

Enhancing Local Feature Learning for 3D Point Cloud Processing using Unary-Pairwise Attention

Haoyi Xiu^{1,2}

xiu.h.aa@m.titech.ac.jp

Xin Liu^{2,†}

xin.liu@aist.go.jp

Weimin Wang^{3,†}

wangweimin@dlut.edu.cn

Kyoung-Sook Kim²

ks.kim@aist.go.jp

Takayuki Shinohara¹

shinohara.t.af@m.titech.ac.jp

Qiong Chang⁴

q.chang@c.titech.ac.jp

Masashi Matsuoka¹

matsuoka.m.ab@m.titech.ac.jp

¹ Department of Architecture and Building Engineering, Tokyo Institute of Technology, Tokyo, Japan

² Artificial Intelligence Research Center, AIST, Tokyo, Japan

³ DUT-RU International School of Information Science and Engineering, Dalian University of Technology, Dalian, China

⁴ Department of Computer Science, Tokyo Institute of Technology, Tokyo, Japan

Abstract

We present a simple but effective attention named the unary-pairwise attention (UPA) for modeling the relationship between 3D point clouds. Our idea is motivated by the analysis that the standard self-attention (SA) that operates globally tends to produce almost the same attention maps for different query positions, revealing difficulties for learning query-independent and query-dependent information jointly. Therefore, we reformulate the SA and propose query-independent (Unary) and query-dependent (Pairwise) components to facilitate the learning of both terms. In contrast to the SA, the UPA ensures query dependence via operating locally. Extensive experiments show that the UPA outperforms the SA consistently on various point cloud understanding tasks including shape classification, part segmentation, and scene segmentation. Moreover, simply equipping the popular PointNet++ method with the UPA even outperforms or is on par with the state-of-the-art attention-based approaches. In addition, the UPA systematically boosts the performance of both standard and modern networks when it is integrated into them as a compositional module.

1 Introduction

3D data has become increasingly available thanks to the advent of modern 3D sensors. The 3D point cloud is one of the simplest shape representations which is typically represented as

† Corresponding author(s)

© 2021. The copyright of this document resides with its authors.

It may be distributed unchanged freely in print or electronic forms.

spatially scattered 3D points. Recently, automatic understanding of 3D point clouds using deep learning [15] has attracted much interest in various applications such as autonomous driving [4, 23] and remote sensing [25, 45].

The irregular nature of 3D point clouds brings challenges on the deep learning-based point cloud analysis because popular methods like convolutional neural network (CNN) only work on the regularly structured data (e.g., 2D and 3D grids). Therefore, 3D point clouds are often projected to regular formats such as voxels [19, 44] and images [14, 26] to enable regular convolutions. Recently, PointNet [21] has triggered the development of methods that directly operate on point clouds [22, 27, 32]. The key to their success is rooted in the use of shared multi-layer perceptrons (MLPs) and symmetric functions (e.g., max-pooling and avg-pooling). Both types of operations ensure permutation invariance, making them a perfect fit for point cloud processing.

On the other hand, the success of the self-attention (SA) [8, 28] in natural language processing has triggered various applications of the SA to 2D vision problems (e.g., image recognition [9], generation [42], and object detection [51]). The SA updates the query features by aggregating features from other positions (keys) based on pairwise relationships. The SA is permutation invariant; therefore, it is directly applicable to 3D point clouds. Recent research shows that the SA can indeed benefit the point cloud processing [11, 16, 38]. However, it is found in the 2D vision domain that the SA often produces almost the same attention maps for very different query positions [3, 40]. Such a finding is crucial for exploring the direction of SA-based research. It is necessary to validate if such a problem exists in the point cloud processing. Meanwhile, another concern of the the SA is its quadratic dependency on the input cardinality [30], which limits its application to substantially down-sampled inputs [31].

In this paper, we first perform qualitative and quantitative analyses to show that the SA tends to attend some fixed positions regardless of different queries. In other words, the attention is biased towards the learning of the query-independent information while suppressing the learning of query-dependent information. Based on this observation, we propose the unary-pairwise attention (UPA) by reformulating the SA to exploit query-independent and query-dependent information simultaneously while ensuring query-dependence. Specifically, given a query point and its nearest neighbors, the unary attention produces attention scores using absolute features, which ensures query independence of produced scores. The pairwise attention, by contrast, calculates attention scores using relative features to encode query dependence. Both attentions are permutation invariant; thus, it is fitting for 3D point cloud processing. The graphical description of UPA is shown in Fig. 1. We show that the




Figure 1: **Overview of UPA.** Given a query point (the blue point), UPA updates the feature by combining the outputs from unary and pairwise attentions. \otimes represents matrix multiplication and \oplus denotes element-wise addition.

UPA outperforms the SA in various tasks through extensive experiments. Furthermore, the UPA brings systematical improvements for both standard and modern networks when it is integrated into networks as a compositional module.

The key contributions are summarized as follows:

- We conduct qualitative and quantitative analyses to show that the SA tends to attend fixed positions regardless of different queries.
- We propose a new form of attention, the unary-pairwise attention (UPA), to enhance local feature learning of 3D point clouds.
- We perform various experiments to demonstrate that the UPA consistently outperforms the SA across a range of tasks. Moreover, as a compositional module, the UPA systematically provides performance improvements for standard and modern networks.

2 Related Work

Deep learning on 3D point clouds. Owing to its unstructured nature, point clouds need to be projected onto regular grids to enable regular convolutions. Some methods convert point clouds into multi-view images [12, 26] while other methods voxelize 3D point clouds [8, 10, 13, 44]. The performance of image-based methods may heavily rely on the choice of projection planes whereas the memory costs of voxel-based methods grow cubically with the resolution. Besides, both types of methods lose fine-grained information due to projections. To overcome these issues, Qi *et al.* propose PointNet [21] that can directly operate on 3D point clouds. Subsequently, PointNet++ [22] is proposed to tackle the local structure by applying PointNets to local subsets of point clouds. Owing to its simplicity and effectiveness, PointNet++ becomes the key building block for recent studies [17, 18, 27, 32, 35, 36]. While their works focus on developing methods based on convolution, this study aims to develop a new operation based on the attention mechanism.

Self-attention. Transformers [28] has revolutionized natural language processing and inspired vision researchers to apply the SA in image processing tasks [2, 9, 9, 20, 31]. To further adapt the SA to specific applications, some works apply it locally [12, 24]; some works make it more expressive [5, 40]; some works improve efficiency [8, 41]. The idea of the SA has also been introduced in point cloud/set processing [16]. PAT [38] develops a parameter-efficient variant whereas PointASNL [67] uses the SA to augment convolution-based networks. Some other works apply channel-wise modulation [13, 29, 43] to exploit fine-grained details. By contrast, motivated by analyses on the SA, our approach aims to enhance the SA by explicitly modeling query-independent and query-dependent information simultaneously. Furthermore, unlike the SA that attends globally, our method operates locally to guarantee query dependence while being able to tackle the voluminous input.

3 Method

This section begins by analyzing the SA and discussing the observed problems. Then, we present the formulation of the proposed UPA, and UPA blocks for seamless integration to existing networks as a compositional module.

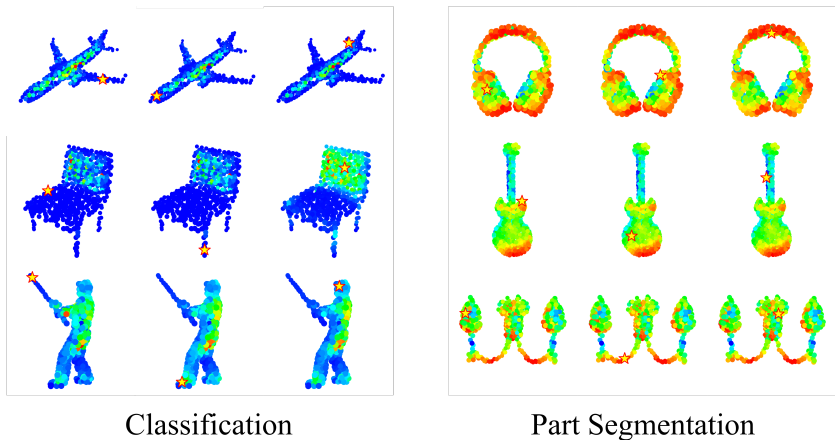


Figure 2: **Attention maps generated by SA layers.** Analyses are performed on ModelNet (classification) and ShapeNet (part segmentation) datasets. Stars indicate query positions. Different query positions produce similar attention maps.

	ModelNet			ShapeNet		
	OA	stage 1	stage 2	mIoU	stage 1	stage 2
PointNet++ [22] + SA	92.1	0.094	0.047	85.8	0.000	0.000
PointNet++ [22] + DNL [41]	92.1	0.162	0.147	85.8	0.000	0.000

Table 1: **Results of the quantitative analysis on SA and DNL layers.** Stage n denotes the inserting position of SA and DNL layers after the n th set abstraction level [22]. Scores for each stage represent $mJSD = \frac{1}{N^2 h} \sum_{i=1}^N \sum_{j=1}^N JSD(\text{AttMap}_i, \text{AttMap}_j)$, which measures the average similarity of attention maps over all query positions. N denotes the number of input points, h is the number of attention heads (8 for the classification and 1 for the segmentation), and AttMap_i and AttMap_j represent attention maps of points i and j .

3.1 Analysis of Self-Attention

Let $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^N$ denote the feature map of a point cloud where N is the total number of points and \mathbf{x}_i indicates a feature vector associated with a point. Then the SA can be defined as:

$$\mathbf{y}_i = \sum_{j=1}^N \text{softmax}_j(\mathbf{q}_i^\top \mathbf{k}_j) \mathbf{v}_j \quad (1)$$

where i and j index query and key elements, respectively. The query $\mathbf{q}_i = W_Q \mathbf{x}_i$, key $\mathbf{k}_j = W_K \mathbf{x}_j$, and value vectors $\mathbf{v}_j = W_V \mathbf{x}_j$ are linear transformations of the query point \mathbf{x}_i and the key point \mathbf{x}_j . W_Q , W_K and $W_V \in \mathbb{R}^{d_{in} \times d_{out}}$ are learned linear projections. $\mathbf{y}_i \in \mathbb{R}^{d_{out}}$ represents the output feature. d_{in} and d_{out} are input and output feature dimensions, respectively. $\text{softmax}_j(\cdot)$ is applied to normalize dot product outputs between the query and the corresponding keys. As a result, the output \mathbf{y}_i is a convex combination of value vectors.

While the SA is effective in point cloud recognition tasks, few studies investigate the behavior of attention maps generated by SA layers; thus, we provide qualitative and quantitative analyses of learned attention maps on classification (ModelNet40 [34]) and part segmentation (ShapeNet [39]). To provide an intuitive understanding of the behavior of the SA,

the qualitative analysis visualizes the attention maps of different query positions. Then, the quantitative analysis is performed to quantify the average similarity of all attention maps. We adopt point-averaged Jensen-Shannon Divergence (mJSD) as the similarity measure. To train SA layers, we adopt PointNet++ [22], which is the key building block for recent developments, as the backbone, and apply an SA layer after each set abstraction level [22]. The qualitative results are shown in Fig. 2. Unexpectedly, attention maps are similar to each other although the query positions are different. Furthermore, as shown in Table 1, mJSD scores in SA layers are generally small, indicating that SA layers are prone to learning the global structure of point clouds by prioritizing query-independent information. In addition, we train DNL [40] layers, in which the dot product is mathematically disentangled into query-independent and query-dependent terms, under the same setting as SA layers. Despite the improvements of mJSD scores in classification, DNL layers also degenerate into query-independent operators in a more challenging part segmentation task, revealing that the query-dependent information needs to be exploited more systematically.

3.2 Unary-Pairwise Attention

Based on the observations from the above analysis, we propose the unary-pairwise attention (UPA) to handle the query-dependent and query-independent information simultaneously while ensuring the query dependence of attentions. We propose two distinct formulations that operate in parallel to optimize each component with minimal mutual interference. Furthermore, we apply the UPA to the local regions of query points because the SA tends to degenerate into a query-independent operator given a global receptive field. In such a manner, attention outputs become query-dependent, which we find beneficial for point cloud processing. In addition, constraining operating scope also reduces the time/space complexity from quadratic to linear to the input cardinality, and thus enables UPA to be scalable to voluminous data.

Formally, the general formulation of the UPA for a query \mathbf{x}_i can be defined as:

$$\mathbf{y}_i = \sum_{\mathbf{x}_j \in \mathcal{N}(\mathbf{x}_i)} \text{softmax}(f(\mathbf{x}_i, \mathbf{x}_j))g(\mathbf{x}_j) \quad (2)$$

where $\mathcal{N}(\mathbf{x}_i)$ is k -nearest neighbors of \mathbf{x}_i . f is a relation function that measures the relationship between two inputs. $g: \mathbb{R}^{d_{in}} \mapsto \mathbb{R}^{d_{in}}$ is a transformation function that is implemented as a simple linear projection.

The exact form of the relation function f is need-dependent. We introduce two distinct instantiations, one is responsible for exploiting the unary relation and the other for the pairwise one:

$$f_u(\mathbf{x}_i, \mathbf{x}_j) = W_u \mathbf{x}_j \quad (3)$$

$$f_e(\mathbf{x}_i, \mathbf{x}_j) = W_e(\mathbf{x}_j - \mathbf{x}_i) \quad (4)$$

The relation functions $f_u: \mathbb{R}^{d_{in}} \mapsto \mathbb{R}$ and $f_e: \mathbb{R}^{d_{in}} \mapsto \mathbb{R}$ map input features to scores for subsequent attention weight calculations. $W_u \in \mathbb{R}^{1 \times d_{in}}$ and $W_e \in \mathbb{R}^{1 \times d_{in}}$ are learned pointwise linear projections. Note that the transformation function g is shared between two components to reduce the complexity.

For the unary relation, each neighbor \mathbf{x}_j individually predicts a score for itself; thus, the generated attention map is independent of the pairwise relationship. On the other hand, the

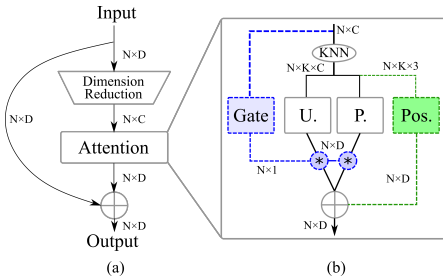


Figure 3: **Attention block designs.** (a) The general attention block used in the experiments. (b) The detail of the UPA block. The green branch is activated when the task is part segmentation whereas the blue one is activated when the task is classification or semantic segmentation.

pairwise relation function f_e maps the relative feature (relative to the query) to the score in which the pairwise interaction between the query and the neighbor is considered.

Like the SA, our formulation can be easily extended to the multi-head setting by arranging the relation function f to predict h scores and perform attention h times in the corresponding subspace of input features. Then the output of each head is concatenated to compose the final outputs $\mathbf{y}_i = \text{Concat}(\mathbf{y}_i^{\text{head}_1}, \dots, \mathbf{y}_i^{\text{head}_h})$, where $\mathbf{y}_i^{\text{head}_h} \in \mathbb{R}^{d_{in}/h}$.

Design of the attention block. The general block design is illustrated in Fig. 3 (a). The block consists of a dimensionality reduction layer, which is an MLP, an attention layer, and a residual connection. The block receives input point clouds and produces the new feature vectors $\mathbf{z}_i \in \mathbb{R}^{d_{out}}$ that are the sum of input features and output of the attention layer. We construct the specific attention block by injecting the selected attention mechanism into the attention layer. The output of an attention mechanism is nonlinearly transformed and added to the input. For instance, in the case of the UPA, the output is calculated as: $\mathbf{z}_i = \mathbf{u}_i + \mathbf{e}_i + \mathbf{x}_i$, where \mathbf{u}_i and \mathbf{e}_i are nonlinearly transformed attention outputs $\alpha(\mathbf{y}_i^{\text{unary}})$ and $\beta(\mathbf{y}_i^{\text{pairwise}})$, respectively. $\alpha, \beta : \mathbb{R}^{d_{in}} \mapsto \mathbb{R}^{d_{out}}$ are component-wise MLPs.

Task-specific UPA block designs. The task-specific designs of the UPA block are illustrated in the Fig. 3 (b). We additionally design the UPA block considering positional information [4, 28] for shape part segmentation tasks in which explicitly encoding 3D layout is found to be beneficial. Specifically, given 3D coordinates of a query point and its neighbor $\mathbf{p}_i, \mathbf{p}_j \in \mathbb{R}^3$, positional features are computed as $\mathbf{x}_{pos} = \delta(\mathbf{p}_j - \mathbf{p}_i)$, where $\delta : \mathbb{R}^3 \mapsto \mathbb{R}^{d_{in}}$ is an MLP consisting of two linear projections with a ReLU activation in between. Then, the positional encoding is produced by following the procedures of the unary attention treating \mathbf{x}_{pos} as the input feature. The position encoding is shown as the green branch in Fig. 3. To tackle shape classification/scene segmentation, a gating mechanism is introduced to adaptively control the amount of information taken from each component. Specifically, each point predicts a score s_i by linearly transforming the input feature \mathbf{x}_i such that $\mathbf{z}_i = \phi(s_i) \cdot \mathbf{u}_i + \varphi(s_i) \cdot \mathbf{e}_i + \mathbf{x}_i$. We expect that an explicit gating is useful for enhancing/suppressing relevant/irrelevant information. For simplicity, we set ϕ and φ as $\text{sigmoid}(s_i)$ and $1 - \text{sigmoid}(s_i)$, respectively. The gating procedure is described graphically in Fig. 3 (the blue branch).

4 Experiments

In this section, we present experimental results on the shape classification, part segmentation, and scene segmentation tasks. The performance of the UPA is compared with recent attention-based networks. Subsequently, we apply the UPA to various backbone networks to investigate its impact on standard and modern networks.




Figure 4: **Task-specific architectures used in this study.** The upper model is used for the classification whereas the bottom one is used for the segmentation task. N represents the number of input points.

Experimental settings. We are particularly interested in comparing the relative performance improvements provided by various attentions. Specifically, experiments are performed by fixing the backbone architecture while altering attention layers. We choose the standard SA and DNL [40] as baselines. Moreover, their localized variants (local-SA and local-DNL) are presented to quantify the direct impact of restricting receptive fields.

We choose PointNet++ [27] as the backbone because of its popularity as a building block in recent studies [17, 18, 53, 56]. The network architectures are illustrated in Fig. 4 and detailed configurations are reported in the supplementary material. Results using other backbones are presented in Sec. 4.4. To compare with recent attention-based methods, we also report the performance of Set Transformer [16], PAT [58], and PointASNL [57]. Other implementation details are reported in the supplementary material.

4.1 Shape Classification

Data. We use ModelNet40 [54] dataset, which contains 9,843 CAD models for training and 2,468 models for testing. We use the preprocessed point cloud data provided by [27] for benchmarking. All inputs are normalized into a unit ball. We augment the input by random anisotropic scaling and random translation. Following [18], the input point number and features are set to 1,024 and 3D coordinates, respectively.

Results. As shown in Table 2, UPA and local-DNL provide the greatest performance gain compared to others, revealing the effectiveness of modeling query-dependent and query-independent information simultaneously. Furthermore, UPA outperforms or is on par with powerful attention-based methods, showing its effectiveness as a compositional module. The usefulness of the restricted operating scope is verified as both local-SA and local-DNL outperform their global counterparts.

4.2 Shape Part Segmentation

Data. We use the ShapeNet Part dataset [59] to evaluate the performance on shape part segmentation. The dataset contains 16,880 models in which 14,006 are used for training and 2,874 for testing. Sixteen shape categories and 50 parts are included, each model being annotated with 2 to 6 parts. We use the data provided by [27] and take randomly sampled 2,048 points with the surface normal as input. The same augmentation strategy as the classification task is used. The voting [27, 56, 57] is applied as a post-processing step.

Method	ModelNet40 [34]	ShapeNet [39]	S3DIS [11]
Set Transformer [16]	89.2	-	-
PAT [38]	91.7	-	60.1
PointASNL [37]	92.9	86.1	62.6
PointNet++ [22]	90.7	85.1	57.3
SA (PointNet++)	92.6 (↑ 1.9)	86.1 (↑ 1.0)	59.6 (↑ 2.3)
Local-SA (PointNet++)	92.7 (↑ 2.0)	86.5 (↑ 1.4)	60.1 (↑ 2.8)
DNL (PointNet++)	92.1 (↑ 1.4)	86.4 (↑ 1.3)	57.9 (↑ 0.6)
Local-DNL (PointNet++)	92.9 (↑ 2.2)	86.3 (↑ 1.2)	61.4 (↑ 4.1)
Ours (PointNet++)	92.9 (↑ 2.2)	86.5 (↑ 1.4)	63.3 (↑ 6.0)

Table 2: **The results of various point cloud understanding tasks.** The performance is measured using overall accuracy (OA), instance average IoU (mIoU), and class averaged IoU (mIoU) for classification, part segmentation, and scene segmentation, respectively.

Results. We use mean instance IoU as the performance metric [21]. As reported in Table 2, UPA and local-SA achieves the best performance. Compared with the ones of the DNL and local-DNL, the UPA achieves the better performance, which verifies the suitability of the proposed formulations. Notably, except the SA, all attention variants outperform PointASNL, in which the standard SA is combined with convolutions, demonstrating their valid improvements over the SA.

4.3 Scene Segmentation

Data. We evaluate our models on Stanford large-scale 3D indoor spaces (S3DIS) [11] dataset for scene segmentation. It contains six indoor environments including 272 rooms. Each point is annotated with one of 13 categories. We follow the data preparation procedure of PointNet [21]. Specifically, each input point is represented by a 9-dim vector (XYZ, RGB, and normalized location as to the room). We train the model for about 50K iterations. We use Area five for testing and the others for training.

Results. As shown in Table 2, the UPA substantially improves the baseline by 6.0 mIoU, outperforming other networks significantly in terms of the relative performance gain. The SA and local-SA successfully enhance the baseline; however, they have achieved lower relative gains. We conjecture that query-dependent information such as smoothness modeled by the pairwise term is especially crucial for scene understanding as many scenes are dominated by flat objects. As shown in Fig. 5, UPA obtains smoother predictions compared with the baseline. UPA successfully provides a greater improvement compared to local-DNL, which further verifies the usefulness of its pairwise term in scene understanding.

4.4 Integration with Modern Architectures

We further investigate the effectiveness of UPA blocks by applying them to a wide range of existing networks. As shown in Table. 3, the UPA provides consistent improvements to all networks. In particular, the UPA successfully enhances PointASNL in which SA layers are used extensively, revealing that UPA is can provide additional benefits beyond the ones of the SA.




Figure 5: **Qualitative results of scene segmentation.** The last column shows the effect of the UPA: **green** indicates the points corrected by the UPA; **red** indicates wrongly modified points; **blue** means unchanged points. The UPA generates smoother predictions compared to the baseline. For challenging scenes with a lot of clutter, the network equipped with UPA blocks is more boundary-aware and able to detect objects that the baseline fails to detect.

Model	ShapeNet			S3DIS		
	Before	After	Δ	Before	After	Δ
PointNet [21]	83.7	85.1	$\uparrow 1.4$	41.1	50.6	$\uparrow 9.5$
PointConv [63]	85.7	86.5	$\uparrow 0.8$	62.8	64.9	$\uparrow 2.1$
RSCNN [18]	86.2	86.5	$\uparrow 0.3$	62.3	63.6	$\uparrow 1.3$
PointASNL [37]	86.1	86.4	$\uparrow 0.3$	62.6	62.9	$\uparrow 0.3$

Table 3: Results of adding UPA blocks to various backbones.

5 Design Analysis

We validate the design choices of the UPA in this section. Note that we do not perform voting in following experiments. PointNet++ is used as the baseline throughout this section.

5.1 Block Component Analysis

As shown in Table 4, the unary attention is more effective in part segmentation while the pairwise one is more effective in scene segmentation. We suspect that query-independent feature is more useful in describing salient part boundaries whereas query-dependent features enforce smoothness in scenes in which flat objects dominate. The best performance is achieved after adding task-specific components to the block designs. However, the optimal design of the block is still an open question, which we leave to future work.

Models	ShapeNet	S3DIS
Baseline	85.1	57.3
Unary	85.9	59.7
Pairwise	85.7	60.2
Unary + Pairwise	85.9	61.8
Unary + Pairwise + Position	86.1	60.8
Unary + Pairwise + Gating	85.7	63.3

Table 4: **Results of the block component analysis.** Task-specific block designs effectively combine both attentions and improve the performance.




Figure 6: **Ablation study of neighborhood size k on ShapeNet.** The performance reaches its peak at 16 and gradually starts to reduce when k gets larger.

(a) Pooling		(b) Stage		(c) Compo. Arrangement	
Baseline	85.1	Baseline	85.1	Baseline	85.1
Mean	85.7	Stage 1	85.5	Unary-Pairwise	86.0
Max	85.3	Stage 2	85.6	Pairwise-Unary	85.4
Attention	86.1	Stage 3	85.8	Parallel	86.1

Table 5: Results of the ablation study.

5.2 Ablation Study

We choose part segmentation as the default task for ablation studies as we think that the task has sufficient complexity. We report the instance mIoU for each experiment.

Neighborhood size k . As shown in Fig. 6, enlarging k from 8 to 16 gradually improves the performance. However, mIoU starts to drop when k gets larger. We conjecture that larger receptive fields contain information that is not helpful or even harmful for targeted tasks, thus complicating the optimization.

Pooling method. Average and Max improve performance, showing that fixed operations wrapped by our block are still beneficial. However, Attention offers more expressiveness as it achieves the best performance.

Stage. We examine the performance gain by adding a UPA block to each stage. Stage n denotes the position of the block after n th set abstraction level. As shown in Table 5, the UPA influences the performance more when it is integrated into deeper stages.

Component arrangement. Here we compare the two arrangements of components: sequential and parallel. As shown in Table 5, the parallel arrangement outperforms all sequential ones, which verifies our design choice.

6 Conclusions

We propose the unary-pairwise attention (UPA) for enhancing 3D point cloud processing. Our analyses show that the standard self-attention (SA), which operates globally, is biased towards query-independent information, leaving query-dependent one not well exploited. As a result, the SA produces similar attention maps for different queries. Therefore, our new attention aims to jointly exploit both information while always being query-dependent by operating locally. Extensive experiments demonstrate that the UPA consistently outperforms the SA and other attentions, especially in the challenging task for which encoding query dependence appears useful. Moreover, equipped with the proposed UPA, the vanilla PointNet++ successfully outperforms or is on par with the state-of-the-art attention-based methods for various tasks. In addition, as a compositional module, the UPA successfully boosts the performance of various modern backbones, demonstrating its wide applicability.

Acknowledgement

This paper is partially supported by a project, JPNP18010, commissioned by the New Energy and Industrial Technology Development Organization (NEDO) and JSPS Grant-in-Aid for Scientific Research (Grant Number 21K12042).

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.
- [2] Irwan Bello, Barret Zoph, Ashish Vaswani, Jonathon Shlens, and Quoc V Le. Attention augmented convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3286–3295, 2019.
- [3] Yue Cao, Jiarui Xu, Stephen Lin, Fangyun Wei, and Han Hu. Gcnet: Non-local networks meet squeeze-excitation networks and beyond. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019.
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020.
- [5] Yunpeng Chen, Yannis Kalantidis, Jianshu Li, Shuicheng Yan, and Jiashi Feng. A²-nets: Double attention networks. *Advances in Neural Information Processing Systems*, 31:352–361, 2018.
- [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.
- [7] Yaodong Cui, Ren Chen, Wenbo Chu, Long Chen, Daxin Tian, and Dongpu Cao. Deep learning for image and point cloud fusion in autonomous driving: A review. *arXiv preprint arXiv:2004.05224*, 2020.
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [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.

- [11] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. Pct: Point cloud transformer. *arXiv preprint arXiv:2012.09688*, 2020.
- [12] Han Hu, Zheng Zhang, Zhenda Xie, and Stephen Lin. Local relation networks for image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3464–3473, 2019.
- [13] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randa-net: Efficient semantic segmentation of large-scale point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11108–11117, 2020.
- [14] Asako Kanezaki, Yasuyuki Matsushita, and Yoshifumi Nishida. Rotationnet: Joint object categorization and pose estimation using multiviews from unsupervised viewpoints. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5010–5019, 2018.
- [15] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553): 436–444, 2015.
- [16] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International Conference on Machine Learning*, pages 3744–3753. PMLR, 2019.
- [17] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on χ -transformed points. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 828–838, 2018.
- [18] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8895–8904, 2019.
- [19] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 922–928. IEEE, 2015.
- [20] Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. Image transformer. In *International Conference on Machine Learning*, pages 4055–4064. PMLR, 2018.
- [21] 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.
- [22] Charles R Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *arXiv preprint arXiv:1706.02413*, 2017.
- [23] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 918–927, 2018.

- [24] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jonathon Shlens. Stand-alone self-attention in vision models. *arXiv preprint arXiv:1906.05909*, 2019.
- [25] Takayuki Shinohara, Haoyi Xiu, and Masashi Matsuoka. Fwnet: Semantic segmentation for full-waveform lidar data using deep learning. *Sensors*, 20(12):3568, 2020.
- [26] 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.
- [27] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6411–6420, 2019.
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017.
- [29] Lei Wang, Yuchun Huang, Yaolin Hou, Shenman Zhang, and Jie Shan. Graph attention convolution for point cloud semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10296–10305, 2019.
- [30] Sinong Wang, Belinda Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.
- [31] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018.
- [32] 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.
- [33] 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.
- [34] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.
- [35] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. *arXiv preprint arXiv:2105.01288*, 2021.
- [36] Mutian Xu, Runyu Ding, Hengshuang Zhao, and Xiaojuan Qi. Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. *arXiv preprint arXiv:2103.14635*, 2021.

- [37] Xu Yan, Chaoda Zheng, Zhen Li, Sheng Wang, and Shuguang Cui. Pointasnl: Robust point clouds processing using nonlocal neural networks with adaptive sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5589–5598, 2020.
- [38] Jiancheng Yang, Qiang Zhang, Bingbing Ni, Linguo Li, Jinxian Liu, Mengdie Zhou, and Qi Tian. Modeling point clouds with self-attention and gumbel subset sampling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3323–3332, 2019.
- [39] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM Transactions on Graphics (ToG)*, 35(6):1–12, 2016.
- [40] Minghao Yin, Zhuliang Yao, Yue Cao, Xiu Li, Zheng Zhang, Stephen Lin, and Han Hu. Disentangled non-local neural networks. In *European Conference on Computer Vision*, pages 191–207. Springer, 2020.
- [41] Kaiyu Yue, Ming Sun, Yuchen Yuan, Feng Zhou, Errui Ding, and Fuxin Xu. Compact generalized non-local network. *arXiv preprint arXiv:1810.13125*, 2018.
- [42] Han Zhang, Ian Goodfellow, Dimitris Metaxas, and Augustus Odena. Self-attention generative adversarial networks. In *International conference on machine learning*, pages 7354–7363. PMLR, 2019.
- [43] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip Torr, and Vladlen Koltun. Point transformer. *arXiv preprint arXiv:2012.09164*, 2020.
- [44] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4490–4499, 2018.
- [45] Xiao Xiang Zhu, Devis Tuia, Lichao Mou, Gui-Song Xia, Liangpei Zhang, Feng Xu, and Friedrich Fraundorfer. Deep learning in remote sensing: A comprehensive review and list of resources. *IEEE Geoscience and Remote Sensing Magazine*, 5(4):8–36, 2017.