

LargeKernel3D: Scaling up Kernels in 3D Sparse CNNs

Yukang Chen^{1*} Jianhui Liu^{2*} Xiangyu Zhang³ Xiaojuan Qi² Jiaya Jia^{1,4}

¹The Chinese University of Hong Kong ²The University of Hong Kong ³MEGVII ⁴SmartMore

Abstract

Recent advance in 2D CNNs has revealed that large kernels are important. However, when directly applying large convolutional kernels in 3D CNNs, severe difficulties are met, where those successful module designs in 2D become surprisingly ineffective on 3D networks, including the popular depth-wise convolution. To address this vital challenge, we instead propose the spatial-wise partition convolution and its large-kernel module. As a result, it avoids the optimization and efficiency issues of naive 3D large kernels. Our large-kernel 3D CNN network, LargeKernel3D, yields notable improvement in 3D tasks of semantic segmentation and object detection. It achieves 73.9% mIoU on the ScanNetv2 semantic segmentation and 72.8% NDS nuScenes object detection benchmarks, ranking 1st on the nuScenes LIDAR leaderboard. The performance further boosts to 74.2% NDS with a simple multi-modal fusion. In addition, LargeKernel3D can be scaled to $17 \times 17 \times 17$ kernel size on Waymo 3D object detection. For the first time, we show that large kernels are feasible and essential for 3D visual tasks. Our code and models is available at github.com/dvlab-research/LargeKernel3D.

1. Introduction

3D Sparse convolutional neural networks (CNNs) have been widely used as feature extractors in 3D tasks, *e.g.*, semantic segmentation [9,24] and object detection [55,65,75]. The advantages of efficiency and convenient usage ensure its important role in various applications, such as autonomous driving and robotics. However, 3D sparse CNNs are recently challenged by transformer-based methods [45,46,79], mainly from the aspect of building effective receptive fields. Both global and local [21,45] self-attention mechanisms are able to capture context information from a large spatial scope. 2D Vision Transformers (ViTs) also emphasize their advantages in modeling long-range dependencies [20,42,51]. In contrast, common 3D sparse CNNs are limited in this regard. It is because the receptive fields

of default 3D sparse CNN are constrained by small kernel sizes and spatial disconnection of sparse features (due to the property of submanifold sparse convolution [25]).

Literature about 2D CNNs [18,43,62] presents a series of methods, combined with large kernels, to enlarge the receptive fields and model capacity. ConvNeXt [43] employs 7×7 depth-wise convolution as a strong design, combing with other training techniques to challenge its Swin Transformer counterpart [42]. RepLKNet [18] pursues extremely large kernel sizes of 31×31 to boost the performance of different tasks. To ensure the effectiveness of RepLKNet [18], additional factors, including depth-wise convolution, are also required. Other work [27] also emphasizes the importance of depth-wise convolution. Due to differences between 3D and 2D tasks, these methods, however, are found not a good solution for 3D sparse CNNs.

We first analyze the difficulties of 3D large-kernel CNN design in two aspects. The first challenge is *efficiency*. It is easy to understand that 3D convolution is with the cubic kernel size and computation increases fast. For example, the model size increases 10+ times when kernels change from $3 \times 3 \times 3$ to $7 \times 7 \times 7$. The second difficulty exists in the *optimization* procedure. 3D datasets may contain only thousands of scenes, which cannot match 2D image benchmarks [15,40] in terms of scales. In addition, 3D point clouds or voxels are sparse, instead of dense images. Thus, it might be insufficient to optimize the proliferated parameters of large kernels and leads to over-fitting.

In this paper, we propose spatial-wise partition convolution as the 3D large-kernel design. It is a new family of group convolution by sharing weights among spatially adjacent locations, rather than depth-wise convolution [29] of channel-level groups. As shown in Fig. 1, spatial-wise partition convolution remaps a large kernel (*e.g.*, 7×7) as a small one (*e.g.*, 3×3) via grouping spatial neighbors, while the absolutely large spatial size remains unchanged. With regard to the *efficiency* issue, it occupies few model sizes to keep parameters the same as those of small kernels. Moreover, it takes less latency, compared with plain large kernel counterparts. As for the *optimization* challenge, weight-sharing among spatial dimensions gives parameters more chance to update and overcome the over-fitting issue.

*Equal Contribution.

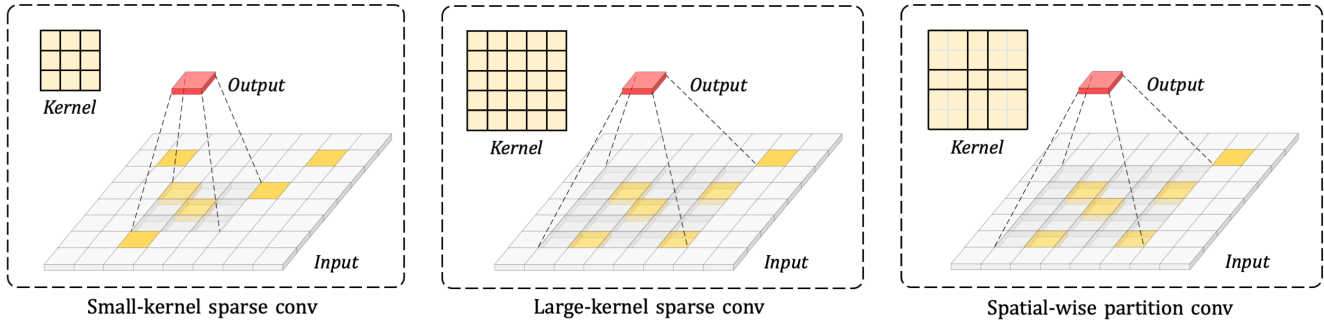


Figure 1. Sparse convolutions with different kernels. *Small-kernel sparse convolution* gathers features in a local area. It is efficient but discards sufficient information flow due to feature disconnection and the small scope. *Large-kernel sparse convolution* is capable of capturing long-range information, at the price of a large number of parameters and computation. Our proposed *spatial-wise partition convolution* uses large kernel sizes, and shares weights among local neighbors for efficiency. We show 2D features for the simplicity sake.

To increase the detail-capturing ability of large kernels [18], we introduce position embeddings for spatial-wise group convolution. It makes notable effects for large kernel sizes. We name the proposed block as spatial-wise large-kernel convolution (*SW-LK Conv*). We compare the efficiency between plain 3D submanifold sparse convolution and ours, as shown in Tab. 1. Both parameters and latency of the baseline increases dramatically, as its kernel size becomes larger, while ours is far more efficient.

SW-LK Conv can readily replace plain convolution layers in existing 3D convolutional networks. We establish large-kernel backbone networks LargeKernel3D on existing 3D semantic segmentation [9] and object detection [16, 75] networks. It achieves notable improvement upon state-of-the-art methods [9, 16, 75], with a small model complexity overhead. Extensive experiments validate our effectiveness on large-scale benchmarks, including ScanNetv2 [13], nuScenes [4], and Waymo [57]. For object detection, LargeKernel3D achieves **72.8%** NDS on nuScenes, ranking **1st** on the nuScenes LIDAR leaderboard. Without bells and whistles, it further improves to **74.2%** NDS in a simple voxel-wise multi-modal fusion manner. More importantly, it is scalable to $17 \times 17 \times 17$ kernel sizes on the large-scale Waymo 3D object detection.

We visualize the *Effective Receptive Fields* (ERFs) of plain 3D CNNs and our LargeKernel3D in Fig. 2. It shows that deep small-kernel networks are also constrained by limited ERFs, since sparse features are spatially disconnected. Note that our large-kernel networks elegantly resolve this issue. For the first time, we show that large-kernel CNN designs become effective on essential 3D visual tasks.

2. Related Work

Convolutional Large-kernel Networks Large-kernel setting has been commonly used in 2D convolutional net-

works [18, 43, 68]. ConvNeXt [43] combines various training techniques and large kernel sizes, to compete with Swin Transformer [42]. RepLKNet [18] is a specific work that focuses on large-kernel designs. It reveals that the large kernel sizes are more beneficial to tasks of object detection and semantic segmentation than image classification. GCNs [47] also show improvement from large-kernel designs on semantic segmentation. There are methods approximating large kernel sizes using implicit techniques, including Fourier domain transformation [52], and continuous functions [53]. Other methods change kernel shapes to enlarge receptive fields using dilated convolutions [5, 76], deformable convolutions [14, 81] and active convolutions [31]. These schemes sparsify kernels, and is hard to capture the original sparse features in 3D deep learning.

3D Feature Extractors A key challenge in 3D tasks [32, 38] is to learn effective representation from the sparse and non-uniformed 3D geometric data, *e.g.*, point clouds. In general, there are two kinds of 3D feature extractors. The first is to learn on point cloud directly, using a series of PointNet networks [48, 49]. The sampling and grouping operations in PointNet++ [49] are also time-consuming. Follow-ups [46, 79] use grid-sampling or scene-partition to save computation in each forward and combine results from multiple runs. The second is to process point clouds with voxelization and apply 3D sparse CNNs. Due to its efficiency advantages, this stream of methods has been widely used in various 3D tasks, such as 3D segmentation [9–11, 33] and backbone networks in 3D object detectors [16, 26, 55, 64].

The spatial group convolution (SGC) [78] is relevant to the proposed spatial-wise partition convolution. SGC is originally designed for the task of 3D scene completion. They both conduct the partition operation spatially. The difference is also essential and clear. Our method partition kernel weighs while SGC splits input features into groups. In terms of targets, SGC focuses more on efficiency improve-

Table 1. Efficiency comparison between plain 3D sparse convolution and our SW-LK Conv. The baseline is a submanifold sparse convolutional layer with input and output channels 16. The input data is a sparse tensor with 80,000 voxels randomly scattered and batch size 1. We test all latencies on a single 2080Ti GPU by an average of 10 times running, after a warm-up start of additional 10 times.

Kernel Size		3×3×3	5×5×5	7×7×7	9×9×9	11×11×11	13×13×13	15×15×15	17×17×17
Plain	Params	6.9 K	32.0 K	87.8 K	186.6 K	340.7 K	562.4 K	864.0 K	1.3 M
	Latency	2.5 ms	4.2 ms	8.9 ms	17.5 ms	31.1 ms	55.1 ms	81.1 ms	106.3 ms
Ours	Params	-	8.9 K	12.4 K	18.6 K	28.2 K	42.1 K	60.9 K	85.5 K
	Latency	-	3.4 ms	3.9 ms	4.8 ms	6.2 ms	8.4 ms	11.4 ms	15.8 ms

ment by changing feature sparsity in each spatial group, while our method is to facilitate large kernels in 3D tasks. Further analysis and comparisons are given in experiments.

Vision Transformer Vision transformers that conduct attention computation in local areas or windows share a similar spirit with large-kernel models [18]. Swin Transformers [42] capture features in shifted windows with sizes from 7 to 12. Its variants [19, 41] show that larger window sizes are beneficial to performance. Focal Transformer [71] captures fine-grained local attention with adaptive patch sizes. SST [21] and Stratified Transformer [34] apply window-based self-attention on 3D detection and segmentation.

Position Encoding Since the self-attention in transformers is permutation-equivalent, position encoding is well-designed for imbuing the network with the position information. It can be grouped into two streams of absolute and relative position encoding. For the first group, positions are encoded in a precise manner. Specifically, encoding [63] is generated with the sinusoidal functions of different frequencies. They are then added to the input. In contrast, relative position encoding [54] takes the relative relationship into consideration. It can naturally handle longer sequences and is more scalable. Most work [3, 50, 66, 73] computes the relative distance between tokens and applies implicit encoding using learnable transformations. Besides, the literature [12] shows that convolution with padding can provide effective relative position information for relative position encoding. In addition to transformers, the recent work [72] also shows that position embedding is beneficial to the CNNs with fourier transformation.

Considering that our spatial-wise partition convolution divides the kernel into different portions in pursuit of higher efficiency. The position information may be blurred due to weight sharing in each portion. We thus use relative position encoding as a bias to supplement the lost position information. More details are provided in Sec. 4.

3. Revisiting 3D Sparse CNNs

Preliminary 3D sparse CNNs consist of 3D sparse convolutions, which are typically regular and submanifold sparse

convolutions [25]. We formulate 3D sparse convolutions in Eq. (1). Given a set of sparse input features $\{x_{p \in P}\}$ with a number of c_{in} channels, the position of each feature $p \in P$ sparsely distributes in 3D space. We process these features by a convolution with kernel weights $w \in \mathbb{R}^{|K| \times c_{in} \times c_{out}}$. For example, in the 3D coordinate space, w contains $c_{in} \times c_{out}$ spatial kernels with size 3 and $|K| = 27$. The convolution process at the output position \bar{p} is represented as

$$y_{\bar{p}} = \sum_{k \in K} w_k \cdot x_{\bar{p}+k}, \tag{1}$$

where k is an 3D offset distance from \bar{p} . $\bar{p}+k$ is the location around center \bar{p} . It enumerates all discrete locations in the kernel space K where sparse features $x_{\bar{p}+k}$ exist.

Convolution Types Regular sparse convolutions enlarge their output positions by dilating each input position $p \in P$ to the kernel shape K . It dramatically decreases feature density and increases the computational burden. Thus, regular sparse convolutions with stride 2 are only employed at the first layer in each stage for down-sampling. Submanifold sparse convolutions, in contrast, fix their output positions \bar{p} identical to input positions $p \in P$, which preserves the computation cost at low levels. Please refer to the original paper [25] for more details.

Due to the efficiency advantage, submanifold sparse convolutions dominate most layers of 3D CNNs, except for the down-sampling layers. However, limited by the small local scope, it misses sufficient information flow for the spatially disconnected features. Increasing kernel sizes becomes a potential solution to relieve this issue.

Obstacles in 3D Large Kernels *Efficiency* is the first issue in 3D large-kernel CNNs. When we increase kernel sizes, the amount of parameters and computational burden grows much faster than those of 2D CNNs. For instance, the parameter amount in one 3D convolutional layer increases from $27 \cdot C_{in \cdot out}$ to $343 \cdot c_{in} \cdot c_{out}$, given its kernel size changing from 3 to 7. It is thus clear that naively enlarging 3D kernels is unreasonable.

Optimization difficulty is the second obstacle. The booming parameters requires sufficient data for learning. How-




Figure 2. *Effective Receptive Fields* (ERF) of plain 3D CNN, plain 3D CNN - $2\times$, and our LargeKernel3D for 3D object detection. The plain 3D CNN backbone has insufficient ERF size, which hardly covers a nearby car. Plain 3D CNN - $2\times$ is a deeper version with its layers doubled in each stage. More layers help little in ERF, as the disconnection between sparse features. Our LargeKernel3D obtains large ERF at both the object centers (left) and edges (right), capturing context information. More illustrations are provided in the *appendix*. The figure is best viewed in color and by zoom-in.

ever, 3D datasets are commonly not in that large scale as 2D benchmarks. For example, ImageNet [15] contains millions of images, while 3D datasets commonly contain only no more than one thousand scenes. The over-fitting issue gradually deteriorates the performance when increasing kernel sizes from $3\times 3\times 3$ to $7\times 7\times 7$ in MinkowskiNet-34 [9].

Previous 2D knowledge is not helpful in 3D large-kernel CNNs. Recent 2D methods [18, 43] show some helpful components in large-kernel CNNs, including depth-wise convolutions [29], layer normalization [1], GELU [28] activation. We examine these popular components and find that they are either ineffective or even harmful in 3D CNNs. To demonstrate this finding, we use MinkowskiNet-34 [9] to benchmark and train on the 3D semantic segmentation dataset of ScanNetv2 [13]. We change the original activation function from ReLU to GELU, batch normalization to layer normalization, and plain convolution to depth-wise operation, respectively. Directly using these components makes the final system yield performance drop on both small and large kernel networks. It requires us to seek other ways to construct 3D large-kernel CNNs.

4. 3D Large-kernel Convolutional Network

4.1. Spatial-wise Partition Convolution

A standard convolution kernel can be viewed as a 3D matrix, which consists of input channels C_{in} , output channels C_{out} , and the spatial kernel dimension $|K|$. Taking kernel size as k , the volume of spatial kernel dimension $|K| = k\times k$ for 2D convolutions and $k\times k\times k$ for 3D convolutions. Depth-wise convolutions share weights along channel dimensions, where group numbers equal to input channels. Point-wise convolutions fix kernel size as 1, which is common as a conjunction layer to depth-wise convolutions to adjust output channels.

Different from 2D methods, which rely on depth-wise convolutions for performance boosting [27, 43] and accuracy-FLOPs trade-offs [18], we empirically find that




Figure 3. Illustration on spatial-wise partition convolution. In the plain large-kernel convolution, due to the sparsity of input features, only a small proportion of weights are involved and updated in each training step. The spatial-wise partition convolution relieves this issue by sharing weights among spatial neighbors. It increases the chance for weights to be optimized. During inference, we reformulate the spatial-wise partition convolutions into small ones, with feature assignment in a large scope.

depth-wise convolutions are not beneficial to 3D tasks, no matter whether point-wise convolutions are included.

We instead propose spatial-wise partition convolution for 3D large-kernel CNNs. It shares weights among spatial dimension K on convolutional kernels, instead of among channel dimension. It is also different from SGC [78], which partitions spatial groups on input features. Specially, in Eq. (1), w_k shares the same values in a local area $k \in K$. We thus group the original large-kernel convolution from 7×7 into 3×3 , by sharing weights among spatial neighbors. Spatial-wise partition convolution is specially designed for 3D tasks and has the following advantages in terms of efficiency and performance.

Regarding *efficiency*, when naively enlarging the kernel size from $3\times 3\times 3$ to $7\times 7\times 7$ on MinkowskiNet-34 [9], both model size and latency have a sharp increase by multiple times. In contrast, our spatial-wise partition convolution avoids increasing parameters and introduces limited latency overhead. The advantage of our method over plain large kernels attributes to the specific implementation of sparse




Figure 4. Structure of spatial-wise large-kernel convolution (SW-LK Conv). It consists of a large-kernel spatial-wise partition convolution and a learnable position embedding. The position embedding is used to make up the detail-capturing ability of large kernels.

convolution. For a spatial-wise partition convolution, we directly use the small-kernel layer during inference, and enlarge its feature-assign areas to the large-kernel scope. As shown in Fig. 3, thanks to the weight-sharing operation, it much saves the multiplication, *e.g.*, from 343 to 27 times. We implement this operation upon the open-sourced *spconv* library and omit the details here. Tab. 4 shows that our implementation on spatial-wise partition convolution is far more efficient than the plain large-kernel baseline. The *optimization* difficulty of large-kernel CNNs is relieved with the help of spatial-wise partition convolutions, which constrains model sizes to a limited number. It avoids requiring a large amount of data for training and the over-fitting issue.

4.2. Kernel-wise Position Encoding

Considering that the spatial-wise partition convolution is designed in a sharing weights manner against the spatial sparsity. Although this design is efficient, there is still an issue: voxels within an area share the same weight, which results in local details blurred. This phenomenon is further amplified with an increase of kernel sizes. To address this problem, we propose the kernel-wise position embedding. We illustrate this design in Fig. 4.

Particularly, we initialize position weights $e \in \mathbb{R}^{K \times c_{in}}$, which corresponds to the convolution kernel. During the convolution process, we let the input features query the position weights of the equivalent position and add them together. This process modifies Eq. (1) to

$$y_{\bar{p}} = \sum_{k \in K} w_k \cdot (x_{\bar{p}+k} + e_k). \quad (2)$$

This simple design essentially adds bias with relative position information to the input features. There is almost no extra calculation and additional parameters. It is a decent solution to the position insensitivity problem caused by sharing weights, especially for the extremely large kernels, *e.g.*, $17 \times 17 \times 17$. For more experimental results and details, please refer to the Sec. 5.

4.3. Large-kernel Architecture

Following the above designs and observations, we now describe the architectures of our large-kernel 3D CNNs. We readily replace the plain 3D submanifold sparse convolution [25] layers with our SW-LK Conv in existing 3D backbones for semantic segmentation and object detection.

3D Semantic Segmentation The architectures of LargeKernel3D for 3D semantic segmentation resemble MinkowskiNet [9], including 1 stem layer and 8 stages. The stem layer is a $5 \times 5 \times 5$ submanifold convolution. The first 4 stages account for down-sampling and the last 4 stages are for up-sampling. As for the U-Net structure, the first 4 features are concatenated into the up-sampling stages. We replace the original layers in MinkowskiNet-34 [9] by SW-LK Convs, where spatial-wise $7 \times 7 \times 7$ convolutions is grouped into $3 \times 3 \times 3$ splits. By default, other hyper-parameters, including channels and block numbers, follow MinkowskiNet-34 [9].

3D Object Detection The typical backbone networks in 3D object detectors [16, 55, 75] consist of 1 stem layer and 4 stages. A $3 \times 3 \times 3$ submanifold convolution [25] layer serves as the stem. In each stage, except for the first one, there is a sparse convolutional layer with stride 2 for down-sampling. We substitute other plain blocks with SW-LK Convs. For models on nuScenes dataset, we find that $7 \times 7 \times 7$ kernel sizes are enough. They both use the $3 \times 3 \times 3$ spatial-wise groups. For models on Waymo dataset, kernel sizes are scalable to $17 \times 17 \times 17$. We remain the last stage of backbone network as the original, because receptive fields have already been sufficiently enlarged in the front stages. We keep these hyper-parameters unchanged by the default settings in baseline detectors [16, 75]. We include these detailed numbers in *appendix*.

5. Experiments

5.1. Experimental Setting

3D Semantic Segmentation ScanNetv2 [13] is a large-scale benchmark in 3D semantic segmentation. It con-

Table 2. Ablations on various techniques and kernel sizes on MinkowskiNet-34 and ScanNetv2. LN denotes the layer normalization. DW Conv denotes the depth-wise convolution.

Kernel	Baseline	+GELU	+LN	+DW Conv	+PW	Ours
3×3×3	71.7	70.2	66.8	70.6	70.7	-
7×7×7	68.6	68.4	65.0	68.7	68.7	73.5

Table 3. Ablations on group convolution on ScanNetv2.

Kernel	Baseline	Group 2	4	8	16	DW	Ours
3×3×3	71.7	71.5	71.2	70.7	70.6	70.6	-
7×7×7	68.6	68.8	69.2	68.8	68.7	68.7	73.5

Table 4. Comparisons with relevant convolutional schemes on ScanNetv2. Baseline is a MinkowskiNet-34 network. All models listed are trained in the same training hyper-parameters.

Method	Params	FLOPs	Latency	mIoU
Baseline	37.9 M	182.8 G	108 ms	71.7
+ Kernel 5×5×5	170.3M	537.5 G	212 ms	70.7
+ Kernel 7×7×7	465.0M	1089.5 G	487 ms	68.6
Dilate conv	37.9M	100.1 G	98 ms	64.6
Pooling + Dilate	37.9M	183.2 G	115 ms	nan
Spatial group [78]	37.9M	127.2 G	96 ms	70.0
Deformable conv	42.5M	250.1 G	238 ms	72.0
LargeKernel3D-T	38.4M	171.4 G	111 ms	72.8
LargeKernel3D	40.2M	240.0 G	145 ms	73.5

Table 5. Results on other backbone networks and ScanNetv2.

Model	Baseline	+Ours	Δ
MinkowskiNet-14 [9]	67.0	69.7	+2.7
MinkowskiNet-18 [9]	69.5	71.5	+2.0
SparseConvNet [24]	69.3	71.2	+1.9

tains 1,201 indoor RGB-D scenes for training, 312 scenes for validation, and 100 scenes for testing. Semantic labels in 20 categories are annotated. We train our models with the default settings in MinkowskiNet [9], including training hyper-parameters and data augmentations. We evaluate models in the mean Intersection over Union (mIoU) metric. Detailed experimental settings, including network structures and training details, are listed in the *appendix*.

3D Object Detection nuScenes, and Waymo are all popular 3D object detection benchmarks. nuScenes is a large dataset and contains 1,000 driving sequences in total. Among them, there are 700 scenes split for training, 150 scenes for validation, and 150 scenes for testing. It contains LIDAR, camera, and radar sources with a complete 360° environment. We evaluate our methods on both LIDAR-only and LIDAR-

Table 6. Effects of the spatial-wise partition manners on SW-LK Conv and ScanNetv2. Center group means the size of central area. Shifting is to translate the center position by 1 in all 3 dimensions.

Center group	1×1×1	1×1×1	3×3×3	5×5×5
Shifting	✗	✓	✗	✗
mIoU	73.5	66.8	70.6	68.2

Table 7. Improvement over various kernel sizes on LargeKernel3D upon CenterPoint, training on 1/5 Waymo training set and evaluate on the validation set. Δ means the improvement from the kernel 3 baseline. L1 and L2 means LEVEL_1 and LEVEL_2 mAP.

Kernel	Vehicle		Pedestrian		Cyclist	
	L1	L2	L1	L2	L1	L2
Baseline	70.9	62.9	71.5	63.5	69.1	66.5
7×7×7	71.9	63.8	71.7	63.7	70.4	67.8
11×11×11	72.2	64.2	71.8	63.9	70.2	68.3
15×15×15	72.7	64.6	73.8	65.8	70.8	68.2
17×17×17	73.2	65.1	74.1	66.0	71.0	68.1
Δ	2.3↑	2.2↑	2.6↑	2.5↑	1.9↑	1.6↑

Table 8. Effects of position embedding (PE) in the mAP of vehicle detection on Waymo validation set.

Kernel	PE	Latency	LEVEL_1	LEVEL_2
3×3×3	-	109 ms	70.9	62.9
7×7×7	✗	120 ms	71.6	63.4
	✓	122 ms	71.9	63.8
17×17×17	✗	158 ms	71.5	63.3
	✓	164 ms	73.2	65.1

RGB fusion settings. The main evaluation metrics are mAP and nuScenes detection score (NDS). Waymo [57] dataset also contains 1,000 sequences in total, with 798 for training and 202 for validation. For Waymo dataset, we use 1/5 training data for the ablation study and full training set for main results. Results in Waymo are evaluated on difficulty LEVEL_1 and LEVEL_2 objects. In experiments, we validate our networks on state-of-the-art detection framework of CenterPoint [75] on nuScenes [4] and Waymo [57].

5.2. Ablation Studies

Usefulness of 2D CNN Techniques We first validate the techniques that have proven effective on 2D CNNs, including layer normalization [1], depth-wise convolution [29, 43], and GeLU [28]. We conduct experiments on MinkowskiNet-34 [9] on the ScanNet [13] semantic segmentation dataset, as shown in Tab. 2. Kernel 7×7×7 means modifying kernel sizes in all stages to 7. All these

Table 9. Comparison with other methods on nuScenes *test* split. † means flipping and rotation testing-time augmentations.

Method	NDS	mAP	Car	Truck	Bus	Trailer	C.V.	Ped	Mot	Byc	T.C.	Bar
PointPillars [35]	45.3	30.5	68.4	23.0	28.2	23.4	4.1	59.7	27.4	1.1	30.8	38.9
3DSSD [74]	56.4	42.6	81.2	47.2	61.4	30.5	12.6	70.2	36.0	8.6	31.1	47.9
CBGS [80]	63.3	52.8	81.1	48.5	54.9	42.9	10.5	80.1	51.5	22.3	70.9	65.7
CenterPoint [75]	65.5	58.0	84.6	51.0	60.2	53.2	17.5	83.4	53.7	28.7	76.7	70.9
HotSpotNet [7]	66.0	59.3	83.1	50.9	56.4	53.3	23.0	81.3	63.5	36.6	73.0	71.6
CVCNET [6]	66.6	58.2	82.6	49.5	59.4	51.1	16.2	83.0	61.8	38.8	69.7	69.7
UVTR-L [37]	69.7	63.9	86.3	52.2	62.8	59.7	33.7	84.5	68.8	41.1	74.7	74.9
VISTA [17]	69.8	63.0	84.4	55.1	63.7	54.2	25.1	82.8	70.0	45.4	78.5	71.4
Focals Conv [8]	70.0	63.8	86.7	56.3	67.7	59.5	23.8	87.5	64.5	36.3	81.4	74.1
TransFusion-L [2]	70.2	65.5	86.2	56.7	66.3	58.8	28.2	86.1	68.3	44.2	82.0	78.2
LargeKernel3D	70.6	65.4	85.5	53.8	64.4	59.5	29.7	85.9	72.7	46.8	79.9	75.5
LargeKernel3D †	72.8	68.7	86.7	58.5	67.7	62.7	31.9	88.5	77.1	54.9	82.3	76.6
LargeKernel3D-F†	74.2	71.2	87.7	60.1	69.3	66.0	34.3	89.4	81.3	60.2	86.7	77.4

techniques bring no obvious benefit to baseline networks. In contrast, SW-LK Conv introduces clear improvement upon these strong networks. We also show the effects of group convolution with different group numbers in Tab. 3.

Comparisons with Relevant Schemes We compare both performance and efficiency in Tab. 4 using various existing schemes to enlarge receptive fields. Dilated Conv [5] contains $3 \times 3 \times 3$ kernel parameters. It dilates to $7 \times 7 \times 7$ spatial sizes. Pooling + Dilated Conv inserts a 3D average pooling layer before the dilated convolutions to enlarge receptive fields. Spatial group conv is implemented, following the settings in the original paper [78]. We extend Deformable Conv [14] to 3D offset prediction, following [61]. Latencies are measured on an NVIDIA 2080ti GPU. It shows that none of these related schemes present better performance than the baseline. One reason for the low performance of dilation and deformable convolutions is that they tend to factorize the original dense kernels into sparse ones, which cause more information loss when applied to sparse 3D features than their 2D counterparts. Our LargeKernel3D achieves the best performance among these methods. In addition, we also build a tiny version, LargeKernel3D-T, which has half channel numbers in the last two stages. It introduces almost no additional computation overhead to baseline, but still has notable performance improvement. Except for the MinkowskiNet-34, we also show our effects on other backbone networks in Tab. 5.

Spatial-wise Group Partition Manner We ablate the effects of different spatial-wise group partition manners in Tab. 6. We study the size of central groups and center shifting. This ablation study is conducted on the LargeKernel3D and the Scannetv2 dataset. We find that both large central

Table 10. Comparison on ScanNetv2 mIoU on 3D semantic segmentation. † Sliding-window testing.

Method	Latency	val	test
PointCNN [36]	-	-	45.8
PointNet++ [49]	-	53.5	55.7
RandLA-Net [30]	-	-	64.5
PointConv [67]	80 ms	61.0	66.6
PointASNL [69]	-	63.5	66.6
KPConv [61]	-	69.2	68.6
FusionNet [77]	-	-	68.8
Point Transformer† [79]	1012 ms	70.6	-
Fast Point Transformer [46]	298 ms	72.1	-
SparseConvNet [24]	-	69.3	72.5
Stratified Transformer† [34]	1624 ms	74.3	73.7
MinkowskiNet-34	108 ms	71.7	
+ LargeKernel3D	145	73.5	73.9

sizes and center shifting hurt the performance. As the center sizes enlarge from $1 \times 1 \times 1$ to $5 \times 5 \times 5$, the accuracy sharply drops. Because submanifold sparse convolution [25] always aligns each individual voxel feature at its kernel center, we suppose that central features should maintain distinct and clean. Therefore, we keep the center size as $1 \times 1 \times 1$ without shifting, as a default setting.

Further Enlarging Kernel Sizes We show that LargeKernel3D is scalable to kernel sizes $17 \times 17 \times 17$ on the outdoor dataset Waymo [57] in Tab. 7 with baseline kernel $3 \times 3 \times 3$. Models in this ablation are trained on 1/5 training data and evaluated on the full validation set. Results with full training data are compared in Tab. 11. It shows that the mAP on all categories increases by around 2%. This setting demon-

Table 11. Performance of vehicle detection on the Waymo *val* split. All models listed take LIDAR-only input and single frames, without any test-time augmentations or model ensemble.

Method	LEVEL_1	LEVEL_2
	3D AP / APH	3D AP / APH
VoTr-SSD [45]	68.99 / 68.39	60.22 / 59.69
PointPillars [75]	72.08 / 71.53	63.55 / 63.06
SECOND [70]	72.27 / 71.69	63.85 / 63.33
RangeDet [23]	72.85 / 72.33	64.03 / 63.57
SST [22]	74.22 / 73.77	65.47 / 65.07
VoTr-TSD [45]	74.95 / 74.25	65.91 / 65.29
RSN [59]	75.1 / 74.6	66.0 / 65.5
Voxel RCNN [16]	75.59 / —	66.59 / —
LIDAR-RCNN [39]	76.0 / 75.5	68.3 / 67.9
SST [22]	76.22 / 75.79	68.04 / 67.64
Pyramid RCNN [44]	76.30 / 75.68	67.23 / 66.68
Part-A2-Net [56]	77.05 / 76.51	68.47 / 67.14
PV-RCNN [55]	77.51 / 76.89	68.98 / 68.41
SWFormer [58]	69.2 / 68.8	80.9 / 72.7
CenterPoint [75]	76.59 / 76.05	68.85 / 68.35
+ LargeKernel3D	78.07 / 77.61	69.81 / 69.38

strates that a large-scale dataset is a key to unlocking the potential and capability of large kernel methods. In addition, we also try to extend the kernel size to $19 \times 19 \times 19$, but obtain no further gains.

Effects of Position Embedding In LargeKernel3D, position embedding is introduced to relieve the feature-blurring issue caused by weight-sharing. We ablate its effects in Tab. 8, in terms of both latency and performance. Position embedding introduces limited latency overhead on both kernel $7 \times 7 \times 7$ and $17 \times 17 \times 17$ models. In terms of performance, its effects on the $7 \times 7 \times 7$ model are a bit marginal, but essential on the $17 \times 17 \times 17$ model, which has larger kernel sizes and suffers from feature blurring.

Shrinking Kernels During Inference For the spatial-wise partition convolution, we can shrink kernels into small ones during inference. As shown in Tab. 4, the latency of LargeKernel3D is 145ms on ScanNetv2. It degrades to 514ms with the same 73.5% mIoU, if we disable shrinking.

5.3. Main Results

3D Semantic Segmentation We make comparisons with other 3D semantic segmentation methods on the ScanNetv2 dataset in Tab. 10. Our method surpasses others on *test* split. MinkowskiNet [9] is already a state-of-the-art method in ScanNetv2. Our SW-LK Conv further improves its performance. Our method is also superior to transformer-based methods [46, 79]. Stratified Transformer [34] uses a sliding-

window testing technique[†] and splits integrated scenes into overlapped parts. The method tests each part one by one and then ensembles results together, taking hours on evaluation. Moreover, our results are better on the *test* split.

3D Object Detection We compare our LargeKernel3D upon CenterPoint [75] with previous methods on the *test* split of the nuScenes [4] dataset in Tab. 9. LargeKernel3D improves CenterPoint [75] to 70.6% and 72.8% NDS, with and without test augmentation, both outperforming other LIDAR methods. The multi-modal modal LargeKernel3D-F further improves to 74.2% NDS and 71.2% mAP. The fusion module directly follows a simple painting-based manner [8]. Results on the *val* split is shown in the appendix.

We also show the effectiveness of LargeKernel3D model on the vehicle detection of Waymo [57] validation set in Tab. 11. The kernel size of LargeKernel3D is $17 \times 17 \times 17$. LargeKernel3D improves the vehicle performance of CenterPoint [75] by around 1.5% LEVEL_1 AP / APH. Models listed take single-frame LIDAR data without any test-time augmentations or ensemble. It also performs better than other methods. We use CenterPoint [75] to keep consistency to the nuScenes benchmark. Note that LargeKernel3D has the potential to achieve further performance if equipped with stronger detection heads [55].

6. Conclusion and Discussion

We have studied large-kernel designs for 3D convolutional networks, which have essential differences from the solutions in 2D CNNs. We present spatial-wise partition convolution (SW Conv) that is specifically designed for 3D large kernels. It effectively resolves the efficiency and optimization issues in plain 3D large-kernel CNNs. Based on this design, we further propose the SW-LK Conv and the corresponding LargeKernel3D for 3D semantic segmentation and object detection. Our 3D large-kernel networks achieve decent improvement on both semantic segmentation and object detection tasks. For the first time, we show that 3D large kernels can be realized efficiently and effectively. We expect our findings to advance further development of 3D networks. Additional experiments and limitation analysis are provided in the *appendix*.

Limitations LargeKernel3D mainly relies on hand-designed spatial kernel sizes in 3D semantic segmentation and object detection benchmarks. These sizes might be sub-optimal for other datasets or tasks, depending on the overall scene sizes and data sparsity. Other search-based techniques [60] might be helpful, which we will try latter.

Acknowledgement. This work is supported by Shenzhen Science and Technology Program (KQTD20210811090149095), Hong Kong Research Grant Council - Early Career Scheme (No. 27209621), and General Research Fund Scheme (No. 17202422).

References

- [1] Lei Jimmy Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. Layer normalization. *CoRR*, abs/1607.06450, 2016. [4](#), [6](#)
- [2] Xuyang Bai, Zeyu Hu, Xinge Zhu, Qingqiu Huang, Yilun Chen, Hongbo Fu, and Chiew-Lan Tai. Transfusion: Robust lidar-camera fusion for 3d object detection with transformers. *CoRR*, abs/2203.11496, 2022. [7](#)
- [3] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 3286–3295, 2019. [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 multimodal dataset for autonomous driving. In *CVPR*, pages 11618–11628, 2020. [2](#), [6](#), [8](#)
- [5] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *CoRR*, abs/1706.05587, 2017. [2](#), [7](#)
- [6] 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 *NeurIPS*, 2020. [7](#)
- [7] Qi Chen, Lin Sun, Zhixin Wang, Kui Jia, and Alan L. Yuille. Object as hotspots: An anchor-free 3d object detection approach via firing of hotspots. In *ECCV*, volume 12366, pages 68–84, 2020. [7](#)
- [8] Yukang Chen, Yanwei Li, Xiangyu Zhang, Jian Sun, and Jiaya Jia. Focal sparse convolutional networks for 3d object detection. *CoRR*, abs/2204.12463, 2022. [7](#), [8](#)
- [9] Christopher B. Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, pages 3075–3084, 2019. [1](#), [2](#), [4](#), [5](#), [6](#), [8](#)
- [10] Ruihang Chu, Yukang Chen, Tao Kong, Lu Qi, and Lei Li. Icm-3d: Instantiated category modeling for 3d instance segmentation. *IEEE Robotics and Automation Letters*, 7(1):57–64, 2021. [2](#)
- [11] Ruihang Chu, Xiaoqing Ye, Zhengzhe Liu, Xiao Tan, Xiaojuan Qi, Chi-Wing Fu, and Jiaya Jia. Twist: Two-way inter-label self-training for semi-supervised 3d instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1100–1109, 2022. [2](#)
- [12] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021. [3](#)
- [13] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017. [2](#), [4](#), [5](#), [6](#)
- [14] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *ICCV*, pages 764–773, 2017. [2](#), [7](#)
- [15] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255, 2009. [1](#), [4](#)
- [16] Jiajun Deng, Shaoshuai Shi, Peiwei Li, Wengang Zhou, Yanyong Zhang, and Houqiang Li. Voxel R-CNN: towards high performance voxel-based 3d object detection. In *AAAI*, pages 1201–1209, 2021. [2](#), [5](#), [8](#)
- [17] Shengheng Deng, Zhihao Liang, Lin Sun, and Kui Jia. VISTA: boosting 3d object detection via dual cross-view spatial attention. In *CVPR*, pages 8438–8447, 2022. [7](#)
- [18] Xiaohan Ding, Xiangyu Zhang, Yizhuang Zhou, Jungong Han, Guiguang Ding, and Jian Sun. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. In *CVPR*, pages 11963–11975, 2022. [1](#), [2](#), [3](#), [4](#)
- [19] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. *CoRR*, abs/2107.00652, 2021. [3](#)
- [20] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. [1](#)
- [21] Lue Fan, Ziqi Pang, Tianyuan Zhang, Yu-Xiong Wang, Hang Zhao, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Embracing single stride 3d object detector with sparse transformer. *CoRR*, abs/2112.06375, 2021. [1](#), [3](#)
- [22] Lue Fan, Ziqi Pang, Tianyuan Zhang, Yu-Xiong Wang, Hang Zhao, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Embracing single stride 3d object detector with sparse transformer. In *CVPR*, pages 8448–8458, 2022. [8](#)
- [23] Lue Fan, Xuan Xiong, Feng Wang, Naiyan Wang, and Zhaoxiang Zhang. Rangedet: In defense of range view for lidar-based 3d object detection. In *ICCV*, pages 2898–2907, 2021. [8](#)
- [24] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, pages 9224–9232, 2018. [1](#), [6](#), [7](#)
- [25] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, pages 9224–9232, 2018. [1](#), [3](#), [5](#), [7](#)
- [26] Tianrui Guan, Jun Wang, Shiyi Lan, Rohan Chandra, Zuxuan Wu, Larry Davis, and Dinesh Manocha. M3DETR: multi-representation, multi-scale, mutual-relation 3d object detection with transformers. In *WACV*, pages 2293–2303, 2022. [2](#)
- [27] Qi Han, ZeJia Fan, Qi Dai, Lei Sun, Ming-Ming Cheng, Jiaying Liu, and Jingdong Wang. On the connection between local attention and dynamic depth-wise convolution. In *ICLR*, 2021. [1](#), [4](#)
- [28] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *CoRR*, abs/1606.08415, 2016. [4](#), [6](#)

- [29] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017. [1](#), [4](#), [6](#)
- [30] 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 *CVPR*, pages 11105–11114, 2020. [7](#)
- [31] Yunho Jeon and Junmo Kim. Active convolution: Learning the shape of convolution for image classification. In *CVPR*, pages 1846–1854, 2017. [2](#)
- [32] 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 *CVPR*, pages 6423–6432, 2021. [2](#)
- [33] Xin Lai, Yukang Chen, Fanbin Lu, Jianhui Liu, and Jiaya Jia. Spherical transformer for lidar-based 3d recognition. In *CVPR*, 2023. [2](#)
- [34] Xin Lai, Jianhui Liu, Li Jiang, Liwei Wang, Hengshuang Zhao, Shu Liu, Xiaojuan Qi, and Jiaya Jia. Stratified transformer for 3d point cloud segmentation. *CoRR*, abs/2203.14508, 2022. [3](#), [7](#), [8](#)
- [35] Alex H. Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *CVPR*, pages 12697–12705, 2019. [7](#)
- [36] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. In *NeurIPS*, pages 828–838, 2018. [7](#)
- [37] Yanwei Li, Yilun Chen, Xiaojuan Qi, Zeming Li, Jian Sun, and Jiaya Jia. Unifying voxel-based representation with transformer for 3d object detection. *CoRR*, abs/2206.00630, 2022. [7](#)
- [38] Yiming Li, Tao Kong, Ruihang Chu, Yifeng Li, Peng Wang, and Lei Li. Simultaneous semantic and collision learning for 6-dof grasp pose estimation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3571–3578. IEEE, 2021. [2](#)
- [39] Zhichao Li, Feng Wang, and Naiyan Wang. Lidar R-CNN: an efficient and universal 3d object detector. In *CVPR*, pages 7546–7555, 2021. [8](#)
- [40] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In *ECCV*, volume 8693, pages 740–755, 2014. [1](#)
- [41] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, Furu Wei, and Baining Guo. Swin transformer V2: scaling up capacity and resolution. *CoRR*, abs/2111.09883, 2021. [3](#)
- [42] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 9992–10002, 2021. [1](#), [2](#), [3](#)
- [43] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CoRR*, abs/2201.03545, 2022. [1](#), [2](#), [4](#), [6](#)
- [44] Jiageng Mao, Minzhe Niu, Haoyue Bai, Xiaodan Liang, Hang Xu, and Chunjing Xu. Pyramid R-CNN: towards better performance and adaptability for 3d object detection. In *ICCV*, 2021. [8](#)
- [45] Jiageng Mao, Yujing Xue, Minzhe Niu, Haoyue Bai, Jiashi Feng, Xiaodan Liang, Hang Xu, and Chunjing Xu. Voxel transformer for 3d object detection. In *ICCV*, 2021. [1](#), [8](#)
- [46] Chungyun Park, Yoonwoo Jeong, Minsu Cho, and Jaesik Park. Fast point transformer. *CoRR*, abs/2112.04702, 2021. [1](#), [2](#), [7](#), [8](#)
- [47] Chao Peng, Xiangyu Zhang, Gang Yu, Guiming Luo, and Jian Sun. Large kernel matters - improve semantic segmentation by global convolutional network. In *CVPR*, pages 1743–1751, 2017. [2](#)
- [48] Charles Ruizhongtai Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 77–85, 2017. [2](#)
- [49] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, pages 5099–5108, 2017. [2](#), [7](#)
- [50] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67, 2020. [3](#)
- [51] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? In *NeurIPS*, pages 12116–12128, 2021. [1](#)
- [52] Yongming Rao, Wenliang Zhao, Zheng Zhu, Jiwen Lu, and Jie Zhou. Global filter networks for image classification. In *NeurIPS*, pages 980–993, 2021. [2](#)
- [53] David W. Romero, Anna Kuzina, Erik J. Bekkers, Jakub M. Tomczak, and Mark Hoogendoorn. Ckconv: Continuous kernel convolution for sequential data. *CoRR*, abs/2102.02611, 2021. [2](#)
- [54] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. *arXiv preprint arXiv:1803.02155*, 2018. [3](#)
- [55] 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*, pages 10526–10535, 2020. [1](#), [2](#), [5](#), [8](#)
- [56] Shaoshuai Shi, Zhe Wang, Jianping Shi, Xiaogang Wang, and Hongsheng Li. From points to parts: 3d object detection from point cloud with part-aware and part-aggregation network. *T-PAMI*, 43(8):2647–2664, 2021. [8](#)
- [57] 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 *CVPR*, pages 2443–2451, 2020. [2](#), [6](#), [7](#), [8](#)

- [58] Pei Sun, Mingxing Tan, Weiyue Wang, Chenxi Liu, Fei Xia, Zhaoqi Leng, and Dragomir Anguelov. Swformer: Sparse window transformer for 3d object detection in point clouds. *CoRR*, abs/2210.07372, 2022. 8
- [59] Pei Sun, Weiyue Wang, Yuning Chai, Gamaleldin Elsayed, Alex Bewley, Xiao Zhang, Cristian Sminchisescu, and Dragomir Anguelov. RSN: range sparse net for efficient, accurate lidar 3d object detection. In *CVPR*, pages 5725–5734, 2021. 8
- [60] 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 *ECCV*, volume 12373, pages 685–702, 2020. 8
- [61] 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*, pages 6410–6419, 2019. 7
- [62] Asher Trockman and J. Zico Kolter. Patches are all you need? *CoRR*, abs/2201.09792, 2022. 1
- [63] 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. 3
- [64] Jun Wang, Shiyi Lan, Mingfei Gao, and Larry S. Davis. Infofocus: 3d object detection for autonomous driving with dynamic information modeling. In *ECCV*, volume 12355, pages 405–420, 2020. 2
- [65] Tianyu Wang, Xiaowei Hu, Zhengzhe Liu, and Chi-Wing Fu. Sparse2Dense: Learning to densify 3d features for 3d object detection. In *NeurIPS*, 2022. 1
- [66] Kan Wu, Houwen Peng, Minghao Chen, Jianlong Fu, and Hongyang Chao. Rethinking and improving relative position encoding for vision transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10033–10041, 2021. 3
- [67] Wenxuan Wu, Zhongang Qi, and Fuxin Li. Pointconv: Deep convolutional networks on 3d point clouds. In *CVPR*, pages 9621–9630, 2019. 7
- [68] Li Xu, Jimmy S. J. Ren, Ce Liu, and Jiaya Jia. Deep convolutional neural network for image deconvolution. In *NeurIPS*, pages 1790–1798, 2014. 2
- [69] Xu Yan, Chaoda Zheng, Zhen Li, Sheng Wang, and Shuguang Cui. Pointasnl: Robust point clouds processing using nonlocal neural networks with adaptive sampling. In *CVPR*, pages 5588–5597, 2020. 7
- [70] Yan Yan, Yuxing Mao, and Bo Li. SECOND: sparsely embedded convolutional detection. *Sensors*, 18(10):3337, 2018. 8
- [71] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. *CoRR*, abs/2107.00641, 2021. 3
- [72] Tao Yang, Haokui Zhang, Wenze Hu, Changwen Chen, and Xiaoyu Wang. Fast-parc: Position aware global kernel for convnets and vits. *CoRR*, abs/2210.04020, 2022. 3
- [73] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Russ R Salakhutdinov, and Quoc V Le. Xlnet: Generalized autoregressive pretraining for language understanding. *Advances in neural information processing systems*, 32, 2019. 3
- [74] Zetong Yang, Yanan Sun, Shu Liu, and Jiaya Jia. 3dssd: Point-based 3d single stage object detector. In *CVPR*, pages 11037–11045, 2020. 7
- [75] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. In *CVPR*, pages 11784–11793, 2021. 1, 2, 5, 6, 7, 8
- [76] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. In Yoshua Bengio and Yann LeCun, editors, *ICLR*, 2016. 2
- [77] Feihu Zhang, Jin Fang, Benjamin W. Wah, and Philip H. S. Torr. Deep fusionnet for point cloud semantic segmentation. In *ECCV*, volume 12369, pages 644–663, 2020. 7
- [78] Jiahui Zhang, Hao Zhao, Anbang Yao, Yurong Chen, Li Zhang, and Hongen Liao. Efficient semantic scene completion network with spatial group convolution. In *ECCV*, pages 733–749, 2018. 2, 4, 6, 7
- [79] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip H. S. Torr, and Vladlen Koltun. Point transformer. In *ICCV*, pages 16239–16248, 2021. 1, 2, 7, 8
- [80] Benjin Zhu, Zhengkai Jiang, Xiangxin Zhou, Zeming Li, and Gang Yu. Class-balanced grouping and sampling for point cloud 3d object detection. *CoRR*, abs/1908.09492, 2019. 7
- [81] Xizhou Zhu, Han Hu, Stephen Lin, and Jifeng Dai. Deformable convnets V2: more deformable, better results. In *CVPR*, pages 9308–9316, 2019. 2