

SA-Det3D: Self-Attention Based Context-Aware 3D Object Detection

Prarthana Bhattacharyya Chengjie Huang Krzysztof Czarnecki
 University of Waterloo, Canada
 {p6bhatta, c.huang, k2czarne}@uwaterloo.ca

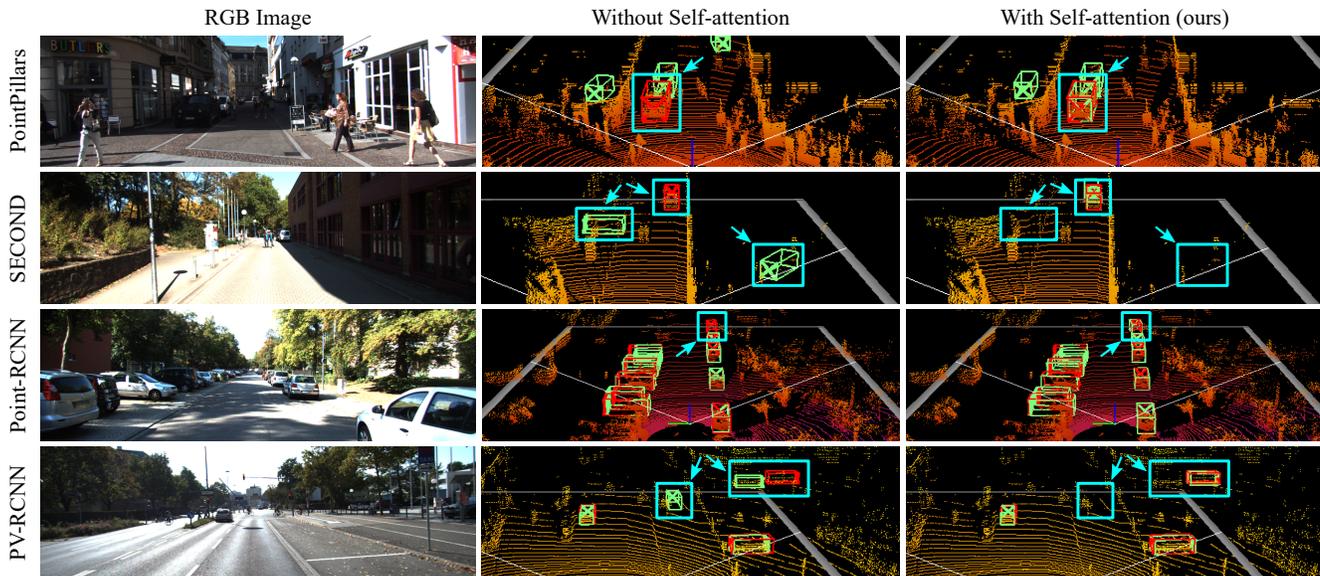


Figure 1: Performance illustrations on KITTI *val*. Red bounding box is ground truth, green is detector outputs. From left to right: (a) RGB images (b) Result of state-of-the-art methods: PointPillars [15], SECOND [46], Point-RCNN [30] and PV-RCNN [29]. (c) Result of our full self-attention (FSA) augmented baselines. Our method identifies missed detections and removes false positives.

Abstract

Existing point-cloud based 3D object detectors use convolution-like operators to process information in a local neighbourhood with fixed-weight kernels and aggregate global context hierarchically. However, non-local neural networks and self-attention for 2D vision have shown that explicitly modeling long-range interactions can lead to more robust and competitive models. In this paper, we propose two variants of self-attention for contextual modeling in 3D object detection by augmenting convolutional features with self-attention features. We first incorporate the pairwise self-attention mechanism into the current state-of-the-art BEV, voxel and point-based detectors and show consistent improvement over strong baseline models of up to 1.5 3D AP while simultaneously reducing their parameter footprint and computational cost by 15-80% and 30-50%, respectively, on the KITTI validation set. We

next propose a self-attention variant that samples a subset of the most representative features by learning deformations over randomly sampled locations. This not only allows us to scale explicit global contextual modeling to larger point-clouds, but also leads to more discriminative and informative feature descriptors. Our method can be flexibly applied to most state-of-the-art detectors with increased accuracy and parameter and compute efficiency. We show our proposed method improves 3D object detection performance on KITTI, nuScenes and Waymo Open datasets. Code is available at <https://github.com/AutoVision-cloud/SA-Det3D>.

1. Introduction

3D object detection has been receiving increasing attention in the computer vision and graphics community, driven by the ubiquity of LiDAR sensors and its widespread ap-

Method	Task	Modality	Context	Scalability	Attention + Convolution Combination	Stage Added
HG-Net [3]	detection	points	global-static	-	gating	Attention modules are added at the end.
PCAN [53]	place-recognition	points	local-adaptive	-	gating	
Point-GNN [31]	detection	points	local-adaptive	-	-	Attention modules fully replace convolution and set-abstraction layers.
GAC [38]	segmentation	points	local-adaptive	-	-	
PAT [47]	classification	points	global-adaptive	randomly sample points subset	-	
ASCN [45]	segmentation	points	global-adaptive	randomly sample points subset	-	
Pointformer [21]	detection	points	global-adaptive	sample points subset and refine	-	
MLCVNet [44]	detection	points	global-static	-	residual addition	Attention modules are inserted into the backbone.
TANet [16]	detection	voxels	local-adaptive	-	gating	
PMPNet [49]	detection	pillars	local-adaptive	-	gated-recurrent-unit	
SCANet [17]	detection	BEV	global-static	-	gating	
A-PointNet [20]	detection	points	global-adaptive	attend sequentially to small regions	gating	
Ours (FSA/DSA)	detection	points, voxels, pillars, hybrid	global-adaptive	attend to salient regions using learned deformations	residual addition	

Table 1: Properties of recent attention-based models for point-clouds

plications in autonomous driving and robotics. Point-cloud based 3D object detection has especially witnessed tremendous advancement in recent years [15, 46, 30, 29, 23, 48, 50, 12, 55, 22]. Grid-based methods first transform the irregular point-clouds to regular representations such as 2D bird’s-eye view (BEV) maps or 3D voxels and process them using 2D/3D convolutional networks (CNNs). Point-based methods sample points from the raw point-cloud and query a local group around each sampled point to define convolution-like operations [25, 35, 41] for point-cloud feature extraction. Both 2D/3D CNNs and point-wise convolutions process a local neighbourhood and aggregate global context by applying feature extractors hierarchically across many layers. This has several limitations: the number of parameters scales poorly with increased size of the receptive field; learned filters are stationary across all locations; and it is challenging to coordinate the optimization of parameters across multiple layers to capture patterns in the data [52].

In addition, point-cloud based 3D object detectors have to deal with missing/noisy data and a large imbalance in points for nearby and faraway objects. This motivates the need for a feature extractor that can learn global point-cloud correlations to produce more powerful, discriminative and robust features. For example, there is a strong correlation between the orientation features of cars in the same lane and this can be used to produce more accurate detections especially for distant cars with fewer points. High-confidence false positives produced by a series of points that resemble a part of an object can be also be eliminated by adaptively acquiring context information at increased resolutions.

Self-attention [36] has recently emerged as a basic building block for capturing long-range interactions. The key idea of self-attention is to acquire global information as a weighted summation of features from all positions to a target position, where the corresponding weight is calculated *dynamically* via a similarity function between the features in an embedded space at these positions. The number of pa-

rameters is independent of the scale at which self-attention processes long-range interactions. Inspired by this idea, we propose two self-attention based context-aware modules to augment the standard convolutional features—Full Self-Attention (FSA) and Deformable Self-Attention (DSA). Our FSA module computes pairwise interactions among all non-empty 3D entities, and the DSA module scales the operation to large point-clouds by computing self-attention on a representative and informative subset of features. Our experiments show that we can improve the performance of current 3D object detectors with our proposed FSA/DSA blocks while simultaneously promoting parameter and compute efficiency.

Contributions

- We propose a generic globally-adaptive context aggregation module that can be applied across a range of modern architectures including BEV [15], voxel [46], point [30] and point-voxel [29] based 3D detectors. We show that we can outperform strong baseline implementations by up to 1.5 3D AP (average precision) while simultaneously reducing parameter and compute cost by 15-80% and 30-50%, respectively, on the KITTI validation set.
- We design a scalable self-attention variant that learns to deform randomly sampled locations to cover the most representative and informative parts and aggregate context on this subset. This allows us to aggregate global context in large-scale point-clouds like nuScenes and Waymo Open dataset.
- Extensive experiments demonstrate the benefits of our proposed FSA/DSA modules by consistently improving the performance of state-of-the-art detectors on KITTI [10], nuScenes [2] and Waymo Open dataset [33].

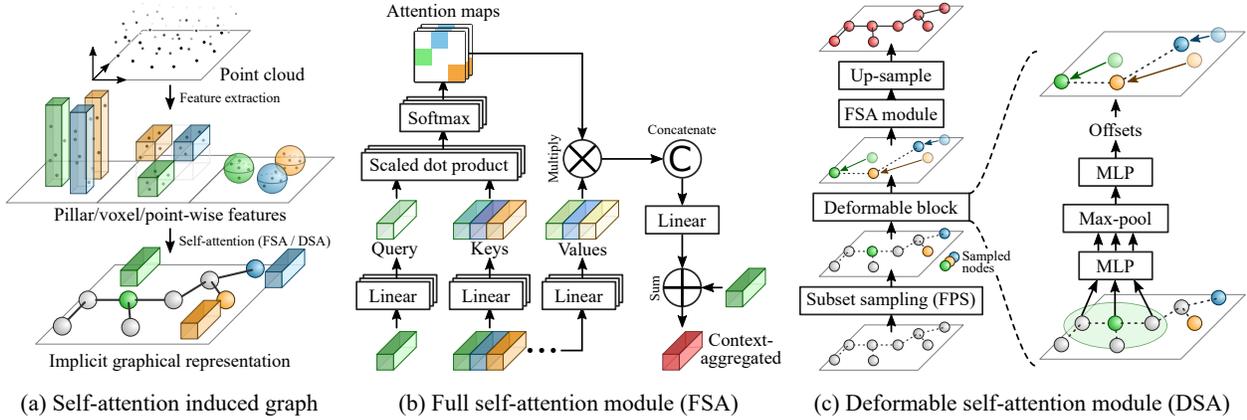


Figure 2: Architectures of the proposed FSA and DSA modules.

2. Related Works

3D Object Detection Current 3D object detectors include BEV, voxel, point or hybrid (point-voxel) methods. *BEV-based* methods like MV3D [5] fuse multi-view representations of the point-cloud and use 2D convolutions for 3D proposal generation. PointPillars [15] proposes a more efficient BEV representation and outperforms most fusion-based approaches while being 2-4 times faster. *Voxel-based* approaches, on the other hand, divide the point-cloud into 3D voxels and process them using 3D CNNs [55]. SECOND [46] introduces sparse 3D convolutions for efficient 3D processing of voxels, and CBGS [56] extends it with multiple heads. *Point-based* methods are inspired by the success of PointNet [24] and PointNet++ [25]. F-PointNet [23] first applied PointNet for 3D detection, extracting point-features from point-cloud crops that correspond to 2D camera-image detections. Point-RCNN [30] segments 3D point-clouds using PointNet++, and uses the segmentation features to better refine box proposals. *Point-Voxel-based* methods like STD [50], PV-RCNN [29] and SA-SSD [12] leverage both voxel and point-based abstractions to produce more accurate bounding boxes.

Relationship to current detectors: Instead of repeatedly stacking convolutions, we propose a simple, scalable, generic and permutation-invariant block called FSA/DSA to adaptively aggregate context information from the entire point-cloud. This allows remote regions to directly communicate and can help in learning relationships across objects. This module is flexible and can be applied in parallel to convolutions within the backbone of modern point-cloud based detector architectures.

Attention for Context Modeling Self-attention [36] has been instrumental to achieving state-of-the-art results in machine translation and combining self-attention with convolutions is a theme shared by recent work in natural language processing [42], image recognition [1], 2D ob-

ject detection [19], activity recognition [39], person re-identification [54] and reinforcement learning [51].

Using self-attention to aggregate global structure in point-clouds for 3D object detection remains a relatively unexplored domain. PCAN [53], TANet [16], Point-GNN [31], GAC [38], PMPNet [49] use local context to learn context-aware discriminative features. However relevant contextual information can occur anywhere in the point-cloud and hence we need global context modeling. HGNet [3], SCANet [17], MLCVNet [44] use global scene semantics to improve performance of object detection, but the global context vector is shared across all locations and channels and does not adapt itself according to the input features leading to a sub-optimal representation. PAT [47], ASCN [45], Pointformer [21] build globally-adaptive point representations for classification, segmentation and 3D detection. But because they use the costly pairwise self-attention mechanism, the self-attention does not scale to the entire point-cloud. Consequently, they process a randomly selected subset of points, which may be sensitive to outliers. To process global context for 3D object detection and scale to large point-clouds, Attentional PointNet [20] uses GRUs [7] to sequentially attend to different parts of the point-cloud. Learning global context by optimizing the hidden state of a GRU is slow and inefficient, however.

In contrast, our method can process context *adaptively* for each location from the entire point-cloud, while also *scaling* to large sets using learned deformations. Since the global context is fused with local-convolutional features, the training is stable and efficient as compared to GRUs or stand-alone attention networks [27]. Table 1 compares our work with recent point-cloud based attention methods.

3. Methods

In this section, we first introduce a Full Self-Attention (FSA) module for discriminative feature extraction in 3D

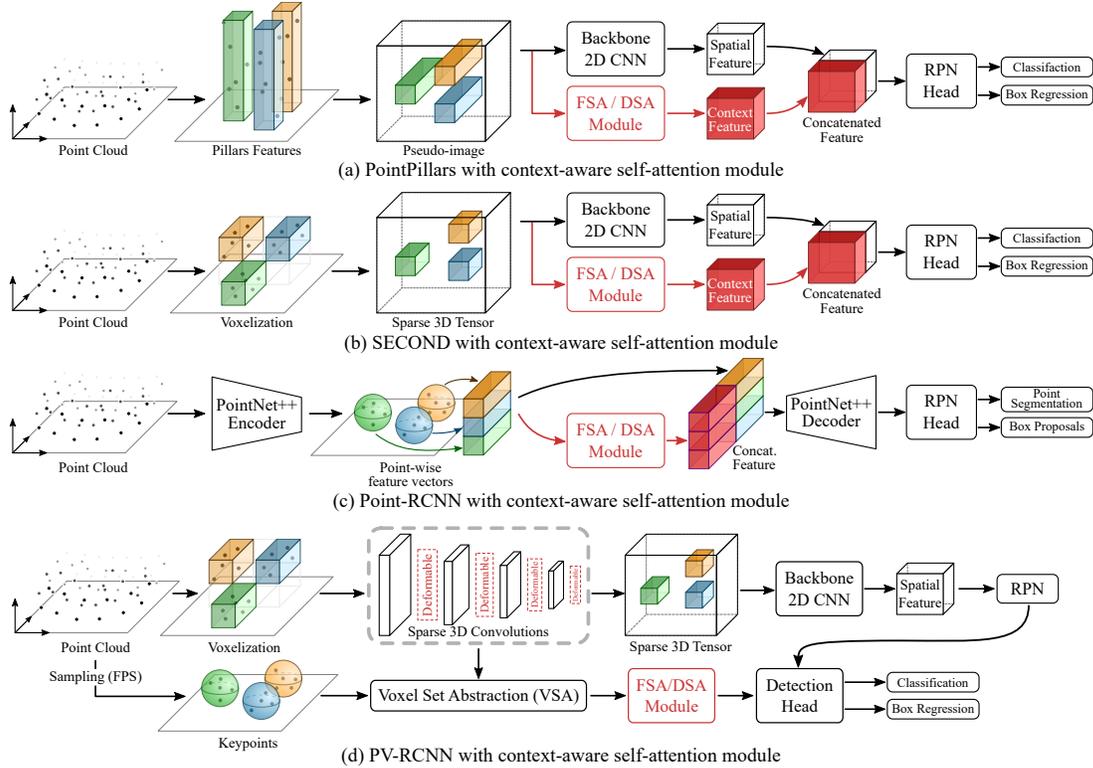


Figure 3: Proposed FSA/DSA module augmented network architectures for different backbone networks.

object detection that aims to produce more powerful and robust representations by exploiting global context. Next, inspired by 2D deformable convolutions [8] we introduce a variant of FSA called Deformable Self-Attention (DSA). DSA can reduce the quadratic computation time of FSA and scale to larger and denser point-clouds. The two proposed modules are illustrated in Figure 2.

3.1. Formulation

For the input set $\mathcal{X} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}$ of n correlated features and $i \in \{1, \dots, n\}$, we propose to use self-attention introduced by Vaswani et al. [36] to exploit the pairwise similarities of the i^{th} feature node with all the feature nodes, and stack them to compactly represent the global structural information for the current feature node.

Mathematically, the set of pillar/voxel/point features and their relations are denoted by a graph $G = (\mathcal{V}, \mathcal{E})$, which comprises the node set $\mathcal{V} = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n \in \mathbb{R}^d\}$, together with an edge set $\mathcal{E} = \{\mathbf{r}_{i,j} \in \mathbb{R}^{N_h}, i = 1, \dots, n \text{ and } j = 1, \dots, n\}$. A self-attention module takes the set of feature nodes, and computes the edges (see Figure 2 (a)). The edge $\mathbf{r}_{i,j}$ represents the relation between the i^{th} node and the j^{th} node, and N_h represents the number of heads (number of attention maps in Figure 2 (b)) in the attention mechanism across d feature input channels as described below. We assume that N_h divides d evenly. The advantage

of representing the processed point-cloud features as nodes in a graph is that now the task of aggregating global context is analogous to capturing higher order interaction among nodes by message passing on graphs for which many mechanisms like self-attention exist.

3.2. Full Self-Attention Module

Our Full Self-Attention (FSA) module projects the features \mathbf{x}_i through linear layers into matrices of query vectors \mathbf{Q} , key vectors \mathbf{K} , and value vectors \mathbf{V} (see Figure 2(b)). The similarities between query \mathbf{q}_i and all keys, $\mathbf{k}_{j=1:n}$, are computed by a dot-product, and normalized into attention weights \mathbf{w}_i , via a softmax function. The attention weights are then used to compute the pairwise interaction terms, $\mathbf{r}_{ij} = w_{ij} \mathbf{v}_j$. The accumulated global context for each node vector \mathbf{a}_i is the sum of these pairwise interactions, $\mathbf{a}_i = \sum_{j=1:n} \mathbf{r}_{ij}$. As we mentioned in our formulation, we also use multiple attention heads, applied in parallel, which can pick up channel dependencies independently. The final output for the node i is then produced by concatenating the accumulated context vectors $\mathbf{a}_i^{h=1:N_h}$ across heads, passing it through a linear layer, normalizing it with group normalization [43] and summing it with \mathbf{x}_i (residual connection).

Advantages: The important advantage of this module is that the resolution at which it gathers context is independent of the number of parameters and the operation is

Method	PointPillars [15]				SECOND [46]				Point-RCNN [30]				PV-RCNN [29]			
	3D	BEV	Param	FLOPs	3D	BEV	Param	FLOPs	3D	BEV	Param	FLOPs	3D	BEV	Param	FLOPs
Baseline	78.39	88.06	4.8 M	63.4 G	81.61	88.55	4.6 M	76.9 G	80.52	88.80	4.0 M	27.4 G	84.83	91.11	12 M	89 G
DSA	78.94	88.39	1.1 M	32.4 G	82.03	89.82	2.2 M	52.6 G	81.80	88.14	2.3 M	19.3 G	84.71	90.72	10 M	64 G
FSA	79.04	88.47	1.0 M	31.7 G	81.86	90.01	2.2 M	51.9 G	82.10	88.37	2.5 M	19.8 G	84.95	90.92	10 M	64.3 G
<i>Improve.</i>	+0.65	+0.41	-79%	-50%	+0.42	+1.46	-52%	-32%	+1.58	-	-37%	-38%	+0.12	-	-16%	-27%

Table 2: Performance comparison for moderate difficulty Car class on KITTI *val* split with 40 recall positions

permutation-invariant. This makes it attractive to replace a fraction of the parameter-heavy convolutional filters at the last stages of 3D detectors with self-attention features for improved feature quality and parameter efficiency.

Complexity: The pairwise similarity calculation is $\mathcal{O}(n^2d)$ in nature. The inherent sparsity of point-clouds and the efficient matrix-multiplication based pairwise computation makes FSA a viable feature extractor in current 3D detection architectures. However, it is necessary to trade accuracy for computational efficiency in order to scale to larger point-clouds. In the next section, we propose our Deformable Self-Attention module to reduce the quadratic computation time of FSA.

3.3. Deformable Self-Attention Module

Our primary idea is to attend to a representative subset of the original node vectors in order to aggregate global context. We then up-sample this accumulated structural information back to all node locations. The complexity of this operation is $\mathcal{O}(m^2d)$, where $m \ll n$ is the number of points chosen in the subset. In order for the subset to be representative, it is essential to make sure that the selected nodes cover the informative structures and common characteristics in 3D geometric space. Inspired by deformable convolution networks [8] in vision, we propose a geometry-guided vertex refinement module that makes the nodes self-adaptive and spatially recomposes them to cover locations which are important for semantic recognition. Our node offset-prediction module is based on vertex alignment strategy proposed for domain alignment [26, 11]. Initially m nodes are sampled from the point-cloud by farthest point sampling (FPS) with vertex features \mathbf{x}_i and a 3D vertex position v_i . For the i^{th} node, the updated position v'_i is calculated by aggregating the local neighbourhood features with different significance as follows:

$$x_i^* = \frac{1}{k} \text{ReLU} \sum_{j \in \mathcal{N}(i)} W_{\text{offset}}(\mathbf{x}_i - \mathbf{x}_j) \cdot (v_i - v_j) \quad (1)$$

$$v'_i = v_i + \tanh(W_{\text{align}} x_i^*) \quad (2)$$

where \mathcal{N}_i gives the i -th node's k -neighbors in the point-cloud and W_{offset} and W_{align} are weights learned end-to-end. The final node features are computed by a non-linear processing of the locally aggregated embedding as follows:

$$\mathbf{x}'_i = \max_{j \in \mathcal{N}(i)} W_{\text{out}} \mathbf{x}_j \quad (3)$$

Next, the m adaptively aggregated features $\{\mathbf{x}'_1, \dots, \mathbf{x}'_m\}$ are then passed into a full self-attention (FSA) module to model relationships between them. This aggregated global information is then shared among all n nodes from the m representatives via up-sampling. We call this module a Deformable Self-Attention (DSA) module as illustrated in Figure 2(c).

Advantages: The main advantage of DSA is that it can scalably aggregate global context for pillar/voxel/points. Another advantage of DSA is that it is trained to collect information from the most informative regions of the point-cloud, improving the feature descriptors.

4. Experiments

4.1. Network Architectures

We train and evaluate our proposed FSA and DSA modules on four state-of-the-art architecture backbones: PointPillars [15], SECOND [46], Point-RCNN [30], and PV-RCNN [29]. The architectures of the backbones are illustrated in Figure 3. The augmented backbones can be trained end-to-end without additional supervision.

For the KITTI dataset, the detection range is within [0,70.4] m, [-40,40] m and [-3,1] m for the XYZ axes, and we set the XY pillar resolution to (0.16, 0.16) m and XYZ voxel-resolution of (0.05, 0.05, 0.1) m. For nuScenes, the range is [-50,50] m, [-50,50] m, [-5,3] m along the XYZ axes and the XY pillar resolution is (0.2, 0.2) m. For the Waymo Open dataset, the detection range is [-75.2, 75.2] m for the X and Y axes and [-2, 4] m for the Z-axis, and we set the voxel size to (0.1, 0.1, 0.15) m. Additionally, the deformation radius is set to 3 m, and the feature interpolation radius is set to 1.6 m with 16 samples. The self-attention feature dimension is 64 across all models. We apply 2 FSA/DSA modules with 4 attention heads across our chosen baselines. For DSA, we use a subset of 2,048 sampled points for KITTI and 4,096 sampled points for nuScenes and Waymo Open Dataset. We use standard data-augmentation for point clouds. For baseline models, we reuse the pre-trained checkpoints provided by OpenPCDet [34].

4.2. Implementation Details

KITTI: KITTI benchmark [10] is a widely used benchmark with 7,481 training samples and 7,518 testing samples. We follow the standard split [5] and divide the training

Model	Car - 3D			Car - BEV			Cyclist - 3D			Cyclist - BEV		
	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard	Easy	Mod.	Hard
MV3D [5]	74.97	63.63	54.00	86.62	78.93	69.80	-	-	-	-	-	-
PointPillars [15]	82.58	74.31	68.99	90.07	86.56	82.81	77.10	58.65	51.92	79.90	62.73	55.58
SECOND [46]	83.34	72.55	65.82	89.39	83.77	78.59	71.33	52.08	45.83	76.50	56.05	49.45
PointRCNN [30]	86.96	75.64	70.70	92.13	87.39	82.72	74.96	58.82	52.53	82.56	67.24	60.28
STD [50]	87.95	79.71	75.09	94.74	89.19	86.42	78.69	61.59	55.30	81.36	67.23	59.35
3DSSD [48]	88.36	79.57	74.55	92.66	89.02	85.86	82.48	64.10	56.90	85.04	67.62	61.14
SA-SSD [12]	88.75	79.79	74.16	95.03	91.03	85.96	-	-	-	-	-	-
TANet [16]	83.81	75.38	67.66	-	-	-	73.84	59.86	53.46	-	-	-
Point-GNN [31]	88.33	79.47	72.29	93.11	89.17	83.90	78.60	63.48	57.08	81.17	67.28	59.67
PV-RCNN [29]	90.25	81.43	76.82	94.98	90.65	86.14	78.60	63.71	57.65	82.49	68.89	62.41
PV-RCNN + DSA (Ours)	88.25	81.46	76.96	92.42	90.13	85.93	82.19	68.54	61.33	83.93	72.61	65.82

Table 3: Performance comparison of 3D detection on KITTI *test* split with AP calculated with 40 recall positions. The **best** and **second-best** performances are highlighted across all datasets.

Model	Mode	mAP	NDS	Vehicle									
				Car	Truck	Bus	Trailer	CV	Ped	Moto	Bike	Tr. Cone	Barrier
PointPillars [15]	Lidar	30.5	45.3	68.4	23.0	28.2	23.4	4.1	59.7	27.4	1.1	30.8	38.9
WYSIWYG [13]	Lidar	35.0	41.9	79.1	30.4	46.6	40.1	7.1	65.0	18.2	0.1	28.8	34.7
PointPillars+ [37]	Lidar	40.1	55.0	76.0	31.0	32.1	36.6	11.3	64.0	34.2	14.0	45.6	56.4
PMPNet [49]	Lidar	45.4	53.1	79.7	33.6	47.1	43.0	18.1	76.5	40.7	7.9	58.8	48.8
SSN [57]	Lidar	46.3	56.9	80.7	37.5	39.9	43.9	14.6	72.3	43.7	20.1	54.2	56.3
Point-Painting [37]	RGB + Lidar	46.4	58.1	77.9	35.8	36.2	37.3	15.8	73.3	41.5	24.1	62.4	60.2
PointPillars + DSA (Ours)	Lidar	47.0	59.2	81.2	43.8	57.2	47.8	11.3	73.3	32.1	7.9	60.6	55.3

Table 4: Performance comparison of 3D detection with PointPillars backbone on nuScenes *test* split. “CV”, ”Ped”, “Moto”, “Bike”, “Tr. Cone” indicate construction vehicle, pedestrian, motorcycle, bicycle and traffic cone respectively. The values are taken from the official evaluation server <https://eval.ai/web/challenges/challenge-page/356/leaderboard/1012>.

Difficulty	Method	Vehicle	
		3D AP	3D APH
L1	StarNet [18]	53.7	-
	PointPillars [15]	56.6	-
	PPBA [6]	62.4	-
	MVF [5]	62.9	-
	AFDet [9]	63.7	-
	CVCNet [4]	65.2	-
	Pillar-OD [40]	69.8	-
	†SECOND [46]	70.2	69.7
	PV-RCNN [29]	70.3	69.7
	SECOND + DSA (Ours)	71.1	70.7
L2	†SECOND [46]	62.5	62.0
	PV-RCNN [29]	65.4	64.8
	SECOND + DSA (Ours)	63.4	63.0

Table 5: Comparison on Waymo Open Dataset *validation* split for 3D vehicle detection. Our DSA model has **52%** fewer parameters and **32%** fewer FLOPs compared to SECOND and **80%** fewer parameters and **41%** fewer FLOPs compared to PV-RCNN. †Re-implemented by [34]

samples into *train* and *val* split with 3,712 and 3,769 samples respectively. All models were trained on 4 NVIDIA Tesla V100 GPUs for 80 epochs with Adam optimizer [14] and one cycle learning rate schedule [32]. We also use the same batch size and learning rates as the baseline models.

nuScenes nuScenes [2] is a more recent large-scale benchmark for 3D object detection. In total, there are 28k, 6k, 6k, annotated frames for training, validation, and testing, respectively. The annotations include 10 classes with a

long-tail distribution. We train and evaluate a DSA model with PointPillars as the backbone architecture. All previous methods combine points from current frame and previous frames within 0.5 s, gathering about 300 k points per frame. FSA does not work in this case since the number of pillars in a point cloud is too large to fit the model in memory. In DSA, this issue is avoided by sampling a representative subset of pillars. The model was trained on 4 NVIDIA Tesla V100 GPUs for 20 epochs with a batch size of 8 using Adam optimizer [14] and one cycle learning rate schedule [32].

Waymo Open Dataset Waymo Open Dataset [33] is currently the largest dataset for 3D detection for autonomous driving. There are 798 training sequences with 158,081 LiDAR samples, and 202 validation sequences with 39,987 LiDAR samples. The objects are annotated in the full 360° field of view. We train and evaluate a DSA model with SECOND as the backbone architecture. The model was trained on 4 NVIDIA Tesla V100 GPUs for 50 epochs with a batch size of 8 using Adam optimizer [14] and one cycle learning rate schedule [32].

5. Results

5.1. 3D Detection on the KITTI Dataset

On KITTI, we report the performance of our proposed model on both *val* and *test* split. We focus on the average

precision for moderate difficulty and two classes: car and cyclist. We calculate the average precision on *val* split with 40 recall positions using IoU threshold of 0.7 for car class and 0.5 for cyclist class. The performance on *test* split is calculated using the official KITTI test server.

Comparison with state-of-the-art: Table 2 shows the results for car class on KITTI *val* split. For all four state-of-the-art models augmented with DSA and FSA, both variants were able to achieve performance improvements over strong baselines with significantly fewer parameters and FLOPs. On KITTI *test* split, we evaluate PV-RCNN+DSA and compare it with the models on KITTI benchmark. The results are shown in Table 3. On the car class DSA shows an improvement of 0.15 3D AP on the hard setting, while for the smaller cyclist class we achieve significantly better performance than all other methods with upto 4.5 3D AP improvement on the moderate setting. Overall, the results consistently demonstrate that adding global contextual information benefits performance and efficiency, especially for the difficult cases with smaller number of points.

5.2. 3D Detection on the nuScenes Dataset

To test the performance of our methods in more challenging scenarios, we evaluate PointPillars with DSA modules on the nuScenes benchmark using the official test server. In addition to average precision (AP) for each class, nuScenes benchmark introduces a new metric called nuScenes Detection Score (NDS). It is defined as a weighted sum between mean average precision (mAP), mean average errors of location (mATE), size (mASE), orientation (mAOE), attribute (mAAE) and velocity (mAVE).

Comparison with state-of-the-art: We first compare our PointPillars+DSA model with PointPillars+ [37], a class-balanced re-sampled version of PointPillars inspired by [56]. DSA achieves about 7% improvement in mAP and 4.2% improvement in NDS compared to PointPillars+, even for some small objects, such as pedestrian and traffic cone. Compared with other attention and fusion-based methods like PMPNet and Point-Painting, DSA performs better in the main categories of traffic scenarios such as Car, Truck, Bus and Trailer etc. Overall, our model has the highest mAP and NDS score compared to state-of-the-art PointPillars-based 3D detectors.

5.3. 3D Detection on the Waymo Open Dataset

We also report performance on the large Waymo Open Dataset with our SECOND+DSA model to further validate its effectiveness. The objects in the dataset are split into two levels based on the number of points in a single object, where LEVEL1 objects have at-least 5 points and the LEVEL2 objects have at-least 1 point inside. For evaluation, the average precision (AP) and average precision weighted by heading (APH) metrics are used. The IoU

threshold is 0.7 for vehicles.

Comparison with the state-of-the-art: Table 5 shows that our method outperforms previous state-of-the-art PV-RCNN with a 0.8%AP and 1%APH gain for 3D object detection while having 80% fewer parameters and 41% fewer FLOPs on LEVEL1. This supports that our proposed DSA is able to effectively capture global contextual information for improving 3D detection performance. Better performance in terms of APH also indicates that context helps to predict more accurate heading direction for the vehicles. On LEVEL2, we outperform the SECOND baseline by 0.9% AP and 1.0% APH. Overall SECOND+DSA provides the better balance between performance and efficiency as compared to PV-RCNN. The experimental results validate the generalization ability of FSA/DSA on various datasets.

5.4. Ablation studies and analysis

Ablation studies are conducted on the KITTI validation split [5] for moderate Car class using AP@R40, in order to validate our design choices.

Model variations In our ablation study with PointPillars backbone in Table 6, we represent the number of 2D convolution filters as $N_{filters}$, self-attention heads as N_h , self-attention layers as N_l , sampled points for DSA as N_{keypts} , deformation radius as r_{def} and the up-sampling radius as r_{up} .

Effect of number of filters: We note that both FSA and DSA outperform not only the models with similar parameters by 0.97% and 0.87% respectively, but also the state-of-the-art models with 80% more parameters by 0.65% and 0.55%. This indicates that our modules are extremely parameter efficient. Finally, we also note that if the number of parameters and compute are kept roughly the same as the baseline(Row-D), DSA outperforms the baseline by a large margin of 1.41%. We also illustrate consistent gains in parameter and computation budget across backbones in Figure 4.

Effect of number of self-attention heads and layers (Row-A): We note that increasing heads from 2 to 4 leads to an improvement of 0.37% for PointPillars. Since increasing number of self-attention layers beyond a certain value can lead to over-smoothing [28], we use 2 FSA/DSA layers in the backbone and 4 heads for multi-head attention.

Effect of number of sampled points (Row-B): For DSA, we also vary the number of keypoints sampled for computation of global context. We note that the performance is relatively robust to the number of sampled points.

Effect of deformation and upsampling radius (Row-C): For DSA, we note that the performance is generally robust to the deformation radius upto a certain threshold, but the up-sampling radius needs to be tuned carefully. Generally an up-sampling radius of 1.6m in cars empirically works well.

Model	$N_{filters}$	N_h	N_l	N_{keypts}	r_{def}	r_{up}	3D AP	Params	FLOPs
baseline	(64,128,256)	-	-	-	-	-	78.39	4.8M	63.4G
	(64,64,128)	-	-	-	-	-	78.07	1.5M	31.5G
(A)	(64,64,64)	2	2	-	-	-	78.67	1.0M	31.3G
		4	1	-	-	-	78.34	1.0M	31.5G
		4	2	-	-	-	79.04	1.0M	31.7G
		4	4	-	-	-	78.56	1.0M	32.0G
(B)	(64,64,64)	4	2	512	3	1.6	78.70	1.1M	32.4G
				1024	-	-	78.95	1.1M	32.4G
				2048	-	-	78.94	1.1M	32.4G
				4096	-	-	78.90	1.1M	32.4G
(C)	(64,64,64)	4	2	2048	2	1.6	78.93	1.1M	32.4G
					1.4	1.6	78.22	1.1M	32.4G
					3	2	78.10	1.1M	32.4G
					3	1	78.96	1.1M	32.4G
(D)	(64,128,256)	4	2	2048	2	1	79.80	5.1M	73.5G

Table 6: Ablation of model components with PointPillars backbone on KITTI moderate Car class of *val* split.

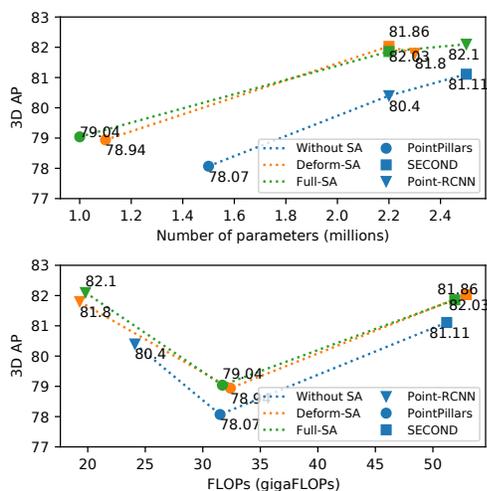


Figure 4: 3D AP on moderate Car class of KITTI *val* split (R40) vs. number of parameters (Top) and GFLOPs (Bottom) for baseline models and proposed baseline extensions with Deformable and Full SA.

Effect of noise on performance We introduce noise points to each object similar to TANet [16], to probe the robustness of representations learned. As shown in Figure 5, self-attention-augmented models are more robust to noise than the baseline. For example, with 100 noise points added, the performance of SECOND and Point-RCNN drops by 3.3% and 5.7% respectively as compared to SECOND-DSA and Point-RCNN-DSA, which suffer a lower drop of 2.7% and 5.1% respectively.

Effect of number of object points on performance We sort the cars based on the numbers of points in them and divide them into 3 groups. Then we calculate the 3D AP across every group. As shown in Figure 6, the effect of the self-attention module becomes apparent as the number of points on the cars decreases. For objects with very

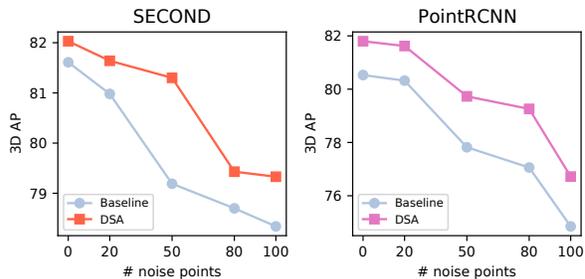


Figure 5: 3D AP of SECOND-DSA (orange) & Point-RCNN-DSA (violet) vs. SECOND & Point-RCNN baseline (light-steel-blue) for noise-points per ground-truth bounding box, varying from 0 to 100 on KITTI *val* moderate

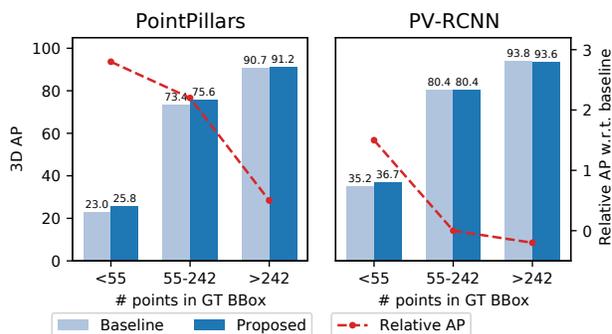


Figure 6: 3D AP of PointPillars-FSA, PV-RCNN-FSA and respective baselines vs. number of points in the ground-truth bounding box on KITTI *val*

few points, FSA can increase the 3D AP for PointPillars by 2.8% and PV-RCNN by 1.5%.

6. Conclusions

In this paper, we propose a simple and flexible self-attention based framework to augment convolutional features with global contextual information for 3D object detection. Our proposed modules are generic, parameter and compute-efficient, and can be integrated into a range of 3D detectors. We explore two forms of self-attention: full (FSA) and deformable (DSA). The FSA module encodes pairwise relationships between all 3D entities, whereas the DSA operates on a representative subset to provide a scalable alternative for global context modeling. Quantitative and qualitative experiments demonstrate that our architecture systematically improves performance of 3D object detectors.

References

- [1] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *ICCV*, 2019.

- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020.
- [3] Jintai Chen, Biwen Lei, Qingyu Song, Haochao Ying, Danny Z. Chen, and Jian Wu. A hierarchical graph network for 3d object detection on point clouds. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 389–398. IEEE, 2020.
- [4] Qi Chen, Lin Sun, Ernest Cheung, and Alan L. Yuille. Every view counts: Cross-view consistency in 3d object detection with hybrid-cylindrical-spherical voxelization. In Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin, editors, *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [5] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *CVPR*, 2017.
- [6] Shuyang Cheng, Zhaoqi Leng, Ekin Dogus Cubuk, Barret Zoph, Chunyan Bai, Jiquan Ngiam, Yang Song, Benjamin Caine, Vijay Vasudevan, Congcong Li, Quoc V. Le, Jonathon Shlens, and Dragomir Anguelov. Improving 3d object detection through progressive population based augmentation. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXI*, volume 12366 of *Lecture Notes in Computer Science*, pages 279–294. Springer, 2020.
- [7] Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014.
- [8] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, 2017.
- [9] Runzhou Ge, Zhuangzhuang Ding, Yihan Hu, Yu Wang, Sijia Chen, Li Huang, and Yuan Li. Afdet: Anchor free one stage 3d object detection. *CoRR*, abs/2006.12671, 2020.
- [10] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013.
- [11] Georgia Gkioxari, Jitendra Malik, and Justin Johnson. Mesh R-CNN. In *ICCV*, 2019.
- [12] Chenhang He, Hui Zeng, Jianqiang Huang, Xian-Sheng Hua, and Lei Zhang. Structure aware single-stage 3d object detection from point cloud. In *CVPR*, 2020.
- [13] Peiyun Hu, Jason Ziglar, David Held, and Deva Ramanan. What you see is what you get: Exploiting visibility for 3d object detection. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 10998–11006. IEEE, 2020.
- [14] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [15] 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*, 2019.
- [16] Zhe Liu, Xin Zhao, Tengpeng Huang, Ruolan Hu, Yu Zhou, and Xiang Bai. TANet: Robust 3d object detection from point clouds with triple attention. *AAAI*, 2020.
- [17] Haihua Lu, Xuesong Chen, Guiying Zhang, Qiuha Zhou, Yanbo Ma, and Yong Zhao. SCANet: Spatial-channel attention network for 3d object detection. In *ICASSP*. IEEE, 2019.
- [18] Jiquan Ngiam, Benjamin Caine, Wei Han, Brandon Yang, Yuning Chai, Pei Sun, Yin Zhou, Xi Yi, Ouais Alsharif, Patrick Nguyen, Zhifeng Chen, Jonathon Shlens, and Vijay Vasudevan. Starnet: Targeted computation for object detection in point clouds. *CoRR*, abs/1908.11069, 2019.
- [19] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention U-Net: Learning where to look for the pancreas, 2018.
- [20] Anshul Paigwar, Özgür Ercent, Christian Wolf, and Christian Laugier. Attentional pointnet for 3d-object detection in point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops, CVPR Workshops 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 1297–1306. Computer Vision Foundation / IEEE, 2019.
- [21] Xuran Pan, Zhuofan Xia, Shiji Song, Li Erran Li, and Gao Huang. 3d object detection with pointformer. *CoRR*, abs/2012.11409, 2020.
- [22] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep Hough voting for 3d object detection in point clouds. In *ICCV*, 2019.
- [23] 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*, 2018.
- [24] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017.
- [25] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *NIPS*, 2017.
- [26] Can Qin, Haoxuan You, Lichen Wang, C-C Jay Kuo, and Yun Fu. PointDAN: A multi-scale 3d domain adaption network for point cloud representation. In *NeurIPS*, 2019.
- [27] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. In *NeurIPS*, volume 32, 2019.
- [28] Yu Rong, Wenbing Huang, Tingyang Xu, and Junzhou Huang. DropEdge: Towards deep graph convolutional networks on node classification. In *ICLR*, 2019.
- [29] Shaoshuai Shi, Chaoxu Guo, Li Jiang, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. PV-RCNN: Point-voxel feature set abstraction for 3d object detection. In *CVPR*, 2020.
- [30] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. PointRCNN: 3d object proposal generation and detection from point cloud. In *CVPR*, 2019.

- [31] Weijing Shi and Raj Rajkumar. Point-gnn: Graph neural network for 3d object detection in a point cloud. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 1708–1716. IEEE, 2020.
- [32] Leslie N Smith. A disciplined approach to neural network hyper-parameters: Part 1—learning rate, batch size, momentum, and weight decay, 2018.
- [33] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, Vijay Vasudevan, Wei Han, Jiquan Ngiam, Hang Zhao, Aleksei Timofeev, Scott Etinger, Maxim Krivokon, Amy Gao, Aditya Joshi, Yu Zhang, Jonathon Shlens, Zhifeng Chen, and Dragomir Anguelov. Scalability in perception for autonomous driving: Waymo open dataset. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 2443–2451. IEEE, 2020.
- [34] OpenPCDet Development Team. Openpcdet: An open-source toolbox for 3d object detection from point clouds. <https://github.com/open-mmlab/OpenPCDet>, 2020.
- [35] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. KPConv: Flexible and deformable convolution for point clouds. In *ICCV*, 2019.
- [36] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NIPS*, 2017.
- [37] Sourabh Vora, Alex H Lang, Bassam Helou, and Oscar Beijbom. Pointpainting: Sequential fusion for 3d object detection. In *CVPR*, 2020.
- [38] Lei Wang, Yuchun Huang, Yaolin Hou, Shenman Zhang, and Jie Shan. Graph attention convolution for point cloud semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16-20, 2019*, pages 10296–10305. Computer Vision Foundation / IEEE, 2019.
- [39] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *CVPR*, 2018.
- [40] Yue Wang, Alireza Fathi, Abhijit Kundu, David A. Ross, Caroline Pantofaru, Thomas A. Funkhouser, and Justin Solomon. Pillar-based object detection for autonomous driving. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXII*, volume 12367 of *Lecture Notes in Computer Science*, pages 18–34. Springer, 2020.
- [41] Wenxuan Wu, Zhongang Qi, and Li Fuxin. PointConv: Deep convolutional networks on 3d point clouds. In *CVPR*, 2019.
- [42] Xin Wu, Yi Cai, Qing Li, Jingyun Xu, and Ho-fung Leung. Combining contextual information by self-attention mechanism in convolutional neural networks for text classification. In *WISE*. Springer, 2018.
- [43] Yuxin Wu and Kaiming He. Group normalization. *Int. J. Comput. Vis.*, 128(3):742–755, 2020.
- [44] Qian Xie, Yu-Kun Lai, Jing Wu, Zhoutao Wang, Yiming Zhang, Kai Xu, and Jun Wang. MLCVNet: Multi-level context votenet for 3d object detection. In *CVPR*, 2020.
- [45] Saining Xie, Sainan Liu, Zeyu Chen, and Zhuowen Tu. Attentional shapecontextnet for point cloud recognition. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pages 4606–4615. IEEE Computer Society, 2018.
- [46] Yan Yan, Yuxing Mao, and Bo Li. SECOND: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018.
- [47] Jiancheng Yang, Qiang Zhang, Bingbing Ni, Linguo Li, Jinxian Liu, Mengdie Zhou, and Qi Tian. Modeling point clouds with self-attention and gumbel subset sampling. In *CVPR*, 2019.
- [48] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3DSSD: Point-based 3d single stage object detector. In *CVPR*, 2020.
- [49] Junbo Yin, Jianbing Shen, Chenye Guan, Dingfu Zhou, and Ruigang Yang. LiDAR-based online 3d video object detection with graph-based message passing and spatiotemporal transformer attention. In *CVPR*, 2020.
- [50] Junbo Yin, Jianbing Shen, Chenye Guan, Dingfu Zhou, and Ruigang Yang. LiDAR-based online 3d video object detection with graph-based message passing and spatiotemporal transformer attention. In *CVPR*, 2020.
- [51] Vinicius Zambaldi, David Raposo, Adam Santoro, Victor Bapst, Yujia Li, Igor Babuschkin, Karl Tuyls, David Reichert, Timothy Lillicrap, Edward Lockhart, et al. Deep reinforcement learning with relational inductive biases. In *ICLR*, 2018.
- [52] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *ICML*. PMLR, 2019.
- [53] Wenxiao Zhang and Chunxia Xiao. PCAN: 3d attention map learning using contextual information for point cloud based retrieval. *CoRR*, abs/1904.09793, 2019.
- [54] Zhizheng Zhang, Cuiling Lan, Wenjun Zeng, Xin Jin, and Zhibo Chen. Relation-aware global attention for person re-identification. In *CVPR*, 2020.
- [55] Yin Zhou and Oncel Tuzel. VoxelNet: End-to-end learning for point cloud based 3d object detection. In *CVPR*, 2018.
- [56] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection, 2019.
- [57] Xinge Zhu, Yuexin Ma, Tai Wang, Yan Xu, Jianping Shi, and Dahua Lin. SSN: shape signature networks for multi-class object detection from point clouds. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, editors, *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part XXV*, volume 12370 of *Lecture Notes in Computer Science*, pages 581–597. Springer, 2020.