

Convolutions for Spatial Interaction Modeling

Zhaoen Su, Chao Wang, David Bradley, Carlos Vallespi-Gonzalez
Carl Wellington, Nemanja Djuric

{suzhaoen, chao.wang, dbradley, cvallespi, cwellington, ndjuric}@aurora.tech

Abstract

In many different fields interactions between objects play a critical role in determining their behavior. Graph neural networks (GNNs) have emerged as a powerful tool for modeling interactions, although often at the cost of adding considerable complexity and latency. In this paper, we consider the problem of spatial interaction modeling in the context of predicting the motion of actors around autonomous vehicles, and investigate alternatives to GNNs. We revisit 2D convolutions and show that they can demonstrate comparable performance to graph networks in modeling spatial interactions with lower latency, thus providing an effective and efficient alternative in time-critical systems. Moreover, we propose a novel interaction loss to further improve the interaction modeling of the considered methods.

1. INTRODUCTION

Interactions or relations between objects are critical for understanding both the individual behaviors and collective properties of many systems. Conceptually, these interactions can be modeled with graph structures that comprise a set of objects (nodes) and their relationships (edges). By applying deep learning techniques, graph neural networks (GNNs) have demonstrated great expressive power in modeling interactions in various fields, including physical science [3, 12, 30, 33], social science [17, 23], knowledge graphs [16], and other research areas [20, 26, 28, 37].

Some of the interactions strongly depend on geometries, such as the Euclidean distance and relative directions between objects, which in this work is referred to as spatial interactions. One problem where spatial interaction is critical is motion forecasting, a key task in the fields of computer vision, robotics in general and autonomous driving (AD) in particular. Specifically, anticipating the future movements of an object requires understanding not only its history, but also the object's interactions with other objects and its environment. These interactions strongly depend on relative spatial features between objects, such as their relative locations, orientations, and velocities.

Graphs have achieved success in modeling spatial interaction [22, 31, 32, 34]. Features of individual objects are typically encoded into attributes of graph nodes, and the graph edges are built by passing node attributes and the relative geometries of the node pair through a mapping function. GNNs follow a message passing scheme, where each node aggregates features of its neighboring nodes to compute its new node attributes. These approaches have two characteristics, as seen in the experimental section: (1) the relative spatial features are not represented implicitly in the graph and need to be handcrafted into the graph edge features; (2) even a single iteration of GNN may be slower than convolutional neural networks (CNNs), which makes GNNs less suitable for applications in fields such as AD where fast inference is safety-critical.

Alternatively, data structures for 2D or 3D convolutional operations are presented in common grid forms, such as through voxelization in 3D, rasterization in 2D bird's-eye view (BEV), or as intermediate CNN features. Importantly, spatial relations are intrinsically represented in these Euclidean space. Thus, they theoretically allow spatial relations between objects to be learned by CNNs with sufficiently large receptive fields [11]. In other words, CNNs have the potential to model spatial interactions. However, even though deep CNN backbones with large receptive fields are widely utilized in trajectory forecasting models, research has shown that adding a GNN after a CNN backbone can still improve interaction modeling [4, 31, 32]. This suggests the CNN backbones often do not fulfill their theoretical potential in modeling spatial interactions between the traffic actors.

In this work, we consider spatial interaction modeling through 2D convolutions and compare them to GNNs within the context of motion forecasting for AD. A key determinant of future motion for other drivers is the avoidance of collisions, which represents a critical interaction that we model explicitly. Collisions can be approximated as geometric overlapping, which provides unambiguous definitions for interaction metrics. We evaluate the methods on large-scale real-world AD data to draw general conclusions. Our contributions are summarized below:

- we identify three components to facilitate modeling spatial interaction with convolutions: (1) large actor-centric interaction region, (2) projecting feature maps into the actor’s frame of reference, and (3) aggregation of per-actor feature maps using convolutions;
- we perform empirical studies to compare interaction modeling using convolutions and graphs, and find that (1) CNNs can perform similarly to or better than GNNs; (2) adding the CNN can considerably improve interaction modeling even when a GNN is used; (3) adding a GNN demonstrates only minor additional gain when the convolutional approach is already used.
- we propose and study a novel interaction loss.

2. RELATED WORK

2.1. Motion forecasting

There exists a significant body of work on forecasting the motion of traffic actors. An input to the forecasting models can be a sequence of past actor states such as positions, headings, or velocities [7, 8, 10, 13, 19, 31], or a sequence of raw sensor data such as LiDAR or radar returns [5, 25] where joint object detection and motion forecasting are performed in an autonomous vehicle’s (AV) frame of reference. While the latter approach may accelerate inference and joint learning by sharing common CNN features among all actors, these single-stage models could benefit from actor-centric features. Two-stage models [4, 9] address this issue by using a first stage to detect the actors and extract features, and then adding a second stage in the frame of reference of detected actors. The two stages are then learned jointly in an end-to-end fashion. The interaction modeling study in this paper adapts a two-stage architecture. Note that the designs used in the study, including rotated region of interest (RROI) [27] and actor-centric design [4, 9, 10], have been developed and applied in previous research in a context different from interaction modeling. However, our empirical study demonstrates that utilizing these ideas allows convolutions to effectively model spatial interaction as well.

2.2. Interaction modeling

GNNs have recently been applied to explicitly express interactions in motion forecasting. NRI [22] models the interaction between actors by using GNNs to infer interactions while simultaneously learning dynamics. Vector-Net [13] and CAR-Net [32] model actor-context interactions. Closely related to our work, SpaGNN [4] is also a two-stage detection-and-forecasting model that builds a graph for vehicles in the second stage to model vehicle-vehicle interaction. The GNN models used for comparison in the study of this paper follow the same design.

Beyond graph models, grid-based spatial relations have been explored using social pooling approaches [1, 7, 15], where pooling is used to capture the impact of surrounding actors in the recurrent architecture. In social-LSTM [1, 15], the LSTM cell receives pooled spatial hidden states from the LSTM cells of neighbors that are embedded into a grid. Besides the parameter-free pooling, convolutional layers have also been explored [7]. By contrast, our proposal is fully convolutional. Moreover, these approaches pool the spatial context of interacting actors while excluding the actor itself, thus the actor-context interaction is not directly modeled in the process.

2.3. Interaction metrics

It is interesting to note that while various techniques have been developed to model spatial interaction, most prior work reports motion forecasting displacement errors. As shown in this study, reducing displacement errors does not necessarily indicate improvement in interaction modeling for a motion forecasting task. An alternative metric that can more explicitly indicate the level of interaction modeling is to measure whether vehicle motion forecasts incorrectly predict overlap with other vehicles [4, 31]. In this work, we also propose vehicle-obstacle overlap rate within motion forecasts as another measure for interaction modeling.

3. METHODOLOGY

In this section we formulate the motion forecasting problem, followed by a discussion of two approaches to interaction modeling: implicitly through 2D convolutions and explicitly through graphs. Fig. 1 illustrates the architectures of the considered end-to-end models that jointly solve tasks of object detection and motion forecasting, taking BEV representation of the sensor data as an input and outputting both object detections and their future trajectories. We emphasize that we purposefully choose a commonly used input representation, neural network design, and loss functions in order to focus on understanding the interaction modeling aspect of these approaches. Moreover, to quantify the analysis we limit our discussion to vehicle actors (see Appendix for analysis of other actor types).

3.1. Problem formulation

Given input data comprising the past and current information of V interacting actors and the environment, a model outputs their current and future states \mathbf{x} represented as $\mathcal{X}_{0:H} = \{\mathbf{x}_t^v, v = 1, \dots, V, t = T_0, \dots, T_H\}$. As mentioned previously, our study considers raw sensor data as an input to the model. Following the joint detection and forecasting architecture [5, 9], we encode the sensor data by voxelizing and stacking a sequence of current and P past

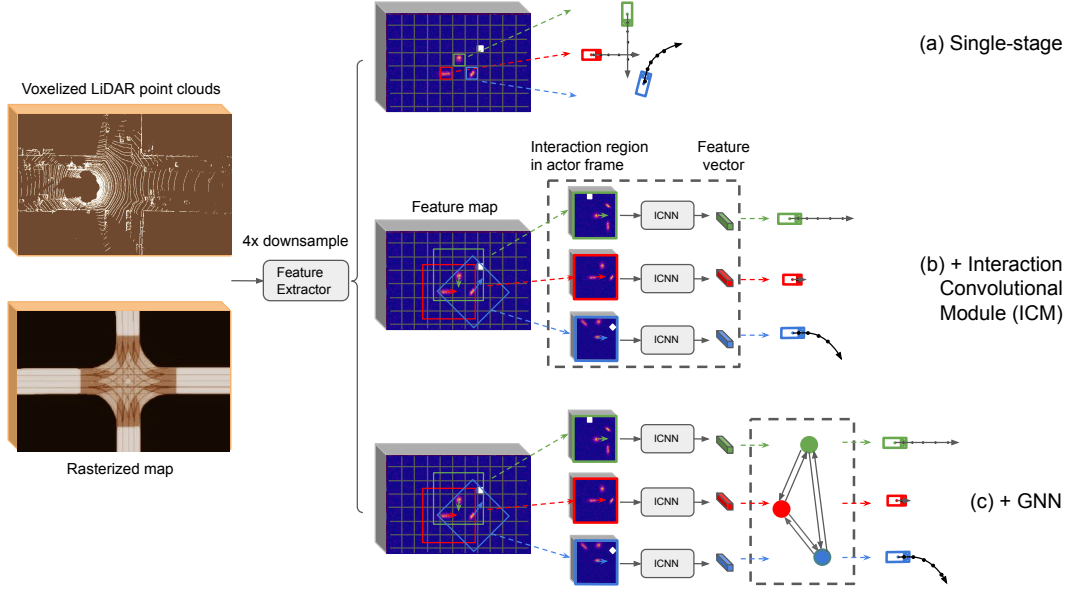


Figure 1. Three model architectures in a scene illustrated with three vehicle actors and one obstacle (denoted by the white spot). All models share the same first-stage design shown from left to middle: input is a BEV raster image comprising past and current point clouds and a semantic map in the AV frame. Through a CNN feature extractor we obtain a $4\times$ downsampled feature map in the AV frame. (a) Single-stage baseline: Object detection and trajectory forecasting are performed at a pixel-level. (b) Adding the proposed Interaction Convolutional Module (ICM). For each actor we define an interaction region (IR) in the actor frame that is used to crop an area from the feature map. Through the weight-sharing interactive CNN (ICNN) a feature vector is aggregated for each actor and then utilized to predict a future trajectory in its frame. (c) Adding GNN into the architecture shown in (b).

LiDAR point clouds around AV at time T_0 in BEV representation, as well as rasterizing semantic map that provides an additional environmental prior, which are used as the model input. The 2D detection at time T_0 for each actor is parameterized by a bounding box represented as $(c_x, c_y, \cos \theta, \sin \theta, w, l)$, denoting the x and y coordinates of the actor’s centroid, the cosine and sine of its heading angle, and the width and length of the box, respectively. Assuming rigid SE2 transformations, future trajectories can be represented as a sequence of tuples $(c_{xt}, c_{yt}, \cos \theta_t, \sin \theta_t)$, with $t \in \{T_1, \dots, T_H\}$ [35].

3.2. Feature extraction and loss functions

As illustrated in Fig. 1a, the first stage of the joint model detects objects and extracts features. From the input BEV raster, a $4\times$ downsampled feature map is extracted by a deep CNN that follows common design (see Appendix for the complete network design). It consists of 3 operations: (1) convolutional block (ConvB) including a convolution (kernel size 3×3), batch normalization, and ReLU optionally; (2) ResNet v2 block (ResB) [18]; and (3) upsampling using bi-linear interpolation. Features are processed at multiple scales to provide larger receptive fields for capturing wider context and past motion of the actors.

Following the computation of the BEV feature map,

classification and regression are performed on the 1D feature vector for each grid cell. Through a fully-connected (FC) layer and a softmax function, we obtain the likelihood p_c of existence of a vehicle actor whose center is located in the cell c . We use focal loss ℓ_f [24] to address the foreground/background imbalance. Through a separate FC layer, the network at the same time regresses the detection bounding boxes \mathcal{X}_0 . The centroid and heading are relative to the cell center and the AV heading, respectively. Then, the first-stage detection loss is given as follows (the hat-notation $\hat{*}$ indicates the ground-truth targets)

$$\begin{aligned} \mathcal{L}_{det} = & \sum_{c \in all} \ell_f(\hat{p}_c, p_c) + \sum_{v \in veh} \left(\ell_1(\hat{l}^v - l^v) \right. \\ & + \ell_1(\hat{w}^v - w^v) + \ell_1(\hat{c}_{x0}^v - c_{x0}^v) + \ell_1(\hat{c}_{y0}^v - c_{y0}^v) \\ & \left. + \ell_1(\cos \hat{\theta}_0^v - \cos \theta_0^v) + \ell_1(\sin \hat{\theta}_0^v - \sin \theta_0^v) \right), \end{aligned} \quad (1)$$

where *all* and *veh* represent all grid cells and vehicle foreground grid cells, respectively, \hat{p}_c equals 1 for foreground cells and 0 otherwise, ℓ_1 is smooth- L_1 loss (with the transition value set to 0.1).

In addition to the detection loss, end-to-end models also optimize for the prediction loss that is only applied to future waypoints of the actors. Moreover, we model the multimodality of the predictions [9] by classifying three modes

for each actor (i.e., turning left, turning right, or going straight), where a separate trajectory is regressed for each mode along with the corresponding mode probability [6] based on the focal loss. In addition, regression loss is applied only to the trajectory mode that is closest to the observed trajectory. Then, the prediction loss is given as

$$\begin{aligned} \mathcal{L}_{pred} = & \sum_{v \in veh} \sum_{m=1}^{M=3} \left(\ell_f(\hat{p}_m^v, p_m^v) \right. \\ & + \frac{1_{\hat{m}=m}}{H} \sum_{t=T_1}^{T_H} \left(\ell_1(\hat{c}_{xt}^v - c_{xmt}^v) + \ell_1(\hat{c}_{yt}^v - c_{ymt}^v) \right. \\ & \left. \left. + \ell_1(\cos \hat{\theta}_t^v - \cos \theta_{mt}^v) + \ell_1(\sin \hat{\theta}_t^v - \sin \theta_{mt}^v) \right) \right), \end{aligned} \quad (2)$$

where p_m^v denotes probability of m -th trajectory mode of actor v , 1_c is an indicator function equaling 1 if the condition c holds and 0 otherwise, and \hat{m} indicates an index of the mode closest to the ground truth. Future centroids and headings are relative to the cell center and the AV heading, respectively (see Fig. 1a), while they are in the actor frames in the two-stage models (see Fig. 1b-c). Then, \mathcal{L}_{det} and \mathcal{L}_{pred} can be optimized together in a joint training.

For single-stage models the detection and prediction values are both optimized in the first stage (Fig. 1a). On the other hand, when the first stage serves as a part of the two-stage architecture (Fig. 1b-c), \mathcal{L}_{det} is optimized as a part of the first-stage output while \mathcal{L}_{pred} is optimized in the second stage, discussed in the remainder of this section.

3.3. Interaction using convolutions implicitly

In the previous section we discussed the first-stage feature extraction, that computes per-actor grid features which are then used as an input to the second-stage models to predict future motion. In this section we discuss how to compute the per-actor features better at capturing interactions:

- To capture relationship to nearby actors for the actor for whom the future trajectories are predicted (called the actor of interest), an input of the forecasting module can be a region covering the interacting actors and objects on the feature map, instead of only using the feature pixel. For the traffic use-case this interaction region (IR) should cover the area within which the objects should be paid attention to. Our results show that for vehicle actors, a large region ahead of the vehicle provides good context to model interaction.
- To overcome rotational variance of convolutions, instead of cropping the IR features in the coordinate frame of the original BEV grid whose orientation is determined by AV, we define IR in the frame of the actor of interest (i.e., actor frame), in which the output trajectories are also defined (commonly referred to

as RROI [27]). Our results confirm the importance of rotational invariance in modeling interactions.

- To effectively propagate non-local information of the interacting actors to the actor of interest, we can use an interactive CNN (ICNN) consisting of a few downsampling convolutional layers that eventually condense an IR comprising the actor of interest itself, its surrounding actors, and the environment, into a feature vector used as the final feature for this actor.

As mentioned, the actor-centric feature map and the RROI techniques have been utilized in a number of applications [4,9,10], where it was found to lower displacement errors in trajectory forecasting tasks. In this paper we demonstrate that, by combining these ideas, convolutions are effective in modeling spatial interactions as well. Moreover, as shown by our experiments, by varying the parameters of these ingredients one can control the level of interaction modeling, providing further evidence that spatial interactions can be effectively captured by convolutions.

The implementation of these three components are illustrated in the dashed box in Fig. 1b, which we refer to as the interaction convolutional module (ICM). For each actor we define a square IR around it, which is then used to crop actor-centric features from the global feature map using bilinear interpolation. We vary the size, orientation, and the position of the actor in the IR to study their effects on the performance of interaction modeling (e.g., in the extreme case where the IR has no area, the cropped feature is just the feature pixel on the feature map). We choose a square IR to simplify the discussion. The length of the square side is referred to as *IR size* in the following discussion. Similarly, the ICNN module always consists of six ConvBs and one ResB to gradually reduce the cropped feature map to a 1D feature vector f_c (e.g., if the crop size is 32×32 , setting the strides of the last five ConvBs to 2 yields a 1D vector; see Appendix for detailed discussion on crop sizes and ICNN design). The final multimodal classification and future trajectory regression in the actor frame are obtained from this 1D vector via a single FC layer, one for each task.

3.4. Interaction using graphs explicitly

The purely convolutional approach described in the previous section provides implicit interaction modeling. To explicitly account for interactions, a common approach is the use of GNNs, discussed in this section. As there exist many variants, we choose one of the more general approaches, the message passing neural network [14,36], which has also been adapted to the motion forecasting problem [4].

Indicated by the dashed box in Fig. 1c, a fully connected graph comprises all of the V actors (represented as nodes), with bi-directional edges between every two actors. The




Figure 2. Schematic of interaction loss. The actor (blue) is approximated with 3 costing circles (green), with minimal distances (black) to an obstacle (grey) and resulting gradients (red).

feature attribute n_i of the i -th node is initialized by

$$n_i^0 = \text{MLP}_{init}(f_{ci}), \quad (3)$$

where f_{ci} is the final feature vector of the i -th actor computed in the previous section. All multi-layer perceptrons (MLPs) in this GNN have two layers. The message passing at the k -th iteration via edge from node j to i is given by

$$m_{j \rightarrow i}^k = \text{MLP}_e^k([n_i^k, \text{rel}_{j \rightarrow i}, n_j^k, \text{rel}_{i \rightarrow j}]), \quad (4)$$

where $[\cdot]$ denotes concatenation. Unlike the implicit convolutional approach in the previous section where the relative spatial relation of actors are intrinsically represented within the crop, spatial relationships are additionally required in a graph representation. The relative geometric feature $\text{rel}_{j \rightarrow i}$ consisting of the coordinates and heading of actor j in the frame of actor i , is computed as

$$\text{rel}_{j \rightarrow i} = \text{MLP}_{\text{rel}}([x_{j \rightarrow i}, y_{j \rightarrow i}, \cos \theta_{j \rightarrow i}, \sin \theta_{j \rightarrow i}]). \quad (5)$$

All of the messages sent to the i -th graph node are aggregated by a max-pooling operation, denoted as

$$m_i^k = \text{Pool}_j(m_{j \rightarrow i}^k). \quad (6)$$

Finally, the node attribute is updated with a Gated Recurrent Unit (GRU) [4, 14, 36] whose hidden state is n_i^k and the input is m_i^k ,

$$n_i^{k+1} = \text{GRU}(n_i^k, m_i^k). \quad (7)$$

In general, the update iterates for K times. Finally, multi-modal classification and future trajectories for the actor are computed from n_i^{k+1} , as discussed in Section 3.3.

3.5. Interaction loss

In this section we introduce a novel interaction loss to improve interaction awareness of the model, which directly penalizes predicted forecast of an actor that overlap with static traffic objects (defined as objects with speed less than 0.2m/s). Traffic objects comprise objects that a vehicle should avoid, including vehicles, cyclists, pedestrians, construction fences, etc. At each prediction horizon, the predicted actor is approximated with 3 inscribed costing circles, as illustrated in Fig. 2. The loss is then computed as

$$\mathcal{L}_{col} = \frac{1}{3VH} \sum_v \sum_n \sum_{t=1}^H \sum_{l=1}^{L=3} \max(0, R_{vtl} - d_{vntl}), \quad (8)$$

where V , N , H , and L are the numbers of actors, non-moving obstacles, prediction time horizons, and costing circles, respectively. R_{vtl} is a radius of a costing circle (determined by the size of a ground-truth bounding box), while d_{vntl} is a signed minimum distance between the l -th costing circle center of the v -th actor and the n -th obstacle bounding box at time t . The distance is negative when the center is inside the obstacle's bounding box.

Note that the loss only considers overlaps between predicted trajectories and the ground-truth bounding boxes of static obstacles. Moving actors may have multimodal trajectory distributions, and it can be unclear when an overlap between the trajectories of two moving actors should be penalized by the loss. In summary, when the costing circles overlap with an obstacle bounding box, the interaction loss would only back-propagate gradients through the predicted centroid and heading. The loss is added to the prediction loss \mathcal{L}_{pred} where it is applied to the \hat{m} -th predicted trajectory, and optimized jointly in the end-to-end training.

4. EXPERIMENTS

Input and output. The considered area is of size $150 \times 100 \times 3.2\text{m}$, centered at AV and discretized as a $960 \times 640 \times 16$ grid into which the LiDAR sweep information is encoded. The input contains 10 LiDAR sweeps collected at 0.1s interval, as well as a semantic HD map from the current timestamp. The models detect the vehicle actors at the current time step and forecast their trajectories at future time horizons $t \in \{0.1, 0.2, \dots, 4.0\text{s}\}$. Non-maximum suppression (NMS) [29] with Intersection over Union (IoU) threshold set at 0.1 is applied in order to eliminate duplicate detections.

Metrics. The studies are focused on prediction accuracy and interaction performance. IoU threshold for object detection matching is set at 0.5. We observe that the detection performance changes little in all of the considered models reported in the paper, with average precision at 94.0 ± 0.4 . Furthermore, we ensure equal numbers of trajectories are considered in the metrics by adjusting the detection probability threshold at a fixed recall of 0.8 [4]. Each actor has 3 predicted trajectory modes, and we assign the trajectory of the most probable mode to the actor in the following metric computation. We use displacement error (DE) at 4s to measure the prediction accuracy, averaged over all actors.

To quantify the interaction performance of the models we consider two overlap metrics in our experiments (additional results of other metrics are provided in Appendix):




Figure 3. Effects of the ICM components: interaction region (IR), interaction CNN (ICNN), and actor frame (AF). **Extractor** is the single-stage model; **+IR+ICNN** represents the two-stage model that defines IRs in the AV frame; **+ICM** (i.e., +IR+ICNN+AF) is the proposed two-stage model that defines IRs in the actor frame. All IRs have a fixed 5:1 front-to-back ratio, unless specified otherwise. Inset: models with a fixed IR size of 60m and varied front-to-back ratios.

- Actor-actor overlap rate is the percentage of predicted trajectories of detected actors overlapping with predicted trajectories of other detected actors.
- Actor-static overlap rate is the percentage of predicted trajectories of detected actors overlapping with ground-truth static traffic objects.

An actor overlap is defined as an intersection-over-obstacle-polygon of more than 0.05 at any point of the 4s-long trajectory, set to this value to eliminate false positive overlaps due to small noise in the labeled bounding boxes.

Data. We conducted an evaluation on a large in-house data set, containing 19,000 scenes of 25s each and collected across several cities in North America with high-quality 10Hz annotations. To mitigate the metric variance of the sparse overlaps, (1) a large split of 5,000 scenes is left out for testing; (2) the test frames in scenes have a temporal spacing of 2s to avoid counting the same overlaps multiple times; and (3) the training and test sets are split geographically to prevent models from memorizing the same static obstacles and environment. Using this larger data set, as opposed to using popular open-sourced data sets that are significantly smaller, allows for lower metric variance and deriving more general conclusions. Lastly, by using the same input, backbone network, loss functions, and training settings, our studies contrast the interaction modeling approaches. This allows us to focus on the relative performance of these approaches, as opposed to comparing independent models where ensuring similar network capacities and equally well-tuned hyper-parameters are typically challenging tasks. In Appendix we provide details on the data sets, metric variance, and comparison to other motion forecasting models on public data sets.

4.1. Results

Interaction using convolutions. The performance of the single-stage model (Fig. 1a) that contains only the feature extractor is shown in Fig. 3 (**Extractor**, black). The

+IR+ICNN (green) curve shows the performance of the two-stage model without rotating the interaction region for each actor into the actor frame. In particular, starting from the 1D per-actor feature map vectors (0m), we increase the IR size to 80m. By cropping larger feature map regions that contain more interacting actors and surrounding context, displacement error and forecasted overlap rates decrease.

We then rotate IRs to match estimated actor orientation instead of using the common AV frame (**+ICM**, blue). For zero IR size (i.e., a cropped feature is still the feature pixel), we observe DE drops significantly compared to the model using the AV frame with zero IR size (green). This has been explained previously as a benefit of a standardized output representation [9]. Although defining the IR in the actor frame reduces rotational variance, the zero-size IR covers no interacting actors and we thus observe little change in the actor overlap rates. Here a lower DE is not less correlated to better interaction modeling. As the IR is increased in size, both DE and the interaction metrics improve dramatically. Crop sizes of beyond 60m show no further improvement, likely because the majority of interacting actors and obstacles are already included within the 60m region.

In all of the IRs above we have fixed the front-to-back ratio to 5:1, meaning an IR of size 60m includes 50m ahead and 10m behind the actor. In Fig. 3 inset we fix the total size at 60m, and vary the front-to-back ratio (blue). As the vast majority of actors are moving forward, we can see that placing more of the IR ahead of the actor improves interaction modeling. It is interesting to note the divergence between DE and overlap rates again: after the front-to-back ratio is above 1:1, the overlap rates continue to drop marginally, while the DE improvement stops. Even for the actor-centered IR (inset, green), not rotating the IR to match the actor orientation yields worse DE and overlap rates, which further confirms the importance of removing rotational variance for interaction modeling using convolutions. From Fig. 3, we observe that by cropping an actor-frame defined region of the feature map and then applying convo-




Figure 4. Performance from adding a GNN on top of the ICM (Fig. 1c) for different ICM interaction region sizes. **+GNN (only attributes)** encodes only node attributes in the graph edges; **+GNN (only relative)** encodes only relative locations and orientations in the graph edges.




Figure 5. Comparison of ICM and GNN (which includes ICM). **+GNN (no edges)** is identical to **+GNN** except the graph edges are cut off. **+IL** represents models trained with the additional interaction loss. Note that comparing **+ICM** at large IR (interaction is modeled by ICM) against **+GNN** at small IR (interaction is modeled by GNN) shows that a pure ICM can outperform a pure GNN in modeling interactions.

lutions improves forecasting and interaction modeling considerably. Strong dependence of overlap rates on IR size provides evidence that convolutions are effectively capturing interactions once other actors are inside the IR.

Interaction using graphs. As illustrated in Fig. 1c, for these experiments we add a GNN after the ICM. Note that, as discussed earlier, setting the IR size to 0m deactivates the ICM while retaining the benefit of reduced rotational variance. For zero IR size (Fig. 4, **+GNN**, red) we see that the GNN indeed improves DE and overlap rates significantly as compared to the models without designated interaction modeling capability in Fig. 3 (**+ICM**, 0m). Notably, even when GNN is utilized, we observe that ICM can still provide additional performance improvements as we gradually increase ICM’s interaction modeling by expanding the IR size. We also examine the benefit of the hand-crafted relative geometries in the graph edges. When the IR is small (i.e., ICM is limited), keeping only the node attributes n_i (blue) or relative geometries $rel_{i,j}$ (green) significantly damages the graph modeling. For large IR sizes, the difference between the three graph models becomes minor, suggesting that with larger feature crops the ICM has effectively compensated for missing GNN features.

The GNNs in the models above are single-iteration. We also evaluated the effect of increasing the GNN iterations to $K = 2$. An additional iteration (i.e., $K = 2$) reduces DE

and overlap rates further by a small amount when the IR size is small, which could be explained by the well-known bottleneck phenomenon of GNNs [2] and the fact that the graph is fully connected. This improvement is negligible for all but the smallest IRs, and no further exploration of additional iterations is provided below.

Convolutions vs. graphs for interaction. In Fig. 5 we compare the implicit ICM (blue) and explicit GNN (red) approaches. With zero IR size (where ICM is effectively off) the gain of adding GNN is significant. However, as the IR grows, we observe that the performance gap steadily narrows. In other words, while turning on ICM (by increasing IR size) can further improve the performance of GNN models, adding a graph to an ICM with a sufficiently large IR provides only minor benefits. To understand the gaps between **+ICM** and **+GNN** with large IR sizes, we study a graph-less model (**+GNN (no edge)**, black) created by removing graph edges in **+GNN**. For large IRs, the graph-less model matches the performance of **+GNN**, suggesting that explicit interaction graph of GNN contributes little to the performance. Thus, the gaps between **+ICM** and **+GNN** for larger IR sizes are mainly due to extra network capacity of GNN. Lastly, we see that comparing **+ICM** at large IR (i.e., interaction modeled by ICM) against **+GNN** at small IR (i.e., interaction modeled by GNN) shows that a pure ICM can outperform a pure GNN in modeling interactions.




Figure 6. Predicted trajectories sampled at 2Hz of baseline (top) and ICM (bottom). red: overlapped obstacles; blue: forecasts of the actors of interest; grey: forecasts of other actors; green: labels (also see attached videos in the Appendix).

Interaction loss. We can also see that adding the interaction loss (Eq. 8) reduces the overlap between actors’ predicted trajectories for both interaction modeling approaches (green and magenta in Fig. 5). The improvement is significant for smaller IRs, which may be due to the fact that the smaller IRs do not provide enough information to model the interactions effectively, benefiting more from this added supervision. On the other hand, when interaction is modeled more effectively through larger IR, the loss is sparser and thus contributes less. Interestingly, the interaction loss does not affect DE results except for ICM models at small IR where the interaction modeling is limited.

Maneuver-specific qualitative results. In Fig. 6 we present a comparison of the baseline ICM model with 0m size (that has no designated interaction modeling) and the ICM model with 60m size on three typical maneuvers observed in interacting scenarios: adaptive cruise control (ACC), turn, and nudging. We note that the 0m model in all cases incorrectly predicts overlapping trajectories. In the ACC case the ICM model correctly predicts that the vehicle would decelerate and queue after others, while in the turn case it outputs a trajectory that follows the lane and avoids overlapping with the vehicles after the turn. In the nudging case the vehicle motion starts with considerable curvature, the forecast correctly reduces the curvature and straightens the trajectory to avoid the parked cars. We also examined the results of adding GNNs on +ICM (60m) on these maneuvers, and observed no significant difference.

Inference time. The baseline model that includes the feature extractor and other parts such as input pre-processing and output post-processing takes 45.6ms per frame. Next we measure the additional time costs of adding the ICM and GNN modules to the baseline model, shown in Table 1. The ICM of zero IR size adds an additional 5.2ms, which includes processing of the feature pixel and computation of the final output. ICM with non-zero size uses convolutions and bilinear feature cropping, operations that have been optimized in current GPU software and hardware. As a result, even the largest 80m ICM is only a few milliseconds slower than the 0m ICM. Lastly, the GNN itself takes

Table 1. Inference times of modules (tested on Nvidia Titan RTX)

Module	IR size [m]	Inference [ms]
ICM	0	5.2
ICM	80	8.1
GNN	-	46.9

46.9ms, multiple times slower than the slowest ICM. This is consistent with earlier results showing GNN inference may be inefficient resulting in higher latency [21]. Coupled with the earlier results showing that modeling interaction using convolutions can give competitive performance compared to GNNs, we see that the convolutional approach represents an efficient and practical alternative to GNNs.

5. LIMITATIONS AND SOCIAL IMPACT

While graphs represent a general approach to modeling various relations, this work shows that convolutions are also effective when it comes to modeling spatial interaction. However, in Euclidean space where CNNs are applied, some information such as driver-driver vocal interactions are not currently represented. Besides, as our studies are limited to the 2D AV application, it is unanswered whether convolutions are still effective in capturing interactions in 3D space (such as for modeling human body motion). Additionally, the overlap rates over trajectories enable us to explicitly evaluate quality of the interaction modeling. However, the metric sparsity leads to requirement for having larger test data to ensure the obtained results are meaningful, which limits a wider use of the metric. Lastly, neither of the considered approaches guarantees zero overlaps over the trajectories, and CNN models also often have lower interpretability than the GNN approaches.

6. CONCLUSION

We compared and contrasted 2D convolutional and graph neural networks for the task of spatial interaction modeling, providing empirical evidence that under certain conditions convolutional networks can reach comparable performance to the state-of-the-art GNNs (e.g., by modifying IR), thus allowing similar motion forecasting accuracy and interaction modeling while maintaining reduced latency and model complexity of the model. We analyzed common components of the interaction approaches, leading to a better understanding of how each benefits the interaction modeling. Moreover, we introduced a novel interaction-aware loss and showed its impact on the considered approaches. Our work presents a basis for wider use of convolutional layers for the task of spatial interaction, providing evidence that the gap between convolutional models and more complex and computationally expensive GNN models may not be as large as previously suspected.

References

- [1] Alexandre Alahi, Kratarth Goel, Vignesh Ramanathan, Alexandre Robicquet, Li Fei-Fei, and Silvio Savarese. Social lstm: Human trajectory prediction in crowded spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 961–971, 2016. [2](#)
- [2] Uri Alon and Eran Yahav. On the bottleneck of graph neural networks and its practical implications. *arXiv preprint arXiv:2006.05205*, 2021. [7](#)
- [3] Peter W Battaglia, Razvan Pascanu, Matthew Lai, Danilo Rezende, and Koray Kavukcuoglu. Interaction networks for learning about objects, relations and physics. *arXiv preprint arXiv:1612.00222*, 2016. [1](#)
- [4] Sergio Casas, Cole Gulino, Renjie Liao, and Raquel Urtasun. Spatially-aware graph neural networks for relational behavior forecasting from sensor data. *arXiv preprint arXiv:1910.08233*, 2019. [1](#), [2](#), [4](#), [5](#)
- [5] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *Conference on Robot Learning*, pages 947–956, 2018. [2](#)
- [6] Henggang Cui, Vladan Radosavljevic, Fang-Chieh Chou, Tsung-Han Lin, Thi Nguyen, Tzu-Kuo Huang, Jeff Schneider, and Nemanja Djuric. Multimodal trajectory predictions for autonomous driving using deep convolutional networks. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 2090–2096. IEEE, 2019. [4](#)
- [7] Nachiket Deo and Mohan M. Trivedi. Convolutional social pooling for vehicle trajectory prediction. *CoRR*, abs/1805.06771, 2018. [2](#)
- [8] Frederik Diehl, Thomas Brunner, Michael Truong-Le, and Alois C. Knoll. Graph neural networks for modelling traffic participant interaction. *CoRR*, abs/1903.01254, 2019. [2](#)
- [9] Nemanja Djuric, Henggang Cui, Zhaoen Su, Shangxuan Wu, Huahua Wang, Fang-Chieh Chou, Luisa San Martin, Song Feng, Rui Hu, Yang Xu, et al. Multinet: Multiclass multi-stage multimodal motion prediction. In *Proceedings of the IEEE Intelligent Vehicles Symposium (IV)*, 2020. [2](#), [3](#), [4](#), [6](#)
- [10] Nemanja Djuric, Vladan Radosavljevic, Henggang Cui, Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, and Jeff Schneider. Short-term motion prediction of traffic actors for autonomous driving using deep convolutional networks. *arXiv preprint arXiv:1808.05819*, 2018. [2](#), [4](#)
- [11] Francis Engelmann, Theodora Kontogianni, and Bastian Leibe. Dilated point convolutions: On the receptive field size of point convolutions on 3d point clouds. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9463–9469. IEEE, 2020. [1](#)
- [12] Alex Fout, Jonathon Byrd, Basir Shariat, and Asa Ben-Hur. Protein interface prediction using graph convolutional networks. In *Advances in neural information processing systems*, pages 6530–6539, 2017. [1](#)
- [13] Jiyang Gao, Chen Sun, Hang Zhao, Yi Shen, Dragomir Anguelov, Congcong Li, and Cordelia Schmid. Vectornet: Encoding hd maps and agent dynamics from vectorized representation. *arXiv preprint arXiv:2005.04259*, 2020. [2](#)
- [14] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. *arXiv preprint arXiv:1704.01212*, 2017. [4](#), [5](#)
- [15] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social GAN: socially acceptable trajectories with generative adversarial networks. *CoRR*, abs/1803.10892, 2018. [2](#)
- [16] Takuo Hamaguchi, Hidekazu Oiwa, Masashi Shimbo, and Yuji Matsumoto. Knowledge transfer for out-of-knowledge-base entities: A graph neural network approach. *arXiv preprint arXiv:1706.05674*, 2017. [1](#)
- [17] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in neural information processing systems*, pages 1024–1034, 2017. [1](#)
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. [3](#)
- [19] Boris Ivanovic and Marco Pavone. Modeling multimodal dynamic spatiotemporal graphs. *CoRR*, abs/1810.05993, 2018. [2](#)
- [20] Elias Khalil, Hanjun Dai, Yuyu Zhang, Bistra Dilkina, and Le Song. Learning combinatorial optimization algorithms over graphs. In *Advances in neural information processing systems*, pages 6348–6358, 2017. [1](#)
- [21] Kevin Kinningham, Christopher Re, and Philip Levis. Grip: A graph neural network accelerator architecture. *arXiv preprint arXiv:2007.13828*, 2020. [8](#)
- [22] Thomas Kipf, Ethan Fetaya, Kuan-Chieh Wang, Max Welling, and Richard Zemel. Neural relational inference for interacting systems. In *International Conference on Machine Learning*, pages 2688–2697. PMLR, 2018. [1](#), [2](#)
- [23] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*, 2016. [1](#)
- [24] T.-Y. Lin, P. Goyal, R. Girshick, K. He, and P. Dollar. Focal loss for dense object detection. In *ICCV*, 2017. [3](#)
- [25] Wenjie Luo, Bin Yang, and Raquel Urtasun. Fast and furious: Real time end-to-end 3d detection, tracking and motion forecasting with a single convolutional net. In *Proc. of the IEEE CVPR*, pages 3569–3577, 2018. [2](#)
- [26] Sindy Löwe, David Madras, Richard Zemel, and Max Welling. Amortized causal discovery: Learning to infer causal graphs from time-series data. *arXiv preprint arXiv:2006.10833*, 2020. [1](#)
- [27] Jianqi Ma, Weiyuan Shao, Hao Ye, Li Wang, Hong Wang, Yingbin Zheng, and Xiangyang Xue. Arbitrary-oriented scene text detection via rotation proposals. *CoRR*, abs/1703.01086, 2017. [2](#), [4](#)
- [28] Eli A. Meir, Haggai Maron, Shie Mannor, and Gal Chechik. How to stop epidemics: Controlling graph dynamics with reinforcement learning and graph neural networks. *arXiv preprint arXiv:2010.05313*, 2020. [1](#)
- [29] Alexander Neubeck and Luc Van Gool. Efficient non-maximum suppression. In *18th International Conference on Pattern Recognition (ICPR'06)*, volume 3, pages 850–855. IEEE, 2006. [5](#)

- [30] Shah Rukh Qasim, Jan Kieseler, Yutaro Iiyama, and Maurizio Pierini. Learning representations of irregular particle-detector geometry with distance-weighted graph networks. *The European Physical Journal C*, 79(7), Jul 2019. [1](#)
- [31] Nicholas Rhinehart, Rowan McAllister, Kris M. Kitani, and Sergey Levine. PRECOG: prediction conditioned on goals in visual multi-agent settings. *CoRR*, abs/1905.01296, 2019. [1](#), [2](#)
- [32] Amir Sadeghian, Ferdinand Legros, Maxime Voisin, Ricky Vesel, Alexandre Alahi, and Silvio Savarese. Car-net: Clairvoyant attentive recurrent network. *CoRR*, abs/1711.10061, 2017. [1](#), [2](#)
- [33] Alvaro Sanchez-Gonzalez, Nicolas Heess, Jost Tobias Springenberg, Josh Merel, Martin Riedmiller, Raia Hadsell, and Peter Battaglia. Graph networks as learnable physics engines for inference and control. *arXiv preprint arXiv:1806.01242*, 2018. [1](#)
- [34] Michael Schlichtkrull, Thomas N Kipf, Peter Bloem, Rianne Van Den Berg, Ivan Titov, and Max Welling. Modeling relational data with graph convolutional networks. In *European semantic web conference*, pages 593–607. Springer, 2018. [1](#)
- [35] Zhaoen Su, Chao Wang, Henggang Cui, Nemanja Djuric, Carlos Vallespi-Gonzalez, and David Bradley. Temporally-continuous probabilistic prediction using polynomial trajectory parameterization. In *IEEE International Conference on Robotics and Automation (IROS)*, 2021. [3](#)
- [36] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. [4](#), [5](#)
- [37] Da Xu, Chuanwei Ruan, Evren Korpeoglu, Sushant Kumar, and Kannan Achan. Inductive representation learning on temporal graphs. *arXiv preprint arXiv:2002.07962*, 2020. [1](#)