

OA-CNNs: Omni-Adaptive Sparse CNNs for 3D Semantic Segmentation

Bohao Peng¹ Xiaoyang Wu² Li Jiang^{3*} Yukang Chen¹
Hengshuang Zhao² Zhuotao Tian^{4*} Jiaya Jia¹
¹CUHK ²HKU ³CUHK, Shenzhen ⁴HIT, Shenzhen

Abstract

The booming of 3D recognition in the 2020s began with the introduction of point cloud transformers. They quickly overwhelmed sparse CNNs and became state-of-the-art models, especially in 3D semantic segmentation. However, sparse CNNs are still valuable networks, due to their efficiency treasure, and ease of application. In this work, we reexamine the design distinctions and test the limits of what a sparse CNN can achieve. We discover that the key credit to the performance difference is **adaptivity**. Specifically, we propose two key components, i.e., adaptive receptive fields (spatially) and adaptive relation, to bridge the gap. This exploration led to the creation of Omni-Adaptive 3D CNNs (OA-CNNs), a family of networks that integrates a lightweight module to greatly enhance the adaptivity of sparse CNNs at minimal computational cost. Without any self-attention modules, OA-CNNs favorably surpass point transformers in terms of accuracy in both indoor and outdoor scenes, with much less latency and memory cost. Notably, it achieves 76.1%, 78.9%, and 70.6% mIoU on ScanNet v2, nuScenes, and SemanticKITTI validation benchmarks respectively, while maintaining at most 5× better speed than transformer counterparts. This revelation highlights the potential of pure sparse CNNs to outperform transformer-related networks. Our code is built upon Pointcept [9], which is available at here ¹.

1. Introduction

3D scene understanding is critical in various practical applications, including robotics, autonomous driving, and augmented reality [13, 16, 19, 20, 30, 55, 70, 71, 74]. In contrast to images, which typically exhibit densely and uniformly arranged pixels [10, 23, 35, 49, 51, 52], 3D point clouds often manifest irregular and scattered distributions. It leads to various feature extractors in 3D scene understanding.

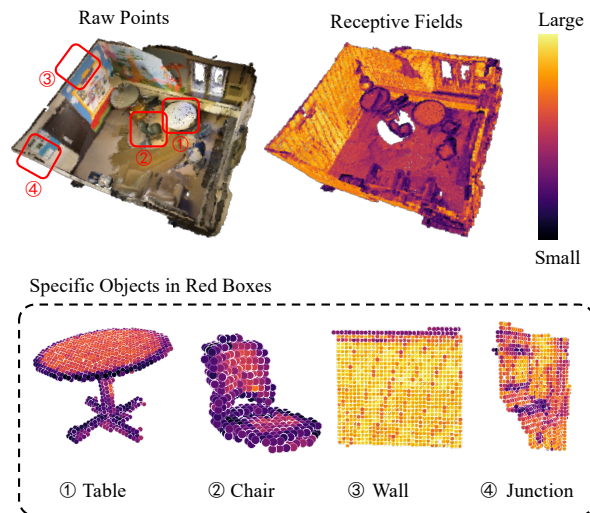


Figure 1. Visualization of 3D scene receptive fields controlled by our proposed adaptive aggregator. Objects’ edges and junctions require smaller receptive fields due to their sophisticated structures, while flat planes and unitary structures require broader fields.

There are two mainstream 3D networks. The first is point-based networks [41, 42], which advocate directly manipulating the unstructured points. Thanks to the flexibility of point-wise operations, point-based methods, particularly those with transformer architectures [12, 32, 36, 39, 50, 53, 54, 73], have gradually become dominant. The second is sparse CNNs [7, 13], where irregular point clouds are converted into voxels during data preprocessing. This allows us to leverage the *locally structured* benefits and facilitate high efficiency. Due to this practical value, sparse CNNs have been widely exploited in existing literature [37, 45, 65, 75]. However, its accuracy is usually inferior to its transformer counterparts [21, 36, 60, 73], especially in 3D scene semantic segmentation.

Given the high potential of sparse CNNs, we carefully examine the inner reasons for the performance gap in this paper. We find that the key distinction between sparse

*Corresponding authors.

¹<https://github.com/Pointcept/Pointcept>

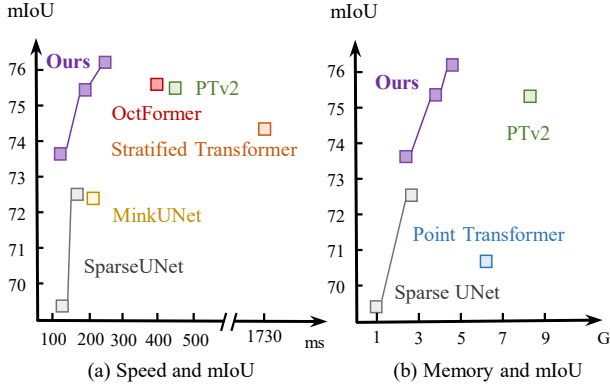


Figure 2. Comparison between various transformer-based [21, 60, 73] and CNN-based [7, 13] within RTX 3090. For OctFormer, we reproduce the official repository and include the cost of building the octree. If a method has multiple versions, they are indicated by different dots.

CNNs and point transformers behind is **adaptivity** – the latter can flexibly adapt to individual contexts while it may not be feasible for the former with static perception. Without degrading efficiency, we bridge this gap via two key components: (1) spatially adaptive receptive fields, and (2) adaptive relations.

Adaptively adjusting receptive fields via attention mechanisms is one of key designs in transformer-based frameworks [54, 73] to achieve top performance. Intuitively, different parts of the 3D scene with various geometric structures and appearances should be catered with different receptive sizes, as visualized in Fig. 1. Flat and sparse regions like the wall and floor need large receptive fields to yield consistent predictions with broader cues, while sophisticated parts like the plane junctions and small objects need smaller ones to screen unnecessary context that may overwhelm the local details. To enable our CNN-based framework to adaptively perceive the contextual information, we partition the 3D scene into non-overlapping pyramid grids. We then utilize the proposed Adaptive Relation Convolution (ARConv) in multiple scales and design a selective aggregator to *adaptively aggregate the multi-scale outputs based on the local characteristics*. Instead of pursuing consistent large receptive fields (like LargeKernel3D [6]), we find that this adaptive manner is sufficient and more efficient.

Adaptive relationships, achieved via self-attention maps, is another key strength over CNNs. To facilitate the establishment of relationships among local contexts, we introduce a multi-one-multi paradigm in ARConv, as depicted in Fig. 6. Specifically, we dynamically generate kernel weights for non-empty voxels based on their correlations with the grid centroid. By adopting this approach, we can maintain a lightweight design[58] with a linear complexity proportional to the voxel quantity, which *effectively expands the receptive fields and achieves optimal efficiency*.

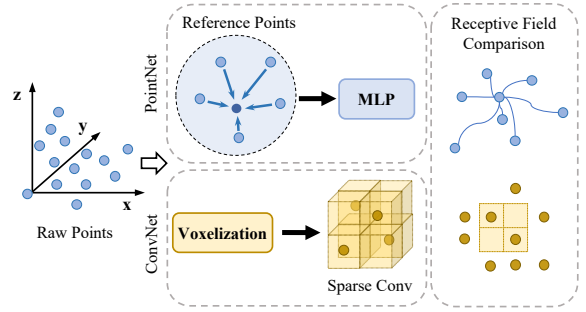


Figure 3. Comparisons between the 3D point-based [41, 73] and convolutional networks [7, 13]. PointNets directly process the raw points and provide more flexible and broader receptive fields. ConvNets handle structural data after additional voxelization pretreatment with higher efficiency and lower consumption.

Extensive experiments validate our approach’s effectiveness, and our designs enable sparse CNNs to outperform state-of-the-art point-based methods with transformer architectures, with little efficiency compromise, as shown in Fig. 2. We conduct the comparisons under the same experimental settings, without any additional pretraining or auxiliary methods. Remarkably, it achieves mIoU scores of 76.1%, 78.9%, and 70.6% on the ScanNet v2 [11], nuScenes [4], and SemanticKITTI [2] validation benchmarks, respectively. It highlights the potential of sparse CNNs over transformer-related models in both performance and efficiency, regardless of indoor or outdoor scenes.

In conclusion, our contributions are listed as follows:

- We analyze and find that adaptivity is the key to bridging the gap between sparse CNNs and point transformers.
- We propose OA-CNNs as solutions, consisting of dynamic receptive fields and adaptive relation mapping.
- Our method outperforms state-of-the-art methods with promising efficiency on popular benchmarks including ScanNet v2, ScanNet200, nuScenes and SemanticKITTI semantic segmentation.

2. Related Work

Point-based learning. Point-based methods advocate directly processing the unstructured raw points without any additional regulation pretreatment [14, 18, 31, 63, 67]. PointNet [41] is pioneering work in this trend which leverages point-wise MLP and permutation invariance operation to obtain the global feature of input points. More details and comparisons are shown in Fig. 3. Several follow-up works [15, 18, 42] continue to strengthen their capabilities through hierarchical multi-scale perception and local-global feature aggregation. Especially with the development of the attention mechanism [54, 68, 69], point-wise perception with the transformer architecture [21, 60, 62, 73] provides long-range dependences and bridges global contexts rela-

tionships. These frameworks have shown outperforming superiority and gradually become dominant. However, attention calculation and point-wise operation suffer from more expensive computation and memory consumption, and the complex architecture also makes them more challenging to deploy.

CNN-based learning. Compared with dense images arranging pixels into a rasterized grid, point cloud directly records the points’ spatial coordinates, which are typically irregular and lack unified metrics. Projection-based [5, 24–26, 29, 46] methods intuitively project the raw 3D points into flat images from various views, and the subsequent operations are logically the same as the 2D pipeline. However, the projection seriously destroyed the point cloud’s geometrical information, especially for the in-door scenes with more stereoscopic structures. An alternative technique is to quantize the 3D scene and transform irregular point clouds into regular voxel representation [3, 37, 38, 45]. 3D convolutions are commonly applied to handle these voxel collections while consuming high computation and memory. Sparse and submanifold convolutions [13] are introduced to alleviate these issues and improve efficiency. Sparse convolution introduces the hash table for the voxels’ indices retrieval, which is convenient and efficient. Moreover, 3D submanifold convolution has made a further restriction only processing the non-empty elements sacrificing some flexibility in change for more efficiency and less consumption. However since the complexity of the kernel size is $O(K^3)$, the receptive fields of sparse convolutions are still limited by the parameter quantity, which seriously restricts the global perception ability. In this work, we explore a lightweight design [58] to expand 3D convolution with an adaptive receptive range [27].

Dynamic convolutions. Regular convolutions optimize the learnable kernel weights during training and fix kernel weights in the inference process. Dynamic convolution [17, 66] proposes to generate the convolution kernel adaptively depending on the specific conditions. Previous works [48, 59, 64] have widely explored introducing dynamic convolution into sparse data processing. However, these works are also based on point-wise methods and typically generate kernel weights depending on the relative position information, which requires expensive computation and memory consumption. In this work, we inherit conditional convolution to propose a lightweight grid convolution with a regular structure. Moreover, we introduce the adaptive aggregator for the multi-scale pyramid aggregation to bridge extended-range contexts efficiently.

3. Omni-Adaptive 3D Sparse CNNs

In this section, we provide a detailed introduction to our designed lightweight modules and their application in con-

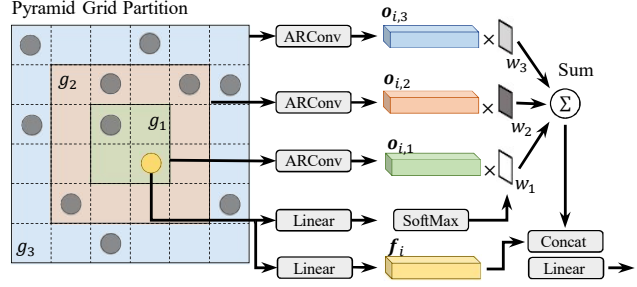


Figure 4. Illustration for the adaptive aggregator, which learns to aggregate various grid contexts under multi-pyramid scales from the voxel’s instinct characteristics.

structing a series of omni-adaptive 3D sparse CNNs (OA-CNNs). It surpasses point transformers in 3D recognition with limited latency/memory overhead. OA-CNNs consist of three design contents, *i.e.*, spatially adaptive receptive fields in Sec. 3.1, Adaptive Relation Convolution (ARConv) in Sec. 3.2, and the overall architecture in Sec. 3.3.

3.1. Spatially adaptive receptive fields

Motivation. Various receptive field sizes are required in distinct positions and objects in one 3D scene. For example, as shown in Fig. 1, regions belonging to the wall and floor are relatively flat and elementary, which require larger receptive fields to yield consistent predictions. However the geometric structures of the plane junction or sophisticated objects are more complex and need smaller receptive fields to retain the local characteristics. Transformer frameworks [21, 54, 73] adjust the perception range by the attention mechanism retrieving the relevance with the surrounding contexts but significantly increasing memory and computing consumption. However, sparse CNNs lack the ability to handle this issue. In OA-CNNs, we overcome this by directly determining the perception size with the aid of the intrinsic voxel features, as illustrated in Fig. 4.

Voxel grid. Expanding the receptive field is necessary for pursuing adaptive perception since the typical 3D convolution kernel size is generally set as $3 \times 3 \times 3$ limited by the parameter quantity. To achieve this, we utilize the voxel grid in our approach. Formally, define $\mathcal{V} = (\mathcal{P}, \mathcal{F})$ as a sparse voxelized 3D scene representation containing a set of voxels $\mathbf{v}_i = (\mathbf{p}_i, \mathbf{f}_i)$, where $\mathbf{p}_i \in \mathbb{R}^3$ represents the positional integer indice and $\mathbf{f}_i \in \mathbb{R}^d$ is the corresponding feature with d channels. The global voxel set \mathcal{V} is then partitioned into N non-overlapping voxel grids $[\mathcal{V}_1, \mathcal{V}_2, \dots, \mathcal{V}_N]$, $\mathcal{V}_i = \{\mathbf{v}_j \mid \mathbf{p}_j \in \Omega(i)\}$, where \mathcal{V}_i indicates i -th voxel grid and $\Omega(i)$ obtains i -th voxel grid’s indices range. The voxel grid size can be considerably larger than that of the typical 3D convolution kernel, such that the

receptive field is effectively expanded.

Pyramid grid partition. Although a sufficiently large grid size can provide a global view, it may not be able to capture intricate details for sophisticated objects. In an effort to prepare the alternative grid sizes for adaptively accommodating different areas, we rasterize the entire 3D scene into pyramid voxel grids. Specifically, let's define $\mathcal{G} = \{g_k\}^K$ as the set of K grid sizes partitioning the 3D space, where K is set as 3 in our experiments. The output $\mathbf{o}_i \in \mathbb{R}^{k \times d}$ of the i -th voxel grid under k -th scale is obtained as:

$$\mathbf{o}_{i,k,:} = \text{Conv}(\{\mathbf{f}_j \mid \mathbf{p}_j \in \Omega(i, g_k)\}), \quad (1)$$

where $\Omega(i, g_k)$ represents the range of voxel indices in the i -th voxel grid in the size g_k , and $\text{Conv}(\cdot)$ indicates the convolution for aggregating voxel features in the voxel grids to get the voxel grid feature. Observing the intolerably heavy parameters associated with the standard sparse 3D convolution $\text{Conv}(\cdot)$ using a large kernel, we introduce the ARConv in Sec. 3.2 as a solution to this issue. The ARConv improves results without sacrificing efficiency and establishes relationships among the voxel grid.

Adaptive aggregator. To achieve a customizable receptive field, we propose an adaptive aggregator that autonomously adjusts the receptive fields based on the intrinsic characteristics and spatial structure of individual voxels, which is illustrated in Fig. 4. Given K multi-scale grid partitions with sizes $\mathcal{G} = \{g_k\}^K$, our proposed adaptive aggregator weights and fuses the multi-scale outputs. We use a learnable function δ_{adp} to predict the preference weights \mathbf{w}_i of K grid sizes as:

$$\mathbf{w}_i = \text{SoftMax}(\delta_{adp}(\mathbf{f}_i)), \quad (2)$$

where $\mathbf{w} \in \mathbb{R}^{N_i \times K}$, N_i denotes the number of voxels inside the i -th voxel grid, and $\delta_{adp} : \mathbb{R}^d \mapsto \mathbb{R}^K$ is a learnable linear layer and $\text{SoftMax}(\cdot)$ denotes the softmax operation over K grid sizes. We subsequently employ the predicted weights to aggregate the convolution outputs, which contain global information, with the original features to enhance them,

$$\mathbf{f}'_i = \delta_{out}(\delta_{proj}(\mathbf{f}_i) \oplus \sum_{k=1}^K \mathbf{w}_{i,k} \cdot \mathbf{o}_{\phi'(i,k,k)}), \quad (3)$$

where $\delta_{out} : \mathbb{R}^{2d} \mapsto \mathbb{R}^d$ and $\delta_{proj} : \mathbb{R}^d \mapsto \mathbb{R}^d$ are two linear layers with normalization and activation, \oplus denotes vector concatenation and $\phi'(i, k)$ reversely returns the voxel grid index containing the i -th voxel under g_k grid size partition.

So far, we have presented a method for constructing the spatially adaptive receptive fields based on individual context, but it is not yet capable of establishing adaptive relationships as the point-based transformer counterparts.

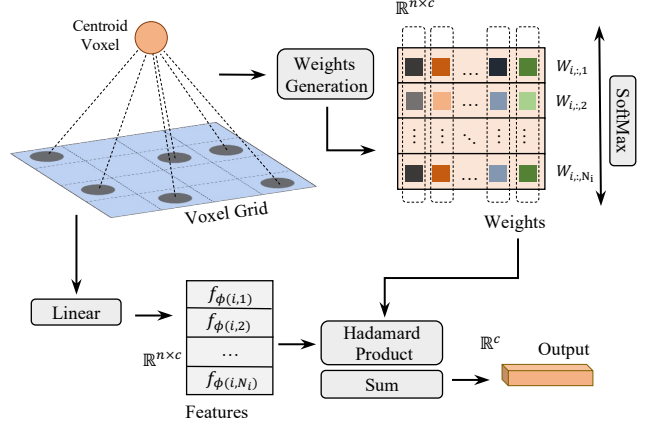


Figure 5. Illustration of the Adaptive Relation Convolution (ARConv). It dynamically generates grid convolution’s kernel weights only for the non-empty voxels with their relationships to the centroid voxel.

3.2. Adaptive relation convolution

Observations. Transformer frameworks [39, 60] have achieved remarkable success and become one of the dominant architectures in 3D semantic segmentation. Their performance superiority is largely related to the ability of relation learning among various local point features. It is achieved by self-attention mechanisms and essentially increases the representation capacity. However, plain sparse CNNs miss this design.

On the other hand, CNNs have verified, via extensive research [6, 33, 57], the importance of large receptive fields to enable a global perception. Unfortunately, 3D convolution struggles to improve perception range by directly expanding the convolution kernel since its complexity is $\mathcal{O}(K^3)$, where K is the kernel size, indicating that the consumption of the large kernel may be unacceptable in practice, especially for the edge devices. To this end, we explore the large-kernel design to be lightweight and propose the Adaptive Relation Convolution (ARConv), which incorporates the aforementioned adaptive relation reasoning into sparse CNNs. More details are illustrated in Fig. 5.

Depthwise convolution. To assemble the framework in a lightweight manner, we could start by considering depthwise convolution for parsing the voxel grid features. In practical applications, it is also found that the depthwise convolution generalizes better [58] and converges faster as shown in our experiments. Compared with regular convolutions performed over multiple input channels, depthwise convolutions independently apply a single convolutional filter for each input channel and keep each channel separate. The output for i -th voxel grid feature $\mathbf{o}_i \in \mathbb{R}^d$ and c -th di-




Figure 6. Comparison of the multi-to-multi paradigm in point transformers with the multi-one-multi paradigm in OA-CNNs.

mension can be precisely described as,

$$\mathbf{o}_{i,c} = \sum_{j=1}^{N_i} \mathbf{W}_{i,c,j} \cdot \mathbf{f}_{\phi(i,j),c}, \quad (4)$$

where N_i is the number of non-empty voxels in the i -th voxel grid \mathcal{V}_i , $\mathbf{W}_i \in \mathbb{R}^{d \times N_i}$ indicates the learnable kernel weight and $\phi(i, j)$ returns the j -th non-empty voxel index in i -th voxel grid.

Adaptive relation kernel. For accomplishing the adaptive relation reasoning, the attention mechanisms [54, 73] adopt the multi-to-multi paradigm, which incorporates a “reference set” [42, 60] for capturing long-range dependencies through multiple queries and keys. However, this approach results in significant inference time and memory demands on GPUs. In contrast, we propose a more efficient multi-one-multi pipeline, generating a single centroid voxel of the grid, which serves as the agent for capturing long-range relationships. This strategy facilitates efficient computation and lowers memory consumption, while still enabling the extraction of complex relationships among the non-empty voxels in the grid. The idea is illustrated in Fig. 6.

Specifically, for the sub voxel grid \mathcal{V}_i , its corresponding centroid voxel feature $\mathbf{f}_i^{ctr} \in \mathbb{R}^d$, where d indicates the number of channels, is formed as:

$$\mathbf{f}_i^{ctr} = \text{AvgPool}(\{\delta_{proj}(\mathbf{f}_j) \mid \mathbf{p}_j \in \Omega(i)\}), \quad (5)$$

where $\text{AvgPool}(\cdot)$ applies 3D average pooling over the input, $\Omega(\cdot)$ indicates the subset’s indices range, and $\delta_{proj} : \mathbb{R}^d \mapsto \mathbb{R}^d$ is a linear projection layer with normalization and activation.

Then the dynamic kernel weight $\mathbf{W}_i \in \mathbb{R}^{d \times N_i}$ of the depthwise convolution for the i -th voxel grid is generated by considering voxels’ feature correlations with the centroid voxel:

$$\mathbf{W}_{i,:,j} = \delta_{weight}(\mathbf{f}_{\phi(i,j)} - \mathbf{f}_i^{ctr}), \quad (6)$$

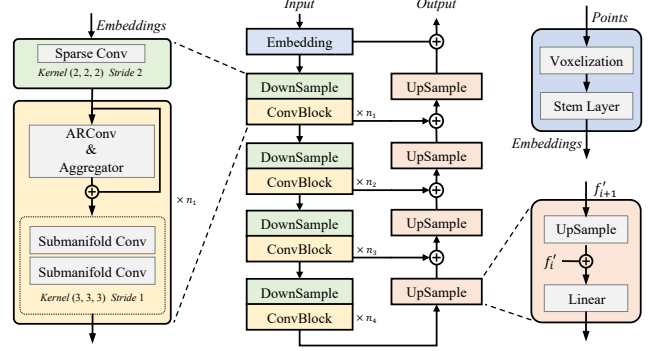


Figure 7. Illustration for the whole architecture and more implementation details.

where $\delta_{weight} : \mathbb{R}^d \mapsto \mathbb{R}^d$ is a linear projection layer, and $\phi(i, j)$ returns the j -th non-empty voxel index in i -th voxel grid.

We normalize the dynamically generated weights $\mathbf{W}_{i,:,j}$ using softmax operation along each channel separately across the whole voxel grid. The normalization enhances the stability of the neural network outputs during training and assigns feature weights based on internal relevance between the specific voxel and the centroid voxel. Mathematically, for the c -th channel,

$$\mathbf{W}'_{i,c,j} = \frac{\exp(\mathbf{W}_{i,c,j} - \text{Max}(\mathbf{W}_{i,:,j}))}{\sum_{k=1}^{N_i} \exp(\mathbf{W}_{i,c,k} - \text{Max}(\mathbf{W}_{i,:,j}))}, \quad (7)$$

where $\text{Max}(\cdot)$ returns the maximum value. We empirically find that the dynamically generated weights were volatile at the early training phase, yielding large values that may cause the exponential function explosion and lead to “inf” outputs. Thus, we adopt an additional operation in Eq. (7) that subtracts the maximum values from the numerator and denominator respectively to prevent the explosion without affecting the output – it is numerically equal to the case without this operation.

In essence, we have introduced an efficient approach named Adaptive Relation Convolution (ARConv) that generates kernel weights only for the non-empty voxels dynamically by considering their correlations to the geometric centroid representatives, thus achieving effectiveness without sacrificing efficiency.

3.3. Architecture

In this section, we provide the architectural details of the OA-CNNs. Fig. 7 depicts the overall structure.

Concretely, the sparse and submanifold voxel modules [13, 36] both process spatially sparse data effectively. The primary difference between them is that submanifold convolution only handles the non-empty voxels in the 3D

scene and strictly preserves the original geometric structure. Differently, sparse convolution can extract features at empty locations and is more flexible. We construct our basic blocks with an ARConv module followed by two submanifold convolutions with necessary normalization and activation layers. Following [40, 42], we adopt the hierarchical structure to the encoder and use a sparse convolution with kernel size and stride that are both set to (2, 2, 2), down-sampling the spacial size to $1/8$ at each time. As for the upsampling process, the up-block only consists of a skip connection and a single linear layer that aligns the feature channel numbers without other components.

4. Experiments

4.1. Implementation details.

Datasets. We conducted experiments using our proposed OA-CNNs on the standard benchmark, ScanNet v2 [11], as well as its recent extension, ScanNet200 [44], and the S3DIS dataset [1] for indoor scenes. ScanNet v2 contains 1,201 training scenes and 312 validation scans reconstructed from RGB-D frames. The model utilizes reconstructed meshes to sample point clouds as input, where each point cloud is attributed a semantic label from a set of 20 categories. ScanNet200 benchmark extends the class categories to 200, an order of magnitude more than the previous. The S3DIS dataset consists of 271 rooms in six areas from three different buildings with 13 categories. Following a standard protocol, area 5 is withheld during training and used for S3DIS testing. As for the outdoor semantic segmentation, we select two popular benchmarks, nuScenes [4] and SemanticKITTI [2]. The nuScenes dataset contains approximately 1000 scenes, with each scene consisting of multiple sensor sweeps captured from a moving vehicle. In contrast, the SemanticKITTI dataset consists of sequences from the raw KITTI dataset, which contains 22 sequences in total. Each sequence includes around 1,000 lidar scans, corresponding to approximately 20,000 individual frames.

Training details. We train our models on 4 RTX 3090 GPUs with the batch size and the number of epochs set to 16 and 100, respectively. With the considerations regarding computational efficiency and memory constraints, the training process leverages a subset of up to 100,000 randomly sampled points from the point cloud. In contrast, the full point cloud is used during validation to ensure an unbiased and rigorous evaluation of the model’s performance. Moreover, we attribute parts of the point-based frameworks’ performance superiority to the modern training strategy with advanced data enhancement [43, 60]. We refer to these strategies to train our models. Specifically, we use the AdamW optimizer [34] for parameter optimization, which is widely used in transformer architectures. The ini-

Method	Input	Val mIoU	Test mIoU
PointNet++ [42]	point	53.5	55.7
PointNeXt-XL [43]	point	71.5	71.2
PointCNN [28]	point	-	45.8
KPConv [48]	point	69.2	68.6
PointConv [59]	point	61.0	66.6
PointTransformer [73]	point	70.6	-
FastPointTransformer [39]	point	72.1	-
Stratified Transformer [21]	point	74.3	73.7
OctFormer [56]	point	75.7	76.6
PTv2 [60]	point	75.4	75.2
SparseUNet [13]	voxel	69.3	72.5
MinkowskiNet [7]	voxel	72.2	73.6
LargeKernel3D [6]	voxel	73.2	73.9
OA-CNNs(ours)	voxel	76.1	75.6

Table 1. We compared semantic segmentation results on ScanNet v2. All the selected methods are under the same experimental settings without the use of additional pretraining or auxiliary methods.

Outdoor Sem. Seg. Method	Benchmarks	
	nuScenes [4]	SemanticKITTI [2]
SparseUNet [13]	73.3	63.8
SPVNAS [47]	77.4	64.7
Cylinder3D [76]	76.1	64.3
SphereFormer [22]	78.4	67.8
OA-CNNs(ours)	78.9	70.6

Table 2. Results on outdoor semantic segmentation benchmarks.

Method	Val				Test
	Head	Comm.	Tail	All	All
MinkowskiNet [7]	48.3	19.1	7.9	25.1	25.3
LGround [44]	51.5	22.7	12.5	28.9	27.2
SparseUNet [61]	-	-	-	28.8	-
OctFormer [56]	-	-	-	-	32.6
PTv2 [61]	-	-	-	29.3	-
OA-CNNs(Ours)	51.3	28.0	17.7	32.3	33.3

Table 3. Results on ScanNet200 for semantic segmentation.

tial learning rate lr is 0.001, and the weight decay is set to 0.02 with the cosine annealing strategy. Following [60] for data preprocessing, we estimate normal vectors for points and add coordinates as additional feature input. As for the data augmentation, we apply random drop, random deformation, and color jitter following [60, 73].

4.2. Comparisons

Performance. We conduct a comprehensive comparison of our proposed OA-CNNs with alternative backbone models on multiple benchmarks, including ScanNet v2, ScanNet200, S3DIS, nuScenes, and SemanticKITTI [1, 2, 4, 11, 44]. All the methods compared in our experiments are evaluated under the same experimental settings, without any ad-

Method	Input	OA	mIoU
PointNet [41]	point	-	41.1
PointTransformer [73]	point	90.8	70.4
PTv2 [60]	point	91.1	71.6
MinkowskiNet [7]	voxel	-	65.4
OA-CNNs(ours)	voxel	90.7	71.1

Table 4. Results on S3DIS area 5 for semantic segmentation.

ID	Enlarge Methods	mIoU (%)
I	Baseline	73.0
II	Multi-head Self-attention	73.5
III	Grouped Vector Attention	74.3
IV	Pyramid Pooling	75.0
V	ARConv	76.1

Table 6. Ablation studies on the different methods commonly used for enlarging receptive fields.

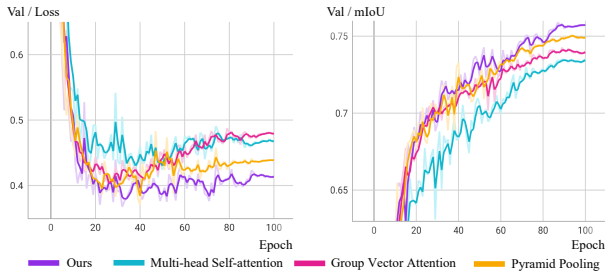


Figure 8. Compared with other classical modules to expand the receptive fields, our proposed method is more stable, has faster convergence during training, and acquires better performance.

ditional pretraining or auxiliary methods. The results are shown in Tabs. 1, 2, 3, 4. Our proposed model exhibits superior performance over prior state-of-the-art point-based frameworks and transformer architectures in both indoor and outdoor scenes. Indeed, these results highlight the superior generalization capability of OA-CNNs, demonstrating their potential to outperform point-based and transformer models in various benchmarks even without any self-attention modules.

4.3. Ablation Study

Efficiency. We also compare our models with various CNN-Based and transformer-based methods [7, 13, 21, 60, 73] regarding accuracy, inference speed, and GPU memory consumption, as shown in Fig. 2. We can observe that, while transformer-based methods have demonstrated impressive performance, they come with a drawback – they require extensive time and memory for frequently querying nearest neighbors, attention computation, and other point-based operations. Differently, thanks to the CNN architec-

ID	Aggregation	Stage Num	mIoU (%)	Δ
I	w/o	1	75.0	+ 0.0
II	Concatenation	3	75.2	+ 0.2
III	Adaptive (ours)	2	75.2	+ 0.2
IV	Adaptive (ours)	3	76.1	+ 1.1

Table 5. Effectiveness of adaptive aggregator and naive concatenation through ablation studies with varying stage numbers.

ID	Conv	Groups	mIoU (%)	ID	Type	mIoU (%)
I	Grouped	[2, 2, 4, 8]	75.0	I	Pos	75.3
II	Grouped	[4, 4, 8, 16]	75.4	II	Pos+Ctr	75.9
III	Depthwise	-	76.1	III	Ctr	76.1

Table 7. Performance Comparison of Depthwise Convolution and Regular Grouped Convolution.

Table 8. Comparison of various weight generation methods.

Type	Blocks	Time (ms)	Mem. (G)	mIoU (%)
OA-CNN (S)	[2, 2, 2, 2]	117	2.1	73.6
OA-CNN (B)	[3, 3, 9, 3]	190	3.3	75.3
OA-CNN (L)	[3, 3, 9, 8]	213	3.6	76.1

Table 9. Comparison between various versions of our proposed models. The channels for each stage are set to [64, 64, 128, 256] and kept the same.

ture that exploits the structural data arrangement and hash acceleration to attain notable efficiency and low memory consumption, our method takes the performance lead but still preserves a superior balance between effectiveness and efficiency.

Receptive field expansion. We verify the effectiveness of our proposed Adaptive Relation Convolution (ARConv) by the comparison with three alternative modules commonly used for receptive field expansion: 1) multi-head self-attention [54]; 2) grouped vector attention [60]; and 3) pyramid pooling [72].

For attention-based modules, we operate the voxels like nearest neighbor finding and grouping following the point transformer [73]. The test results are shown in Tab. 6, where our ARConv outperforms other competitors. Moreover, Fig. 8 presents the comparison of the validation loss/mIoU during the training process, and ARConv exhibits a superior capacity for mitigating overfitting than the others, as evidenced by the lack of considerable deterioration in validation loss during the later period of training.

Aggregation methods. We validate the effectiveness and superiority of the pyramid grid partition and proposed adaptive aggregator, and the experimental results are shown in

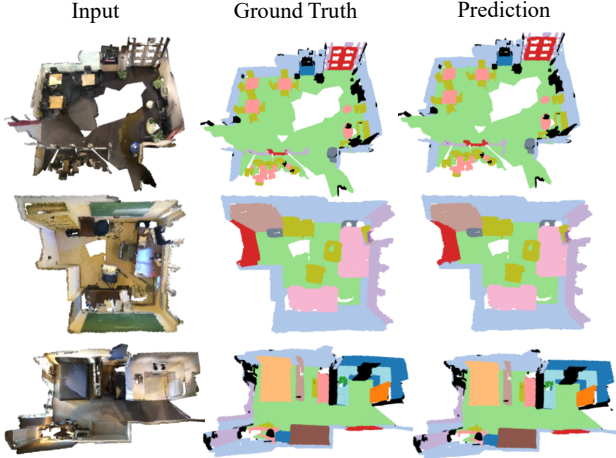


Figure 9. Visualization of segmentation results on ScanNet v2.

Tab. 5. The first row shows the result of the model with the single-scale partition, where the additional aggregation is not necessary. The second experiment, which adopts direct concatenation for aggregation, leads to a marginal improvement in performance. Then, by introducing our proposed adaptive aggregator that adjusts the receptive field of each voxel based on its intrinsic properties, we observed a significant improvement in performance as compared to using concatenation. Also, we investigate the effects of the number of pyramid stages and find that the three stages acquire the best result, and all experiments follow this configuration without otherwise specified.

Depthwise convolution. In contrast to regular convolution that applies one filter $\mathbf{W} \in \mathbb{R}^{c \times l \times c}$, where c represents the channels and l represents the inputs' length, across all input channels, depthwise convolution applies a single filter for each input channel independently. Initially, we attempted to implement regular convolution with the proposed dynamic kernel weights but found it to be unstable and non-convergent, particularly during the early training stages. Consequently, we replaced it with both grouped convolution [8] and depthwise convolution. The outcomes are presented in Tab. 7. Our adoption of depthwise convolution with dynamically generated weights $\mathbf{W} \in \mathbb{R}^{l \times c}$ yields linear complexity to input channels, showcasing the dual benefits of efficiency and performance.

Dynamic kernel weights. Previous point-based methods, such as [48, 64], have also explored dynamically generated kernel weights. However, their approaches aim to incorporate geometric information from points rather than local semantic relationships, mimicking convolutional operations while still following the PointNet paradigm. Differently, our work is based on sparse convolution networks. We assess the effectiveness of our design for dynamic kernel weight generation by comparing it with other alternatives in Tab. 8, where ctr represents our adaptive relation kernel

and pos denotes kernel weight generation using relative positions. Our adaptive relation kernel demonstrates superior performance to other methods.

Multiple versions. We present multiple versions of OA-CNNs, achieved by adjusting the number of blocks in each stage while keeping other configurations consistent. In all models, the number of channels per block is set to [64, 64, 128, 256]. The impact on performance and efficiency is demonstrated in Tab. 9, where all models are evaluated on a single RTX 3090 to ensure a fair comparison.

4.4. Visual Analysis

Predictions. The qualitative results of point cloud semantic segmentation are presented in Fig. 9. Our model exhibits exceptional predictive accuracy on the ScanNet v2 dataset, with results that demonstrate high consistency with the ground truth.

Receptive fields. Fig. 1 visualizes the varying receptive field sizes of different objects and parts with distinct geometric structures and appearances within a 3D indoor scene. We calculate the sizes of the receptive fields as follows:

$$r_i = \sum_{k=1}^K w_{i,k} g_k, \quad (8)$$

where g_k represents the k -th grid size and $w \in \mathbb{R}^{n \times K}$ indicates preference weights predicted by the learnable adaptive aggregator in Eq. (2). We then map different sizes r_i to the corresponding colors. Fig. 1 confirms our intuition that 3D scenes' flatter areas with simplistic structures, such as walls and floors, require larger receptive fields. Conversely, smaller objects and more intricate areas, such as edges and junctions, need smaller ones. Additionally, we observe that the floor generally requires a smaller receptive field than the wall and ceiling, since it is necessary to exploit more local contexts to distinguish itself from the objects placed on the floor. More visual comparisons of the receptive field are put in the supplementary materials.

5. Conclusion

This study highlights the potential of sparse convolution networks to surpass transformer architectures in both efficiency and performance. To achieve this, we introduce omni-adaptive 3D CNNs (OA-CNNs), which consist of two key components: spatially dynamic receptive fields and adaptive relation convolution. As for limitations, the current pyramid grid sizes are set empirically, highlighting the need for future research to develop more scientifically and logically grounded search algorithms.

Acknowledgements. This work receives partial support from the Shenzhen Science and Technology Program under No. KQTD20210811090149095.

References

- [1] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1534–1543, 2016. **6**
- [2] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9297–9307, 2019. **2, 6**
- [3] Yizhak Ben-Shabat, Michael Lindenbaum, and Anath Fischer. 3dmfv: Three-dimensional point cloud classification in real-time using convolutional neural networks. *IEEE Robotics and Automation Letters*, 3(4):3145–3152, 2018. **3**
- [4] 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. **2, 6**
- [5] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3d object detection network for autonomous driving. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 1907–1915, 2017. **3**
- [6] Yukang Chen, Jianhui Liu, Xiaojuan Qi, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Scaling up kernels in 3d cnns. *CoRR*, abs/2206.10555, 2022. **2, 4, 6**
- [7] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3075–3084, 2019. **1, 2, 6, 7**
- [8] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016. **8**
- [9] Pointcept Contributors. Pointcept: A codebase for point cloud perception research. <https://github.com/Pointcept/Pointcept>, 2023. **1**
- [10] Jiequan Cui, Shu Liu, Zhuotao Tian, Zhisheng Zhong, and Jiaya Jia. Reslt: Residual learning for long-tailed recognition. *TPAMI*, 2023. **1**
- [11] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. **2, 6**
- [12] 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. *arXiv preprint arXiv:2010.11929*, 2020. **1**
- [13] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9224–9232, 2018. **1, 2, 3, 5, 6, 7**
- [14] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11108–11117, 2020. **2**
- [15] Qianguo Huang, Weiyue Wang, and Ulrich Neumann. Recurrent slice networks for 3d segmentation of point clouds. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2626–2635, 2018. **2**
- [16] Zhening Huang, Xiaoyang Wu, Xi Chen, Hengshuang Zhao, Lei Zhu, and Joan Lasenby. Openins3d: Snap and lookup for 3d open-vocabulary instance segmentation. *arXiv preprint arXiv:2309.00616*, 2023. **1**
- [17] Xu Jia, Bert De Brabandere, Tinne Tuytelaars, and Luc V Gool. Dynamic filter networks. *Advances in neural information processing systems*, 29, 2016. **3**
- [18] Li Jiang, Hengshuang Zhao, Shu Liu, Xiaoyong Shen, Chi-Wing Fu, and Jiaya Jia. Hierarchical point-edge interaction network for point cloud semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10433–10441, 2019. **2**
- [19] Li Jiang, Shaoshuai Shi, Zhuotao Tian, Xin Lai, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *ICCV*, 2021. **1**
- [20] Xin Lai, Zhuotao Tian, Li Jiang, Shu Liu, Hengshuang Zhao, Liwei Wang, and Jiaya Jia. Semi-supervised semantic segmentation with directional context-aware consistency. In *CVPR*, 2021. **1**
- [21] Xin Lai, Jianhui Liu, Li Jiang, Liwei Wang, Hengshuang Zhao, Shu Liu, Xiaojuan Qi, and Jiaya Jia. Stratified transformer for 3d point cloud segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8500–8509, 2022. **1, 2, 3, 6, 7**
- [22] Xin Lai, Yukang Chen, Fanbin Lu, Jianhui Liu, and Jiaya Jia. Spherical transformer for lidar-based 3d recognition. In *CVPR*, 2023. **6**
- [23] Xin Lai, Zhuotao Tian, Yukang Chen, Yanwei Li, Yuhui Yuan, Shu Liu, and Jiaya Jia. LISA: reasoning segmentation via large language model. *arXiv preprint arXiv:2308.00692*, 2023. **1**
- [24] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. **3**
- [25] Felix Järemo Lawin, Martin Danelljan, Patrik Tosteberg, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Deep projective 3d semantic segmentation. In *Computer Analysis of Images and Patterns: 17th International Conference, CAIP 2017, Ystad, Sweden, August 22-24, 2017, Proceedings, Part I 17*, pages 95–107. Springer, 2017.
- [26] Bo Li, Tianlei Zhang, and Tian Xia. Vehicle detection from 3d lidar using fully convolutional network. *arXiv preprint arXiv:1608.07916*, 2016. **3**

- [27] Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 510–519, 2019. 3
- [28] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems*, 31, 2018. 6
- [29] Yiqun Lin, Zizheng Yan, Haibin Huang, Dong Du, Ligang Liu, Shuguang Cui, and Xiaoguang Han. Fpconv: Learning local flattening for point convolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4293–4302, 2020. 3
- [30] Jianhui Liu, Yukang Chen, Xiaoqing Ye, Zhuotao Tian, Xiao Tan, and Xiaojuan Qi. Spatial pruned sparse convolution for efficient 3d object detection. In *NeurIPS*, 2022. 1
- [31] Yongcheng Liu, Bin Fan, Gaofeng Meng, Jiwen Lu, Shiming Xiang, and Chunhong Pan. Densepoint: Learning densely contextual representation for efficient point cloud processing. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 5239–5248, 2019. 2
- [32] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 1
- [33] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11976–11986, 2022. 4
- [34] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6
- [35] Xiaoliu Luo, Zhuotao Tian, Taiping Zhang, Bei Yu, Yuan Yan Tang, and Jiaya Jia. Pfenet++: Boosting few-shot semantic segmentation with the noise-filtered context-aware prior mask. *TPAMI*, 2024. 1
- [36] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel transformer for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3164–3173, 2021. 1, 5
- [37] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 922–928. IEEE, 2015. 1, 3
- [38] Hsien-Yu Meng, Lin Gao, Yu-Kun Lai, and Dinesh Manocha. Vv-net: Voxel vae net with group convolutions for point cloud segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 8500–8508, 2019. 3
- [39] Chunghyun Park, Yoonwoo Jeong, Minsu Cho, and Jaesik Park. Fast point transformer. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16949–16958, 2022. 1, 4, 6
- [40] Bohao Peng, Zhuotao Tian, Xiaoyang Wu, Chengyao Wang, Shu Liu, Jingyong Su, and Jiaya Jia. Hierarchical dense correlation distillation for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23641–23651, 2023. 6
- [41] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 1, 2, 7
- [42] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 1, 2, 5, 6
- [43] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Abed Al Kader Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *arXiv preprint arXiv:2206.04670*, 2022. 6
- [44] David Rozenberszki, Or Litany, and Angela Dai. Language-grounded indoor 3d semantic segmentation in the wild. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 6
- [45] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1746–1754, 2017. 1, 3
- [46] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE international conference on computer vision*, pages 945–953, 2015. 3
- [47] Haotian* Tang, Zhijian* Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3d architectures with sparse point-voxel convolution. In *European Conference on Computer Vision*, 2020. 6
- [48] 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 *Proceedings of the IEEE/CVF international conference on computer vision*, pages 6411–6420, 2019. 3, 6, 8
- [49] Zhuotao Tian, Michelle Shu, Pengyuan Lyu, Ruiyu Li, Chao Zhou, Xiaoyong Shen, and Jiaya Jia. Learning shape-aware embedding for scene text detection. In *CVPR*, 2019. 1
- [50] Zhuotao Tian, Xin Lai, Li Jiang, Shu Liu, Michelle Shu, Hengshuang Zhao, and Jiaya Jia. Generalized few-shot semantic segmentation. In *CVPR*, 2022. 1
- [51] Zhuotao Tian, Hengshuang Zhao, Michelle Shu, Zhicheng Yang, Ruiyu Li, and Jiaya Jia. Prior guided feature enrichment network for few-shot segmentation. *TPAMI*, 2022. 1
- [52] Zhuotao Tian, Pengguang Chen, Xin Lai, Li Jiang, Shu Liu, Hengshuang Zhao, Bei Yu, Ming-Chang Yang, and Jiaya Jia. Adaptive perspective distillation for semantic segmentation. *TPAMI*, 2023. 1

- [53] Zhuotao Tian, Jiequan Cui, Li Jiang, Xiaojuan Qi, Xin Lai, Yixin Chen, Shu Liu, and Jiaya Jia. Learning context-aware classifier for semantic segmentation. In *AAAI*, 2023. [1](#)
- [54] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. [1](#), [2](#), [3](#), [5](#), [7](#)
- [55] Chengyao Wang, Li Jiang, Xiaoyang Wu, Zhuotao Tian, Bohao Peng, Hengshuang Zhao, and Jiaya Jia. Groupcontrast: Semantic-aware self-supervised representation learning for 3d understanding. *arXiv preprint arXiv:2403.09639*, 2024. [1](#)
- [56] Peng-Shuai Wang. Octformer: Octree-based transformers for 3d point clouds. *arXiv preprint arXiv:2305.03045*, 2023. [6](#)
- [57] Sanghyun Woo, Shoubhik Debnath, Ronghang Hu, Xinlei Chen, Zhuang Liu, In So Kweon, and Saining Xie. Convnext v2: Co-designing and scaling convnets with masked autoencoders. *arXiv preprint arXiv:2301.00808*, 2023. [4](#)
- [58] Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. Pay less attention with lightweight and dynamic convolutions. *arXiv preprint arXiv:1901.10430*, 2019. [2](#), [3](#), [4](#)
- [59] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, pages 9621–9630, 2019. [3](#), [6](#)
- [60] Xiaoyang Wu, Yixing Lao, Li Jiang, Xihui Liu, and Hengshuang Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. *arXiv preprint arXiv:2210.05666*, 2022. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#)
- [61] Xiaoyang Wu, Xin Wen, Xihui Liu, and Hengshuang Zhao. Masked scene contrast: A scalable framework for unsupervised 3d representation learning. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 2023. [6](#)
- [62] Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong He, and Hengshuang Zhao. Point transformer v3: Simpler, faster, stronger. In *CVPR*, 2024. [2](#)
- [63] Xiaoyang Wu, Zhuotao Tian, Xin Wen, Bohao Peng, Xihui Liu, Kaicheng Yu, and Hengshuang Zhao. Towards large-scale 3d representation learning with multi-dataset point prompt training. In *CVPR*, 2024. [2](#)
- [64] Mutian Xu, Runyu Ding, Hengshuang Zhao, and Xiaojuan Qi. Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3173–3182, 2021. [3](#), [8](#)
- [65] Yan Yan, Yuxing Mao, and Bo Li. Second: Sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. [1](#)
- [66] Brandon Yang, Gabriel Bender, Quoc V Le, and Jiquan Ngiam. Condconv: Conditionally parameterized convolutions for efficient inference. *Advances in Neural Information Processing Systems*, 32, 2019. [3](#)
- [67] Honghui Yang, Sha Zhang, Di Huang, Xiaoyang Wu, Haoyi Zhu, Tong He, Shixiang Tang, Hengshuang Zhao, Qibo Qiu, Binbin Lin, Xiaofei He, and Wanli Ouyang. Unipad: A universal pre-training paradigm for autonomous driving. In *CVPR*, 2024. [2](#)
- [68] Senqiao Yang, Tianyuan Qu, Xin Lai, Zhuotao Tian, Bohao Peng, Shu Liu, and Jiaya Jia. An improved baseline for reasoning segmentation with large language model. *arXiv preprint arXiv:2312.17240*, 2023. [2](#)
- [69] Senqiao Yang, Jiarui Wu, Jiaming Liu, Xiaoqi Li, Qizhe Zhang, Mingjie Pan, Yulu Gan, Zehui Chen, and Shanghang Zhang. Exploring sparse visual prompt for domain adaptive dense prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 16334–16342, 2024. [2](#)
- [70] Yunhan Yang, Xiaoyang Wu, Tong He, Hengshuang Zhao, and Xihui Liu. Sam3d: Segment anything in 3d scenes. *arXiv preprint arXiv:2306.03908*, 2023. [1](#)
- [71] Dong Zhang, Yi Lin, Hao Chen, Zhuotao Tian, Xin Yang, Jinhui Tang, and Kwang-Ting Cheng. Deep learning for medical image segmentation: Tricks, challenges and future directions. *arXiv preprint arXiv:2209.10307*, 2022. [1](#)
- [72] Hengshuang Zhao, Jianping Shi, Xiaojuan Qi, Xiaogang Wang, and Jiaya Jia. Pyramid scene parsing network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2881–2890, 2017. [7](#)
- [73] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [74] Zhisheng Zhong, Jiequan Cui, Yibo Yang, Xiaoyang Wu, Xiaojuan Qi, Xiangyu Zhang, and Jiaya Jia. Understanding imbalanced semantic segmentation through neural collapse. 2023. [1](#)
- [75] Haoyi Zhu, Honghui Yang, Xiaoyang Wu, Di Huang, Sha Zhang, Xianglong He, Tong He, Hengshuang Zhao, Chunhua Shen, Yu Qiao, and Wanli Ouyang. Ponderv2: Pave the way for 3d foundation model with a universal pre-training paradigm. *arXiv preprint arXiv:2310.08586*, 2023. [1](#)
- [76] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. *arXiv preprint arXiv:2011.10033*, 2020. [6](#)