

Height-Fidelity Dense Global Fusion for Multi-modal 3D Object Detection

Hanshi Wang^{1,2,3,5*}, Jin Gao^{1,2,5†}, Weiming Hu^{1,2,5,6}, Zhipeng Zhang^{3,4†}

¹State Key Laboratory of Multimodal Artificial Intelligence Systems (MAIS), CASIA

²School of Artificial Intelligence, University of Chinese Academy of Sciences

³School of Artificial Intelligence, Shanghai Jiao Tong University ⁴Anyverse Intelligence

⁵Beijing Key Laboratory of Super Intelligent Security of Multi-Modal Information

⁶School of Information Science and Technology, ShanghaiTech University

{hanshi.wang.cv, zhipeng.zhang.cv}@outlook.com, {wmhu, jin.gao}@nlpr.ia.ac.cn

Abstract

We present the first work demonstrating that a **pure Mamba block** can achieve **efficient Dense Global Fusion**, meanwhile guaranteeing top performance for camera-LiDAR multi-modal 3D object detection. Our motivation stems from the observation that existing fusion strategies are constrained by their inability to simultaneously achieve efficiency, long-range modeling, and retaining complete scene information. Inspired by recent advances in state-space models (SSMs) [8] and linear attention [35, 43], we leverage their linear complexity and long-range modeling capabilities to address these challenges. However, this is non-trivial since our experiments reveal that simply adopting efficient linear-complexity methods does not necessarily yield improvements and may even degrade performance. We attribute this degradation to the loss of height information during multi-modal alignment, leading to deviations in sequence order. To resolve this, we propose height-fidelity LiDAR encoding that preserves precise height information through voxel compression in continuous space, thereby enhancing camera-LiDAR alignment. Subsequently, we introduce the Hybrid Mamba Block, which leverages the enriched height-informed features to conduct local and global contextual learning. By integrating these components, our method achieves state-of-the-art performance with the top-tire NDS score of 75.0 on the nuScenes [2] validation benchmark, even surpassing methods that utilize high-resolution inputs. Meanwhile, our method maintains efficiency, achieving faster inference speed than most recent state-of-the-art methods. Code is available at <https://github.com/AutoLab-SAI-SJTU/MambaFusion>

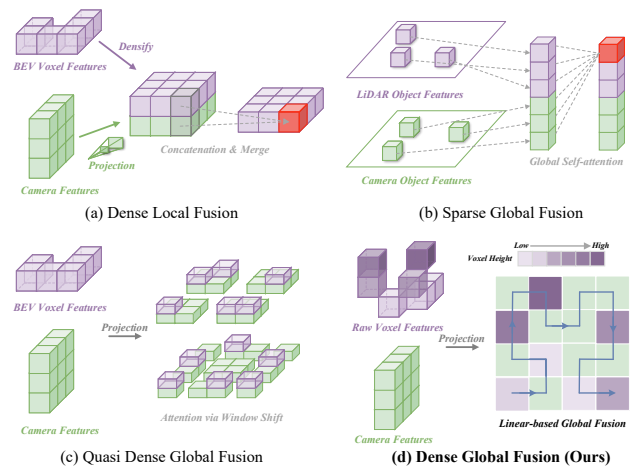


Figure 1. Illustration of the LiDAR and camera fusion framework: a) Dense Local Fusion, b) Sparse Global Fusion, c) Quasi Dense Global Fusion, d) **Dense Global Fusion (Ours)**.

1. Introduction

3D object detection is a foundational task in autonomous driving, enabling vehicles to accurately interpret their surroundings. As the complexity of driving environments escalates, the demand for robust detection methods becomes increasingly critical. LiDAR and camera are two commonly employed sensors in autonomous vehicles, each possessing distinct advantages. Specifically, LiDAR excels in delivering precise spatial and depth information, while cameras provide rich visual textures and color details. Their complementary nature offers a valuable opportunity for multi-sensor fusion, which can enhance detection accuracy and robustness across diverse and challenging scenarios.

Fig. 1 (a), (b), and (c) summarize three types of camera-LiDAR fusion strategies present in existing frameworks, categorized by the range of context utilized in representation learning. More specifically, Dense Local Fusion [26, 32] (Fig. 1(a)) efficiently integrates LiDAR and visual features from aligned local fields in Bird’s Eye View (BEV)

*This work was completed during Hanshi’s remote internship at SJTU.

†Corresponding author.

space. Despite the simplicity, its capability is constrained by the lack of comprehensive contextual information from the surrounding environment (**efficient but short-sighted**). To leverage long-range dependencies, Sparse Global Fusion [57, 59] (Fig. 1(b)) employs self-attention to facilitate interaction among all object features extracted from the predicted bounding boxes of both modalities. While this approach effectively models long-range global dependency, it inadvertently discards scene information that, although not belonging to predefined categories, is crucial for accurately localizing targets (**efficient and far-sighted, but un-omniscient**¹). Quasi Dense Global Fusion [47, 61] (Fig. 1(c)) represents a compromise between the aforementioned frameworks, drawing inspiration from Swin Transformer [44] to shift the local fused BEV window. However, the inability to directly interact with global features hinders comprehensive scene understanding (**efficient but short-sighted and un-omniscient**). We summarize the characteristics of these frameworks in Tab. 1. Consequently, a critical question naturally arises: *Is an efficient, far-sighted, and omniscient long-range fusion framework possible?*

Before answering the question, we first discuss one key requirement for camera-LiDAR global fusion: efficiency. Efficiency is crucial, as marginal performance improvements that incur substantial computational costs are nearly insignificant for autonomous driving applications. This rationale leads us to forgo the straightforward approach of applying self-attention to the multimodal features, given the quadratic complexity associated with this operator. Fortunately, current linear complexity methods (*e.g.* Mamba and linear attention) may offer a solution. However, the adaptation is non-trivial. For instance, when we substitute the fusion module in recent method UniTR [47], we observe a degradation in performance on nuScenes (see Tab. 2). Then what challenges prevent their effective use?

We believe that the core challenge lies in how to effectively arrange the feature elements from the two modalities in a unified space. One intuitive approach is to use the BEV space as the unified coordinate system, a strategy that has been shown to be effective in previous Transformer-based local fusion models [32]. However, directly compressing the features into 3D space may reduce the efficacy of fusion, as it leads to the loss of height information that is crucial for global interaction. To address this issue, we propose using the image plane as the unified space, projecting the raw voxel features derived from the point clouds onto this space. Our study further reveals that the current voxel compression and merging techniques in LiDAR backbones introduce deviations in height information due to quantization errors in discrete space. This, in turn, leads to inevitable errors during multi-modality alignment (see Fig. 5). To mitigate this, we propose a height-fidelity LiDAR encoding

Framework	Efficient	Far-sighted	Omniscient
Dense Local	✓	✗	✗
Sparse Global	✓	✓	✗
Quasi Dense Global	✓	✗	✗
Dense Global (Ours)	✓	✓	✓

Table 1. Analysis of different fusion strategies.

method, which performs voxel compression in continuous space, thereby retaining more accurate and reliable height information. With the constructed unified space and height-aware representation, we introduce a Hybrid Mamba Block (HMB) that simultaneously and efficiently learns mutual information at both local and global scales.

By incorporating the proposed method into a new framework, our model achieves state-of-the-art performances with NDS score of 75.0 on nuScenes [2] validation dataset. Moreover, our framework demonstrates a 1.5× faster speed compared to recent top method IS-FUSION [53].

In summary, our main contributions are as follows: 1) We introduce the Hybrid Mamba Block to enable efficient dense global multi-modality fusion, which is compatible with different linear attention. 2) We propose a height-fidelity LiDAR encoding technique that retains more reliable height information, thereby improving the accuracy of camera-LiDAR alignment in a unified space. 3) For the first time, we demonstrate that a framework utilizing pure linear operations for fusion can achieve SOTA performance in camera-LiDAR multi-modal 3D object detection.

2. Related Work

3D Object Detection with Cameras-LiDAR Fusion. 3D object detection using cameras and LiDAR is critical for autonomous vehicles. Early work focused on independent methods for each sensor type [12, 16, 22, 23, 27, 28, 30, 36, 37, 48, 49]. To improve driving safety, researchers explored fusion strategies that leverage both modalities’ complementary strengths. Early fusion [18, 39, 45, 46, 54] augments raw point clouds with aligned image features. Middle fusion [26, 32] extracts mutual information in a high-dimensional representation space [32, 47, 53], enabling richer feature learning. Late fusion [57, 59] aligns and merges features from predicted boxes. Revisiting these methods reveals that efficient long-range context learning remains a major bottleneck under strict latency constraints.

Long Range Modeling. Capturing long-range dependencies is essential for learning robust representations. In the context of language and vision foundation models, Transformers [44] and large convolutional kernels [6, 29, 38] have been proposed to address this challenge. However, the direct application of these operators to 3D detection framework presents significant difficulties. Specifically, conventional convolutional kernels struggle with the sparse nature of LiDAR data, while the global attention mechanism in

¹“Omniscient”: each position can interact with all scene information.

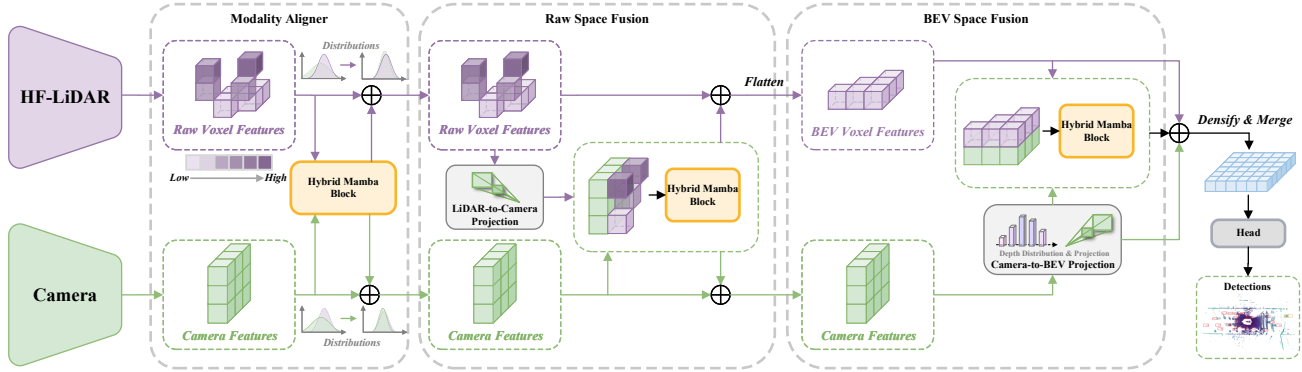


Figure 2. **Overview of our framework.** We propose a novel height-fidelity LiDAR encoding method, which preserves more accurate height information within voxel features, thereby enhancing the precision of multi-modal alignment. Then, we introduce the Hybrid Mamba Block (HMB), designed to: 1) align the latent-space distributions across modalities, 2) facilitate cross-modal information exchange between raw voxel and camera features, and 3) perform efficient multi-modal representation fusion and compression within BEV space.

Transformers incurs a prohibitively high quartic computational cost. To mitigate these challenges, window-based methods [47, 53] have been introduced, which perform global fusion exclusively among window features, thereby reducing latency. However, this approach inevitably sacrifices some detail during long-range feature modeling.

Linear Attention. To mitigate the quadratic complexity of vanilla attention, linear attention [20] was proposed and has shown strong gains [35, 43]. State space models (SSMs), a variant of linear attention [13], demonstrate superior performance. They have been applied in vision tasks [31, 60], point cloud processing [25, 33, 56, 58], and RGB-T fusion [7]. Yet fusing 2D camera and 3D LiDAR remains a significant challenge due to inherent disparities between these modalities. DRAMA [55] uses Mamba to merge LiDAR and camera BEV features but compresses point clouds into BEV maps, sacrificing height information and relying on CNN backbones. In this work, we for the first time adapt pure linear attention to capture long and short-term dependencies in multi-modal 3D detection for autonomous driving, which enables flow of long-range information from single modality representation learning to multi-modal fusion.

3. Method

3.1. Analysis of Different Linear Attention

In point cloud processing, non-uniform spatial intervals arising from occlusion and varying distances pose significant challenges for conventional linear attention methods. Some linear attention methods (e.g., Criss-Cross Attention [19]) require continuous 2D inputs, making them unsuitable for sparse point clouds, while others (e.g., RWKV [35]) rely on fixed space/time intervals. In contrast, Mamba leverages HiPPO [9] to capture long-range dependencies without relying on explicit space/time priors, which is an essential feature for handling variable spacing in sparse point clouds. As illustrated in Tab. 2, replac-

Components	vanilla	improvements	+ours	improvements
RetNet [43]	64.6/68.3	-6.0/-4.6	70.7/73.1	+0.1/+0.2
RWKV [35]	64.9/68.8	-5.7/-4.1	70.8/73.3	+0.2/+0.4
Mamba [8]	65.8/70.4	-4.8/-2.5	71.9/74.3	+1.3/+1.4

Table 2. Comparison of different linear attention methods on 50% nuScenes v.s. UniTR [47] with 70.6/72.9 (mAP/NDS).

ing the window-based transformer in UniTR [47] with linear attention results in a noticeable performance drop. Although Mamba also experiences some decline, it maintains a smaller performance drop compared to other methods, demonstrating its relative robustness. Based on the above analysis, these findings underscore Mamba’s advantages, motivating its adoption for our main experiments. The following sections detail our Mamba-based approach.

3.2. Overall Architecture

Fig. 2 illustrates the architecture of our proposed framework, designed for dense global fusion in multimodal 3D detection. Firstly, we extract visual and sparse point cloud features from multi-view images and point clouds. In this work, we propose the height-fidelity encoding strategy to preserve more detailed height information of point cloud features. This strategy is specifically designed to enhance the quality of multi-modal feature alignment during the fusion of LiDAR and camera features, as illustrated in Fig. 4 and Sec. 3.4. Secondly, the modality aligner harmonizes the feature distributions of the two modalities, stabilizing the multi-modal fusion process. Thirdly, to learn the mutual information between the modalities, we project their features into unified coordinate systems in two ways: 1) height-fidelity LiDAR features are projected into the image space and reordered with the raw visual feature embeddings; 2) visual features are projected into the BEV space and reordered with the point cloud features. These reordered features are fused using our proposed Hybrid Mamba Block (see Sec. 3.3). Finally, the detection head predicts the results using the fused features from the BEV space.

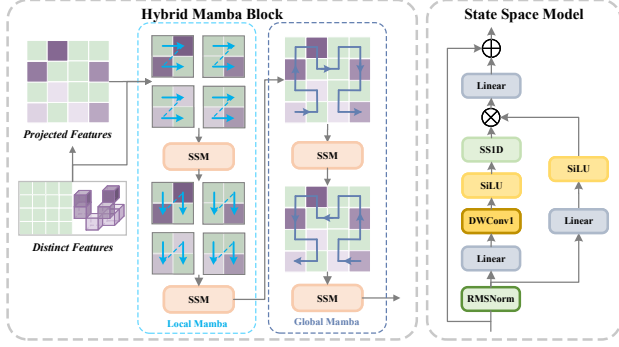


Figure 3. **Illustration of the Hybrid Mamba Block.** The block first employs the **Local Mamba** to enable localized feature interaction around each spatial position. Subsequently, the **Global Mamba** facilitates interaction across the entire space.

3.3. Hybrid Mamba Block

To achieve dense global fusion, we introduce the Hybrid Mamba Block, which combines local and global Mamba modules to integrate information across all modalities in the entire scenario. The local module captures fine-grained features, while the global module enhances the understanding of the broader driving environment. Notably, the proposed Hybrid Mamba Block can either independently process the features from each modality or jointly fuse their information within a unified coordinate system, offering flexibility to adapt to various fusion requirements.

Preliminaries. Rooted in linear systems theory [14], the State Space Model (SSM) [5, 8, 10, 11] provides a robust framework for representing dynamical systems. The continuous-time SSM is:

$$\mathbf{h}'(t) = \mathbf{A}\mathbf{h}(t) + \mathbf{B}\mathbf{x}(t), \quad (1)$$

$$\mathbf{y}(t) = \mathbf{C}^\top \mathbf{h}(t) + \mathbf{D}\mathbf{x}(t), \quad (2)$$

where $\mathbf{h}(t) \in \mathbb{R}^C$ is the hidden state, $\mathbf{y}(t) \in \mathbb{R}^L$ is the output, \mathbf{A} governs system dynamics, \mathbf{B} insert the input $\mathbf{x}(t)$ to the state, \mathbf{C} projects the state to output, and \mathbf{D} allows residual connection. \mathbf{A} based on HiPPO [9] in Mamba effectively captures long-range dependencies by transforming global features into a compressed representation.

Global Mamba. It is well-established that Mamba was originally designed to process 1D sequential data. To extend its applicability to 2D camera/BEV space and raw 3D space, we adopt the Hilbert curve [15], as used in prior work [24], to serialize the feature tokens from both modalities. Specifically, the Hilbert curve traverses the feature embeddings of both the camera and LiDAR modalities, serving as a space-filling curve renowned for preserving spatial locality, as illustrated in Fig. 3. This approach ensures efficient traversal and seamless integration of heterogeneous information in the entire 3D space.

More precisely, given the embeddings from single- or multi-modalities feature \mathbf{F} with their corresponding coordi-

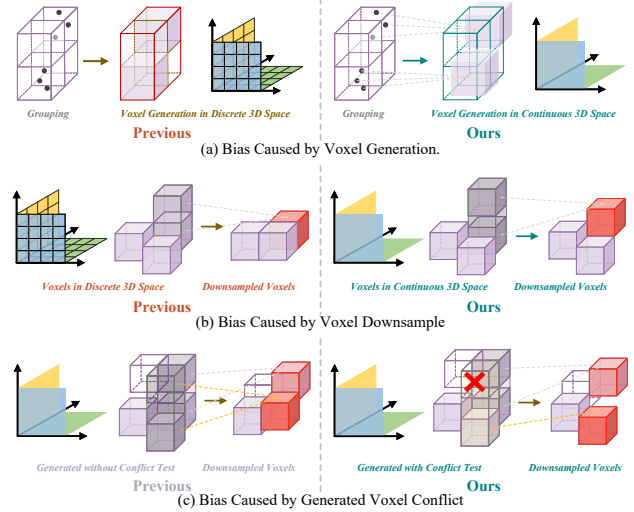


Figure 4. **Height-Fidelity LiDAR Encoding.** (a-b) Height information loss caused by quantization error in discrete space. (c) Height deviation introduced by the generated voxels in LION [33].

ates $\mathcal{C} = \{(x_i, y_i, z_i)\}_{i=1}^{N_{\text{feat}}}$, we compute the Hilbert indices $\mathcal{P}_h = \{h_i\}_{i=1}^{N_{\text{feat}}}$ for each embedding following [40],

$$h_i = \text{HilbertIndex}(x_i, y_i, z_i), \quad (3)$$

where $\text{HilbertIndex}(\cdot)$ maps the 3D coordinate to a 1D Hilbert curve index, so that we can reorder the embeddings in features \mathbf{F} . Besides, we use positional embedding to preserve information about the original coordinates,

$$\mathbf{F}' \leftarrow \mathbf{F} + \text{PosEmbedding}(\mathcal{C}). \quad (4)$$

Since the receptive field of Mamba is unilateral, we process the reordered embeddings in a bi-directional way,

$$\mathbf{F}_{\text{global}} = \text{Mamba}(\mathbf{F}', \mathcal{P}_h). \quad (5)$$

In this manner, each token comprehensively integrates information from all available features, effectively capturing long-range dependencies across the entire space.

Local Mamba. Representation learning in local regions has been extensively studied, which is demonstrated to be a valuable complement to global features. Therefore, we also adopt Mamba in small regions to capture local structure. As depicted in Fig. 3, we segment the input features into non-overlapping regions of size $w \times w$. For each feature embedding $\mathbf{f}_i \in \mathbb{R}^C$ located at coordinates (x_i, y_i, z_i) , we compute its region index r_i and its positions (x'_i, y'_i) in the local region as follows:

$$\begin{cases} r_i = \lfloor \frac{x_i}{w} \rfloor \times \lfloor \frac{y_i}{w} \rfloor + \lfloor \frac{y_i}{w} \rfloor, \\ (x'_i, y'_i) = (x_i \bmod w, y_i \bmod w). \end{cases} \quad (6)$$

By combining the region index and the in-region positions, we obtain the new positions $\mathcal{P}_r = \{r_i, x'_i, y'_i\}_{i=1}^{N_{\text{feat}}}$ for each

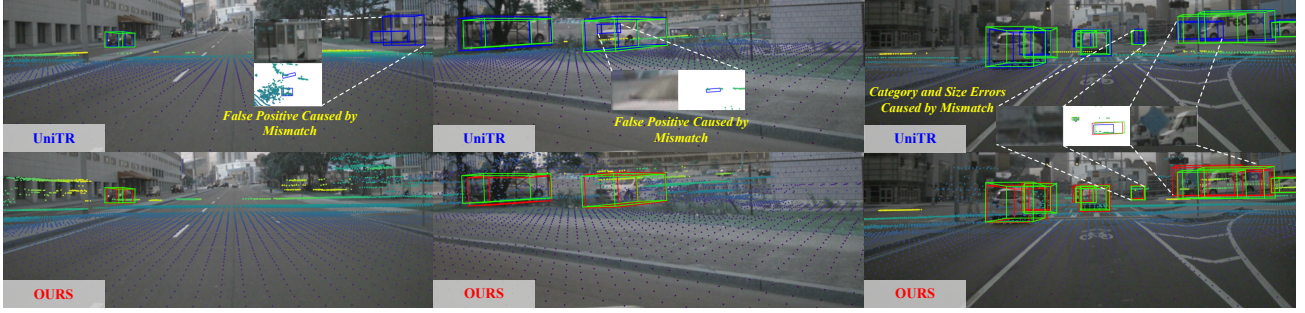


Figure 5. **Voxel projection and detection results on images.** Comparison of our projection method (bottom), enhanced by our Height-Fidelity LiDAR encoding, with previous methods (top). The figure demonstrates that projection errors introduce false positives, category errors, and unsatisfactory scales. The **green, blue** and **red** boxes indicate ground truth, UniTR’s output, and our results, respectively.

feature embedding. We then apply the Mamba to aggregate information within each region using these new positions, yielding locally enriched features $\mathbf{F}_{\text{local}}$ as follows:

$$\mathbf{F}_{\text{local}} = \text{LocalMamba}(\mathbf{F}, \mathcal{P}_r). \quad (7)$$

To facilitate comprehensive information exchange, we perform fusion along both the x and y directions (see Fig. 3).

3.4. Height-Fidelity LiDAR Encoding

As shown in Fig. 4(a), during voxel generation in the LiDAR backbone, each voxel’s position is represented solely by the 3D coordinates of its centroid, disregarding the internal point distribution within the voxel. Conventional voxel feature integration repeatedly downscales features along the height dimension, which determines the 3D position of new voxels by averaging the coordinates of the merged ones. This process results in further loss of height information, as illustrated in Fig. 4(b). Fig. 5 shows the negative effects of height information loss during multi-modality alignment, which eventually degrade detection accuracy and quality.

To mitigate the height information loss caused by quantization errors in discrete space, we calculate the voxel coordinates directly in continuous 3D space.

$$\mathcal{C}_{P_0} = \text{ScatterMean}(\mathbf{c}, \mathbf{v}_0), \quad (8)$$

where \mathbf{v}_0 represents the voxel assigned for each point, $\mathbf{c} \in \mathbb{R}^{N_{\text{point}} \times 3}$ is the coordinates of point clouds. Then, we also perform the downsampling operation using raw coordinates in continuous 3D space. Specifically, as depicted in Fig. 4(b), during voxel merging at stage s , we compute the coordinates \mathcal{C}_{P_s} of the corresponding voxels in raw space based on the contained point clouds,

$$\mathcal{C}_{P_s} = \text{ScatterMean}(\mathcal{C}_{P_{s-1}}, \mathbf{v}_{s-1}). \quad (9)$$

In addition to biases introduced by discretized coordinates, voxel generation also introduces deviation in our baseline LiDAR backbone LION [33]. Specifically, it selects the top k salient features with coordinates

$\{x_i, y_i, z_i\}_{i=1}^k$, and generate new voxels at positions $\{(x_i \pm 1, y_i \pm 1, z_i)\}_{i=1}^k$. However, as illustrated in Fig. 4(c), this approach may introduce errors during the downsampling process due to overlapping or conflicting voxels. To address this issue, we exclude voxels that are likely to merge with surrounding voxels, ensuring that only distinct, non-overlapping voxels are retained,

$$\left\{ \left(x_i \pm 1, y_i \pm 1, \left\lfloor \frac{z_i}{m} \right\rfloor \right) \right\}_{i=1}^k \notin \mathcal{C}_{P_s}, \quad (10)$$

where m is the downsampling scale.

3.5. Multi-modality Representation Learning

This section provides a detailed description of the development of a fusion framework incorporating the proposed modules. First, we align the feature distributions of point cloud and image modalities by passing them through a shared HMB. This process is akin to the behavior of “normalization” layers. After achieving consistent feature distributions, we perform fusion in both the raw 3D space and the BEV space. In the raw space, we project the 3D voxel features onto the image planes, facilitating the fusion of the image and point cloud features. In the BEV space, we use the Lift-Splat-Shoot (LSS) transformation to project image features into the BEV space and fuse them with the point cloud features, creating a unified multi-modal representation. This fusion process is carried out efficiently by downsampling BEV features and preserving only relevant information in the point cloud features, ultimately providing a rich representation for the detection task.

4. Experiments

4.1. Experimental Setup

Datasets and Evaluation Metrics. We evaluate our framework on the nuScenes [2] dataset, which comprises 40,157 annotated samples collected from six cameras providing a full 360-degree field of view and a 32-beam LiDAR sensor. For evaluation, we employ Mean Average Precision (mAP) and the nuScenes Detection Score (NDS).

Methods	Present at	Resolution	Modality	NDS (val)	mAP (val)	NDS (test)	mAP (test)	FPS
PointPainting [45]	CVPR'20	-	C+L	69.6	65.8	-	-	-
PointAugmenting [46]	CVPR'21	1600 × 900	C+L	-	-	71.0	66.8	-
MVP [54]	NeurIPS'21	1600 × 900	C+L	70.0	66.1	70.5	66.4	-
TransFusion [1]	CVPR'22	800 × 448	C+L	71.3	67.5	71.6	68.9	-
AutoAlignV2 [4]	ECCV'22	1440 × 800	C+L	71.2	67.1	72.4	68.4	-
UVTR [21]	NeurIPS'22	1600 × 900	C+L	70.2	65.4	71.1	67.1	-
BEVFusion (PKU) [26]	NeurIPS'22	704 × 256	C+L	71.0	67.9	71.8	69.2	0.8
DeepInteraction [51]	NeurIPS'22	800 × 448	C+L	72.6	69.9	73.4	70.8	2.6
BEVFusion (MIT) [32]	ICRA'23	704 × 256	C+L	71.4	68.5	72.9	70.2	4.2
CMT [50]	ICCV'23	1600 × 640	C+L	72.9	70.3	74.1	72.0	2.8
SparseFusion [59]	ICCV'23	704 × 256	C+L	72.8	70.4	73.8	72.0	5.3
ObjectFusion [3]	ICCV'23	704 × 256	C+L	72.3	69.8	73.3	71.0	-
UniTR [47]	ICCV'23	704 × 256	C+L	73.3	70.5	74.5	70.9	4.9
IS-FUSION [53]	CVPR'24	1056 × 384	C+L	74.0	72.8	75.2	73.0	3.2
DAL [17]	ECCV'24	1056 × 384	C+L	74.0	71.5	74.8	72.0	4.3
SparseLIF [57]	ECCV'24	1600 × 640	C+L	74.6	71.2	-	-	2.9
MambaFusion-Lite	Ours	704 × 256	C+L	74.0	71.6	75.0	72.0	5.4
MambaFusion-Base	Ours	704 × 256	C+L	75.0	72.7	75.9	73.2	4.7

Table 3. Comparison on the nuScenes validation and test datasets [2]. The modalities used are Camera (C) and LiDAR (L). The highest scores are in **red**, while the second-highest are in **blue**. Note that SparseLIF does not provide test set results for single-frame models.

	Baseline-LC	HMB-M	HMB-R	HMB-B	HFL	mAP	NDS
①	✓					67.7	70.9
②	✓	✓				68.6	71.6
③	✓		✓			66.9	70.2
④	✓			✓		69.6	72.0
⑤	✓		✓		✓	70.3	72.7
⑥	✓	✓		✓		70.6	73.1
⑦	✓	✓	✓		✓	71.4	73.6
⑧	✓		✓	✓	✓	71.2	73.7
⑨	✓	✓	✓	✓	✓	71.9	74.3

Table 4. Ablation study of each module. 'Baseline-LC' refers to a simplified version of the model, 'HFL' denotes the height-fidelity LiDAR encoding, and 'HMB-X' represents the Hybrid Mamba Block applied at different stages, as described in Sec. 3.5.

Implementation Details. Following UniTR [47], we use a voxel resolution of [0.3, 0.3, 0.25]. The spatial range is $[-54.0m, 54.0m]$ for x and y axes, and $[-5.0m, 3.0m]$ for z axis. Multi-view images are used at a resolution of 704×256 . The LiDAR and camera tokenizers are partly adapted from VMamba [31] and LION [33]. Adam optimizer and one-cycle learning rate policy [41] are used during training. Our model is trained for 10 epochs with a maximum learning rate of 1.5×10^{-3} on four A100 GPUs.

4.2. Comparison with State-of-the-Art Methods

As shown in Tab. 3, we present results on the nuScenes validation and test sets, obtained without test-time augmentations or model ensembling. Our Lite model outperforms the previous quasi-dense global fusion method UniTR [47] by 1.1 mAP and 0.7 NDS with 10% faster inference speed. Compared with the recent state of the arts, our Base model achieves 75.0 NDS on the validation dataset, surpassing all compared methods. Moreover, compared to the most recent

IS-Fusion [53] and SparseLIF [57], our method achieves a 50% and 62% increase in inference speed. Overall, our method strikes an impressive balance between performance and efficiency. Detailed per-category and other criteria like mASE are provided in the supplementary material.

4.3. Ablation Studies

This section analyzes the impact of each proposed module. To expedite validation during ablation experiments, we use only half of the dataset, constructed by selecting every alternate frame. Evaluation is performed on the nuScenes validation set unless otherwise specified.

4.3.1. Component-wise Ablation

As presented in Tab. 4, we conducted an ablation study to evaluate each component's contribution. The Baseline-LC model, integrating point cloud and image features following BEVFusion [26], achieved an mAP of 67.7 and an NDS of 70.9 (①). Incorporating Hybrid Mamba Block (HMB) at various stages generally enhanced performance (②④). However, when fusing multi-modal features in the raw 3D space (③), the performance degrades due to projection misalignments between modalities (Fig. 5). Since Mamba is sequential, alignment errors would be amplified compared to Transformers. Introducing the height-fidelity LiDAR encoding (HFL) solves this issue (⑤ vs. ③). Finally, integrating HMB at different stages along with HFL further enhanced performance, achieving 71.9 mAP and 74.3 NDS.

4.3.2. Analysis of Hybrid Mamba Block

In this section, we study the effect of detailed design in the proposed Hybrid Mamba Block. As shown in Tab. 5(a), we find that only using local Mamba leads to suboptimal

Component	mAP	NDS
Local only	70.5	72.1
Global only	71.5	73.7
Local and Global	71.9	74.3

(a) Local and global fusion.

Sorting Paradigm	mAP	NDS
Coordinate Order	71.2	73.5
Z-Order Curve [34]	71.5	74.0
Hilbert Curve [15]	71.9	74.3

(b) Different sorting paradigm in global Mamba.

Component	mAP	NDS
Baseline	69.5	72.2
+ xy in local	70.9	73.0
+ bidirectional in global	71.9	74.3

(c) Ablation on serialization combination.

Table 5. Design choices of hybrid Mamba block. We explore the impact of each detailed components.

	Backbone	VoxelVFE	Conflict Test	mAP	NDS
①				70.3	73.1
②	✓			71.2	73.8
③	✓	✓		71.5	73.9
④	✓	✓	✓	71.9	74.3

Table 6. “Backbone” denotes the phase of voxel integration (Fig. 4(a)). “VoxelVFE” refers to the stage when voxels are generated (Fig. 4(b)). “Conflict Test” refers to that in Fig. 4(c).

results, while only employing the global Mamba yields significantly better performance with 1.0 gains in mAP and 1.6 gains in NDS. Combining local and global Mamba components further enhances the results, demonstrating their complementary benefits. Tab. 5(b) presents the impact of different 3D space-filling curves, where the Hilbert Curve outperforms others in capturing spatial relationships. Lastly, as shown in Tab. 5(c), we investigate different serialization combinations within the Hybrid Mamba Block, where the bidirectional strategy demonstrates better performances.

4.3.3. Impact of Height-Fidelity LiDAR Encoding

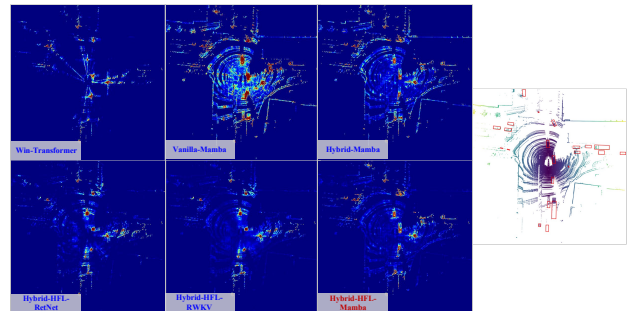
As shown in Fig. 5, previous methods that omit height information fail to effectively project corresponding point clouds to regions where objects are presented, resulting in false positives, misclassifications, and erroneous estimations of object scale. Quantitative analysis in Tab. 6 demonstrates that incorporating height information at different stages enhances performance. Specifically, adding height-aware encoding in the backbone improves the mAP by 0.9 and NDS by 0.7 compared to the baseline (② vs. ①). ③ and ④ further improve performance by preserving height information during voxel generation ((Fig. 4(b))) and applying the conflict test strategy (Fig. 4(c)), where ④ obtains an mAP of 71.9 and NDS of 74.3. These findings confirm the effectiveness of our height-fidelity LiDAR encoding in improving feature alignment between LiDAR and image data, thereby enhancing overall detection performance.

4.3.4. Comparison with Transformer

In this section, we analyze the performance and efficiency of the window-based Transformer compared to our proposed model. As shown in Tab. 7 ①, we replicate the method from UniTR [47]. Initially, we replace all Transformer layers with Mamba. While this substitution accelerates computation, it results in a significant decrease in performance (from 68.9 to 65.8 mAP in ⑤). To address this issue, we change the fusion space to Image+BEV (⑤) and

	Component	Alignment	Space	mAP	NDS	FPS
①	Win-Transformer	Discrete	Image+Frustum	68.9	71.9	4.9
②	Win-Transformer	Discrete	Image+BEV	69.5	72.2	3.2
③	Win-Transformer	Continuous	Image+BEV	70.6	72.9	3.1
④	Vanilla Mamba	Discrete	Image+Frustum	65.8	70.4	6.9
⑤	Vanilla Mamba	Discrete	Image+BEV	68.0	70.7	6.4
⑥	Hybrid Mamba	Discrete	Image+BEV	70.3	73.1	4.8
⑦	Hybrid Mamba	Continuous	Image+BEV	71.9	74.3	4.7

Table 7. Comparison with window-based Transformer. “Vanilla Mamba” denotes directly replacing the Transformer with Mamba. “Discrete” means using coordinates in discrete space, while “Continuous” refers to the continuous space employed in the height-fidelity LiDAR encoding. “Image+Frustum” corresponds to the coordinate space utilized in UniTR [47], whereas “Image+BEV” pertains to the coordinate space described in Sec. 3.5.

Figure 6. **The effective receptive fields (ERFs).** “Win” denotes window partition, “Hybrid” means Hybrid Mamba (ResNet, RWKV) Block, “HFL” refers to Height-Fidelity LiDAR Encoding.

incorporate the Hybrid Mamba Block along with the height-fidelity LiDAR encoding (⑥ and ⑦). These modifications lead to a substantial improvement in model performance. Finally, we apply the approaches from ⑤ and ⑦ to the window-based Transformer to evaluate the generalizability of our proposed method. However, the window-based Transformer incurs higher computational costs and obtains lower performance due to its limited global perception capability (② and ③). Moreover, our proposed height-fidelity LiDAR encoding significantly enhances performance with minimal additional computational overhead (⑦ vs. ⑥).

4.3.5. ERFs and Analysis of Long-rang Modeling

As shown in Fig. 6, we visualize the effective receptive fields (ERFs) to investigate how each approach captures long-range dependencies. Specifically, we randomly select tokens from ground-truth boxes, compute their ERFs, and merge the results by taking the maximum value across voxel

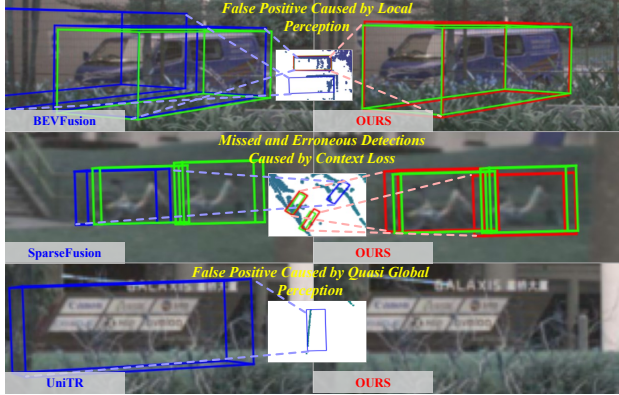


Figure 7. **Qualitative comparison of different fusion strategies.** The **green**, **blue** and **red** boxes indicate ground truth, output of previous methods, and our results, respectively.

locations. We observe that UniTR [47] with a window-based transformer captures local information, limiting its coverage of larger objects and overall environmental context, while projection errors introduce additional noise. Vanilla Mamba provides long-range modeling and global receptive field but falls in capturing information effectively. The Hybrid Mamba Block improves spatial modeling, while misaligned modalities shift attention to irrelevant areas. The inclusion of HFL mitigates this misalignment. Next, we compare various linear attention mechanisms on ERFs. Mamba demonstrates stronger coverage and focuses more effectively on key objects, whereas RWKV [35] and RetNet [43] exhibit additional noise and fail to attend to certain crucial targets. Finally, Tab. 2 shows that although all linear attentions exceed the baseline, Mamba achieves the highest performance, reaffirming our choice of Mamba.

4.3.6. Visualization and Analysis.

In this section, we present a visual analysis of three prominent fusion strategies, highlighting the limitations that our method effectively addresses. As shown in Fig. 7, Dense Local Fusion (BEVFusion [32]) struggles with long-range dependency modeling. The localized receptive field and errors in projecting image features to BEV cause a single object to disperse across multiple BEV regions, resulting in false positives. Sparse Global Fusion (SparseFusion [59]) faces issues when background context is crucial for object differentiation. The loss of background information causes misalignment between image and LiDAR features, resulting in missed or incorrect detections. Similarly, Quasi-Dense Global Fusion (UniTR [47]) struggles to capture objects requiring long-range dependencies, especially those that don't fit within a single window or rely on background differentiation. This limited observation makes it difficult to accurately determine the object's category, leading to false positives. In contrast, our method addresses these limitations by employing a Dense Global Fusion framework.

Approach	Clean		Missing F		Preserve F		Stuck	
	mAP	NDS	mAP	NDS	mAP	NDS	mAP	NDS
PointAugmenting [46]	46.9	55.6	42.4	53.0	31.6	46.5	42.1	52.8
MX-Net [39]	61.0	66.1	47.8	59.4	17.5	41.7	43.8	58.8
TransFusion [1]	66.9	70.9	65.3	70.1	64.4	69.3	65.9	70.2
BEVFusion [26]	67.9	71.0	65.9	70.7	65.1	69.9	66.2	70.3
UniTR [47]	70.5	73.3	68.5	72.4	66.5	71.2	68.1	71.8
Ours (full dataset)	72.7	75.0	70.6	74.1	67.9	72.3	71.3	74.2

Table 8. Evaluating the robustness on degradation conditions. “F” is the front camera, and Stuck means unsynchronized timestamp between LiDAR and cameras.

Method	DE _{avg} ↓	CR _{avg} ↓
FusionAD [52]	0.81	0.12
+MambaFusion (ours)	0.70	0.10

Table 9. End-to-end performance on nuScenes dataset.

4.4. Robustness Studies

Following the BEVFusion [26] evaluation protocols, we evaluate our method on different degradation conditions, including missing the front camera (“Missing F” in Tab. 8), only preserving the front camera (“Preserve F” in Tab. 8) and unsynchronized LiDAR and camera timestamps (“Stuck” in Tab. 8). As shown in Tab. 8, our model demonstrates superior robustness compared to other methods. Notably, our method performs particularly well in the “Stuck” setting, suggesting that global perception is advantageous in handling unsynchronized scenarios.

4.5. Generality: MambaFusion for End-to-end AD

Our approach aims to achieve generalizable multi-modality long-range modeling. As shown in Tab. 9, integrating it into FusionAD [52] yields improvements in motion planning, reducing average Displacement Error (DE_{avg}) and Collision Rate (CR_{avg}). We further validate its effectiveness on Waymo [42] and in BEV segmentation on nuScenes (see supplementary material). Collectively, these results underscore the robustness and versatility of our method.

More Experiments. We conduct experiments on different backbones and resolutions in supplementary material.

5. Conclusion

In this paper, we introduce the first dense global fusion framework for multi-modal 3D object detection. The core components contributing to our success are the efficient Hybrid Mamba Block and the effective height-fidelity LiDAR encoding. The former efficiently aggregates global information and the latter enables better feature alignment between modalities. Our method achieves superior performance on the competitive nuScenes benchmark, demonstrating its effectiveness in real-world scenarios. Analytical experiments validate the efficacy of our proposed approach and provide insights into the reasons behind its performance gains. We hope that our work will inspire further research in multi-modal 3D object detection and encourage the exploration of Mamba in related fields.

Acknowledgements

This work was supported in part by the Natural Science Foundation of China (Grant No. 62422317, U22B2056, 62036011, 62192782, U2441241), the Beijing Natural Science Foundation (Grant No. JQ22014, L223003), the Project of Beijing Science and Technology Committee (Project No. Z231100005923046).

References

- [1] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. In *CVPR*, pages 1090–1099, 2022. 6, 8
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, pages 11621–11631, 2020. 1, 2, 5, 6
- [3] Qi Cai, Yingwei Pan, Ting Yao, Chong-Wah Ngo, and Tao Mei. Objectfusion: Multi-modal 3d object detection with object-centric fusion. In *ICCV*, pages 18067–18076, 2023. 6
- [4] Zehui Chen, Zhenyu Li, Shiquan Zhang, Liangji Fang, Qin-hong Jiang, and Feng Zhao. Deformable feature aggregation for dynamic multi-modal 3d object detection. In *ECCV*, pages 628–644. Springer, 2022. 6
- [5] Tri Dao and Albert Gu. Transformers are ssms: Generalized models and efficient algorithms through structured state space duality. *arXiv preprint arXiv:2405.21060*, 2024. 4
- [6] Xiaohan Ding, Xiangyu Zhang, Jungong Han, and Guiguang Ding. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *CVPR*, pages 11963–11975, 2022. 2
- [7] Xiangbo Gao, Asiegbu Miracle Kanu-Asiegbu, and Xiaoxiao Du. Mambast: A plug-and-play cross-spectral spatial-temporal fuser for efficient pedestrian detection. *arXiv preprint arXiv:2408.01037*, 2024. 3
- [8] Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2023. 1, 3, 4
- [9] Albert Gu, Tri Dao, Stefano Ermon, Atri Rudra, and Christopher Ré. Hippo: Recurrent memory with optimal polynomial projections. *NIPS*, 33:1474–1487, 2020. 3, 4
- [10] Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured state spaces. *arXiv preprint arXiv:2111.00396*, 2021. 4
- [11] Albert Gu, Isys Johnson, Karan Goel, Khaled Saab, Tri Dao, Atri Rudra, and Christopher Ré. Combining recurrent, convolutional, and continuous-time models with linear state space layers. *NeurIPS*, 34:572–585, 2021. 4
- [12] Chunrui Han, Jinrong Yang, Jianjian Sun, Zheng Ge, Runpei Dong, Hongyu Zhou, Weixin Mao, Yuang Peng, and Xiangyu Zhang. Exploring recurrent long-term temporal fusion for multi-view 3d perception. 2024. 2
- [13] Dongchen Han, Ziyi Wang, Zhuofan Xia, Yizeng Han, Yifan Pu, Chunjiang Ge, Jun Song, Shiji Song, Bo Zheng, and Gao Huang. Demystify mamba in vision: A linear attention perspective. *NIPS*, 37:127181–127203, 2025. 3
- [14] Joao P Hespanha. *Linear systems theory*. Princeton university press, 2018. 4
- [15] David Hilbert and David Hilbert. Über die stetige abbildung einer linie auf ein flächenstück. *Dritter Band: Analysis-Grundlagen der Mathematik- Physik Verschiedenes: Nebst Einer Lebensgeschichte*, pages 1–2, 1935. 4, 7
- [16] Junjie Huang, Guan Huang, Zheng Zhu, Yun Ye, and Dalong Du. Bevdet: High-performance multi-camera 3d object detection in bird-eye-view. *arXiv preprint arXiv:2112.11790*, 2021. 2
- [17] Junjie Huang, Yun Ye, Zhujin Liang, Yi Shan, and Dalong Du. Detecting as labeling: Rethinking lidar-camera fusion in 3d object detection. In *ECCV*, pages 439–455. Springer, 2025. 6
- [18] Tengting Huang, Zhe Liu, Xiwu Chen, and Xiang Bai. Epnnet: Enhancing point features with image semantics for 3d object detection. In *ECCV*, pages 35–52. Springer, 2020. 2
- [19] Zilong Huang, Xinggang Wang, Lichao Huang, Chang Huang, Yunchao Wei, and Wenyu Liu. Ccnet: Criss-cross attention for semantic segmentation. In *ICCV*, pages 603–612, 2019. 3
- [20] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *IMCL*, pages 5156–5165. PMLR, 2020. 3
- [21] Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based representation with transformer for 3d object detection. *NeurIPS*, 35:18442–18455, 2022. 6
- [22] Yin hao Li, Zheng Ge, Guanyi Yu, Jinrong Yang, Zengran Wang, Yukang Shi, Jianjian Sun, and Zeming Li. Bevdepth: Acquisition of reliable depth for multi-view 3d object detection. In *AAAI*, pages 1477–1485, 2023. 2
- [23] Zhiqi Li, Wenhai Wang, Hongyang Li, Enze Xie, Chonghao Sima, Tong Lu, Yu Qiao, and Jifeng Dai. Bevformer: Learning bird’s-eye-view representation from multi-camera images via spatiotemporal transformers. In *ECCV*, pages 1–18. Springer, 2022. 2
- [24] Zhuoyuan Li, Yubo Ai, Jiahao Lu, ChuXin Wang, Jiacheng Deng, Hanzhi Chang, Yanzhe Liang, Wenfei Yang, Shifeng Zhang, and Tianzhu Zhang. Mamba24/8d: Enhancing global interaction in point clouds via state space model. *arXiv preprint arXiv:2406.17442*, 2024. 4
- [25] Ding kang Liang, Xin Zhou, Xinyu Wang, Xingkui Zhu, Wei Xu, Zhikang Zou, Xiaoqing Ye, and Xiang Bai. Pointmamba: A simple state space model for point cloud analysis. *arXiv preprint arXiv:2402.10739*, 2024. 3
- [26] Tingting Liang, Hongwei Xie, Kaicheng Yu, Zhongyu Xia, Zhiwei Lin, Yongtao Wang, Tao Tang, Bing Wang, and Zhi Tang. Bevfusion: A simple and robust lidar-camera fusion framework. *NeurIPS*, 35:10421–10434, 2022. 1, 2, 6, 8
- [27] Xuewu Lin, Tianwei Lin, Zixiang Pei, Lichao Huang, and Zhizhong Su. Sparse4d v2: Recurrent temporal fusion with sparse model. *arXiv preprint arXiv:2305.14018*, 2023. 2

- [28] Haisong Liu, Yao Teng, Tao Lu, Haiguang Wang, and Limin Wang. Sparsebev: High-performance sparse 3d object detection from multi-camera videos. In *ICCV*, pages 18580–18590, 2023. 2
- [29] Shiwei Liu, Tianlong Chen, Xiaohan Chen, Xuxi Chen, Qiao Xiao, Boqian Wu, Tommi Kärkkäinen, Mykola Pechenizkiy, Decebal Mocanu, and Zhangyang Wang. More convnets in the 2020s: Scaling up kernels beyond 51x51 using sparsity. *arXiv preprint arXiv:2207.03620*, 2022. 2
- [30] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. Petr: Position embedding transformation for multi-view 3d object detection. In *ECCV*, pages 531–548. Springer, 2022. 2
- [31] Yue Liu, Yunjie Tian, Yuzhong Zhao, Hongtian Yu, Lingxi Xie, Yaowei Wang, Qixiang Ye, and Yunfan Liu. Vmamba: Visual state space model. *arXiv preprint arXiv:2401.10166*, 2024. 3, 6
- [32] Zhijian Liu, Haotian Tang, Alexander Amini, Xinyu Yang, Huizi Mao, Daniela L Rus, and Song Han. Bevfusion: Multi-task multi-sensor fusion with unified bird’s-eye view representation. pages 2774–2781. IEEE, 2023. 1, 2, 6, 8
- [33] Zhe Liu, Jinghua Hou, Xinyu Wang, Xiaoqing Ye, Jingdong Wang, Hengshuang Zhao, and Xiang Bai. Lion: Linear group rnn for 3d object detection in point clouds. *arXiv preprint arXiv:2407.18232*, 2024. 3, 4, 5, 6
- [34] Guy M Morton. A computer oriented geodetic data base and a new technique in file sequencing. 1966. 7
- [35] Bo Peng, Eric Alcáide, Quentin Anthony, Alon Albalak, Samuel Arcadinho, Stella Biderman, Huanqi Cao, Xin Cheng, Michael Chung, Matteo Grella, et al. Rwkv: Reinventing rnns for the transformer era. *arXiv preprint arXiv:2305.13048*, 2023. 1, 3, 8
- [36] Jonah Philion and Sanja Fidler. Lift, splat, shoot: Encoding images from arbitrary camera rigs by implicitly unprojecting to 3d. In *ECCV*, pages 194–210. Springer, 2020. 2
- [37] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30, 2017. 2
- [38] Yongming Rao, Wenliang Zhao, Yansong Tang, Jie Zhou, Ser Nam Lim, and Jiwen Lu. Hornet: Efficient high-order spatial interactions with recursive gated convolutions. *NeurIPS*, 35:10353–10366, 2022. 2
- [39] Vishwanath A Sindagi, Yin Zhou, and Oncel Tuzel. Mvxnet: Multimodal voxelnet for 3d object detection. pages 7276–7282. IEEE, 2019. 2, 8
- [40] John Skilling. Programming the hilbert curve. In *AIP Conference Proceedings*, pages 381–387. American Institute of Physics, 2004. 4
- [41] Leslie N Smith. Cyclical learning rates for training neural networks. pages 464–472. IEEE, 2017. 6
- [42] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, pages 2446–2454, 2020. 8
- [43] Yutao Sun, Li Dong, Shaohan Huang, Shuming Ma, Yuqing Xia, Jilong Xue, Jianyong Wang, and Furu Wei. Retentive network: A successor to transformer for large language models. *arXiv preprint arXiv:2307.08621*, 2023. 1, 3, 8
- [44] A Vaswani. Attention is all you need. *NeurIPS*, 2017. 2
- [45] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *CVPR*, pages 4604–4612, 2020. 2, 6
- [46] Chunwei Wang, Chao Ma, Ming Zhu, and Xiaokang Yang. Pointaugmenting: Cross-modal augmentation for 3d object detection. In *CVPR*, pages 11794–11803, 2021. 2, 6, 8
- [47] Haiyang Wang, Hao Tang, Shaoshuai Shi, Aoxue Li, Zhenguo Li, Bernt Schiele, and Liwei Wang. Unitr: A unified and efficient multi-modal transformer for bird’s-eye-view representation. In *ICCV*, pages 6792–6802, 2023. 2, 3, 6, 7, 8
- [48] Tai Wang, Xinge Zhu, Jiangmiao Pang, and Dahua Lin. Fcos3d: Fully convolutional one-stage monocular 3d object detection. In *ICCV*, pages 913–922, 2021. 2
- [49] Yue Wang, Vitor Campagnolo Guizilini, Tianyuan Zhang, Yilun Wang, Hang Zhao, and Justin Solomon. Detr3d: 3d object detection from multi-view images via 3d-to-2d queries. pages 180–191. PMLR, 2022. 2
- [50] Junjie Yan, Yingfei Liu, Jianjian Sun, Fan Jia, Shuailin Li, Tiancai Wang, and Xiangyu Zhang. Cross modal transformer via coordinates encoding for 3d object detection. *arXiv preprint arXiv:2301.01283*, 2(3):4, 2023. 6
- [51] Zeyu Yang, Jiaqi Chen, Zhenwei Miao, Wei Li, Xiatian Zhu, and Li Zhang. Deepinteraction: 3d object detection via modality interaction. *NeurIPS*, 35:1992–2005, 2022. 6
- [52] Tengju Ye, Wei Jing, Chunyong Hu, Shikun Huang, Lingping Gao, Fangzhen Li, Jingke Wang, Ke Guo, Wencong Xiao, Weibo Mao, et al. Fusionad: Multi-modality fusion for prediction and planning tasks of autonomous driving. *arXiv preprint arXiv:2308.01006*, 2023. 8
- [53] Junbo Yin, Jianbing Shen, Runnan Chen, Wei Li, Ruigang Yang, Pascal Frossard, and Wenguan Wang. Is-fusion: Instance-scene collaborative fusion for multimodal 3d object detection. In *CVPR*, pages 14905–14915, 2024. 2, 3, 6
- [54] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Multi-modal virtual point 3d detection. *NeurIPS*, 34:16494–16507, 2021. 2, 6
- [55] Chengran Yuan, Zhanqi Zhang, Jiawei Sun, Shuo Sun, Zefan Huang, Christina Dao Wen Lee, Dongen Li, Yuhang Han, Anthony Wong, Keng Peng Tee, et al. Drama: An efficient end-to-end motion planner for autonomous driving with mamba. *arXiv preprint arXiv:2408.03601*, 2024. 3
- [56] Guowen Zhang, Lue Fan, Chenhang He, Zhen Lei, Zhaoxiang Zhang, and Lei Zhang. Voxel mamba: Group-free state space models for point cloud based 3d object detection. *arXiv preprint arXiv:2406.10700*, 2024. 3
- [57] Hongcheng Zhang, Liu Liang, Pengxin Zeng, Xiao Song, and Zhe Wang. Sparselif: High-performance sparse lidar-camera fusion for 3d object detection. In *ECCV*, pages 109–128. Springer, 2025. 2, 6
- [58] Tao Zhang, Xiangtai Li, Haobo Yuan, Shunping Ji, and Shuicheng Yan. Point could mamba: Point cloud learning via state space model. *arXiv preprint arXiv:2403.00762*, 2024. 3

- [59] Zhizhuo Zhou and Shubham Tulsiani. Sparsefusion: Distilling view-conditioned diffusion for 3d reconstruction. In *CVPR*, pages 12588–12597, 2023. [2](#), [6](#), [8](#)
- [60] Lianghai Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint arXiv:2401.09417*, 2024. [3](#)
- [61] Yutao Zhu, Xiaosong Jia, Xinyu Yang, and Junchi Yan. Flatfusion: Delving into details of sparse transformer-based camera-lidar fusion for autonomous driving. *arXiv preprint arXiv:2408.06832*, 2024. [2](#)