

# Voxel or Pillar: Exploring Efficient Point Cloud Representation for 3D Object Detection

Yuhao Huang, Sanping Zhou, Junjie Zhang, Jinpeng Dong, Nanning Zheng\*

National Key Laboratory of Human-Machine Hybrid Augmented Intelligence,  
National Engineering Research Center for Visual Information and Applications,  
and Institute of Artificial Intelligence and Robotics, Xi'an Jiaotong University  
{hyh950623@stu, spzhou@mail, hooz1009@stu, djp1235a@stu, nnzheng@mail}.xjtu.edu.cn

## Abstract

Efficient representation of point clouds is fundamental for LiDAR-based 3D object detection. While recent grid-based detectors often encode point clouds into either voxels or pillars, the distinctions between these approaches remain under-explored. In this paper, we quantify the differences between the current encoding paradigms and highlight the limited vertical learning within. To tackle these limitations, we introduce a hybrid Voxel-Pillar Fusion network (VPF), which synergistically combines the unique strengths of both voxels and pillars. Specifically, we first develop a sparse voxel-pillar encoder that encodes point clouds into voxel and pillar features through 3D and 2D sparse convolutions respectively, and then introduce the Sparse Fusion Layer (SFL), facilitating bidirectional interaction between sparse voxel and pillar features. Our efficient, fully sparse method can be seamlessly integrated into both dense and sparse detectors. Leveraging this powerful yet straightforward framework, VPF delivers competitive performance, achieving real-time inference speeds on the nuScenes and Waymo Open Dataset.

## Introduction

LiDAR-based 3D object detection methods have been widely adopted in autonomous driving and robot navigation systems, as point clouds from LiDAR sensors reflect geometric information explicitly and are rarely affected by weather conditions. Unlike 2D images, point clouds are sparse and non-uniformly distributed, which poses two challenges for 3D object detection: constructing efficient detection network and establishing robust object representation from points with varying distributions.

There are several paradigms for point cloud representations. Range-view-based methods (Meyer et al. 2019; Bewley et al. 2021; Fan et al. 2021) convert point clouds into compact 2.5D range images and apply well-studied 2D detectors (Ren et al. 2015; Lin et al. 2017a) to predict 3D boxes. While efficient, these methods may distort the geometry of 3D point clouds, and introduce the scale variance problems (Li et al. 2019). Alternatively, point-based detectors (Qi et al. 2019; Shi, Wang, and Li 2018; Yang et al. 2020; Zhang et al. 2022) extract point-wise features with

\*Corresponding author.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

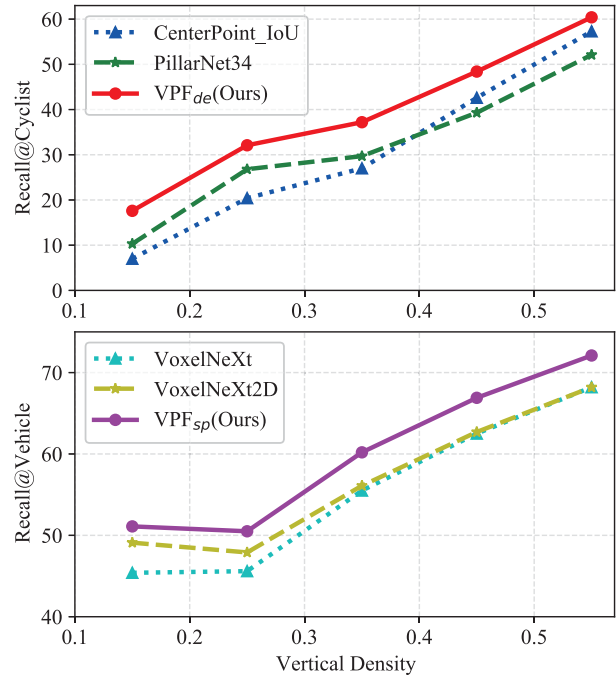


Figure 1: Recall vs. Vertical Density comparison. For both dense and sparse detectors (Yin, Zhou, and Krahenbuhl 2021; Guangsheng Shi 2022; Chen et al. 2023), pillar-based representations show enhanced recall under low vertical densities, while voxel-based representations tend to excel in high-density scenarios. Notably, our hybrid representation offers consistent improvements across different situations.

PointNet series directly (Qi et al. 2017a,b), and benefit from flexible receptive fields. Nevertheless, they suffer from the time-consuming spherical query and aggressive downsampling strategy. Finally, grid-based methods (Yin, Zhou, and Krahenbuhl 2021; Hu et al. 2022; Guangsheng Shi 2022; Shi et al. 2020a; He et al. 2022) quantize the point cloud into regular voxels or pillars and generate bounding boxes in 2D Bird's Eye View (BEV), achieving notable balance in both performance and efficiency.

While both the voxel and pillar are prevalent representations in grid-based methods, their distinctions are rarely

discussed. Given their different behavior in the vertical direction, we conduct an exploratory experiment on Waymo Open Dataset (Sun et al. 2020) to analyze the consequential effects. We first split the ground truth (GT) according to its vertical point distribution: Given each GT box, we uniformly divide it into 10 bins in the vertical direction, then calculate the vertical density as  $S_Z = N_{bin}/10$ , where  $N_{bin}$  denotes the number of non-empty bins. Next, we count the recall vs. vertical density curve of several grid-based detectors<sup>1</sup> (Yin, Zhou, and Krahenbuhl 2021; Guangsheng Shi 2022; Chen et al. 2023). As illustrated in Fig. 1, Voxel-based detectors (i.e., CenterPoint\_IoU and VoxelNeXt) show consistent better performance when the vertical density is high, while pillar-based methods (i.e., PillarNet and VoxelNeXt2D) excel in situations with low density.

This observation underscores the limited vertical representation in the current voxel and pillar-based methods. Specifically, detectors utilizing voxels typically employ the 3D submanifold convolution (Graham and van der Maaten 2017), which confines feature diffusion from non-empty voxels to empty ones, thereby restricting the receptive field. On the contrary, pillar-based methods discretized point clouds into vertical volumes, each of which encodes all its neighboring points in a certain X-Y coordinate. While this paradigm obtains the full-ranged vertical receptive field, it encounters challenges in capturing fine-grained features and is prone to significant information loss, particularly when processing areas with high point density. Given the limitations of both voxels and pillars, we are inspired to develop a method that **synergistically combines them for a more robust representation, especially in the vertical direction, while maintaining computational efficiency.**

In this paper, we propose Voxel-Pillar Fusion (VPF), a hybrid point cloud representation designed to synergistically harness the strengths of both voxels and pillars. We begin by crafting a sparse voxel-pillar encoder that segments point clouds into voxels and pillars, subsequently encoding these sparse volume features through 3D and 2D sparse convolutions. To enlarge the vertical receptive field of voxels and enrich the fine-grained information in pillars, we present the Sparse Fusion Layer (SFL) to establish the voxel-pillar bidirectional interaction. Specifically, SFL aggregates voxel features vertically and broadcasts pillar features to their corresponding vertical voxels. It then integrates these aggregated and broadcasted features with the original pillar and voxel features. Moreover, our method, both computationally efficient and fully sparse, can be seamlessly incorporated into both dense and sparse detectors. (Yin, Zhou, and Krahenbuhl 2021; Guangsheng Shi 2022; Chen et al. 2023). We also conduct comprehensive experiments on nuScenes (Caesar et al. 2020) and Waymo Open Dataset (Sun et al. 2020). The results indicate that our method achieves state-of-the-art performance with real-time inference speed, and enhances the vertical representation from both voxels and pillars as showcased in Fig. 1.

Our contributions are summarized as follows:

- We highlight the limitations in vertical representation learning of current grid-based methods and introduce a hybrid point cloud representation that synergistically harnesses the strengths of both voxels and pillars.
- We propose the Sparse Fusion Layer, which facilitates the voxel-pillar bidirectional interaction to enlarge the vertical receptive field of voxels and enrich the fine-grained information in pillars.
- We devise both dense and sparse detectors based on our proposed hybrid representation. Through comprehensive experiments on large-scale datasets, we validate the significance and practicality of the voxel-pillar fusion.

## Related Work

### Grid-based 3D Object Detection

Grid-based methods primarily utilize voxel or pillar representations. Voxel-based detectors divide point clouds into 3D voxels and deploy 3D convolutions for voxel-wise feature extraction. Pioneering work VoxelNet (Zhou and Tuzel 2018) replaces the hand-crafted representation with the voxel feature encoding layer, enabling the end-to-end training procedure. Then, SECOND (Yan, Mao, and Li 2018) introduces the 3D sparse convolution (Graham 2015; Graham and van der Maaten 2017) to avoid the redundant computation on empty voxels. Two-stage detectors (Deng et al. 2021; Yin, Zhou, and Krahenbuhl 2021) adopt the coarse-to-fine pattern in 2D detection (Ren et al. 2015; Lin et al. 2017a). For instance, Voxel R-CNN (Deng et al. 2021) proposes the voxel RoI pooling to extract RoI features from the voxels. As a contrast, single-stage methods (Zheng et al. 2021; Hu et al. 2022) directly generate bounding boxes in one stage. Some recent works (He et al. 2022; Wang et al. 2023; Yang et al. 2023) have been investigating the powerful transformer architecture (Vaswani et al. 2017; Dosovitskiy et al. 2020; Kirillov et al. 2020) in 3D object detection, while FSD (Fan et al. 2022b) and VoxelNeXt (Chen et al. 2023) exploring the fully sparse detection framework. For Pillar-based methods (Yang, Luo, and Urtasun 2018; Lang et al. 2019; Guangsheng Shi 2022), point clouds are encoded into 2D volumes (pillars). PointPillars (Lang et al. 2019) converts point clouds to pillars and deploys PointNet (Qi et al. 2017a) for pillar-wise feature extraction. PillarNet (Guangsheng Shi 2022) proposes the 2D sparse backbone for efficient pillar encoding, achieving a favorable trade-off between performance and inference speed. Then, SST (Fan et al. 2022a) introduces a single-stride transformer to enhance performance in the challenging pedestrian category. Though voxel or pillar representations are widely used, their differences are rarely discussed.

### Multi-Source Feature Fusion

Methods of multi-source fusion (Yang et al. 2019; Pang, Morris, and Radha 2020; Zhou et al. 2020a) are proposed to combine distinct information (Zhou et al. 2023) from different sources. For instance, point-voxel join representation (Yang et al. 2019; Shi et al. 2020b,a) is presented to integrate the flexible receptive field with efficient feature learning schemes. Besides, multi-view fusion (Chen

<sup>1</sup>We adopt single-stage voxel-encoded CenterPoint with the same detection head as PillarNet, as detailed in Ablation Studies.




Figure 2: The framework of VPF. Point clouds are first processed by the sparse voxel-pillar encoder, which extracts correlated sparse voxel and pillar features. The subsequent Sparse Fusion Layer facilitates bidirectional interaction, capturing supplementary information from both types of sparse features. Together, these components form a hybrid backbone capable of integrating with both dense and sparse detectors.

et al. 2016; Zhou et al. 2020b; Wang et al. 2020; Xin et al. 2019) is another paradigm for information supplementation. MVF (Zhou et al. 2020b) utilizes the complementary information from both BEV and perspective view, while Pillar-OD (Wang et al. 2020) applies pillar encoder in BEV and cylindrical view. Recently, multi-modal fusion (Liang et al. 2019; Pang, Morris, and Radha 2020; Li et al. 2022) has achieved remarkable progress. CLOCS (Pang, Morris, and Radha 2020) proposes a late fusion strategy to exploit the geometric and semantic consistencies between 2D and 3D predictions. DeepFusion (Li et al. 2022) presents InverseAug and LearnableAlign to align multi-modal features in the late stage. While multi-source fusion provides additional information from different sources, it also raises a problem. As the separate sources usually vary, effective alignment of multi-source features becomes challenging and sophisticated. In our work, we design a straightforward yet efficient structure for voxel-pillar fusion.

## Methodology

In pursuit of enhancing the vertical point cloud representation, we propose the Voxel-Pillar Fusion (VPF), a hybrid point cloud representation harnessing both 2D and 3D volume information. As shown in Fig. 2, point clouds are initially quantized into voxels and pillars with the same resolution in the X-Y plane. Then, we adopt the sparse Voxel-Pillar encoder with four intermediate steps, each of which features a sparse convolution block, with 3D and 2D sparse convolutions for voxel and pillar feature extraction, respectively. To enrich the local context for the pillar branch and infuse the vertical semantics into the voxel branch, we deploy the Sparse Fusion Layer (SFL) at the end of each step, creating a bidirectional interaction between the voxel and pillar features. Finally, we design both dense and sparse detectors, *i.e.*,  $VPF_{de}$  and  $VPF_{sp}$ , which are equipped with our presented hybrid representation backbone.

### Sparse Voxel-Pillar Encoder

**Consistent Voxel-Pillar Encoding.** While voxels and pillars represent point clouds differently, at the same resolution in the X-Y plane, the vertical collection of voxels contains the same point clouds as the corresponding pillar’s under specific X-Y coordinates. We exploit this trait to construct a consistent voxel-pillar encoder.

Given a point cloud  $\mathcal{S} = \{s_i \in \mathbb{R}^4\}_{i=1}^N$ , where  $N$  is the number of points. We first divide the 3D space into voxels and pillars with spatial resolution of  $L \times W \times H$  and  $L \times W$ , respectively. Next, dynamic voxelization (Zhou et al. 2020b) is deployed for initial sparse voxel and pillar feature generation. In the voxel branch,  $\mathcal{S}$  is quantized based on a pre-defined voxel size. This process yields point-to-voxel indices, represented as  $\mathcal{V} = \{v_i = (l_i, w_i, h_i) \in \mathbb{N}^3\}_{i=1}^N$ . Subsequently, the initial voxel features are derived by computing the mean of all points in each non-empty voxel with a sparse mean-pooling (Fey and Lenssen 2019). In the pillar branch, due to the consistency of voxels and pillars in Bird’s Eye View (BEV), the point-to-pillar indices,  $\mathcal{P} = \{p_i = (l_i, w_i) \in \mathbb{N}^2\}_{i=1}^N$  are obtained by removing the vertical indices. Finally, we apply PointNet (Lang et al. 2019) with sparse max-pooling to form the initial pillar features.

**Sparse Conv Block.** We extract sparse voxel and pillar features via dual-branch sparse convolution (conv) blocks, which include both 3D and 2D sparse convolutions. Similar to previous single branch framework (Yan, Mao, and Li 2018; Guangsheng Shi 2022), for each block, the voxel branch consists of a 3D regular sparse convolution followed by multiple 3D submanifold sparse convolutional layers, while the pillar branch is composed of its 2D equivalents. The regular sparse convolutions are applied for spatial downsampling, and submanifold sparse convolutions ensure that the output locations are identical to the input locations to optimize efficiency. As depicted in Fig. 3, for both 2D and 3D regular sparse convolutions, we equalize the kernel size, stride, and padding operations in the horizontal dimen-

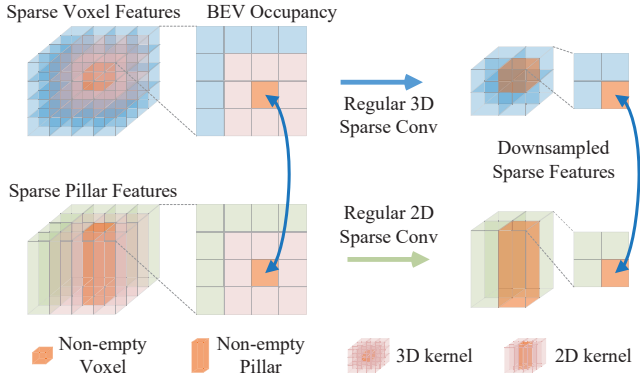


Figure 3: Consistent voxel-pillar downsampling process. In the downsampling procedure, by equalizing the kernel sizes, strides, and padding operations of 2D and 3D regular sparse convolutions in X-Y dimensions, the consistent BEV occupancy is preserved for sparse voxel and pillar features.

sion (*i.e.*,  $L$  and  $W$  dimensions). This ensures that voxel and pillar features expand to the same X-Y plane location. In such manner, non-empty voxels and pillars possess consistent occupancy in BEV, *i.e.*, each non-empty pillar has several corresponding non-empty voxels ( $\geq 1$ ) at the same X-Y coordinates. This consistency facilitates the fusion of sparse voxel and pillar features.

Overall, the sparse voxel-pillar encoder consists of 4 intermediate steps, where sparse conv blocks are deployed to sequentially generate sparse voxel and pillar features with  $1\times$ ,  $2\times$ ,  $4\times$ , and  $8\times$  downsampling sizes.

### Sparse Fusion Layer

Lateral connections are often used to merge different levels of semantics (Lin et al. 2017a) or dual-branch network (Feichtenhofer et al. 2018). Since the sparse voxel and pillar features are extracted separately via the 3D and 2D sparse convolutions, we introduce the Sparse Fusion Layer (SFL) to establish a bidirectional lateral connection for the element-wise fusion between the sparse voxel and pillar features.

**Sparse Pooling and Broadcasting.** We apply the sparse pooling and broadcasting operations to match the size and dimension of sparse features before the fusion. Denote the sparse voxel features as voxel indices  $\{v_i = (l_i, w_i, h_i) \in \mathbb{N}^3\}_{i=1}^{N_v}$  with corresponding features  $\{f_i^v \in \mathbb{R}^{D_v}\}_{i=1}^{N_v}$ , where  $N_v$  and  $D_v$  refer to the number of non-empty voxels and voxel feature dimension. Likewise, sparse pillar indices and features are formed as  $\{p_j = (l_j, w_j) \in \mathbb{N}^2\}_{j=1}^{N_p}$  and  $\{f_j^p \in \mathbb{R}^{D_p}\}_{j=1}^{N_p}$  respectively.  $N_v$  is larger than  $N_p$  due to many-to-one voxel-to-pillar correspondence.

With the consistent occupancy in BEV assured, for a non-empty pillar in certain X-Y coordinate, we find the corresponding voxels by matching the horizontal vector of voxel indices  $(l_i, w_i)$  with pillar indices  $p_j = (l_j, w_j)$ , and forms the  $N_v \times N_p$  voxel-pillar index matrix  $C$ , where each element  $c_{ij}$ ,  $i = 1, \dots, N_v$ ,  $j = 1, \dots, N_p$  is defined as,

$$c_{ij} = \begin{cases} 1, & l_i = l_j, w_i = w_j \\ 0, & \text{otherwise} \end{cases}. \quad (1)$$

Hence, the voxel-to-pillar sparse feature pooling and its inverse operation, feature broadcasting, are defined as,

$$\text{Pool: } f_j^{v \rightarrow p} = \text{Pool}\{f_i^v | c_{ij} = 1, \forall i\}, \quad (2)$$

$$\text{BroadCast: } f_i^{p \rightarrow v} = \{f_j^p | c_{ij} = 1, \forall j\}, \quad (3)$$

where  $\text{Pool}\{\cdot\}$  denotes element-wise max-pooling and pillar features  $f_j^p$  are identically projected onto the non-empty voxels in broadcasting.

**Sparse Voxel-Pillar Fusion.** In addition to sparse pooling and broadcasting, we further introduce the Sparse Fusion Layer (SFL), as illustrated in Fig. 2. Given the pair of sparse voxel and pillar features, SFL first calculates the voxel-pillar index matrix  $C$ . Then, sparse features engage in bidirectional interactions. In the voxel-to-pillar connection, voxel features undergo vertical aggregation by sparse max-pooling, detailed in Equ. 2, resulting in pillar-wise features with the shape  $(N_p, D_v)$ . These pillar-wise features are spatially consistent with original pillar features, while distinct in feature dimension. Next, a 2D submanifold convolution is applied to produce the pillar-wise pooled features  $\{f_j^{v \rightarrow p} \in \mathbb{R}^{D_p}\}_{j=1}^{N_p}$  with dimensions  $(N_p, D_p)$ . For pillar-to-voxel branch, pillar features are first transformed via the 2D submanifold convolution, and then broadcasted to form the voxel-wise features  $\{f_i^{p \rightarrow v} \in \mathbb{R}^{D_v}\}_{i=1}^{N_v}$  with the shape  $(N_v, D_v)$ , in Equ. 3. Notably, we adopt this asymmetrical structure with an emphasis on efficiency. Given that pillar’s feature dimension is typically several times larger than that of the voxel, it’s not cost-effective to directly transfer voxel features to the pillar dimension using the 3D submanifold convolution. We then carry out element-wise summations for feature aggregation,

$$\begin{aligned} f_i^v &= f_i^v + f_i^{p \rightarrow v}, i = 1, \dots, N_v, \\ f_j^p &= f_j^p + f_j^{v \rightarrow p}, j = 1, \dots, N_p. \end{aligned} \quad (4)$$

The SFL is incorporated after each intermediate step of the sparse Voxel-Pillar encoder to foster multi-level interactions between voxels and pillars. We highlight two primary benefits of the SFL. First, SFL offers a more straightforward approach for feature alignment by preserving the horizontal consistency in the sparse voxel-pillar encoder. Second, due to its dynamic nature, it avoids the random drops or paddings of voxels and pillars, thereby facilitating a transition to the fully sparse detection framework.

### Detection Framework

Our method serves as a fully sparse backbone that can be seamlessly incorporated into different types of detectors. We introduce  $\text{VPF}_{\text{de}}$  and  $\text{VPF}_{\text{sp}}$ , as variants of dense and sparse detectors respectively.

**VPF<sub>de</sub>.** For the dense grid-based methods, sparse voxel or pillar features are first converted to dense BEV feature maps and then processed by detection head. In this paper, we construct a dense detector named as VPF<sub>de</sub>. Given the sparse features from the proposed backbone, we first design the Dense Fusion Neck (DFN) to combine the dense feature maps from both voxel and pillar branches. DFN follows the common hierarchical structure (Lin et al. 2017a; Yan, Mao, and Li 2018) for multi-scale feature aggregation. As shown in Fig. 2, we apply convolution blocks  $\text{Block}(M, D)$  to extract dense features with  $8\times$  and  $16\times$  downsampling sizes in voxel/pillar branch separately, where  $M$  and  $D$  denote the number of convolution layers and output dimension. Next, dense voxels and pillar features with the same scale are fused by element-wise summation. Finally, we combine the different scale features via upsampling and concatenation as (Yan, Mao, and Li 2018). For the detection head, we adopt the usual center-based head (Yin, Zhou, and Krahenbuhl 2021) with IoU-Aware rectification (Hu et al. 2022) to incorporate the regression accuracy with the classification score. The final predicted score is calculated by the rectification function (Hu et al. 2022),

$$S_{\text{pred}} = S_{\text{cls}}^{1-\alpha} * \text{IoU}_{\text{pred}}^{\alpha}, \quad (5)$$

where  $S_{\text{cls}}$  is the classification score and  $\text{IoU}_{\text{pred}}$  is the IoU prediction,  $\alpha$  is the hyperparameter to balance the two.

**VPF<sub>sp</sub>.** Sparse detection frameworks are presented to avoid redundant computation (Chen et al. 2023; Fan et al. 2022a) and support long-range detection (Fan et al. 2022b). Since the proposed voxel-pillar encoder and SFL are fully sparse architectures, we could easily incorporate our hybrid representation backbone into the existing sparse detector, and therefore we present VPF<sub>sp</sub>. As demonstrated in Fig. 2, we apply additional down-sampling layers (Chen et al. 2023) for voxel and pillar branches, which obtain the  $16\times$  and  $32\times$  downsampled sparse features. These multi-scale sparse features are aggregated via the sparse height compression in each separate branch, and then combined with element-wise summation. Ultimately, the sparse head predicts objects from the sparse pillar-wise features in BEV space, and the box regression and IoU prediction paradigm stay the same as the center-based head and Equ. 5.

**Training Loss.** In our loss function, we use focal loss (Lin et al. 2017b) and L1 loss as the classification and box regression loss, noted  $\mathcal{L}_{\text{cls}}$  and  $\mathcal{L}_{\text{reg}}$  respectively. The IoU head is supervised by L1 loss and encoded by  $(2 * \text{IoU} - 0.5) \in [-1, 1]$ . We also use the Distance-IoU (DIOU) loss (Zheng et al. 2020) to further optimize the object center regression,

$$\mathcal{L}_{\text{diou}} = 1 - \text{IoU}(b, b^{gt}) + \frac{c^2(b, b^{gt})}{d^2}, \quad (6)$$

where  $\text{IoU}(b, b^{gt})$  refers to the 3D IoU between predicted box  $b$  and corresponding ground truth  $b^{gt}$ ,  $c$  denotes the center offset of  $b$  and  $b^{gt}$ , and  $d$  refers to the diagonal distance of minimum enclosing cuboid covering both  $b$  and  $b^{gt}$ .

Combined with the IoU prediction loss  $\mathcal{L}_{\text{iou}}$  and DIOU Loss  $\mathcal{L}_{\text{diou}}$ , the overall loss is formed as,

$$\mathcal{L} = \mathcal{L}_{\text{cls}} + \mathcal{L}_{\text{iou}} + \gamma(\mathcal{L}_{\text{diou}} + \mathcal{L}_{\text{reg}}), \quad (7)$$

where  $\gamma$  is the loss weight similar to (Guangsheng Shi 2022).

## Experiments

### Datasets

**Waymo Open Dataset (WOD)** (Sun et al. 2020) consists of 798 training, 202 validation, and 150 testing sequences with 200K annotated frames. The detection range covers the area of  $[-75m, -75m, 75m, 75m]$ , and evaluation metrics are average precision (AP) and average precision weighted by heading (APH). Test samples are split into two difficulties, LEVEL 1 for objects with more than 5 inside points and LEVEL 2 for objects with at least 1 point.

**nuScenes Dataset** (Caesar et al. 2020) contains 700 training, 150 validation, and 150 testing sequences, with 34K annotated keyframes and 10 categories. The evaluation metric is nuScenes detection score (NDS), which is a weighted sum of mAP and other true positive metrics.

### Implementation Details

**Network Architecture.** Our VPF builds on the sparse backbone, which consists of the sparse voxel-pillar encoder with 4 convolution blocks and Sparse Fusion Layers (SFL). In the VPF<sub>de</sub> model, the feature dimensions for the voxel and pillar branches in each convolution block are set at  $[16, 32, 64, 64]$  and  $[32, 64, 128, 256]$ , respectively. Meanwhile, for the VPF<sub>sp</sub> model, these dimensions are  $[16, 32, 64, 128]$  and  $[32, 64, 128, 256]$ .

Additionally, for the Dense Fusion Neck (DFN) applied in our dense detectors, the number of convolution layers  $M$  and the output dimension  $D$  are set as 5, 128. As for our sparse detectors, feature dimensions in additional downsampled layers (*i.e.*,  $16\times$  and  $32\times$  convolution blocks) are set as  $[128, 128]$ ,  $[256, 256]$  for the voxel and pillar branches. Note that the voxel features are upsampled to  $256D$  during the final voxel-pillar feature fusion.

**Model Setting.** In data preprocessing, point clouds are divided into voxels and pillars with the size of  $[0.1m, 0.1m, 0.15m]$ ,  $[0.1m, 0.1m]$  for WOD, and  $[0.075m, 0.075m, 0.2m]$ ,  $[0.075m, 0.075m]$  for nuScenes dataset. We apply the IoU thresholds (0.8, 0.55, 0.55) and  $\alpha$  in Equ. 5 as (0.68, 0.71, 0.65) for Vehicle, Pedestrian, and Cyclist in WOD. In nuScenes, the IoU and  $\alpha$  are set as 0.2 and 0.5 for all categories.

**Training Details.** All models are trained on 4 NVIDIA 3090 GPUs with Adam optimizer. The learning rates are set as  $1e-3$  and  $3e-3$  for nuScenes and Waymo dataset, separately. We use the common data augmentation strategies (Shi et al. 2022) and ground-truth sampling fade strategy as (Wang et al. 2023). For WOD, we train VPF<sub>de</sub> and VPF<sub>sp</sub> with batch size 16 for 15 epochs and 12 epochs separately. For nuScenes, our models are trained with the same batch size for 20 epochs.

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

**Waymo Open Dataset.** We compare our VPF variants with existing single-frame methods on WOD *val* set. As shown in Tab. 1, for the single-stage framework, our method

Methods	Stages	LEVEL 2		Veh. (LEVEL1/LEVEL2)				Ped. (LEVEL1/LEVEL2)				Cyc. (LEVEL1/LEVEL2)			
		mAP/mAPH	AP	APH	AP	APH	AP	APH	AP	APH	AP	APH	AP	APH	
SECOND (2018)	Single	61.0/57.2	72.3	71.7	63.9	63.3	68.7	58.2	60.7	51.3	60.6	59.3	58.3	57.0	
MVF (2020b)	Single	-	62.9	-	-	-	65.3	-	-	-	-	-	-	-	
PointPillars (2019)	Single	62.8/57.8	72.1	71.5	63.6	63.1	70.6	56.7	62.8	50.3	64.4	62.3	61.9	59.9	
Pillar-OD (2020)	Single	-	69.8	-	-	-	72.5	-	-	-	-	-	-	-	
AFDetV2-Lite (2022)	Single	71.0/68.8	77.6	77.1	69.7	69.2	80.2	74.6	72.2	67.0	73.7	72.7	71.0	70.1	
IA-SSD (2022)	Single	62.3/58.1	70.5	69.7	61.6	61.0	69.4	58.5	60.3	50.7	67.7	65.3	65.0	62.7	
CenterFormer (2022)	Single	71.1/68.9	75.0	74.4	69.9	69.4	78.6	73.0	73.6	68.3	72.3	71.3	69.8	68.8	
Centerpoint <sup>†</sup> (2021)	Single	71.7/69.4	77.7	77.1	69.7	69.2	80.9	75.1	72.6	67.2	75.6	74.4	72.8	71.7	
PillarNet (2022)	Single	71.0/68.5	79.1	78.6	70.9	70.5	80.6	74.0	72.3	66.2	72.3	71.2	69.7	68.7	
VoxelNeXt2D (2023)	Single	70.9/68.2	77.9	77.5	69.7	69.2	80.2	73.5	72.2	65.9	73.3	72.2	70.7	69.6	
VoxelNeXt <sub>K3</sub> (2023)	Single	72.2/70.1	78.2	77.7	69.9	69.4	81.5	76.3	73.5	68.6	76.1	74.9	73.3	72.2	
GD-MAE (2023)	Single	72.9/70.4	79.4	78.9	70.9	70.5	82.2	75.9	74.8	68.8	75.8	74.8	73.0	72.0	
DSVT (2023)	Single	73.2/71.0	79.3	78.8	70.9	70.5	82.8	77.0	<b>75.2</b>	69.8	76.4	75.4	73.6	72.7	
VPF <sub>sp</sub> -Lite	Single	73.1/70.9	79.1	78.6	71.1	70.6	82.1	76.7	74.5	69.4	76.5	75.5	73.6	72.7	
VPF <sub>de</sub> -Lite		73.6/71.4	80.2	79.7	71.9	71.5	82.5	76.9	74.8	69.4	77.1	76.0	74.2	73.2	
VPF <sub>sp</sub>		<b>73.6/71.6</b>	<b>80.2</b>	<b>79.8</b>	<b>71.7</b>	<b>71.3</b>	<b>82.5</b>	<b>77.5</b>	74.9	<b>70.1</b>	<b>77.2</b>	<b>76.2</b>	<b>74.3</b>	<b>73.3</b>	
VPF <sub>de</sub>		<b>73.9/71.7</b>	<b>80.5</b>	<b>80.0</b>	<b>72.3</b>	<b>71.9</b>	<b>82.8</b>	<b>77.3</b>	<b>75.1</b>	<b>69.9</b>	<b>77.2</b>	<b>76.0</b>	<b>74.3</b>	<b>73.2</b>	
Voxel RCNN (2021)		Two	-	75.6	-	66.6	-	-	-	-	-	-	-	-	-
PartA <sup>2</sup> (2020b)	Two	66.9/63.8	77.1	76.5	68.5	68.0	75.2	66.9	66.2	58.6	68.6	67.4	66.1	64.9	
LiDAR R-CNN (2021)	Two	65.8/61.3	76.0	75.5	68.3	67.9	71.2	58.7	63.1	51.7	68.6	66.9	66.1	64.4	
RSN (2021)	Two	-	75.1	74.6	66.0	65.5	77.8	72.7	68.3	63.7	-	-	-	-	
PV-RCNN (2020a)	Two	66.8/63.3	77.5	76.9	69.0	68.4	75.0	65.6	66.0	57.6	67.8	66.4	65.4	64.0	
SST-TS-1f (2022a)	Two	-	76.2	75.8	68.0	67.6	81.4	74.0	72.8	65.9	-	-	-	-	
PV-RCNN++ (2022)	Two	71.7/69.5	79.3	78.8	70.6	70.2	81.3	76.3	73.2	68.0	73.7	72.7	71.2	70.2	
FSD <sub>spconv</sub> (2022b)	Two	72.9/70.8	79.2	78.8	70.5	70.1	82.6	77.3	73.9	69.1	77.1	76.0	74.4	73.3	
GD-MAE (2023)	Two	74.1/71.6	80.2	79.8	72.4	72.0	83.1	76.7	75.5	69.4	77.2	76.2	74.4	73.4	
DSVT (2023)	Two	74.3/72.1	80.2	79.7	72.0	71.6	83.7	78.0	76.1	70.7	<b>77.8</b>	<b>76.8</b>	<b>74.9</b>	<b>73.9</b>	
VPF <sub>sp</sub> -TS	Two	<b>74.7/72.5</b>	<b>81.4</b>	<b>80.9</b>	<b>73.1</b>	<b>72.7</b>	<b>83.7</b>	<b>78.5</b>	<b>76.2</b>	<b>71.3</b>	77.3	76.3	74.4	73.5	
VPF <sub>de</sub> -TS		<b>75.1/72.9</b>	<b>81.4</b>	<b>80.9</b>	<b>73.6</b>	<b>73.2</b>	<b>84.0</b>	<b>78.6</b>	<b>76.6</b>	<b>71.3</b>	<b>77.8</b>	<b>76.7</b>	<b>75.2</b>	<b>74.1</b>	

Table 1: Single-frame performance comparison on the WOD *val* set, without test-time augmentation or model ensemble. <sup>†</sup>: single-stage with IoU prediction setting. Top-2 results are highlighted in bold for different stages.

achieves promising performance among all single-stage methods. Specifically, both VPF<sub>sp</sub> and VPF<sub>de</sub> are superior to dense or sparse baselines (i.e., Centerpoint, PillarNet, and VoxelNeXts) with remarkable margins. Moreover, our method outperforms recent transformer-based GD-MAE (Yang et al. 2023) and DSVT (Wang et al. 2023) with more efficient training schedule (4 3090 vs. 8 A100 GPUs) or MAE pre-training) and inference speed. We also provide lightweight versions of our framework, VPF<sub>sp</sub>-Lite and VPF<sub>de</sub>-Lite, which reduce the convolution layer and channel dimension for a lower computing budget. They achieve competitive results while running at 60 ( $\pm 5$ ) ms and 58 ( $\pm 4$ ) ms, as demonstrated in Fig. 4. To compare with two-stage detectors, we simply adopt CT3D (Sheng et al. 2021) as the second stage, obtaining competitive performance as well. While the performance of VPF<sub>sp</sub> is inferior to VPF<sub>de</sub> in the Vehicle category since it’s harder to estimate the object center estimation by sparse architecture, VPF<sub>sp</sub> is more computationally efficient in long-range detection.

**nuScenes Dataset.** We evaluate and compare VPF with LiDAR-based methods on the nuScenes *test* set. As depicted in Tab. 2, our approach surpasses previous methods in both mAP and NDS metrics with a considerable improvement. More specifically, VPF<sub>de</sub> achieves 67.0 mAP and 72.7 NDS in the single model setting, with the NDS boosts to 73.8 using double-flip test. In addition, VPF improves the perfor-




Figure 4: Performance vs. inference latency on WOD *val* set. Tested on a single 3090 GPU with batch size 1.

mance of vertical sensitive categories (e.g., Pedestrian, Bicycle) by a large margin while maintaining competitive results on other classes (e.g., Truck). We attribute this to our robust vertical representation. Given that slender objects, like pedestrians, present richer vertical information than horizontal, our hybrid voxel-pillar encoding paradigm is helpful in establishing effective representations for such entities.

## Ablation Studies

In this section, we conduct experiments on WOD *val* set to validate the effectiveness of voxel-pillar hybrid representation. Moreover, we investigate the impact of varying point cloud representations and architectures on 3D detection. All

Method	NDS	mAP	Car	Truck	Bus	Trailer	C.V.	Ped	Mot	Byc	T.C.	Bar
PointPillars (2019)	45.3	30.5	68.4	23.0	28.2	23.4	4.1	59.7	27.4	1.1	30.8	38.9
3DSSD (2020)	56.4	42.6	81.2	47.2	61.4	30.5	12.6	70.2	36.0	8.6	31.1	47.9
CBGS (2019)	63.3	52.8	81.1	48.5	54.9	42.9	10.5	80.1	51.5	22.3	70.9	65.7
HotSpotNet (2020b)	66.0	59.3	83.1	50.9	56.4	53.3	23.0	81.3	63.5	36.6	73.0	71.6
CVCNET (2020a)	66.6	58.2	82.6	49.5	59.4	51.1	16.2	83.0	61.8	38.8	69.7	69.7
CenterPoint (2021)	65.5	58.0	84.6	51.0	60.2	53.2	17.5	83.4	53.7	28.7	76.7	70.9
CenterPoint <sup>†</sup> (2021)	67.3	60.3	85.2	53.5	63.6	56.0	20.0	84.6	59.5	30.7	78.4	71.1
AFDetV2-Lite (2022)	68.5	62.4	86.3	54.2	62.5	58.9	26.7	85.8	63.8	34.3	80.1	71.0
VISTA-OHS <sup>†</sup> (2022)	69.8	63.0	84.4	55.1	63.7	54.2	25.1	84.6	70.0	45.4	78.5	71.1
Focals Conv (2022a)	70.0	63.8	<u>86.7</u>	<u>56.3</u>	<b>67.7</b>	59.5	23.8	87.5	64.5	36.3	81.4	74.1
LargeKernel3D (2022b)	70.5	65.3	85.9	55.3	66.2	60.2	26.8	85.6	72.5	46.6	80.0	74.3
PillarNet-34 <sup>†</sup> (2022)	71.4	66.0	<b>87.6</b>	57.5	63.6	63.1	27.9	87.3	70.1	42.3	83.3	77.2
VoxelNeXt <sup>†</sup> (2023)	71.4	66.2	85.3	55.7	66.2	57.2	29.8	86.5	75.2	48.8	80.7	76.1
LinK (Lu et al. 2023)	71.0	66.3	86.1	55.7	65.7	62.1	30.9	85.8	73.5	47.5	80.4	75.5
VPF <sub>de</sub>	<u>72.7</u>	<u>67.0</u>	85.8	55.1	63.5	<u>62.1</u>	<u>33.3</u>	<u>87.6</u>	<u>72.5</u>	<u>48.6</u>	<u>82.9</u>	<b>78.2</b>
VPF <sub>de</sub> <sup>†</sup>	<b>73.8</b>	<b>68.6</b>	86.3	<b>56.8</b>	66.1	<b>64.5</b>	<b>34.6</b>	<b>88.3</b>	<b>75.8</b>	<b>51.9</b>	<b>84.2</b>	77.6

Table 2: Performance comparison for 3D object detection on the nuScenes *test* set. <sup>†</sup> indicates the flipping test is used. We highlight the best results in bold and the best non-ensemble results with the underline.

Voxel	Pillar	DFN	SFL	Lat.	LEVEL 2 APH			
					Veh.	Ped.	Cyc.	mAPH
✓				51ms	67.8	66.5	68.2	67.5
	✓			57ms	68.5	64.1	66.2	66.3
✓	✓			102ms	69.1	66.7	69.2	68.3
✓	✓	✓		69ms	69.2	66.0	68.9	68.0
✓	✓	✓	✓	75ms	<b>70.2</b>	<b>68.2</b>	<b>70.7</b>	<b>69.7</b>

Table 3: Effects of hybrid encoding paradigm. CenterPoint and PillarNet are used as voxel/pillar-based baselines.

Exp.	Param.	LEVEL 2 APH			
		Veh.	Ped.	Cyc.	mAPH
Scaled PillarNet	92.4M	69.0	64.8	67.8	67.2
Scaled CenterPoint	93.2M	69.2	67.0	70.1	68.8
Scaled CenterPoint <sup>†</sup>	96.8M	69.2	67.2	70.5	69.0
Voxel-Voxel	67.1M	68.9	66.2	69.5	68.2
Pillar-Pillar	111.1M	69.3	64.5	67.1	66.9
VPF <sub>de</sub>	92.4M	<b>70.2</b>	<b>68.2</b>	<b>70.7</b>	<b>69.7</b>

Table 4: Effects of architecture and scalability. <sup>†</sup>: deploy the vertical residual for feature aggregation.

models are trained on the full train set for 7 epochs, with comparisons based on APH of LEVEL 2 difficulty.

**Effect of Voxel-Pillar Encoding Paradigm.** We take VPF<sub>de</sub> as the framework to illustrate the effects of hybrid representation. For voxel-/pillar-based baselines, we deploy the single-stage Centerpoint (Yin, Zhou, and Krahenbuhl 2021) and PillarNet (Guangsheng Shi 2022). To ensure a fair comparison, we add the IoU prediction to CenterPoint as (Guangsheng Shi 2022). First, we compare the single representations with the naive model ensemble, in which the predictions from CenterPoint and PillarNet are combined via NMS, as shown in the 3<sup>rd</sup> row. The ensemble strategy consistently improves the detection accuracy. Then, as presented in the 4<sup>th</sup> row, we deploy the proposed voxel-pillar encoder with Dense Fusion Neck (DFN). This late fusion

Exp.	Connections	LEVEL 2 APH			
		Veh.	Ped.	Cyc.	mAPH
1	Voxel-to-Pillar	70.0	67.3	70.4	69.2
2	Pillar-to-Voxel	69.4	66.3	68.7	68.1
3	Bidirectional	<b>70.2</b>	<b>68.2</b>	<b>70.7</b>	<b>69.7</b>

Table 5: Ablation on different interaction paradigm.

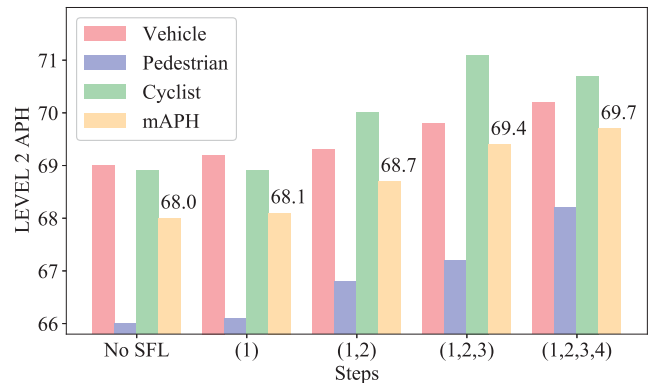


Figure 5: Ablation on deploying steps for SFL.

strategy achieves comparable results while being more efficient due to the unified encoding process of voxels and pillars. Finally, we deploy the Sparse Fusion Layer (SFL), obtaining 1.7 mAPH improvement with a minor increase in latency, which demonstrates the importance of voxel-pillar interaction within the sparse backbone.

**Effects of Model Scale and Architecture.** Since our model has a different scale and dual-branch pipeline, we ablate the effectiveness of model scale and architecture in Tab. 4. First, we construct the scaled PillarNet and CenterPoint by compositely increasing the number of layers and feature dimensions while keeping the BEV map resolution

unchanged, which makes them comparable with VPF in parameter scales. As shown in the 1<sup>st</sup> – 2<sup>nd</sup> rows, our hybrid voxel-pillar encoding paradigm is significantly prior to the single representation counterpart. Then, we explore the validity of vertical feature aggregation in the 3<sup>rd</sup> row. For each stage within the voxel backbone, we introduce a vertical residual. This residual is constructed by aggregating voxel features vertically via sparse pooling. The gathered features serve as residuals and are added to the original voxel features. This vertical residual leads to an improvement on slender objects, *i.e.*, Pedestrian and Cyclist, while it still has a large gap compared with our approach. In addition, we investigate different dual-branch pipelines to explore the impact from wider or multi-branch networks (Chen et al. 2017; Xie et al. 2017). The results of voxel-only or pillar-only frameworks in 4<sup>th</sup> – 5<sup>th</sup> rows of Tab. 4 indicates although multi-branch architecture could boost detection performance, the improvements in our work are mainly from the synergy of voxel and pillar.

**Sparse Fusion Layer.** We study the effects of the SFL from two perspectives. First, we verify the bidirectional connection of SFL. Given its voxel-to-pillar and pillar-to-voxel connections, we compare the performance of the interaction paradigm in Tab. 5. It indicates that each individual connection could obtain clear improvement, and applying the bidirectional connection achieves the best results. Second, we analyze the deploying intermediate steps of SFL for the voxel-pillar encoder in Fig. 5. The performance increases with the module stacking in each step. As SFL amplifies the local context for the pillar branch and infuses the vertical semantics into the voxel branch, it primarily functions in the latter steps of the sparse backbone, helping to enrich the vertical information during sparse convolutions.

## Conclusion

Inspired by the distinctions and constraints of voxel and pillar encoding paradigms, we propose the VPF, a hybrid detection framework that combines the strengths of both. VPF utilizes a sparse voxel-pillar encoder for consistent dual-branch feature extraction. Further, our Sparse Fusion Layer facilitates bidirectional interaction between the sparse voxel and pillar features, jointly enhancing vertical representation learning. Notably, our proposed components can be seamlessly integrated into both dense and sparse detectors, yielding promising performance and real-time inference speeds.

**Limitations.** VPF relies on consistent occupancy in the bird’s eye view for element-wise fusion, requiring equal horizontal partitioning for voxels and pillars. For future work, we will explore the asynchronous voxel-pillar fusion strategies (e.g., varying resolution or multi-stride fusion) for more flexible point cloud representation learning.

## Acknowledgements

This paper is supported by the National Natural Science Foundation of China (No. 62088102). We would like to thank four anonymous reviewers for their constructive comments.

## References

- Bewley, A.; Sun, P.; Mensink, T.; Anguelov, D.; and Sminchisescu, C. 2021. Range Conditioned Dilated Convolutions for Scale Invariant 3D Object Detection. In *CoRL*.
- 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 *CVPR*.
- Chen, Q.; Sun, L.; Cheung, E.; and Yuille, A. L. 2020a. Every view counts: Cross-view consistency in 3d object detection with hybrid-cylindrical-spherical voxelization. In *Neurips*.
- Chen, Q.; Sun, L.; Wang, Z.; Jia, K.; and Yuille, A. 2020b. Object as hotspots: An anchor-free 3d object detection approach via firing of hotspots. In *ECCV*.
- Chen, X.; Ma, H.; Wan, J.; Li, B.; and Tian, X. 2016. Multi-View 3D Object Detection Network for Autonomous Driving. In *CVPR*.
- Chen, Y.; Li, J.; Xiao, H.; Jin, X.; Yan, S.; and Feng, J. 2017. Dual path networks. In *NeurIPS*.
- Chen, Y.; Li, Y.; Zhang, X.; Sun, J.; and Jia, J. 2022a. Focal sparse convolutional networks for 3d object detection. In *CVPR*.
- Chen, Y.; Liu, J.; Qi, X.; Zhang, X.; Sun, J.; and Jia, J. 2022b. Scaling up kernels in 3d cnns. *arXiv preprint*.
- Chen, Y.; Liu, J.; Zhang, X.; Qi, X.; and Jia, J. 2023. VoxelNeXt: Fully Sparse VoxelNet for 3D Object Detection and Tracking. In *CVPR*.
- Deng, J.; Shi, S.; Li, P.; Zhou, W.; Zhang, Y.; and Li, H. 2021. Voxel R-CNN: Towards High Performance Voxel-based 3D Object Detection. In *AAAI*.
- Deng, S.; Liang, Z.; Sun, L.; and Jia, K. 2022. Vista: Boosting 3d object detection via dual cross-view spatial attention. In *CVPR*.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houlsby, N. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. In *ICLR*.
- Fan, L.; Pang, Z.; Zhang, T.; Wang, Y.-X.; Zhao, H.; Wang, F.; Wang, N.; and Zhang, Z. 2022a. Embracing single stride 3d object detector with sparse transformer. In *CVPR*.
- Fan, L.; Wang, F.; Wang, N.; and Zhang, Z. 2022b. Fully Sparse 3D Object Detection. In *NeurIPS*.
- Fan, L.; Xiong, X.; Wang, F.; long Wang, N.; and Zhang, Z. 2021. RangeDet: In Defense of Range View for LiDAR-based 3D Object Detection. In *ICCV*.
- Feichtenhofer, C.; Fan, H.; Malik, J.; and He, K. 2018. Slow-Fast Networks for Video Recognition. In *ICCV*.
- Fey, M.; and Lenssen, J. E. 2019. Fast Graph Representation Learning with PyTorch Geometric. In *ICLRW*.
- Graham, B. 2015. Sparse 3D convolutional neural networks. In *BMVC*.
- Graham, B.; and van der Maaten, L. 2017. Submanifold Sparse Convolutional Networks. *arXiv preprint*.

- Guangsheng Shi, C. M., Ruifeng Li. 2022. PillarNet: Real-Time and High-Performance Pillar-based 3D Object Detection. In *ECCV*.
- He, C.-H.; Li, R.; Li, S.; and Zhang, L. 2022. Voxel Set Transformer: A Set-to-Set Approach to 3D Object Detection from Point Clouds. In *CVPR*.
- Hu, Y.; Ding, Z.; Ge, R.; Shao, W.; Huang, L.; Li, K.; and Liu, Q. 2022. Afdetv2: Rethinking the necessity of the second stage for object detection from point clouds. In *AAAI*.
- Kirillov, A.; Usunier, N.; Carion, N.; Zagoruyko, S.; Synnaeve, G.; Massa, F.; Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-End Object Detection with Transformers. In *ECCV*.
- Lang, A. H.; Vora, S.; Caesar, H.; Zhou, L.; Yang, J.; and Beijbom, O. 2019. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*.
- Li, Y.; Chen, Y.; Wang, N.; and Zhang, Z. 2019. Scale-Aware Trident Networks for Object Detection. In *ICCV*.
- Li, Y.; Yu, A. W.; Meng, T.; Caine, B.; Ngiam, J.; Peng, D.; Shen, J.; Wu, B.; Lu, Y.; Zhou, D.; Le, Q. V.; Yuille, A.; and Tan, M. 2022. DeepFusion: Lidar-Camera Deep Fusion for Multi-Modal 3D Object Detection. In *CVPR*.
- Li, Z.; Wang, F.; and Wang, N. 2021. LiDAR R-CNN: An Efficient and Universal 3D Object Detector. In *CVPR*.
- Liang, M.; Yang, B.; Chen, Y.; Hu, R.; and Urtasun, R. 2019. Multi-Task Multi-Sensor Fusion for 3D Object Detection. In *CVPR*.
- Lin, T.-Y.; Dollár, P.; Girshick, R. B.; He, K.; Hariharan, B.; and Belongie, S. J. 2017a. Feature Pyramid Networks for Object Detection. In *CVPR*.
- Lin, T.-Y.; Goyal, P.; Girshick, R.; He, K.; and Dollár, P. 2017b. Focal loss for dense object detection. In *ICCV*.
- Lu, T.; Ding, X.; Liu, H.; Wu, G.; and Wang, L. 2023. LinK: Linear Kernel for LiDAR-based 3D Perception. In *CVPR*.
- Meyer, G. P.; Laddha, A.; Kee, E.; Vallespi-Gonzalez, C.; and Wellington, C. K. 2019. Lasernet: An efficient probabilistic 3d object detector for autonomous driving. In *CVPR*.
- Pang, S.; Morris, D. H.; and Radha, H. 2020. CLOCs: Camera-LiDAR Object Candidates Fusion for 3D Object Detection. In *IROS*.
- Qi, C.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *CVPR*.
- Qi, C.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. In *NeurIPS*.
- Qi, C. R.; Litany, O.; He, K.; and Guibas, L. J. 2019. Deep Hough Voting for 3D Object Detection in Point Clouds. In *ICCV*.
- Ren, S.; He, K.; Girshick, R. B.; and Sun, J. 2015. Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks. *IEEE TPAMI*.
- Sheng, H.; Cai, S.; Liu, Y.; Deng, B.; Huang, J.; Hua, X.; and Zhao, M. 2021. Improving 3D Object Detection with Channel-wise Transformer. In *ICCV*.
- Shi, S.; Guo, C.; Jiang, L.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020a. Pv-rcnn: Point-voxel feature set abstraction for 3d object detection. In *CVPR*.
- Shi, S.; Jiang, L.; Deng, J.; Wang, Z.; Guo, C.; Shi, J.; Wang, X.; and Li, H. 2022. PV-RCNN++: Point-voxel feature set abstraction with local vector representation for 3D object detection. *IJCV*.
- Shi, S.; Wang, X.; and Li, H. 2018. PointRCNN: 3D Object Proposal Generation and Detection from Point Cloud. In *CVPR*.
- Shi, S.; Wang, Z.; Shi, J.; Wang, X.; and Li, H. 2020b. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE TPAMI*.
- Sun, P.; Kretschmar, H.; Dotiwalla, X.; Chouard, A.; Patnaik, V.; Tsui, P.; Guo, J.; Zhou, Y.; Chai, Y.; Caine, B.; et al. 2020. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*.
- Sun, P.; Wang, W.; Chai, Y.; Elsayed, G.; Bewley, A.; Zhang, X.; Sminchisescu, C.; and Anguelov, D. 2021. Rsn: Range sparse net for efficient, accurate lidar 3d object detection. In *CVPR*.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L.; and Polosukhin, I. 2017. Attention is All you Need. In *NeurIPS*.
- Wang, H.; Shi, C.; Shi, S.; Lei, M.; Wang, S.; He, D.; Schiele, B.; and Wang, L. 2023. Dsvt: Dynamic sparse voxel transformer with rotated sets. In *CVPR*.
- Wang, Y.; Fathi, A.; Kundu, A.; Ross, D. A.; Pantofaru, C.; Funkhouser, T.; and Solomon, J. 2020. Pillar-based object detection for autonomous driving. In *ECCV*.
- Xie, S.; Girshick, R.; Dollár, P.; Tu, Z.; and He, K. 2017. Aggregated residual transformations for deep neural networks. In *CVPR*.
- Xin, X.; Wang, J.; Xie, R.; Zhou, S.; Huang, W.; and Zheng, N. 2019. Semi-supervised person re-identification using multi-view clustering. *PR*.
- Yan, Y.; Mao, Y.; and Li, B. 2018. SECOND: Sparsely Embedded Convolutional Detection. *Sensors*.
- Yang, B.; Luo, W.; and Urtasun, R. 2018. PIXOR: Real-time 3D Object Detection from Point Clouds. In *CVPR*.
- Yang, H.; He, T.; Liu, J.; Chen, H.; Wu, B.; Lin, B.; He, X.; and Ouyang, W. 2023. GD-MAE: generative decoder for MAE pre-training on lidar point clouds. In *CVPR*.
- Yang, Z.; Sun, Y.; Liu, S.; and Jia, J. 2020. 3DSSD: Point-based 3D Single Stage Object Detector. In *CVPR*.
- Yang, Z.; Sun, Y.; Liu, S.; Shen, X.; and Jia, J. 2019. STD: Sparse-to-Dense 3D Object Detector for Point Cloud. In *ICCV*.
- Yin, T.; Zhou, X.; and Krahenbuhl, P. 2021. Center-based 3d object detection and tracking. In *CVPR*.
- Zhang, Y.; Hu, Q.; Xu, G.; Ma, Y.; Wan, J.; Guo, Y.; and Prediction, C. 2022. Not All Points Are Equal: Learning Highly Efficient Point-based Detectors for 3D LiDAR Point Clouds. In *CVPR*.

- Zheng, W.; Tang, W.; Chen, S.; Jiang, L.; and Fu, C.-W. 2021. CIA-SSD: Confident IoU-Aware Single-Stage Object Detector From Point Cloud. In *AAAI*.
- Zheng, Z.; Wang, P.; Liu, W.; Li, J.; Ye, R.; and Ren, D. 2020. Distance-IoU loss: Faster and better learning for bounding box regression. In *AAAI*.
- Zhou, S.; Wang, J.; Wang, L.; Wan, X.; Hui, S.; and Zheng, N. 2023. Inverse Adversarial Diversity Learning for Network Ensemble. *IEEE TNNLS*.
- Zhou, S.; Wang, J.; Wang, L.; Zhang, J.; Wang, F.; Huang, D.; and Zheng, N. 2020a. Hierarchical and interactive refinement network for edge-preserving salient object detection. *IEEE TIP*.
- Zhou, Y.; Sun, P.; Zhang, Y.; Anguelov, D.; Gao, J.; Ouyang, T.; Guo, J.; Ngiam, J.; and Vasudevan, V. 2020b. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *CoRL*.
- Zhou, Y.; and Tuzel, O. 2018. VoxelNet: End-to-End Learning for Point Cloud Based 3D Object Detection. In *CVPR*.
- Zhou, Z.; Zhao, X.; Wang, Y.; Wang, P.; and Foroosh, H. 2022. Centerformer: Center-based transformer for 3d object detection. In *ECCV*.
- Zhu, B.; Jiang, Z.; Zhou, X.; Li, Z.; and Yu, G. 2019. Class-balanced grouping and sampling for point cloud 3d object detection. *arXiv preprint*.