

BEVDilation: LiDAR-Centric Multi-Modal Fusion for 3D Object Detection

Guowen Zhang¹, Chenhang He¹, Liyi Chen¹, Lei Zhang^{1*}

¹ The Hong Kong Polytechnic University

{guowen.zhang, liyi0308.chen}@connect.polyu.hk, chenhang.he@polyu.edu.hk, cslzhang@comp.polyu.edu.hk

Abstract

Integrating LiDAR and camera information in the bird’s eye view (BEV) representation has demonstrated its effectiveness in 3D object detection. However, because of the fundamental disparity in geometric accuracy between these sensors, indiscriminate fusion in previous methods often leads to degraded performance. In this paper, we propose BEVDilation, a novel LiDAR-centric framework that prioritizes LiDAR information in the fusion. By formulating image BEV features as implicit guidance rather than naive concatenation, our strategy effectively alleviates the spatial misalignment caused by image depth estimation errors. Furthermore, the image guidance can effectively help the LiDAR-centric paradigm to address the sparsity and semantic limitations of point clouds. Specifically, we propose a Sparse Voxel Dilation Block that mitigates the inherent point sparsity by densifying foreground voxels through image priors. Moreover, we introduce a Semantic-Guided BEV Dilation Block to enhance the LiDAR feature diffusion processing with image semantic guidance and long-range context capture. On the challenging nuScenes benchmark, BEVDilation achieves better performance than state-of-the-art methods while maintaining competitive computational efficiency. Importantly, our LiDAR-centric strategy demonstrates greater robustness to depth noise compared to naive fusion.

Code — <https://github.com/gwenzhang/BEVDilation>

Introduction

3D object detection plays an important role in applications of autonomous driving, virtual reality, and robots. To ensure precise and robust perception, LiDAR-camera fusion (Huang et al. 2024; Liu et al. 2023b; Liang et al. 2022; Yang et al. 2022; Wang et al. 2023a; Li et al. 2024a) is employed to integrate complementary information. In specific, LiDAR point clouds can offer accurate location and geometry information, while RGB images provide rich semantic and context information. However, significant view discrepancies and inherent modality limitations make robust and effective 3D multi-modal detection particularly challenging. To tackle this challenge, recent works focus on establishing cross-model correspondences to enable unified represen-

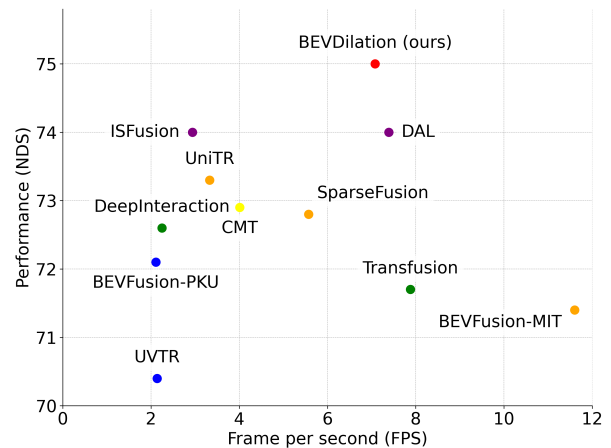


Figure 1: 3D object detection performance (NDS) vs speed (FPS) on nuScenes validation set.

tation learning. Input-based (Vora et al. 2020; Wang et al. 2021; Yin, Zhou, and Krähenbühl 2021) and query-based methods (Yang et al. 2022; Li et al. 2024a; Yan et al. 2023; Bai et al. 2022) enhance LiDAR point clouds or object proposals by integrating semantic information from RGB images. BEV-based methods (Liu et al. 2023b; Liang et al. 2022; Yin et al. 2024; Jiao et al. 2023) align and fuse multi-modal features in a unified BEV representation.

However, the fundamental disparity in geometric and localization accuracy between LiDAR and cameras implies that indiscriminate fusion can lead to degraded performance (Huang et al. 2024). Specifically, cameras lack dense depth priors, and depth estimation methods (e.g., LSS (Phillion and Fidler 2020)) are prone to errors and overfitting (Li et al. 2023; Huang et al. 2024). In contrast, point clouds provide accurate localization information. The superior performance of LiDAR-based detectors (Zhang et al. 2024c; Liu et al. 2025) over image-based detectors (Li et al. 2024b; Liu et al. 2022) further demonstrates this inherent disparity. Thus, as shown in Fig 2 (a), naively aligning point cloud features with those derived from estimated depth constitutes a suboptimal fusion paradigm. Given the localization accuracy gap between point clouds and images, LiDAR information should be prioritized. However, this prioritization in

*Corresponding author

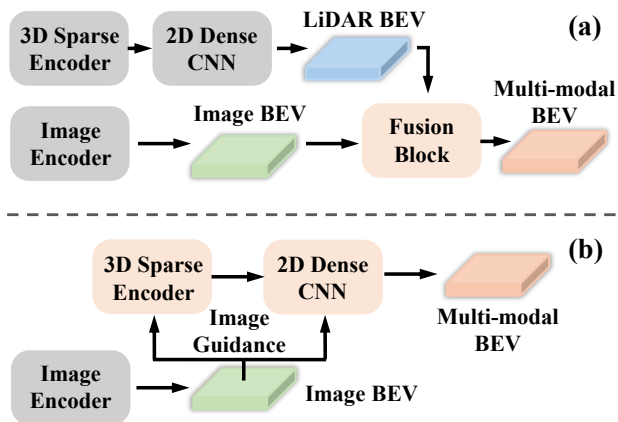


Figure 2: Comparison of (a) indiscriminate fusion and (b) our LiDAR-centric strategy.

multi-modal fusion is non-trivial. First, prioritizing LiDAR also highlights its own limitations. Point clouds are highly sparse and unevenly distributed, leading to incomplete geometric information and constrained representation capacity. Besides, in the LiDAR modality, object centers are usually empty, which weakens the representation power of point clouds. To address center feature missing (Fan et al. 2022), previous methods (Zhang et al. 2023; Zhou and Tuzel 2018; Yan, Mao, and Li 2018) often use feature diffusion (Fan et al. 2022) (e.g., stacked 2D CNN), which struggles to capture context information and distinguish noise due to insufficient semantic information. Second, when integrating image information, it is challenging to design effective fusion mechanisms to preserve the primacy of LiDAR data, thereby reducing the noise from image depth estimation.

To overcome the aforementioned limitation, in this paper we propose a novel LiDAR-centric fusion paradigm, termed BEVDilation, for multi-modal 3D object detection. As illustrated in Fig 2 (b), by formulating the image features as guidance, the fused multi-modal features can primarily derive their geometric and localization information from LiDAR point clouds. Specifically, image features provide semantic and contextual cues to drive sparse-to-dense feature diffusion in the LiDAR detector, rather than being naively aligned with point cloud features. This strategy reduces the impact of image depth estimation noise and improves fusion robustness. To address the inherent limitations of point clouds, we further propose two modules with it. Inspired by voxel generation (Liu et al. 2025; Sun et al. 2022; Zhang et al. 2024a) in LiDAR-based detectors, which do not rely on strict geometric information, we introduce the **Sparse Voxel Dilation Block (SVDB)** to densify foreground features and enhance features representation in highly sparse point clouds. It employs complementary multi-modal features to identify the foreground regions in BEV space and fills LiDAR’s empty foreground areas with new learnable voxels. This incorporates the image instance information into LiDAR and improves the completeness of point cloud features. To capture long-range context dependencies and refine point fea-

tures, we propose the **Semantic-Guided BEV Dilation Block (SBDB)** for LiDAR feature diffusion. In specific, following deformable convolution (Dai et al. 2017; Zhu et al. 2019b), we add offsets to the regular convolution sampling position to improve feature diffusion efficiency and adapt to geometric variation. Furthermore, SBDB uses modulation scalars to filter out noise features based on image semantic information. All deformations are conditioned on multi-modal BEV features but operate solely on LiDAR features. In this way, better LiDAR representations can be achieved under our LiDAR-centric design.

In summary, the major contributions of our work are:

- We propose BEVDilation, a LiDAR-centric framework for multi-modal 3D detection. BEVDilation formulates image features as guidance while prioritizing LiDAR information, achieving effective and robust fusion.
- To address the limitations of point clouds in the LiDAR-centric paradigm, we propose the Sparse Voxel Dilation Block (SVDB) and Semantic-Guided BEV Dilation Block (SBDB) to reduce sparsity and augment semantic information in BEVDilation.
- Extensive experiments on the nuScenes (Caesar et al. 2020) dataset show that our BEVDilation yields a new state-of-the-art for multi-modal 3D object detection while maintaining competitive computational efficiency.

Related Work

LiDAR-based 3D Object Detection. LiDAR sensors are the most widely adopted sensor in autonomous driving due to the precise geometric information provided by point clouds. Current LiDAR-based 3D object detection methods rely on two major representations: point clouds and voxels. Inspired by PointNet (Qi et al. 2017), point-based methods (Shi, Wang, and Li 2019; Qi et al. 2019) directly extract geometric features from local regions of raw point clouds. However, these methods suffer from low inference efficiency and insufficient context modeling due to the irregular and sparse nature of point clouds. Voxel-based methods (He et al. 2024; Zhang et al. 2024c; Liu et al. 2025; Zhang et al. 2024b; Mao et al. 2021; Zhao et al. 2024) transform raw point clouds into regular grids and well balance performance and efficiency. Voxel-based methods currently dominate the field of LiDAR-based 3D object detection. They often adopt a sparse backbone (Zhang et al. 2024c; Yan, Mao, and Li 2018) to extract features from sparse voxels in 3D space. These features are then converted to dense feature maps in BEV and processed by a 2D dense backbone (Yan, Mao, and Li 2018; Zhang et al. 2023) to enlarge the receptive field and refine features. However, LiDAR-based detectors relying solely on sparse and irregular point clouds exhibit degraded performance in challenging scenarios, such as long-range and geometrically similar object (e.g., pedestrians and vertical poles) detection. In this paper, we integrate additional image information to address the limitation of point cloud data by densifying sparse features and incorporating semantic guidance.

Multi-Modality 3D Object Detection. To achieve accurate and robust perceptions, existing approaches (Li et al.

2024a; Wang et al. 2023a; Huang et al. 2024; Yin et al. 2024; Yan et al. 2023; Liu et al. 2023b) explore capturing complementary information from LiDAR and cameras. Despite the substantial performance gap between camera-based (Liu et al. 2022, 2023a; Li et al. 2024b, 2023) and LiDAR-based detectors (Zhang et al. 2024c; Liu et al. 2025), the rich semantic information in images can further enhance LiDAR representations. Current multi-modal 3D detection approaches can be basically classified into three categories: input-based (Vora et al. 2020; Wang et al. 2021; Yin, Zhou, and Krähenbühl 2021), BEV-based (Liu et al. 2023b; Liang et al. 2022; Yin et al. 2024; Huang et al. 2024) and query-based (Yan et al. 2023; Li et al. 2024a; Yang et al. 2022). Pioneering input fusion methods (Vora et al. 2020; Yin, Zhou, and Krähenbühl 2021) decorate or densify LiDAR points with the semantic information from images, which is more sensitive to calibration errors. BEV-based fusion (Liu et al. 2023b; Liang et al. 2022; Huang et al. 2024) methods project LiDAR and image features into a unified BEV representation with estimated depth. Nevertheless, the uncertain depth of images introduces spatial misalignment and reduces the robustness of BEVFusion. Query-based methods (Li et al. 2024a; Yan et al. 2023; Cai et al. 2023) extract sparse instance-level representations and refine them with an attention operation. However, this strategy leaves several fundamental limitations of individual modalities unaddressed. In contrast, our proposed BEVDilation resolves LiDAR sparsity and semantic deficiency while alleviating spatial misalignment issues.

Method

In this section, we present BEVDilation, a LiDAR-centric 3D backbone that can be applied for multi-modal 3D detection. First, we introduce the overall architecture of BEVDilation. Then, we describe in detail the fundamental components of BEVDilation, including the Sparse Voxel Dilation Block and Semantic-Guided BEV Dilation Block.

Overall Architecture

An overview of our proposed BEVDilation is shown in Fig 3. As in previous works (Bai et al. 2022; Li et al. 2024a; Yin et al. 2024), BEVDilation takes LiDAR point clouds and corresponding multi-view images as input, which are fed into two individual backbones to extract multi-modal BEV features. Unlike previous multi-modal methods that indiscriminately fuse modality-specific features (Liu et al. 2023b; Liang et al. 2022; Li et al. 2024a; Yan et al. 2023), BEVDilation introduces a LiDAR-centric framework. It uses image features as guidance to mitigate the inherent limitations of point clouds, including sparsity, irregular distributions, and semantic ambiguities. Specifically, a Sparse Voxel Dilation Block working on the sparse voxel features is proposed, which increases voxel density in foreground instances and completes features in LiDAR-occluded areas. Image and LiDAR features are fused to generate a BEV foreground mask, which subsequently determines the regions for voxel dilation. To mitigate the semantic ambiguity and improve geometric accuracy, BEVDilation introduces a Semantic-Guided BEV Dilation Block that integrates multi-modal

guidance in feature diffusion, transforming sparse LiDAR BEV features into geometrically consistent and semantically enriched dense LiDAR representations.

Sparse Voxel Dilation Block

Given the challenges of high sparsity of point clouds and information degradation in LiDAR modality, we propose a sparse voxel dilation strategy that leverages the autoregressive capacity and global receptive field of Mamba (Liu et al. 2025) to address these issues.

Image Encoder. For each image I , a shared image backbone is used to extract multi-scale feature maps that are subsequently fused into a single-scale representation $\mathbf{F}_I \in \mathbb{R}^{C_i \times H \times W}$ by FPN (Lin et al. 2017), where H and W denote the image resolution and C_i represents the number of feature channels. Then, the multi-view image features are transformed from the perspective view to BEV representation $\mathbf{F}_I^{bev} \in \mathbb{R}^{C_{ib} \times 1 \times Y \times X}$ using the well-established view transformation (VT.) algorithm Lift-Splat-Shot (LSS) (Phillion and Fidler 2020), where X and Y denote the size of the BEV feature map and C_{ib} is the number of image BEV feature channels.

Voxel Sparse Encoder. For a fair comparison, we use the popular sparse backbone VoxelNet (Zhou and Tuzel 2018) to extract sparse voxel representations. The irregular point clouds P are first transformed into sparse voxels by a voxel feature encoding. Then, we extract sparse voxel features $\mathbf{F}_P \in \mathbb{R}^{N \times C_p}$ along with their spatial coordinates $\mathcal{V} \in \mathbb{R}^{N \times 3}$ using a series of sparse convolutional blocks. N indicates the number of voxels and C_p represents the number of voxel feature channels.

Sparse Voxel Dilation Block. As shown in Fig 3, to enhance the accuracy of dilation regions, we use multi-modal information to predict the foreground mask. The prediction head takes the image BEV features \mathbf{F}_I^{bev} and LiDAR sparse BEV features \mathbf{F}_P^{bev} as inputs. Specifically, the foreground mask predictor is formulated as follows:

$$\mathbf{P}_{fg} = \sigma(f_{conv}((\mathbf{F}_I^{bev}; \mathbf{F}_P^{bev}))), \quad (1)$$

where f_{conv} denotes two-layer convolutions, σ represents the sigmoid activation, and $(;)$ indicates concat operation. The binary foreground mask $\mathcal{M} \in \{0, 1\}^{Y \times X}$ is obtained:

$$\mathcal{M}_{i,j} = \mathbb{I}[\mathbf{P}_{fg}(i, j) > \tau], \quad (2)$$

where (i, j) denotes the indices within the BEV grid and τ is the threshold. Regions corresponding to positive mask values are designated for padding new voxels. During training, we treat BEV grid cells within object bounding boxes as the ground truth foreground to supervise \mathcal{M} . The dilated voxels can effectively compensate for LiDAR-occluded regions and densify the foreground instances. Besides, the dilated features in previous methods (Sun et al. 2022; Zhang et al. 2024a; Liu et al. 2025) are often assigned using K-NN approaches or zero initialization. These handcrafted strategies tend to yield sub-optimal results due to their limited adaptability to varying scene contexts. Therefore, as shown in the Fig 4, all newly dilated features are generated as learnable embeddings \mathbf{F}_{ebd} . Inspired by the recent success of Mamba

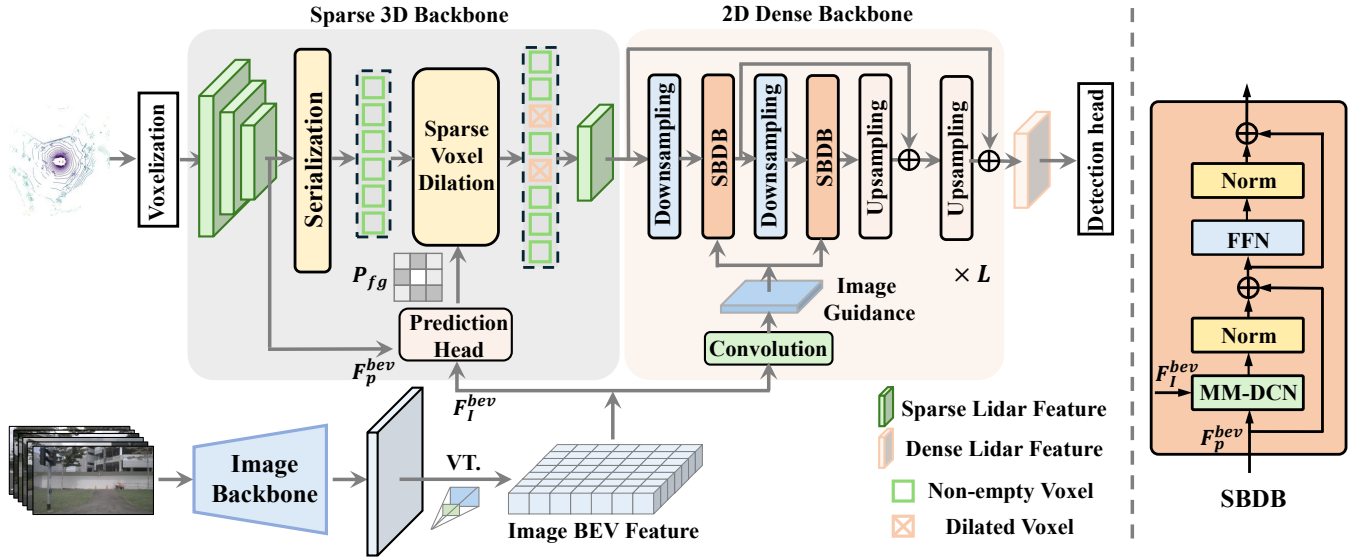


Figure 3: The overall architecture of our proposed **BEVDilation**. Given the point clouds and multi-view images, we take two individual backbones to extract multi-modal BEV features. For the LiDAR branch, we enhance the LiDAR BEV features with our proposed Sparse Voxel Dilation Block and Semantic-Guided BEV Dilation Block.




Figure 4: An illustration of SVDB. The newly padded and original voxels are merged with global receptive fields.

in LiDAR detectors (Zhang et al. 2024c; Liu et al. 2025), we utilize its auto-regressive property to refine voxel features with global receptive fields. In specific, we merge and sort the newly dilated and original voxels using Hilbert serialization (Hilbert and Hilbert 1935), and then refine the features with a single group-free Mamba layer (Zhang et al. 2024c) as following:

$$\mathbf{F}_P^{bev} = \text{Mamba}(\text{HS}([\mathbf{F}_P^{bev}; \mathbf{F}_{ebd}])), \quad (3)$$

where HS denotes the operation that sorts the voxel sequence using a Hilbert curve (Zhang et al. 2024c). Overall, SVDB effectively tackles point cloud sparsity and LiDAR blind spots while preserving LiDAR’s geometric primacy.

Semantic-Guided BEV Dilation Block

Though SVDB mitigates foreground sparsity, the BEV feature remains sparse, thereby demanding an efficient feature

diffusion (Fan et al. 2022). The feature diffusion in existing LiDAR detectors exhibits two major limitations: semantic ambiguity and the rigid geometric constraints in standard convolutions. Therefore, we introduce the Semantic-Guided BEV Dilated Block to build a semantic-guided feature diffusion and enhance LiDAR’s ability to model geometric variations through multi-modal deformation.

As shown in Fig. 3, the SBDB is designed with a multi-modal deformable convolution (MM-DCN), a feed-forward network (FFN), layer normalization, and residual connections. To enlarge the receptive field and improve computational efficiency, it operates on downsampled BEV features. In MM-DCN, the deformations, specifically the offset and modulation scalar, are conditioned on the multi-modal BEV information. Moreover, to reduce spatial misalignment from uncertain depth estimation, while the deformations (offsets and modulation scalars) in MM-DCN are conditioned on multi-modal BEV features, the deformable convolutions operate solely on LiDAR features. Specifically, image BEV features are first downsampled using a lightweight convolutional encoder to match the resolution of LiDAR BEV features. These downsampled image BEV features are shared across different stages. Subsequently, the SBDB takes these processed image BEV features \mathbf{F}_I^{bev} and the corresponding LiDAR BEV features \mathbf{F}_P^{bev} as input, defined as:

$$\begin{aligned} \tilde{\mathbf{F}}_P^{bev} &= \text{LN}(\text{MM-DCN}([\mathbf{F}_P^{bev}, \mathbf{F}_I^{bev}])) + \mathbf{F}_P^{bev}, \\ \mathbf{F}_P^{bev} &= \text{LN}(\text{MLP}(\tilde{\mathbf{F}}_P^{bev})) + \tilde{\mathbf{F}}_P^{bev}. \end{aligned} \quad (4)$$

Here, $\text{LN}(\cdot)$ stands for Layer Normalization. $\text{MM-DCN}(\cdot)$ represents the deformable convolutions used in the first residual branch, where deformations (offsets Δp_k and modulation scalars \mathbf{m}_k) are conditioned on multi-modal features.

The convolution is applied solely to the LiDAR features \mathbf{F}_P^{bev} . Formally, the **MM-DCN** is denoted as:

$$\hat{\mathbf{F}}_P^{bev}(p_0) = \sum_{k=1}^M \mathbf{w}_k \mathbf{m}_k \cdot \mathbf{F}_P^{bev}(p_0 + p_k + \Delta p_k), \quad (5)$$

where $\mathbf{F}_P^{bev}(p)$ denotes the features at location p from LiDAR BEV features, and \mathbf{w}_k are the standard convolutional kernel weights for the k -th location. M represents the total number of sampling points. For clarity, groups in DCN (Wang et al. 2023b) are omitted. By adapting multi-modal deformable convolutions with the LiDAR-centric strategy, SBDB achieves precise and robust feature diffusion, enhancing the LiDAR representations.

Integration with Detection

With the proposed SVDB and SBDB, we build BEVDilation, a LiDAR-centric multi-modal fusion paradigm. The architecture of BEVDilation is illustrated in Fig 3. The SVDB is placed subsequent to the VoxelNet sparse backbone. Our 2D dense backbone, comprising L stages, directly replaces the dense backbone of the baseline. For the detection head and loss function, we adopt the same settings as DAL (Huang et al. 2024), adding a focal loss to supervise the binary foreground mask.

Experiments

Datasets and Evaluation Metrics

Following previous works (Li et al. 2024a; Yan et al. 2023; Bai et al. 2022), we evaluate our method on the nuScenes dataset (Caesar et al. 2020). This dataset comprises 750 training scenes, 150 validation scenes, and 150 testing scenes. It provides point clouds from a 32-beam LiDAR and 6 images from multi-view cameras. For 3D object detection, nuScenes employs the mean Average Precision (mAP), nuScenes detection scores (NDS), mean Average Translation Error (mATE), mean Average Scale Error (mASE), mean Average Orientation Error (mAOE), mean Average Velocity Error (mAVE) and mean Average Attribute Error (mAEE) to measure model performance.

Implementation Detail

Our method is implemented based on the open-source framework MMDetection3d (Contributors 2020). For the camera branch, we use ResNet50 (He et al. 2016) as the image backbone and initialized it with weights pre-trained on nuImage (Caesar et al. 2020), following current approaches (Yang et al. 2022; Yin et al. 2024; Li et al. 2024a). The input image resolution is 800×448 . For the LiDAR branch, we follow DAL (Huang et al. 2024) and use the voxel size $(0.075m, 0.075m, 0.2m)$, with the point cloud range defined as $[-54m, 54m]$ along the XY-axes and $[-5m, 3m]$ for the Z-axis. The initial BEV feature map is of size 180×180 . We adopt the sparse 3D backbone from VoxelNet (Zhou and Tuzel 2018), consistent with previous works (Li et al. 2024a; Huang et al. 2024; Yin et al. 2024). For the 2D dense CNN, we stack eight SBDB blocks distributed over four stages. The τ in SVDB is set to 0.4. The

number of groups in the multi-modal deformable convolutions (Wang et al. 2023b) is set to 16.

Training and Inference. Following previous training schemes (Yin et al. 2024; Huang et al. 2024), BEVDilation is trained in an end-to-end manner with one stage. We optimize the model using AdamW optimizer with weight decay 0.01, one-cycle learning rate policy, max learning rate 0.0001, and batch size 24 for 10 epochs. The class-balanced sampling strategy from CBGS (Zhu et al. 2019a) and multi-modal data augmentation from DAL (Huang et al. 2024) are adopted during training. We use the CenterPoint (Yin, Zhou, and Krahenbuhl 2021) detection head. All the models are trained on 8 RTX A6000 GPUs. During evaluation, image features are only used for searching proposal candidates, while regression relies entirely on LiDAR features as in (Huang et al. 2024).

Comparison with State-of-the-art Methods

As shown in Table 1, we compare BEVDilation with state-of-the-art (SOTA) methods on both the validation and test set of the nuScenes dataset without any test-time augmentations or model ensembles. Our proposed BEVDilation achieves impressive results with 75.0 NDS and 73.0 mAP on the validation set, which is at least +1.0 and +1.5 higher than the SOTA method, including those with more powerful backbones. Our method also exhibits the best mAP and NDS on the test set. Additionally, as shown in Fig 1, BEVDilation strikes an effective balance between speed and accuracy. It outperforms DAL +1.0 NDS in detection accuracy on the nuScenes validation while achieving comparable speed. Some methods, like BEVFusion, are faster than BEVDilation; however, their accuracy is substantially lower. A more detailed performance analysis, as well as additional experimental results, is provided in the supplementary material.

Following common model scaling practices (Huang et al. 2024; Li et al. 2024a), we enhance BEVDilation with SwinTiny (Liu et al. 2021) and a higher input resolution. Table 2 presents a comparison between the scaled BEVDilation and other scaled SOTA methods. Despite Mamba-fusion employing more advanced LiDAR (LION-Mamba) and image (VMamba) backbones, BEVDilation nonetheless demonstrates significant advantages, surpassing it by +0.3 NDS and +0.7 mAP while achieving faster inference speeds. BEVDilation also maintains superior computational efficiency compared to most existing approaches.

Ablation Studies

To better investigate the effectiveness of BEVDilation, we conduct a set of ablation studies by using the nuScenes validation set in this subsection. All the models are trained end-to-end for 10 epochs.

Effectiveness of each component. To more clearly demonstrate the effectiveness of the SVDB and SBDB in BEVDilation, we conduct experiments by adding each of them to a multi-modal baseline. The baseline takes VoxelNet (Zhou and Tuzel 2018) and SECOND (Yan, Mao, and Li 2018) dense backbone as the LiDAR backbone and ResNet50 (He et al. 2016) as the image backbone. The detection head follows the design in DAL (Huang et al. 2024). As evidenced

Methods	Modality	LiDAR Backbone	Camera Backbone	<i>validation set</i>		<i>test set</i>		FPS
				NDS	mAP	NDS	mAP	
PETrv2 (Liu et al. 2023a)	C	-	ResNet-101	52.4	42.1	55.3	45.6	-
BEVFormer (Li et al. 2024b)	C	-	ResNet-101	51.7	41.6	53.5	44.5	-
SECOND (Yan, Mao, and Li 2018)	L	VoxelNet	-	63.0	52.6	63.3	52.8	-
TransFusion-L (Bai et al. 2022)	L	VoxelNet	-	70.1	65.1	70.2	65.5	-
FocalFormer3D-L (Chen et al. 2023b)	L	VoxelNet	-	70.9	66.4	72.6	68.7	-
MVP (Yin, Zhou, and Krähenbühl 2021)	L+C	VoxelNet	CenterNet	70.0	66.1	70.5	66.4	-
GraphAlign (Song et al. 2023)	L+C	VoxelNet	DeepLabv3	-	-	70.6	66.5	-
FUTR3D (Chen et al. 2023a)	L+C	VoxelNet	ResNet-101	68.3	64.5	-	-	-
UVTR (Li et al. 2022)	L+C	VoxelNet	ResNet-101	70.2	65.4	71.1	67.1	2.14
TransFusion (Bai et al. 2022)	L+C	VoxelNet	ResNet-50	71.3	67.5	71.6	68.9	7.88
AutoAlignV2 (Chen et al. 2022)	L+C	VoxelNet	CSPNet	71.2	67.1	72.4	68.4	-
BEVFusion (Liang et al. 2022)	L+C	VoxelNet	Dual-Swin-T	72.1	69.6	73.3	71.3	2.11
BEVFusion (Liu et al. 2023b)	L+C	VoxelNet	Swin-T	71.4	68.5	72.9	70.2	11.60
DeepInteraction (Yang et al. 2022)	L+C	VoxelNet	ResNet-50	72.6	69.9	73.4	70.8	2.25
CMT (Yan et al. 2023)	L+C	VoxelNet	VoV-99	71.9	69.4	73.0	70.4	4.01
FocalFormer3D (Chen et al. 2023b)	L+C	VoxelNet	ResNet-50	73.1	70.5	73.9	71.6	-
MSMDFusion (Jiao et al. 2023)	L+C	VoxelNet	ResNet-50	-	-	74.0	71.5	2.57
SparseFusion (Xie et al. 2023)	L+C	VoxelNet	ResNet-50	72.8	70.4	73.8	72.0	5.57
UniTR (Wang et al. 2023a)	L+C	DSVT	DSVT	73.3	70.5	74.5	70.9	3.33 [†]
ObjectFusion (Cai et al. 2023)	L+C	VoxelNet	Swin-T	72.3	69.8	73.3	71.0	-
GraphBEV (Song et al. 2024)	L+C	VoxelNet	Swin-T	72.9	70.1	73.6	71.7	3.12
DAL (Huang et al. 2024)	L+C	VoxelNet	ResNet-50	74.0	71.5	74.8	72.0	7.39
BEVDilation (ours)	L+C	VoxelNet	ResNet-50	75.0	73.0	75.4	73.1	7.08

Table 1: Comparison with existing methods on nuScenes **validation** and **test** set. ‘L’ denotes the LiDAR and ‘C’ indicates the camera. Symbol † means UniTR without cache acceleration. The inference speeds are evaluated on an NVIDIA A6000 GPU.

Methods	Image Backbone	Input Resolution	NDS	mAP	FPS
CMT	VoVNet-99	1600×640	72.9	70.3	4.01
SparseFusion	Swin-T	800×448	73.1	71.0	5.22
DAL	ResNet-50	1056×384	74.0	71.5	7.39
IS-Fusion	Swin-T	1056×384	74.0	72.8	2.94
MambaFusion	VMamba	704×256	75.0	72.7	3.84
BEVDilation	Swin-T	1056×384	75.3	73.4	4.13

Table 2: Results on nuScenes **validation** set with stronger image backbone and higher input resolution.

in Table 3, both SVDB (+1.2 mAP) and SBDB (+1.9 mAP) can significantly improve the accuracy over the baseline, which validates the feasibility of our LiDAR-centric strategy. Combining SVDB and SBDB further improves performance, suggesting that densifying sparse BEV foreground enhances feature diffusion, leading to better completion of LiDAR occluded regions.

Robustness of our LiDAR-centric paradigm. To validate the robustness and effectiveness of our LiDAR-centric paradigm against indiscriminate fusion, we conducted experiments between our BEVDilation and BEVFusion (Liu et al. 2023b). We introduce three types of degradation dur-

Baseline-LC	SVDB	SBDB	mAP	NDS	FPS
✓			70.6	73.3	9.12
✓	✓		71.8	74.0	8.62
✓		✓	72.5	74.8	7.51
✓	✓	✓	73.0	75.0	7.08

Table 3: Ablations on the nuScenes validation split. Baseline-LC denotes our multi-modal baseline method.

ing inference: depth random noise, depth one-hot noise (Li et al. 2023), and spatial misalignment (Dong et al. 2023). As shown in Table 4, BEVDilation exhibits lower performance drops than BEVFusion under all degradations, demonstrating its superior robustness. Furthermore, under noisy conditions, BEVDilation consistently achieves higher mATE than standard BEVFusion, indicating that prioritizing LiDAR information helps maintain accurate geometric perception.

Effectiveness of multimodal features. To validate the efficacy of image semantic information in our LiDAR-centric framework, we perform experiments on some potential alternatives to SVDB and SBDB. As shown in Table 5, removing image guidance for SVDB foreground mask prediction leads to substantial performance drops, -0.5 mAP and -0.4 NDS, demonstrating the critical role of images in re-

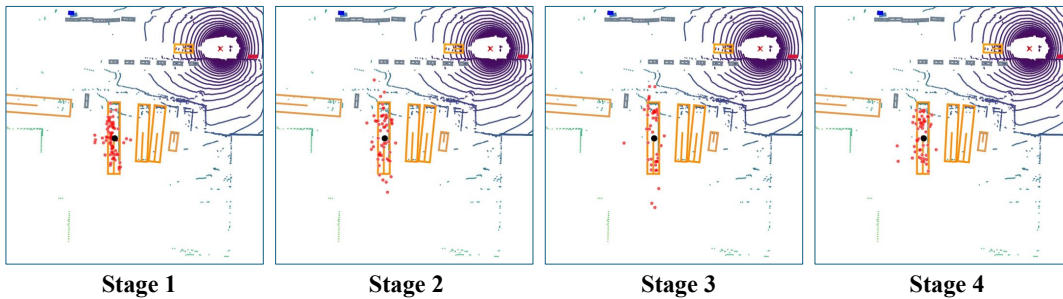


Figure 5: Visualization of sampling locations of **SBDB** at different stages. The black dot indicates the object center, and these red dots denote the sampling locations of this center query point in SBDB.

Methods	BEVDilation			BEVFusion		
	mAP \uparrow	NDS \uparrow	mATE \downarrow	mAP \uparrow	NDS \uparrow	mATE \downarrow
Baseline	73.0	75.0	26.9	68.5	71.4	28.7
One-hot Noise	70.3 \downarrow 2.7	73.6 \downarrow 1.4	27.2 \uparrow 0.3	63.1 \downarrow 5.4	68.4 \downarrow 3.0	29.4 \uparrow 0.7
Random Noise	70.0 \downarrow 3.0	73.4 \downarrow 1.6	27.3 \uparrow 0.4	62.3 \downarrow 6.2	68.0 \downarrow 3.4	29.6 \uparrow 0.9
Spatial Misalignment	70.4 \downarrow 2.6	73.6 \downarrow 1.4	27.1 \uparrow 0.2	61.2 \downarrow 7.3	67.3 \downarrow 4.1	29.8 \uparrow 1.1

Table 4: Robustness to image depth estimation noise on nuScenes validation set.

Methods	Modality	NDS	mAP
Baseline	LC	73.3	70.6
+ SVDB (w/o image)	LC	73.6	71.3
+ SVDB (zero init)	LC	72.7	70.4
+ SVDB (ours)	LC	74.0	71.8
+ SBDB (w/o image)	LC	74.2	71.7
+ SBDB (fusion)	LC	74.4	71.8
+ SBDB (ours)	LC	74.8	72.5

Table 5: Ablation studies for our proposed components on the nuScenes **validation** set. ‘w/o image’ denotes the removal of image priors in the current block. ‘fusion’ means using DCN on multimodal features in SBDB.

covering LiDAR-occluded regions. Furthermore, replacing our learnable embedding for padded empty locations with a fixed zero value (Liu et al. 2025; Sun et al. 2022) causes a significant performance drop of 1.2 mAP. This indicates that handcrafted features fail to adapt to scene-specific contexts and result in suboptimal feature representations. For SBDB, removing image-conditioned deformations in DCN leads to a significant decline in performance, which highlights the necessity of semantic and context guidance from images. We also compare our LiDAR-centric fusion within SBDB against a direct fusion baseline that uses standard deformable convolutions on concatenated multi-modal BEV features. The decreased performance confirms that our approach can effectively reduce the impact of noise from image depth estimation.

Visualization of BEVDilation. To validate the effectiveness of our semantic-guided BEV dilation, we visualize

its sampling locations across different stages. We project the multi-modal conditioned sampling locations from all groups (Wang et al. 2023b) onto a unified BEV map and filter out those locations with modulation scalars below 0.01. As shown in Fig 5, compared to rigid grid patterns in standard convolutions, SBDB dynamically adapts its sampling locations to geometric variations. For example, as shown in stage 3, SBDB can directly extract features at object boundaries, reducing the need for deep CNN stacks that struggle with off-grid points. This clearly demonstrates SBDB’s ability to handle the irregular distributions in point clouds. Furthermore, the progressively expanding receptive fields in the deeper stages highlight SBDB’s capability to aggregate long-range contextual information, which is important to address the Center Feature Missing issue. We further include visualizations of the predicted foreground occupancy from SVDB in the supplementary material.

Conclusion

In this paper, we proposed BEVDilation, a LiDAR-centric backbone for multi-modal 3D object detection. We first analyzed the degradation from current naive fusion. Thus, we proposed a LiDAR-centric paradigm to reduce the influence of image depth estimation noise in fusion. In specific, the SVDB block is used to densify foreground voxel features and fill in the LiDAR-occluded regions with image priors. Then, the SBDB equipped with multi-modal conditioned deformations diffuses sparse voxel features into dense BEV features with geometry-aware adaptability and enhanced semantic context. Experiments demonstrated that BEVDilation achieved state-of-the-art results on the nuScenes dataset. BEVDilation presented a new LiDAR-centric solution for multi-modal 3D detection.

References

- Bai, X.; Hu, Z.; Zhu, X.; Huang, Q.; Chen, Y.; Fu, H.; and Tai, C.-L. 2022. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1090–1099.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2020. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11621–11631.
- Cai, Q.; Pan, Y.; Yao, T.; Ngo, C.-W.; and Mei, T. 2023. Objectfusion: Multi-modal 3d object detection with object-centric fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18067–18076.
- Chen, X.; Zhang, T.; Wang, Y.; Wang, Y.; and Zhao, H. 2023a. Futr3d: A unified sensor fusion framework for 3d detection. In *proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 172–181.
- Chen, Y.; Yu, Z.; Chen, Y.; Lan, S.; Anandkumar, A.; Jia, J.; and Alvarez, J. M. 2023b. Focalformer3d: focusing on hard instance for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8394–8405.
- Chen, Z.; Li, Z.; Zhang, S.; Fang, L.; Jiang, Q.; and Zhao, F. 2022. Deformable feature aggregation for dynamic multi-modal 3D object detection. In *European conference on computer vision*, 628–644. Springer.
- Contributors, M. 2020. MMDetection3D: OpenMMLab next-generation platform for general 3D object detection. <https://github.com/open-mmlab/mmdetection3d>.
- Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; and Wei, Y. 2017. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 764–773.
- Dong, Y.; Kang, C.; Zhang, J.; Zhu, Z.; Wang, Y.; Yang, X.; Su, H.; Wei, X.; and Zhu, J. 2023. Benchmarking robustness of 3d object detection to common corruptions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1022–1032.
- Fan, L.; Wang, F.; Wang, N.; and Zhang, Z.-X. 2022. Fully sparse 3d object detection. *Advances in Neural Information Processing Systems*, 35: 351–363.
- He, C.; Li, R.; Zhang, G.; and Zhang, L. 2024. ScatterFormer: Efficient Voxel Transformer with Scattered Linear Attention. *arXiv preprint arXiv:2401.00912*.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Hilbert, D.; and Hilbert, D. 1935. Über die stetige Abbildung einer Linie auf ein Flächenstück. *Dritter Band: Analysis- Grundlagen der Mathematik- Physik Verschiedenes: Nebst Einer Lebensgeschichte*, 1–2.
- Huang, J.; Ye, Y.; Liang, Z.; Shan, Y.; and Du, D. 2024. Detecting as labeling: Rethinking lidar-camera fusion in 3d object detection. In *European Conference on Computer Vision*, 439–455. Springer.
- Jiao, Y.; Jie, Z.; Chen, S.; Chen, J.; Ma, L.; and Jiang, Y.-G. 2023. Msmdfusion: Fusing lidar and camera at multiple scales with multi-depth seeds for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 21643–21652.
- Li, Y.; Chen, Y.; Qi, X.; Li, Z.; Sun, J.; and Jia, J. 2022. Unifying voxel-based representation with transformer for 3d object detection. *Advances in Neural Information Processing Systems*, 35: 18442–18455.
- Li, Y.; Ge, Z.; Yu, G.; Yang, J.; Wang, Z.; Shi, Y.; Sun, J.; and Li, Z. 2023. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 1477–1485.
- Li, Y.; Li, H.; Huang, Z.; Chang, H.; and Wang, N. 2024a. SparseFusion: Efficient Sparse Multi-Modal Fusion Framework for Long-Range 3D Perception. *arXiv preprint arXiv:2403.10036*.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Yu, Q.; and Dai, J. 2024b. Bevformer: learning bird’s-eye-view representation from lidar-camera via spatiotemporal transformers. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Liang, T.; Xie, H.; Yu, K.; Xia, Z.; Lin, Z.; Wang, Y.; Tang, T.; Wang, B.; and Tang, Z. 2022. Bevfusion: A simple and robust lidar-camera fusion framework. *Advances in Neural Information Processing Systems*, 35: 10421–10434.
- Lin, T.-Y.; Dollár, P.; Girshick, R.; He, K.; Hariharan, B.; and Belongie, S. 2017. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2117–2125.
- Liu, Y.; Wang, T.; Zhang, X.; and Sun, J. 2022. Petr: Position embedding transformation for multi-view 3d object detection. In *European conference on computer vision*, 531–548. Springer.
- Liu, Y.; Yan, J.; Jia, F.; Li, S.; Gao, A.; Wang, T.; and Zhang, X. 2023a. Petrv2: A unified framework for 3d perception from multi-camera images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3262–3272.
- Liu, Z.; Hou, J.; Wang, X.; Ye, X.; Wang, J.; Zhao, H.; and Bai, X. 2025. Lion: Linear group rnn for 3d object detection in point clouds. *Advances in Neural Information Processing Systems*, 37: 13601–13626.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, 10012–10022.
- Liu, Z.; Tang, H.; Amini, A.; Yang, X.; Mao, H.; Rus, D. L.; and Han, S. 2023b. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. In *2023 IEEE international conference on robotics and automation (ICRA)*, 2774–2781. IEEE.

- Mao, J.; Xue, Y.; Niu, M.; Bai, H.; Feng, J.; Liang, X.; Xu, H.; and Xu, C. 2021. Voxel transformer for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3164–3173.
- Philion, J.; and Fidler, S. 2020. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XIV 16*, 194–210. Springer.
- Qi, C. R.; Litany, O.; He, K.; and Guibas, L. J. 2019. Deep hough voting for 3d object detection in point clouds. In *proceedings of the IEEE/CVF International Conference on Computer Vision*, 9277–9286.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 652–660.
- Shi, S.; Wang, X.; and Li, H. 2019. Pointcnn: 3d object proposal generation and detection from point cloud. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 770–779.
- Song, Z.; Wei, H.; Bai, L.; Yang, L.; and Jia, C. 2023. GraphAlign: Enhancing accurate feature alignment by graph matching for multi-modal 3D object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3358–3369.
- Song, Z.; Yang, L.; Xu, S.; Liu, L.; Xu, D.; Jia, C.; Jia, F.; and Wang, L. 2024. Graphbev: Towards robust bev feature alignment for multi-modal 3d object detection. In *European Conference on Computer Vision*, 347–366. Springer.
- Sun, P.; Tan, M.; Wang, W.; Liu, C.; Xia, F.; Leng, Z.; and Angelov, D. 2022. Swformer: Sparse window transformer for 3d object detection in point clouds. In *European Conference on Computer Vision*, 426–442. Springer.
- Vora, S.; Lang, A. H.; Helou, B.; and Beijbom, O. 2020. Pointpainting: Sequential fusion for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4604–4612.
- Wang, C.; Ma, C.; Zhu, M.; and Yang, X. 2021. Pointaugmenting: Cross-modal augmentation for 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11794–11803.
- Wang, H.; Tang, H.; Shi, S.; Li, A.; Li, Z.; Schiele, B.; and Wang, L. 2023a. Unitr: A unified and efficient multi-modal transformer for bird’s-eye-view representation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6792–6802.
- Wang, W.; Dai, J.; Chen, Z.; Huang, Z.; Li, Z.; Zhu, X.; Hu, X.; Lu, T.; Lu, L.; Li, H.; et al. 2023b. Internimage: Exploring large-scale vision foundation models with deformable convolutions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14408–14419.
- Xie, Y.; Xu, C.; Rakotosaona, M.-J.; Rim, P.; Tombari, F.; Keutzer, K.; Tomizuka, M.; and Zhan, W. 2023. Sparsefusion: Fusing multi-modal sparse representations for multi-sensor 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 17591–17602.
- Yan, J.; Liu, Y.; Sun, J.; Jia, F.; Li, S.; Wang, T.; and Zhang, X. 2023. Cross modal transformer: Towards fast and robust 3d object detection. In *Proceedings of the IEEE/CVF international conference on computer vision*, 18268–18278.
- Yan, Y.; Mao, Y.; and Li, B. 2018. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10): 3337.
- Yang, Z.; Chen, J.; Miao, Z.; Li, W.; Zhu, X.; and Zhang, L. 2022. Deepinteraction: 3d object detection via modality interaction. *Advances in Neural Information Processing Systems*, 35: 1992–2005.
- Yin, J.; Shen, J.; Chen, R.; Li, W.; Yang, R.; Frossard, P.; and Wang, W. 2024. Is-fusion: Instance-scene collaborative fusion for multimodal 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14905–14915.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 11784–11793.
- Yin, T.; Zhou, X.; and Krähenbühl, P. 2021. Multimodal virtual point 3d detection. *Advances in Neural Information Processing Systems*, 34: 16494–16507.
- Zhang, G.; Chen, J.; Gao, G.; Li, J.; Liu, S.; and Hu, X. 2024a. Safdnet: A simple and effective network for fully sparse 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14477–14486.
- Zhang, G.; Fan, J.; Chen, L.; Zhang, Z.; Lei, Z.; and Zhang, L. 2024b. General Geometry-aware Weakly Supervised 3D Object Detection. *arXiv preprint arXiv:2407.13748*.
- Zhang, G.; Fan, L.; He, C.; Lei, Z.; ZHANG, Z.-X.; and Zhang, L. 2024c. Voxel mamba: Group-free state space models for point cloud based 3d object detection. *Advances in Neural Information Processing Systems*, 37: 81489–81509.
- Zhang, G.; Junnan, C.; Gao, G.; Li, J.; and Hu, X. 2023. Hednet: A hierarchical encoder-decoder network for 3d object detection in point clouds. *Advances in Neural Information Processing Systems*, 36: 53076–53089.
- Zhao, H.; Zhang, Q.; Zhao, S.; Chen, Z.; Zhang, J.; and Tao, D. 2024. Simdistill: Simulated multi-modal distillation for bev 3d object detection. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, 7460–7468.
- Zhou, Y.; and Tuzel, O. 2018. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4490–4499.
- Zhu, B.; Jiang, Z.; Zhou, X.; Li, Z.; and Yu, G. 2019a. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint arXiv:1908.09492*.
- Zhu, X.; Hu, H.; Lin, S.; and Dai, J. 2019b. Deformable convnets v2: More deformable, better results. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9308–9316.