

# StreamingFlow: Streaming Occupancy Forecasting with Asynchronous Multi-modal Data Streams via Neural Ordinary Differential Equation

Yining Shi<sup>1,3\*</sup> Kun Jiang<sup>1†</sup> Ke Wang<sup>2</sup> Jiushi Li<sup>1</sup> Yunlong Wang<sup>1</sup> Mengmeng Yang<sup>1</sup> Diange Yang<sup>1†</sup>  
<sup>1</sup> School of Vehicle and Mobility, Tsinghua University <sup>2</sup> KargoBot, Inc <sup>3</sup> DiDi Chuxing

## Abstract

Predicting the future occupancy states of the surrounding environment is a vital task for autonomous driving. However, current best-performing single-modality methods or multi-modality fusion perception methods are only able to predict uniform snapshots of future occupancy states and require strictly synchronized sensory data for sensor fusion. We propose a novel framework, StreamingFlow, to lift these strong limitations. StreamingFlow is a novel BEV occupancy predictor that ingests asynchronous multi-sensor data streams for fusion and performs streaming forecasting of the future occupancy map at any future timestamps. By integrating neural ordinary differential equations (N-ODE) into recurrent neural networks, StreamingFlow learns derivatives of BEV features over temporal horizons, updates the implicit sensor’s BEV features as part of the fusion process, and propagates BEV states to the desired future time point. It shows good zero-shot generalization ability of prediction, reflected in the interpolation of the observed prediction time horizon and the reasonable inference of the unseen farther future period. Extensive experiments on two large-scale datasets, nuScenes [2] and Lyft L5 [14], demonstrate that StreamingFlow significantly outperforms previous vision-based, LiDAR-based methods, and shows superior performance compared to state-of-the-art fusion-based methods.

## 1. Introduction

Occupancy grid is gaining more traction in the self-driving community, due to its versatilities in downstream tasks, e.g., irregularly shaped object representation, robot navigation, etc. With the help of modern deep learning technologies, occupancy grid maps (OGMs) have evolved significantly. For example, not just LiDAR point clouds, camera images can now be used to construct high-quality occupancy grid maps [26]. In addition to its fine-grained geometric infor-

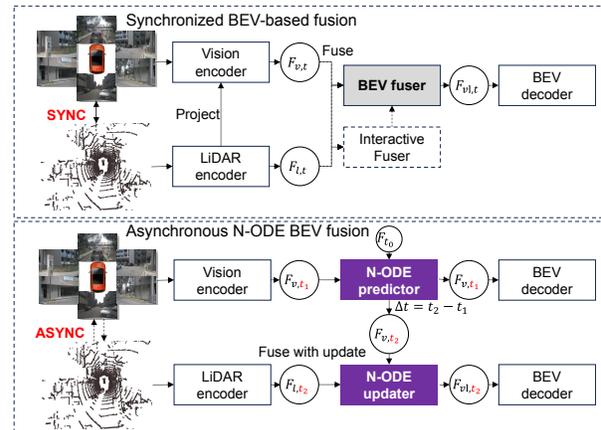


Figure 1. Comparison between conventional synchronized BEV fusion (top) and our asynchronous BEV fusion (bottom). We formulate fusion with an update-predict-update approach.

mation, occupancy grids can also represent rich semantic and motion cues [1, 9–11]. Such cues grant occupancy networks the ability to predict the occupancy state of the environment at desired future timestamps. While trajectory forecasting only contains object-level information, the occupancy grid conveys a much more comprehensive understanding of the surrounding environment, which is the key to achieving better intelligence and driving safety.

Despite the great potential of occupancy-based prediction, we notice that existing methods can only predict at fixed frequencies [1, 9–11, 33]. This comes from the fact that sensors sample the environment at predetermined frequencies, e.g., motors in mechanical LiDARs spin at 10Hz, and many cameras capture images at 30FPS. Thus, previous algorithms predict at the same intervals as they are trained with ground-truth labels of future intervals. If future occupancy can be predicted at any given timestamps in a continuous manner, self-driving algorithms can have shorter latency and more rapid responses. Streaming forecasting can also relax the sensor synchronization constraints, as the continuous forecasting model can update the prediction as new sensory data come in. Compared with discrete snapshot-based prediction, streaming motion forecast-

\*Work done during internship at DiDi Chuxing.

† Corresponding author: Kun Jiang, Diange Yang.

ing [24] naturally has safety advantages.

We noticed several challenges in fulfilling such a streaming forecasting paradigm. Firstly, sensory data and ground truth annotations are collected at a fixed frequency. Streaming forecasting aims to achieve temporally dense and continuous prediction, with data only sparsely and uniformly sampled in the temporal domain. Secondly, mainstream network architectures are not designed for such streaming requirements. Forecasting at any given timestamp requires delicate modeling of the temporal dynamics, but how to embed the temporal dynamics into the widely adopted BEV feature representation remains under-explored.

To this end, we propose to use the neural ordinary differential equation (N-ODE) framework for such temporal dynamics modeling. N-ODE is originally proposed for learning continuous implicit processes in dynamic systems. In our application, by representing the continuous motion trends of the driving scenarios into a time series of implicit spatial BEV features, the neural ODE framework not only allows streaming occupancy predictions on any temporal horizon with one fixed model, but also relaxes the sensor synchronization constraints.

With such flexibility, we propose a novel method, StreamingFlow, that performs asynchronous multi-modal sensor fusion to achieve streaming occupancy forecasting. Toward the goal of streaming modeling of temporal BEV features, StreamingFlow integrates neural ordinary differential equations (N-ODE) onto wrapped-based recurrent neural networks. It learns derivatives of BEV features over temporal horizons, updates the implicit sensor’s BEV feature as part of the fusion process, and propagates BEV states to the desired future time point. A graphical illustration of fusion strategies between conventional BEV-based fusion and ours is illustrated in Fig. 1.

Our contributions can be summarized as follows:

1. To the best of our knowledge, we propose the first streaming occupancy flow prediction framework that supports asynchronous multi-sensor fusion;
2. We design a novel temporal feature propagation strategy that achieves high performance on dense continuous occupancy prediction with only temporally sparse labels;
3. Our proposed method achieves state-of-the-art performance on the widely used nuScenes [2] and Lyft L5 [14] datasets, validating the effectiveness and robustness of our proposed algorithms.

## 2. Related Works

### 2.1. Occupancy forecasting

Occupancy flow prediction, first proposed in [23], is a valuable supplement to trajectory-set prediction [3, 22]. The advantages of occupancy flow are a non-parametric distribution output and fine-grained grid flow. Recent LiDAR-

based methods [21, 30, 32] attach displacement vectors to non-empty BEV grids to describe short-term motion in future 1.0s. These networks typically have a BEV-based spatial-temporal backbone and several shared heads that indicate the classification, motion, and displacement of each grid. STPN [32] and BE-STI [30] are effective BEV-based spatial-temporal feature extractors for binary-occupied grids which consist of only 2D convolutional blocks. Vision-centric pipelines usually formulate the prediction task as future instance segmentation on BEV grids. FIERY [9] is the first vision-centric stochastic future prediction method. StretchBEV [1] uses a variational autoencoder for decoupling learning of temporal dynamics and BEV decoder. End-to-end driving frameworks [10, 11, 33, 34] conduct planning based on occupancy prediction. UniAD [11] and TBP-Former [6] models temporal future with video transformers. One major bottleneck of video transformers is prediction of long-term sequence length [28]. Position encoding enables video transformer to interpolate shorter sequences but few video transformers can generalize to the unseen longer sequence length. Unlike all prior works that predict discrete near-future snapshots, We focus on the streaming flow prediction which adapt the predictor to a more distant future and finer granularity.

### 2.2. Multi-sensor Spatial-temporal Fusion

Spatial multi-modal fusion is widely investigated in the feature level. BEVFusion and its variants [7, 8, 20] combine multi-modal input features in a shared BEV space. [4, 16] support feature interaction in 3D voxel space. Temporal fusion is a common aggregator for perception tasks. LiDAR-based networks [30, 32] typically employ data-level temporal fusion, which transforms multi-frame point clouds to the current ego’s coordinate as inputs via ego pose or registration methods. Most vision-centric methods align the BEV feature maps from multiple time steps based on ego pose, and then concatenate them for fusion [9, 12, 17, 25] or support their interactions of multi-frame BEV features through attention [13, 18, 27]. Prior spatial fusion methods usually assume strict synchronization. CoBEVFlow [31] is most similar to our approach in that they predict BEV flow to interpolate asynchronous data timestamps from roadside unit to vehicle system, but short-term flow information cannot support long-term prediction. To this end, we provide a flexible spatio-temporal fusion framework which predicts the future frames and triggers update process as fusion when asynchronous sensor data arrives.

## 3. Methodology

### 3.1. Framework

The pipeline of StreamingFlow is depicted in Fig. 2. Given multi-modal data streams from the previous few seconds,

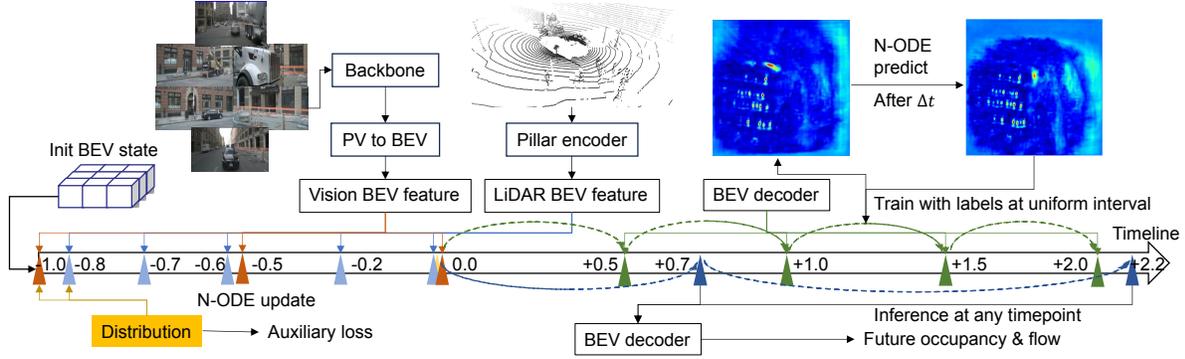


Figure 2. The framework of StreamingFlow. Raw data streams are encoded to BEV features, respectively. The SpatialGRU-ODE process operates on the timeline with two stages split by the present timestamp, asynchronous multi-sensor deep feature via the SpatialGRU-ODE update process and continuous occupancy flow prediction via SpatialGRU-ODE predict process.

the task is defined as predicting the future several seconds’ instance occupancy and flow on uniformly-shaped BEV grids. The framework includes three phases: 1) BEV encoders for LiDAR and camera branches, respectively, 2) asynchronous multi-sensor fusion, and 3) streaming occupancy prediction. As the data stream runs, modality-specific encoders map raw data to BEV features. The BEV state starts from zeroed features and fuses every incoming BEV feature in a trigger mode instead of a matching mode in conventional fusion. The BEV state propagates until a new BEV feature is received. The incoming BEV feature is fused to the BEV state and state continues to propagate until next incoming observation.

The SpatialGRU-ODE is the core for streaming prediction. The SpatialGRU-ODE iterates for ODE steps, updates when a new measurement occurs within a certain ODE step, and performs future prediction. In the training stage, the loss comes from the supervision of decoded BEV grids at the timestamps with labels, as well as an auxiliary probabilistic loss from Kullback-Leibler divergence (KLD) between updated BEV features and latent observations. In the inference process, SpatialGRU-ODE predicts via a variable ODE step only related to the required timestamps for evaluation and application.

The BEV decoder, loss functions, and post-processing techniques follow standard practices, and we leave the details in the supplementary materials.

### 3.2. BEV Encoders

**LiDAR branch.** We adopt the popular pillar representations for point cloud features. We observed that a very lightweight pillar-net encoder is sufficient for high accuracy. Heavier backbones consisting of sparse convolutional modules or voxel transformers can be trivially integrated for performance at the cost of additional computation.

**Camera branch.** As a baseline for algorithms aligned with previous work, we adopt the Lift-splat-shoot

(LSS [26]) settings with depth supervision for depth-based view transformation. We adopt accelerated BEV pool operators to speed up the inference stage. Other BEV mapping methods are trivially compatible with our framework.

### 3.3. SpatialGRU-ODE

Gate Recurrent Unit (GRU) consists of reset gate ( $r_t$ ), update gate ( $z_t$ ), and update vector ( $g_t$ ). The update process follows the formula,

$$\begin{aligned} r_t &= \sigma(W_r x_t + U_r h_{t-1} + b_r) \\ z_t &= \sigma(W_z x_t + U_z h_{t-1} + b_z) \\ g_t &= \tanh(W_h x_t + U_h (r_t \odot h_{t-1}) + b_h) \end{aligned} \quad (1)$$

Neural-ODE on GRU blocks focuses N-ODE’s ability to model continuous time series given sporadic sensor data coming from non-uniform time intervals. SpatialGRU-ODE is highly inspired by GRU-ODE-Bayes [5], which serves as the foundation component of the proposed spatiotemporal fusion and flow predictor. Eq. (1) formulates the general process of GRU blocks. With regard to  $h_t$  as a BEV feature, SpatialGRU has its state variable  $h_t$  with shape  $[B, C, H, W]$  where  $B, C, H, W$  are the batch size, embedded dims, height, and width.  $\sigma$  is the sigmoid function, and  $\tanh$  denotes a CNN block in SpatialGRU settings.  $x_t$  is the input state at timestamp  $t$  with the same shape as  $h_t$ . For simplicity, the update gate  $W_r = U_r$ , the reset gate  $W_z = U_z$ , and the GRU init bias  $b_r = b_z$ .  $r_t \odot h_{t-1}$  is the matrix multiplication between reset gate and state. In summary, the update state  $h_t$  is calculated as,

$$h_t = z_t \cdot h_{t-1} + (1 - z_t) \cdot g_t = GRU(h_{t-1}, g_t) \quad (2)$$

The amount of change of the state variable is calculated as:

$$\begin{aligned} \Delta h_t &= h_t - h_{t-1} = z_t \cdot h_{t-1} + (1 - z_t) \cdot g_t - h_{t-1} \\ &= (1 - z_t) \cdot (g_t - h_{t-1}) \end{aligned} \quad (3)$$

The derivatives of the state variable are calculated as:

$$\frac{\Delta h(t)}{dt} = (1 - z(t)) \cdot (g(t) - h(t)) \quad (4)$$

---

**Algorithm 1** The pseudo-code of SpatialGRU-ODE
 

---

**Input:** state-timestamp pairs  $[t, m[t]], t \in [t_0, \dots, t_{present}]$ , ODE step  $\Delta t$

**Output:**  $loss_{KLD}$ , state-timestamp pairs  $[t, h[t]], t \in [t_0, \dots, t_{present}, \dots, t_{future}]$ .

- 1: Initialize current time  $T = t_0$ , KL divergence loss  $loss_{KLD} = 0$ , current state feature  $h = h_0$
  - 2: **for**  $t_{obs}$  in  $[t_0, \dots, t_{present}]$  **do**
  - 3:   **while**  $T < t_{obs}$  **do**
  - 4:      $h = \text{ODE\_Predict}(h, \Delta t)$
  - 5:      $T+ = \Delta t$
  - 6:   **end while**
  - 7:    $h, loss = \text{ODE\_Update}(h, m[t_{obs}])$
  - 8:    $loss_{KLD} + = loss$
  - 9:   save  $h[t_{obs}] = h$
  - 10: **end for**
  - 11: **for**  $t_{predict}$  in  $[t_{present+1}, \dots, t_{future}]$  **do**
  - 12:   **while**  $T < t_{predict}$  **do**
  - 13:      $h = \text{ODE\_Predict}(h, \Delta t)$
  - 14:      $T+ = \Delta t$
  - 15:   **end while**
  - 16:   save  $h[t_{predict}] = h$
  - 17: **end for**
- 

GRU-ODE-Bayes has proven the convergence of this process, as is also applicable to the initial state of SpatialGRU-ODE. If  $h_0 \in [-1, 1]$ , then  $h_j(t) \in [-1, 1]$ , where  $j$  is the index of elements in  $h(t)$ , as  $\left. \frac{dh(t)_j}{dt} \right|_{t:h(t)_j=1} \leq 0$  and  $\left. \frac{dh(t)_j}{dt} \right|_{t:h(t)_j=-1} \geq 0$ .

N-ODE has several nice properties when implemented on GRU models. The learned parametric derivative makes training not constrained by supervision from fixed future timestamps states, which greatly increases flexibility. On the other hand, unlike RNNs, which require emission intervals for updating, the GRU-ODE continuously defined dynamic model can naturally combine data observed at any given time.

The pseudo-code for SpatialGRU-ODE is shown in Algorithm 1. Given state-timestamp pairs  $[t, m[t]]$  from previous multisensor features, The expected output is the states  $[t, h[t]]$  of each measured past timestamp and anticipated future timestamp  $t \in [t_0, \dots, t_{present}, \dots, t_{future}]$ . The ode step may be variable or constant. During the update procedure, a KL divergence loss  $loss_{KLD}$  is gathered to quantify the similarity between updated state feature  $h$  and observed feature  $m[t_{obs}]$ .

### 3.4. Asynchronous Modality-agnostic Fusion

Asynchronous modality-agnostic fusion attempts to fuse LiDAR and camera BEV features in a way similar to the

Bayesian predict-then-update process. The proposed strategy is able to eliminate two crucial limits for multi-sensor fusion, which are strict multi-modal synchronizations and fixed interval data flow.

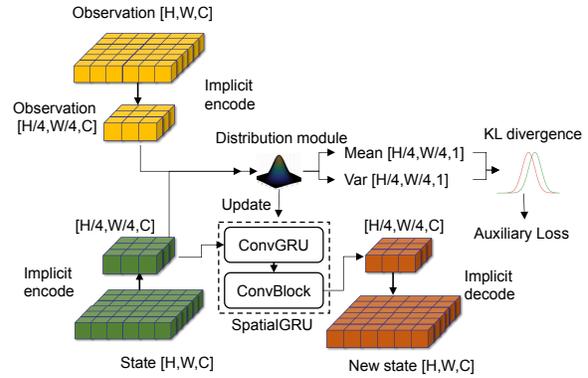


Figure 3. Illustration of the measurement update process of SpatialGRU-ODE in temporal-agnostic fusion.

Fig. 3 presents the network structure for the SpatialGRU-ODE update process. After raw data is encoded to BEV maps by different BEV encoders, BEV features from different modalities are agnostic on the source, ordered chronologically, and formulated as multiple timestamp-BEV-feature pairs  $[t, h_t]$ . The output retains the same shape as the input, but the hidden features are as  $\frac{1}{4}$  size of the input for the sake of eliminating memory storage. In order not to degrade the ultimate performance, both state and observation in BEV feature space with shape  $[B, C, H, W]$  are implicitly encoded into a smaller feature space  $[B, C, H/4, W/4]$  and finally implicit decoded in the same way for occupancy flow decoder. The temporal propagator runs the subsequent ODE step till the arrival of the next observation. When a new observation occurs, the predicted state and new observation are processed into the distribution module with CNN blocks to obtain the mean and variance of BEV features both with the size  $[1, H/4, W/4]$ . Means and variances of both features are estimated using KL divergence, an auxiliary loss that reflects the similarity between predicted and measured characteristics. Probabilistic auxiliary loss is calculated as follows:

$$loss_{KLD} = \sum_{j=1}^{H \times W} D_{KL}(p_{state,j} || p_{meas,j}) \quad (5)$$

The core module for the updated BEV feature is a dual-pathway SpatialGRU module, which first propagates the observed and predicted features respectively, then mixes the two hidden states distributions and finally undergoes the weighted summation of two features after the trusting gate and softmax function to generate the final new state.

### 3.5. Streaming Occupancy Flow Predictor

The occupancy flow predictor handles multi-frame BEV features from past and present timestamp-agnostic fusion as input and propagates to future steps in an ODE-enabled variational recurrent neural network. Particularly, the granularity of temporal prediction is decoupled from the training granularity and only related to the minimal ODE step in this method. If we set the variable ODE step, the prediction enables arbitrary granularity.

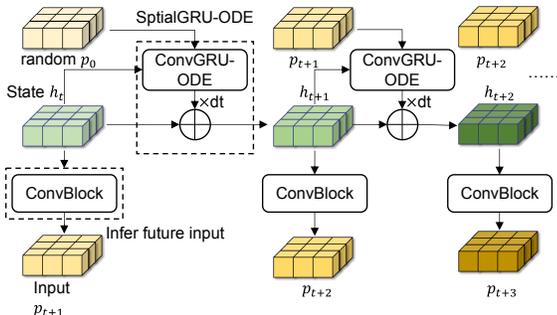


Figure 4. Illustration of streaming prediction process of SpatialGRU-ODE.

Fig. 4 illustrates the SpatialGRU-ODE procedure. Similar to the update process, the states are implicitly encoded into a smaller feature space. Except for the initial input which is the first BEV feature, the input with shape  $[H/4, W/4, C]$  is inherited from the last ODE step. The core module, SpatialGRU-ODE, computes the derivatives of BEV features in unit time intervals via a recurrent block specified in Eq. (4). SpatialGRU-ODE may employ either Euler or Midpoint solver for the update and fixed or variable ODE time step. To obtain the projected states, the derivatives are multiplied by delta time  $dt$  and added to the initial states. A CNN block is used to infer the next input from the current state in order to make the subsequent prediction. Only states near to the needed timestamps with supervision (closer than half of the minimum ode step) are preserved and decoded in the BEV decoder for supervision. Thus, the prediction time step  $dt$  is not always constant and uniform, and supervision signals are not required for each future timestamp.

## 4. Experiments

In this section, we seek to answer two questions with the following experiments: (1) Whether the model works as well as published state-of-the-art algorithms for standard occupancy forecasting tasks. (2) Whether we are able to generalize to streaming forecasting settings with one model trained with temporally-sparse occupancy labels.

### 4.1. Datasets and Metrics

The NuScenes [2] dataset is a public large-scale autonomous driving dataset collected by Motional. NuScenes dataset provides a full sensor suite including 1 top LiDAR, 6 cameras, 5 radars, GPS, and IMU, with  $360^\circ$  coverage of the surroundings. It contains 1000 scenes, each lasting 20 seconds, of which annotations of 850 scenes at 2Hz are available. We use the training and validation splits as same as previous works [1, 9, 10, 30, 34], with 28130 samples for training and 6019 samples for validation.

The Lyft L5 dataset [14] is another large-scale autonomous dataset provided by Lyft. This dataset contains 180 scenes, each lasting 25-45s in length and annotated at 5Hz. We use 6 ring cameras and the top LiDAR from its sensor suite, which has a  $360^\circ$  field of view. We use training and validation splits from the practice of FIERY [9], which include 16000 training samples and 4000 validation samples.

The occupancy flow prediction is formulated as a video panoptic prediction task [15]. We use Intersection over Union (IoU) to measure the semantic segmentation quality of each frame, and video panoptic quality (PQ), recognition quality (RQ), and segmentation quality (SQ) to measure both the accuracy and consistency of instance segmentation. VPQ is computed as:

$$\begin{aligned} \text{VPQ} &= \sum_{t=0}^H \frac{\sum_{(p_t, q_t) \in TP_t} \text{IoU}(p_t \cdot q_t)}{|TP_t| + \frac{1}{2}|FP_t| + \frac{1}{2}|FN_t|} \\ &= \sum_{t=0}^H \frac{|TP_t|}{|TP_t| + \frac{1}{2}|FP_t| + \frac{1}{2}|FN_t|} = \text{SQ} \cdot \text{RQ} \end{aligned} \quad (6)$$

### 4.2. Implementation Details

**Task details.** The standard occupancy prediction task is to take the past 1.0s sensor data for input, which is 3 keyframes in nuScenes (5 keyframes in Lyft), predict the motion in the future 2s, which is 4 keyframes in nuScenes and 10 keyframes in Lyft.

**Data details.** The BEV occupancy labels are from 3D bounding boxes labels projected to BEV. The perception range is set as  $[100m, 100m]$  and the grid resolution is set as  $0.5m$ . Images are resized to  $224 \times 480$  pixels for each frame. Each point cloud is densified by aggregating point clouds of adjacent sweeps. Point clouds are voxelized to  $200 \times 200 \times 13$  voxels. StreamingFlow adopts an asynchronous data stream as input, with LiDAR streams every 0.2s and camera streams every 0.5s as input. For larger models fairly compared with the state-of-the-art, we use images with size  $320 \times 800$  and aggregated LiDAR streams, each with 10 sweeps.

**Model details.** The framework is implemented on Pytorch 1.10.2 and Pytorch Lightning 1.2.5. The optimizer is AdamW with a learning rate of  $1e-4$ . The learning strategy

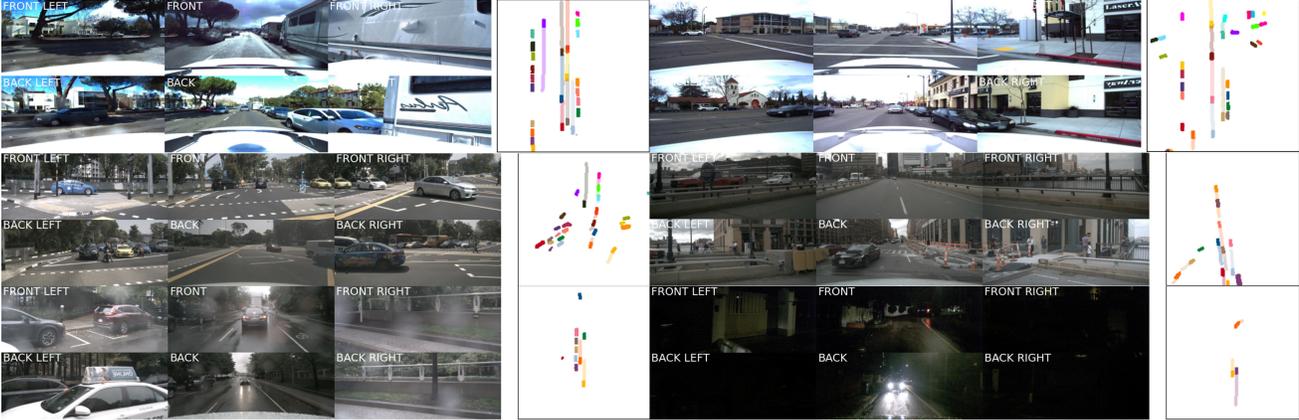


Figure 5. Visualization of StreamingFlow for diverse driving scenarios. Different colors represent different instances of the agents, and lighter colors represent the future occupancy of the agents. (top): samples from Lyft dataset, highway (left), and urban (right). (middle and bottom): samples from nuScenes dataset, sunny (middle left), overcast after rain (middle right), rainy (bottom left), and night (bottom right). StreamingFlow works well in all challenging driving scenarios.

is cosine annealing with a weight decay of 0.01. All models are trained on 8 A6000 GPUs for 20 epochs with a total batch size 16. The image backbone is EfficientNet-b4 [29] for tiny and base model and Effi-b7 for small model.

### 4.3. Comparison with the State-of-the-art

Comparisons of methods implemented in the nuScenes dataset are shown in Tab. 1. All methods are trained for a standard instance flow segmentation and prediction task. Fig. 5 illustrates some examples covering diverse scenarios.

StreamingFlow achieves competitive performance while significantly outperforming other single-modality methods. StreamingFlow surpasses the best vision-based BEVerse [34] with a larger backbone Swin-small [19] by +7.0 IoU and +10.0 VPQ, and the best LiDAR-based BE-STI [30] +7.7 IoU and 5.1VPQ. Compared to FusionAD, StreamingFlow-base uses the same LiDAR backbone(SpConv), similar image backbone(Effi-B4 vs. ResNet101) but a half smaller image size( $800 \times 320$  vs.  $1600 \times 900$ ), and is +1.6 VPQ and 2.4 IoU higher than the latest fusion-based FusionAD [33].

Comparisons of methods implemented in the Lyft L5 dataset are shown in Tab. 2. StreamingFlow achieves the best of all metrics while surpassing baseline methods by a large margin: +20.6 IoU and +23.5 VPQ against ST-P3 [10]. Moreover, it surpasses our reproduced BEVFusion with lightweight encoders by +2.3 mIoU and 0.2 VPQ.

### 4.4. Streaming Forecasting Results

We design three experiments on nuScenes dataset for demonstration of the streaming occupancy prediction: (1) Extension of the prediction temporal horizons to unseen future; (2) Prediction at any desired future time point; (3) Data streams input at different frame rates. Experiments are con-

ducted on models loaded from the checkpoints trained on standard instance forecasting tasks.

**Extension of the temporal horizons of unseen future.** Tab. 3 shows the prediction result of StreamingFlow which extends the temporal horizons from 2.0s to 8.0s. Baselines reported in [1] are trained for different lengths of temporal horizons. StreamingFlow shows excellent performance in variable prediction length without retraining. Similar to prior arts which extend prediction horizons to farther future, StreamingFlow shows a similar gradual performance decay. We then compare zero-shot Streamingflow with the model trained with the longest future 8s groundtruth labels in Tab. 4. The zero-shot long-term prediction results are only  $-0.9$  IoU and  $-1.8$  VPQ less than fully supervised models, which demonstrates great generalization ability of foreseeing the unseen future.

**Streaming prediction at any desired time point.** Since nuScenes do not provide labels for non-keyframes, we adopt the practice of MotionNet and interpolate the box instances between two adjacent keyframes. We randomly choose to predict at a list of time points and report the results in Tab. 5. To simply illustrate the capability, we designed the main task to predict uniformly spaced temporal horizons, but this framework can predict uneven and arbitrary future snapshots. The prediction results are stable generally. There is a slight decline as the forecast becomes denser.

We provide demo videos as supplementary materials to show the effectiveness of streaming prediction. We provide a finer granularity of streaming perception at any future time. In the video, we visualize the prediction results every 0.05s from the current to the future 2s to 3s, with a total of 40 to 60 frames.

**Asynchronous multi-modal data streams as input.**

Method	Modality	Image backbone	LiDAR backbone	Future Semantic Seg.	Future Instance Seg.		
				IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
Static*	C	Effi-B4	-	32.2	27.6	70.1	39.1
FIERY*[9]	C	Effi-B4	-	37.0	30.2	70.2	42.9
StretchBEV[1]	C	Effi-B4	-	37.1	29.0	-	-
BEVerse[34]	C	Swin-tiny	-	38.7	33.3	70.6	47.2
BEVerse[34]	C	Swin-small	-	40.8	36.1	70.7	51.1
ST-P3 Gaus.[10]	C	Effi-B4	-	38.6	31.7	70.2	45.2
ST-P3 Ber.[10]	C	Effi-B4	-	38.9	32.1	70.4	45.6
UniAD[11]	C	Effi-B4	-	40.2	33.5	-	-
MotionNet[32]	L	-	STPN	37.2	38.9	75.6	51.4
BE-STI[30]	L	-	BESTI-STPN	40.1	41.0	75.5	54.3
StreamingFlow-tiny	LC	Effi-B4	PillarNet	47.8	46.1	75.8	60.8
StreamingFlow-small	LC	Effi-B7	PillarNet	51.1	48.0	75.1	64.0
FusionAD[33]	LC	ResNet-101	VoxelNet	51.5	51.1	-	-
StreamingFlow-base	LC	Effi-B4	VoxelNet	53.9	52.7	78.7	66.9

Table 1. Comparison with the state-of-the-art methods for instance-aware occupancy flow prediction on nuScenes[2] validation set. 'C' denotes vision-only methods, 'L' denotes LiDAR-based methods, and 'LC' denotes a LiDAR-camera fusion methods. \*: reported in [9].

Method	IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
Static*	24.1	20.7	-	-
Extrapolation model*	24.8	21.2	-	-
FIERY*[9]	36.3	27.2	-	-
ST-P3[10]	36.3	32.4	71.1	45.5
BEVFusion-tiny $\dagger$	54.6	55.7	78.0	71.4
StreamingFlow-tiny	56.9	55.9	78.1	71.6

Table 2. Comparison with the state-of-the-art methods for future instance segmentation (2.0s) on Lyft L5 AV [14] validation set. \*: results reported in [9].  $\dagger$ : results of BEVFusion with LSS-based image encoder and pillar-based encoder reproduced by us.

Method	Pred	IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
FIERY	2s	35.8	29.0	-	-
	4s	30.1	23.6	-	-
	6s	26.7	20.9	-	-
StretchBEV	2s	37.1	29.0	-	-
	4s	32.5	23.8	-	-
	6s	28.4	21.0	-	-
StreamingFlow	2s	47.8	46.1	75.8	60.8
	3s	44.4	41.1	74.1	55.4
	4s	41.7	38.4	73.6	52.1
	5s	39.1	35.8	73.2	48.9
	6s	36.7	33.6	72.9	46.0
	8s	32.5	29.8	72.6	40.1

Table 3. Comparison of occupancy flow prediction for any temporal horizons.

We validate asynchronous data streams as input and standard occupancy forecasting as output. We test multiple settings of data stream frame rate and report the results in Tab.

Schedule	2s	3s	4s	6s	8s
ZS.	47.8/46.1	44.4/41.1	41.7/38.4	36.7/33.6	32.5/29.8
Sup.	48.2/46.9	46.0/45.2	41.9/38.7	37.8/34.2	33.6/31.6

Table 4. ZS: zero-shot inference of StreamingFlow. Sup: StreamingFlow trained with the longest prediction horizons 8s. Each value is given in the format of 'IoU/PQ'.

Prediction Interval	IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
0.5	47.8	46.1	75.8	60.8
0.25	43.4	40.1	74.1	54.1
0.6	45.6	44.3	75.2	58.9

Table 5. Prediction results at any desired intervals. For example, prediction at the interval of 0.6 means prediction at [0.6, 1.2, 1.8].

6. The model performs the best when the inference uses the training configuration. Denser inputs lead to similar forecasting accuracy.

LiDAR stream	Cam stream	IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
5	2	47.8	46.1	75.8	60.8
10	2	47.6	45.8	75.6	60.5
10	4	47.1	45.4	75.1	60.4

Table 6. Prediction results with multi-modal data stream of different frame rates. Frame rates are shown in Hz.

## 4.5. Ablation Study

**Effect of different fusion strategies.** We compare SpatialGRU-ODE with fusion strategies at different stages and the implementation details are in supplementary

material. As shown in Tab. 7, asynchronous fusion is inferior to synchronous fusion and SpatialGRU-ODE for both datasets. For nuScenes dataset, the ordinary temporally-synchronized fusion approach performs the best. SpatialGRU-ODE performs  $-2.4$  IoU and  $-0.9$  VPQ, slightly less than BEVFusion [20]. For Lyft dataset, SpatialGRU-ODE outperforms synchronized fusion. It performs  $+2.3$  IoU and  $+0.2$  VPQ more than BEVFusion [20].

Fusion mode	IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
nuScenes dataset				
Spatial then temporal	50.2	47.0	75.6	63.0
Temporal then spatial	44.7	42.7	74.6	57.2
SpatialGRU-ODE	47.8	46.1	75.8	60.8
Lyft L5 AV dataset				
Spatial then temporal	54.6	55.7	78.0	71.4
Temporal then spatial	50.4	50.2	75.2	66.7
SpatialGRU-ODE	56.9	55.9	78.1	71.9

Table 7. Comparison with the baseline fusion methods for future instance segmentation (2.0s) on nuScenes validation set. StreamingFlow-ODE is compared with different fusion modes. Spatial fusion is conducted in the same way as BEVFusion [20]. ‘Spatial Then Temporal’ is for synchronous mode, and ‘Temporal Then Spatial’ is for asynchronous mode.

**Effect of ODE solvers.** ODE solvers are vital tools for a N-ODE framework. We test two basic ode solvers, Euler, and midpoint solver as compared in Tab. 8 on two datasets. In both cases, the midpoint solver yields superior VPQ, surpassing the Euler solver by 3.2 points on nuScenes dataset and 0.8 points on Lyft dataset. On nuScenes dataset, the midpoint solver shows similar IoU with the Euler solver, but on Lyft dataset, the midpoint solver performs worse semantic segmentation by  $-1.9$ .

ODE Solver	IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
nuScenes dataset				
Euler	47.8	42.9	73.5	58.6
Midpoint	47.8	46.1	75.8	60.8
Lyft L5 AV dataset				
Euler	56.9	55.9	78.1	71.5
Midpoint	55.0	56.7	78.2	72.4

Table 8. Comparison with prediction using different ODE solvers.

**Effect of ODE update intervals.** We ablate different ODE update intervals which determine the minimum granularity of occupancy prediction. The ODE step is set to 0.05, 0.1, 0.5 and variable for comparison. Variable means the ODE step is the next timestamp minus the current timestamp. In general, the finer a single update is divided during the ODE update process, the higher the final prediction ac-

curacy. It is also intuitive that a higher update frequency may eliminate continuous errors. As shown in Tab. 9, the predictor with interval 0.05s surpasses interval 0.1s by 0.1 IoU and 3.1 VPQ, and interval 0.5s by 0.5 IoU and 3.1 VPQ. Remarkably, the variable time step update also shows nice performance, which is  $+0.5$  IoU and  $-2.2$  VPQ when compared to the predictor with ODE step 0.05s. A variable time step updater strikes a good balance between precision and inference speed, so it is more favorable for long-term prediction to 8s future.

Step	IoU $\uparrow$	PQ $\uparrow$	SQ $\uparrow$	RQ $\uparrow$
nuScenes dataset				
0.05	47.8	46.1	75.8	60.8
0.1	47.7	43.0	73.7	58.4
0.5	47.3	42.1	73.3	57.4
Variable	48.2	43.9	73.9	58.8
Lyft L5 dataset				
0.05	56.9	55.9	78.1	71.9
Variable	54.0	53.5	77.4	69.7

Table 9. Comparison of prediction under different ode steps.

## 5. Conclusion

We present StreamingFlow as the first practice for fusion-based streaming occupancy forecasting. Starting from the motivation of fusing and predicting from non-ideal data streams, the streaming feature brings a more flexible temporal scene understanding and generalizes occupancy forecasting well to any temporal horizon. The advantage of SpatialGRU-ODE is that it decouples the supervision in the training process and outputs in the inference process, by modeling the derivatives of BEV grids and propagating future grid states to required timestamps. We hope that this work will inspire more Neural-ODE applications and insights into the streaming automotive perception.

## 6. Acknowledgement

This work was supported in part by the National Natural Science Foundation of China under Grants (52372414, U22A20104, 52102464). This work was also sponsored by Tsinghua University-DiDi Joint Research Center for Future Mobility and Tsinghua University-Zongmu Technology Joint Research Center.

## References

- [1] Adil Kaan Akan and Fatma Güney. Stretchbev: Stretching future instance prediction spatially and temporally. In *European Conference on Computer Vision*, pages 444–460. Springer, 2022. 1, 2, 5, 6, 7
- [2] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Gi-

- ancarlo Baldan, and Oscar Beijbom. nuscenec: A multi-modal dataset for autonomous driving. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11618–11628, 2020. 1, 2, 5, 7
- [3] Sergio Casas, Wenjie Luo, and Raquel Urtasun. Intentnet: Learning to predict intention from raw sensor data. In *Conference on Robot Learning*, pages 947–956. PMLR, 2018. 2
- [4] Zehui Chen, Zhenyu Li, Shiquan Zhang, Liangji Fang, Qin-hong Jiang, and Feng Zhao. Deformable feature aggregation for dynamic multi-modal 3d object detection. In *European conference on computer vision*, pages 628–644. Springer, 2022. 2
- [5] Edward De Brouwer, Jaak Simm, Adam Arany, and Yves Moreau. Gru-ode-bayes: Continuous modeling of sporadically-observed time series. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2019. 3
- [6] Shaoheng Fang, Zi Wang, Yiqi Zhong, Junhao Ge, and Siheng Chen. Tbp-former: Learning temporal bird’s-eye-view pyramid for joint perception and prediction in vision-centric autonomous driving. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1368–1378, 2023. 2
- [7] Adam W Harley, Zhaoyuan Fang, Jie Li, Rares Ambrus, and Katerina Fragkiadaki. A simple baseline for bev perception without lidar. *arXiv preprint arXiv:2206.07959*, 2022. 2
- [8] Noureldin Hendy, Cooper Sloan, Feng Tian, Pengfei Duan, Nick Charchut, Yuesong Xie, Chuang Wang, and James Philbin. Fishing net: Future inference of semantic heatmaps in grids. *arXiv preprint arXiv:2006.09917*, 2020. 2
- [9] Anthony Hu, Zak Murez, Nikhil Mohan, Sofia Dudas, Jeffrey Hawke, Vijay Badrinarayanan, Roberto Cipolla, and Alex Kendall. FIERY: Future Instance Prediction in Bird’s-Eye View from Surround Monocular Cameras. *Proceedings of the IEEE International Conference on Computer Vision*, pages 15253–15262, 2021. 1, 2, 5, 7
- [10] Shengchao Hu, Li Chen, Penghao Wu, Hongyang Li, Junchi Yan, and Dacheng Tao. St-p3: End-to-end vision-based autonomous driving via spatial-temporal feature learning. In *European Conference on Computer Vision*, pages 533–549. Springer, 2022. 2, 5, 6, 7
- [11] Yihan Hu, Jiazhi Yang, Li Chen, Keyu Li, Chonghao Sima, Xizhou Zhu, Siqi Chai, Senyao Du, Tianwei Lin, Wenhai Wang, Lewei Lu, Xiaosong Jia, Qiang Liu, Jifeng Dai, Yu Qiao, and Hongyang Li. Planning-oriented autonomous driving. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17853–17862, 2023. 1, 2, 7
- [12] Junjie Huang and Guan Huang. Bevdet4d: Exploit temporal cues in multi-camera 3d object detection. *arXiv preprint arXiv:2203.17054*, 2022. 2
- [13] Yanqin Jiang, Li Zhang, Zhenwei Miao, Xiatian Zhu, Jin Gao, Weiming Hu, and Yu-Gang Jiang. Polarformer: Multi-camera 3d object detection with polar transformer. In *Proceedings of the AAAI conference on Artificial Intelligence*, pages 1042–1050, 2023. 2
- [14] R. Kesten, M. Usman, J. Houston, T. Pandya, K. Nadhamuni, A. Ferreira, M. Yuan, B. Low, A. Jain, P. Ondruska, S. Omari, S. Shah, A. Kulkarni, A. Kazakova, C. Tao, L. Platinsky, W. Jiang, and V. Shet. Level 5 perception dataset 2020. <https://level-5.global/level5/data/>, 2019. 1, 2, 5, 7
- [15] Dahun Kim, Sanghyun Woo, Joon-Young Lee, and In So Kweon. Video panoptic segmentation. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9856–9865, 2020. 5
- [16] Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based representation with transformer for 3d object detection. *Advances in Neural Information Processing Systems*, 35:18442–18455, 2022. 2
- [17] Yinhao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1477–1485, 2023. 2
- [18] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *European conference on computer vision*, pages 1–18. Springer, 2022. 2
- [19] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 6
- [20] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela L Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *2023 IEEE international conference on robotics and automation (ICRA)*, pages 2774–2781. IEEE, 2023. 2, 8
- [21] Chenxu Luo, Xiaodong Yang, and Alan Yuille. Self-Supervised Pillar Motion Learning for Autonomous Driving. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 3182–3191, 2021. 2
- [22] 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 conference on Computer Vision and Pattern Recognition*, pages 3569–3577, 2018. 2
- [23] Reza Mahjourian, Jinkyu Kim, Yuning Chai, Mingxing Tan, Ben Sapp, and Dragomir Anguelov. Occupancy flow fields for motion forecasting in autonomous driving. *IEEE Robotics and Automation Letters*, 7(2):5639–5646, 2022. 2
- [24] Ziqi Pang, Deva Ramanan, Mengtian Li, and Yu-Xiong Wang. Streaming motion forecasting for autonomous driving. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7407–7414, 2023. 2
- [25] Jinyoung Park, Chenfeng Xu, Shijia Yang, Kurt Keutzer, Kris M Kitani, Masayoshi Tomizuka, and Wei Zhan. Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection. In *The Eleventh International Conference on Learning Representations*, 2022. 2

- [26] Jonah Philion and Sanja Fidler. Lift, Splat, Shoot: Encoding Images from Arbitrary Camera Rigs by Implicitly Unprojecting to 3D. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 12359 LNCS:194–210, 2020. [1](#), [3](#)
- [27] Zequn Qin, Jingyu Chen, Chao Chen, Xiaozhi Chen, and Xi Li. Unifusion: Unified multi-view fusion transformer for spatial-temporal representation in bird’s-eye-view. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8690–8699, 2023. [2](#)
- [28] Javier Selva, Anders S. Johansen, Sergio Escalera, Kamal Nasrollahi, Thomas B. Moeslund, and Albert Clapés. Video transformers: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11):12922–12943, 2023. [2](#)
- [29] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019. [6](#)
- [30] Yunlong Wang, Hongyu Pan, Jun Zhu, Yu-Huan Wu, Xin Zhan, Kun Jiang, and Diange Yang. Be-sti: Spatial-temporal integrated network for class-agnostic motion prediction with bidirectional enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17093–17102, 2022. [2](#), [5](#), [6](#), [7](#)
- [31] Sizhe Wei, Yuxi Wei, Yue Hu, Yifan Lu, Yiqi Zhong, Siheng Chen, and Ya Zhang. Asynchrony-robust collaborative perception via bird’s eye view flow. *Advances in Neural Information Processing Systems*, 36, 2024. [2](#)
- [32] Pengxiang Wu, Siheng Chen, and Dimitris N. Metaxas. MotionNet: Joint Perception and Motion Prediction for Autonomous Driving Based on Bird’s Eye View Maps. *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pages 11382–11392, 2020. [2](#), [7](#)
- [33] Tengju Ye, Wei Jing, Chunyong Hu, Shikun Huang, Lingping Gao, Fangzhen Li, Jingke Wang, Ke Guo, Wencong Xiao, Weibo Mao, et al. Fusionad: Multi-modality fusion for prediction and planning tasks of autonomous driving. *arXiv preprint arXiv:2308.01006*, 2023. [1](#), [2](#), [6](#), [7](#)
- [34] Yunpeng Zhang, Zheng Zhu, Wenzhao Zheng, Junjie Huang, Guan Huang, Jie Zhou, and Jiwen Lu. Beverse: Unified perception and prediction in birds-eye-view for vision-centric autonomous driving. *arXiv preprint arXiv:2205.09743*, 2022. [2](#), [5](#), [6](#), [7](#)