

# OcTr: Octree-based Transformer for 3D Object Detection

Chao Zhou<sup>1,2</sup>, Yanan Zhang<sup>1,2</sup>, Jiaxin Chen<sup>2</sup>, Di Huang<sup>1,2,3\*</sup>

<sup>1</sup>State Key Laboratory of Software Development Environment, Beihang University, Beijing, China

<sup>2</sup>School of Computer Science and Engineering, Beihang University, Beijing, China

<sup>3</sup>Hangzhou Innovation Institute, Beihang University, Hangzhou, China

{zhouchaobeing, zhangyanan, jiaxinchen, dhuang}@buaa.edu.cn

## Abstract

A key challenge for LiDAR-based 3D object detection is to capture sufficient features from large scale 3D scenes especially for distant or/and occluded objects. Albeit recent efforts made by Transformers with the long sequence modeling capability, they fail to properly balance the accuracy and efficiency, suffering from inadequate receptive fields or coarse-grained holistic correlations. In this paper, we propose an **Octree-based Transformer**, named **OcTr**, to address this issue. It first constructs a dynamic octree on the hierarchical feature pyramid through conducting self-attention on the top level and then recursively propagates to the level below restricted by the octants, which captures rich global context in a coarse-to-fine manner while maintaining the computational complexity under control. Furthermore, for enhanced foreground perception, we propose a hybrid positional embedding, composed of the semantic-aware positional embedding and attention mask, to fully exploit semantic and geometry clues. Extensive experiments are conducted on the Waymo Open Dataset and KITTI Dataset, and OcTr reaches newly state-of-the-art results.

## 1. Introduction

3D object detection from point clouds has received extensive attention during the past decade for its ability to provide accurate and stable recognition and localization in autonomous driving perception systems. In this task, feature learning plays a very fundamental and crucial role; yet it is rather challenging due to not only the disordered and sparse nature of data sampling, but also to insufficient acquisition under occlusion or at a distance. To address this issue, many methods have been proposed, which can be taxonomized into two major classes, *i.e.* grid-based and point-based. The former first regularize point clouds into multi-view images or voxels and then apply 2D or 3D CNNs to build shape rep-

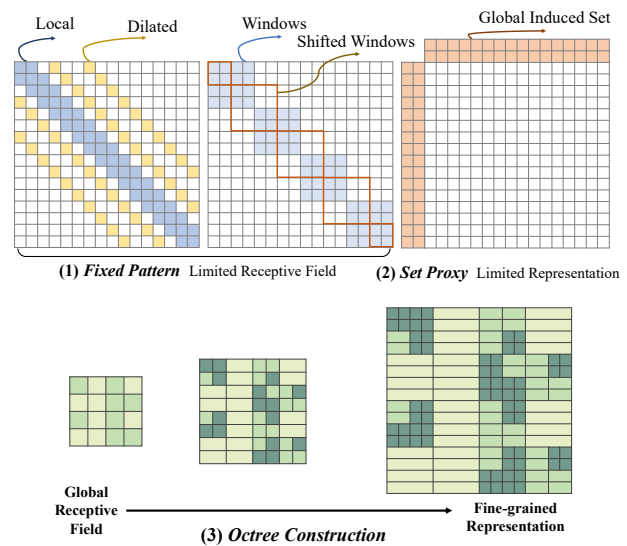


Figure 1. Illustration of three sparsification strategies of attention matrices. Fixed pattern (1) narrows receptive fields and set proxy (2) discards elaborate correlations. The proposed octree construction (3) keeps the global receptive field in a coarse-grained manner while maintaining fine-grained representations.

resentations [4, 52], while the latter directly conduct MLP based networks such as PointNet++ [33] and DGCNN [50] on original points for geometry description [32, 40, 42, 60]. Unfortunately, they fail to capture necessary context information through the small receptive fields in the deep models, leading to limited results.

Witnessing the recent success of Transformers in NLP, many studies have investigated and extended such architectures for 3D vision [24, 29, 58, 61]. Transformers are reputed to model long-range dependencies, delivering global receptive fields, and to be suitable for scattered inputs of arbitrary sizes. Meanwhile, in contrast to those static weights that are learned in convolutions, Transformers dynamically aggregate the input features according to the relationships

\*indicates the corresponding author.

between tokens. Regarding the case in 3D object detection, compared to point-based Transformers [11, 61], voxel-based ones show the superiority in efficiency. However, they tend to suffer heavy computations when dealing with large scale scenes because of the quadratic complexity of Transformers, with the underlying dilemma between the grid size and the grid amount in voxelization. Taking the KITTI dataset as an example, it is unrealistic for Transformers to operate on the feature map with the spatial shape of  $200 \times 176 \times 5$ , which is commonly adopted in most of the detection heads [38, 46, 52, 55].

More recently, there have appeared an influx of efficient self-attention model variants that attempt to tackle long sequences as input. They generally sparsify the attention matrix by fixed patterns [7, 23, 34], learned patterns [21, 45] or a combination of them [1, 56]. Fixed patterns chunk the input sequence into blocks of local windows or dilation windows, whilst learned patterns determine a notion of token relevance and eliminate or cluster outliers. Specific to 3D object detection from point clouds, VoTr [24] modifies self-attention with pre-defined patterns including local windows and stride dilation ones in a sparse query manner, and the dilation mechanism enlarges the receptive field by sampling attending tokens in a radius. SST [9] splits input tokens into non-overlapping patterns in a block-wise way and enables window shifting to capture cross-window correlation. Despite some improvements reported, they both only achieve bigger local receptive fields rather than the expected global ones, and computations still increase rapidly with the expansion of receptive fields.

Another alternative on self-attention is to take advantage of a proxy memory bank which has the access to the entire sequence tokens [1, 2, 56]. By using a small number of induced proxies to compress the whole sequence, it diffuses the global context efficiently. VoxSet [12] adapts Set Transformer [19] to 3D object detection and exploits an induced set to model a set-to-set point cloud translation. With the help of the compressed global proxies and Conv-FFN, it obtains a global receptive field; nevertheless, as they admit, it is sub-optimal to set only a few latent codes as proxies for a large 3D scene, prone to impairing the representation of different point cloud structures and their correlations. Therefore, there remains space for a stronger solution.

In this paper, we present a novel Transformer network, namely **Octree-based Transformer (OcTr)**, for 3D object detection. We firstly devise an octree-based learnable sparse pattern, *i.e.* *OctAttn*, which meticulously and efficiently encodes point clouds of scenes as shown in Fig. 1. The *OctAttn* module constructs a feature pyramid by gathering and applies self-attention to the top level of the feature pyramid to select the most relevant tokens, which are deemed as the octants to be divided in the subsequent. When propagating to the level below, the key/value inputs are restricted by the

octants from the top. Through recursively conducting this process, *OctAttn* captures rich global context features by a global receptive field in a coarse-to-fine manner while reducing the quadratic complexity of vanilla self-attention to the linear complexity. In addition, for better foreground perception, we propose a hybrid positional embedding, which consists of the semantic-aware positional embedding and attention mask, to fully exploit geometry and semantic clues. Thanks to the designs above, OcTr delivers a competitive trade-off between accuracy and efficiency.

Our contribution is summarized in three-fold:

1. We propose OcTr for voxel-based 3D object detection, which efficiently learns enhanced representations by a global receptive field with rich contexts.
2. We propose an octree-based learnable attention sparsification scheme (*OctAttn*) and a hybrid positional embedding combining geometry and semantics.
3. We carry out experiments on the Waymo Open Dataset (WOD) and the KITTI dataset and report state-of-the-art performance with significant gains on far objects.

## 2. Related Work

### 2.1. 3D Object Detection from Point Clouds

There exist two prevailing point-cloud representations in 3D object detection, *i.e.* point-based and voxel-based.

The point-based methods [31, 40, 42, 60] directly process raw point clouds in the irregular 3D space. As a pioneering attempt, F-Pointnet [32] employs instance segmentation in frustums to extract proposals. VoteNet [31] clusters objects from the surface in a deep Hough voting manner. PointRCNN [40] generates 3D RoIs with foreground segmentation and applies an RCNN-style [35] two-stage refinement. Different from PointRCNN, some existing methods [3, 53, 59] build a lightweight and efficient single stage 3D object detection framework. However, the current point-based methods still suffer from a large computation burden, which is not suitable for large-scale point cloud scenes.

The voxel-based ones [4, 6, 13, 52, 55, 62, 63] conduct voxelization on entire point clouds to construct regular grids. VoxelNet [63] exploits the voxel feature encoding layer with 3D convolutions to extract the feature of each voxel. SECOND [52] improves the model with sparse 3D convolutions, significantly increasing the speed of both training and inference. Pointpillars [18] compacts point clouds into vertical columns and encodes them with 2D CNNs. Several recent methods [13, 38, 54] also explore merging point-based and voxel-based networks into one framework for complementary features from different representations of point clouds. Unfortunately, they all use small convolution kernels with limited receptive fields, which are not competent to capture global context that is important to 3D detection.

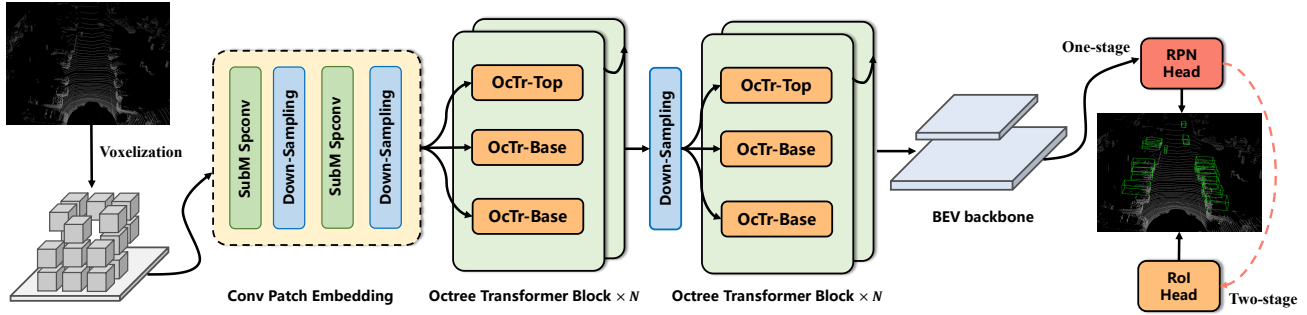


Figure 2. Framework overview of the proposed Octree-based Transformer (OcTr) model.

## 2.2. Transformer in 3D Vision

Inspired by the great success of the self-attention mechanism in NLP [47] and CV [8, 23], Transformers have been adapted to 3D vision for their ability to capture long-range dependencies. For instance, Point Transformer [61] brings in vector attention that modulates individual feature channels for point cloud classification and segmentation; PCT [11] presents offset-attention with the implicit Laplace operator and normalization refinement which is more suitable for point cloud learning. To address the high latency, some methods [30, 57] adopt voxels or patches for acceleration.

For 3D detection, 3DETR [28] treats and predicts bounding boxes as sequences in an end-to-end manner. CT3D [37] leverages a channel-wise Transformer architecture to refine the RoI head. To learn context-aware representations, some studies [9, 12, 24, 29] introduce Transformers into a point- or voxel-based encoder. Pointformer [29] stacks local, global and local-global Transformers based on the point-based encoder; VoTr [24] exploits dilated attention with fast query to enlarge receptive fields; VoxSet [12] builds an induced point set as proxies of global context and applies point-to-point translation using voxels as mediums; and SST [9] embraces single strides without down-sampling and conducts window attention combining with its shifted version. Even though they all expand receptive fields by diverse Transformer variants, global context is not adequately involved or efficiently utilized. In contrast, we propose Octree-based Transformer (OcTr) for voxel-based 3D object detection, achieving a true global receptive field that balances accuracy and efficiency.

## 3. Method

### 3.1. Framework

This sub-section describes the overall framework of the proposed OcTr model as shown in Fig. 2. Specifically, as inspired by [51], we first voxelize the point cloud into regular grids and adopt the sparse 3D convolution for patch embedding, where the grids are regarded as the “tokens” and are passed through the Octree Transformer Blocks

(OTB). Compared with the vanilla Transformer block, the self-attention module is substituted by the proposed octree-attention *OctAttn*, which encodes global context in a more efficient way. After applying a hybrid semantic embedding on multi-scale features, we sequentially stack two OTBs, tailed by a down-sampling layer. The voxel features are then projected into the BEV view by point-wise convolutions and are passed through a multi-scale dense 2D backbone. Ultimately, an anchor-based or anchor-free RPN head is used for 3D proposal generation, and the RoI head is optional for refinement. Note that our OcTr can be adopted for most of the voxel-based detection frameworks by simply altering 3D backbones.

### 3.2. Self-attention Revisit

According to [47], Transformer encoder blocks typically include a multi-head self-attention (MHSA) mechanism, a feed-forward network (FFN), a normalization function, and the residual connections [14]. Given an input sequence  $X$ , the principle part of MHSA is formulated as

$$\text{MHSA}(X) = \sum_{h=1}^H W_h \left[ \sigma \left( \frac{XW_q W_k^T X^T}{\sqrt{d}} \right) \cdot XW_v \right], \quad (1)$$

where  $h$  denotes the index of the head, and  $H$ ,  $\sigma$ ,  $W$  and  $d$  are the amount of heads, softmax function, learnable weight and feature dimension, respectively. The subscripts of  $q$ ,  $k$  and  $v$  indicate query, key and value. The inputs and outputs of the MHSA module are connected by residual connectors and normalization layers. The MLP-based FFN connects its inputs/outputs in a similar manner.

### 3.3. Octree Attention

An octree is a multi-scale asymmetric and efficient representation for unstructured 3D data such as point clouds. To build an octree for the input point cloud, we recursively sub-divide it in the breadth-first order until the pre-defined octree depth is reached. Whether to sub-divide an octant is determined by the occupancy [36, 48], the surface approximation [44, 49] or a learning algorithm [26].

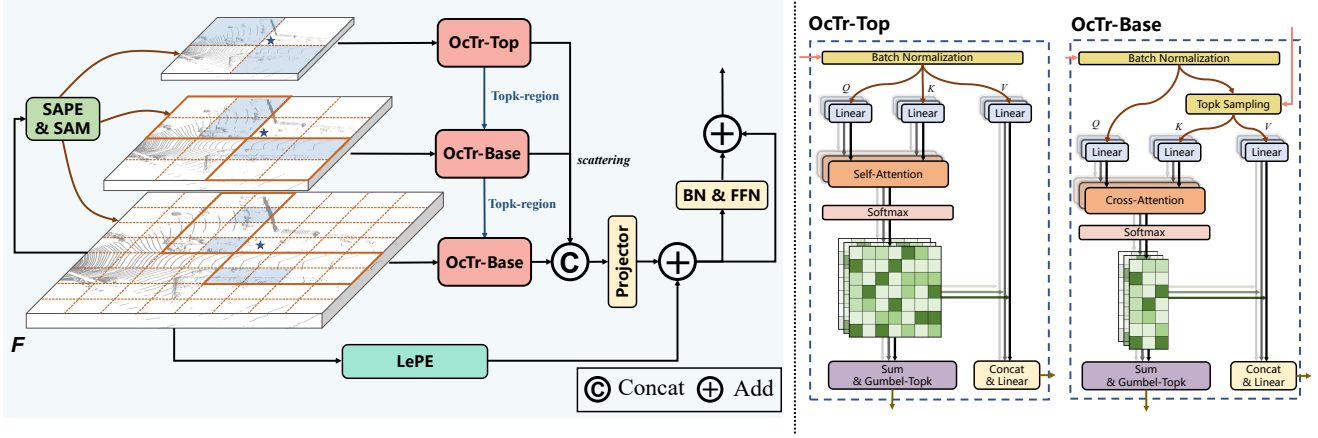


Figure 3. **Octree Attention** (*OctAttn*). The *OctAttn* module constructs a dynamic octree on a hierarchical feature pyramid with a partitive criterion of attention scores. For illustration, an octree representation is individually pruned from the pyramid for each grid in  $F$ . For example, the blue grids are the divided octants for the grids with stars. The detailed structures of **OcTr-Top** and **OcTr-Base** are shown at the right side.

Despite of being resource-friendly, the octree representation is discrete and non-differentiable, which motivates us to ameliorate previous cumbersome pre-processing division by introducing a novel octree-based attention mechanism. As depicted in [21, 45], the attention matrices calculated by self-attention imply the relevance of input tokens and guide feature selection. We thus prune the dense attention matrices of the multi-scale feature pyramid to sparse octree attention in an adaptive and parallel manner, namely *OctAttn*.

As shown in Fig. 3, let the output feature map and coordinates of convolutional patch embedding be  $F_0 \in \mathbb{R}^{M \times d}$  and  $I_0 \in \mathbb{R}^{M \times 3}$ ;  $M$  and  $d$  indicate the amount of the non-empty grids in a batch and feature dimension, respectively. Based on  $F_0$ , we generate a multi-scale feature pyramid as

$$C = \{F_n, I_n\}^N, n \in [0, N], \quad (2)$$

where

$$I_n = \lfloor \frac{I_0}{2^n} \rfloor, F_n = \text{BN}(\mathcal{S}_{\max}(F_0, I_n)), \quad (3)$$

$n$ ,  $N$ ,  $\text{BN}$  and  $\mathcal{S}_{\max}$  are the index of the level of the multi-scale feature pyramid, the height of the pyramid, batch normalization [16] and the max scatter function, respectively.

The pruning begins from the top of the pyramid. The top feature map, i.e.  $F_{N-1}$ , is reorganized to dense input tokens  $\bar{F}_{N-1} \in \mathbb{R}^{B \times m_{N-1} \times d}$ , where  $B$  and  $m_{N-1}$  are the batch size and maximum number of non-empty voxels per batch, respectively. The voxel is padded with 0 if it is empty. As shown in Eq. (4), the MHSA takes  $\bar{F}_{N-1}$  as input and outputs the attention score matrices  $\mathcal{A}_{N-1} \in \mathbb{R}^{B \times m_{N-1} \times m_{N-1}}$  and attentive features  $\bar{F}'_{N-1}$  as

$$\begin{aligned} \mathcal{A}_{N-1} &= \sum_{h=1}^H \sigma\left(\frac{\bar{F}_{N-1} W_q W_k^T \bar{F}_{N-1}^T}{\sqrt{d}}\right), \\ \bar{F}'_{N-1} &= \sum_{h=1}^H W_h \left[\sigma\left(\frac{\bar{F}_{N-1} W_q W_k^T \bar{F}_{N-1}^T}{\sqrt{d}}\right) \cdot \bar{F}_{N-1} W_v\right]. \end{aligned} \quad (4)$$

For each query token, we select the top  $k$  attention scores as its most relevant token group in a row-wise way, denoted by  $O_{N-1} \in \mathbb{Z}^{B \times m_{N-1} \times k}$ .

When propagating from top to bottom through the pyramid and reaching the  $n$ -th level, we uniformly sample limited  $K$  attending octants from the selected regions with features in  $\bar{F}_n$  and top  $k$  indices in  $O_{n+1}$ . We conduct cross-attention instead of self-attention with the dense query sequence of the  $n$ -th level,  $\bar{Q}_n \in \mathbb{R}^{B \times m_n \times d}$ , and the compact sampled key/value sequence of the  $n$ -th level,  $\bar{K}_n/\bar{V}_n \in \mathbb{R}^{B \times m_n \times K \times d}$ , which is formulated as below

$$\begin{aligned} \mathcal{A}_n &= \sum_{h=1}^H \sigma\left(\frac{\bar{Q}_n W_q W_k^T \bar{K}_n^T}{\sqrt{d}}\right), \\ \bar{F}'_n &= \sum_{h=1}^H W_h \left[\sigma\left(\frac{\bar{Q}_n W_q W_k^T \bar{K}_n^T}{\sqrt{d}}\right) \cdot \bar{V}_n W_v\right]. \end{aligned} \quad (5)$$

Backing off the sampling, this can be treated as an attention mask on self-attention matrices. The above process is run recursively until reaching the bottom level of the pyramid.

Furthermore, as the top  $k$  selection is a hard decision that disables gradient back-propagation, we adopt the Gumbel-top  $k$  technique [17] to perform a differentiable and continuous approximation by replacing the vanilla top  $k$  selection.



The normalized scores used in top $k$  are derived from the distribution in Eq. (6) during training, maintaining the original ones during inference.  $g$ ,  $\tau$ ,  $m_n$  denote the noise sampled from the Gumbel distribution, the temperature and the amount of non-empty voxels in layer  $n$ , respectively.

$$p_i = \frac{\exp((\mathcal{A}_n^i + g_i)/\tau)}{\sum_i \exp((\mathcal{A}_n^i + g_i)/\tau)} \in [0, 1]. \quad (6)$$

In order to leverage the multi-scale features in distinct spatial shapes, we concatenate them by upsampling, which is implemented by inverse indices of the scatter function, followed by a linear projection layer for aligning the input feature dimension.

As the local context is generally critical for object detection, inspired by [7], we additionally introduce a Locally enhanced Positional Embedding (LePE) which enables local neighbor interactions on the value sequence. With sub-manifold sparse convolutions, we replace the residual connections in the attention mechanism with LePE.

Finally, OTB is formulated as follows

$$\begin{aligned} \tilde{F} &= \text{FC}(\{F'_{N-1}|F'_{N-2}|\dots|F'_0\}) + \text{LePE}(F_0), \\ \tilde{F}' &= \text{BN}(\text{FFN}(\tilde{F})) + \tilde{F}, \end{aligned} \quad (7)$$

where  $F'_n \in \mathbb{R}^{m_0 \times d}$  denotes the compact tensor of the up-sampled  $F'_n$ ,  $|$  indicates concatenation, and FC denotes the fully-connected layer.

Besides, we analyze the time complexity of the octree attention as below

$$\begin{aligned} &\mathcal{O}\left(\left(\frac{M}{\omega^{N-1}}\right)^2 + \sum_{n=0}^{N-2} \frac{KM}{\omega^n}\right) \\ &= \mathcal{O}\left(\left(\frac{M}{\omega^{N-1}}\right)^2 + \frac{\omega}{\omega-1} KM(1 - \omega^{1-N})\right), \end{aligned} \quad (8)$$

where  $\omega$  is the average down-sampling ratio in the sparse voxel representation.

### 3.4. Semantic Positional Embedding

Due to the large proportion of background grids in point clouds, the attention matrices are dominated by background grid pairs, leading to a sub-optimal solution. To fully leverage the local 3D shape patterns and original voxel coordinates, we propose a hybrid positional embedding to capture both the geometry and semantic clues as displayed in Fig. 4.

Specifically, we first segment foreground grids using the supervision from the ground truth. Segmentation scores are predicted by a sub-manifold sparse convolution branch with a sigmoid function; and the focal loss [22] is applied to balance the foreground and background.

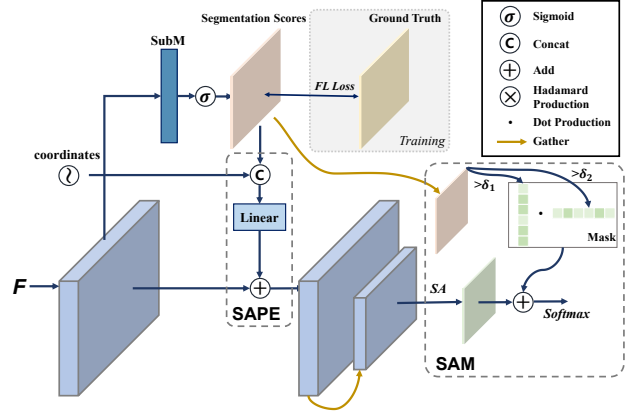


Figure 4. Illustration of semantic positional embedding, where an extra foreground segmentation branch is adopted with supervision and semantic scores are concatenated as absolute positional embeddings and serve as relative masks on attention matrices.

We concatenate the center coordinate  $(x, y, z)$  and the semantic *score* with the feature  $f$  in a grid-wise manner, followed by a linear projection without the bias as below

$$\begin{aligned} \text{SAPE}(X) &= \text{FC}_{d+4 \rightarrow d}(\{x, y, z, \text{score}|f\}) \\ &= \text{FC}_{d \rightarrow d}(f) + \text{FC}_{4 \rightarrow d}(x, y, z, \text{score}). \end{aligned} \quad (9)$$

Eq. (9) is equivalent to the absolute positional embedding (APE), thus being denoted as the Semantic APE (SAPE). The scatter function such as the mean, max and batch normalization naturally provide position and semantic information of the downsampled voxel grids, making it applicable for the multi-scale feature pyramid.

Besides the semantic clues implicitly used in SAPE, we employ the Semantic Attention Mask (SAM) based on the segmentation scores and we mask the attention matrices to address correlations between inferior queries and superior keys in a simple yet effective way. Given a scalar attention matrix  $\mathcal{A} \in \mathbb{R}^{N_q \times N_k}$  before softmax and the segmentation scores of query and key/value  $S_q \in \mathbb{R}^{N_q}$  and  $S_k \in \mathbb{R}^{N_k}$ , we formulate the attention matrix after softmax as

$$\mathcal{A}' = \sigma(-\Gamma \cdot [1 - (S_q \geq \delta_q)(S_k \geq \delta_k)]^T) + \mathcal{A}, \quad (10)$$

where  $N_q$ ,  $\delta_q$ ,  $N_k$ ,  $\delta_k$ ,  $\sigma$  and  $\Gamma$  are the length and threshold of the query sequence, length and threshold of the key sequence, softmax function and infinite scalar, respectively.

Finally, we broadcast the semantic mask  $\mathcal{A}'$  to all heads.

## 4. Experiments

We evaluate the proposed OcTr network on the Waymo Open Dataset (WOD) and KITTI dataset, both of which are popular in 3D object detection. In this section, we introduce the benchmarks and implementation details, make comparison to the previous state-of-the-art counterparts, and ablate the key designs of OcTr.

Model	Vehicle (L1) mAP/mAPH	Vehicle (L2) mAP/mAPH	Pedes. (L1) mAP/mAPH	Pedes. (L2) mAP/mAPH	Cyclist (L1) mAP/mAPH	Cyclist (L2) mAP/mAPH
SECOND [52]	70.96/70.34	62.58/62.02	65.23/54.24	57.22/47.49	57.13/55.62	54.97/53.53
PointPillar [18]	70.43/69.83	62.18/61.64	66.21/46.32	58.18/40.64	55.26/51.75	53.18/49.80
PartA <sup>2</sup> Net [41]	74.82/74.32	65.88/65.42	71.76/63.64	62.53/55.30	67.35/66.15	65.05/63.89
PVRCNN [38]	75.41/74.74	67.44/66.80	71.98/61.24	63.70/53.95	65.88/64.25	63.39/61.82
CenterPoint [55]	71.33/70.76	63.16/62.65	72.09/65.49	64.27/58.23	68.68/67.39	66.11/64.87
LiDAR-RCNN [20]	73.5/73.0	64.7/64.2	71.2/58.7	63.1/51.7	68.6/66.9	66.1/64.4
Voxel-RCNN [6]	75.59/-	66.59/-	-/-	-/-	-/-	-/-
PVRCNN++ [39]	77.82/77.32	69.07/68.62	77.99/71.36	69.92/63.74	71.80/70.71	69.31/68.26
SST <sup>†</sup> [9]	76.22/75.79	68.04/67.64	81.39/74.05	72.82/65.93	-/-	-/-
PDV [15]	76.85/76.33	69.30/68.81	74.19/65.96	65.85/58.28	68.71/67.55	66.49/65.36
Ours	<b>78.12/77.63</b>	<b>69.79/69.34</b>	<b>80.76/74.39</b>	<b>72.48/66.52</b>	<b>72.58/71.50</b>	<b>69.93/68.90</b>

Table 1. Performance on WOD with 202 *validation* sequences for vehicle (IoU=0.7), pedestrian (IoU=0.5) and cyclist (IoU=0.5), using 20% samples for training. All the results are achieved by the models simultaneously trained for 3 classes on single frames, except the ones of the model marked by †, which is only trained for a single class. Refer to *Supp. B* for the results trained with 100% samples.

Model	mAP <sub>3D</sub> (L1)@Vehicle			
	Overall	0-30m	30m-50m	50m-inf
PV-RCNN [38]	70.30	91.92	69.21	42.17
Voxel-RCNN [6]	75.59	92.49	74.09	53.15
VoTR-TSD [24]	74.95	92.28	73.36	51.09
CT3D [37]	76.30	92.51	75.07	55.36
Pyramid_PV [25]	76.30	92.67	74.91	54.54
PDV [15]	76.85	<b>93.13</b>	75.49	54.75
VoxSeT [12]	77.82	92.78	77.21	54.41
Ours	<b>78.82</b>	92.99	<b>77.66</b>	<b>58.02</b>

Model	mAP <sub>3D</sub> (L2)@Vehicle			
	Overall	0-30m	30-50m	50m-inf
PV-RCNN [38]	65.36	91.58	65.13	36.46
Voxel-RCNN [6]	66.59	91.74	67.89	40.80
CT3D [37]	69.04	91.76	68.93	42.60
PDV [15]	69.30	<b>92.41</b>	69.36	42.16
VoxSeT [12]	70.21	92.05	70.10	43.20
Ours	<b>70.50</b>	91.78	<b>71.28</b>	<b>45.46</b>

Table 2. Results on the WOD *validation* set in different ranges for vehicle detection.

#### 4.1. Datasets and Implementation Details

WOD [43] is a large dataset of autonomous driving scenes. It totally contains 798 training sequences with around 160K LiDAR samples and 202 validation sequences with 40K LiDAR samples, with the mean Average Precision (mAP) and mAP weighted by heading accuracy (mAPH) as evaluation metrics. There are also two levels of difficulty describing the sparsity in each bounding box, and LEVEL\_1 (L1) and LEVEL\_2 (L2) denote more than 5 points and 1-5 points, respectively. For detection performance along distance, it provides mAP/mAPH on 0-30m, 30m-50m and 50m-inf.

KITTI [10] is a widely used benchmark for 3D object detection, which includes 3,712, 3,769 and 7,518 frames for training, validation and testing, respectively. mAP is used as the official metric with 11 recall points for the *val* set and

40 for the *test* set, and the IoU thresholds are set to 0.7, 0.5, and 0.5 for car, pedestrian and cyclist. We use the official setting in all experiments.

**Implementation Details** The total loss for optimizing the overall two-stage detection is formulated as Eq. (11), where  $\mathcal{L}_{rcnn}$  can be omitted if there is no RoI head. Refer to *Supp. A.1* for more information.

$$\mathcal{L}_{det} = \mathcal{L}_{rpn} + \mathcal{L}_{rcnn} + \mathcal{L}_{seg}. \quad (11)$$

#### 4.2. Results on WOD

The results on the *validation* set are displayed in Tab. 1, and it can be seen that we achieve new state-of-the-art performance on all the three classes. In particular, for pedestrian, we outperform the baseline model PV-RCNN++ [39] by 2.77%/2.56% in terms of both L1 and L2 mAP, which indicates the effectiveness of the proposed model in handling hard examples.

In comparison with other Transformer-based models, we focus on vehicle since the counterparts only report the performance on it. As Tab. 2 shows, our OcTr achieves the best mAP among all these convolution- and Transformer-based backbones. It also outperforms the Transformer-based detection head network CT3D [37] by 2.52% and 1.46% in L1 and L2 mAP. Regarding the accuracies at different distances, OcTr ranks the first place in the range of 30m-50m and 50m-inf, which surpasses the previous best by 0.45%, 2.66% in L1 mAP and 1.18%, 2.26% in L2 mAP respectively. It clearly illustrates that OcTr has the advantage in capturing long-range fine-grained context, which facilitates dealing with objects far away. Indeed, far objects generally have much more sparse points than near ones and heavily rely on context for detection, thus benefiting more from OcTr. Refer to Fig. 5 for visualization.

Model	mAP <sub>3D</sub> @Car			mAP <sub>3D</sub> @Pedestrian			mAP <sub>3D</sub> @Cyclist		
	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
SECOND [52]	<b>88.61</b>	<b>78.62</b>	<b>77.22</b>	56.55	52.98	47.73	80.58	67.15	63.10
PointPillars [18]	88.46	77.28	74.65	57.75	52.29	47.90	80.04	62.61	59.52
VoTR [24]	87.86	78.27	76.93	-	-	-	-	-	-
VoxSeT [12]	88.45	78.48	77.07	60.62	54.74	50.39	84.07	68.11	65.14
Ours	88.43	78.57	77.16	<b>61.49</b>	<b>57.17</b>	<b>52.35</b>	<b>85.29</b>	<b>70.44</b>	<b>66.17</b>

Table 3. Results of the single-stage models on the KITTI *val* set. All the models adopt the same anchor-based region proposal network as the detection head. “Mod.” denotes the moderate difficulty level.

Model	mAP <sub>3D</sub> @Car on test				mAP <sub>3D</sub> @Car on val			
	Easy	Mod.	Hard	Mean	Easy	Mod.	Hard	Mean
SECOND [52]	83.34	72.55	65.82	73.90	88.61	78.62	77.22	81.48
PointPillars [18]	82.58	74.31	68.99	75.29	86.62	76.06	68.91	77.20
STD [54]	87.95	79.71	75.09	80.92	89.70	79.80	<b>79.30</b>	82.93
SA-SSD [13]	88.75	79.79	74.16	80.90	<b>90.15</b>	79.91	78.78	82.95
3DSSD [53]	88.36	79.57	74.55	80.83	89.71	79.45	78.67	82.61
PV-RCNN [38]	90.25	81.43	76.82	82.83	89.35	83.69	78.70	83.91
Voxel-RCNN [6]	<b>90.90</b>	81.62	77.06	83.19	89.41	84.52	78.93	84.29
CT3D [37]	87.83	81.77	77.16	82.25	89.54	86.06	78.99	<u>84.86</u>
VoTR-TSD [24]	89.90	82.09	<b>79.14</b>	<u>83.71</u>	89.04	84.04	78.68	83.92
VoxSeT [12]	88.53	82.06	77.46	82.68	89.21	<u>86.71</u>	78.56	84.83
Focals Conv [4]	90.55	<u>82.28</u>	77.59	83.47	89.52	84.93	79.18	84.54
Ours	<u>90.88</u>	<b>82.64</b>	<u>77.77</u>	<b>83.76</b>	<u>89.80</u>	<b>86.97</b>	<u>79.28</u>	<b>85.35</b>

Table 4. Comparison to the state-of-the-art models on the KITTI *test* and *val* sets. “Mod.” and “Mean” denote the moderate difficulty level and the average mAP for the three levels, respectively. The best results are **bolded** and the second best ones are underlined.

Detector	Veh. mAP (L1/L2)	Pedes. mAP (L1/L2)
SECOND [52]	70.96/62.58	65.23/57.22
Ours	<b>73.28/65.05</b>	<b>68.08/60.36</b>
PV-RCNN [38]	75.41/67.44	71.98/63.70
Ours	<b>76.77/68.31</b>	<b>73.22/64.30</b>
PV-RCNN++ [39]	77.82/69.07	77.99/69.92
Ours	<b>78.01/69.60</b>	<b>80.75/72.45</b>

Table 5. Results of extensions to different representative detectors on the WOD *validation* set.

### 4.3. Results on KITTI

The performance of the single-stage detectors is shown in Tab. 3. We take SECOND [52], a commonly used anchor-based model, as the baseline, and compare OcTr with another two advanced Transformer-based variants VoxSet [12] and VoTR [24]. We can see that OcTr achieves comparable results with SECOND and VoxSeT in car, while it significantly outperforms all the counterparts and reports the best performance both in pedestrian and cyclist. Benefiting from the large receptive field and fine-grained global context, it exceeds SECOND by 4.19% and 3.29% for pedestrians and cyclists respectively, where hard samples often appear.

We summarize the performance of the two-stage models on the KITTI *test* set in Tab. 4. With the help of the

Attention	Veh. mAP (L1/L2)	Pedes. mAP (L1/L2)
Ours ( <i>OctAttn</i> )	<b>73.3/65.1</b>	<b>68.1/60.4</b>
Performer [5]	71.4/63.6	65.7/57.9
ACT [27]	71.7/63.5	64.3/56.1
VoTr [24]	69.4/61.5	65.0/57.0
Nearest <i>K</i>	68.2/59.8	64.9/56.7

Table 6. Ablation on various attention mechanisms and sampling patterns on the WOD *validation* set.

multi-scale backbone features and rich global context, OcTr reaches a leading mAP in car at the moderate level, surpassing the state-of-the-art Focals-Conv [4] by 0.36%. We also evaluate OcTr on the KITTI *val* set, and OcTr again delivers the best performance in the average score, outperforming the second-best by 0.49%. One can observe that we rank the best or the second best in all the cases.

### 4.4. Ablation study

**Scalability on various detectors** As summarized in Tab. 5, we conduct experiments on three different and representative detectors, SECOND (single-stage, anchor-based), PV-RCNN (two-stage, anchor-based) and PV-RCNN++ (two-stage, anchor-free). Regardless of the number of stages or region proposal network, we acquire sound improvements

LEPE	SAPE	SAM	Veh. mAP (L1/L2)	Pedes. mAP (L1/L2)
			71.35/63.30	65.75/57.89
✓			72.34/64.32	66.56/58.62
✓	✓		72.64/64.46	66.62/58.83
✓		✓	72.86/64.40	67.79/59.90
✓	✓	✓	<b>73.28/65.05</b>	<b>68.08/60.36</b>

Table 7. Ablation on semantic positional embedding on the WOD validation set.

top $k$ number	Veh. mAP (L1/L2)	Pedes. mAP (L1/L2)
1	70.38/62.20	64.19/56.43
4	72.58/64.42	66.21/58.42
8	<b>73.28/65.05</b>	<b>68.08/60.36</b>
16	73.25/65.01	67.89/60.10

Table 8. Results of different  $k$  values on the WOD validation set.

Method	#Param. (M)	Latency (ms)	Memory (GB)
SECOND [52]	5.3	48	<b>2.3</b>
VoTR-SSD [24]	4.8	67	3.0
VoxSeT-SSD [12]	3.0	<b>37</b>	3.6
OcTr-SSD	<b>2.9</b>	64	2.5

Table 9. Resource costs of different backbones with single-stage detectors on the KITTI dataset, test on GTX2080Ti.

compared to the baselines of the sparse convolution backbones, highlighting its scalability.

**Ablation on *OctAttn*** We carry out additional experiments to make apple-to-apple comparison with several representative linear Transformer methods, including Performer [5] (kernel-based linear attention), ACT [27] (cluster-based linear attention), VoTr [24] (fixed patterns) and the Nearest- $K$  strategy. Tab. 6 lists the results, and our *OctAttn* clearly performs the best, showing its ability.

**Ablation on semantic positional embedding** We individually evaluate the contributions of LePE, SAPE and SAM with SECOND as the baseline detector on WOD in Tab. 7. By incorporating LePE, the L1/L2 performance is boosted by 0.99%/1.02% and 0.81%/0.73% on vehicle and pedestrian, illustrating its necessity. Furthermore, we separately verify the validity of SAPE and SAM. With semantic clues, we observe that SAPE increases by 0.3% L1 mAP on vehicle, while SAM provides an L1/L2 mAP improvement of 0.52%/0.08% and 1.23%/1.28% on vehicle and pedestrian. Finally, we simultaneously apply SAPE and SAM and construct the full model of OcTr, which gains 0.94%/0.73% and 1.52%/1.74% L1/L2 mAP on vehicle and pedestrian, showing its impact.

**Influence by top $k$**  Tab. 8 shows the performance of OcTr with various values of  $k$ . As we can see, the performance improves when  $k$  becomes larger but quickly saturates. Due to the redundancy of the scanned scenes, we argue that only a few tokens with high relevance need to be subdivided to

embrace fine-grained features.

**Analysis on model complexity** We compare OcTr with two recent Transformer-based models in terms of resource cost by keeping the same detection head. Tab. 9 shows that with the learnable octree attention mechanism, OcTr consistently maintains less model parameters and less memory occupancies than the counterparts. Regarding the inference speed, VoxSeT runs faster, but it should be noted that VoxSeT inputs with pillars which discard the height dimension, leading to inferior results. As for VoTR, we deliver a mild improvement in efficiency while bringing a large gain in performance.

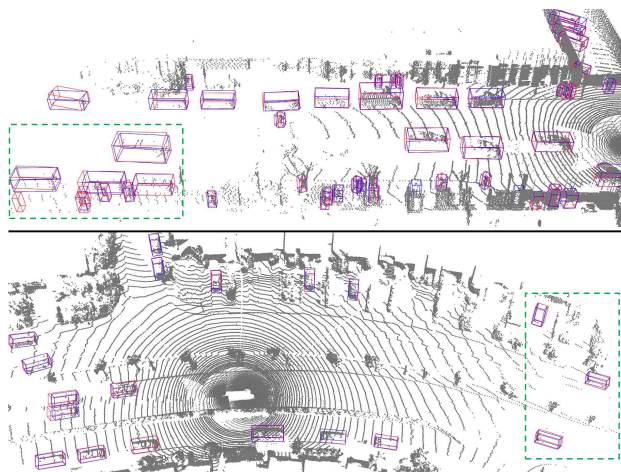


Figure 5. Visualization of results by OcTr on the WOD validation split. Blue/red indicates predicted/ground-truth bounding boxes

## 5. Conclusion

This paper proposes a novel voxel-based approach to 3D object detection, namely OcTr. It aims to balance the fine-grained global representation and efficiency with acceptable resource costs. To this end, we propose a learned sparsification attention mechanism, *OctAttn*, which adaptively prunes the octants from the multi-scale feature pyramid in a top-to-bottom manner. Furthermore, we adopt a hybrid semantic-aware positional embedding based on foreground segmentation. Extensive experiments are conducted on WOD and KITTI, and OcTr reaches the state-of-the-art performance, validating its effectiveness.

## Acknowledgment

This work is partly supported by the National Natural Science Foundation of China (No. 62022011 and No. 62202034), the Research Program of State Key Laboratory of Software Development Environment (SKLSDE-2021ZX-04), and the Fundamental Research Funds for the Central Universities.



## References

- [1] Joshua Ainslie, Santiago Ontanon, Chris Alberti, Vaclav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. Etc: Encoding long and structured inputs in transformers. In *EMNLP*, 2020. [2](#)
- [2] Iz Beltagy, Matthew E Peters, and Arman Cohan. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020. [2](#)
- [3] Chen Chen, Zhe Chen, Jing Zhang, and Dacheng Tao. Sasa: Semantics-augmented set abstraction for point-based 3d object detection. In *AAAI*, volume 1, 2022. [2](#)
- [4] Yukang Chen, Yanwei Li, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Focal sparse convolutional networks for 3d object detection. In *CVPR*, pages 5428–5437, 2022. [1](#), [2](#), [7](#)
- [5] Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamás Szepesvári, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J. Colwell, and Adrian Weller. Rethinking attention with performers. In *ICLR*, 2021. [7](#), [8](#)
- [6] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel r-cnn: Towards high performance voxel-based 3d object detection. In *AAAI*, volume 35, pages 1201–1209, 2021. [2](#), [6](#), [7](#)
- [7] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *CVPR*, pages 12124–12134, 2022. [2](#), [5](#)
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. [3](#)
- [9] Lue Fan, Ziqi Pang, Tianyuan Zhang, Yu-Xiong Wang, Hang Zhao, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Embracing single stride 3d object detector with sparse transformer. In *CVPR*, pages 8458–8468, 2022. [2](#), [3](#), [6](#)
- [10] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *CVPR*, pages 3354–3361. IEEE, 2012. [6](#)
- [11] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *CVM*, 7(2):187–199, 2021. [2](#), [3](#)
- [12] Chenhang He, Ruihuang Li, Shuai Li, and Lei Zhang. Voxel set transformer: A set-to-set approach to 3d object detection from point clouds. In *CVPR*, pages 8417–8427, 2022. [2](#), [3](#), [6](#), [7](#), [8](#)
- [13] Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure aware single-stage 3d object detection from point cloud. In *CVPR*, pages 11873–11882, 2020. [2](#), [7](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [3](#)
- [15] Jordan SK Hu. Point density-aware voxels for lidar 3d object detection. In *CVPR*, pages 8469–8478, 2022. [6](#)
- [16] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, pages 448–456. PMLR, 2015. [4](#)
- [17] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. In *ICLR*, 2016. [4](#)
- [18] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, pages 12697–12705, 2019. [2](#), [6](#), [7](#)
- [19] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosior, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *ICML*, pages 3744–3753. PMLR, 2019. [2](#)
- [20] Zhichao Li, Feng Wang, and Naiyan Wang. Lidar r-cnn: An efficient and universal 3d object detector. In *CVPR*, pages 7546–7555, 2021. [6](#)
- [21] Youwei Liang, Chongjian Ge, Zhan Tong, Yibing Song, Jue Wang, and Pengtao Xie. Not all patches are what you need: Expediting vision transformers via token reorganizations. In *ICLR*, 2022. [2](#), [4](#)
- [22] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *ICCV*, pages 2980–2988, 2017. [5](#)
- [23] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 10012–10022, 2021. [2](#), [3](#)
- [24] Jiageng Mao. Voxel transformer for 3d object detection. In *ICCV*, pages 3164–3173, 2021. [1](#), [2](#), [3](#), [6](#), [7](#), [8](#)
- [25] Jiageng Mao, Minzhe Niu, Haoyue Bai, Xiaodan Liang, Hang Xu, and Chunjing Xu. Pyramid r-cnn: Towards better performance and adaptability for 3d object detection. In *ICCV*, pages 2723–2732, 2021. [6](#)
- [26] Julien NP Martel, David B Lindell, Connor Z Lin, Eric R Chan, Marco Monteiro, and Gordon Wetzstein. Acorn: Adaptive coordinate networks for neural scene representation. In *ACM SIGGRAPH*, 2021. [3](#)
- [27] Zheng Minghang, Gao Peng, Zhang Renrui, Li Kunchang, Li Hongsheng, and Dong Hao. End-to-end object detection with adaptive clustering transformer. In *BMVC*, page 226, 2021. [7](#), [8](#)
- [28] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *ICCV*, pages 2906–2917, 2021. [3](#)
- [29] Xuran Pan, Zhuofan Xia, Shiji Song, Li Erran Li, and Gao Huang. 3d object detection with pointformer. In *CVPR*, pages 7463–7472, 2021. [1](#), [3](#)
- [30] Chungyun Park, Yoonwoo Jeong, Minsu Cho, and Jaesik Park. Fast point transformer. In *CVPR*, pages 16949–16958, 2022. [3](#)
- [31] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *ICCV*, pages 9277–9286, 2019. [2](#)
- [32] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *CVPR*, pages 918–927, 2018. [1](#), [2](#)

- [33] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, volume 30, 2017. 1
- [34] Jiezhong Qiu, Hao Ma, Omer Levy, Scott Wen-tau Yih, Sinong Wang, and Jie Tang. Blockwise self-attention for long document understanding. *arXiv preprint arXiv:1911.02972*, 2019. 2
- [35] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NeurIPS*, volume 28, 2015. 2
- [36] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, pages 3577–3586, 2017. 3
- [37] Hualian Sheng, Sijia Cai, Yuan Liu, Bing Deng, Jianqiang Huang, Xian-Sheng Hua, and Min-Jian Zhao. Improving 3d object detection with channel-wise transformer. In *ICCV*, pages 2743–2752, 2021. 3, 6, 7
- [38] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rnn: Point-voxel feature set abstraction for 3d object detection. In *CVPR*, pages 10529–10538, 2020. 2, 6, 7
- [39] Shaoshuai Shi, Li Jiang, Jiajun Deng, Zhe Wang, Chaoxu Guo, Jianping Shi, Xiaogang Wang, and Hongsheng Li. Pv-rnn++: Point-voxel feature set abstraction with local vector representation for 3d object detection. *IEEE TPAMI*, 2021. 6, 7
- [40] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3d object proposal generation and detection from point cloud. In *CVPR*, pages 770–779, 2019. 1, 2
- [41] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *IEEE TPAMI*, 43(8):2647–2664, 2020. 6
- [42] Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *CVPR*, pages 1711–1719, 2020. 1, 2
- [43] Pei Sun, Henrik Kretschmar, 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. 6
- [44] Jia-Heng Tang, Weikai Chen, Jie Yang, Bo Wang, Songrun Liu, Bo Yang, and Lin Gao. Octfield: Hierarchical implicit functions for 3d modeling. In *NeurIPS*, 2021. 3
- [45] Shitao Tang, Jiahui Zhang, Siyu Zhu, and Ping Tan. Quadtree attention for vision transformers. In *ICLR*, 2022. 2, 4
- [46] OpenPCDet Development Team. Openpcdet: An open-source toolbox for 3d object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020. 2
- [47] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, volume 30, 2017. 3
- [48] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM TOG*, 36(4):1–11, 2017. 3
- [49] Peng-Shuai Wang, Chun-Yu Sun, Yang Liu, and Xin Tong. Adaptive o-cnn: A patch-based deep representation of 3d shapes. *ACM TOG*, 37(6):1–11, 2018. 3
- [50] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM TOG*, 38(5):1–12, 2019. 1
- [51] Tete Xiao, Mannat Singh, Eric Mintun, Trevor Darrell, Piotr Dollár, and Ross Girshick. Early convolutions help transformers see better. In *NeurIPS*, volume 34, pages 30392–30400, 2021. 3
- [52] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 1, 2, 6, 7, 8
- [53] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *CVPR*, pages 11040–11048, 2020. 2, 7
- [54] Zetong Yang, Yanan Sun, Shu Liu, Xiaoyong Shen, and Jiaya Jia. Std: Sparse-to-dense 3d object detector for point cloud. In *CVPR*, pages 1951–1960, 2019. 2, 7
- [55] Tianwei Yin, Xingyi Zhou, and Philipp Krahenbuhl. Center-based 3d object detection and tracking. In *CVPR*, pages 11784–11793, 2021. 2, 6
- [56] Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. Big bird: Transformers for longer sequences. In *NeurIPS*, volume 33, pages 17283–17297, 2020. 2
- [57] Cheng Zhang, Haocheng Wan, Xinyi Shen, and Zizhao Wu. Patchformer: An efficient point transformer with patch attention. In *CVPR*, pages 11799–11808, 2022. 3
- [58] Yanan Zhang, Jiabin Chen, and Di Huang. Cat-det: Contrastively augmented transformer for multi-modal 3d object detection. In *CVPR*, pages 908–917, 2022. 1
- [59] Yifan Zhang, Qingyong Hu, Guoquan Xu, Yanxin Ma, Jianwei Wan, and Yulan Guo. Not all points are equal: Learning highly efficient point-based detectors for 3d lidar point clouds. In *CVPR*, pages 18953–18962, 2022. 2
- [60] Yanan Zhang, Di Huang, and Yunhong Wang. Pc-rgnn: Point cloud completion and graph neural network for 3d object detection. In *AAAI*, volume 35, pages 3430–3437, 2021. 1, 2
- [61] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *ICCV*, pages 16259–16268, 2021. 1, 2, 3
- [62] Wu Zheng, Weiliang Tang, Sijin Chen, Li Jiang, and Chi-Wing Fu. Cia-ssd: Confident iou-aware single-stage object detector from point cloud. In *AAAI*, volume 35, pages 3555–3562, 2021. 2
- [63] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, pages 4490–4499, 2018. 2