

PointMixer: MLP-Mixer for Point Cloud Understanding

Jaesung Choe^{†1}, Chungyun Park^{†2},
Francois Rameau¹, Jaesik Park², and In So Kweon¹

¹ Korea Advanced Institute of Science and Technology (KAIST), South Korea

² Pohang University of Science and Technology (POSTECH), South Korea

<https://github.com/LifeBeyondExpectations/ECCV22-PointMixer>

{jaesung.choe, frameau, iskweon77}@kaist.ac.kr

{p0125ch, jaesik.park}@postech.ac.kr

Abstract. MLP-Mixer has newly appeared as a new challenger against the realm of CNNs and Transformer. Despite its simplicity compared to Transformer, the concept of channel-mixing MLPs and token-mixing MLPs achieves noticeable performance in image recognition tasks. Unlike images, point clouds are inherently sparse, unordered and irregular, which limits the direct use of MLP-Mixer for point cloud understanding. To overcome these limitations, we propose **PointMixer**, a universal point set operator that facilitates information sharing among unstructured 3D point cloud. By simply replacing token-mixing MLPs with Softmax function, PointMixer can “mix” features within/between point sets. By doing so, PointMixer can be broadly used for intra-set, inter-set, and hierarchical-set mixing. We demonstrate that various channel-wise feature aggregation in numerous point sets is better than self-attention layers or dense token-wise interaction in a view of parameter efficiency and accuracy. Extensive experiments show the competitive or superior performance of PointMixer in semantic segmentation, classification, and reconstruction against Transformer-based methods.

1 Introduction

3D scanning devices, such as LiDAR or RGB-D sensors, are widely used to capture a scene as 3D point clouds. Unlike images, point clouds are inherently sparse, unordered, and irregular. These properties make standard neural network architectures [21,62] hardly applicable. To tackle these challenges, there have been numerous ad hoc solutions, such as sparse convolution networks [8,15], graph neural networks [57,61,74,75,78,84,6], and point convolution networks [48,67,76,83]. Despite their structural differences, these techniques have all been designed to extract meaningful feature representation from point clouds [18]. Among existing solutions, Transformer [7,17,73,96,54] appears to be particularly beneficial to extract features from point clouds. Indeed, the self-attention layer that encompasses the dense token-wise relations is specifically relevant in the context of processing irregular and unordered 3D points.

[†] Both authors have equally contributed to this work.

Beyond the well established realm of CNNs and Transformer, MLP-Mixer [68] proposes a new architecture that exclusively uses MLPs. Based on the pioneering study [68], concurrent works [5,16,23,35,41,69,88,89] address the locality issue of MLPs [5,11,23,39,88,89] and discuss the necessity of self-attention [41,65,92]. More recently, Yu *et al.* [90] claim that the general architecture formulation is more important than the specific token-wise interaction strategies, such as self-attention and token-mixing MLPs. Despite an increasing interest, the MLP-like architectures for point clouds have not yet been fully explored.

In this paper, we introduce the *PointMixer* that newly extends the philosophy of MLP-like architectures to point cloud analysis. Specifically, we demonstrate the dense token-wise interaction are not essential factors in this context. Instead of using token-mixing MLPs, we extend the usage of channel-mixing MLPs into numerous point sets. As illustrated in Fig. 1 and Table 1, PointMixer layer shares and mixes point features (1) within grouped points (intra-set), (2) between point sets (inter-set), or (3) points in different hierarchical sets. In particular, we newly introduce the concepts of the inter-set mixing and the hierarchical-set mixing, which is clearly different from previous studies [47,56,96] that only focus on the intra-set mixing. To this end, PointMixer layer is a universal point set operator that can propagate point responses into various point sets. We claim that various channel MLPs on numerous point sets can outperform self-attention layers or token-mixing MLPs for point clouds. Moreover, the PointMixer network is a general architecture that has symmetric encoder-decoder blocks fully equipped with PointMixer layers.

We conduct extensive experiments on various 3D tasks. These large-scale experiments demonstrate that our method achieves compelling performance among Transformer-based studies [51,96]. Our contributions are summarized as below:

- PointMixer as a universal point set operator that facilitates mixing of point features through various point sets: intra-set, inter-set, and hierarchical-set.
- Highly accurate and parameter-efficient block design that purely consists of channel-mixing MLPs without token-mixing MLPs.
- Symmetric encoder-decoder network to propagate hierarchical point responses through PointMixer layer, instead of trilinear interpolation.
- Extensive experiments in various 3D point cloud tasks that highlight the efficacy of PointMixer network against recent Transformer-based studies.

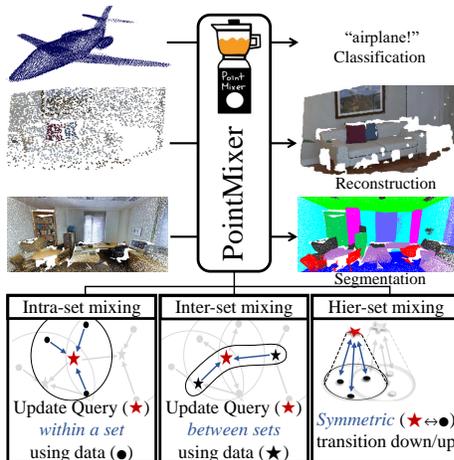


Fig. 1. We present a MLP-like architecture for various point cloud processing, which considers numerous point sets with larger receptive fields to “mix” information.

Table 1. Technical comparisons. Locality represents the local feature aggregation among sampled points. We split the function of the set operator as “intra”-set mix, “inter”-set mix, and “hierarchical”-set mix. Also, we present the symmetric property of encoder-decoder architecture of related work [55,56,96].

Method	Locality	Set operator			Symmetric pyramid arch	Token-mix
		Intra	Inter	Hier		
PointNet [55]	✗	✗	✗	✗	✗	Pooling
PointNet++ [56]	✓	✓	✗	✗	✗	Pooling
PointTrans [96]	✓	✓	✗	✗	✗	Self-attn
PointMLP [47]	✓	✓	✗	✗	✗	Affine
PointMixer (ours)	✓	✓	✓	✓	✓	Softmax

2 Related Work

In this section, we revisit previous approaches for point cloud understanding, and then briefly introduce Transformers and MLP-like architectures.

Deep learning on point clouds. Point clouds are naturally sparse, unordered, and irregular, which makes it difficult to design a deep neural network for point cloud understanding. To handle such complex data structures, two distinct philosophies have been investigated: voxel-based and point-based methods.

Voxel-based methods [34,8,14,50,63,98] first quantize an irregular point cloud into the regular voxel grids, which makes it efficient to search neighbor voxels. However, the voxelization process inevitably loses the geometric details of the original point cloud. This issue often leads to infer inaccurate predictions though several recent methods [45,64,93] try to alleviate the quantization artifacts. Therefore, point-based methods have been actively studied.

PointNet [55] is a pioneering paper that processes an unordered and irregular points in deep neural architectures. Based on this seminal work, PointNet++ [56] presents the ways of involving feature hierarchy as well as points’ locality. In details, this paper adopts k -Nearest Neighbor (k NN) for local neighborhood sampling and the Farthest Point Sampling algorithm (FPS) for feature hierarchy (*e.g.*, transition downsampling). This pyramid encoder-decoder network largely influences on the MLP-based methods [58,95,28,42,25,31,36,26,82,85], point convolution studies [33,38,40,43,48,79,83], and graph-based networks [57,61,78,75,84]. However, as stated in Table 1, PointNet++ is limited to capture local point responses within the scope of intra-set. Moreover, we found that this vanilla architecture is asymmetric between transition down layers and transition up layers. While downsampling layers adopt pooling with k NN and FPS, upsampling layers re-compute k NN and trilinear interpolation, which brings asymmetric feature propagation in the encoder-decoder architecture. More recently, Transformer-based study for point cloud processing [96] still suffers from the same issue in the pyramid architecture design.

Our work unifies a local feature aggregation layer, a downsampling layer, and an upsampling layer into *an universal set operator*, named PointMixer layer. This novel layer brings the symmetric and learnable down/upsampling architecture for various 3D perception tasks such as object shape classification [80], semantic segmentation [2] and point cloud reconstruction tasks [7].

Transformers and MLP-like architectures. Transformer-based architecture has recently become a game changer in both natural language processing [10,59,65,73] and computer vision [3,12,19,44,49,60,70,96,91].

Vision Transformer (ViT) [12] has opened the applicability of Transformers on visual recognition tasks. Because of the quadratic runtime of the self-attention layers in Transformers, ViT adopts tokenized inputs that divide the image into small region of patches [71]. The idea of patch embeddings is widely used in the following studies that focus on various issues in ViT [30,66]: locality [60,72,13,24] and hierarchy [77,81,44,4]. With those self-attention layers [60,94], Transformer-based point cloud studies [96,17,51] demonstrate accurate predictions in both 3D shape classification [80] and semantic segmentation [2]. However, a general-purpose layer for both local feature learning and down/upsampling has drawn little attention in handling 3D points.

Recently, there exist trials to go beyond the hegemony of CNNs and Transformer by introducing MLP-like architectures. The pioneering paper, MLP-Mixer [68], presents a MLP-like network constituted of token-mixing MLPs and channel-mixing MLPs for image classification task. Especially in computer vision, this MLP-like architectures appear to be a new paradigm with their simple formulation and superior performance given large-scale training data. Subsequent papers raise issues and develop potentials in MLP-like architectures: (1) can MLPs handle position-sensitive information or locality [39,23,88,5,11]? and (2) does self-attention is truly needed [65,41,92]? Though these issues are still controversial, more recent paper, Metaformer [90], addresses the importance of general architecture formulation instead of the specific dense token-wise interaction strategies such as self-attention layers [73,12] or token-mixing MLPs [68]. Simply, by replacing complicated token-mixing operators with the average pooling layer, MetaFormer achieves remarkable performance against the recent MLP-based and Transformer-based studies.

Despite their success, modern MLP-like approaches have not yet been fully exploited to point clouds. In contrast to recent Transformer-based point cloud studies [96,51,17,54], our PointMixer network is a general-purpose architecture that symmetrically upsamples/downsamples points’ responses and truly exploits the strength of MLPs to operate mixing within/beyond sets of points. By doing so, we successfully conduct various tasks, 3D semantic segmentation, point cloud reconstruction, and object classification tasks [2,7,80].

3 Method

In this section, we describe the details of our PointMixer design. For the sake of clarity, we compare the general formulation of MLP-Mixer with representative point-based approaches, such as PointNet++ and Point Transformer (Sec. 3.1). Then, we examine whether MLP-Mixer is of relevance to a point set operator (Sec. 3.2). Finally, we introduce our PointMixer layer (Sec. 3.3) that is adopted in our entire network (Sec. 3.4).

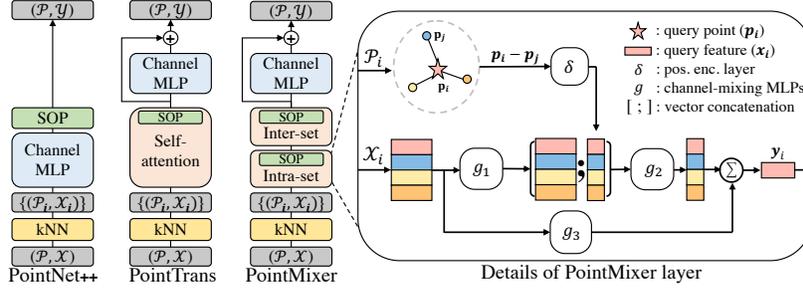


Fig. 2. Block design comparison and PointMixer layer details. Note that SOP means symmetric operation, such as pooling and summation.

3.1 Preliminary

Let's assume a point cloud $\mathcal{P} = \{\mathbf{p}_i\}_{i=1}^N$ and its corresponding point features $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^N$ where $\mathbf{p}_i \in \mathbb{R}^3$ is the position of the i -th point and $\mathbf{x}_i \in \mathbb{R}^C$ is its corresponding feature. The objective is to learn a function $f: \mathcal{X} \rightarrow \mathcal{Y}$ to produce the output point features $\mathcal{Y} = \{\mathbf{y}_i\}_{i=1}^N$. Instead of processing the entire point cloud directly, most approaches treat data locally based on points' proximity. For this purpose, k -Nearest Neighbor (k NN) [52,29,97] is widely used to get an index map of neighbor points, which is denoted as:

$$\mathcal{M}_i = k\text{NN}(\mathcal{P}, \mathbf{p}_i), \quad (1)$$

where \mathcal{M}_i is an index map of the K closest points for a query point $\mathbf{p}_i \in \mathbb{R}^3$. In other words, given a query index i , the index map \mathcal{M}_i is defined as a set of K nearest neighbor indices ($K = |\mathcal{M}_i|$). Accordingly, k NN can be understood as a directional graph that links each query point \mathbf{p}_i to its K closest neighbor points $\mathcal{P}_i = \{\mathbf{p}_j \in \mathcal{P} | j \in \mathcal{M}_i\}$ and the corresponding features $\mathcal{X}_i = \{\mathbf{x}_j \in \mathcal{X} | j \in \mathcal{M}_i\}$. **PointNet++** [56] addresses the problem of PointNet [55] that has difficulty in capturing local responses. To cope with this problem, it utilizes k NN and Farthest Point Sampling algorithm and builds asymmetric encoder-decoder network. Instead of dealing with the entire point cloud directly, PointNet++ aggregates set-level responses locally as follows:

$$\mathbf{y}_i = \text{maxpool}_{j \in \mathcal{M}_i} \left(\text{MLP}([\mathbf{x}_j; \mathbf{p}_i - \mathbf{p}_j]) \right), \quad (2)$$

where \mathbf{y}_i is the output feature vector for the i -th point and $[\cdot]$ denotes vector concatenation. By adopting this grouping and sampling strategy, MLPs can capture local responses from unordered 3D points.

Point Transformer [96] adopts vector subtraction attention as a similarity measurement for token-wise communication. Following PointNet++, Point Transformer also uses k NN to compute local responses. Given a query point feature \mathbf{x}_i and its neighbor feature set \mathcal{X}_i , Point Transformer operates self-attention layers to densely relate token interaction as:

$$\mathbf{y}_i = \sum_{j \in \mathcal{M}_i} \text{softmax} \left(\psi(\mathbf{W}_1 \mathbf{x}_i - \mathbf{W}_2 \mathbf{x}_j + \delta(\mathbf{p}_i - \mathbf{p}_j)) \right) \left(\mathbf{W}_3 \mathbf{x}_j + \delta(\mathbf{p}_i - \mathbf{p}_j) \right), \quad (3)$$

where \mathbf{W} indicates a linear transformation matrix, $\psi(\cdot)$ denotes an MLPs to calculate vector similarities, $\delta(\mathbf{p}_i - \mathbf{p}_j)$ is a positional encoding vector to embed local structures of 3D points, and $\mathbf{p}_i - \mathbf{p}_j$ is the relative distance between a query point \mathbf{p}_i and its neighbor point \mathbf{p}_j .

3.2 MLP-Mixer as Point Set Operator

MLP-Mixer [68] has achieved remarkable success by only using MLPs for image classification. However, when dealing with sparse and unordered points, the direct application of the MLP-Mixer network is restricted. Let us revisit MLP-Mixer to ease the understanding of our PointMixer.

The MLP-Mixer layer³ consists of token-mixing MLPs and channel-mixing MLPs. MLP-Mixer takes K tokens having C -dimensional features, denoted as $\mathbf{X} \in \mathbb{R}^{K \times C}$, where tokens are features from image patches. It begins with token-mixing MLPs that transposes the spatial axis and channel axis to mix spatial information. Then, it continues with channel-mixing MLPs so that input tokens are mixed in spatial and channel dimensions.

$$\mathbf{X}' = \mathbf{X} + (\mathbf{W}_2 \rho(\mathbf{W}_1 (\text{LayerNorm}(\mathbf{X}))^\top))^\top, \quad (4)$$

$$\mathbf{Y} = \mathbf{X}' + \mathbf{W}_4 \rho(\mathbf{W}_3 \text{LayerNorm}(\mathbf{X}')), \quad (5)$$

where \mathbf{W} is the weight matrix of a linear function and ρ is GELU [22]. By Eq. (4), token-mixing MLPs are sensitive to the order of the input tokens, which is permutation-variant property. Due to its property, positional encoding is not required in the vanilla MLP-Mixer as stated in the paper [68]. However, this property is not desirable for processing irregular and unordered point clouds, which is different characteristics of uniform and ordered pixels in the image. To cope with this issue, previous point-based layers [56,96] are independent to orders of input point, *i.e.*, permutation-invariance⁴. Moreover, as a point set operator, vanilla MLP-Mixer layer only computes *intra-set* relations as PointNet++ and Point Transformer do. From this analysis, we observe room for improvement in point cloud understanding. We propose PointMixer layer that is permutation-invariant and can also be used for a learnable upsampling in Sec. 3.3.

3.3 PointMixer: Universal Point Set Operator

We introduce an approach to embed geometric relations between points' features into the MLP-Mixer's framework. As illustrated in Fig. 2, the PointMixer layer takes a point set $\mathcal{P}_i = \{\mathbf{p}_j\}$ and its associated point features set $\mathcal{X}_i = \{\mathbf{x}_j\}$ as inputs in order to compute the output feature vector \mathbf{y}_i . For a point \mathbf{p}_i , we first compute a score vector $\mathbf{s} = [s_1, \dots, s_K]$ to aggregate \mathcal{X}_i as follows:

$$s_j = g_2\left([g_1(\mathbf{x}_j); \delta(\mathbf{p}_i - \mathbf{p}_j)]\right) \quad \text{where } \forall j \in \mathcal{M}_i, \quad (6)$$

³ We set the relation of terminologies as layer \subset block \subset network.

⁴ To deal with unordered points, layers are permutation-invariant ($f_{\text{layer}} : \mathcal{X}_i \rightarrow \mathbf{y}_i$) and blocks are permutation-equivariant ($f_{\text{block}} : \mathcal{X} \rightarrow \mathcal{Y}$).

where $g(\cdot)$ is the channel-mixing MLPs, $\delta(\cdot)$ is the positional encoding MLPs, and \mathbf{x}_j is a j -th element of the feature vector set \mathcal{X}_i . Note that we follow the relative positional encoding scheme [56,96] to deal with unstructured 3D points. As a result, we obtain the score vector $\mathbf{s} \in \mathbb{R}^K$. Finally, the features of the K adjacent points are gathered to produce a new feature vector \mathbf{y}_i as:

$$\mathbf{y}_i = \sum_{j \in \mathcal{M}_i} \text{softmax}(s_j) \odot g_3(\mathbf{x}_j), \quad (7)$$

where $\text{softmax}(\cdot)$ is the Softmax function that normalizes the spatial dimension, and \odot indicates an element-wise multiplication. Note that this symmetric operation (SOP) in the proposed PointMixer is different from both the average pooling in MLP-Mixer and the max pooling in PointNet++, as described in Table 1.

As a set operator, PointMixer layer has different characteristics compared to both MLP-Mixer layer [68]. First, PointMixer layer sees relative positional information $\delta(\mathbf{p}_i - \mathbf{p}_j)$ to encode the local structure of a point set. Second, the vanilla MLP-Mixer layer does not have the Softmax function. Last, PointMixer layer does not have token-mixing MLPs for scalability to arbitrary number of neighbor points and for permutation-invariance to deal with unordered points. Let us explain the reasons behind these differences.

No token-mixing MLPs. There are two reasons that we do not put token-mixing MLPs into PointMixer layer. First, token-mixing MLPs are permutation-variant, which makes it incompatible with unordered point clouds. As stated in Point Transformer [96], permutation-invariant property is a necessary condition, also for PointMixer layer. Second, while a given pixel in an image systematically admits 8 adjacent pixels, each 3D point does not have pre-defined number of neighbors, which is determined by the clustering algorithms, such as k NN [78,56,96], radius neighborhood [67], or hash table [8,7]. Since token-mixing MLPs can only take fixed number of input points⁵, token-mixing MLPs are not suitable for handling various cardinality of point sets.

Inspired by Synthesizer [65], we alleviate this problem by replacing the token-mixing MLPs with the Softmax function. We conjecture that the Softmax function weakly binds token-wise information in a non-parametric manner (*i.e.*, normalization). By doing so, PointMixer layer can calculate arbitrary cardinality of point sets and have a permutation-invariant property. As described in Fig. 3 and in Table 1, PointMixer layer can be used as a universal point operator for mixing various types of point sets: intra-set, intra-set, and hierarchical-set.

Intra-set mixing. Given a point set \mathcal{P}_i and its corresponding feature set \mathcal{X}_i , intra-set mixing aims to compute point-wise interaction within each set. Usually, k NN is widely used to cluster neighbor points into groups. For example in Fig. 3-(a), we apply k NN on each query point where $K=3$. As a result, a red point (\star) has three neighbor points (\bullet). Based on the index map \mathcal{M}_i from k NN, the PointMixer layer updates a query feature \mathbf{x}_i using its neighbor feature set \mathcal{X}_i , as in Eq. (6) and Eq. (7). While intra-set mixing is proven to be useful in various methods [18], the receptive field is bounded within a set, as depicted in Fig. 3-(a). To overcome this restriction, we propose the inter-set mixing operation.

⁵ Please refer to the supplementary material for further details.

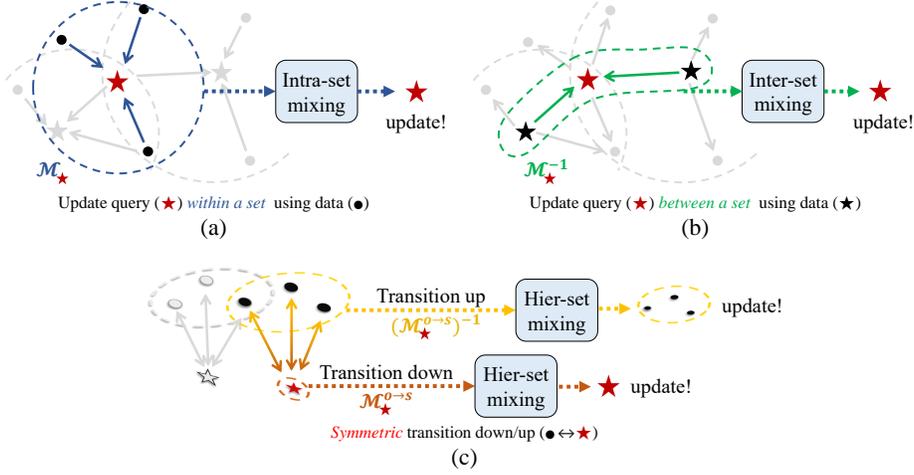


Fig. 3. PointMixer is a universal point set operator: (a) intra-set mixing, (b) inter-set mixing, and (c) hierarchical-set mixing.

Inter-set mixing. This is a new concept of spreading point features between different sets. Using the index mapping \mathcal{M}_i , we can trace back to find another set \mathcal{P}_j that includes a query point \mathbf{p}_i as their neighbors. This process can be viewed as the inverse version of k NN, and we define the inverse mapping of \mathcal{M}_i as $\mathcal{M}_i^{-1} = \{j | i \in \mathcal{M}_j\}$. For example in Fig. 3-(b), given index mapping $\mathcal{M}_{★}$, we compute inverse index mapping $\mathcal{M}_{★}^{-1}$. Then, we can find the two adjacent sets whose query points are black points (★). It implies the red point is included in two adjacent sets, as drawn in Fig. 3-(b)⁶. In shorts, inverse mapping \mathcal{M}_i^{-1} finds the set index j that includes a point \mathbf{p}_i . By doing so, inter-set mixing can aggregate point responses between neighbor sets \mathcal{P}_j into the query point \mathbf{p}_i .

Hierarchical-set mixing. The PointMixer layer is universally applicable for the transition down/up layers as shown in Fig. 3-(c). For instance, let's prepare a point set $\mathcal{P}_s = \{\mathbf{p}_j\}_{j=1}^{N'}$ that is sampled from original point set $\mathcal{P}_o = \{\mathbf{p}_i\}_{i=1}^N$ ($\mathcal{P}_s \subset \mathcal{P}_o$). Using a point $\mathbf{p}_j \in \mathcal{P}_s$, we calculate its neighbors from \mathcal{P}_o and obtain index mapping $\mathcal{M}_j^{o \rightarrow s}$. By putting $\mathcal{M}_j^{o \rightarrow s}$ in Eq. (6) and Eq. (7), we readily pass the feature from \mathcal{P}_o to \mathcal{P}_s (*i.e.*, point *downsampling*), which is computed as:

$$\mathcal{M}_j^{o \rightarrow s} = k\text{NN}(\mathcal{P}_o, \mathbf{p}_j) \quad \text{where } \forall \mathbf{p}_j \in \mathcal{P}_s. \quad (8)$$

For point *upsampling*, we notice that conventional U-Net in both PointNet++ and Point Transformer is *not symmetric* in terms of downsampling and upsampling. This is because the spatial grouping is performed asymmetrically as visualized in Fig. 4. In details, conventional approaches [56,96,95] build another k NN map $\mathcal{M}_i^{s \rightarrow o}$ for the upsampling as below:

$$\mathcal{M}_i^{s \rightarrow o} = k\text{NN}(\mathcal{P}_s, \mathbf{p}_i) \quad \text{where } \forall \mathbf{p}_i \in \mathcal{P}_o. \quad (9)$$

⁶ There is a chance to collect variable number of points after an inverse mapping \mathcal{M}_i^{-1} .

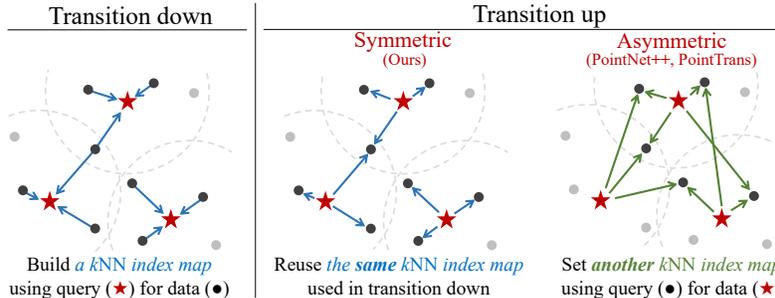


Fig. 4. Comparison of transition layer: symmetric (ours), asymmetric [56,95,96].

However, this is not symmetric because *nearest neighbor is not a symmetric function*: even if \mathbf{p}_i 's nearest neighbor is \mathbf{p}_j , \mathbf{p}_j 's nearest neighbor may not be \mathbf{p}_i .

Instead of creating a new index map $\mathcal{M}_i^{s \rightarrow o}$, our PointMixer layer can use $(\mathcal{M}_j^{o \rightarrow s})^{-1}$ for upsampling (Fig. 3(b)) by re-using the original index mapping from the downsampling $\mathcal{M}_j^{o \rightarrow s}$. See Fig. 4 for the symmetric upsampling. The benefit of this approach is that it does not introduce additional k -Nearest Neighbor search. Furthermore, we can propagate point responses in different hierarchy based on their scores computed by our PointMixer, instead of using trilinear interpolation [56,96]. These technical differences result in higher performance than that of asymmetric design on dense prediction tasks (see Sec. 4.4).

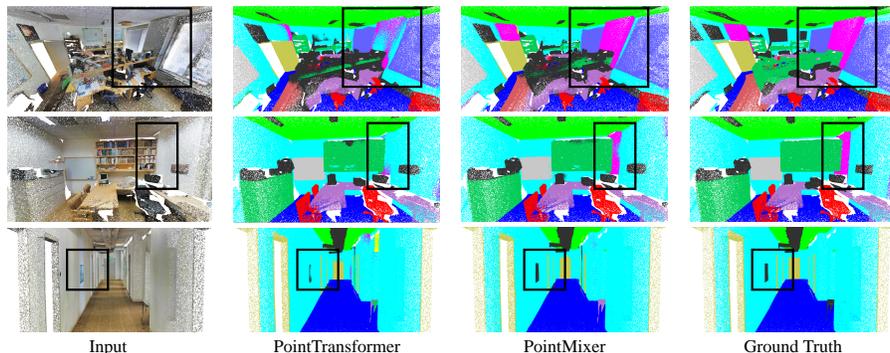
As illustrated in Fig. 5(b) of the transition down block, we use Farthest Point Sampling algorithm to produce \mathcal{P}_s from \mathcal{P}_o . Then, we utilize k NN to sample \mathcal{P}_s from \mathcal{P}_o . The resulting point locations are used to calculate the relative distance $\mathbf{p}_i - \mathbf{p}_j$. In the transition up block, we keep using the index map $\mathcal{M}_j^{o \rightarrow s}$ calculated in the transition down block. To this end, we apply the PointMixer layer in transition up/down while maintaining the symmetric relation between the sampled point cloud \mathcal{P}_s and the original points \mathcal{P}_o . We empirically prove that this symmetric upsampling layer helps the network to predict dense point-level representations accurately in Sec. 4.

3.4 Network Architecture

In this section, we describe the details of our MLP-like encoder-decoder architecture as shown in Fig. 5. Our network is composed of several MLP blocks, such as the transition down blocks, transition up blocks, and Mixer blocks. For a fair comparison, our network mainly follows the proposed hyper-parameters in Point Transformer [96] for network composition. Overall, our network takes a deep pyramid-style network that progressively downsamples points to obtain global features. For dense prediction tasks, such as semantic segmentation or point reconstruction, we include upsampling blocks for per-point estimation. Finally, our header block is designed for task-specific solutions. For classification, we take fully-connected layers, dropout layers, and global pooling layer. For semantic segmentation and point reconstruction, the header block consists of MLPs without pooling layers.

Table 2. Semantic segmentation results on S3DIS Area 5 test dataset [2].

Method	Param.	mAcc	mIoU	ceiling	floor	wall	column	window	door	table	chair	sofa	book	board	clutter
PACConv [83]	-	73.0	66.6	94.6	98.6	82.4	26.4	58.0	60.0	80.4	89.7	69.8	74.3	73.5	57.7
KPConv <i>deform</i> [67]	-	72.8	67.1	92.8	97.3	82.4	23.9	58.0	69.0	81.5	91.0	75.4	75.3	66.7	58.9
MinkowskiNet [8]	37.9M	71.7	65.4	91.8	98.7	86.2	34.1	48.9	62.4	81.6	89.8	47.2	74.9	74.4	58.6
PointTrans [96]	7.8M	76.5	70.4	94.0	98.5	86.3	38.0	63.4	74.3	89.1	82.4	74.3	80.2	76.0	59.3
FastPointTrans [54]	37.9M	77.9	70.3	94.2	98.0	86.0	53.8	61.2	77.3	81.3	89.4	60.1	72.8	80.4	58.9
PointMixer (ours)	6.5M	77.4	71.4	94.2	98.2	86.0	43.8	62.1	78.5	90.6	82.2	73.9	79.8	78.5	59.4

**Fig. 6.** Qualitative results in semantic segmentation on S3DIS Area 5 test dataset [2].

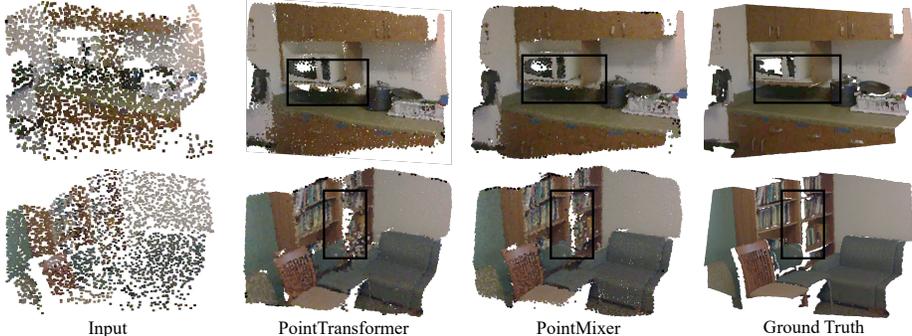
4.2 Point Cloud Reconstruction

To highlight the versatility of our approach, we propose a large-scale assessment for the newly introduced task of point cloud reconstruction [7] where an inaccurate point cloud is jointly denoised, densified, and completed. This experiment is also particularly interesting to evaluate the generalization capabilities of networks since it allows us to test methods in unmet environments. Specifically, we train on the synthetic objects of ShapeNetPart [87] and evaluate the reconstruction accuracy on unmet indoor scenes from ScanNet [9] (real reconstruction) and ICL-NUIM [20] (synthetic data). Note that in [7], the point reconstruction is performed in two stages: 1) point upsampling via a sparse hourglass network, 2) denoising and refinement via Transformer network. For this evaluation, we propose to replace the second stage of this pipeline with various architectures (*i.e.*, PointMixer network and previous studies [55,56,96]) to compare their performances. Under the same data augmentation and data pre-processing as in [7,37], we train PointMixer and previous studies [55,56,96]. For evaluation, we utilize the Chamfer distance (CD) to measure the distance between the predictions and the ground truth point clouds. Additionally, we use the accuracy (Acc.), completeness (Cp.), and F1 score to measure the performance in occupancy aspects [27,1,86,32].

Though our network is solely trained in a synthetic/object dataset, our network can generalize towards unmet scenes including real-world 3D scans [9] and synthetic/room-scale point clouds [20] as in Table 3 and Fig. 7. Moreover, our method compares favorably to previous Transformer-based [96,7] and MLP-based [55,56] studies. In particular, the performance gap between ours and previous studies become larger in the ScanNet [9] and ICL-NUIM dataset [20], which indicates better generalization performance.

Table 3. Point cloud reconstruction results.

Method	ShapeNet-Part [87]				ScanNet [9]				ICL-NUIM [20]			
	CD(↓)	Acc.(↑)	Cp.(↑)	F1(↑)	CD(↓)	Acc.(↑)	Cp.(↑)	F1(↑)	CD(↓)	Acc.(↑)	Cp.(↑)	F1(↑)
PointNet [55]	1.33	63.2	38.8	48.2	3.05	37.5	27.8	32.6	2.98	46.9	33.2	38.1
PointNet++ [56]	1.25	65.1	39.0	50.1	2.97	38.3	29.5	33.4	2.88	48.8	35.8	39.9
PointRecon [7]	1.19	81.0	40.4	53.4	2.86	40.4	30.2	34.1	2.78	54.1	38.1	43.6
PointTrans [96]	1.12	75.9	40.9	52.7	2.79	41.1	32.1	35.6	2.57	51.1	36.4	41.6
PointMixer (ours)	1.11	77.1	41.5	53.7	2.74	42.1	33.5	37.8	2.43	56.5	38.2	44.7

**Fig. 7.** Qualitative results in point reconstruction on ScanNet dataset [9].

4.3 Shape Classification

The ModelNet40 [80] dataset has 12K CAD models with 40 object categories. We follow the official train/test splits to train and evaluate ours and the previous studies. For fair comparison, we follow the data augmentation and pre-processing as proposed in PointNet++ [56], which is also adopted in the recent studies [96,48,67,83]. For evaluation, we adopt the mean accuracy within each category (mAcc), and the overall accuracy over all classes (OA) with the same evaluation protocol with previous approaches [56,96,83]. The results presented in Table 4 show that our approach outperforms the recently proposed techniques. Especially, PointMixer achieves the highest mAcc with outperforming other methods by a large margin as 0.8 mAcc, using 1024 points without normals. Moreover, our network shows this competitive performance with less parameters (**3.9M**) against previous studies (Point Transformer⁷ 5.3M, KPConv 15.2M). Based on these comparisons with Point Transformer and related work, we conclude that PointMixer network effectively and efficiently aggregates various point responses through intra-set mixing, inter-set mixing, and downsampling layers, even for shape classification task.

Table 4. Shape classification results on ModelNet40 dataset [80].

Method	Param.	mAcc	OA
PointNet [55]	-	86.2	89.2
PointNet++ [56]	1.4M	-	90.7
PointConv [79]	18.6M	-	92.5
DGCNN [78]	-	90.2	92.9
KPConv <i>rigid</i> [67]	15.2M	-	92.9
PointTrans [96]	5.3M	90.6	93.7
PointMixer (ours)	3.9M	91.4	93.6

⁷ Since there is no official release of codes, we use the best implementation of Point Transformer available in the public domain, which contains the official code provided by the authors of Point Transformer and reproduces the reported accuracy.

4.4 Ablation Study

We conduct an extensive ablation study about the proposed PointMixer in semantic segmentation on the S3DIS dataset [2] and in point cloud reconstruction on the ShapeNet-Part dataset [87]. First, we proceed the case study of our PointMixer (Table 5, top-right one). Second, we analyze the influence of positional encoding on PointMixer (Table 5, left one). Last, we compare PointMixer that is free from token-mixing MLPs (Table 5, bottom-right one).

Universal point set operator. PointMixer can be applicable to act as intra-set mixing, inter-set mixing, and hierarchical-set mixing. Technically speaking, intra-set mixing and hierarchical-set mixing for point downsampling utilizes \mathcal{M}_i , but inter-set mixing and hierarchical-set mixing for point upsampling use \mathcal{M}_i^{-1} . This ablation study aims to validate the inverse mapping as well as a functionality of PointMixer layer. As in Table 5, especially, when we combine the usage of hierarchical-set mixing and inter-set mixing, the synergy brings large performance improvement in both two tasks, 2.2 mIoU in semantic segmentation and 0.02 Chamfer distance in point reconstruction. It implies that the various ways of sharing points’ responses are beneficial for point cloud understanding.

Unnecessary token-mixing MLPs. We validate our claims about the role of the Softmax function, *i.e.*, weakly binding tokens. For this purpose, we replace the Softmax functions with token-mixing MLPs, as proposed in the vanilla MLP-Mixers. Table 5 demonstrates that even without explicit use of token-mixing MLPs, our PointMixer successfully achieves similar accuracy in semantic segmentation and point reconstruction. Moreover, it takes 3 days for training the PointMixer with token-mixing MLPs while our original PointMixer requires a day (twice faster) with less parameter consumption (18% less). These results are consistently support to Metaformer [90] and Synthesizer [65] in that dense token-wise interaction is not an essential choice.

Positional encoding. MLP-Mixer does not use positional encoding since token-mixing MLPs are sensitive to the order of tokens. However, relative position information is an important factor to handle unstructured 3D points. Also, our PointMixer layer is free from token-mixing MLPs to obtain permutation-invariant property to function as a universal point set operator. Without any modification on the layer except positional encoding layers $\delta(\mathbf{p}_i - \mathbf{p}_j)$, we experiment the effectiveness of positional encoding in all of these usages. As in Table 5, there is large performance drop when we intentionally omit positional encodings in our PointMixer layer. Different from MLP-Mixer, we claim the relative positional encoding is a necessary condition to deal with 3D points.

Symmetric architecture. When we think of the convolutional neural networks for image recognition, it is natural to design symmetric encoder network and decoder network. For instance in the previous paper [53], this paper verify the importance of symmetric transposed convolution layer design for semantic segmentation task. In contrast, point-based studies [96,95] dominantly rely on the asymmetric PointNet++ architecture, which is similar to FCN [46].

The top-right sub-table in Table 5 supports our claim. When we do not use hierarchical-set mixing, the architecture become asymmetric, and transition

Table 5. Ablation study about positional encoding (left), types of mixing (top-right) and token-mixing MLPs (bottom-right). In the bottom-right table, we denote channel-mixing MLPs and token-mixing MLPs as C-MLP and T-MLP, respectively. Note that number of parameters (Param.) and training time (Time) is measured under semantic segmentation task on S3DIS dataset [2].

Preserve (✓)			Semantic seg		Point recon	
Pos. enc. in			mIoU	mAcc	CD(↓)	F1(↑)
Intra	Inter	Hier				
			64.7	71.8	1.13	53.0
		✓	65.9	72.5	1.13	53.0
	✓		64.0	70.5	1.13	53.0
	✓	✓	66.1	72.9	1.12	53.6
✓			69.6	76.3	1.12	53.6
✓		✓	70.2	76.8	1.13	52.5
✓	✓		69.9	76.5	1.13	53.5
✓	✓	✓	71.4	77.4	1.11	53.7

Preserve (✓)			Semantic seg		Point recon	
Intra	Inter	Hier	mIoU	mAcc	CD(↓)	F1(↑)
✓			69.2	75.8	1.13	53.5
✓	✓		69.1	75.7	1.12	53.5
✓		✓	69.3	76.6	1.11	53.5
✓	✓	✓	71.4	77.4	1.11	53.7

C-MLP	T-MLP	Softmax	Param.	Time	mIoU	CD(↓)
✓			6.5M	40h	58.3	1.23
✓		✓	6.5M	44h	71.4	1.11
✓	✓		7.4M	88h	71.1	1.12
✓	✓	✓	7.4M	95h	71.1	1.11

down and up layers are identical to that in PointNet++ [56] and Point Transformer [96]. It turns out that the asymmetric architecture (the second row) degrades the performance of the network (the last row) in terms of both mIoU and Chamfer distance by 2.3 absolute percentage and 0.01, respectively. Moreover, when we apply inter-set mixing, the performance gap further increases in both semantic segmentation and point reconstruction tasks.

5 Conclusion

We propose a MLP-like architecture for point cloud understanding, which focuses on sharing point responses in numerous and diverse point sets through a universal point set operator, *PointMixer*. Regardless of the cardinality of a point set, PointMixer layer can “mix” point responses in intra-set, inter-set, and hierarchical-set. Moreover, we present a point-based general architecture that involves symmetric encoder-decoder blocks for propagating information through hierarchical point sets. Extensive experiments validate the efficacy of our PointMixer network with superior or compelling performance compared to Transformer-based studies. Through out this paper, we claim that dense token-wise calculation, such as self-attention layers or token-mixing MLPs, is not an essential choice for point cloud processing. Instead, we emphasize the importance of information sharing toward various point sets.

Acknowledgements

(1) IITP grant funded by the Korea government(MSIT) (No.2019-0-01906, Artificial Intelligence Graduate School Program(POSTECH)) and (2) Institute of Information and communications Technology Planning and Evaluation (IITP) grant funded by the Korea government(MSIT) (No.2021-0-02068, Artificial Intelligence Innovation Hub)

References

1. Aanæs, H., Jensen, R.R., Vogiatzis, G., Tola, E., Dahl, A.B.: Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision (IJCV)* **120**(2), 153–168 (2016)
2. Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., Savarese, S.: 3d semantic parsing of large-scale indoor spaces. In: *Proceedings of the IEEE International Conference on Computer Vision and Pattern Recognition* (2016)
3. Bello, I.: Lambdanetworks: Modeling long-range interactions without attention. In: *International Conference on Learning Representations* (2020)
4. Chen, C.F.R., Fan, Q., Panda, R.: Crossvit: Cross-attention multi-scale vision transformer for image classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 357–366 (2021)
5. Chen, S., Xie, E., Ge, C., Liang, D., Luo, P.: Cyclemlp: A mlp-like architecture for dense prediction. In: *International Conference on Learning Representations (ICLR)* (2022)
6. Choe, J., Joo, K., Imtiaz, T., Kweon, I.S.: Volumetric propagation network: Stereolidar fusion for long-range depth estimation. *IEEE Robotics and Automation Letters* **6**(3), 4672–4679 (2021)
7. Choe, J., Joung, B., Rameau, F., Park, J., Kweon, I.S.: Deep point cloud reconstruction. In: *International Conference on Learning Representations (ICLR)* (2022)
8. Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 3075–3084 (2019)
9. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5828–5839 (2017)
10. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. In: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*. pp. 4171–4186 (2019)
11. Ding, X., Xia, C., Zhang, X., Chu, X., Han, J., Ding, G.: Repmlp: Re-parameterizing convolutions into fully-connected layers for image recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2022)
12. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
13. d’Ascoli, S., Touvron, H., Leavitt, M.L., Morcos, A.S., Biroli, G., Sagun, L.: Convit: Improving vision transformers with soft convolutional inductive biases. In: *International Conference on Machine Learning*. pp. 2286–2296. PMLR (2021)
14. Graham, B., Engelcke, M., van der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. *CVPR* (2018)
15. Graham, B., van der Maaten, L.: Submanifold sparse convolutional networks. *arXiv preprint arXiv:1706.01307* (2017)
16. Guo, J., Tang, Y., Han, K., Chen, X., Wu, H., Xu, C., Xu, C., Wang, Y.: Hire-mlp: Vision mlp via hierarchical rearrangement. *arXiv preprint arXiv:2108.13341* (2021)

17. Guo, M.H., Cai, J.X., Liu, Z.N., Mu, T.J., Martin, R.R., Hu, S.M.: Pct: Point cloud transformer. *Computational Visual Media* **7**(2), 187–199 (2021)
18. Guo, Y., Wang, H., Hu, Q., Liu, H., Liu, L., Bennamoun, M.: Deep learning for 3d point clouds: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)* (2020)
19. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. *arXiv preprint arXiv:2103.00112* (2021)
20. Handa, A., Whelan, T., McDonald, J., Davison, A.J.: A benchmark for rgb-d visual odometry, 3d reconstruction and slam. In: *IEEE Int. Conf. on Robotics and Automation. (ICRA)* (2014)
21. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
22. Hendrycks, D., Gimpel, K.: Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415* (2016)
23. Hou, Q., Jiang, Z., Yuan, L., Cheng, M.M., Yan, S., Feng, J.: Vision permutator: A permutable mlp-like architecture for visual recognition. *arXiv preprint arXiv:2106.12368* (2021)
24. Hu, H., Zhang, Z., Xie, Z., Lin, S.: Local relation networks for image recognition. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 3464–3473 (2019)
25. Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: Randla-net: Efficient semantic segmentation of large-scale point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 11108–11117 (2020)
26. Huang, Q., Wang, W., Neumann, U.: Recurrent slice networks for 3d segmentation of point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
27. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanæs, H.: Large scale multi-view stereopsis evaluation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2014)
28. Jiang, L., Zhao, H., Liu, S., Shen, X., Fu, C.W., Jia, J.: Hierarchical point-edge interaction network for point cloud semantic segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 10433–10441 (2019)
29. Johnson, J., Douze, M., Jégou, H.: Billion-scale similarity search with gpus. *IEEE Transactions on Big Data* (2019)
30. Khan, S., Naseer, M., Hayat, M., Zamir, S.W., Khan, F.S., Shah, M.: Transformers in vision: A survey. *ACM Computing Surveys (CSUR)* (2021)
31. Klokov, R., Lempitsky, V.: Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2017)
32. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)* (2017)
33. Komarichev, A., Zhong, Z., Hua, J.: A-cnn: Annularly convolutional neural networks on point clouds. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7421–7430 (2019)
34. Lee, J., Choy, C., Park, J.: Putting 3d spatially sparse networks on a diet. *arXiv preprint arXiv:2112.01316* (2021)
35. Li, J., Hassani, A., Walton, S., Shi, H.: Convmlp: Hierarchical convolutional mlps for vision. *arXiv preprint arXiv:2109.04454* (2021)

36. Li, J., Chen, B.M., Lee, G.H.: So-net: Self-organizing network for point cloud analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2018)
37. Li, R., Li, X., Heng, P.A., Fu, C.W.: Point cloud upsampling via disentangled refinement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2021)
38. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems* **31** (2018)
39. Lian, D., Yu, Z., Sun, X., Gao, S.: AS-MLP: An axial shifted MLP architecture for vision. In: International Conference on Learning Representations (2022)
40. Lin, Y., Yan, Z., Huang, H., Du, D., Liu, L., Cui, S., Han, X.: Fpconv: Learning local flattening for point convolution. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4293–4302 (2020)
41. Liu, H., Dai, Z., So, D., Le, Q.V.: Pay attention to MLPs. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems (NeurIPS)* (2021)
42. Liu, Y., Fan, B., Meng, G., Lu, J., Xiang, S., Pan, C.: Densepoint: Learning densely contextual representation for efficient point cloud processing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5239–5248 (2019)
43. Liu, Y., Fan, B., Xiang, S., Pan, C.: Relation-shape convolutional neural network for point cloud analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8895–8904 (2019)
44. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10012–10022 (2021)
45. Liu, Z., Tang, H., Lin, Y., Han, S.: Point-voxel cnn for efficient 3d deep learning. *Advances in Neural Information Processing Systems (NeurIPS)* **32** (2019)
46. Long, J., Shelhamer, E., Darrell, T.: Fully convolutional networks for semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3431–3440 (2015)
47. Ma, X., Qin, C., You, H., Ran, H., Fu, Y.: Rethinking network design and local geometry in point cloud: A simple residual MLP framework. In: International Conference on Learning Representations (ICLR) (2022)
48. Mao, J., Wang, X., Li, H.: Interpolated convolutional networks for 3d point cloud understanding. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1578–1587 (2019)
49. Mao, J., Xue, Y., Niu, M., Bai, H., Feng, J., Liang, X., Xu, H., Xu, C.: Voxel transformer for 3d object detection. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 3164–3173 (2021)
50. Maturana, D., Scherer, S.: Voxnet: A 3d convolutional neural network for real-time object recognition. In: *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 922–928. IEEE (2015)
51. Mazur, K., Lempitsky, V.: Cloud transformers: A universal approach to point cloud processing tasks. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 10715–10724 (October 2021)
52. Muja, M., Lowe, D.G.: Fast approximate nearest neighbors with automatic algorithm configuration. *VISAPP (1)* **2**(331-340), 2 (2009)

53. Noh, H., Hong, S., Han, B.: Learning deconvolution network for semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 1520–1528 (2015)
54. Park, C., Jeong, Y., Cho, M., Park, J.: Fast point transformer. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 16949–16958 (2022)
55. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 652–660 (2017)
56. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. arXiv preprint arXiv:1706.02413 (2017)
57. Qi, X., Liao, R., Jia, J., Fidler, S., Urtasun, R.: 3d graph neural networks for rgbd semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 5199–5208 (2017)
58. Qian, G., Hammoud, H., Li, G., Thabet, A., Ghanem, B.: Assanet: An anisotropic separable set abstraction for efficient point cloud representation learning. *Advances in Neural Information Processing Systems (NeurIPS)* **34** (2021)
59. Radford, A., Narasimhan, K., Salimans, T., Sutskever, I.: Improving language understanding by generative pre-training (2018)
60. Ramachandran, P., Parmar, N., Vaswani, A., Bello, I., Levskaya, A., Shlens, J.: Stand-alone self-attention in vision models. In: Proceedings of the 33rd International Conference on Neural Information Processing Systems. pp. 68–80 (2019)
61. Simonovsky, M., Komodakis, N.: Dynamic edge-conditioned filters in convolutional neural networks on graphs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3693–3702 (2017)
62. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
63. Su, H., Jampani, V., Sun, D., Maji, S., Kalogerakis, E., Yang, M.H., Kautz, J.: Splatnet: Sparse lattice networks for point cloud processing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2530–2539 (2018)
64. Tang, H., Liu, Z., Zhao, S., Lin, Y., Lin, J., Wang, H., Han, S.: Searching efficient 3d architectures with sparse point-voxel convolution. In: European Conference on Computer Vision (ECCV). pp. 685–702. Springer (2020)
65. Tay, Y., Bahri, D., Metzler, D., Juan, D., Zhao, Z., Zheng, C.: Synthesizer: Rethinking self-attention in transformer models. In: ICML (2021)
66. Tay, Y., Deghani, M., Bahri, D., Metzler, D.: Efficient transformers: A survey. arXiv preprint arXiv:2009.06732 (2020)
67. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 6411–6420 (2019)
68. Tolstikhin, I.O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al.: Mlp-mixer: An all-mlp architecture for vision. *Advances in Neural Information Processing Systems (NeurIPS)* **34** (2021)
69. Touvron, H., Bojanowski, P., Caron, M., Cord, M., El-Nouby, A., Grave, E., Izacard, G., Joulin, A., Synnaeve, G., Verbeek, J., et al.: Resmlp: Feedforward networks for image classification with data-efficient training. arXiv preprint arXiv:2105.03404 (2021)

70. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning. pp. 10347–10357. PMLR (2021)
71. Trockman, A., Kolter, J.Z.: Patches are all you need? arXiv preprint arXiv:2201.09792 (2022)
72. Vaswani, A., Ramachandran, P., Srinivas, A., Parmar, N., Hechtman, B., Shlens, J.: Scaling local self-attention for parameter efficient visual backbones. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 12894–12904 (2021)
73. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: Advances in Neural Information Processing Systems (NeurIPS) (2017)
74. Wang, C., Samari, B., Siddiqi, K.: Local spectral graph convolution for point set feature learning. In: Proceedings of the European conference on computer vision (ECCV). pp. 52–66 (2018)
75. Wang, L., Huang, Y., Hou, Y., Zhang, S., Shan, J.: Graph attention convolution for point cloud semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10296–10305 (2019)
76. Wang, S., Suo, S., Ma, W.C., Pokrovsky, A., Urtasun, R.: Deep parametric continuous convolutional neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2589–2597 (2018)
77. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 568–578 (2021)
78. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. *ACM Transactions of Graphics (ToG)* **38**(5), 1–12 (2019)
79. Wu, W., Qi, Z., Fuxin, L.: Pointconv: Deep convolutional networks on 3d point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9621–9630 (2019)
80. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1912–1920 (2015)
81. Xie, E., Wang, W., Yu, Z., Anandkumar, A., Alvarez, J.M., Luo, P.: Segformer: Simple and efficient design for semantic segmentation with transformers. *Advances in Neural Information Processing Systems (NeurIPS)* **34** (2021)
82. Xu, M., Zhou, Z., Qiao, Y.: Geometry sharing network for 3d point cloud classification and segmentation. In: Association for the Advancement of Artificial Intelligence (AAAI) (2020)
83. Xu, M., Ding, R., Zhao, H., Qi, X.: Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3173–3182 (2021)
84. Xu, Q., Sun, X., Wu, C.Y., Wang, P., Neumann, U.: Grid-gcn for fast and scalable point cloud learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5661–5670 (2020)
85. Yang, Z., Sun, Y., Liu, S., Qi, X., Jia, J.: Cn: Channel normalization for point cloud recognition. In: European Conference on Computer Vision (ECCV). pp. 600–616. Springer (2020)

86. Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: European Conference on Computer Vision (ECCV) (2018)
87. Yi, L., Kim, V.G., Ceylan, D., Shen, I.C., Yan, M., Su, H., Lu, C., Huang, Q., Sheffer, A., Guibas, L.: A scalable active framework for region annotation in 3d shape collections. *ACM Transactions on Graphics (ToG)* **35**(6), 1–12 (2016)
88. Yu, T., Li, X., Cai, Y., Sun, M., Li, P.: S²-mlp: Spatial-shift mlp architecture for vision. arXiv preprint arXiv:2106.07477 (2021)
89. Yu, T., Li, X., Cai, Y., Sun, M., Li, P.: S²-mlpv2: Improved spatial-shift mlp architecture for vision. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV) (2022)
90. Yu, W., Luo, M., Zhou, P., Si, C., Zhou, Y., Wang, X., Feng, J., Yan, S.: Metaformer is actually what you need for vision. In: CVPR (2022)
91. Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.H., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 558–567 (2021)
92. Zhang, D.J., Li, K., Chen, Y., Wang, Y., Chandra, S., Qiao, Y., Liu, L., Shou, M.Z.: Morphmlp: A self-attention free, mlp-like backbone for image and video. arXiv preprint arXiv:2111.12527 (2021)
93. Zhang, F., Fang, J., Wah, B., Torr, P.: Deep fusionnet for point cloud semantic segmentation. In: European Conference on Computer Vision (ECCV). pp. 644–663. Springer (2020)
94. Zhao, H., Jia, J., Koltun, V.: Exploring self-attention for image recognition. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10076–10085 (2020)
95. Zhao, H., Jiang, L., Fu, C.W., Jia, J.: Pointweb: Enhancing local neighborhood features for point cloud processing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5565–5573 (2019)
96. Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 16259–16268 (2021)
97. Zhou, Q.Y., Park, J., Koltun, V.: Open3d: A modern library for 3d data processing. arXiv preprint arXiv:1801.09847 (2018)
98. Zhou, Y., Tuzel, O.: Voxelnet: End-to-end learning for point cloud based 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4490–4499 (2018)