information from an HD map in a learning-based manner.

Grid-based approaches [13], [14], [15], [16], [17], [18] rasterize the information of the map, including lanes, into a multilayer Bird’s-Eye View (BEV) grid. In order to maintain the direction of lanes, the heading is most commonly color coded using the HSV color space. Convolutional neural network backbones, such as ResNet [15], [16] or MobileNetV2 [14], [19], are then used to extract information from this grid. The main drawback of grid-based approaches is their dependency on rasterization, which always comes with a loss of information [2].

Vectorization-based approaches do not rely on rasterization. A distinction between two groups of vectorization-based approaches can be made. The first group operates on lane segments. Lane segments are coherent sections of lanes. Different encoders, such as graph neural networks [1], [4] or a combination of 1D convolutional and recurrent neural networks [20], are used to form embeddings of each lane segment’s geometry. These embeddings are then included during the subsequent decoding process to obtain motion predictions. The second group dispenses the forming of lane segment embeddings by subsampling the centerline of each lane segment to a vector-level [2], [3], [21]. Embeddings of these vectors are then again included during the subsequent decoding process.

Motion prediction benchmarks show that vectorization-based approaches typically outperform grid-based approaches by achieving lower prediction errors [12], [22].

### B. Motion Prediction with Navigation Maps

To the best of our knowledge, there are no publications in the field of autonomous driving that extensively evaluate motion prediction with navigation maps and compare it to HD map-reliant approaches. CRAT-Pred [23] is a model specifically developed for trajectory prediction if there is no map available at all. Using freely accessible navigation maps, which is the main idea of our work, is located between map-free approaches and approaches that rely on HD maps.

In the field of behavior planning, there are approaches that include navigation maps. Ort et al. [24], [25] describe a system to plan the future motion of the autonomous vehicle, given a navigation map from OSM and LiDAR scans. The approach mainly relies on handcrafted rules and neglects the influence of surrounding dynamic objects, such as pedestrians and vehicles. Xu et al. [26] describe a transformer-based approach to plan the future motion of the autonomous vehicle. Although they refer to this as trajectory prediction, we believe that their approach is an imitation learning-based planner. Input to the transformer is a BEV representation of the autonomous vehicle’s LiDAR scan and the BEV representation of the route to follow. They annotate the route to follow by hand, but claim that this can also be extracted from navigation maps, such as OSM.

Our approach is not intended to replace existing, sophisticated behavior planning algorithms, but rather to provide high-quality motion predictions of surrounding agents to these algorithms.

### C. Knowledge Distillation for Motion Prediction

Knowledge distillation is the process of transferring the knowledge of one or multiple teacher models (teacher) into a student model (student) [27], [28]. In its original use case, large unwieldy teachers are used for training a small student, achieving comparable or even better performance. However, there are also approaches studying knowledge distillation not from the perspective of model compression, but for improving model performance [29]. For this use-case, teachers and students share the same architecture.

There is few related work regarding knowledge distillation for motion prediction. Neitz et al. [30] apply principles of knowledge distillation to the field of motion prediction, allowing to combine advantages of model-based and model-free prediction techniques. Monti et al. [31] use a teacher and student of the same architecture to obtain a model that is able to predict human motion with a limited amount of observed timesteps. A higher amount of input observations are used for the teacher than for the student. The underlying idea of using superior input representations for the teachers has its origin in the field of video processing [32], [33].

Our approach shares this idea but applies it to a new problem. We exploit an HD map-aware teacher in order to distill knowledge to an architecturally identical student that is only provided a navigation map.

### III. APPROACH

This section describes our approach of obtaining navigation maps and integrating them into learning-based motion prediction models. To exploit HD maps during training, our method of knowledge distillation is used.

#### A. Obtaining Navigation Maps

We obtain navigation maps from OSM [11], which is a crowdsourced open geographic database. A navigation map forms a directed graph  $\mathcal{G} = \{\mathcal{V}, \mathcal{E}\}$  consisting of nodes  $v_i \in \mathcal{V}$  and edges  $e_{j,i} = (v_j, v_i) \in \mathcal{E}$ . The feature vector  $\mathbf{v}_i = (\text{lat}_i, \text{lon}_i)$  of node  $v_i$  consists of its latitude and longitude in global spherical WGS84 coordinates. Edge  $e_{j,i}$  represents a road segment from node  $v_j$  to node  $v_i$ . We limit our navigation maps to road segments that have one of the following car-accessible types: *motorway*, *trunk*, *primary*, *secondary*, *tertiary*, *unclassified*, *residential*, *motorway\_link*, *trunk\_link*, *primary\_link*, *secondary\_link*, *tertiary\_link*, *living\_street*.

Due to OSM being crowdsourced, the available features can vary for each road segment. To show that prediction models benefit from navigation maps even if there are no road segment features available, we do not include any features for road segments in our experiments. Adding such features is part of future work. The geometric information of each road segment  $e_{j,i}$  is given by the coordinates of its source  $v_j$  and destination node  $v_i$ .

#### B. Learning from Navigation Maps

The graph structure of our navigation map representation allows for straightforward adaptation of already existing HD map-reliant prediction models. To integrate a navigation map into such a model, the following steps are required:

- Transform node feature vectors  $\mathbf{v}_i$  (spherical coordinates) of the global navigation map to the local coordinate system the prediction model operates on (most commonly Cartesian coordinates), resulting in  $\tilde{\mathbf{v}}_i = (x_i, y_i)$ .
- Replace queries for lane segments in the HD map with queries for road segments (road segments are represented by edges in the graph  $\mathcal{G}$ ) in the navigation map.
- Replace queries for succeeding and preceding lane segments with queries for succeeding and preceding road segments.

Our API, written for Argoverse and used in our experiments, offers these functionalities and mimics the original API for HD maps, allowing to include navigation maps with ease.

#### C. Knowledge Distillation with HD Map-Aware Teacher

In some cases, HD maps might be available during training but unavailable during inference. This applies to many real-world use-cases, as the recordings of datasets are often limited to a small selection of regional areas, but the prediction models trained with them are to be applied in areas that go beyond this small selection. Obtaining HD maps for the small areas of the recordings is usually feasible, but obtaining and

maintaining them with a similar level of quality globally is way more challenging.

We therefore propose a model-agnostic knowledge distillation method to improve prediction performance with navigation maps if HD maps are available during training. Model-agnostic means that this method can be applied to any already existing and upcoming learning-based prediction model.

Our method consists of a teacher and a student prediction model. An overview is given in Fig. 2. Teacher and student share the same architecture but are provided with different input information. The teacher is provided with information of the HD map and the student is provided with information of the navigation map.

Every learning-based prediction model that relies on map information has a stage where information of the agent to be predicted and the map are fused together into one latent embedding  $\xi$ . Ways to perform this fusion step include attention (e.g., [2], [3]) or concatenation (e.g., [13], [14], [17]). Our method exploits this property of prediction models and uses the latent embedding  $\xi' \in \mathcal{R}^{d_t}$  generated by the fully trained teacher to guide the latent embedding  $\xi \in \mathcal{R}^d, d \geq d_t$  of the student during training. To obtain this behavior, we define the distillation loss as

$$\mathcal{L}_{\text{dist}}(\xi', \xi) = \frac{1}{d_t} \sum_{i=0}^{d_t-1} (\xi'_i - \xi_i)^2. \quad (1)$$

During the training of the student, the total loss is then calculated as

$$\mathcal{L}_{\text{total}} = \alpha \cdot \mathcal{L}_{\text{model}} + \beta \cdot \mathcal{L}_{\text{dist}}, \quad (2)$$

with  $\mathcal{L}_{\text{model}}$  referring to the original loss of the prediction model.  $\alpha$  and  $\beta$  are scaling factors.

By introducing the distillation loss, the student is encouraged to form map-aware embeddings  $\xi$  that are similar to the one obtained by the teacher that has access to the HD map. Accordingly, the exploitation of the HD map during training is carried out directly in the latent space and requires no handcrafted preparation of the input data.

## IV. EXPERIMENTS

This section describes our evaluation on the publicly available Argoverse Motion Forecasting Dataset [12], subsequently referred to as Argoverse. We evaluate our OSM-based approach with two state-of-the-art trajectory prediction models.

#### A. Dataset

In contrast to many other datasets, including the recent Argoverse 2 dataset [34], the Argoverse dataset contains information about the origins of its local Cartesian coordinate systems. As described in this section, this allows to transform the spherical coordinates of navigation maps into the local coordinate system used by Argoverse.




Fig. 2. Overview of our model-agnostic method of using knowledge distillation to exploit HD maps during training: By means of the distillation loss  $\mathcal{L}_{\text{dist}}$ , the latent embeddings  $\xi'$  of the HD map-aware teacher model are used to guide the latent embeddings  $\xi$  of the navigation map-aware student model.

1) *Dataset Properties*: Argoverse is a trajectory prediction dataset consisting of sequences recorded in Miami and Pittsburgh. Each sequence has a duration of five seconds, is sampled with 10 Hz and contains the position of tracked vehicles in the surroundings of an autonomous vehicle. 201k sequences form the training split and 39k form the validation split. Given the first two seconds of a sequence, the goal is to predict the trajectory of the subsequent three seconds of one predefined target vehicle. Two HD maps, one covering relevant areas of Miami and one covering relevant areas of Pittsburgh, are provided in the dataset.

2) *Coordinate System*: All coordinates are given in local Cartesian city coordinates. The following transformations are applied to transform global spherical coordinates  $\mathbf{v}_i = (\text{lat}_i, \text{lon}_i)$  of a navigation map  $\mathcal{G}$  into local city coordinates  $\tilde{\mathbf{v}}_i = (x_i, y_i)$ :

- Transform global spherical coordinates into Universal Transverse Mercator (UTM) coordinates.
- Subtract the origin of the local city coordinate system of the resulting UTM coordinates.

Origins of the local city coordinate systems of Miami and Pittsburgh are given in the original publication [12].<sup>1</sup>

### B. Metrics

We follow the standard protocol for the evaluation of trajectory prediction models and use the minimum Average Displacement Error (minADE), minimum Final Displacement Error (minFDE) and Miss Rate (MR) for single mode ( $k = 1$ ) and multi-modal predictions ( $k = 6$ ).

<sup>1</sup>Miami: 580560.0088 Easting, 2850959.9999 Northing, Zone 17; Pittsburgh: 583710.0070 Easting, 4477259.9999 Northing, Zone 17

minFDE is the lowest Euclidean distance of the  $k$  predicted endpoints and the ground-truth endpoint. minADE is the average Euclidean distance of the trajectory selected for the minFDE calculation and the ground-truth trajectory. MR is the ratio of sequences where none of the  $k$  predicted endpoints is within a radius of 2 m to the ground-truth endpoint. Metrics are averaged over the corresponding dataset split.

### C. Model Implementation Details

We evaluate our approach using two state-of-the-art trajectory prediction models with publicly available code. Model architectures and training protocols are the same as in the original publications. To obtain the geometry of a road segment  $e_{j,i}$ , we linearly interpolate between  $\tilde{\mathbf{v}}_j$  and  $\tilde{\mathbf{v}}_i$  with a step size of 2 m.

We evaluate two different variants of our teacher-student method per prediction model. The first variant uses  $d = d_t$ , meaning that the student is encouraged to approximate the full HD map-aware embeddings  $\xi'$  of the teacher. The second variant uses  $d = 1.5 \cdot d_t$ . This results in a shared embedding:  $d_t$  features are used to approximate  $\xi'$  and  $d_t \cdot 0.5$  features are unguided. Using additional unguided features allows the student to also extract navigation map-specific information. All experiments are performed with  $\alpha = \beta = 1$ .

1) *LaneGCN*: LaneGCN [2] originally uses graph-based encoding of lane geometry. Four attention-based fusion mechanisms, namely Actor-Lane, Lane-Lane, Lane-Actor and Actor-Actor, aggregate information into one latent embedding per agent in a scene. Subsequently, linear residual layers are used to predict multiple trajectories with corresponding confidences.

We use our developed Argoverse API for navigation maps and adapt LaneGCN with our approach described in

TABLE I  
RESULTS OF OUR APPROACH APPLIED TO LANE GCN

Model Variant	HD Map		Nav. Map		Embed. Size		minADE	$k = 1$		$k = 6$		
	Train	Val	Train	Val	$d_t$	$d$		minFDE	MR	minADE	minFDE	MR
Original (Map-Free)					-	128	1.52	3.44	0.54	0.78	1.28	0.14
Ours (w/ Nav. Map, Small Embedding)			✓	✓	-	128	1.43	3.17	0.51	0.76	1.21	<b>0.12</b>
Ours (w/ Nav. Map, Large Embedding)			✓	✓	-	192	1.42	3.16	0.51	0.76	1.21	0.13
Ours (w/ Nav. Map + Teacher-Student, $d = d_t$ )	✓		✓	✓	128	128	1.40	3.12	<b>0.50</b>	0.75	1.20	<b>0.12</b>
Ours (w/ Nav. Map + Teacher-Student, $d = 1.5 \cdot d_t$ )	✓		✓	✓	128	192	<b>1.38</b>	<b>3.07</b>	<b>0.50</b>	<b>0.74</b>	<b>1.17</b>	<b>0.12</b>
Original (w/ HD Map)	✓	✓			-	128	1.34	2.96	0.49	0.71	1.08	0.10

TABLE II  
RESULTS OF OUR APPROACH APPLIED TO HiVT

Model Variant	HD Map		Nav. Map		Embed. Size		minADE	$k = 1$		$k = 6$		
	Train	Val	Train	Val	$d_t$	$d$		minFDE	MR	minADE	minFDE	MR
Original (Map-Free)					-	64	1.48	3.27	0.52	0.76	1.24	0.14
Ours (w/ Nav. Map, Small Embedding)			✓	✓	-	64	1.43	3.14	0.50	0.73	1.15	0.12
Ours (w/ Nav. Map, Large Embedding)			✓	✓	-	96	<b>1.41</b>	3.09	0.49	<b>0.71</b>	1.11	<b>0.11</b>
Ours (w/ Nav. Map + Teacher-Student, $d = d_t$ )	✓		✓	✓	64	64	1.44	3.16	0.49	0.73	1.15	0.12
Ours (w/ Nav. Map + Teacher-Student, $d = 1.5 \cdot d_t$ )	✓		✓	✓	64	96	<b>1.41</b>	<b>3.08</b>	<b>0.48</b>	<b>0.71</b>	<b>1.09</b>	<b>0.11</b>
Original (w/ HD Map)	✓	✓			-	64	1.36	2.98	0.48	0.69	1.03	0.10




Fig. 3. Normalized histogram of the  $\text{minFDE}_{@k=6}$  of different LaneGCN variants.

Section III-B. This results in the graph-based encoding of lane geometry being replaced by a graph-based encoding of road geometry. For our teacher-student method described in Section III-C, we use the agent-wise embeddings  $\xi$  resulting from the Lane-Actor fusion. In LaneGCN, this corresponds to the stage where information of the agent to be predicted and the map are fused together.

2) *HiVT*: HiVT [3] originally uses cross-attention to fuse information of agents with geometric information of surrounding lanes. Predictions and confidences are generated by multilayer perceptrons.

Again, we use our developed Argoverse API for navigation maps and adapt HiVT with our proposed approach, resulting in geometry on a lane-level being replaced by geometry on a road-level. For our teacher-student method, the agent-wise embeddings  $\xi$  resulting from the fusion of agent information with map information via cross-attention are used.

#### D. Quantitative Results

Table I and Table II show the quantitative results of our approach applied to LaneGCN and HiVT on the Argoverse validation set. For both models, the results of our approach are compared with the original models trained map-free and with HD map.

The results of LaneGCN and HiVT both show that integrating navigation maps significantly improves performance compared to the map-free models. Our proposed model-agnostic teacher-student method that uses knowledge distillation to exploit HD maps during training improves performance even more. Especially for LaneGCN, the  $d = 1.5 \cdot d_t$  variant outperforms all other navigation map-reliant variants. The same statement, although with slightly less improvement, applies to HiVT.

For HiVT, it is observable that the difference between the variant without ( $d = 96$ ) and with our teacher-student method ( $d = 1.5 \cdot d_t, d_t = 64$ ) is small. However, the performance of both is close to the original HD map-reliant variant.

In general, the results of the navigation map-reliant variants are close to the results of the original HD map-reliant variant. For instance, the  $\text{minFDE}_{@k=6}$  difference between our best navigation map-reliant LaneGCN variant and the original HD map-reliant variant for a 3s prediction horizon is 9 cm on average. For HiVT, our teacher-student method even results in an identical  $\text{MR}_{@k=1}$  as the original HD map-reliant variant.

Fig. 3 shows a normalized histogram of the  $\text{minFDE}_{@k=6}$  of LaneGCN beyond the Miss Rate distance of 2m. Kernel density estimation is used to obtain a smoothed trend. One distinct observation is that our navigation map-reliant variant ( $d = 1.5 \cdot d_t$ ) consistently leads to less sequences with high

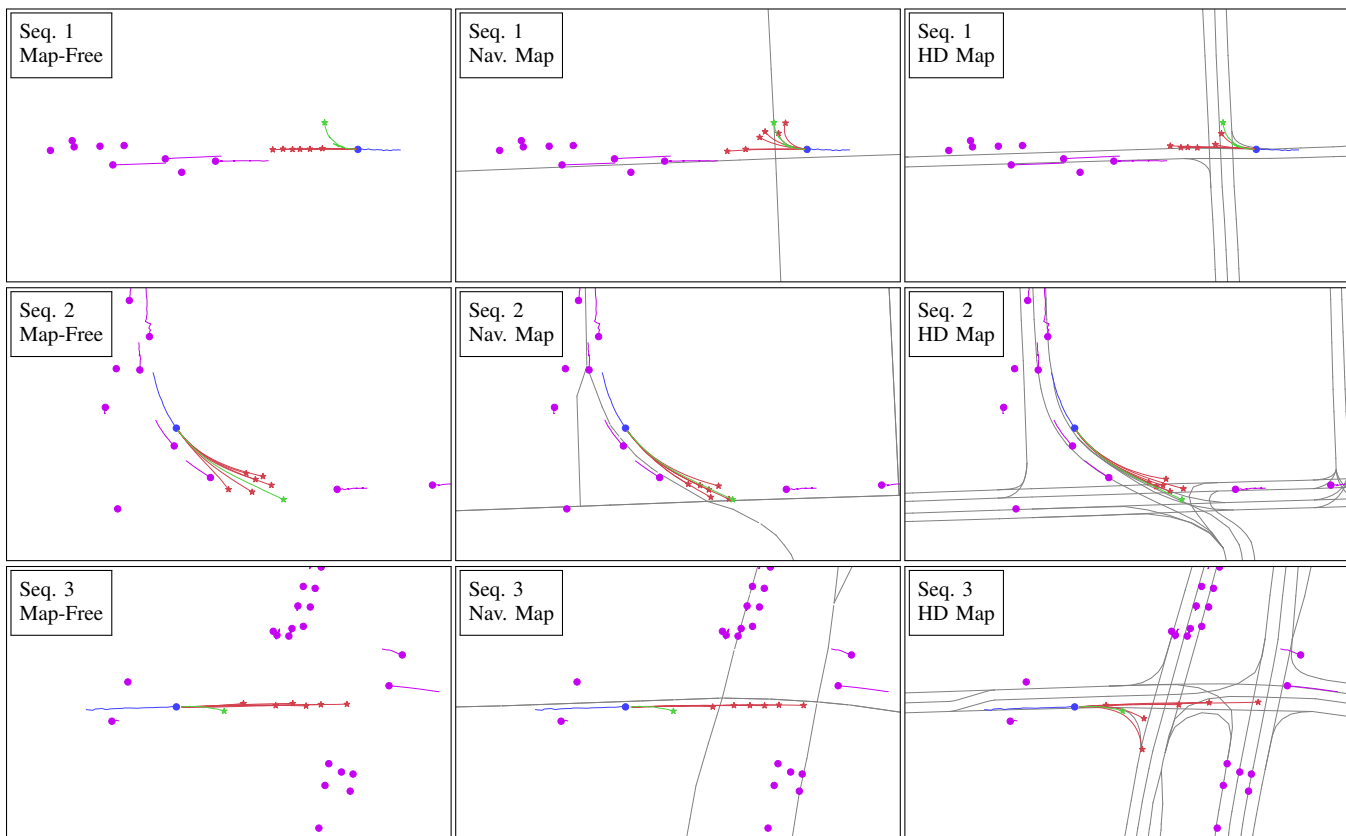


Fig. 4. Exemplary sequences of LaneGCN on the Argoverse validation set, trained and validated map-free (left), with our navigation map-aware teacher-student method  $d = 1.5 \cdot d_t$  (center) and with the HD map (right). Three sequences are shown. The past observed trajectory of the vehicle to be predicted is colored in blue, the ground-truth future trajectory in green. Predictions are colored in red. Past trajectories of other vehicles are colored in purple. The map-information that is used by the corresponding model is colored in gray.

prediction errors than the map-free variant. This suggests that the integration of navigation maps reduces sequences in which the intent of the predicted agent was not correctly predicted at all (e.g., failing to predict a turn or a lane change maneuver).

### E. Qualitative Results

Our qualitative results are generated with different LaneGCN variants on the Argoverse validation set. We reviewed over 400 sequences and selected three interesting ones for visualization in Fig. 4. While there are many positive examples, we only observed two clear negative examples. One of them is illustrated as the third sequence in the figure. For each sequence, the results of the map-free variant (left), the adaptation with our navigation map-reliant teacher-student method (center) and the original HD map-reliant variant (right) are shown. Color codes are given in the caption of the figure.

The first sequence shows a turn maneuver. While the map-free variant fails to predict this turn maneuver, both map-reliant variants are able to correctly predict both possible intents of the vehicle. The second sequence is another positive example. Both map-reliant variants are able to precisely predict the movement along the curved road. Due to having no information about the underlying road or lane topology

and geometry, the predictions of the map-free variant are less precise. The third sequence illustrates a negative example. Both, the map-free and the navigation map-reliant variant, fail to predict the turn maneuver. One possible explanation for failure of the navigation map-reliant variant is the acute angle of the road into which the turn is made.

## V. CONCLUSION

In this paper, we explore the use of navigation maps as an alternative to HD maps for motion prediction. By implementing the proposed approach for two learning-based trajectory prediction models, we prove a significant improvement in prediction performance compared to using no map. In combination with our model-agnostic knowledge distillation method that exploits HD maps during training, results that are close to the HD map-reliant counterparts are achieved. The publicly available source code of our navigation map API for Argoverse enables other researchers to develop and evaluate their own navigation map-based approaches for motion prediction with ease.

It remains to be investigated whether similar results are achievable using navigation maps in other HD map-reliant application areas beyond motion prediction, for instance traffic scene reasoning [35].

## REFERENCES

- [1] J. Gao, C. Sun, H. Zhao, Y. Shen, D. Anguelov, C. Li, and C. Schmid, "Vectormet: Encoding hd maps and agent dynamics from vectorized representation," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 11 522–11 530.
- [2] M. Liang, B. Yang, R. Hu, Y. Chen, R. Liao, S. Feng, and R. Urtasun, "Learning lane graph representations for motion forecasting," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 541–556.
- [3] Z. Zhou, L. Ye, J. Wang, K. Wu, and K. Lu, "Hivt: Hierarchical vector transformer for multi-agent motion prediction," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 8823–8833.
- [4] J. Wang, T. Ye, Z. Gu, and J. Chen, "Ltp: Lane-based trajectory prediction for autonomous driving," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 17 134–17 142.
- [5] L. Mi, H. Zhao, C. Nash, X. Jin, J. Gao, C. Sun, C. Schmid, N. Shavit, Y. Chai, and D. Anguelov, "Hdmapgen: A hierarchical graph generative model of high definition maps," in *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021, pp. 4227–4236.
- [6] S. He and H. Balakrishnan, "Lane-level street map extraction from aerial imagery," in *2022 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2022, pp. 1496–1505.
- [7] A. Joshi and M. R. James, "Generation of accurate lane-level maps from coarse prior maps and lidar," *IEEE Intelligent Transportation Systems Magazine*, vol. 7, no. 1, pp. 19–29, 2015.
- [8] N. Homayounfar, J. Liang, W. Ma, J. Fan, X. Wu, and R. Urtasun, "Dagmapper: Learning to map by discovering lane topology," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 2911–2920.
- [9] Y. Zhou, Y. Takeda, M. Tomizuka, and W. Zhan, "Automatic construction of lane-level hd maps for urban scenes," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 6649–6656.
- [10] G. Mátyus, S. Wang, S. Fidler, and R. Urtasun, "Hd maps: Fine-grained road segmentation by parsing ground and aerial images," in *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 3611–3619.
- [11] OpenStreetMap contributors, "Planet dump retrieved from <https://planet.osm.org>," <https://www.openstreetmap.org>, 2017.
- [12] M.-F. Chang, J. Lambert, P. Sangkloy, J. Singh, S. Bak, A. Hartnett, D. Wang, P. Carr, S. Lucey, D. Ramanan, and J. Hays, "Argoverse: 3d tracking and forecasting with rich maps," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 8740–8749.
- [13] N. Djuric, V. Radosavljevic, H. Cui, T. Nguyen, F.-C. Chou, T.-H. Lin, N. Singh, and J. Schneider, "Uncertainty-aware short-term motion prediction of traffic actors for autonomous driving," in *2020 IEEE Winter Conference on Applications of Computer Vision (WACV)*, 2020, pp. 2084–2093.
- [14] J. Strohbeck, V. Belagiannis, J. Müller, M. Schreiber, M. Herrmann, D. Wolf, and M. Buchholz, "Multiple trajectory prediction with deep temporal and spatial convolutional neural networks," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1992–1998.
- [15] T. Phan-Minh, E. C. Grigore, F. A. Boulton, O. Beijbom, and E. M. Wolff, "Covernet: Multimodal behavior prediction using trajectory sets," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020, pp. 14 062–14 071.
- [16] Y. Chai, B. Sapp, M. Bansal, and D. Anguelov, "Multipath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction," in *Proceedings of the Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, vol. 100. PMLR, 2020, pp. 86–99.
- [17] T. Gilles, S. Sabatini, D. Tsishkou, B. Stanculescu, and F. Moutarde, "Home: Heatmap output for future motion estimation," in *2021 IEEE International Intelligent Transportation Systems Conference (ITSC)*, 2021, pp. 500–507.
- [18] A. Kamenev, L. Wang, O. B. Bohan, I. Kulkarni, B. Kartal, A. Molchanov, S. Birchfield, D. Nistér, and N. Smolyanskiy, "Predictionnet: Real-time joint probabilistic traffic prediction for planning, control, and simulation," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 8936–8942.
- [19] 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 *2019 IEEE International Conference on Robotics and Automation (ICRA)*, 2019, pp. 2090–2096.
- [20] T. Gilles, S. Sabatini, D. Tsishkou, B. Stanculescu, and F. Moutarde, "Gohome: Graph-oriented heatmap output for future motion estimation," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 9107–9114.
- [21] W. Zeng, M. Liang, R. Liao, and R. Urtasun, "Lanercnn: Distributed representations for graph-centric motion forecasting," in *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2021, pp. 532–539.
- [22] S. Ettinger, S. Cheng, B. Caine, C. Liu, H. Zhao, S. Pradhan, Y. Chai, B. Sapp, C. R. Qi, Y. Zhou, Z. Yang, A. Chouard, P. Sun, J. Ngiam, V. Vasudevan, A. McCauley, J. Shlens, and D. Anguelov, "Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset," in *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2021, pp. 9710–9719.
- [23] J. Schmidt, J. Jordan, F. Gritschneider, and K. Dietmayer, "Crat-pred: Vehicle trajectory prediction with crystal graph convolutional neural networks and multi-head self-attention," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 7799–7805.
- [24] T. Ort, L. Paull, and D. Rus, "Autonomous vehicle navigation in rural environments without detailed prior maps," in *2018 IEEE International Conference on Robotics and Automation (ICRA)*, 2018, pp. 2040–2047.
- [25] T. Ort, K. Murthy, R. Banerjee, S. K. Gottipati, D. Bhatt, I. Gilitschenski, L. Paull, and D. Rus, "Maplite: Autonomous intersection navigation without a detailed prior map," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 556–563, 2020.
- [26] J. Xu, L. Xiao, D. Zhao, Y. Nie, and B. Dai, "Trajectory prediction for autonomous driving with topometric map," in *2022 IEEE International Conference on Robotics and Automation (ICRA)*, 2022, pp. 8403–8408.
- [27] C. Bucilua, R. Caruana, and A. Niculescu-Mizil, "Model compression," in *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, ser. KDD '06. New York, NY, USA: Association for Computing Machinery, 2006, pp. 535–541.
- [28] G. Hinton, O. Vinyals, and J. Dean, "Distilling the knowledge in a neural network," 2015.
- [29] T. Furlanello, Z. C. Lipton, M. Tschannen, L. Itti, and A. Anandkumar, "Born-again neural networks," in *Proceedings of the 35th International Conference on Machine Learning, ICML 2018*, ser. Proceedings of Machine Learning Research, J. G. Dy and A. Krause, Eds., vol. 80. PMLR, 2018, pp. 1602–1611.
- [30] A. Neitz, G. Parascandolo, and B. Schölkopf, "A teacher-student framework to distill future trajectories," in *9th International Conference on Learning Representations, ICLR 2021*, 2021.
- [31] A. Monti, A. Porrello, S. Calderara, P. Coscia, L. Ballan, and R. Cucchiara, "How many observations are enough? knowledge distillation for trajectory forecasting," in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022, pp. 6553–6562.
- [32] S. Bhardwaj, M. Srinivasan, and M. M. Khapra, "Efficient video classification using fewer frames," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 354–363.
- [33] A. Porrello, L. Bergamini, and S. Calderara, "Robust re-identification by multiple views knowledge distillation," in *Computer Vision – ECCV 2020*, A. Vedaldi, H. Bischof, T. Brox, and J.-M. Frahm, Eds. Cham: Springer International Publishing, 2020, pp. 93–110.
- [34] B. Wilson, W. Qi, T. Agarwal, J. Lambert, J. Singh, S. Khandelwal, B. Pan, R. Kumar, A. Hartnett, J. K. Pontes, D. Ramanan, P. Carr, and J. Hays, "Argoverse 2: Next generation datasets for self-driving perception and forecasting," in *Proceedings of the Neural Information Processing Systems Track on Datasets and Benchmarks (NeurIPS Datasets and Benchmarks 2021)*, 2021.
- [35] T. Monninger, J. Schmidt, J. Rupperecht, D. Raba, J. Jordan, D. Frank, S. Staab, and K. Dietmayer, "Scene: Reasoning about traffic scenes using heterogeneous graph neural networks," *IEEE Robotics and Automation Letters*, vol. 8, no. 3, pp. 1531–1538, 2023.