

# KPConvX: Modernizing Kernel Point Convolution with Kernel Attention

Hugues Thomas<sup>†</sup>, Yao-Hung Hubert Tsai<sup>†</sup>, Timothy D. Barfoot<sup>‡</sup>, Jian Zhang<sup>†</sup>  
<sup>†</sup>Apple, <sup>‡</sup>University of Toronto

<https://github.com/apple/ml-kpconvx>

## Abstract

In the field of deep point cloud understanding, KPConv is a unique architecture that uses kernel points to locate convolutional weights in space, instead of relying on Multi-Layer Perceptron (MLP) encodings. While it initially achieved success, it has since been surpassed by recent MLP networks that employ updated designs and training strategies. Building upon the kernel point principle, we present two novel designs: KPConvD (depthwise KPConv), a lighter design that enables the use of deeper architectures, and KPConvX, an innovative design that scales the depthwise convolutional weights of KPConvD with kernel attention values. Using KPConvX with a modern architecture and training strategy, we are able to outperform current state-of-the-art approaches on the ScanObjectNN, Scannetv2, and S3DIS datasets. We validate our design choices through ablation studies and release our code and models.

## 1. Introduction

The field of 3D point cloud understanding has experienced significant growth in the past decade. This growth can be attributed to the availability of advanced 3D sensors and the increasing use of deep learning in various research domains. The evolution of this field has known different phases. Until 2017, most of the proposed approaches relied on projection in images or 3D grids [3, 17, 27, 34, 35, 44]. However, after 2017, point-based methods gained dominance [2, 13, 19, 20, 31, 32, 41, 47, 50], leading to a rapid expansion of the field. More recently, 3D point cloud understanding followed the trend of attention and transformer networks, which are widely adopted in other deep learning fields [15, 25, 29, 49, 52, 54]. This diverse range of approaches is a response to the challenge posed by the unstructured and continuous nature of point clouds. In this dynamic field, we explore the potential of Kernel Point Convolution (KPConv) [41], one of the most successful point-based methods, when enhanced with modern techniques and attention mechanisms.

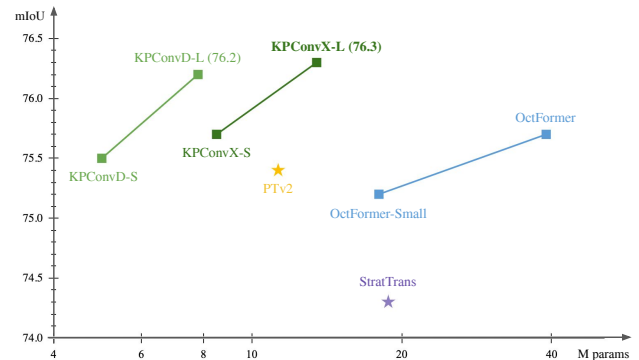


Figure 1. KPConvD and KPConvX using small (S) and large (L) architectures outperform other state-of-the-art architectures on ScanNetv2 dataset using a relatively small number of parameters.

KPConv proposed a convolution design that utilizes kernel points to store network weights in specific positions within the space. In this study, we revisit this design and introduce a new, lighter operator called KPConvD. Additionally, we enhance the operator by incorporating "geometric attention", a novel self-attention mechanism applied to the kernel points, resulting in KPConvX. Apart from the convolution operator itself, we also reconsider the network architectures built using this operator. Due to its lightweight nature, we are able to construct deeper architectures. Overall, we observe that our networks are faster than the original KPConv and achieve state-of-the-art performance.

In Sec. 2, we will provide a comprehensive overview of the different approaches to deep point cloud understanding. Our focus will be on point-based methods, and we will compare KPConvX to other similar designs. Sec. 3 is dedicated to the definition of our convolution design. As depicted in Fig. 2, we utilize kernel points similar to KPConv, but we introduce two major changes: depthwise weights, which combine neighbors' features using the Hadamard product, and a new concept called **kernel attention** that scales the kernel weights based on the current input features. This approach, which defines attention geometrically rather than topologically, is analogous to image involution [18] and is tailored for kernel points, which have a fixed order and

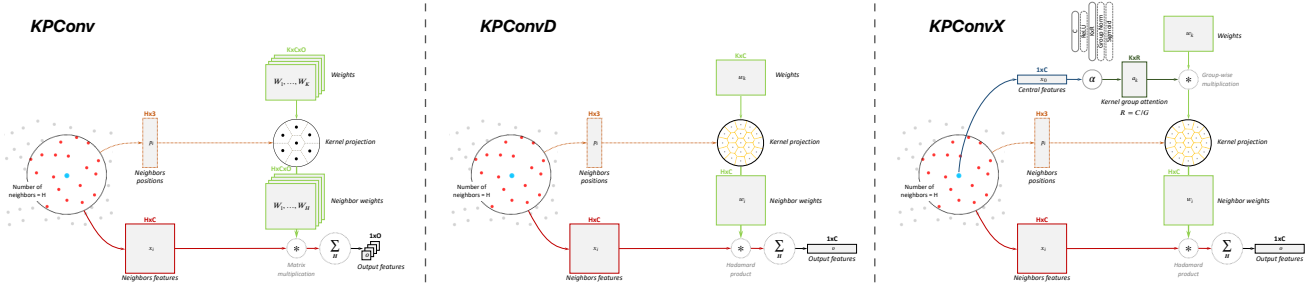


Figure 2. Illustration of our new designs compared to the original KPConv operator. KPConvD adopts a lighter depthwise design and KPConvX includes kernel attention.

spatial disposition, as opposed to neighbor points. Additionally, three other significant modifications are not illustrated in Fig. 2. When projecting kernel weights onto neighbors, we simply use the nearest kernel point instead of summing the influences of all kernel points. This nearest-kernel implementation allows for a larger number of kernel points without significantly increasing computations. Furthermore, we define the positions of kernel points using two shells, as proposed in [21]. Lastly, we also limit the number of neighbors to a fixed value, enabling faster computations.

This lightweight operator allows for much deeper networks. As suggested in other studies that modernize deep networks [15, 23, 33], we increase the depth (number of layers) and width (number of channels) of our networks and use a KPConv stem as the first layer. We also employ partition-based pooling for downsampling layers and incorporate convolution blocks in the decoder upsampling, as described in [49]. Finally, we adopt a scaling factor of  $\sqrt{2}$  (instead of 2) for the number of channels from one layer to the next. The overall network architecture is illustrated in Fig. 3.

The impact of these choices is evaluated in our experiments in Sec. 4. Ablation studies demonstrate that our contributions result in better performance. Our networks outperform the state of the art on object classification and semantic segmentation tasks, ranking first on the ScanObjectNN [42], S3DIS-Area5 [1], and Scannetv2 [7] benchmarks. Fig. 1 highlights the superiority of our approach over other competing methods on Scannetv2. KPConvX-L is our best-performing network, while KPConvD-L offers a great compromise between performance and efficiency.

In short, our contributions are as follows:

- We introduce **KPConvD**, a novel kernel point depthwise convolution operator.
- We introduce **KPConvX**, a version of KPConvD augmented with kernel point attention.
- We design modern **architectures** showing state-of-the-art performances using our operators.

We also aim to provide a comprehensive open-source code, allowing the reproduction of our results and the use of KPConvX as a building block for other networks and tasks.

## 2. Related Work

**Deep point cloud understanding:** As explained in the introduction, point clouds are fundamentally different from images. They are unstructured yet permutation invariant, and their data is continuous but sparse. Due to this complex nature, numerous approaches have been proposed for point cloud understanding.

One of the early ideas for processing point clouds with deep networks was to project points onto a 3D grid [27, 35]. The use of sparse structures like hash-maps or octrees allowed these approaches to scale to larger grids and achieve improved performance [4, 6, 10, 34, 36, 44]. Alternatively, points can be projected onto 2D image planes and processed with image networks [3, 16, 17, 35, 38]. However, quantizing irregular point clouds onto regular grids inevitably results in data loss, and grids lack flexibility.

**Point-wise methods:** To address these problems, it is possible to directly extract features from the point clouds instead of using projections [31]. Typically, features are mixed using linear layers, and geometric patterns are extracted using a local operator. The definition of this local operator is crucial and often the main distinguishing factor between different methods. Many approaches utilize Multi-Layer Perceptrons (MLP) as feature extractors, operating directly on point coordinates [5, 8, 13, 19, 20, 26, 32, 33, 46, 47, 50, 53]. MLP extractors can be viewed as combinations of hyperplanes, and with a sufficient number of parameters (planes), they can represent any geometric pattern. However, complex patterns require a large number of parameters.

**3D transformers:** Transformer-based networks utilize attention as their primary feature extraction mechanism. Whether they employ voxel grouping [15, 29, 43] or local point neighborhoods [25, 30, 48, 49, 52, 54], they always employ MLP to encode point features and geometric patterns. As a result, they face the same limitation as MLP-based convolutions. Additionally, their attention mechanism is based on the combination of neighbors’ keys with central queries, whereas we adopt a simpler central-generated kernel attention similar to [18].

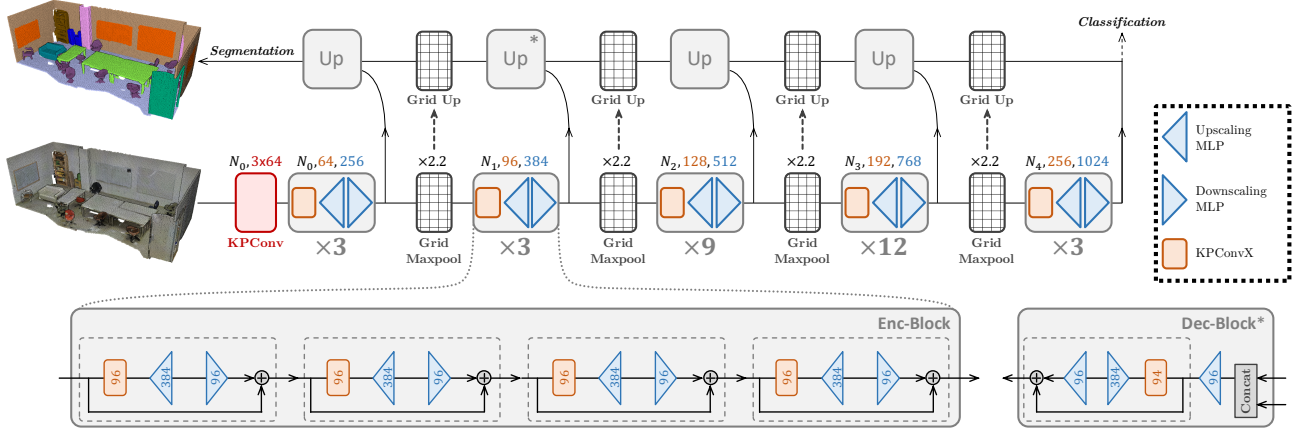


Figure 3. Illustration of our network architecture KPConv-X-L. It can be used for semantic segmentation or shape classification. It has a total of 44 encoder blocks, plus one stem KPConv and 4 decoder blocks. We use inverted bottleneck blocks and grid subsampling from one layer to the next.

**Geometric convolutions:** Eventually, local feature extractors can be defined geometrically [2, 13, 21, 41]. This reduces the complexity of the convolution operator without losing data through quantization. We use this approach and incorporate an attention mechanism to the convolution weights, similar to the modulations used in the deformable version of KPConv [41]. However, instead of a single modulation per kernel point, we use group modulations as described in [18].

### 3. Our Designs

This section describes our novel convolution designs called KPConvD and KPConvX, illustrated in Fig. 2. We first discuss the modifications made to the convolution itself, followed by an explanation of the kernel attention mechanism that we added to the operator. Finally, we provide an overview of our new network architectures (Fig. 3) and training strategies.

#### 3.1. KPConvD: Depthwise Kernel Point Convolution

In essence, our KPConvD operator can be understood as a depthwise version of KPConv. However, aside from some hyperparameter adjustments, the core implementation has been optimized for efficiency. We provide a detailed description of the modifications, using notations similar to those used in the original KPConv paper, where  $K$  represents the number of kernel points and  $H$  represents the number of neighbors:

- input points:  $x_i \in \mathbb{R}^3$  with  $i < H$
- input features:  $f_i \in \mathbb{R}^C$  with  $i < H$
- kernel points:  $\tilde{x}_k \in \mathbb{R}^3$  with  $k < K$
- kernel weights:  $W_k \in \mathbb{R}^{C \times O}$  with  $k < K$

Therefore, a standard KPConv can be written as

$$(\mathcal{F} * g)(x) = \sum_{x_i \in \mathcal{N}_x} \sum_{k < K} h(x_i - x, \tilde{x}_k) W_k f_i, \quad (1)$$

where  $\mathcal{N}_x$  is a radius neighborhood with  $r \in \mathbb{R}$ , and  $h$  is the influence (or correlation) function. Note that in fact, the implementation of KPConv uses a radius neighborhood truncated by a maximum number of neighbors for efficiency. It is therefore equivalent to a KNN neighborhood, where all points further than  $r$  are ignored, which is how we implement it in our work. We thus have a fixed number of neighbors, named  $H$ . We use the same influence function as KPConv:

$$h_{ik} = \max\left(0, 1 - \frac{\|x_i - x - \tilde{x}_k\|}{\sigma}\right), \quad (2)$$

where  $\sigma$  is the influence distance of the kernel points. We use the notation  $h_{ik}$  in place of  $h(x_i - x, \tilde{x}_k)$  for convenience, but it also represents an optimization in our implementation. All the blocks of the same layer can share the same kernel points and therefore share the same influences. We thus compute the matrix of influences  $(h_{ik}) \in \mathbb{R}^{H \times K}$  only once per layer and share the values  $h_{ik}$  for the rest of the layer

With these new notations, the depthwise equivalent of KPConv can be defined as

$$(\mathcal{F} * g)(x) = \sum_{i < H} \sum_{k < K} h_{ik} w_k \odot f_i, \quad (3)$$

where  $w_k \in \mathbb{R}^C$  is a single vector of weights instead of a full matrix and  $\odot$  is the Hadamard product. We observe that the influence matrices usually contain a majority of zeros. Indeed, the kernel points only apply their weights to the neighbors that are in their area of influence. However, in this definition, the influence of all kernel points is summed




Figure 4. Illustration of kernel dispositions in 2D with one or two shells. Red circles are kernel points. The shells highlighted in green are placed regularly along the radius (left). The nearest-neighbor area of each kernel point is shown in yellow (right).

for each neighbor, which is a waste of computations. We reduce the number of operations by choosing a single kernel point for each neighbor instead:

$$(\mathcal{F} * g)(x) = \sum_{i < H} h_{ik^*} w_{k^*} \odot f_i, \quad (4)$$

where  $k^*$  is the index of the nearest kernel point to the neighbor  $i$ . With this nearest-kernel implementation, we are able to increase the number of kernel points and thus the descriptive power of the convolution without affecting the number of operations.

Because of this ability to increase the number of kernel points, we adopt a shell definition of the kernel point positions as proposed in [21], but keep the convolution definition defined in Eq. (4). The kernel point locations are found with a similar optimization process to the one used in KPConv, but with an equality constraint added for the radius of the points of each shell. More details can be found in the supplementary material. We note the number of kernel points as a list  $[1, N_1, \dots, N_s]$ , where  $s$  is the number of shells. In the rest of the paper, we use a two-shell kernel disposition [1, 14, 28]. Examples of 2D kernel points' disposition and their corresponding nearest-kernel areas are shown in Fig. 4.

Eventually, we choose a radius of  $r = 2.1$  times the sub-sampling grid size, slightly smaller than the original KPConv. This value was found empirically (see additional experiments in the supplementary material). The chosen radius does not affect memory consumption as the number of neighbors is fixed. Because of our nearest-kernel implementation, the area of influence of each kernel point does not overlap, and we use a large influence radius  $\sigma = r$ . We note that this nearest-kernel design was proposed in the thesis [40], but never implemented.

### 3.2. KPConvX: Adding Kernel Attention

Local self-attention, as it is commonly used in transformers for point clouds [49, 54], is defined with respect to the neighbor features. It is a topological operator, which combines features without extracting geometric information. This is why most of the self-attention designs for point clouds incorporate some kind of geometric encoding.

On the contrary, the kernel attention we propose is geometric in nature. The attention weights are generated for local parts of the space, instead of being generated for neighbors depending on their features. This design gives a geometric structure to our attention mechanism, and the ability to focus on geometric patterns, as shown in Fig. 5. Therefore, our operator does not need additional position encodings to capture geometric information.

With the previous notations, we define KPConvX as

$$(\mathcal{F} * g)(x) = \sum_{i < H} h_{ik^*} (m_{k^*} \odot w_{k^*}) \odot f_i, \quad (5)$$

where  $m_{k^*}$  is a vector of modulations generated for the  $k^{\text{th}}$  kernel point and  $\odot$  is a grouped version of the Hadamard product. Indeed,  $m_{k^*} \in \mathbb{R}^{C_g}$  where  $C_g = C/G$  and  $G$  is the number of groups. We note that with  $G = C$ , our design is very similar to the modulations of deformable KPConv [41], and with  $G = 1$ , every channel gets its own modulation. In the following, we use  $G = 8$  unless otherwise specified. The modulations are generated all together from the feature of the central point:

$$(m_k)_{k < K} = \alpha(x), \quad (6)$$

where  $\alpha$  is defined as a two-layer MLP, with  $C$  hidden channels,  $K \times C_g$  output channels, and a sigmoid as final activation. We note that with this attention definition, augmenting the number of kernel points will affect the number of operations. However, thanks to our nearest-kernel implementation, the increase in operations is still reduced compared to the full-summation kernel.

As stated above, from this general definition we can regress to a depthwise convolution, by removing the mod-

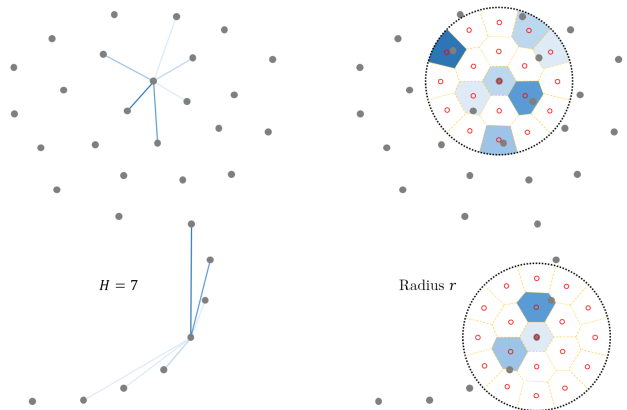


Figure 5. Illustration of our kernel attention principle, where chunks of space are weighted (right) compared to standard point self-attention where the neighbors are weighted depending on their feature instead of their position (left). In our design, no position encoding is needed, as the attention itself is a position encoding.

ulations, as defined in Eq. (4). But it is also possible to remove the convolution weights and only use modulations. Similarly to the image involution [18], we would then define a kernel point involution:

$$(\mathcal{F} \circledast g)(x) = \sum_{i < H} h_{ik^*} (m_{k^*} \odot f_i) . \quad (7)$$

We named this design KPIInv, but it did not match the performances of KPConvX, and we do not use it in our experiments.

KPConvX design would not work without kernel points or a similar structure. Why? The modulations are generated with  $\alpha$  in a certain order, which corresponds to the order of the kernel points. If we were to produce modulations directly for neighbors, as a naive re-implementation of the image involution for point clouds would do, they would be applied in this order to the neighbors. However, even if they are ordered by coordinates or distance to the center, the neighbors’ order remains highly unstable, and therefore the modulations would be applied randomly to different neighbors every time. This instability undermines the network’s ability to extract geometric patterns and we observed weaker performances in the few tests we ran.

Furthermore, predicting weights from the central point forces the network to include contextual information in its features. This contextual information is necessary for deciding where to focus attention spatially in ensuing layers. We believe this is one of the crucial properties that make KPConvX a better descriptor.

### 3.3. Modern Architecture and Training

In this work, we also focused on the design of new network architectures. We follow the path of other works that advocate the use of modern techniques for deep networks in images [23] or point clouds [33]. Firstly, we define a small (S) and large (L) architectures, with respectively [2, 2, 2, 8, 2] blocks per layer and [3, 3, 9, 12, 3] blocks per layer. We choose an initial width of 64 channels for both networks as shown in Fig. 3.

For the stem of the network, we choose a full standard KPConv, to extract strong initial features as advocated in [15]. The cost of this layer is negligible compared to the rest of the network.

Each network block is designed as an inverted bottleneck similarly to [23, 33]. Compared to the standard ResNet bottleneck design [11] that starts with a downsampling MLP, follows with a local extractor, and finishes with an upsampling MLP, the inverted design places the local extractor at the beginning of the block and has two MLPs that upsample and then downsample the features. We use a normalization layer (Batch Normalization) and an activation layer (Leaky ReLU), after KPConvX and after each MLP. For the final

MLP of the block, the activation layer is moved after the addition.

Each layer has its own fixed number of neighbors for the local operator. In all our experiments, we use [12, 16, 20, 20, 20] neighbors for each layer. The number of neighbors is lower in the early layers because the point clouds are more sparse before being subsampled.

Instead of using heavy operations for the pooling layers, we follow [49] and implement partition-based pooling with a grid of increasing size. In the same fashion, we use a ratio of 2.2 for the grid scaling between layers, which reduces the overall memory cost of the network. The grid upsampling operation projects the features back to the points that were in the same grid cell, without any modification. We also append an additional block in each decoder layer to improve the performance.

Furthermore, a common network regularization scheme, stochastic depth [14], is implemented with DropPath operations. It consists of randomly skipping network blocks for entire elements of a batch. The implementation of DropPath for our network is not trivial because we use variable batch size as in the original KPConv [41]. We refer to the supplementary material and our open-source implementation for more details.

Interestingly, we notice that the number of channels in the deeper layers does not need to be doubled to maintain good performances. Instead of keeping a ratio of 2 between the widths of two consecutive layers, we use a ratio of  $\sqrt{2}$ . This reduces the size and memory consumption of our network, without affecting its performance.

**Segmentation Architecture:** After the last upsampling layer, we use a standard segmentation head, which consists of a two-layer MLP with 64 hidden channels, and  $n_{\text{class}}$  output channels, followed by a softmax layer. At training, we use a standard cross-entropy loss.

**Classification Architecture:** The classification architecture only uses the encoder part of the network. The features are aggregated with a global average pooling and processed by the classification head, which consists of a two-layer MLP with 256 hidden channels, and  $n_{\text{class}}$  output channels, followed by a softmax layer. We use a cross-entropy loss with label smoothing [37] for training.

Finally, we train our networks with the more recent AdamW optimizer [24] and we use up-to-date data augmentation strategies [33]. More details and other training parameters can be found in the supplementary material.

## 4. Experiments

To assess the effectiveness of the proposed method, we perform experimental evaluations on three datasets: Stanford Large-Scale 3D Indoor Spaces (S3DIS) [1], ScanNetv2 [7], and ScanObjectNN [42]. We begin by presenting our setup and comparing our results to the state of the art on these

Table 1. Semantic segmentation results on S3DIS Area 5 evaluated under mIoU (%), mAcc (%), and OA (%) metrics.

Method	mIoU	mAcc	OA
KPConv [41]	67.1	72.8	-
PTv1 [54]	70.4	76.5	90.8
SPoTr [30]	70.8	76.4	90.7
PointNeXt [33]	70.5	76.8	90.6
PointMixer [5]	71.4	77.4	-
PTv2 [49]	71.6	77.9	91.1
StratTrans [15]	72.0	78.1	91.5
PointVector [8]	72.3	78.1	91.0
PointMetaBase [22]	72.3	-	91.3
KPConvX-L (ours)	<b>73.5</b>	<b>78.7</b>	<b>91.7</b>

datasets. Next, we conduct several ablation studies to confirm the effectiveness of our contributions.

### 4.1. Semantic Segmentation

**Data and metrics.** S3DIS [1] is a challenging benchmark that consists of 6 large-scale indoor areas, with a total of 271 rooms spread across three different buildings. The points are densely sampled on the mesh surfaces and annotated with 13 semantic categories. Scannetv2 [7], on the other hand, is relatively larger and includes 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 adopt the standard experimental setup [32, 39, 41, 54], using the fifth area as our test set for S3DIS and following the official train/evaluation split for ScanNetv2. The standard metric for these datasets is the mean class-wise intersection-over-union (mIoU). Additionally, we may also present the mean of class-wise accuracy (mAcc) and the overall point-wise accuracy (OA) for S3DIS.

**Experimental setup.** For both datasets, during training, we randomly select a point from a random room as the center

Table 2. Semantic segmentation results on ScanNetv2 validation set evaluated under mIoU (%) and number of parameters.

Method	params	mIoU
KPConv [41]	15.0M	69.2
PTv1 [54]	7.8M	70.6
FastPtTrans [29]	37.9M	72.1
MinkUNet [6]	37.9M	72.2
LargeKernel [4]	40.2M	73.2
BPNet [12]	-	73.9
StratTrans [15]	18.8M	74.3
EQ-Net [51]	-	75.3
PTv2 [49]	11.3M	75.4
OctFormer [43]	39.0M	75.7
KPConvX (ours)	13.5M	<b>76.3</b>

of a sphere. We then sample all the points within a radius of 2.5m to create an input point cloud. Our 5-channel input feature includes a constant one feature, RGB values, and the  $z$ -coordinate of the points (before augmentation) for S3DIS. For ScanNetv2, we add 3 values for the normals. Following the approach in [41], we stack these variable-length point clouds to create stacked batches. We use a target batch size of 4 and accumulate 6 forward passes before propagating the gradient backward. This effectively gives us a batch size of 24, while keeping the memory consumption of a batch size of 4.

We subsample the input point clouds using a voxel size of 0.04m for S3DIS and 0.02m for ScanNetv2. Similar to recent methods, we do not use spheres on the test set. Instead, we evaluate our networks on the entire rooms. More details on the training parameters and augmentations can be found in the supplementary material.

In contrast to [33], we advocate for the use of voting when evaluating network predictions. We argue that if all the votes come from the same network with fixed weights, it cannot be considered an ensemble method. Moreover,




Figure 6. Visualization of Semantic segmentation results on in rooms of S3DIS Area 5.

when a network is trained with data augmentation, its results may vary depending on how it encounters the test data. For instance, as shown in [15], the results can unexpectedly be high when using a certain room orientation. By utilizing voting, we reduce the variance in the results and provide smoother predictions.

**Results.** We compare our best network performance to state-of-the-art results using different versions of the dataset. The results are summarized in Tab. 2. Our KPConvX-L model achieves the highest mIoU score in both cases, outperforming the current state of the art by +1.2 mIoU on S3DIS and by +0.6 mIoU on Scannet. We also provide qualitative results of point cloud semantic segmentation in Fig. 6.

## 4.2. Shape Classification

**Data and metrics.** ScanObjectNN is a recent classification benchmark that includes approximately 15,000 real scanned objects and 2,902 unique object instances [42]. Due to occlusions and noise, it is considered more challenging than other 3D shape classification benchmarks. Following the standard protocol [33, 52], we utilize the hardest variant of the dataset, PB\_T50\_RS. The performances are evaluated without pre-training using the mean of class-wise accuracy (mAcc) and overall point-wise accuracy (OA).

**Experimental setup.** We adopt the same point resampling strategy as [33, 52]: during training, we randomly sample 1,024 points for each shape, while during testing, points are uniformly sampled. Although the input points are not sub-sampled with a grid, we still need to choose a fake initial grid size (0.02cm) to scale our convolution radius and define the next layer grids. Our batch size is set to 32 point clouds, with an accumulation of 2 forward passes, resulting in an effective batch size of 64. Similarly, we employ voting during testing due to the aforementioned reasons.

**Results.** Tab. 3 displays our network performances compared to other competing methods. We present both the average and best scores from ten runs. Our top-performing network, KPConvX-L, outperforms PointVector [8] by a significant margin (1.1% mAcc and OA on average). It is important to note that we did not compare against methods that employ a pre-training strategy.

## 4.3. Ablation: From KPConv to KPConvX

In this experiment, we applied a complete series of changes to transform KPConv into KPConvX. We conducted the experiment on S3DIS dataset under controlled settings, using the same training parameters and data augmentation as in our other experiments. This is why our version of KPConv achieved a slightly better score than the original paper. In the previous experiments, we provided the results of our best model and plan to share these best weights. However, in all the following ablation studies, for each version,

Table 3. Shape classification on ScanObjectNN **without** pre-training. We report our average and best scores after retraining the networks ten times.

Method	OA		mAcc	
	mean±std	best	mean±std	best
PointNet [31]	-	68.2	-	63.4
PointNet++ [32]	-	77.9	-	75.4
DGCNN [45]	-	78.1	-	73.6
PointCNN [20]	-	78.5	-	75.1
Simpleview [9]	-	80.5	-	-
PointBERT [52]	-	83.1	-	-
Point-MAE [28]	-	85.2	-	-
PointMLP [26]	85.4 ± 0.3	-	83.9 ± 0.5	-
PointStack [46]	86.9 ± 0.3	87.2	85.8 ± 0.3	86.2
PointNeXt [33]	87.7 ± 0.4	88.2	85.8 ± 0.6	86.6
PointVector [8]	87.8 ± 0.4	88.6	86.2 ± 0.5	86.8
SPoTr [30]	-	88.6	-	86.8
KPNeXt-S (ours)	88.3 ± 0.4	<u>89.0</u>	86.7 ± 0.5	87.4
KPNeXt-L (ours)	<b>88.9 ± 0.3</b>	<b>89.3</b>	<b>87.3 ± 0.5</b>	<b>88.1</b>

we trained and tested 10 identical models and got an average score. This is more reliable to provide fair comparison studies. Tab. 4 shows the average and standard deviation of the 10 test mIoUs on S3DIS Area5, in addition to the best mIoU obtained. For each version, we also provide the number of million parameters, the network throughput (TP) in instances per second (ins/s), and the GPU memory consumption in GB during inference. This allows a better understanding of the impact of each change. For comparison purposes, the throughput and GPU consumption are measured using the same parameters as [33]: batches of 16 input point clouds, each containing 15,000 points, were fed to the network. Each model was trained and tested on a single Nvidia Tesla V100 32GB GPU.

First, we transform KPConv into a depthwise operation, reduce the radius, reduce the number of neighbors to fixed values, and add one kernel shell ([1, 14] to [1, 14, 28]). This results in a +0.5 mIoU improvement.

Next, we add more layers from KP-FCNN [2, 2, 2, 2, 2] to our large architecture [3, 3, 9, 12, 3]. We also reduce the

Table 4. Ablation study from KPConv to KPConvX. Best mIoU is highlighted in **bold** and mIoUs within 1% of the best one are underlined.

Method	mIoU (10 tries)	TP	GPU	params
	mean±std	ins/s	GB	M
KPConv (our impl)	68.5 ± 0.2	35.4	12.7	14.1
+ depthwise	68.0 ± 0.3	38.3	7.1	7.3
↘ radius & neighbors	68.2 ± 0.3	51.4	4.9	7.3
+ shell (1 14 28)	69.0 ± 0.6	41.8	11.2	7.4
+ new architecture	70.9 ± 0.7	36.6	10.1	10.1
+ inv-blocks	<u>72.2 ± 0.4</u>	41.3	8.1	7.8
+ nearest influence	<u>72.3 ± 0.4</u>	46.5	5.2	7.8
+ shareKP (KPConvD-L)	<u>72.2 ± 0.7</u>	64.1	4.6	7.8
+ attention (KPConvX-L)	<b>72.4 ± 0.9</b>	47.7	6.8	13.5

Table 5. Ablation study of the architectures. Best mIoU is highlighted in **bold** and mIoUs within 1% of the best one are underlined.

Architecture	mIoU (10-try avg)	TP	GPU	params
	mean±std	ins/s	GB	M
KPConvX-L	<b>72.4</b> ± 0.9	47.7	6.8	13.5
KPConvX-S	71.1 ± 0.5	62.3	6.8	8.5
KPConvD-L	<u>72.2</u> ± 0.7	64.1	4.6	7.8
KPConvD-S	70.2 ± 0.8	75.7	4.6	5.0

channel scaling to  $\sqrt{2}$  and add decoder layers and Drop-Path. With these changes, the mIoU improves by 1.9%, while maintaining similar GPU memory consumption and throughput. To finish the architecture changes, we use inverted bottleneck blocks for an additional 1.9% mIoU gain.

We then demonstrate the effectiveness of our nearest influence design, which is 13% faster (TP↗) and 56% lighter (GPU↘). Additionally, our kernel-point-sharing strategy makes the model 38% faster (TP↗) and 13% lighter (GPU↘) without impacting the results.

At this point, we have our **KPConvD-L** architecture, which improves performances of the original KPConv by **+3.7 mIoU** on average, runs **80% faster**, and uses only **20% GPU memory** compared to the original KPConv.

Finally, we add our attention mechanism to create the **KPConvX-L** architecture. Compared to the original KPConv, it improves performances by **+3.9 mIoU** on average, runs **35% faster**, and uses only **half GPU memory**.

#### 4.4. Other Ablation Studies

In the previous ablation study, we demonstrated the improvements brought by our contributions. In this section, we provide more insight into our method by studying the effect of some parameter changes.

First, we compare our different architectures on S3DIS Area 5. From Tab. 5, we observe that our best network is KPConvX-L. However, it also has the highest computational cost. Using KPConvD-L is a good option for more efficiency, this model being 35% faster, while only sacrificing 0.2% mIoU on average. Overall, our throughput is highly competitive and comparable to PointNeXt-XL [33], which presented a throughput of 46 ins/s, similar to KPConvX-L in the same setup, even though their performances are only 70.5 mIoU on average on S3DIS Area 5.

Next, we assess the impact of the number of groups in KPConvX on the performance of KPConvX-L. As shown in Tab. 6, when there are fewer groups, the network has a greater number of parameters to learn. Conversely, with more groups, the modulations become less effective. We observe that the network performs really well with 8 groups, which is the value we use in the rest of the paper. This value provides a balance between enhancing descriptive power and avoiding an excessive parameter burden.

Table 6. Ablation study of the KPConvX groups. Best mIoU is highlighted in **bold** and mIoUs within 1% of the best one are underlined.

Number of groups	mIoU (10-try avg)	TP	GPU	params
	mean±std	ins/s	GB	M
$G = 1$	<u>72.5</u> ± 0.3	32.1	9.3	47.0
$G = 4$	<b>72.5</b> ± 0.4	48.2	7.2	18.3
$G = 8^*$	<u>72.4</u> ± 0.6	47.7	6.8	13.5
$G = 16$	72.1 ± 0.5	56.8	6.7	11.1
$G = C$	<u>72.2</u> ± 0.3	70.2	5.4	8.9

Our goal is to help the reader gain a deeper understanding of the mechanisms that contribute to the performance and efficiency of our models. We thus provide more ablation and parameter studies in the supplementary material, including the convolution radius and the number of kernel points and shells.

## 5. Conclusion

We present KPConvX, an efficient feature extractor for point clouds that combines depthwise convolution and self-attention. Additionally, we introduce KPConvX-L, a new deep architecture for semantic segmentation and shape classification. KPConvX-L is trained using the latest strategies and achieves state-of-the-art performance on several benchmarks for 3D semantic segmentation and 3D shape classification.

**Limitations and future work.** Deep learning architectures utilize local feature extractors to generate new information from local neighborhoods. However, the topological or geometric nature of these local feature extractors is often understudied. Transformers and some 3D graph architectures have a topological nature, as they are based on neighbor-feature relationships. Without incorporating MLP geometric encodings, they would overlook the geometry of the point clouds. On the other hand, structured convolutions like KPConv or voxel networks are inherently geometric encodings, merging features solely based on their locations, without considering neighbor features. With KPConvX, we introduced attention into a geometric operator, in contrast to recent point transformers that incorporated geometric encodings into a topological operator. Nevertheless, in both cases, the feature extraction operators still generate new information, either topologically or geometrically. It is important to conduct thorough studies to understand how the topological or geometric nature of local feature extractors impacts the learning process, and whether these two ideas can be combined in a single architecture to separate topological and geometric features.



## 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. <http://buildingparser.stanford.edu/dataset.html>. 2, 5, 6
- [2] Matan Atzmon, Haggai Maron, and Yaron Lipman. Point convolutional neural networks by extension operators. *ACM Transactions on Graphics (TOG)*, 37(4):71, 2018. 1, 3
- [3] Alexandre Boulch, Bertrand Le Saux, and Nicolas Audebert. Unstructured point cloud semantic labeling using deep segmentation networks. In *Proceedings of the Workshop on 3D Object Retrieval (3DOR)*, 2017. 1, 2
- [4] Yukang Chen, Jianhui Liu, Xiangyu Zhang, Xiaojuan Qi, and Jiaya Jia. Largekernel3d: Scaling up kernels in 3d sparse cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13488–13498, 2023. 2, 6
- [5] Jaesung Choe, Chunghyun Park, Francois Rameau, Jaesik Park, and In So Kweon. Pointmixer: Mlp-mixer for point cloud understanding. In *European Conference on Computer Vision*, pages 620–640. Springer, 2022. 2, 6
- [6] 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. 2, 6
- [7] 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. [http://kaldir.vc.in.tum.de/scannet\\_benchmark](http://kaldir.vc.in.tum.de/scannet_benchmark). 2, 5, 6
- [8] Xin Deng, WenYu Zhang, Qing Ding, and XinMing Zhang. Pointvector: A vector representation in point cloud analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9455–9465, 2023. 2, 6, 7
- [9] Ankit Goyal, Hei Law, Bowei Liu, Alejandro Newell, and Jia Deng. Revisiting point cloud shape classification with a simple and effective baseline. In *International Conference on Machine Learning*, pages 3809–3820. PMLR, 2021. 7
- [10] 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. 2
- [11] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 5
- [12] Wenbo Hu, Hengshuang Zhao, Li Jiang, Jiaya Jia, and Tien-Tsin Wong. Bidirectional projection network for cross dimension scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14373–14382, 2021. 6
- [13] Binh-Son Hua, Minh-Khoi Tran, and Sai-Kit Yeung. Pointwise convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 984–993, 2018. 1, 2, 3
- [14] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer, 2016. 5
- [15] 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, 5, 6, 7
- [16] 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. 2
- [17] Felix Järemo Lawin, Martin Danelljan, Patrik Tosteberg, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Deep projective 3d semantic segmentation. In *International Conference on Computer Analysis of Images and Patterns*, pages 95–107. Springer, 2017. 1, 2
- [18] Duo Li, Jie Hu, Changhu Wang, Xiangtai Li, Qi She, Lei Zhu, Tong Zhang, and Qifeng Chen. Involution: Inverting the inherence of convolution for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12321–12330, 2021. 1, 2, 3, 5
- [19] Jiaxin Li, Ben M. Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9397–9406, 2018. 1, 2
- [20] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. In *Advances in Neural Information Processing Systems*, pages 820–830, 2018. 1, 2, 7
- [21] Yuyan Li, Chuanmao Fan, Xu Wang, and Ye Duan. Spnet: Multi-shell kernel convolution for point cloud semantic segmentation. In *International Symposium on Visual Computing*, pages 366–378. Springer, 2021. 2, 3, 4
- [22] Haojia Lin, Xiawu Zheng, Lijiang Li, Fei Chao, Shanshan Wang, Yan Wang, Yonghong Tian, and Rongrong Ji. Meta architecture for point cloud analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17682–17691, 2023. 6
- [23] 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. 2, 5
- [24] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 5
- [25] Dening Lu, Kyle Gao, Qian Xie, Linlin Xu, and Jonathan Li. 3DPCT: 3D Point Cloud Transformer with Dual Self-attention. *arXiv preprint arXiv:2209.11255*, 2022. 1, 2
- [26] Xu Ma, Can Qin, Haoxuan You, Haoxi Ran, and Yun Fu. Rethinking network design and local geometry in point

- cloud: A simple residual mlp framework. *arXiv preprint arXiv:2202.07123*, 2022. 2, 7
- [27] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *Intelligent Robots and Systems (IROS), 2015 IEEE/RSJ International Conference on*, pages 922–928. IEEE, 2015. 1, 2
- [28] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. *arXiv preprint arXiv:2203.06604*, 2022. 7
- [29] 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, 2, 6
- [30] Jinyoung Park, Sanghyeok Lee, Sihyeon Kim, Yunyang Xiong, and Hyunwoo J Kim. Self-positioning point-based transformer for point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21814–21823, 2023. 2, 6, 7
- [31] 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
- [32] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, pages 5099–5108, 2017. 1, 2, 6, 7
- [33] 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. 2, 5, 6, 7, 8
- [34] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017. 1, 2
- [35] 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. 1, 2
- [36] Hang Su, Varun Jampani, Deqing Sun, Subhransu Maji, Evangelos Kalogerakis, Ming-Hsuan Yang, and Jan Kautz. Splatnet: Sparse lattice networks for point cloud processing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2530–2539, 2018. 2
- [37] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2818–2826, 2016. 5
- [38] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent convolutions for dense prediction in 3d. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3887–3896, 2018. 2
- [39] Lyne Tchapmi, Christopher Choy, Iro Armeni, JunYoung Gwak, and Silvio Savarese. Segcloud: Semantic segmentation of 3d point clouds. In *2017 International Conference on 3D Vision (3DV)*, pages 537–547. IEEE, 2017. 6
- [40] Hugues Thomas. *Learning new representations for 3D point cloud semantic segmentation*. PhD thesis, Université Paris sciences et lettres, 2019. 4
- [41] 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 International Conference on Computer Vision*, pages 6411–6420, 2019. 1, 3, 4, 5, 6
- [42] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1588–1597, 2019. 2, 5, 7
- [43] Peng-Shuai Wang. Octformer: Octree-based transformers for 3D point clouds. *ACM Transactions on Graphics (SIG-GRAPH)*, 42(4), 2023. 2, 6
- [44] Peng-Shuai Wang, Yang Liu, Yu-Xiao Guo, Chun-Yu Sun, and Xin Tong. O-cnn: Octree-based convolutional neural networks for 3d shape analysis. *ACM Transactions on Graphics (TOG)*, 36(4):72, 2017. 1, 2
- [45] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):1–12, 2019. 7
- [46] Kevin Tirta Wijaya, Dong-Hee Paek, and Seung-Hyun Kong. Advanced Feature Learning on Point Clouds using Multi-resolution Features and Learnable Pooling. *arXiv preprint arXiv:2205.09962*, 2022. 2, 7
- [47] 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. 1, 2
- [48] Wenxuan Wu, Qi Shan, and Li Fuxin. Pointconvformer: Revenge of the point-based convolution. *arXiv preprint arXiv:2208.02879*, 2022. 2
- [49] Xiaoyang Wu, Yixing Lao, Li Jiang, Xihui Liu, and Hengshuang Zhao. Point transformer v2: Grouped vector attention and partition-based pooling. In *NeurIPS*, 2022. 1, 2, 4, 5, 6
- [50] 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. 1, 2
- [51] Zetong Yang, Li Jiang, Yanan Sun, Bernt Schiele, and Ji-aya Jia. A unified query-based paradigm for point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8541–8551, 2022. 6
- [52] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud

transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19313–19322, 2022. [1](#), [2](#), [7](#)

- [53] Hengshuang Zhao, Li Jiang, Chi-Wing Fu, and Jiaya Jia. Pointweb: Enhancing local neighborhood features for point cloud processing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5565–5573, 2019. [2](#)
- [54] 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](#), [4](#), [6](#)