

MVFuseNet: Improving End-to-End Object Detection and Motion Forecasting through Multi-View Fusion of LiDAR Data

Ankit Laddha, Shivam Gautam, Stefan Palombo, Shreyash Pandey, Carlos Vallespi-Gonzalez
Aurora Innovation

aladdha, sgautam, spalombo, spandey, cvallespi@aurora.tech

Abstract

In this work, we propose *MVFuseNet*, a novel end-to-end method for joint object detection and motion forecasting from a temporal sequence of LiDAR data. Most existing methods operate in a single view by projecting data in either range view (RV) or bird’s eye view (BEV). In contrast, we propose a method that effectively utilizes both RV and BEV for spatio-temporal feature learning as part of a temporal fusion network as well as for multi-scale feature learning in the backbone network. Further, we propose a novel sequential fusion approach that effectively utilizes multiple views in the temporal fusion network. We show the benefits of our multi-view approach for the tasks of detection and motion forecasting on two large-scale self-driving data sets, achieving state-of-the-art results. Furthermore, we show that *MVFuseNet* scales well to large operating ranges while maintaining real-time performance.

1. Introduction

Object detection and motion forecasting are of paramount importance for autonomous driving. Object detection entails recognizing and localizing objects in the scene, whereas motion forecasting entails predicting the future trajectory of the detected objects. Traditionally, cascaded approaches treat detection and motion forecasting as two separate tasks, which enables classical methods for detection and motion forecasting to be used. However, these methods optimize for these tasks separately, suffering from cascading errors and missing the opportunity to share learned features for both tasks [1]. To overcome these issues, multiple end-to-end methods have been proposed [2, 3, 4, 5] for jointly solving both detection and motion forecasting. These methods have demonstrated excellent performance [1] while operating in real-time. Following the end-to-end paradigm, we propose a novel method for jointly detecting objects and predicting their future trajectories using time-series LiDAR data.

The input LiDAR data is natively captured in the per-

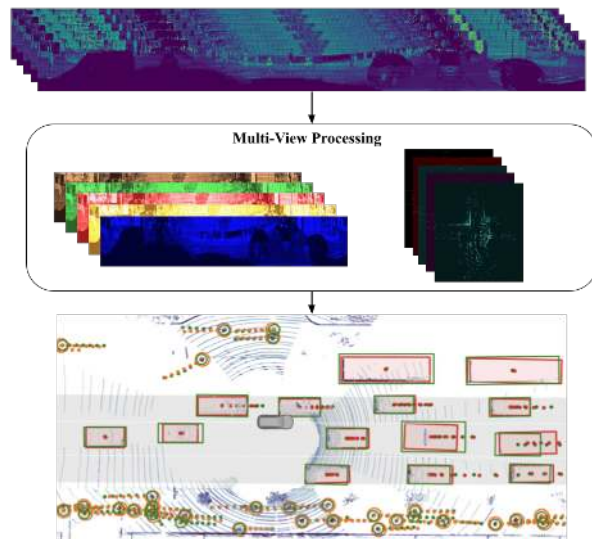


Figure 1: The input to our method is a temporal sequence of 3D native range view images from LiDAR (top) and the output is object detections and motion predictions in the Cartesian bird’s eye view (bottom). In contrast to previous single view methods, we propose to process the sequence in both views (middle).

spective range view (RV). However, since most planning algorithms operate in the Cartesian bird’s eye view (BEV) space, the object detections and their forecasts need to also be in the same Cartesian space (see Figure 1). Therefore, every method converts perspective RV information to a Cartesian BEV at some stage during its processing. Most existing methods lie on the extreme ends of the spectrum with respect to when they perform this conversion during their processing, and most use a single view entirely. On one hand, methods such as [3, 6, 7] process LiDAR data exclusively in RV and only convert their final output to BEV during post processing. These methods are efficient for processing large spatial regions due to the compact size of the input image and offer state-of-the-art performance in the detection of small objects (e.g., pedestrians, bikes) and far

away objects. On the other hand, methods such as [2, 4, 8] project the LiDAR data in the BEV first, with minimal or no pre-processing in RV, and perform most of the processing in BEV. The Cartesian BEV has the advantage of a strong prior due to range invariance of object shape and motion. This provides an edge to existing BEV methods on motion forecasting tasks; however, their scalability to operate in large areas remains a challenge. There has been some recent work on using multiple views for detection [9, 10], but the space of models that can efficiently use multiple views for end-to-end detection and motion forecasting remains largely unexplored.

Given the complementary benefits of processing in both views, we posit that effectively combining both of them can lead to improved performance in both detection and forecasting. Therefore, in this work we propose *MVFuseNet*, a novel end-to-end joint object detection and motion forecasting method which achieves state-of-the-art results on two large scale data sets and has real-time performance when processing a large spatial region. To accomplish this, we propose a novel sequential multi-view (MV) fusion network to aggregate a temporal sequence of LiDAR data for learning spatio-temporal features. We further propose a multi-view backbone network to process the spatio-temporal features for detection and forecasting. We demonstrate the effectiveness of multiple views over a single view on multiple data sets with different characteristics and LiDAR resolutions.

2. Related Work

In this section, we first discuss the existing literature on LiDAR representation, and then look at various approaches for motion forecasting.

2.1. LiDAR representation

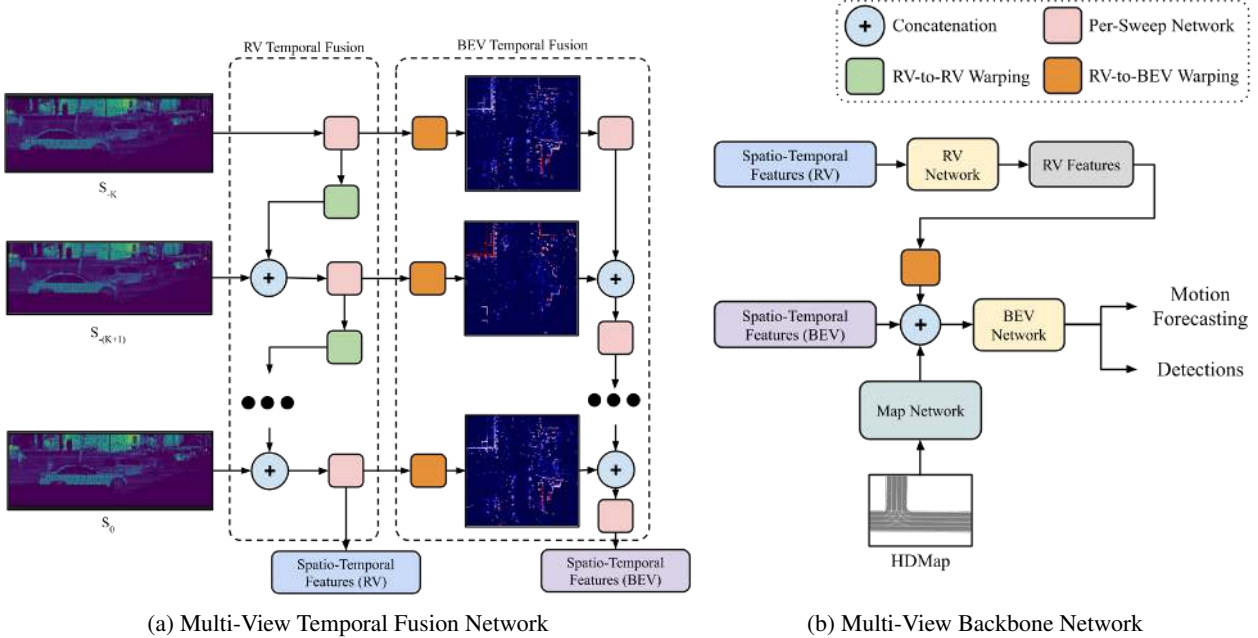
A spinning LiDAR captures data as a multi-channel image of range measurements. In the literature, these range measurements have been represented in various ways for processing: unstructured 3D point clouds [11, 12], 3D voxels [13, 14], a 2D BEV grid [8, 15, 16] and the native 2D RV grid [7, 17, 18, 19]. The point cloud and voxel based methods are computationally expensive and do not scale well to highly dynamic and crowded outdoor scenes. In comparison, 2D BEV or RV grid based methods are efficient but only use a single view (either BEV or RV) for processing LiDAR data. Recent work has investigated the use of multiple views [20, 21, 22, 23, 24] and shown that the complementary benefits of both views improve performance. However, these methods use only one frame of LiDAR data and only solve perception tasks such as object detection and semantic segmentation. In contrast, we propose a method which aggregates data from multiple frames to jointly solve both detection and motion forecasting in an end-to-end method by utilizing both the BEV and RV.

Recently, [10] proposed a multi-view approach for the joint task. In this method, the authors proposed fusing a single-frame RV projection with multiple frames of BEV projection, which improves object detection performance. This method, however, limits the temporal fusion of LiDAR data to BEV and only employs RV features of a single sweep, missing high resolution motion cues. In contrast, our proposed method performs spatio-temporal fusion of both BEV and RV features for multiple frames of LiDAR data. To the best of our knowledge, this is the first method that exploits multiple views for both temporal fusion and multi-scale backbone feature learning. We show that this leads to better detection and motion forecasting performance.

2.2. Motion Forecasting

Traditional learning-based motion forecasting approaches [25, 26, 27, 28] use temporal sequences of detections [7, 8, 21, 29] to learn spatio-temporal features for each object. Recent work in traditional motion forecasting has focused on improving the modeling of uncertainty through multi-modality [30, 31, 32, 28, 33, 34, 35, 36] and interactions between actors and the scene [26, 27, 37, 38, 39, 40]. In contrast, we look at the complementary problem of learning better spatio-temporal object features for forecasting using sensor data. Our proposed method can also benefit from many of the recent advances in the motion forecasting literature. However, to simplify the experimentation, we leave their incorporation to future work. These traditional methods are successful in capturing complex relationships and generating realistic longer-term forecasts, but they suffer from cascading error issues [1] and lose out on the rich features learned from sensor data. These methods also work on a per-object basis, which makes them hard to scale to dense, urban environments.

To address the issues with traditional forecasting approaches, the seminal work by [41] proposed to jointly solve both object detection and motion forecasting. [2] improved upon [41] by incorporating scene information using a semantic and geometric HDMaP. Approaches such as [4] and [1] build on top of [2] by adding an object-centric sub-network to refine future trajectories. These methods show that recent work on multi-modal predictions and the use of interaction graphs to model complex relationships can be easily extended to the framework of joint object detection and motion forecasting. [5] and [10] are recent multi-sensor methods that build on top of [4] by using radar and camera inputs respectively. These methods, by virtue of operating in BEV, lose out on high-resolution point information and are often limited by range of operation. RV based methods such as [6] and [3] overcome the limitation on operating range but are outperformed in the motion forecasting task by recent BEV based methods. In this work, we improve the joint framework by including multi-view representation



(a) Multi-View Temporal Fusion Network

(b) Multi-View Backbone Network

Figure 2: **MVFuseNet Overview:** We propose a novel approach for (a) multi-view temporal fusion of LiDAR data in RV and BEV to learn spatio-temporal features. We sequentially aggregate sweeps by projecting the data from one sweep to the next sweep in the temporal sequence. (b) These multi-view spatio-temporal features are further processed by a multi-view backbone to combine them with map features and learn multi-scale features for final detection and motion forecasting.

in multiple parts of the network and achieve state-of-the-art performance on both object detection and motion forecasting while scaling to large areas of operation in real-time.

3. MV Detection and Motion Forecasting

Figure 2 shows an overview of our proposed approach. Our main contribution is an end-to-end object detection and motion forecasting method that processes the time-series LiDAR data in both range view and bird’s eye view. We first describe prerequisite information related to the input and view-projections in Section 3.1. We then discuss our contribution of using multiple views for temporal fusion of a sequence of LiDAR data in Section 3.2. In Section 3.3, we discuss our contribution of a multi-view backbone network to extract per-cell features. Finally, we present our method for joint detection and motion forecasting using the per-cell features in section 3.4, followed by the loss functions used to train the model in Section 3.5.

3.1. Preliminaries

Input: Let us assume that we are given a time-series of $K + 1$ sweeps, where each sweep contains all the LiDAR points from a full 360° rotation of a LiDAR sensor. This time series LiDAR data can be denoted by $\{S_k\}_{k=-K}^0$, where $k = 0$ is the most recent sweep and $-K \leq k \leq 0$ are the past sweeps. We term the most recent sweep as the *reference sweep*. Each LiDAR sweep contains N_k range measurements, which can be transformed into a set of 3D

points, $S_k = \{p_k^i\}_{i=1}^{N_k}$, using the pose (viewpoint) of the sensor \mathcal{P}_k at the end of sweep capture. We assume that pose for each sweep is provided by an onboard localization system. Therefore, we can calculate the transformation of points from one viewpoint to another. We denote the k -th sweep transformed into the n -th sweep’s coordinate frame as, $S_{k,n} = \{p_{k,n}^i\}_{i=1}^{N_k}$, where each point $p_{k,n}^i$ is represented by its 3D coordinates, $[x_{k,n}^i, y_{k,n}^i, z_{k,n}^i]^T$. In spherical coordinates the same point $p_{k,n}^i$ can be represented using the radial distance $r_{k,n}^i$, azimuth $\theta_{k,n}^i$ and elevation $\phi_{k,n}^i$. Note that $p_{k,n}^i$ represents the same LiDAR return as p_k^i , only transformed into a different frame.

Projections: For each point p_k^i captured at pose \mathcal{P}_k , the range view projection at pose \mathcal{P}_n is defined by discretizing the azimuth and elevation angles of $p_{k,n}^i$. Similarly, the bird’s eye projection at pose \mathcal{P}_n is the x and y coordinates of $p_{k,n}^i$.

Per-Point Features: For each point p_k^i in S_k , we define a set of associated features as concatenation of its coordinates in original viewpoint, $[x_{k,k}^i, y_{k,k}^i, z_{k,k}^i]^T$, coordinates in most recent viewpoint, $[x_{k,0}^i, y_{k,0}^i, z_{k,0}^i]^T$ and the remission or intensity e_k^i of the LiDAR return.

3.2. Multi-View Temporal Fusion Network

The goal of the temporal fusion sub-network is to aggregate a time-series of LiDAR data in order to learn spatio-temporal features. The most straightforward approach, as

employed by many previous works [2, 4, 6], is the *one-shot* approach where all the data is accumulated in a single frame. All points are first transformed into the frame defined by the reference pose and then the aggregation is done by projecting them in either BEV or RV. For multiple views this can be trivially extended by projecting the points in both BEV and RV for aggregation. However, directly projecting all the past LiDAR data into the RV of the most recent sweep leads to significant performance degradation due to heavy data loss in the projection step [3]. Therefore, instead of previous approaches that focus on *one-shot* projection, we propose a novel sequential multi-view fusion approach to effectively aggregate the temporal LiDAR data.

Figure 2a shows our proposed fusion approach. We assume that the input is a time-series of multi-channel RV images in their original capture pose. These images contain the per point features, f_k^i , as defined in Section 3.1. We sequentially fuse the LiDAR sweeps from one time-step to the next in both views. At each time-step we warp the previous time-step’s RV features to the current time-step’s frame (green box), and then use a sub-network (see Figure 3a) to learn spatio-temporal features for each cell in RV (pink box). These learned features are then projected into the BEV (orange box) and concatenated with the BEV features from the previous time-step. Similarly to RV, a sub-network is then used to learn spatio-temporal features for each cell in BEV. The feature learning networks (pink box) in each view and time-step are independent and no weights are shared across time or view. It is important to note that unlike previous methods that project raw point-features to the BEV, our method projects learned RV features to be used in the BEV. We further discuss the methods used to warp features from one RV to another and for projecting the RV features to BEV.

RV-to-RV Feature Warping: Let us assume that we would like to warp the RV feature map $\mathcal{R}_{k,k}$ of k th sweep to the RV feature map $\mathcal{R}_{k,n}$ at viewpoint on n th sweep. We assume that the point p_k^i is projected to location $l_{k,k}^i$ in $\mathcal{R}_{k,k}$ and $l_{k,n}^i$ in $\mathcal{R}_{k,n}$. Therefore, we define the feature warping by copying the features from one RV to another such that $\mathcal{R}_{k,n}(l_{k,n}^i) = \mathcal{R}_{k,k}(l_{k,k}^i)$. Similar to [3, 7], if more than one point project into the same cell location $l_{k,n}^i$, we pick the closest point for feature rendering.

RV-to-BEV Feature Warping: Let us assume that we would like to warp the RV feature map $\mathcal{R}_{k,k}$ of k th sweep to the BEV feature map $\mathcal{B}_{k,0}$. We also assume that a point p_k^i in \mathcal{S}_k can be projected in $\mathcal{R}_{k,k}$ to extract a learned feature g_k^i . We calculate the features of cell $l_{k,0}^i$ in $\mathcal{B}_{k,0}$ by aggregating the features of all the points $\mathcal{A}_k^i = \{p_k^i, i = 1, \dots, M\}$ that are projected into that cell location. Similarly to [8, 21], for each point in a cell, we calculate its feature vector h_k^i by concatenating the difference between the coordinates of the point and the cell $\Delta c = [x_k^i - l_{x,k,0}^i, y_k^i - l_{y,k,0}^i]$, and the

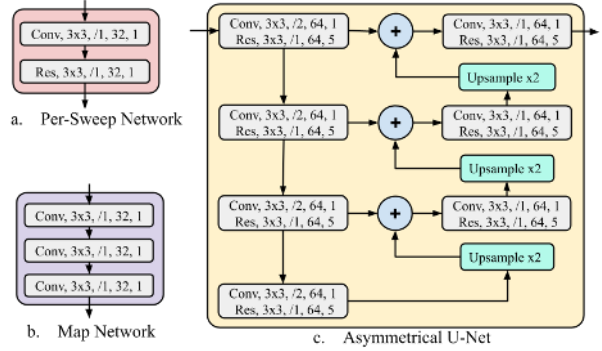


Figure 3: **Network Components:** (a) We use the depicted per-sweep network to process each sweep during temporal fusion in both views. Note that no weights are shared across time and views during temporal fusion. (b) The HDMap is processed with the depicted network to learn local map-only features which are combined with the LiDAR features. (c) The asymmetric U-Net network is used to extract and combine multi-scale features in BEV. In RV, only the width dimension is down-sampled and the first convolutional layer is not strided. Each layer in the networks is represented as $B, k \times k, /s, C, N$, where B is the block name, k is the kernel size, s is the stride, C is the number of channels and N is the number of repetitions of the block. *Conv* denotes a convolutional layer followed by batch normalization and ReLU. *Res* denotes a residual block as defined in [42]. Finally, we up-sample using bi-linear interpolation.

RV features of the point g_k^i . For aggregating the features of all the points in the cell we use:

$$\mathcal{B}_{k,0}(l_{k,0}^i) = \frac{1}{M} \sum_{i=0}^M \text{MLP}(h_k^i), \quad (1)$$

where MLP is a linear layer followed by batch normalization and ReLU.

3.3. Multi-View Backbone Network

The goal of the backbone is to process the spatio-temporal features and combine them with map features to learn per-cell features that can be used for object detection and motion forecasting. As shown in Figure 2b, our backbone network processes the spatio-temporal features in both views. We first take the spatio-temporal features in RV as input and learn multi-scale RV features by extracting and combining features using an asymmetrical U-Net (see Figure 3c). These RV features are then projected to BEV using the same technique as in Section 3.2 and concatenated with learned map features and the spatio-temporal BEV features (see Figure 2b). We rasterize the map features in BEV [2, 4] and learn high level features using a convolutional neural network (see Figure 3b). Similar to RV, this multi-view, multi-sensor feature volume is further processed by another

asymmetrical U-Net to learn multi-scale features in BEV (see Figure 3c).

3.4. Output Prediction

Given the per cell features from the backbone network, our goal is to detect objects observed in the current sweep \mathcal{S}_0 and predict their trajectory. We use a dense, single-stage convolutional header for detecting objects using the per-cell features. Similarly to [4, 43], we first predict if a cell contains the center of an object for some class. For each center cell, we then predict an associated bounding box and use non-maximum suppression to remove duplicates. For motion forecasting of large objects such as vehicles, we extract a rotated region of interest (RROI) [1, 4] of $60 \times 60m$ centered at the object to learn actor-centric features which are then used to predict the trajectory. However, for smaller objects such as pedestrians and bicycles, we use the features of the center cell to predict the trajectory since we empirically found that this leads to better results.

3.5. End-to-End Training

Similarly to [4, 5], we train the proposed method end-to-end using a multi-task loss incorporating both detection and trajectory loss: $\mathcal{L}_{total} = \mathcal{L}_{det} + \mathcal{L}_{traj}$.

Detection Loss (\mathcal{L}_{det}) is a multi-task loss defined as a weighted sum of classification and regression loss: $\mathcal{L}_{det} = \mathcal{L}_{det}^{cls} + \lambda \mathcal{L}_{det}^{reg}$. We use focal loss [44] for classifying if a BEV cell is at the center of an object class. For each center cell, we use smooth L1 loss to learn parameters of the object bounding box relative to that cell. We parameterize each box i by its center (x_i, y_i) , orientation (θ_i) and size (w_i, h_i) . The orientation is further parameterized as $(\cos(\theta_i), \sin(\theta_i))$.

Trajectory Loss (\mathcal{L}_{traj}) is defined as an average of per future time-step loss: $\mathcal{L}_{traj} = 1/T \sum_{t=1}^T \mathcal{L}_t^{KL}$ [4]. We consider each waypoint at time t of a trajectory j to be a 2D Laplace distribution parameterized by its position (x_j^t, y_j^t) and scale $(\sigma_{j,x}^t, \sigma_{j,y}^t)$. We use the KL divergence [45] between the ground truth and predicted distribution as loss \mathcal{L}_t^{KL} to learn the per waypoint distribution.

4. Experiments

4.1. Data set and Metrics

We report results on two autonomous driving data sets, with different LiDAR resolutions and characteristics, to show the efficacy of our proposed approach. In particular, we use the publicly available nuScenes [46] data set, and a much larger internal data set. The nuScenes data set consists of 1k snippets. It has a low resolution LiDAR which generates $\sim 30k$ points per sweep and a square region of interest (ROI) of length 100m, centered on the self-driving vehicle (SDV). On the other hand, our internal data set consists of 17k snippets. It has a higher resolution LiDAR which generates $\sim 130k$ points per sweep and uses a ROI of a square

of 200m length. On both data sets, we report results on three major classes of traffic participants: vehicles, pedestrians and bikes.

Following previous works [1, 4, 3], we use average precision (AP) with intersection over union (IoU) based association between ground truth and the detected object. Furthermore, we use L_2 displacement error at multiple time horizons to evaluate motion forecasting. We compute L_2 as the Euclidean distance between the center of the predicted true positive box and the associated ground truth box. Note that the official nuScenes leaderboard evaluates the task of detection and state estimation, whereas in this work we solve the joint task of detection and motion forecasting. Therefore, we use the same metrics as used in previous work [1, 2] instead of the official leaderboard metrics.

4.2. Implementation Details

We use the PyTorch [47] library for implementing the proposed approach. On nuScenes, the input RV is chosen to be of size 32×1024 based on the LiDAR characteristics. Furthermore, the input BEV feature map is chosen to be 400×400 and the backbone output is chosen to be 200×200 , to balance runtime and resolution. This results in an input resolution of 25cm and an output resolution of 50cm. On our internal data set, the input RV is 64×2048 and both the input BEV and output BEV feature map are of the size 400×400 . Due to the large ROI, this results in a resolution of 50cm at both input and output. For both data sets, we use the LiDAR data from the past 0.5 seconds as input and predict the trajectory for 3 seconds into the future, sampled at 10Hz. Since nuScenes is much smaller than our internal data set, we use data augmentation during training. Specifically, we generate labels at non-key frames by linearly interpolating the labels at adjacent key frames. We further randomly augment each frame by applying translation ($\pm 1m$ for the x - and y -axes and $\pm 0.2m$ for z axis) and rotation (between $\pm 45^\circ$ along the z -axis) to both the point clouds and labels.

We train with a batch size of 64 distributed over 32 GPUs. We first pre-train the network without rotated ROI for 20 epochs and then warm start the model with the pre-trained weights and train for 6 more epochs. We train the network using a cosine learning rate schedule with a starting rate of 1×10^{-3} and an end rate of 2×10^{-5} . We set the gamma in focal loss to 2 and the loss weight λ in the detection loss to 0.2.

4.3. Comparison to the State-of-the-Art

In this section, we compare our method to existing end-to-end methods using the evaluation setting of [4, 10]. As shown in Table 1, our novel multi-view method significantly outperforms all other methods, on both detection and forecasting tasks for all evaluated classes.

We see significant improvements on both detection and

Table 1: **nuScenes**: Comparison of proposed *MVFuseNet*, with existing end-to-end methods. The reported L_2 is at 3s.

Method	Vehicle		Pedestrian		Bikes	
	AP (%) \uparrow	L_2 (cm) \downarrow	AP (%) \uparrow	L_2 (cm) \downarrow	AP (%) \uparrow	L_2 (cm) \downarrow
SpAGNN [1]	-	145	-	-	-	-
Laserflow [6]	56.1	143	-	-	-	-
RVFuseNet [3]	59.9	120	-	-	-	-
LiRANet [5]	63.7	102	-	-	-	-
IntentNet [2]	60.3	118	63.4	84	31.8	173
MultiXNet [4]	60.6	105	66.1	80	32.6	203
L-MV [10]	61.1	107	71.0	82	38.2	187
LC-MV [10]	62.9	107	71.4	80	39.8	179
MVFuseNet (Ours)	67.8	99	76.4	75	44.5	138

motion forecasting when we compare our method to the best RV-based method *RVFuseNet* [3], and the state-of-the-art BEV-based method *MultiXNet* [4]. Notably, our method shows a $\sim 15\%$ improvement on pedestrian detection, a $\sim 40\%$ improvement on bike detection, and a $\sim 30\%$ improvement on motion forecasting of bikes, as compared to the best BEV-only *MultiXNet*. Next, we compare our method to another recent multi-view method *L-MV* [10]. As shown in Table 1, our method outperforms *L-MV* [10] on all classes by a large margin on both detection and forecasting. Building on top of *MultiXNet*, *L-MV* only improved the detection performance by incorporating a single sweep in RV. In contrast, we are able to utilize the temporal sequence in RV to improve both detection and motion forecasting performance. This demonstrates that our proposed method can leverage multiple views much more effectively than previous multi-view end-to-end methods. Finally, we show that our method, with only LiDAR information, is able to outperform multi-sensor methods like *LiRANet* [5] (which uses RADAR in addition to LiDAR) and *LC-MV* [10] (which uses camera images in addition to LiDAR).

4.4. Ablation Studies

In this section, we analyze the impact of individual components of our multi-view fusion model. We study the individual effect of using RV and BEV information in the temporal fusion network, as well as in the backbone network. Further, we study the efficacy of our sequential fusion method for fusing multiple time-step information.

4.4.1 Views in Temporal Fusion Network

First, we study the use of multiple views in temporal fusion, as compared to only using a single view. The RV-only baseline is created by removing the *BEV Temporal Fusion* block in Figure 2a. Similarly the BEV-only baseline is created by removing the *RV Temporal Fusion* block in Figure 2a and directly warping the input RV features into BEV without any temporal fusion in RV. For a fair comparison, we keep the same number of parameters between the single-view and multi-view experiments by moving the additional convolutions from one view to another. The results of the

comparison are shown in Table 2. We observe that BEV-only fusion significantly outperforms the RV-only fusion in the task of motion forecasting. We believe this is due to the strong prior that the BEV representation provides to motion forecasting. However, after combining both views we get a large performance improvement over the BEV-only fusion model. This suggests that there is relevant information to tasks of detection and motion forecasting that is unique to each view. We also observe that the relative performance improvement on smaller objects, such as bikes (20%) and pedestrians (6%), is larger than on bigger objects such as vehicles (2%), suggesting that the network is able to utilize the higher resolution information present in RV.

4.4.2 Views in Backbone Network

Next, we analyze the impact of including multiple views in the backbone network. We perform this ablation with the best performing multi-view, temporal fusion model from Section 4.4.1. The RV-only and BEV-only baselines are created by removing the *BEV Network* and *RV Network* respectively in Figure 2b. For the RV-only baseline we extend the detector to include two extra convolutional layers to aggregate some BEV context. Similar to the previous study, we keep the same number of parameters between single-view and multi-view experiments by moving the convolutions from one view to another. As we can see from the results in Table 3, using both RV and BEV in the backbone improves performance over any single-view method on both data sets. We further observe that the relative improvement on our internal data set is larger than on nuScenes. We believe this can be attributed to the better utilization of the 2x higher resolution LiDAR in the RV.

4.4.3 Strategies for Multi-View Temporal Fusion

Finally, we compare our proposed *sequential* fusion approach to the naive *one-shot* approach. In contrast to *sequential* warping, the *one-shot* approach warps the raw per-point features from all the past sweeps directly into the RV and BEV of the reference pose. The temporal aggregation in each view is then performed by concatenating the features

Table 2: Comparison of Views in Temporal Fusion Network

View	Vehicle					Pedestrian					Bikes				
	AP (%) \uparrow		L_2 (cm) \downarrow			AP (%) \uparrow		L_2 (cm) \downarrow			AP (%) \uparrow		L_2 (cm) \downarrow		
	0.5 IoU	0.7 IoU	0.0 s	1.0 s	3.0 s	0.1 IoU	0.3 IoU	0.0 s	1.0 s	3.0 s	0.1 IoU	0.3 IoU	0.0 s	1.0 s	3.0 s
nuScenes															
RV	80.3	61.8	46.5	87.4	193.2	64.8	63.1	17.5	93.9	273.2	36.2	31.8	32.5	103.5	244.6
BEV	83.2	65.1	41.5	57.5	122.4	70.8	69.0	16.6	33.7	84.8	42.5	37.8	31.1	58.7	140.9
Both	85.1	67.2	38.8	53.7	115.9	73.5	71.9	16.2	33.2	84.4	48.0	43.1	28.7	52.6	125.1
Internal data set															
RV	85.2	70.0	34.2	44.2	73.4	65.4	67.3	18.5	46.6	121.3	48.9	42.8	26.8	53.2	107.0
BEV	88.3	75.0	29.6	34.4	55.9	71.8	69.9	17.6	31.6	76.4	48.3	42.6	26.1	33.3	56.1
Both	89.6	76.7	27.8	32.4	53.4	75.6	73.7	16.9	30.0	73.4	57.9	51.4	24.5	31.7	54.0

Table 3: Comparison of Views in the Backbone Network

View	Vehicle					Pedestrian					Bikes				
	AP (%) \uparrow		L_2 (cm) \downarrow			AP (%) \uparrow		L_2 (cm) \downarrow			AP (%) \uparrow		L_2 (cm) \downarrow		
	0.5 IoU	0.7 IoU	0.0 s	1.0 s	3.0 s	0.1 IoU	0.3 IoU	0.0 s	1.0 s	3.0 s	0.1 IoU	0.3 IoU	0.0 s	1.0 s	3.0 s
nuScenes															
RV	84.8	66.67	39.7	55.12	120.0	76.1	74.4	15.6	31.3	80.3	50.9	47.2	27.4	51.8	128.3
BEV	85.1	67.2	38.8	53.7	115.9	73.5	71.9	16.2	33.2	84.4	48.0	43.1	28.7	52.6	125.1
Both	85.5	67.8	38.2	53.1	115.0	76.4	74.6	15.9	31.6	79.9	49.5	44.5	28.9	54.3	131.6
Internal data set															
RV	90.2	77.4	27.0	31.8	53.3	79.1	77.1	16.3	29.6	73.7	63.9	56.4	23.2	32.9	62.9
BEV	89.6	76.7	27.8	32.4	53.4	75.6	73.7	16.9	30.0	73.4	57.9	51.4	24.5	31.7	54.0
Both	90.8	78.4	26.1	30.6	51.4	79.7	77.8	16.1	28.8	71.6	64.5	57.9	22.7	30.2	53.2

Table 4: Comparison of Multi-View Temporal Fusion Strategies

Strategy	Vehicle					Pedestrian					Bikes				
	AP (%) \uparrow		L_2 (cm) \downarrow			AP (%) \uparrow		L_2 (cm) \downarrow			AP (%) \uparrow		L_2 (cm) \downarrow		
	0.5 IoU	0.7 IoU	0.0 s	1.0 s	3.0 s	0.1 IoU	0.3 IoU	0.0 s	1.0 s	3.0 s	0.1 IoU	0.3 IoU	0.0 s	1.0 s	3.0 s
nuScenes															
One Shot	84.3	66.3	40.1	56.0	120.3	74.5	72.7	16.1	33.6	86.1	46.6	42.2	29.3	58.4	142.4
Sequential	85.5	67.8	38.2	53.1	115.0	76.4	74.6	15.9	31.6	79.9	49.5	44.5	28.9	54.3	131.6
Internal data set															
One Shot	90.6	78.1	26.5	31.4	52.2	78.8	76.9	16.2	29.7	74.0	62.5	56.2	22.8	33.6	63.0
Sequential	90.8	78.4	26.2	30.6	51.4	79.7	77.8	16.1	28.8	71.6	64.5	57.9	22.7	30.2	53.2

from all the warped views which are then used to learn independent spatio-temporal BEV and RV features and are then fed as input to the backbone network (Figure 2b). The major difference lies in the absence of the sequential fusion in both views. For a fair comparison, we ensure that the number of parameters in each per-view network is same as the total parameters in the corresponding *sequential* temporal fusion. As we can see from Table 4, our *sequential* approach can better utilize multiple views for fusion of the temporal sequence of LiDAR data. We note that our approach has a larger relative improvement on nuScenes as compared to the internal data set. We attribute this to the fact that in nuScenes the information loss resulting from the temporal fusion stage in RV [3] has higher impact than when using higher resolution LiDAR which provides more redundancy.

4.5. Run-time Analysis

We report the run-time results using a Titan RTX GPU. Our method can process the operating range of 50m on nuScenes in ~ 30 ms and the range of 100m on our internal data set in ~ 55 ms. In contrast, the previous BEV-only method [4] runs on the shorter range of 50m in ~ 38 ms [5].

BEV-only methods do not scale well with range and have not reported numbers on larger operating ranges of 100m.

RV-only methods [3, 6] have shown the ability to scale better with larger ranges than BEV-only methods. These methods reportedly process the range of 100m in ~ 60 ms. As compared to them, we achieve faster runtime of ~ 55 ms. Therefore, our method combines the runtime advantages that RV-only methods enjoy, with better detection and motion forecasting performance of BEV-only methods. Finally, our method can finish processing all the data and produce output for each sweep before the arrival of the next sweep at 10Hz. Therefore, it is suitable for real-time on-board operations as it exhibits no latency related loss of data.

4.6. Qualitative Analysis

We present a qualitative comparison of our proposed multi-view model with a single-view BEV-only method, in Figure 4. While detection of vehicles is similar between the two methods, *MVFuseNet* more accurately detects pedestrians. Also, we show a few cases where our method is able to improve the motion prediction of vehicles and pedestrians

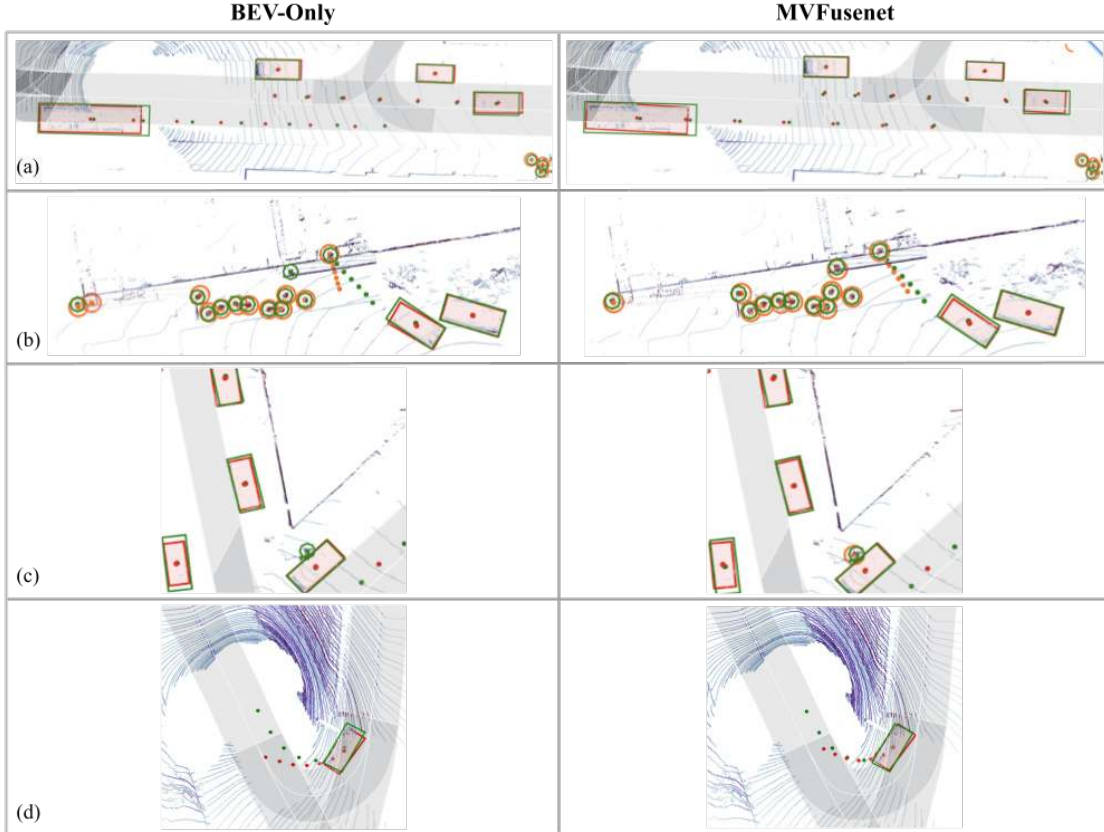


Figure 4: Qualitative comparison of proposed *MVFusenet* with the *BEV-only* model which uses BEV in both temporal fusion and the backbone. Model outputs for detections and trajectories are depicted in **red** for vehicles and in **orange** for pedestrians. The ground truth is depicted in **green**. In (a), the *MVFusenet* model produces better quality motion forecasts as compared to the BEV-only method for a moving vehicle (middle left). In (b), the BEV-only method exhibits multiple failure modes for pedestrians pertaining to a false positive (left), a false negative (top middle), and an inaccurate trajectory for the moving pedestrian (top right), while *MVFusenet* exhibits the correct behavior. The example in (c) shows the BEV-only model failing to detect a pedestrian adjacent to the vehicle. Finally, in (d) we see that both models fail to accurately predict the position at 3s for a vehicle turning sharply, but the proposed model more accurately predicts the turning behaviour.

over single-view BEV-only method.

5. Conclusion and Future Work

We presented a novel multi-view model for end-to-end object detection and motion forecasting. We introduced a unique method for multi-view temporal fusion, as well as a novel multi-view backbone network. We proved the effectiveness of our approach as compared to existing single-view and multi-view fusion methods on two large-scale data sets. We showed that the proposed method can leverage the complementary information in the RV and BEV and improve accuracy on both detection and motion forecasting tasks, while maintaining low latency and scaling to larger operating ranges. In particular, we demonstrated that incorporating both views in temporal fusion and using a sequential fusion approach significantly improves perfor-

mance over only using BEV. Finally, we established a new state-of-the-art result on the publicly available nuScenes data set for joint detection and forecasting.

In addition, we have demonstrated that the presented LiDAR-only approach outperforms multi-sensor approaches which rely on RADAR or camera. However, as future work, we plan to incorporate these additional sensors to improve the robustness of the proposed approach. Additionally, we used a simple uncertainty representation and forecasting method to simplify the experimentation. In the future, we plan to incorporate recent advances in multi-modal motion forecasting and actor-scene interactions.

References

- [1] 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, 5, 6
- [2] Sergio Casas, Wenjie Luo, and Raquel Urtasun. IntentNet: Learning to predict intention from raw sensor data. In *Proceedings of the Conference on Robot Learning (CoRL)*, 2018. 1, 2, 4, 5, 6
- [3] Ankit Laddha, Shivam Gautam, Gregory P Meyer, and Carlos Vallespi-Gonzalez. Rv-fusenet: Range view based fusion of time-series lidar data for joint 3d object detection and motion forecasting. *arXiv preprint arXiv:2005.10863*, 2020. 1, 2, 4, 5, 6, 7
- [4] Nemanja Djuric, Henggang Cui, Zhaoen Su, Shangxuan Wu, Huahua Wang, Fang-Chieh Chou, Luisa San Martin, Song Feng, Rui Hu, Yang Xu, Alyssa Dayan, Sidney Zhang, Brian C. Becker, Gregory P. Meyer, Carlos V-Gonzalez, and Carl K. Wellington. Multixnet: Multiclass multistage multimodal motion prediction, 2020. 1, 2, 4, 5, 6, 7
- [5] Meet Shah, Zhiling Huang, Ankit Laddha, Matthew Langford, Blake Barber, Sidney Zhang, Carlos Vallespi-Gonzalez, and Raquel Urtasun. Liranet: End-to-end trajectory prediction using spatio-temporal radar fusion. *arXiv preprint arXiv:2010.00731*, 2020. 1, 2, 5, 6, 7
- [6] Gregory P Meyer, Jake Charland, Shreyash Pandey, Ankit Laddha, Carlos Vallespi-Gonzalez, and Carl K Wellington. Laserflow: Efficient and probabilistic object detection and motion forecasting. *arXiv preprint arXiv:2003.05982*, 2020. 1, 2, 4, 6, 7
- [7] Gregory P. Meyer, Ankit Laddha, Eric Kee, Carlos Vallespi-Gonzalez, and Carl K. Wellington. LaserNet: An efficient probabilistic 3D object detector for autonomous driving. In *Proceedings of the IEEE CVPR*, 2019. 1, 2, 4
- [8] Alex H. Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. PointPillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE CVPR*, 2010. 2, 4
- [9] Yin Zhou, Pei Sun, Yu Zhang, Dragomir Anguelov, Jiyang Gao, Tom Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *Conference on Robot Learning*, 2020. 2
- [10] Sudeep Fadadu, Shreyash Pandey, Darshan Hegde, Yi Shi, Fang-Chieh Chou, Nemanja Djuric, and Carlos Vallespi-Gonzalez. Multi-view fusion of sensor data for improved perception and prediction in autonomous driving. *arXiv preprint arXiv:2008.11901*, 2020. 2, 5, 6
- [11] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN: 3D object proposal generation and detection from point cloud. In *Proceedings of the IEEE CVPR*, 2019. 2
- [12] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. STD: Sparse-to-dense 3D object detector for point cloud. In *Proceedings of the IEEE ICCV*, 2019. 2
- [13] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-end learning for point cloud based 3D object detection. In *Proceedings of the IEEE CVPR*, 2018. 2
- [14] Yan Yan, Yuxing Mao, and Bo Li. SECOND: Sparsely embedded convolutional detection. *Sensors*, 2018. 2
- [15] Bin Yang, Wenjie Luo, and Raquel Urtasun. PIXOR: Real-time 3D object detection from point clouds. In *Proceedings of the IEEE CVPR*, 2018. 2
- [16] Yang Zhang, Zixiang Zhou, Philip David, Xiangyu Yue, Zerong Xi, Boqing Gong, and Hassan Foroosh. Polarnet: An improved grid representation for online lidar point clouds semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9601–9610, 2020. 2
- [17] Bo Li, Tianlei Zhang, and Tian Xia. Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916*, 2016. 2
- [18] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. Rangenet++: Fast and accurate lidar semantic segmentation. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4213–4220. IEEE, 2019. 2
- [19] Deyvid Kochanov, Fatemeh Karimi Nejadasl, and Olaf Booij. Kprnet: Improving projection-based lidar semantic segmentation. *arXiv preprint arXiv:2007.12668*, 2020. 2
- [20] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3D object detection network for autonomous driving. In *Proceedings of the IEEE CVPR*, 2017. 2
- [21] Yin Zhou, Pei Sun, Yu Zhang, Dragomir Anguelov, Jiyang Gao, Tom Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *Conference on Robot Learning*, pages 923–932. PMLR, 2020. 2, 4
- [22] Ke Chen, Ryan Oldja, Nikolai Smolyanskiy, Stan Birchfield, Alexander Popov, David Wehr, Ibrahim Eden, and Joachim Pehserl. Mvlidarnet: Real-time multi-class scene understanding for autonomous driving using multiple views. *arXiv preprint arXiv:2006.05518*, 2020. 2
- [23] Venice Erin Liong, Thi Ngoc Tho Nguyen, Sergi Widjaja, Dhananjai Sharma, and Zhuang Jie Chong. Amvnet: Assertion-based multi-view fusion network for lidar semantic segmentation. *arXiv preprint arXiv:2012.04934*, 2020. 2
- [24] Zhidong Liang, Ming Zhang, Zehan Zhang, Xian Zhao, and Shiliang Pu. Rangercnn: Towards fast and accurate 3d object detection with range image representation. *arXiv preprint arXiv:2009.00206*, 2020. 2
- [25] Nemanja Djuric, Vladan Radosavljevic, Henggang Cui, Thi Nguyen, Fang-Chieh Chou, Tsung-Han Lin, and Jeff Schneider. Motion prediction of traffic actors for autonomous driving using deep convolutional networks. *arXiv preprint arXiv:1808.05819*, 2018. 2

- [26] 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 CVPR*, 2016. 2
- [27] Nachiket Deo and Mohan M Trivedi. Convolutional social pooling for vehicle trajectory prediction. In *Proceedings of the IEEE CVPR Workshops (CVPRW)*, 2018. 2
- [28] Namhoon Lee, Wongun Choi, Paul Vernaza, Christopher B Choy, Philip HS Torr, and Manmohan Chandraker. DESIRE: Distant future prediction in dynamic scenes with interacting agents. In *Proceedings of the IEEE CVPR*, 2017. 2
- [29] Gregory P. Meyer, Jake Charland, Darshan Hegde, Ankit Laddha, and Carlos Vallespi-Gonzalez. Sensor fusion for joint 3D object detection and semantic segmentation. In *Proceedings of the IEEE CVPR Workshops (CVPRW)*, 2019. 2
- [30] Karttikeya Mangalam, Harshayu Girase, Shreyas Agarwal, Kuan-Hui Lee, Ehsan Adeli, Jitendra Malik, and Adrian Gaidon. It is not the journey but the destination: Endpoint conditioned trajectory prediction. *arXiv preprint arXiv:2004.02025*, 2020. 2
- [31] Hang Zhao, Jiyang Gao, Tian Lan, Chen Sun, Benjamin Sapp, Balakrishnan Varadarajan, Yue Shen, Yi Shen, Yuning Chai, Cordelia Schmid, et al. Tnt: Target-driven trajectory prediction. *arXiv preprint arXiv:2008.08294*, 2020. 2
- [32] Yuning Chai, Benjamin Sapp, Mayank Bansal, and Dragomir Anguelov. MultiPath: Multiple probabilistic anchor trajectory hypotheses for behavior prediction. *arXiv preprint arXiv:1910.05449*, 2019. 2
- [33] Tung Phan-Minh, Elena Corina Grigore, Freddy A Boulton, Oscar Beijbom, and Eric M Wolff. Covnet: Multimodal behavior prediction using trajectory sets. In *Proceedings of the IEEE CVPR*, 2020. 2
- [34] Nicholas Rhinehart, Rowan McAllister, Kris Kitani, and Sergey Levine. Precog: Prediction conditioned on goals in visual multi-agent settings. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2821–2830, 2019. 2
- [35] Charlie Tang and Russ R Salakhutdinov. Multiple futures prediction. In *Proceedings of Advances in Neural Information Processing Systems (NIPS)*, 2019. 2
- [36] 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)*. IEEE, 2019. 2
- [37] 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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11525–11533, 2020. 2
- [38] Cunjun Yu, Xiao Ma, Jiawei Ren, Haiyu Zhao, and Shuai Yi. Spatio-temporal graph transformer networks for pedestrian trajectory prediction. In *European Conference on Computer Vision*, pages 507–523. Springer, 2020. 2
- [39] Agrim Gupta, Justin Johnson, Li Fei-Fei, Silvio Savarese, and Alexandre Alahi. Social gan: Socially acceptable trajectories with generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2255–2264, 2018. 2
- [40] Ming Liang, Bin Yang, Rui Hu, Yun Chen, Renjie Liao, Song Feng, and Raquel Urtasun. Learning lane graph representations for motion forecasting. *arXiv preprint arXiv:2007.13732*, 2020. 2
- [41] 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 *Proceedings of the IEEE CVPR*, 2018. 2
- [42] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE CVPR*, 2016. 4
- [43] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. *arXiv preprint arXiv:2006.11275*, 2020. 5
- [44] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE ICCV*, 2017. 5
- [45] Gregory P Meyer and Niranjan Thakurdesai. Learning an uncertainty-aware object detector for autonomous driving. *arXiv preprint arXiv:1910.11375*, 2019. 5
- [46] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. NuScenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*, 2019. 5
- [47] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *arXiv preprint arXiv:1912.01703*, 2019. 5