
Generalizing Convolution to Point Clouds

Davide Bacciu¹ Francesco Landolfi¹

Abstract

Convolution, a fundamental operation in deep learning for structured grid data like images, cannot be directly applied to point clouds due to their irregular and unordered nature. Many approaches in literature that perform convolution on point clouds achieve this by designing a convolutional operator from scratch, often with little resemblance to the one used on images. We present two point cloud convolutions that naturally follow from the convolution in its standard definition popular with images. We do so by relaxing the indexing of the kernel weights with a “soft” dictionary that resembles the attention mechanism of the transformers. Finally, experimental results demonstrate the effectiveness of the proposed relaxations on two benchmark point cloud classification tasks.

1. Introduction

Convolutional neural networks (CNNs) [1] have revolutionized numerous computer vision tasks due to their ability to extract hierarchical features from structured grid data like images. However, applying CNNs directly to 3D point clouds, which represent objects as unordered sets of points, proves challenging. Unlike images, point clouds lack inherent structure and ordering, making standard convolution operations inapplicable. This hurdle has motivated researchers to develop specialized convolution approaches for point cloud data. The pioneering work of Qi, Su, Mo, *et al.* [2] paved the way for deep learning on point clouds. Their approach bypassed the need for traditional convolutions altogether, proposing instead an invariant operator based on set abstraction. Many follow-up works, such as Pointnet++ [3] and DGCNN [4], also proposed different convolution operators on point clouds that share little with the one performed on images or voxel grids. In this paper, we

¹Department of Computer Science, Università di Pisa, Pisa, Italy. Correspondence to: Francesco Landolfi <francesco.landolfi@phd.unipi.it>.

Published at the 2nd Differentiable Almost Everything Workshop at the 41st International Conference on Machine Learning, Vienna, Austria. July 2024. Copyright 2024 by the author(s).

derive instead two relaxations—continuously differentiable approximations—of the convolution operation on grid data that is applicable to unordered sets of points. We achieve this by substituting the indexing of the kernel weights performed by the convolution with a “soft” dictionary that resembles the attention mechanism of the transformers [5]. We test both relaxations and their depthwise variants on two standard point cloud classification datasets, on which they obtain comparable results with respect to common point cloud convolution baselines from literature.

2. Generalizing Convolution to Point Clouds

Given the tensors $\mathbf{X} \in \mathbb{R}^{b_1 \times \dots \times b_d \times f}$, $\mathbf{Z} \in \mathbb{R}^{(b_1 - c_1 + 1) \times \dots \times (b_d - c_d + 1) \times g}$, and $\mathbf{K} \in \mathbb{R}^{c_1 \times \dots \times c_d \times f \times g}$, a d -dimensional multi-channel *convolution* [1] can be defined as¹

$$\begin{aligned} \mathbf{Z}_{i_1, \dots, i_d, k} &= [\text{conv}(\mathbf{X}, \mathbf{K})]_{i_1, \dots, i_d, k} \\ &= \sum_{j_1, \dots, j_d, h} \mathbf{X}_{i_1 + j_1, \dots, i_d + j_d, h} \cdot \mathbf{K}_{j_1, \dots, j_d, h, k}, \end{aligned} \quad (1)$$

where \mathbf{X} , \mathbf{Z} , and \mathbf{K} , represent, respectively, the *input*, the *output*, and the *kernel* tensors. Using a vectorial notation of the indices (all tensor indices start from 0), where $\mathbf{X}_{\mathbf{i}} = \mathbf{X}_{i_1, \dots, i_d}$ with $\mathbf{i} = [i_1, \dots, i_d]^T$, we can rewrite Equation (1) as

$$\mathbf{Z}_{\mathbf{i}} = [\text{conv}(\mathbf{X}, \mathbf{K})]_{\mathbf{i}} = \sum_{\mathbf{j} \in C} \mathbf{X}_{\mathbf{i} + \mathbf{j}} \cdot \mathbf{K}_{\mathbf{j}}, \quad (2)$$

where $C = [c_1] \times \dots \times [c_d]$, with $[c] = \{0, 1, \dots, c - 1\}$. Notice that $\mathbf{Z}_{\mathbf{i}} \in \mathbb{R}^g$, $\mathbf{X}_{\mathbf{i} + \mathbf{j}} \in \mathbb{R}^f$, and $\mathbf{K}_{\mathbf{j}} \in \mathbb{R}^{f \times g}$, for all $\mathbf{j} \in C$. Another way to see Equation (2) is by obtaining the output entries by means of the neighbouring ones, that is,

$$\mathbf{Z}_{\mathbf{i}} = [\text{conv}(\mathbf{X}, \mathbf{K})]_{\mathbf{i}} = \sum_{\mathbf{h} \in N(\mathbf{i})} \mathbf{X}_{\mathbf{h}} \cdot \mathbf{K}_{\mathbf{h} - \mathbf{i}}, \quad (3)$$

where $N(\mathbf{i}) = \{\mathbf{i} + \mathbf{j} \mid \mathbf{j} \in C\}$ denotes the set of neighbouring indices of \mathbf{i} .

From Equations (2) and (3) we can derive two relaxations, that allow us to generalize convolution to (possibly irregular, non-grid-like) point clouds.

¹Equation (1) should be called “cross-correlation” instead, but in a learning setting it makes no difference whether the kernel is flipped or not.

First relaxation. Equation (2) can be also represented as a *point cloud convolution*, as follows. Let $V = [n]$ with $n = \prod_{i=1}^d b_i$ and assume w.l.o.g. that the convolution is “centered”, that is, that the convolution size is odd and uniform across all the dimensions, i.e., $C = \{-c, -c+1, \dots, 0, \dots, c-1, c\}^d$ for some constant $c \in \mathbb{N}$, and that the input is properly zero-padded (all these assumption can be replaced by a using a suitable metric and a proper transformation of the indices). Instead of using vector indices, we can vectorize (or “flatten”) the tensors up to their d -th dimensions, obtaining $\mathbf{X} \in \mathbb{R}^{n \times f}$, $\mathbf{Z} \in \mathbb{R}^{n \times g}$, and $\mathbf{K} \in \mathbb{R}^{|C| \times f \times g}$, and store instead the d -dimensional vector indices in the auxiliary matrix $\mathbf{P} \in \mathbb{R}^{n \times d}$, such that $\mathbf{X}_{\mathbf{P}_v} = \mathbf{X}_v$ for all $v \in V$. This is a common scenario in point cloud tasks, where \mathbf{P} is used to store the geometric information of the points, also referred as *coordinate matrix*, while \mathbf{X} contains instead other attributes of the points and is generally referred as *feature matrix*. In our setting, similarly to the coordinate matrix, we also define the *offset matrix* $\mathbf{C} \in \mathbb{R}^{|C| \times d}$ containing all the elements of C . We can now reformulate Equation (2) as follows,

$$\mathbf{Z}_v = \sum_{i=0}^{|C|-1} \mathbf{X}_{u^*} \mathbf{K}_i \quad \text{s.t.} \quad \mathbf{P}_{u^*} = \mathbf{P}_v + \mathbf{C}_i. \quad (4)$$

We still cannot apply straightforwardly Equation (4) to a generic point cloud, since $\mathbf{P}_u + \mathbf{C}_i$ may not correspond to any point in \mathbf{P} . Hence, we can look for the nearest point available, as follows

$$\begin{aligned} \mathbf{Z}_v &= \sum_{i=0}^{|C|-1} \mathbf{X}_{u^*} \mathbf{K}_i \\ \text{s.t.} \quad u^* &= \arg \min_u \delta(\mathbf{P}_v + \mathbf{C}_i, \mathbf{P}_u), \end{aligned} \quad (5)$$

where $\delta : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ is any distance function, such as $\delta(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|$. We can further relax Equation (5) by making it smoothly differentiable by substituting the *argmin* with a *softmax*, as follows,

$$\mathbf{Z}_v = \sum_{i=0}^{|C|-1} \sum_{u \in V} \frac{e^{-\delta(\mathbf{P}_v + \mathbf{C}_i, \mathbf{P}_u)}}{\sum_w e^{-\delta(\mathbf{P}_v + \mathbf{C}_i, \mathbf{P}_w)}} \mathbf{X}_u \mathbf{K}_i. \quad (6)$$

In this way, $\mathbf{C} \in \mathbb{R}^{|C| \times n}$ can also be parametrical, allowing the offsets to be learned adaptively from data. With this last step, we completely replaced the “indexing” performed by convolutions in Equation (2) with an attention mechanism [5], where the *keys* are formed by the point coordinates \mathbf{P} , the *values* are the point features \mathbf{X} , and the *queries* are instead the neighboring points defined by $\mathbf{P}_v + \mathbf{C}_i$. To avoid computing an attention matrix of size $O(n^2)$, which prevents this formulation to scale to large point clouds, we can “mask” distant points and compute the *softmax* only on

the k nearest ones, as follows,

$$\mathbf{Z}_v = \sum_{i=0}^{|C|-1} \sum_{u \in N[v]} \frac{e^{-\delta(\mathbf{P}_v + \mathbf{C}_i, \mathbf{P}_u)}}{\sum_w e^{-\delta(\mathbf{P}_v + \mathbf{C}_i, \mathbf{P}_w)}} \mathbf{X}_u \mathbf{K}_i, \quad (7)$$

where $N[v] = k\text{-NN}(v, \mathbf{P})$ is the k -nearest neighbors of v in \mathbf{P} , containing v itself.

Second relaxation. We can derive a similar relaxation by starting instead from Equation (3), as follows. First, we reformulate the definition of neighboring nodes, by adopting instead a closed ball around the “central” node,

$$\mathbf{Z}_v = \sum_{u \in B_c[v]} \mathbf{X}_u \mathbf{K}_{i^*} \quad \text{s.t.} \quad \mathbf{C}_{i^*} = \mathbf{P}_u - \mathbf{P}_v, \quad (8)$$

where $B_c[v] = \{u \in \mathbb{N} \mid \|\mathbf{P}_u - \mathbf{P}_v\|_\infty \leq c\}$. Notice that we adopted the infinity norm inside the ball to mimic the “(hyper-)cubic” shape of the convolution’s receptive field, but we could also use any other form of neighbourhood such as a 2-norm ball or the k -nearest neighbors. Notice also that, again, we still cannot apply straightforwardly Equation (8) to a generic point cloud, since the offset $\mathbf{P}_u - \mathbf{P}_v$ may not lie in C . We can then apply a first relaxation to Equation (8) by assigning every neighbor displacement to the nearest available offset, as follows

$$\begin{aligned} \mathbf{Z}_v &= \sum_{u \in B_c[v]} \mathbf{X}_u \mathbf{K}_{i^*} \\ \text{s.t.} \quad i^* &= \arg \min_i \delta(\mathbf{P}_u - \mathbf{P}_v, \mathbf{C}_i). \end{aligned} \quad (9)$$

We can further relax Equation (9) by making it smoothly differentiable, as follows,

$$\mathbf{Z}_v = \sum_{u \in B_c[v]} \sum_{i=0}^{|C|-1} \frac{e^{-\delta(\mathbf{P}_u - \mathbf{P}_v, \mathbf{C}_i)}}{\sum_j e^{-\delta(\mathbf{P}_u - \mathbf{P}_v, \mathbf{C}_j)}} \mathbf{X}_u \mathbf{K}_i. \quad (10)$$

In this way, we allow again $\mathbf{C} \in \mathbb{R}^{|C| \times n}$ to be parametrical, so that the offsets are learned adaptively from data. In this case, differently from the previous relaxation, the “indexing” performed by convolutions in Equation (3) is replaced by an attention mechanism, where the *keys* are formed by the offsets in \mathbf{C} , the *values* are the filters \mathbf{K} , and the *queries* are instead the neighbors’ offsets $\mathbf{P}_u - \mathbf{P}_v$.

Depthwise convolution. The two relaxation only affect the way the kernels are averaged together but not do affect instead how they are applied to the feature matrix, i.e., the $\mathbf{X}_u \mathbf{K}_i$ part remains untouched throughout the relaxations. This allows us to generalize the two relaxations also to other kinds of operations, such as the *depthwise convolution* [6]. In the depthwise convolution, defined as

$$\begin{aligned} \mathbf{z}_{i_1, \dots, i_d, k} &= [\text{dw-conv}(\mathbf{X}, \mathbf{K})]_{i_1, \dots, i_d, k} \\ &= \sum_{j_1, \dots, j_d} \mathbf{x}_{i_1 + j_1, \dots, i_d + j_d, k} \cdot \mathbf{K}_{j_1, \dots, j_d, k}, \end{aligned} \quad (11)$$

the input and output tensors have the same number of features f , and the kernel tensor has shape $\mathbf{K} \in \mathbb{R}^{c_1 \times \dots \times c_d \times f}$. In the point cloud setting, we have $\mathbf{K} \in \mathbb{R}^{|C| \times f}$ and, if we go over the first relaxation steps again, we obtain

$$\mathbf{z}_v = \sum_{i=0}^{|C|-1} \sum_{u \in N[v]} \frac{e^{-\delta(\mathbf{P}_v + \mathbf{C}_i, \mathbf{P}_u)}}{\sum_w e^{-\delta(\mathbf{P}_v + \mathbf{C}_i, \mathbf{P}_w)}} \mathbf{X}_u \circ \mathbf{K}_i, \quad (12)$$

where \circ denotes the element-wise product. A similar formulation can be obtained also for the second relaxation.

3. Related Works

The seminal work of PointNet by Qi, Su, Mo, *et al.* [2] introduced a novel approach to directly process point cloud data without the need for voxelization or other preprocessing steps. PointNet’s architecture leverages symmetric functions to ensure permutation invariance, making it highly effective for classification and segmentation tasks. Building upon PointNet, PointNet++ [3] improved the ability to capture local structures by using a hierarchical approach, combining local and global features to better handle varying point densities. Graph-based methods have also gained traction for point cloud processing. Wang, Sun, Liu, *et al.* [4] proposed Dynamic Graph CNN (DGCNN), which dynamically constructs graphs in each layer of the network to capture local geometric relationships.

To bridge the gap between traditional convolutional networks and point cloud processing, researchers have developed convolution-like operations for point clouds that closely resemble our second relaxation. In Point-wise CNN [7], the authors proposed a convolution operator that binned the neighboring points into kernel cells having different kernel weights. FeaSTNet (Feature-Steered Graph Convolution) [8] generalized the convolution by averaging the kernel weights through a (soft-maxed) linear transformation of the local feature vectors. Atzmon, Maron, and Lipman [9] proposed instead a point convolution based on extension and restriction operators, mapping point functions to continuous volumetric functions which are then convolved and restricted back to the point space. Wu, Qi, and Fuxin [10] introduced PointConv, which extends convolution operations to irregular point sets by defining density-aware convolution kernels. Specifically, the authors average the kernel parameters using a kernel density estimation. Similarly, Thomas, Qi, Deschaud, *et al.* [11] proposed KPConv (Kernel Point Convolution), which defines convolutional operations directly in the Euclidean space, averaging the kernels via a linear correlation. In a similar fashion, PACConv (Position Adaptive Convolution) [12] holds the kernel weights in a specialized matrix (which they called it “weight bank”) and assigns a kernel to every neighbor point by means of a neural network (called “score net”).

Driven by the success of transformer architectures in natural language processing, the use of the attention mechanism can also be found in the point cloud processing literature. Notable examples are the graph attention network [13], [14] and the point transformer, in all of its variants [15]–[18].

On a final note, the idea of trainable offsets, even if applied to images, can be dated back to deformable convolution [19]. In their paper, the authors proposed a convolution with variable spacings, similar to our first relaxation, but where the features were aggregated via a bilinear interpolation kernel. A similar work was proposed also by Hassani, Pellegrini, and Masquelier [20].

4. Experiments

We tested both relaxations on two standard point clouds classification tasks from literature, ModelNet40 [21] and ScanObjectNN [22]. In the experiments, we used DGCNN [4] as a backbone classifier, where every of its dynamic edge convolution layer was substituted by a relaxed convolution having the same number of channels. The training was performed following the SimpleView [23] protocol: random translation and scaling of the training samples, cross-entropy loss with label smoothing (0.2), and fixed-number of points per sample (1024, with no resampling). As in SimpleView, the final models were obtained after a pre-fixed number of epochs (60), with no feed-back from the test set. To speed up convergence, we used cosine annealing with a warm restart after 20 epochs. In both relaxations we used a k -nearest neighbors with $k = 8$ and we set the number of offsets to 8. These settings were chosen after a preliminary evaluation on a validation split (10% of the training dataset). In Table 1 we report the average accuracy and balanced accuracy obtained by our relaxations on 10 train/test runs for every model. The relaxations were also tested in their depthwise variant (DW), and in the “inverted residual” variant (Mobile), as proposed in MobileNetV2 [6], where the depthwise convolution is preceded by an expansion layer ($6 \times$ channels) and followed by a reduction one. We also restate the accuracies of four common baselines (SimpleView, PointNet/++, DGCNN) as reported in the original paper of the SimpleView training protocol [23]. Missing values were not reported in the paper.

We can see from the results in Table 1 that our relaxations produced results comparable to the given baselines. Our relaxation is superior to the lowest performing one (PointNet) almost everywhere, while obtaining a similar result to the backbone model (DGCNN) only when using the first relaxation with the inverted residual setting. Depthwise convolution alone seem to worsen the performance of both relaxations, while it obtained the best results in the mobile setting when combined with the first relaxation. Overall, most baselines obtained better results with respect to our re-

	ModelNet40		ScanObjectNN		Param.
	Accuracy	Bal. Acc.	Accuracy	Bal. Acc.	
SimpleView	93.0 ± 0.4	90.5 ± 0.8	79.5 ± 0.5	— ± —	0.8M
PointNet	89.2 ± 0.9	85.1 ± 0.6	68.2 ± —	— ± —	3.5M
PointNet++	92.7 ± 0.1	90.0 ± 0.3	77.9 ± —	— ± —	1.7M
DGCNN	91.9 ± 0.3	89.1 ± 0.3	78.1 ± —	— ± —	1.8M
First Relaxation	91.40 ± 0.31	88.15 ± 0.58	72.18 ± 0.81	66.62 ± 1.04	2.1M
First Rel. (DW)	91.32 ± 0.25	87.57 ± 0.43	71.93 ± 0.75	66.48 ± 1.20	1.7M
First Rel. (Mobile)	91.77 ± 0.25	88.58 ± 0.48	74.13 ± 0.92	69.28 ± 1.10	2.2M
Second Relaxation	90.82 ± 0.20	86.53 ± 0.52	67.41 ± 1.11	63.29 ± 1.31	2.1M
Second Rel. (DW)	90.36 ± 0.35	85.51 ± 0.54	67.92 ± 0.51	64.08 ± 0.91	1.7M
Second Rel. (Mobile)	90.18 ± 0.28	84.85 ± 0.51	65.16 ± 0.76	57.18 ± 0.98	2.2M

Table 1. Classification accuracy, balanced accuracy (mean ± std) and number of parameters of each model.

laxations, which may indicate that their informed inductive bias plays favorably in point clouds tasks compared to one of the general classical convolution.

5. Conclusions

This paper proposes two significant relaxations of the standard convolution operation, specifically tailored for unordered point cloud data. These relaxations address the inherent limitations of classical convolutions, which struggle with the irregular structure of point clouds. The core idea lies in substituting the rigid indexing of kernel weights with a soft attention mechanism. This enables the convolutions to adaptively focus on relevant neighboring points or offsets between points, leading to more robust feature extraction. The experimental evaluation shows that the proposed relaxations achieve comparable performance to existing point cloud convolution methods. This indicates that the relaxations offer a viable alternative framework for applying convolutions to unordered data, while potentially opening doors for further exploration of generalizable convolution operations on such data structures.

References

- [1] I. Goodfellow, Y. Bengio, and A. Courville, *Deep Learning*. MIT Press, 2016. [Online]. Available: <http://www.deeplearningbook.org>.
- [2] C. R. Qi, H. Su, K. Mo, and L. J. Guibas, “PointNet: Deep learning on point sets for 3d classification and segmentation,” presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2017, pp. 652–660. [Online]. Available: https://openaccess.thecvf.com/content_cvpr_2017/html/Qi_PointNet_Deep_Learning_CVPR_2017_paper.html (visited on 10/26/2020).
- [3] C. R. Qi, L. Yi, H. Su, and L. J. Guibas, “PointNet++: Deep hierarchical feature learning on point sets in a metric space,” in *Advances in Neural Information Processing Systems 30*, I. Guyon, U. V. Luxburg, S. Bengio, et al., Eds., Curran Associates, Inc., 2017, pp. 5099–5108. [Online]. Available: <http://papers.nips.cc/paper/7095-pointnet-deep-hierarchical-feature-learning-on-point-sets-in-a-metric-space.pdf> (visited on 10/16/2020).
- [4] Y. Wang, Y. Sun, Z. Liu, S. E. Sarma, M. M. Bronstein, and J. M. Solomon, “Dynamic graph CNN for learning on point clouds,” *ACM Transactions on Graphics*, vol. 38, no. 5, pp. 146:1–146:12, Oct. 10, 2019, ISSN: 0730-0301. DOI: 10.1145/3326362. [Online]. Available: <https://doi.org/10.1145/3326362> (visited on 11/07/2020).
- [5] A. Vaswani, N. Shazeer, N. Parmar, et al., “Attention is All you Need,” in *Advances in Neural Information Processing Systems*, vol. 30, Curran Associates, Inc., 2017. (visited on 07/12/2023).
- [6] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L.-C. Chen, “MobileNetV2: Inverted residuals and linear bottlenecks,” presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 4510–4520. [Online]. Available: https://openaccess.thecvf.com/content_cvpr_2018/html/Sandler_MobileNetV2_Inverted_Residuals_CVPR_2018_paper.html (visited on 05/03/2024).
- [7] B.-S. Hua, M.-K. Tran, and S.-K. Yeung, “Pointwise convolutional neural networks,” presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 984–993. [Online]. Available: https://openaccess.thecvf.com/content_cvpr_2018/html/Hua_Pointwise_Convolutional_Neural_CVPR_2018_paper.html (visited on 06/14/2023).
- [8] N. Verma, E. Boyer, and J. Verbeek, “FeaStNet: Feature-steered graph convolutions for 3d shape analysis,” presented at the Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2018, pp. 2598–2606. [Online]. Available: https://openaccess.thecvf.com/content_cvpr_2018/html/Verma_FeaStNet_Feature-Steered_Graph_CVPR_2018_paper.html (visited on 05/10/2024).

- [9] M. Atzmon, H. Maron, and Y. Lipman, "Point convolutional neural networks by extension operators," *ACM Transactions on Graphics*, vol. 37, no. 4, 71:1–71:12, Jul. 30, 2018, ISSN: 0730-0301. DOI: 10.1145/3197517.3201301. [Online]. Available: <https://dl.acm.org/doi/10.1145/3197517.3201301> (visited on 05/10/2024).
- [10] W. Wu, Z. Qi, and L. Fuxin, "PointConv: Deep convolutional networks on 3d point clouds," presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2019, pp. 9621–9630. [Online]. Available: https://openaccess.thecvf.com/content_CVPR_2019/html/Wu_PointConv_Deep_Convolutional_Networks_on_3D_Point_Clouds_CVPR_2019_paper.html (visited on 11/07/2020).
- [11] H. Thomas, C. R. Qi, J.-E. Deschaud, B. Marcotegui, F. Goulette, and L. J. Guibas, "KPConv: Flexible and deformable convolution for point clouds," presented at the Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 6411–6420. [Online]. Available: https://openaccess.thecvf.com/content_ICCV_2019/html/Thomas_KPConv_Flexible_and_Deformable_Convolution_for_Point_Clouds_ICCV_2019_paper.html (visited on 07/06/2023).
- [12] M. Xu, R. Ding, H. Zhao, and X. Qi, "PAConv: Position adaptive convolution with dynamic kernel assembling on point clouds," presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 3173–3182. [Online]. Available: https://openaccess.thecvf.com/content/CVPR2021/html/Xu_PAConv_Position_Adaptive_Convolution_With_Dynamic_Kernel_Assembling_on_Point_CVPR_2021_paper.html (visited on 07/06/2023).
- [13] P. Veličković, G. Cucurull, A. Casanova, A. Romero, P. Liò, and Y. Bengio, "Graph attention networks," in *International Conference on Learning Representations*, Feb. 15, 2018. [Online]. Available: <https://openreview.net/forum?id=rJXmpikCZ> (visited on 01/31/2021).
- [14] L. Wang, Y. Huang, Y. Hou, S. Zhang, and J. Shan, "Graph attention convolution for point cloud semantic segmentation," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, ISSN: 2575-7075, Jun. 2019, pp. 10 288–10 297. DOI: 10.1109/CVPR.2019.01054.
- [15] X. Pan, Z. Xia, S. Song, L. E. Li, and G. Huang, "3d object detection with pointformer," presented at the Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 7463–7472. [Online]. Available: https://openaccess.thecvf.com/content/CVPR2021/html/Pan_3D_Object_Detection_With_Pointformer_CVPR_2021_paper.html (visited on 05/27/2024).
- [16] H. Zhao, L. Jiang, J. Jia, P. H. S. Torr, and V. Koltun, "Point transformer," presented at the Proceedings of the IEEE/CVF International Conference on Computer Vision, 2021, pp. 16 259–16 268. [Online]. Available: https://openaccess.thecvf.com/content/ICCV2021/html/Zhao_Point_Transformer_ICCV_2021_paper.html (visited on 05/27/2024).
- [17] X. Wu, Y. Lao, L. Jiang, X. Liu, and H. Zhao, "Point transformer v2: Grouped vector attention and partition-based pooling," *Advances in Neural Information Processing Systems*, vol. 35, pp. 33 330–33 342, Dec. 6, 2022. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2022/hash/d78ece6613953f46501b958b7bb4582f-Abstract-Conference.html (visited on 05/27/2024).
- [18] X. Wu, L. Jiang, P.-S. Wang, *et al.*, *Point transformer v3: Simpler, faster, stronger*, version: 2, Mar. 25, 2024. DOI: 10.48550/arXiv.2312.10035. arXiv: 2312.10035[cs]. [Online]. Available: <http://arxiv.org/abs/2312.10035> (visited on 05/27/2024).
- [19] J. Dai, H. Qi, Y. Xiong, *et al.*, "Deformable convolutional networks," presented at the Proceedings of the IEEE International Conference on Computer Vision, 2017, pp. 764–773. [Online]. Available: https://openaccess.thecvf.com/content_iccv_2017/html/Dai_Deformable_Convolutional_Networks_ICCV_2017_paper.html (visited on 05/08/2024).
- [20] I. K. Hassani, T. Pellegrini, and T. Masquelier, "Dilated convolution with learnable spacings," presented at the The Eleventh International Conference on Learning Representations, Sep. 29, 2022. [Online]. Available: <https://openreview.net/forum?id=Q3-1vRh3HOA> (visited on 08/04/2023).
- [21] Z. Wu, S. Song, A. Khosla, *et al.*, "3d ShapeNets: A deep representation for volumetric shapes," in *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, ISSN: 1063-6919, Jun. 2015, pp. 1912–1920. DOI: 10.1109/CVPR.2015.7298801. [Online]. Available: <https://ieeexplore.ieee.org/document/7298801> (visited on 05/05/2024).
- [22] M. A. Uy, Q.-H. Pham, B.-S. Hua, T. Nguyen, and S.-K. Yeung, "Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data," presented at the Proceedings of the IEEE/CVF International Conference on Computer Vision, 2019, pp. 1588–1597. [Online]. Available: https://openaccess.thecvf.com/content_ICCV_2019/html/Uy_Revisiting_Point_Cloud_Classification_A_New_Benchmark_Dataset_and_Classification_ICCV_2019_paper.html (visited on 10/24/2021).
- [23] A. Goyal, H. Law, B. Liu, A. Newell, and J. Deng, "Revisiting point cloud shape classification with a simple and effective baseline," in *Proceedings of the 38th International Conference on Machine Learning*, ISSN: 2640-3498, PMLR, Jul. 1, 2021, pp. 3809–3820. [Online]. Available: <https://proceedings.mlr.press/v139/goyal21a.html> (visited on 10/24/2021).