

# 000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 ASYNCBEV: CROSS-MODAL FLOW ALIGNMENT IN ASYNCHRONOUS 3D OBJECT DETECTION

Anonymous authors

Paper under double-blind review

## ABSTRACT

In autonomous driving, multi-modal perception tasks like 3D object detection typically rely on well-synchronized sensors, both at training and inference. However, despite the use of hardware- or software-based synchronization algorithms, perfect synchrony is rarely guaranteed: Sensors may operate at different frequencies, and real-world factors such as network latency, hardware failures, or processing bottlenecks often introduce time offsets between sensors. Such asynchrony degrades perception performance, especially for dynamic objects. To address this challenge, we propose AsyncBEV, a trainable lightweight and generic module to improve the robustness of 3D Birds’ Eye View (BEV) object detection models against sensor asynchrony. Inspired by scene flow estimation, AsyncBEV first estimates the 2D flow from the BEV features of two different sensor modalities, taking into account the known time offset between these sensor measurements. The predicted feature flow is then used to warp and spatially align the feature maps, which we show can easily be integrated into different current BEV detector architectures (e.g., BEV grid-based and token-based). Extensive experiments demonstrate AsyncBEV improves robustness against both small and large asynchrony between LiDAR or camera sensors in both the token-based CMT (Yan et al., 2023) and grid-based UniBEV (Wang et al., 2024), especially for dynamic objects. We significantly outperform the ego motion compensated CMT and UniBEV baselines, notably by 16.6% and 11.9% NDS on dynamic objects in the worst-case scenario of a 0.5s time offset. Code will be released upon acceptance.

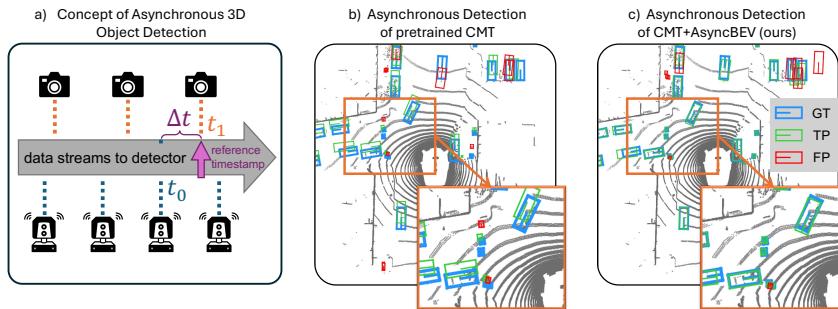


Figure 1: **Multi-modal 3D object detection under sensor asynchrony.** a) shows a general concept of asynchronous 3D object detection. b) shows multi-modal 3D object detection under sensor asynchrony of a pretrained CMT (Yan et al., 2023). Predictions have a large translation error. c) shows our AsyncBEV improves CMT’s robustness against sensor asynchrony and mostly corrects the spatial misalignment of predicted boxes.

## 1 INTRODUCTION

Autonomous vehicles are typically equipped with multiple sensors such as LiDAR, camera, and radar to support perception tasks (Caesar et al., 2020; de Groot et al., 2025). Temporal synchronization among these sensors is crucial for consistent multi-sensor fusion, and autonomous driving

systems (Sun et al., 2020; Wilson et al., 2021; Palfy et al., 2022) therefore employ various hardware- and software-based synchronization strategies (Sivrikaya & Yener, 2004; Gurumadaiah et al., 2025; Yuan et al., 2022; Faizullin et al., 2022). However, perfect synchronization is hard to guarantee in practice. First, sensors may operate at different sampling frequencies or cannot be triggered synchronously with other sensors (e.g. many radar sensors), hence time offsets between them cannot be avoided. Second, competition for computational resources among different parallel components of the processing pipeline can lead to spurious delays, leading to delayed or dropped frames (Fan et al., 2025; Becker et al., 2020). Finally, common sensor failures, such as sensor crash or even adversarial attacks, may further exacerbate misalignment (Xie et al., 2023; Shahriar et al., 2025).

However, downstream multi-modal 3D object detectors are usually trained on well-synchronized and well-curated datasets, and do not consider the possibility of asynchronous inputs. Figure 1 illustrates this: when one modality is delayed, a detector pretrained on synchronized data produces spatial errors in bounding box predictions, which can compromise safety in autonomous driving.

Ego Motion Compensation (EMC) is widely used in autonomous driving to aggregate the temporal information to improve detection performance (Caesar et al., 2020; Li et al., 2022b; Wang et al., 2023b; Park et al., 2023; Cai et al., 2023; Huang & Huang, 2022; Yin et al., 2021). EMC can compensate the spatial offset between sensor measurements at two time frames arising from the known movement of the ego vehicle, which can align cross-modal features of static objects, but it fails on dynamic objects at longer time intervals. Scene flow methods, on the other hand, estimate the motion of dynamic objects (Zhang et al., 2025a), but they require two point clouds as input and assume fixed time offsets, which limits their applicability to asynchronous multi-modal settings.

To address sensor asynchrony, We propose **AsyncBEV**, a lightweight and generic module that can be seamlessly integrated into diverse architectures of multi-modal 3D object detectors. AsyncBEV predicts the additional dynamic object motion not yet captured by EMC to align the asynchronous sensor data to a reference timestamp. To this end, we define a new task,  $\Delta$ -BEVFlow estimation, which aims to predict the flow in the BEV space between multi-modal BEV features given the known time offset between the features. Its offset-aware design facilitates the motion prediction for time offsets that can vary during operation, without significant performance loss in the synchronization case. Since AsyncBEV computes flow from BEV features, it does not make assumptions on the sensor modalities, nor which sensor is considered the reference. To summarize, our main contributions are as follows:

- We introduce AsyncBEV, a trainable and lightweight module that can be integrated into diverse multi-modal 3D object detector frameworks, and improves a detector’s robustness against small, moderate, and severe time differences between sensors, especially for dynamic objects. We integrate AsyncBEV in both a token-based and a dense grid-based detector architecture.
- AsyncBEV utilizes  $\Delta$ -BEVFlow, a novel task that predicts flow in the BEV feature space. Unlike prior work on scene flow,  $\Delta$ -BEVFlow operates on multi-modal feature maps rather than point clouds, and is conditioned on the known time interval between the sensors.
- We study two different flow formulations: Motion-based, which regresses the flow between BEV features directly, whereas velocity-based flow predicts object velocities, and afterward multiplies these with the known time interval to obtain motion. The velocity-based formulation regularizes the task, resulting in better performance than motion-based flow regression.

## 2 RELATED WORK

**Multi-modal 3D Object Detection.** To leverage the complementary information of different modalities (e.g., LiDAR and cameras), multi-modal 3D object detection has drawn wide attention in both academia and industry (Song et al., 2024b; Wang et al., 2023c; Huang et al., 2022). Since cameras and LiDAR capture data in different coordinate systems, object detectors map the multi-modal features into a shared reference system, usually LiDAR coordinates. They thus have to explicitly or implicitly project multi-view image features into the LiDAR (BEV) space. Grid-based methods (Liu et al., 2023; Cai et al., 2023; Wang et al., 2024; Ge et al., 2023; Li et al., 2022a) explicitly project image features into a BEV grid for unified representations, while token-based methods (Yan et al., 2023; Chen et al., 2022; Li et al., 2022a) rely on object queries to implicitly perform coordinate transformation under the guidance of their positional encodings, which represent the mapping

108 between each token and its 3D coordinate information to perform the coordinate transformation  
 109 implicitly. These methods require strict cross-sensor synchronization so that the images and point  
 110 clouds align in the 3D coordinate system, though 3D object detection under sensor asynchrony re-  
 111 mains underexplored.

112 **Scene Flow Estimation.** Scene flow is the 3D motion field of all points in a scene, describing  
 113 the per-point motion in 3D space between consecutive point clouds. Recent scene flow estimators  
 114 deploy feed-forward neural networks (FNN) to achieve high accuracy and real-time inference (Jund  
 115 et al., 2021; Zhang et al., 2024; 2025b; Kim et al., 2024; Lin et al., 2025; Zhang et al., 2025a).  
 116 Motion cues of dynamic objects, which scene flow provides, could help re-align information from  
 117 the asynchronous time frames to a reference time frame. However, typical scene flow methods  
 118 require two adjacent point clouds and assume a fixed time offset, and are therefore not suited for  
 119 aligning measurements from different sensor modalities at dynamic time offsets.

120 **Asynchronous Fusion in Collaborative Perception.** Asynchronous fusion is a crucial problem  
 121 in collaborative perception, as this task needs to fuse information from multiple agents, but distinct  
 122 agents and roadside units have independent clocks and are hard to synchronize with each other.  
 123 CoBEVFlow (Wei et al., 2023) first generates object proposals from distinct agent features and then  
 124 applies a computation-intensive attention-based module to associate the proposals from different  
 125 agents. Its flow is not directly learned but calculated based on the associated object proposals.  
 126 UniV2X (Yu et al., 2025) employs an agent query-based flow predictor to generate the potential  
 127 motion of the query, which takes the associated agent from multiple frames as input. However, these  
 128 flow calculation methods have two main shortcomings. First, their performance of flow generation  
 129 heavily relies on the quality of the generated object proposals. However, in multi-modal (on-board)  
 130 3D object detection, the detection characteristics for different sensors can differ drastically. For  
 131 instance, camera detectors usually perform worse than LiDAR detectors. Second, even assuming  
 132 correct object proposals, in these two works, all agents send homogeneous features from the same  
 133 modality (LiDAR features in CoBEVFlow (Wei et al., 2023) and image features in UniV2X (Yu  
 134 et al., 2025)). The distinct features in multi-modal (on-board) 3D object detection pose further  
 135 challenges in associating object proposals generated from different sensors or predicting query flows  
 136 with a single MLP module. The effectiveness of their temporal alignment method has not been  
 137 validated in a multi-modal setup.

138 **Asynchronous Fusion in On-board Perception.** Asynchronous fusion in on-board perception  
 139 shares a common concept with temporal feature aggregation, i.e., to align features from another  
 140 timestamp to the reference timestamp. Early works (Caesar et al., 2020; Huang & Huang, 2022;  
 141 Park et al., 2023) pioneered the use of EMC to align the static objects, while following works (Li  
 142 et al., 2022b; Wang et al., 2023b; Cai et al., 2023) further apply attention modules to encode the  
 143 object motion implicitly. StreamingFlow (Shi et al., 2024) and TimeAlign (Song et al., 2024a) pio-  
 144 neered autonomous driving perception tasks with asynchronous input. However, both works deploy  
 145 a computation-heavy RNN-based module to predict features at the reference timestamp, taking the  
 146 long sequence history data as input. StreamingFlow (Shi et al., 2024) takes streaming multi-modal  
 147 features as a set of observations and treats these features uniformly, which requires a unified feature  
 148 representation, such as grid-like BEV features, for both modalities. Hayes et al. (2024) leverages  
 149 the velocity information provided by radars to update the point positions from the asynchronous  
 150 timestamp to the reference timestamp, which shares the same spirit as scene flow estimation. Fan  
 151 et al. (2025) proposes directly integrating time features into per-point features to make the detector  
 152 time-aware. None of the above methods explicitly uses the motion information from multi-modal  
 153 features. Hence, to address the sensor asynchrony problem, we propose AsyncBEV, a light-weight,  
 154 generic, interpretable and effective solution, which can work with arbitrary time offsets, asynchrony  
 155 of different sensors, and distinct 3D object detectors.

### 3 METHODOLOGY

156 Most multi-modal 3D object detectors follow a similar structure, which we describe here first under  
 157 the common simplifying assumption that all sensors are synchronous. This paper mainly considers  
 158 camera and LiDAR fusion, though the same principles can be applied to include radar, too.

159 As input, the detector receives images  $\mathbb{I} \in \mathbb{R}^{V \times H_{img} \times W_{img} \times 3}$  from  $V$  camera views, and a point  
 160 cloud  $\mathbb{P} \in \mathbb{R}^{(N^{t_0}, 3)}$ . First,  $\mathbb{I}$  and  $\mathbb{P}$  are fed into two separate backbones to generate modal-specific

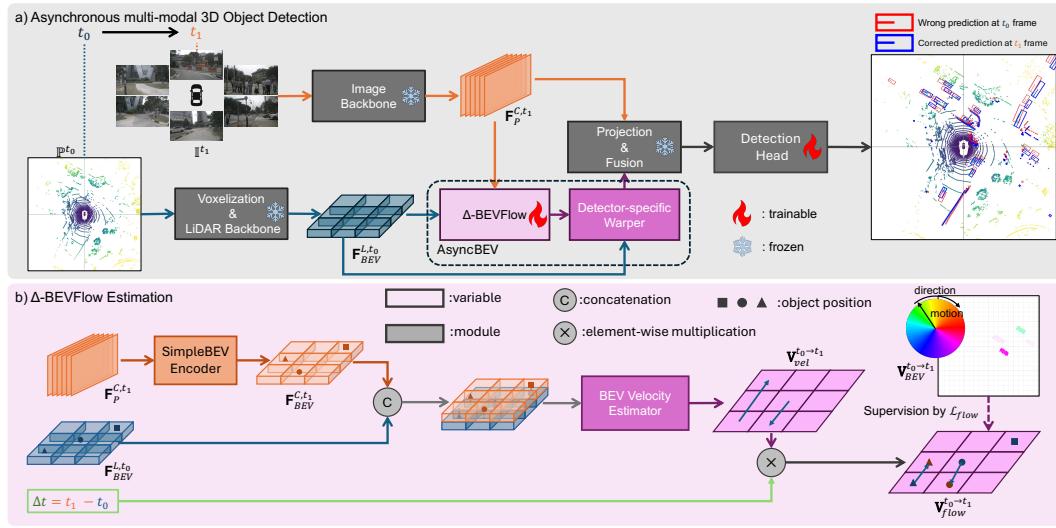


Figure 2: **Overview of AsyncBEV.** a) demonstrates an example pipeline of asynchronous multi-modal 3D object detection. In this setup, images are sampled at the reference timestamp  $t_1$ , while LiDAR point clouds were obtained at an earlier timestamp  $t_0$ . b) shows the  $\Delta$ -BEVFlow estimation, the core module of AsyncBEV. It takes asynchronous BEV features as input and predicts the velocity of each BEV cell. The BEV flow is calculated by multiplying the velocity by  $\Delta t$  and supervised by the ground truth dense flow representation. Finally, the predicted flow is used to spatially align the features from  $t_0$  to  $t_1$  with the detector-specific warper.

features in their native coordinate, resulting in multi-view image features in perspective view  $\mathbf{F}_P^C \in \mathbb{R}^{V \times H_C \times W_C \times D_C}$  and LiDAR features in bird's eye view  $\mathbf{F}_{BEV}^L \in \mathbb{R}^{H_L \times W_L \times D_L}$ , where  $H_C$ ,  $W_C$ ,  $D_C$  and  $H_L$ ,  $W_L$ ,  $D_L$  represent the sizes and dimensions of perspective view (PV) image features and bird's eye view (BEV) LiDAR features, respectively. After getting the multi-modal features in their native coordinate systems  $\mathbf{F}_P^C$  and  $\mathbf{F}_{BEV}^L$ , different detectors deploy distinct strategies to map sensor features from sensor-specific coordinates to a shared 3D coordinate space. As LiDAR is already defined in the BEV space, these methods mainly differ in how to process camera features  $\mathbf{F}_P^C$  to corresponding 3D coordinates  $\mathbb{C}_{3D}^C$ . We shall distinguish two main categories of detectors, and show how to integrate AsyncBEV with each: grid-based and token-based methods.

*Grid-based methods* (Liu et al., 2023; Liang et al., 2022; Wang et al., 2024) learn to project  $\mathbf{F}_P^C$  into a BEV grid with  $\mathbb{C}_{3D}^C$ , resulting in a dense grid-like camera BEV feature map  $\mathbf{F}_{BEV}^C$ , similar to the LiDAR feature map  $\mathbf{F}_{BEV}^L$ .

*Token-based methods* (Yan et al., 2023; Wang et al., 2023a) use a sparse set of object queries, which are defined in the 3D space, to retrieve information from features  $\mathbf{F}_P^C$  through a set of 3D positional encoding generated from  $\mathbb{C}_{3D}^C$ . This achieves the perspective to BEV projection implicitly.

**Asynchrony.** In practice, sensor input can be asynchronous, thus the timestamps associated with data from each sensor will differ. We assume that one sensor, e.g. the camera, is considered the reference that triggers inference of the multi-modal detector at timestamp  $t_1$ <sup>1</sup>. From the other sensor modality, e.g. LiDAR, the most up-to-date data from timestamp  $t_0$  is taken, which will be delayed w.r.t. the reference ( $t_0 < t_1$ ). In this example, we would refer to the cameras as *the synchronous sensor* and to the LiDAR *the asynchronous sensor*, and to  $\Delta t = t_1 - t_0$  as the asynchrony time offset. In both token-based and grid-based frameworks, the asynchrony between sensors results in a misalignment between  $\mathbb{C}_{3D}^{C,t_1}$  and  $\mathbb{C}_{3D}^{L,t_0}$ , which can lead to severe detection degradation, especially for important dynamic objects. Note that we from here on indicate the sensor timestamp  $t_0$  or  $t_1$  in the superscript notation of symbols.

<sup>1</sup>In this work, we focus only on cross-modality asynchrony, while asynchrony within multiple cameras is considered out of scope.

### 3.1 ASYNCBEV

We propose a **AsyncBEV** to predict the flow (movement) of dynamic objects in the BEV space and then warp the misaligned BEV spatial locations before the detector's fusion. Figure 2 a) demonstrates how AsyncBEV can be integrated in a generic 3D object detection framework.

**Flow estimation.** A basic step to align sensor coordinate frames and feature maps is *Ego-Motion Compensation* (EMC), which represents the known motion of the ego-vehicle between time  $t_0$  and  $t_1$ , and thus aligns static objects only. In the flow formulation, EMC can be represented as a 2D flow field  $\mathbf{M}^{t_0 \rightarrow t_1} \in \mathbb{R}^{(N^t, 2)}$  which expresses the relative spatial displacement for each BEV grid cell due to ego-motion. Alternatively, scene flow methods aim to predict non-static motion too as a 3D point-level translation  $\mathbb{V}^{t \rightarrow t + \Delta t} \in \mathbb{R}^{(N^t, 3)}$  from two subsequent point clouds  $\mathbb{P}^t \in \mathbb{R}^{(N^t, 3)}$  and  $\mathbb{P}^{t + \Delta t} \in \mathbb{R}^{(N^{t + \Delta t}, 3)}$ . Therefore, current scene flow estimators can be generally formulated as a function  $\mathbb{V}^{t \rightarrow t + \Delta t} = \theta_{\text{SceneFlow}}(\mathbb{P}^t, \mathbb{P}^{t + \Delta t})$ . Since existing scene flow methods require two point clouds as input and assume a fixed  $\Delta t$ , they cannot be directly used to address asynchronous multi-modal 3D object detection where  $\Delta t$  can vary continuously.

To handle dynamic objects, AsyncBEV adapts the scene flow concept to a new task, **Δ-BEVFlow**, which should capture the motion missing from EMC. Δ-BEVFlow differs from prior scene flow works in several ways: 1) The flow is also conditioned on  $\Delta t$ , 2) it works with multi-modal BEV features instead of raw point clouds, 3) we only require 2D flow in BEV space. The task is thus to learn a function:

$$\mathbf{V}_{\text{BEV}}^{t_0 \rightarrow t_1} = \theta_{\Delta\text{-BEVFlow}}(\mathbf{F}_{\text{BEV}}^{m_0, t_0}, \mathbf{F}_{\text{BEV}}^{m_1, t_1}, \Delta t), \quad (1)$$

where  $\mathbf{V}_{\text{BEV}}^{t_0 \rightarrow t_1} \in \mathbb{R}^{(H_{\text{BEV}} \times W_{\text{BEV}}, 2)}$  represents the 2D flow field defined in the BEV space with the size  $H_{\text{BEV}} \times W_{\text{BEV}}$ , and  $\mathbf{F}_{\text{BEV}}^{m_0, t_0}$  and  $\mathbf{F}_{\text{BEV}}^{m_1, t_1}$  denote the BEV features of modality  $m_0$  and  $m_1$  sampled from asynchronous timestamps  $t_0$  and  $t_1$ , respectively. We consider two possible formulations for Equation 1: *motion estimation* and *velocity estimation*.

The *motion estimation (BEV-ME)* formulation directly regresses the dynamic motion from  $\Delta t$ :

$$\mathbf{E}_{me} = \phi(cat(\mathbf{F}_{\text{BEV}}^{C,t_1}, \mathbf{F}_{\text{BEV}}^{L,t_0}, \Delta t)), \quad \mathbf{V}_{\text{BEV}}^{t_0 \rightarrow t_1} = \psi(\mathbf{E}_{me}). \quad (2)$$

Here  $\phi : \mathbb{R}^{D_L+D_C+1} \rightarrow \mathbb{R}^{D'}$  is an encoder which compresses the (still misaligned) concatenated BEV features to a lower-dimensional BEV feature map, and  $\Delta t$ , reducing memory and computational requirements. Following common practice in scene flow works,  $\psi$  is a U-Net (Ronneberger et al., 2015) which transform the encoded BEV features into a flow field. We note (Fan et al., 2025) recently also proposed to add  $\Delta t$  as additional input dimensions such that a detection network may learn to compensate for dynamic motion, although without an explicit flow formulation.

We here propose an alternative formulation, *velocity estimation (BEV-VE)*. Instead of regressing from  $\Delta t$ , this first predict the velocity of each cell independent of  $\Delta t$ , and afterwards multiply predicted velocities with  $\Delta t$  to obtain the requested motion, exploiting basic knowledge of the relation between motion, velocity and time. This design regularizes the flow at low  $\Delta t$  for near synchronous sensors, since it forces the dynamic motion to become zero as  $\Delta t$  approaches zero, and simplifies the learning task:

$$\mathbf{E}_{ve} = \phi(cat(\mathbf{F}_{\text{REV}}^{C,t_1}, \mathbf{F}_{\text{REV}}^{L,t_0})), \quad \mathbf{V}_{\text{vel}}^{t_0 \rightarrow t_1} = \psi(\mathbf{E}_{ve}), \quad \mathbf{V}_{\text{REV}}^{t_0 \rightarrow t_1} = \mathbf{V}_{\text{vel}}^{t_0 \rightarrow t_1} \times \Delta t. \quad (3)$$

As our ablation study will validate, velocity estimation works better than motion estimation, and it will be our default  $\Delta$ -BEVFlow formulation in the other experiments.

**Lightweight image BEV encoder.**  $\Delta$ -BEVFlow requires BEV features from all sensor modalities. Since token-based 3D object detectors do not explicitly build image BEV features, we first generate the BEV features directly from the image features, ensuring AsyncBEV can work with different detection frameworks. We adopt SimpleBEV (Harley et al., 2023) to generate  $\mathbf{F}_{\text{BEV}}^{C,t_1}$ . SimpleBEV directly projects image features  $\mathbf{F}_P^{C,t_1}$  into BEV space with the cameras' extrinsic and intrinsic parameters, avoiding demanding operations such as depth estimation (Liu et al., 2023; Liang et al., 2022; Huang et al., 2021) and deformable attention (Li et al., 2022b; Wang et al., 2024).

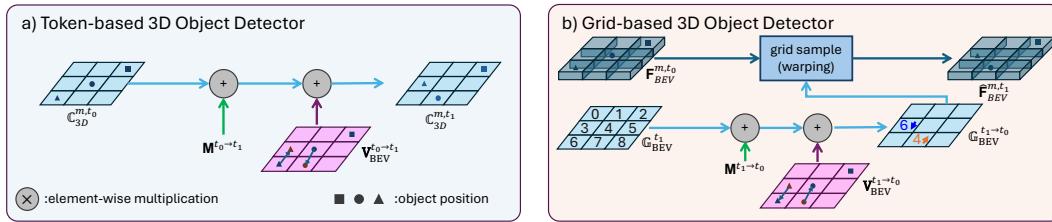


Figure 3: Detector-specific Warper. a) Token-based methods rely on 3D coordinates of tokens to provide spatial guidance. Generated  $\Delta$ -BEVFlow is used to correct the corresponding 3D coordinates of tokens from the asynchronous sensor. b) Grid-based methods project the multi-modal features into predefined BEV grids. Generated  $\Delta$ -BEVFlow is used to generate a look-up table from the reference timestamp to the asynchronous timestamp to warp the asynchronous BEV features to the reference timestamp.

**Detector-specific warper.** Finally, the ego-motion flow  $\mathbf{M}^{t_0 \rightarrow t_1}$  and  $\Delta$ -BEVFlow  $\mathbf{V}_{\text{BEV}}^{t_0 \rightarrow t_1}$  need to be applied to align the features from the different sensor modalities. The exact implementation of this step depends on the aforementioned detector architecture, token-based or grid-based.

*Token-based methods* encode spatial locations of tokens (sensor features) in their 3D positional encoding. Therefore, we directly adjust the 3D coordinates  $\mathbb{C}_{3D}^{m,t_0}$  of the asynchronous sensor  $m \in \{L, C\}$  with the predicted flow  $\mathbf{V}_{\text{BEV}}^{t_0 \rightarrow t_1}$  to obtain the tokens representing timestamp  $t_1$ , as shown in Figure 3 a). Specifically, we first apply the predicted dense flow  $\mathbf{V}_{\text{BEV}}^{t_0 \rightarrow t_1}$  to the position of each sparse token’s positional embedding. We do the same operation for the EMC flow, also updating the token positions to compensate for the vehicle ego motion. The final 3D token coordinates  $\mathbb{C}_{3D}^{m,t_1}$  are used to generate new positional encodings.

*Grid-based methods* implicitly encode spatial positions in the grid indices. Hence, here we can adopt a *grid sampling* operation to look-up features in the BEV feature map at  $t_0$  for each cell in the BEV grid at  $t_1$ . As shown in figure 3 b), we first define a standard BEV grid  $\mathbb{G}_{\text{BEV}}^{t_1}$ , representing the 3D space at  $t_1$ , which should have the same size with BEV features  $\mathbf{F}_{\text{BEV}}^{m,t_0}$ . Then we simply add the ego-motion flow  $\mathbf{M}^{t_1 \rightarrow t_0}$  onto  $\mathbb{G}_{\text{BEV}}^{t_1}$  to compensate the ego-motion, resulting in  $\mathbb{G}_{\text{BEV}}^{t_1, \text{EMC}}$ . We further calculate the final look-up table  $\mathbb{G}_{\text{BEV}}^{t_1 \rightarrow t_0} = \mathbb{G}_{\text{BEV}}^{t_1, \text{EMC}} + \mathbf{V}_{\text{BEV}}^{t_1 \rightarrow t_0}$ . Note that in this setup, we in fact learn flow  $\mathbf{V}_{\text{flow}}^{t_1 \rightarrow t_0}$  instead of  $\mathbf{V}_{\text{flow}}^{t_0 \rightarrow t_1}$ , as was the case for the token-based models. With  $\mathbb{G}_{\text{BEV}}^{t_1 \rightarrow t_0}$ , we obtain pseudo BEV features of modality  $m$  at  $t_1$ ,  $\hat{\mathbf{F}}_{\text{BEV}}^{m,t_1}$ , with the ‘grid\_sample’ function, i.e.,  $\hat{\mathbf{F}}_{\text{BEV}}^{m,t_1} = f_{\text{grid\_sample}}(\mathbf{F}_{\text{BEV}}^{m,t_0}, \mathbb{G}_{\text{BEV}}^{t_1 \rightarrow t_0})$ .

**Loss functions.** AsyncBEV can be trained by back-propagating through an existing frozen object detector framework, using the regular object detection losses. Following the common practice of 3D object detection losses, we also adopt focal loss for classification  $L_{\text{cls}}$  and L1 Loss for 3D bounding box regression  $L_{\text{reg}}$ .

A benefit of the explicit flow formulation is that we can (optionally) also directly supervise it to facilitate training. For this, we adopt the flow loss  $\mathcal{L}_{\text{flow}}$  from DeFlow (Zhang et al., 2024), using generated ground truth BEVFlow from annotated 3D object bounding boxes, as elaborated in Appendix A.5. This flow loss sums the mean L2-distance between the predicted and ground truth flows for three motion speed categories, i.e.  $v \leq 0.4m/s$ ,  $0.4m/s < v \leq 1m/s$  and  $v > 1m/s$ . The overall training objective is then a weighted sum of the three components,  $\mathcal{L}_{\text{total}} = \omega_1 \cdot \mathcal{L}_{\text{flow}} + \omega_2 \cdot \mathcal{L}_{\text{cls}} + \omega_3 \cdot \mathcal{L}_{\text{reg}}$ , where  $\omega_1$ ,  $\omega_2$ , and  $\omega_3$  are hyperparameters that balance the contributions of the loss terms.

## 4 EXPERIMENTS

**Dataset.** We conduct experiments on the nuScenes (Caesar et al., 2020) dataset. nuScenes is a large-scale autonomous driving dataset collected in Boston and Singapore. It has 1000 scenes

324	325	Method	All Objects			Static Objects			Dynamic Objects			FPS ↑
			0s	0.25s	0.5s	0s	0.25s	0.5s	0s	0.25s	0.5s	
326	CMT	NDS ↑	<b>72.9</b>	49.6	43.2	<b>69.0</b>	48.7	44.2	<b>47.5</b>	32.2	26.1	<b>6.7</b>
327		mAP ↑	<b>70.3</b>	35.9	24.8	<b>61.0</b>	32.2	22.8	<b>36.8</b>	16.4	9.0	
328	CMT + EMC	NDS ↑	<b>72.9</b>	66.8	63.3	<b>69.0</b>	<b>68.3</b>	<b>67.5</b>	<b>47.5</b>	34.5	26.8	<b>6.7</b>
329		mAP ↑	<b>70.3</b>	60.6	54.1	<b>61.0</b>	59.8	<b>58.6</b>	<b>36.8</b>	21.1	11.9	
330	CMT + DA	NDS ↑	<b>71.5</b> <small>↓ 1.4</small>	<b>70.2</b> <small>↑ 3.4</small>	<b>67.6</b> <small>↑ 4.3</small>	<b>68.2</b> <small>↓ 0.8</small>	<b>67.2</b> <small>↓ 1.1</small>	<b>65.4</b> <small>↓ 2.1</small>	<b>45.8</b> <small>↓ 1.7</small>	<b>44.9</b> <small>↑ 0.4</small>	<b>41.5</b> <small>↑ 14.7</small>	<b>6.7</b>
331	Fan et al. (2025)	mAP ↑	<b>68.4</b> <small>↓ 1.9</small>	<b>66.7</b> <small>↑ 6.1</small>	<b>63.0</b> <small>↑ 8.9</small>	<b>60.2</b> <small>↓ 0.8</small>	<b>58.7</b> <small>↓ 1.1</small>	<b>56.0</b> <small>↓ 2.6</small>	<b>34.5</b> <small>↓ 2.3</small>	<b>33.8</b> <small>↑ 12.7</small>	<b>30.0</b> <small>↑ 18.1</small>	
332	CMT+AsyncBEV	NDS ↑	<b>70.6</b> <small>↓ 2.3</small>	<b>70.0</b> <small>↑ 3.2</small>	<b>66.2</b> <small>↑ 2.9</small>	<b>67.7</b> <small>↓ 1.3</small>	<b>67.1</b> <small>↓ 1.2</small>	<b>64.7</b> <small>↓ 2.8</small>	<b>43.5</b> <small>↓ 4.0</small>	<b>43.9</b> <small>↑ 9.4</small>	<b>37.9</b> <small>↑ 11.1</small>	<b>6.7</b>
333		mAP ↑	<b>67.9</b> <small>↓ 2.4</small>	<b>66.8</b> <small>↑ 6.2</small>	<b>61.3</b> <small>↑ 7.2</small>	<b>59.1</b> <small>↓ 1.9</small>	<b>58.4</b> <small>↓ 1.4</small>	<b>54.8</b> <small>↓ 3.8</small>	<b>34.7</b> <small>↓ 2.1</small>	<b>33.1</b> <small>↑ 12.0</small>	<b>27.1</b> <small>↑ 15.2</small>	
334	CMT+AsyncBEV (FD)	NDS ↑	<b>72.5</b> <small>↓ 0.4</small>	<b>71.5</b> <small>↑ 4.7</small>	<b>70.0</b> <small>↑ 6.7</small>	<b>68.7</b> <small>↓ 0.3</small>	<b>68.3</b>	<b>67.6</b> <small>↑ 0.1</small>	<b>47.1</b> <small>↓ 0.4</small>	<b>46.1</b> <small>↑ 11.6</small>	<b>43.4</b> <small>↑ 16.6</small>	6.3
335		mAP ↑	<b>69.8</b> <small>↓ 0.5</small>	<b>68.3</b> <small>↑ 7.7</small>	<b>66.4</b> <small>↑ 12.3</small>	<b>60.7</b> <small>↓ 0.3</small>	<b>60.0</b> <small>↓ 0.2</small>	<b>59.1</b> <small>↓ 0.5</small>	<b>36.8</b>	<b>35.4</b> <small>↑ 14.3</small>	<b>32.3</b> <small>↑ 20.4</small>	
336	UniBEV	NDS ↑	<b>73.0</b> <small>↑ 0.1</small>	<b>71.4</b> <small>↑ 4.6</small>	<b>68.7</b> <small>↑ 5.4</small>	<b>69.2</b> <small>↑ 0.2</small>	<b>68.2</b> <small>↓ 0.1</small>	<b>67.0</b> <small>↓ 0.5</small>	<b>47.7</b> <small>↑ 0.2</small>	<b>45.0</b> <small>↑ 0.5</small>	<b>39.6</b> <small>↑ 12.8</small>	6.3
337		mAP ↑	<b>70.5</b> <small>↑ 0.2</small>	<b>68.1</b> <small>↑ 7.5</small>	<b>64.4</b> <small>↑ 10.3</small>	<b>61.4</b> <small>↑ 0.4</small>	<b>59.9</b> <small>↓ 0.1</small>	<b>58.3</b> <small>↓ 0.3</small>	<b>36.9</b> <small>↑ 0.1</small>	<b>34.1</b> <small>↑ 13.0</small>	<b>27.8</b> <small>↑ 15.9</small>	
338	UniBEV+EMC	NDS ↑	<b>66.7</b>	45.7	39.7	<b>63.5</b>	45.3	41.2	<b>42.4</b>	29.6	25.3	<b>2.8</b>
339		mAP ↑	<b>62.0</b>	32.1	22.2	<b>52.2</b>	27.6	19.0	<b>32.9</b>	15.0	8.9	
340	UniBEV+EMC	NDS ↑	<b>66.7</b>	61.7	58.9	<b>63.5</b>	<b>62.9</b>	<b>62.2</b>	<b>42.4</b>	31.2	25.9	<b>2.8</b>
341		mAP ↑	<b>62.0</b>	53.8	48.2	<b>52.2</b>	<b>51.1</b>	<b>50.0</b>	<b>32.9</b>	21.1	11.4	
342	UniBEV + DA	NDS ↑	<b>60.2</b> <small>↓ 6.5</small>	<b>58.3</b> <small>↓ 3.4</small>	<b>54.5</b> <small>↓ 4.4</small>	<b>58.5</b> <small>↓ 5.0</small>	<b>56.8</b> <small>↓ 6.1</small>	<b>53.8</b> <small>↓ 8.4</small>	<b>36.6</b> <small>↓ 5.8</small>	<b>36.0</b> <small>↑ 4.8</small>	<b>33.6</b> <small>↑ 7.7</small>	<b>2.8</b>
343		mAP ↑	<b>53.6</b> <small>↓ 8.4</small>	<b>50.9</b> <small>↓ 2.9</small>	<b>45.9</b> <small>↓ 2.3</small>	<b>45.2</b> <small>↓ 7.0</small>	<b>42.9</b> <small>↓ 8.2</small>	<b>38.3</b> <small>↓ 11.7</small>	<b>24.9</b> <small>↓ 8.0</small>	<b>24.0</b> <small>↑ 2.9</small>	<b>20.4</b> <small>↑ 9.0</small>	
344	StreamingFlow (Shi et al., 2024)	NDS ↑	<b>64.2</b> <small>↓ 2.5</small>	<b>58.8</b> <small>↓ 2.9</small>	<b>54.1</b> <small>↓ 4.8</small>	<b>61.8</b> <small>↓ 1.7</small>	<b>56.8</b> <small>↓ 6.1</small>	<b>53.1</b> <small>↓ 9.1</small>	<b>41.3</b> <small>↓ 1.1</small>	<b>38.5</b> <small>↑ 7.3</small>	<b>34.4</b> <small>↑ 8.5</small>	1.0
345		mAP ↑	<b>60.3</b> <small>↓ 1.7</small>	<b>53.0</b> <small>↓ 0.8</small>	<b>47.4</b> <small>↓ 0.8</small>	<b>49.9</b> <small>↓ 2.3</small>	<b>43.2</b> <small>↓ 7.9</small>	<b>37.3</b> <small>↓ 12.7</small>	<b>31.5</b> <small>↓ 1.4</small>	<b>27.4</b> <small>↓ 6.3</small>	<b>21.6</b> <small>↑ 10.2</small>	
346	UniBEV+AsyncBEV	NDS ↑	<b>65.7</b> <small>↓ 1.0</small>	<b>64.8</b> <small>↑ 3.1</small>	<b>63.3</b> <small>↑ 4.4</small>	<b>62.7</b> <small>↓ 0.8</small>	<b>62.0</b> <small>↓ 0.9</small>	<b>61.5</b> <small>↓ 0.7</small>	<b>41.0</b> <small>↓ 1.4</small>	<b>40.2</b> <small>↑ 9.0</small>	<b>37.8</b> <small>↑ 11.9</small>	2.7
347		mAP ↑	<b>60.8</b> <small>↓ 1.2</small>	<b>59.4</b> <small>↑ 5.6</small>	<b>57.1</b> <small>↑ 8.9</small>	<b>50.8</b> <small>↓ 1.4</small>	<b>49.7</b> <small>↓ 1.4</small>	<b>49.1</b> <small>↓ 0.9</small>	<b>31.0</b> <small>↓ 1.9</small>	<b>29.8</b> <small>↑ 8.7</small>	<b>26.6</b> <small>↑ 15.2</small>	
348	UniBEV+AsyncBEV (FD)	NDS ↑	<b>66.7</b>	<b>65.2</b> <small>↑ 3.5</small>	<b>62.1</b> <small>↑ 3.2</small>	<b>63.5</b>	<b>62.9</b>	<b>62.0</b> <small>↓ 0.2</small>	<b>42.5</b> <small>↑ 0.1</small>	<b>39.6</b> <small>↑ 8.4</small>	<b>33.8</b> <small>↑ 7.9</small>	2.7
349		mAP ↑	<b>62.1</b> <small>↑ 0.1</small>	<b>59.6</b> <small>↑ 5.8</small>	<b>54.7</b> <small>↑ 6.5</small>	<b>52.3</b> <small>↑ 0.1</small>	<b>51.2</b> <small>↑ 0.1</small>	<b>49.9</b> <small>↓ 0.1</small>	<b>32.9</b>	<b>29.1</b> <small>↑ 8.0</small>	<b>20.7</b> <small>↑ 9.3</small>	

**Table 1: Evaluation on the nuScenes val set of asynchronous multi modal 3D object detection (Best, Second Best).** LiDAR is asynchronous. AsyncBEV is only trained once with each model and evaluated on different time offsets. FD: freeze the whole detector and only train AsyncBEV. Performance deltas are measured relative to the EMC baseline. Inference speed in Frames Per Second (FPS) is measured with a single A100 GPU.

in total and is divided into train/val/test sets with 750/150/150 scenes. nuScenes contains multi-modal data from 6 cameras, 1 LiDAR, and 5 radars. In this work, we mainly focus on LiDAR and camera fusion for 3D object detection. Following common practice in nuScenes, we report both the nuScenes Detection Score (NDS) and mean Average Precision (mAP) as metrics. To understand the motion compensation performance of our approach, we further report NDS and mAP for dynamic and static objects. We adopt 0.2 m/s as a threshold to distinguish dynamic and static objects in both predictions and ground truth, then calculate the metrics respectively.

**Baselines.** We integrate AsyncBEV in both the CMT (Yan et al., 2023) and UniBEV (Wang et al., 2024) architectures, which are representative for token-based and grid-based multi-modal 3D object detectors. Compared to its original end-to-end training schema, we retrained UniBEV with a two-stage training protocol to accelerate the training process. In the default setting, we use frozen pre-trained backbones and finetune only the AsyncBEV module and the detection heads. We also train a variant *AsyncBEV (FD)* of each architecture with *frozen* detection heads, so only the AsyncBEV module is trained. We adopt a learning rate of  $2 \times 10^{-5}$  for finetuning in both architectures. Other hyperparameters are kept the same as their original training setups. For both CMT and UniBEV architectures, we also test an EMC-only variation without AsyncBEV, which has no additional learnable components. Data augmentation (DA) is another simple baseline for both networks, which directly finetunes the vanilla detectors with the same asynchronous training scheme as AsyncBEV.

Furthermore, we also deploy the two asynchronous fusion methods Fan et al. (2025) and StreamingFlow (Shi et al., 2024) as baselines. Fan et al. (2025) proposed a temporal information propagation method for point clouds. Modern LiDAR encoders usually accumulate multiple historic sweeps of point clouds with the latest one as input point cloud, and append their time offsets relative to the latest one as an additional temporal feature alongside its geometric dimensions. Following Fan et al. (2025), we simply replace the original temporal feature with the asynchronous time offsets  $\Delta t$  measured relative to the reference timestamp. We integrate Fan et al. (2025) with CMT for asynchronous LiDAR. Please note, Fan et al. (2025) can only be applied to the point cloud input, as images usually have no time dimension. A reasonable extension is to concatenate the  $\Delta t$  to the intermediate features, which is similar to our motion-based flow estimation variant. StreamingFlow utilizes a GRU-ODE to model the temporal dynamics among grid-like features from different timestamps. To adapt StreamingFlow to 3D object detection task, we therefore integrate the GRU-ODE module into our grid-based baseline, UniBEV. For all the trainable baselines, we also adopt the same training

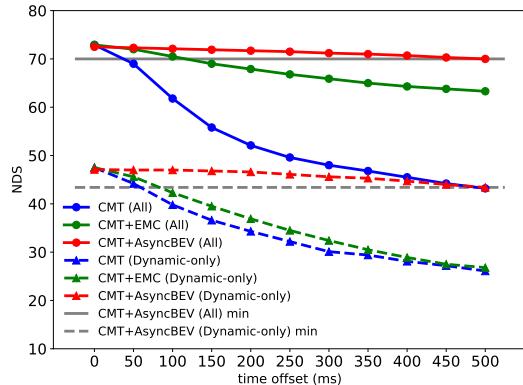
378 scheme with AsyncBEV, i.e., keep the pretrained backbone frozen and finetune the detection head  
 379 and other components.  
 380

381 **Asynchronous training setup.** While our method is generic to any dataset, we follow nuScenes  
 382 conventions here for simplicity. As is common practice, we train and evaluate our method on an-  
 383notated nuScenes keyframes at 2 Hz. Considering the keyframe as the reference timestamp, one  
 384sensor is synchronous while the other sensor is asynchronous. Since the nuScenes dataset is artifi-  
 385cially curated to discard scenes with bad synchronization, we cannot use the natural data, but have to  
 386impute asynchronous data by using older sensor data for the asynchronous sensor. During training  
 387we vary the time offset  $\Delta t$  uniformly at random between 0s and a maximum of 0.5s. We consider  
 3880.5s the worst-case scenario for our robustness test, which could be uncommon but high-impact.  
 389For the asynchronous sensor we use the sensor data that precedes the reference timestamp by the  
 390specified time offset. The cameras have a frequency of 12 Hz and the LiDAR operates at 20 Hz.  
 391Since we cannot realistically infer the sensor data at an arbitrary point in time, we instead use the  
 392closest asynchronous sensor data to the specified time offset. Therefore we cannot evaluate arbi-  
 393trarily small time offsets, but only multiples of the sensor period (83ms and 50ms). nuScenes uses  
 3946 cameras which are triggered when the LiDAR scans across the center of each camera’s Field of  
 395View (FoV). Therefore each camera has a different timestamp. We treat all cameras the same way  
 396and assign them the timestamp of the back-left camera, which is closest to the reference timestamp  
 397of the corresponding lidar sweep.  
 398

#### 398 4.1 ASYNCHRONOUS MULTI-MODAL 3D OBJECT DETECTION

400 Table 1 presents the main evaluation results for large time intervals (severe asynchrony), using the  
 401 camera as the reference time frame while LiDAR is asynchronous with it. A similar table where  
 402 LiDAR is synchronous and cameras asynchronous can be found in Appendix A.1, demon-  
 403 strating the generality of AsyncBEV. Since both setups show similar trends, we focus here on the camera as  
 404 the reference setup only. Figure 4 demonstrates the NDS for All Objects, and for Dynamic Objects  
 405 only, at lower time offset ranges, representing more common less severe cases. A similar figure for  
 406 mAP is shown in Appendix A.8. Note that the last column of Table 1 showcases the FPS of each  
 407 method, which confirms our AsyncBEV only introduces marginal latency to the vanilla models.  
 408

409 **Vanilla Models.** When the asynchrony be-  
 410tween images and point cloud increases to  
 4110.5s, the performance of both 3D object de-  
 412tectors, CMT and UniBEV, decreases dramati-  
 413cally, confirming sensor asynchrony leads to  
 414performance degradation if not properly ac-  
 415counted for. Especially, the performance of  
 416Dynamic Objects declines significantly from  
 41747.5% NDS to 26.1% for CMT and 42.4%  
 418NDS to 25.3% for UniBEV with 0.5s time off-  
 419set. Even with a mild sensor asynchrony with  
 42050ms time interval, its performance for All  
 421Objects and Dynamic Objects drops substan-  
 422tially, as shown in Figure 4.



423 **Figure 4: Performance** of different methods with  
 424 increasing time offsets under LiDAR asynchrony.

425 **Ego Motion Compensation (EMC) only.** Table 1 confirms only performing EMC already  
 426 significantly improves the NDS with 0.5s asyn-  
 427chrony for CMT and UniBEV. Notably, the per-  
 428formance of static objects at 0.5s asynchrony  
 429improves by 23.3% NDS and to 67.5% for CMT and by 21.0% NDS to 62.2%, which reaches a  
 430comparable performance of static objects in the synchronous case with 69.0% and 63.5% for CMT  
 431and UniBEV. However, EMC does not bring a performance improvement for dynamic objects, as  
 432expected. Similar observations can be made in Figure 4: Though the performance of EMC improves  
 433the performance for All Objects over the vanilla CMT across all time offsets, its performance for  
 434Dynamic Objects barely increases.

432 **Data Augmentation.** Finetuning the vanilla model with asynchronous inputs as a data augmentation (DA) strategy also improves robustness to sensor asynchrony, but its effectiveness differs between CMT and UniBEV. Compared to CMT+EMC, the performance of CMT+DA in the synchronous case declines by 1.4% NDS but increases in asynchronous scenarios. In contrast, UniBEV+DA underperforms UniBEV+EMC in both cases. Notably, DA consistently improves the performance for Dynamic Objects of both models over the EMC baselines.

433  
434  
435  
436  
437  
438  
439 **AsyncBEV.** In Table 1, the CMT variants overall outperform UniBEV variants, but AsyncBEV  
440 helps improve asynchrony robustness for both types of architectures. At the time offset is  
441 0.5s, CMT+AsyncBEV and UniBEV+AsyncBEV reach 70.0% and 63.3% NDS respectively, sur-  
442 passing their EMC counterparts with a margin of 6.7% NDS and 12.3% mAP. Notably, both  
443 CMT+AsyncBEV and UniBEV+AsyncBEV achieve strong performance for Dynamic Objects,  
444 namely 43.4% and 37.8% NDS, which significantly outperforms their counterparts with EMC with  
445 a margin of 16.6% and 11.9% NDS, and the vanilla models with a considerable margin of 17.3%.  
446 For static objects, CMT+AsyncBEV performs comparable to CMT+EMC, as expected.

447 We also observe that the AsyncBEV can introduce a slight performance degradation in the synchro-  
448 nized case (0s) for both models, in its default setting where we also fine-tune the detector head. NDS  
449 for all objects shows a modest drop of 0.4% for CMT and 1.0% for UniBEV. However, AsyncBEV  
450 (FD) maintains the performance of the vanilla models in this synchronized setting. In this variant  
451 all original detector components were frozen, and the velocity-based formulation ensures that for  
452  $\Delta t = 0s$  no motion is predicted, thus all BEV features remain unaffected. Still, AsyncBEV (FD)  
453 underperforms compared to default AsyncBEV on Dynamic Objects in asynchronous cases, due to  
454 fewer trainable parameters. Fine-tuning the detector head when training AsyncBEV thus controls a  
455 trade-off between these cases. Given the modest effect on synchronous cases and strong improve-  
456 ments on asynchronous cases, we keep AsyncBEV with fine-tuning the detector heads as the default.

457 Figure 4 confirms AsyncBEV improves the robustness of the vanilla CMT also for smaller time  
458 intervals, showing flatter curves for both categories. Notably, the performance for All Objects of  
459 CMT+AsyncBEV in an extreme sensor asynchrony case, i.e., 0.5s time offset outperforms vanilla  
460 CMT with a mild 50ms time offset.

461 **Other Baselines.** Fan et al. (2025) propagates temporal features, which expects the network to  
462 align the features implicitly according to the given asynchronous time offset. As shown in Table 1,  
463 this strategy mitigates some of the misalignment for higher time offsets compared to the EMC base-  
464 line. On the other hand, it also exhibits a significant drop in performance for the 0s case compared  
465 to vanilla CMT. Overall its performance is inferior to our proposed AsyncBEV for all time offsets,  
466 with the performance gap widening as the time offset increases.

467 StreamingFlow takes a sequence of past images or LiDAR frames as input, and predicts features  
468 for a reference timestamp. Despite its performance in synchronous scenarios drops 2.5% NDS,  
469 StreamingFlow improves performance for asynchronous scenarios compared to the vanilla UniBEV,  
470 Nevertheless, in both scenarios StreamingFlow still underperforms our AsyncBEV, which only uses  
471 the latest asynchronous observations as input. Furthermore, due to the need to process every ob-  
472 servation in the input sequence with the GRU-ODE module, StreamingFlow operates significantly  
473 slower than AsyncBEV in terms of FPS.

474  
475 **Qualitative Results.** A benefit of our flow prediction approach is that such visualizations im-  
476 prove the interpretability of AsyncBEV. Figure 5 showcases the qualitative results of CMT and  
477 CMT+AsyncBEV under the extreme LiDAR asynchrony case with a 0.5s time offset. In Figure  
478 5 a), the object in the highlight box cannot be detected by the original CMT due to large spatial  
479 misalignment in the asynchronous sensor information. Predicted  $\Delta$ -BEVFlow in Figure 5 c) aligns  
480 with the GT flow shown in Figure 5 d) and matches the spatial offsets between the predicted box  
481 and the ground truth box. Thanks to the motion cue provided by  $\Delta$ -BEVFlow, CMT+AsyncBEV  
482 successfully detected the object with a small translation error, as shown in Figure 5 b).

483 To sum up, AsyncBEV brings significant performance gains for 3D object detectors under both  
484 severe and light sensor asynchrony, especially for dynamic objects. The approach is generic and  
485 light-weight, improving detector robustness for arbitrary time offsets during operation, supporting  
different sensor modalities, and distinct 3D object detection architectures.

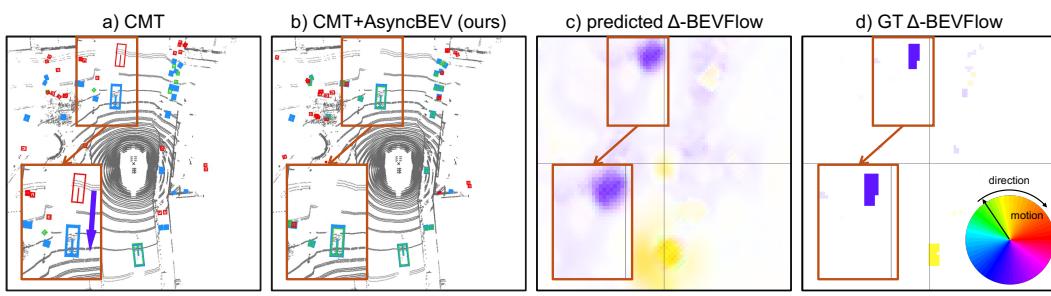


Figure 5: **Qualitative Results of AsyncBEV.** The left two columns demonstrate the 3D object detection performance of CMT and CMT+AsyncBEV under 0.5s LiDAR asynchrony. The right two columns show the output of our predicted  $\Delta$ -BEVFlow and the ground truth  $\Delta$ -BEVFlow.

## 4.2 ABLATION STUDY OF $\Delta$ -BEVFLOW STRATEGY

In this ablation study, we focus on UniBEV only, as performance trends on UniBEV and CMT are generally similar. Table 2 compares the motion-based formulation to the (default) velocity-based formulation of  $\Delta$ -BEVFlow, and the impact of exposing the model to varying delays during training. The original UniBEV performs well in a synchronized case, but fails catastrophically when a serious asynchrony occurs. Based on UniBEV+EMC, we finetune UniBEV with the asynchronous data (EMC + DA) brings significant improvement for the asynchronous case, but also degrades the performance in the synchronized case. Motion-based flow estimation (BEV-ME) helps compensate for the dynamic objects. Velocity-based flow estimation (BEV-VE) still further prevents the performance drop in the synchronization case, and strengthens the results under extreme sensor asynchrony, validating our choice for this formulation as default.

$\Delta$ -BEVFlow design			All Objects (0s)		All Objects (0.5s)	
EMC+DA	Motion	Velocity	NDS	mAP	NDS	mAP
✗	✗	✗	<b>66.7</b>	<b>62.0</b>	39.7	22.2
✓	✗	✗	63.1 $\downarrow$ 3.6	57.6 $\downarrow$ 4.4	61.1 $\uparrow$ 21.4	54.9 $\uparrow$ 32.7
✓	✓	✗	64.2 $\downarrow$ 2.5	58.6 $\downarrow$ 3.4	62.5 $\uparrow$ 22.8	56.1 $\uparrow$ 33.9
✓	✗	✓	65.7 $\downarrow$ 1.0	60.8 $\downarrow$ 1.2	<b>63.3</b> $\uparrow$ 23.6	<b>57.1</b> $\uparrow$ 34.9

Table 2: **Ablation study of  $\Delta$ -BEVFlow design.** This table showcases the core design choice for the new task  $\Delta$ -BEVFlow. **EMC+DA:** EMC + asynchronous data augmentation training. **Motion:** BEV-ME. **Velocity:** BEV-VE, which is our default design.

## 5 CONCLUSIONS

We have proposed AsyncBEV, a light-weight and generic module to improve the robustness of a multi-modal 3D object detector against sensor asynchrony. Generalizing the scene flow concept, we first propose a novel task,  $\Delta$ -BEVFlow, which predicts the flow in the BEV space given an arbitrary time offset for multi-modal BEV features. The predicted flow can then be applied to spatially align the feature maps of the asynchronous sensor to the reference frame. Our experiments on nuScenes show AsyncBEV significantly improves robustness under sensor asynchrony for both token-based and grid-based detectors, with only little computational overhead. Under an extreme LiDAR synchrony case with a 0.5 s time offset, AsyncBEV improves the performance for All Objects of vanilla methods with a considerable margin of 26.8% NDS on token-based CMT and 23.6% NDS on grid-based UniBEV. Particularly, AsyncBEV outperforms EMC for Dynamic Objects with 16.6% NDS on CMT and 11.9% NDS on UniBEV. Similar trends were observed for smaller time offsets. We also showed a velocity-based formulation for the  $\Delta$ -BEVFlow estimator helps regularize the predicted motion of dynamic objects over directly predicting the motion.

**Limitations and Future Work.** In this work, we assume that one sensor is always synchronized with the reference frame, while the other may have time offsets. However, in the real autonomous driving application, there can be more than two sensors, of which multiple can be asynchronous at once. Please refer to Appendix A.9 for more details. In the future, we will develop a full-streaming framework to always fuse the most up-to-date data from multiple sensors, including radar.

540 REFERENCES  
541

542 Pedro HE Becker, Jose Maria Arnau, and Antonio González. Demystifying power and performance  
543 bottlenecks in autonomous driving systems. In *2020 IEEE International Symposium on Workload  
544 Characterization (IISWC)*, 2020.

545 Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush  
546 Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multimodal dataset for  
547 autonomous driving. In *CVPR*, 2020.

548 Hongxiang Cai, Zeyuan Zhang, Zhenyu Zhou, Ziyin Li, Wenbo Ding, and Jiahua Zhao. BEV-  
549 Fusion4D: Learning lidar-camera fusion under bird's-eye-view via cross-modality guidance and  
550 temporal aggregation. *arXiv preprint arXiv:2303.17099*, 2023.

552 Xuanyao Chen, Tianyuan Zhang, Yue Wang, Yilun Wang, and Hang Zhao. Futr3d: A unified sensor  
553 fusion framework for 3d detection. In *CVPR*, 2022.

555 Oscar de Groot, Alberto Bertipaglia, Hidde Boekema, Vishrut Jain, Marcell Kegl, Varun Kotian,  
556 Ted Lentsch, Yancong Lin, Chrysovalantou Messiou, Emma Schippers, et al. A vehicle system for  
557 navigating among vulnerable road users including remote operation. In *IV*, 2025.

558 Marel Faizullin, Anastasiia Kornilova, and Gonzalo Ferrer. Open-source lidar time synchronization  
559 system by mimicking gnss-clock. In *ISPCS*, 2022.

561 Meng Fan, Yifan Zuo, Patrick Blaes, Harley Montgomery, and Subhasis Das. Robust sensor fusion  
562 against on-vehicle sensor staleness. In *CVPRW*, 2025.

563 Chongjian Ge, Junsong Chen, Enze Xie, Zhongdao Wang, Lanqing Hong, Huchuan Lu, Zhenguo  
564 Li, and Ping Luo. MetaBEV: Solving sensor failures for bev detection and map segmentation. In  
565 *CVPR*, 2023.

567 Ajay Kumar Gurumadaiah, Jaehyeong Park, Jin-Hee Lee, JeSeok Kim, and Soon Kwon. Precise  
568 synchronization between lidar and multiple cameras for autonomous driving: An adaptive ap-  
569 proach. *IEEE T-IV*, 2025.

570 Adam W. Harley, Zhaoyuan Fang, Jie Li, Rares Ambrus, and Katerina Fragkiadaki. Simple-BEV:  
571 What really matters for multi-sensor bev perception? In *ICRA*, 2023.

573 Seamie Hayes, Sushil Sharma, and Ciarán Eising. Velocity driven vision: asynchronous sensor  
574 fusion birds eye view models for autonomous vehicles. In *IET Conference Proceedings CP887*,  
575 2024.

577 Junjie Huang and Guan Huang. BEVDet4D: Exploit temporal cues in multi-camera 3d object de-  
578 tection. *arXiv preprint arXiv:2203.17054*, 2022.

590 Junjie Huang, Guan Huang, Zheng Zhu, Ye Yun, and Dalong Du. BEVDet: High-performance  
591 multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021.

592 Keli Huang, Botian Shi, Xiang Li, Xin Li, Siyuan Huang, and Yikang Li. Multi-modal sensor fusion  
593 for auto driving perception: A survey. *arXiv preprint arXiv:2202.02703*, 2022.

594 Philipp Jund, Chris Sweeney, Nichola Abdo, Zhifeng Chen, and Jonathon Shlens. Scalable scene  
595 flow from point clouds in the real world. *RA-L*, 2021.

597 Jaeyeul Kim, Jungwan Woo, Utkcheol Shin, Jean Oh, and Sunghoon Im. Flow4d: Leveraging 4d  
598 voxel network for lidar scene flow estimation. *RA-L*, 2024.

599 Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying Voxel-based  
600 Representation with Transformer for 3D Object Detection. In *NeurIPS*, 2022a.

602 Zhiqi Li, Wenhui Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng  
603 Dai. BEVFormer: Learning bird's-eye-view representation from multi-camera images via spa-  
604 tiotemporal transformers. In *ECCV*, 2022b.

594 Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang,  
 595 Bing Wang, and Zhi Tang. BEVFusion: A simple and robust lidar-camera fusion framework. In  
 596 *NeurIPS*, 2022.

597 Yancong Lin, Shiming Wang, Liangliang Nan, Julian Kooij, and Holger Caesar. VoteFlow: Enforc-  
 598 ing local rigidity in self-supervised scene flow. In *CVPR*, 2025.

600 Zhijian Liu, Haotian Tang, Alexander Amini, Xingyu Yang, Huizi Mao, Daniela Rus, and Song  
 601 Han. BEVFusion: Multi-task multi-sensor fusion with unified bird's-eye view representation. In  
 602 *ICRA*, 2023.

603 Andras Palffy, Ewoud Pool, Srimannarayana Baratam, Julian Kooij, and Dariu Gavrila. Multi-class  
 604 Road User Detection with 3+1D Radar in the View-of-Delft Dataset. *RA-L*, 2022.

605 Jinhyung Park, Chenfeng Xu, Shijia Yang, Kurt Keutzer, Kris Kitani, Masayoshi Tomizuka, and Wei  
 606 Zhan. Time will tell: New outlooks and a baseline for temporal multi-view 3d object detection.  
 607 In *ICLR*, 2023.

608 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomed-  
 609 ical image segmentation. In *MICCAI*, 2015.

610 Md Hasan Shahriar, Md Mohaimin Al Barat, Harshavardhan Sundar, Naren Ramakrishnan,  
 611 Y Thomas Hou, and Wenjing Lou. On the fragility of multimodal perception to temporal mis-  
 612 alignment in autonomous driving. *arXiv preprint arXiv:2507.09095*, 2025.

613 Yining Shi, Kun Jiang, Ke Wang, Jiusi Li, Yunlong Wang, Mengmeng Yang, and Diange Yang.  
 614 Streamingflow: Streaming occupancy forecasting with asynchronous multi-modal data streams  
 615 via neural ordinary differential equation. In *CVPR*, 2024.

616 Fikret Sivrikaya and Bülent Yener. Time synchronization in sensor networks: a survey. *IEEE*  
 617 *network*, 2004.

618 Zhihang Song, Lihui Peng, Jianming Hu, Danya Yao, and Yi Zhang. Timealign: A multi-modal  
 619 object detection method for time misalignment fusing in autonomous driving, 2024a.

620 Ziying Song, Lin Liu, Feiyang Jia, Yadan Luo, Caiyan Jia, Guoxin Zhang, Lei Yang, and Li Wang.  
 621 Robustness-aware 3d object detection in autonomous driving: A review and outlook. *IEEE Trans-*  
 622 *actions on Intelligent Transportation Systems*, 2024b.

623 Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui,  
 624 James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for au-  
 625 tonomous driving: Waymo open dataset. In *CVPR*, 2020.

626 Haiyang Wang, Hao Tang, Shaoshuai Shi, Aoxue Li, Zhenguo Li, Bernt Schiele, and Liwei Wang.  
 627 Unitr: A unified and efficient multi-modal transformer for bird's-eye-view representation. In  
 628 *ICCV*, 2023a.

629 Shihao Wang, Yingfei Liu, Tiancai Wang, Ying Li, and Xiangyu Zhang. Exploring object-centric  
 630 temporal modeling for efficient multi-view 3d object detection. In *ICCV*, 2023b.

631 Shiming Wang, Holger Caesar, Liangliang Nan, and Julian FP Kooij. Unibev: Multi-modal 3d  
 632 object detection with uniform bev encoders for robustness against missing sensor modalities. In  
 633 *IV*, 2024.

634 Yingjie Wang, Qiuyu Mao, Hanqi Zhu, Jiajun Deng, Yu Zhang, Jianmin Ji, Houqiang Li, and Yany-  
 635 ong Zhang. Multi-modal 3d object detection in autonomous driving: a survey. *IJCV*, 2023c.

636 Sizhe Wei, Yuxi Wei, Yue Hu, Yifan Lu, Yiqi Zhong, Siheng Chen, and Ya Zhang. Asynchrony-  
 637 robust collaborative perception via bird's eye view flow. In *NeurIPS*, 2023.

638 Benjamin Wilson, William Qi, Tanmay Agarwal, John Lambert, Jagjeet Singh, Siddhesh Khandel-  
 639 wal, Bowen Pan, Ratnesh Kumar, Andrew Hartnett, Jhony Kaesemel Pontes, Deva Ramanan,  
 640 Peter Carr, and James Hays. Argoverse 2: Next generation datasets for self-driving perception  
 641 and forecasting. In *NeurIPS*, 2021.

648 Shaoyuan Xie, Lingdong Kong, Wenwei Zhang, Jiawei Ren, Liang Pan, Kai Chen, and Ziwei Liu.  
649 Robobev: Towards robust bird’s eye view perception under corruptions. *TPAMI*, 2023.  
650

651 Junjie Yan, Yingfei Liu, Jianjian Sun, Fan Jia, Shuailin Li, Tiancai Wang, and Xiangyu Zhang.  
652 Cross modal transformer via coordinates encoding for 3d object detection. *arXiv preprint*  
653 *arXiv:2301.01283*, 2023.

654 Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking.  
655 In *CVPR*, 2021.

656

657 Haibao Yu, Wenxian Yang, Jiaru Zhong, Zhenwei Yang, Siqi Fan, Ping Luo, and Zaiqing Nie. End-  
658 to-end autonomous driving through v2x cooperation. In *AAAI*, 2025.

659 Kaiwen Yuan, Li Ding, Mazen Abdelfattah, and Z Jane Wang. Licas3: A simple lidar–camera  
660 self-supervised synchronization method. *IEEE T-RO*, 2022.

661

662 Qingwen Zhang, Yi Yang, Heng Fang, Ruoyu Geng, and Patric Jensfelt. Deflow: Decoder of scene  
663 flow network in autonomous driving. In *ICRA*, 2024.

664

665 Qingwen Zhang, Ajinkya Khoche, Yi Yang, Li Ling, Sina Sharif Mansouri, Olov Andersson, and  
666 Patric Jensfelt. Himo: High-speed objects motion compensation in point clouds. *T-RO*, 2025a.

667

668 Qingwen Zhang, Yi Yang, Peizheng Li, Olov Andersson, and Patric Jensfelt. Seflow: A self-  
669 supervised scene flow method in autonomous driving. In *ECCV*, 2025b.

670

671

672

673

674

675

676

677

678

679

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A APPENDIX**703 **A.1 MULTI-MODAL 3D OBJECTION WITH ASYNCHRONOUS CAMERAS**

707 Method		708 All Objects			709 Static Objects			710 Dynamic Objects		
		711 0s	712 0.25s	713 0.5s	714 0s	715 0.25s	716 0.5s	717 0s	718 0.25s	719 0.5s
720 CMT	NDS $\uparrow$	<b>72.9</b>	70.7	67.5	<b>69.0</b>	66.9	63.8	<b>47.5</b>	<b>46.9</b>	45.8
	mAP $\uparrow$	<b>70.3</b>	66.7	61.2	<b>61.0</b>	57.9	52.4	<b>36.8</b>	<b>36.0</b>	34.2
721 CMT + EMC	NDS $\uparrow$	<b>72.9</b>	71.4	68.8	<b>69.0</b>	<b>67.6</b>	<b>65.1</b>	<b>47.5</b>	46.9	45.9
	mAP $\uparrow$	<b>70.3</b>	67.8	63.5	<b>61.0</b>	<b>58.9</b>	<b>54.9</b>	<b>36.8</b>	36.0	34.1
722 CMT+AsyncBEV	NDS $\uparrow$	72.7 $\downarrow 0.2$	<b>72.7</b> $\uparrow 1.3$	<b>72.3</b> $\uparrow 3.5$	68.8 $\downarrow 0.2$	<b>68.8</b> $\uparrow 1.2$	<b>68.4</b> $\uparrow 3.3$	<b>48.1</b> $\uparrow 0.6$	<b>48.1</b> $\uparrow 1.2$	<b>48.0</b> $\uparrow 2.1$
	mAP $\uparrow$	69.8 $\downarrow 0.5$	<b>69.7</b> $\uparrow 1.9$	<b>69.1</b> $\uparrow 5.6$	60.8 $\downarrow 0.3$	<b>60.6</b> $\uparrow 1.7$	<b>60.2</b> $\uparrow 5.3$	<b>37.0</b> $\uparrow 0.1$	<b>36.9</b> $\uparrow 0.9$	<b>36.7</b> $\uparrow 2.6$
723 UniBEV	NDS $\uparrow$	<b>66.7</b>	64.8	62.7	<b>63.5</b>	61.9	60.0	<b>42.4</b>	41.8	40.9
	mAP $\uparrow$	<b>62.0</b>	59.3	56.1	<b>52.2</b>	50.1	47.1	<b>32.9</b>	<b>32.3</b>	30.6
724 UniBEV+EMC	NDS $\uparrow$	<b>66.7</b>	<b>66.5</b>	<b>66.2</b>	<b>63.5</b>	<b>63.5</b>	<b>63.2</b>	<b>42.4</b>	<b>41.9</b>	41.0
	mAP $\uparrow$	<b>62.0</b>	61.4	60.8	<b>52.2</b>	<b>52.0</b>	<b>51.6</b>	<b>32.9</b>	<b>32.2</b>	30.9
725 UniBEV+AsyncBEV	NDS $\uparrow$	<b>66.9</b> $\uparrow 0.2$	<b>66.8</b> $\uparrow 0.3$	<b>66.7</b> $\uparrow 0.5$	63.4 $\downarrow 0.1$	63.3 $\downarrow 0.2$	<b>63.2</b>	42.0 $\downarrow 0.4$	<b>42.1</b> $\uparrow 0.2$	<b>42.2</b> $\uparrow 1.2$
	mAP $\uparrow$	<b>62.0</b>	<b>61.8</b> $\uparrow 0.4$	<b>61.5</b> $\uparrow 0.7$	51.7 $\downarrow 0.5$	51.6 $\downarrow 0.4$	51.5 $\downarrow 0.1$	32.3 $\downarrow 0.6$	32.3 $\uparrow 0.1$	32.3 $\uparrow 1.4$

717 Table III: Evaluation on the nuScenes **val** set of asynchronous multi modal 3D object detection. **(Best, Second Best)**. Camera is asynchronous. Trained with one model and evaluated on different time offsets. Performance deltas are measured relative to the EMC baseline.

718 Table III demonstrates an opposite asynchrony, i.e., LiDAR is on the reference timestamp, but cameras are asynchronous. Because the dominating sensor, LiDAR, is still functioning, the performance 719 degradation is not as significant as in Table 1. We can still observe that NDS for All Objects 720 decreases 5.4% to 67.5% for CMT and 5.0% to 62.7% for UniBEV with a 0.5s asynchronous time 721 offset, respectively. AsyncBEV improves the performance of CMT and UniBEV to 72.3% and 722 66.7% NDS, respectively. Table III validates the generality of our proposed AsyncBEV.

730 **A.2 COMPARISON WITH THE SENSOR-FALLBACK STRATEGY**

733 Method	734 All Objects		735 Static Objects		736 Dynamic Objects	
	NDS $\uparrow$	mAP $\uparrow$	NDS $\uparrow$	mAP $\uparrow$	NDS $\uparrow$	mAP $\uparrow$
CMT (0.5s L)	43.2	25.8	44.2	22.8	26.1	9.0
CMT_C	44.7 $\uparrow 1.5$	38.3 $\uparrow 12.5$	47.4 $\uparrow 3.2$	29.5 $\uparrow 6.7$	28.3 $\uparrow 2.2$	14.7 $\uparrow 5.7$
CMT+AsyncBEV (0.5s L)	70.0 $\uparrow 25.3$	<b>66.4</b> $\uparrow 28.1$	67.6 $\uparrow 20.2$	59.1 $\uparrow 29.6$	43.4 $\uparrow 15.1$	32.3 $\uparrow 17.6$
UniBEV (0.5s L)	39.7	22.2	41.2	19.0	25.3	8.9
UniBEV_C	42.6 $\uparrow 2.9$	35.4 $\uparrow 13.2$	46.1 $\uparrow 4.9$	27.1 $\uparrow 8.1$	28.2 $\uparrow 2.9$	13.1 $\uparrow 4.2$
UniBEV+AsyncBEV (0.5s L)	63.3 $\uparrow 20.7$	57.1 $\uparrow 21.7$	61.5 $\uparrow 15.4$	49.1 $\uparrow 22.0$	37.8 $\uparrow 9.6$	26.6 $\uparrow 13.5$
CMT (0.5s C)	67.5	61.2	63.8	52.4	45.8	34.2
CMT_L	68.1 $\uparrow 0.6$	61.7 $\uparrow 0.5$	64.1 $\uparrow 0.3$	53.0 $\uparrow 0.6$	46.7 $\uparrow 0.9$	34.5 $\uparrow 0.3$
CMT+AsyncBEV (0.5s C)	72.3 $\uparrow 4.2$	<b>69.1</b> $\uparrow 7.4$	68.4 $\uparrow 4.3$	60.2 $\uparrow 7.2$	48.0 $\uparrow 1.3$	36.7 $\uparrow 2.2$
UniBEV (0.5s C)	62.7	56.1	60.0	47.1	40.9	30.6
UniBEV_L	63.7 $\uparrow 1.0$	56.1	60.5 $\uparrow 0.5$	47.7 $\uparrow 0.6$	42.2 $\uparrow 1.3$	31.7 $\uparrow 1.1$
UniBEV+AsyncBEV (0.5s C)	66.7 $\uparrow 3.0$	61.5 $\uparrow 5.4$	63.2 $\uparrow 2.7$	51.5 $\uparrow 3.8$	42.2	32.3 $\uparrow 0.6$

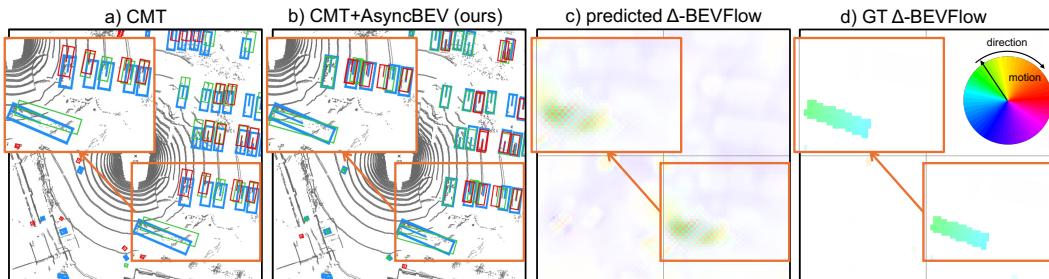
747 Table IV: Performance of single-modality performance of CMT and UniBEV.

748 In our 0.25s and 0.5s delay setups, discarding stale data ( $> 100ms$ ) is equivalent to single-modality 749 inference. Table IV compares our AsyncBEV to single-modality. As shown in Table IV, when 750 LiDAR is delayed and a sensor fallback strategy is adopted, camera-only CMT\_C achieves 44.7% 751 NDS while camera-only UniBEV\_C performs at 42.6% NDS. Both outperform the vanilla CMT 752 (43.2% NDS) and UniBEV (39.7% NDS) under the 0.5s LiDAR asynchrony. This shows that the 753 sensor-fallback strategy can serve as a backup option in cases of severe sensor asynchrony, even 754 though the performance gains remain minor. However, our CMT+AsyncBEV achieves 70.0% NDS, 755

756 and UniBEV achieves 63.3% NDS even with an extreme 0.5s time offset for LiDAR, outperforming  
 757 the single-modality performance significantly.  
 758

759 We make similar observations for the camera-delayed scenario. With a 0.5s time offset for cameras,  
 760 CMT+AsyncBEV achieves 72.3% NDS, while UniBEV+AsyncBEV achieves 66.7% NDS, both  
 761 surpassing the LiDAR-only performance of CMT\_L (68.1% NDS) and UniBEV\_L (63.7% NDS).  
 762 This shows that the asynchronous modality still contains rich information for perception tasks. In  
 763 conclusion, effectively leveraging temporally misaligned information is a better option than com-  
 764 pletely dropping them.  
 765

### 766 A.3 MORE QUALITATIVE RESULTS



778 **Figure 6: Qualitative Results of AsyncBEV.** The left two columns demonstrate the 3D object  
 779 detection performance of CMT and CMT+ AsyncBEV under LiDAR asynchrony with 0.5s offset.  
 780 The right two columns show the output of our BEV-VE module and the ground truth flow in the  
 781 BEV space.

782 This sample shows a driving scenario in a parking lot, in which most of the objects are static.  
 783 detection of CMT introduces large translation errors for predictions due to the ego vehicle motion.  
 784 CMT+AsyncBEV corrects most of the static boxes and re-aligns the moving long bus. The color  
 785 and saturation of the predicted flow align with the GT flow, which validates the interpretability of  
 786 our proposed AsyncBEV.  
 787

### 788 A.4 DENSE GT FLOW REPRESENTATION

791 The supervised scene flow task takes the ground truth flow as supervision to calculate a per-point  
 792 flow for the source point cloud. The ground truth flow is usually generated by calculating the motion  
 793 between bounding boxes (Jund et al., 2021; Zhang et al., 2024). Hence, this flow representation in-  
 794 herits the sparsity of point clouds. As point clouds usually only capture partial objects, this per-point  
 795 sparse flow representation cannot fully present the object motion in a dense BEV grid. Therefore,  
 796 we propose a dense ground truth flow representation, dense flow, which only relies on bounding  
 797 boxes. By decoupling from point clouds, dense flow facilitates generating ground truth flow for  
 798 camera BEV features.

799 The process to generate dense flow is shown in Algorithm 1. The core changes compared with  
 800 the generation of scene flow ground truth are that we use a dense standard grid  $\mathbb{G}_{BEV}$  to represent  
 801 the dense and evenly distributed point cloud instead of using the real point cloud. Then we adopt  
 802 a FlowEmbedder to voxelize  $\mathbb{P}^{t_0}$  into the voxel space and scatter  $\mathbb{V}_{BEV}^{t_0 \rightarrow t_1}$  with the voxelization  
 803 indexes of  $\mathbb{P}^{t_0}$  into the corresponding voxel cell, resulting a dense flow representation  $\mathbf{V}_{BEV}^{t_0 \rightarrow t_1}$ , which  
 804 spatially aligns with camera and LiDAR BEV features.  $\mathbf{V}_{BEV}^{t_0 \rightarrow t_1}$  is used to supervise the predicted  
 805 flow  $\mathbf{V}_{flow}^{t_0 \rightarrow t_1}$ .

### 806 A.5 ALGORITHM TO GENERATE DENSE FLOW REPRESENTATION

807 Algorithm 1 summarizes the steps to generate a dense ground truth flow representation from bound-  
 808 ing boxes, which serves as supervision for  $\Delta$ -BEVFlow.  
 809

**Algorithm 1** Generation of dense ground truth flow representation

---

**Require:** Bounding boxes at time  $t_0$ ,  $\mathbb{B}^{t_0}$ ; Bounding boxes at  $t_1$ ,  $\mathbb{B}^{t_1}$ ; ego motion  $\mathbf{M}^{t_0 \rightarrow t_1}$ ; standard grid  $\mathbb{G}_{BEV}$

1: **for** each box  $b_i^{t_0} \sim \mathbb{B}^{t_0}$  **do**

2:   **for** each box  $b_j^{t_1} \sim \mathbb{B}^{t_1}$  **do** ▷ Boxes are defined in the global coordinate.

3:     **if**  $b_i^{t_0}.\text{ID} == b_j^{t_1}.\text{ID}$  **then**

4:        $T_{box2g}^{t_1} \leftarrow b_j^{t_1}.\text{center}, b_j^{t_1}.\text{orientation}$

5:        $T_{box2g}^{t_0} \leftarrow b_i^{t_0}.\text{center}, b_i^{t_0}.\text{orientation}$

6:        $T_{l2e}^{t_0}, T_{e2g}^{t_0} \leftarrow \mathbf{M}^{t_0 \rightarrow t_1}$

7:        $T_{box2t_0}^{t_1} \leftarrow (T_{l2e}^{t_0})^{-1} \cdot (T_{e2g}^{t_0})^{-1} \cdot T_{box2g}^{t_1}$  ▷ Transfer the box at  $t_1$  to LiDAR frame at  $t_0$

8:        $T_{box2t_0}^{t_0} \leftarrow (T_{l2e}^{t_0})^{-1} \cdot (T_{e2g}^{t_0})^{-1} \cdot T_{box2g}^{t_0}$  ▷ Transfer the box at  $t_0$  to LiDAR frame at  $t_0$

9:        $R_{t_02t_1} \leftarrow T_{box2t_0}^{t_1} \cdot (T_{box2t_0}^{t_0})^{-1}$  ▷ Calculate the motion of boxes from  $t_0$  to  $t_1$

10:       $\mathbb{P}_{b_i^{t_0}} \leftarrow f_{pts\_in\_boxes}(b_i^{t_0}, \mathbb{G}_{BEV})$  ▷ Get grid points in the box at  $t_0$

11:       $\mathbb{V}_{b_i^{t_0}}^{t_0 \rightarrow t_1} \leftarrow R_{t_02t_1} \cdot \mathbb{P}_{b_i^{t_0}} - \mathbb{P}_{b_i^{t_0}}$  ▷ Calculate the scene flow for the points within box

12:

13:     **end if**

14:   **end for**

15: **end for**

16: **return**  $\mathbb{V}_{BEV}^{t_0 \rightarrow t_1} \leftarrow \{\mathbb{V}_{b_i^{t_0}}^{t_0 \rightarrow t_1}, b_i^{t_0} \in \mathbb{B}^{t_0}\}, \mathbb{P}^{t_0} \leftarrow \{\mathbb{P}_{b_i^{t_0}}, b_i^{t_0} \in \mathbb{B}^{t_0}\}$

## A.6 ABLATION OF BEV-VE DESIGN.

Loss	MM	All Objects (0.5s)		Dynamic Objects (0.5s)	
		NDS	mAP	NDS	mAP
✗	✗	63.0	56.8	37.0	25.4
✓	✗	63.2 $\uparrow$ 0.2	57.0 $\uparrow$ 0.2	37.6 $\uparrow$ 0.6	26.0 $\uparrow$ 0.6
✓	✓	63.3 $\uparrow$ 0.3	57.1 $\uparrow$ 0.3	37.8 $\uparrow$ 0.8	26.6 $\uparrow$ 1.2

Table V: Ablation of BEV-VE inputs. MM: Multi-modal features, L: Flow loss.

Table V ablates the effect of whether to use the multi-modal features (MM) and the flow loss (L) with the performance of all objects and dynamic objects under the 0.5s time offset. The basic setting in the first row also applies the concept of BEV-VE to explicitly predict the dynamic object flow and then warp the feature-coordinate mapping, but only takes the asynchronous features as input and is supervised by the detection loss in an end-to-end manner. The basic setting already achieves good performance. Using multi-modal features (MM) and the flow loss (Loss) for supervision further improves the performance under asynchrony with a small margin.

## A.7 ABLATION OF IMAGE BEV FEATURES.

Are the BEV features from SimpleBEV still necessary for a grid-based 3D object detector, which already builds BEV features from the multi-view image features? Table VI ablates the effect of the separate BEV encoder, which provides performance on All Objects and Dynamic Objects of UniBEV under LiDAR asynchrony with a  $0.5s$  time offset. The first row shows the basic setting, in which our BEV-VE directly takes the BEV features for 3D object detection as input, while the second row showcases that an extra SimpleBEV encoder is used specially for the  $\Delta$ -BEVFlow estimation. Despite the basic setting already achieving good performance with 62.7% NDS for All Objects and 37.3% NDS for Dynamic Objects, its counterpart with SimpleBEV consistently improves the performance with a small margin.

Loss	MM	All Objects (0.5s)		Dynamic Objects (0.5s)	
		NDS	mAP	NDS	mAP
✗	✗	63.0	56.8	37.0	25.4
✓	✗	63.2 $\uparrow 0.2$	57.0 $\uparrow 0.2$	37.6 $\uparrow 0.6$	26.0 $\uparrow 0.6$
✓	✓	63.3 $\uparrow 0.3$	57.1 $\uparrow 0.3$	37.8 $\uparrow 0.8$	26.6 $\uparrow 1.2$

Table V: Ablation of BEV-VE inputs. MM: Multi-modal features. L: Flow loss.

SimpleBEV	All Objects (0.5s)		Dynamic Objects(0.5s)	
	NDS	mAP	NDS	mAP
✗	62.7	56.1	37.3	25.7
✓	<b>63.3</b> $\uparrow$ 0.6	<b>57.1</b> $\uparrow$ 1.0	<b>37.8</b> $\uparrow$ 0.5	<b>26.6</b> $\uparrow$ 0.9

Table VI: Ablation of SimpleBEV encoder

864 A.8 PERFORMANCE DEGRADATION FOR INCREASING TIME OFFSETS  
865

866 Figure 7 further demonstrates the curves  
867 of mAP with increasing time offsets for  
868 three methods, vanilla CMT, CMT+EMC and  
869 CMT+AsyncBEV, which shares the same trend  
870 with Figure 4. Performance on all objects  
871 and dynamic objects is reported in the figure.  
872 CMT+AsyncBEV has the most robust perfor-  
873 mance against sensor asynchrony. Its per-  
874 formance in the extreme case with  $0.5s$  is better  
875 than vanilla CMT with one-frame ( $50ms$ ) asyn-  
876 chrony. Despite CMT+EMC delivering im-  
877 proved performance for All Objects, its per-  
878 formance for Dynamic Objects still shares the  
879 same degradation with vanilla CMT, which  
880 proves the necessity of a module to further  
881 compensate for the dynamic object motion.

883 A.9 EXTENSION TO MULTI-SENSOR ASYNCHRONY  
884

885 The AsyncBEV framework can be extended to multi-sensor asynchrony scenarios. We presume  
886 to have one reference sensor (e.g., Lidar) and treat all other (asynchronous) sensors (e.g., cameras  
887 and radars) as paired to the LiDAR reference only. This approach requires one more lightweight  
888 AsyncBEV module for each additional asynchronous sensor modality, since network weights can  
889 be shared between sensors of the same modality. Possibly, multiple sensors of the same modality  
890 could be batched (with a vector containing all their  $\Delta t$ ) in a single forward pass to compute their  
891  $\Delta$ -BEVFlow in parallel.

892 For the multi-sensor asynchrony scenario, the ad-  
893 dditional runtime and memory cost of AsyncBEV is  
894 negligible compared to the dominant cost from pro-  
895 cessing the extra sensor inputs. As shown in Ta-  
896 ble 1, AsyncBEV reduces the inference speed of  
897 CMT by only 0.4 FPS, and this overhead drops  
898 to 0.1 FPS when applied to a heavier baseline  
899 such as UniBEV. Furthermore, Table VII shows  
900 that the camera branch has  $232\times$  the parameters  
901 of AsyncBEV, while the LiDAR branch has  $23\times$ .  
902 Therefore, when scaling to setups with more sensors,  
903 the backbone becomes substantially heavier, and the  
904 marginal runtime and memory overhead introduced by AsyncBEV becomes negligible compared to  
905 the total computation.

906 Finally, we note our experiments also use all six cameras from nuScenes, treating them as a single  
907 sensor. Indeed, performance might even improve further if they are handled all independently, but  
908 our results already show the benefit of our approach even by grouping related asynchronous sensors.

909 A.10 THE USE OF LARGE LANGUAGE MODELS (LLMs)  
910

911 Large Language Models (LLMs) were used only to polish the grammar and phrasing of sentences  
912 and to search for potentially relevant prior works. All conceptual development, experimental design,  
913 implementation, analysis, and initial writing of the paper were carried out by the authors.

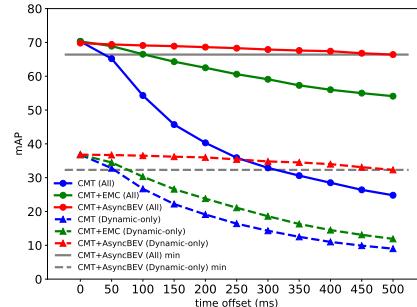


Figure 7: **Performance of CMT, CMT+EMC and CMT+AsyncBEV for different time offsets under LiDAR asynchrony.**

Module	# Params	$\times$ AsyncBEV
Camera branch	70,572,864	232
LiDAR branch	6,975,184	23
AsyncBEV	304,114	1

Table VII: # Params of CMT+AsyncBEV’s modules.  $\times$  AsyncBEV is the parameter ratio of each module relative to AsyncBEV.