diffConv: Analyzing Irregular Point Clouds with an Irregular View

Manxi Lin[™] and Aasa Feragen[™]

Technical University of Denmark, Kongens Lyngby, Denmark {manli,afhar}@dtu.dk

Abstract. Standard spatial convolutions assume input data with a regular neighborhood structure. Existing methods typically generalize convolution to the irregular point cloud domain by fixing a regular "view" through e.g. a fixed neighborhood size, where the convolution kernel size remains the same for each point. However, since point clouds are not as structured as images, the fixed neighbor number gives an unfortunate inductive bias. We present a novel graph convolution named Difference Graph Convolution (diffConv), which does not rely on a regular view. diffConv operates on spatially-varying and density-dilated neighborhoods, which are further adapted by a learned masked attention mechanism. Experiments show that our model is very robust to the noise, obtaining state-of-the-art performance in 3D shape classification and scene understanding tasks, along with a faster inference speed. The code is publicly available at https://github.com/mmmmimic/diffConvNet.

Keywords: 3D point cloud, local aggregation operator, attention

1 Introduction

The availability and affordability of 3D sensors are rapidly increasing. As a result, the point cloud — one of the most widely-employed representations of 3D objects — is used extensively in computer vision applications such as autonomous driving, medical robots, and virtual reality [19,3,1]. The industrial demand has prompted a strong call for exploiting the underlying geometric information in points, fuelling deep learning models that rely on learning effective local point features [19,35,33,11].

Local feature learning schemes [19,38,35] typically consist of two main steps, namely *point grouping* and *feature aggregation*. Point grouping gathers the neighbors of key points, to be encoded and fused in the subsequent feature aggregation.

Unlike grid data, points are unordered and unstructured, posing great challenges to the processing; indeed, standard convolutional neural networks (CNNs) [7,23] only work with grid data in a regular *view*, or neighborhood. Point grouping and feature aggregation are often designed to imitate the regular view, or fixed neighborhood size and structure, that we know from convolutions on images. This imitation imposes inductive biases on the point cloud structure.



Fig. 1. The blue dots refer to key points, green dots denote neighbors, orange circles indicate neighborhoods of key points. The color scale in (d) denotes the neighbor importance of the key point on the airplane body, given by masked attention score.

Fig. 1 illustrates problems associated with these inductive biases for the two most common regular point groupings, namely the k nearest neighbor (KNN) grouping [4] and ball query [19]. KNN fixes its view to the k nearest neighbors of each key point, and the resulting groupings are sensitive to noise and outliers. These are automatically grouped with far away points, sometimes belonging to the object of interest, strongly affecting their downstream features (Fig. 1a).

Ball query, conversely, constrains its view to a pre-defined constant radius ball. We see from (Fig. 1b) that using the same radius, or view, at every location also causes problems: A small radius is needed to avoid grouping noise with far away points as we saw with KNN, but this means that the less dense parts of the airplane, such as its tail, easily become disconnected and mistaken for noise.

We thus see that the inductive bias imposed by regular views leads to problems. Can we process the unstructured point cloud without regular groupings? In this paper, we analyze the irregular point cloud with an *irregular view*. We present a novel graph convolution operator, named Difference Graph Convolution (diffConv), which takes an irregular *view* of the point cloud. In diffConv, we improve ball query with density-dilated neighborhoods where the radius for each point depends on its kernel density. This induces asymmetric neighborhood graphs, and in Fig. 1c) this e.g. helps the plane hold on to its tail.

The dilated ball query still treats all the neighbors inside the ball equally. We further adopt masked attention to introduce an additional, task-specific *learned* irregularity to the neighborhood. As illustrated in Fig 1d, the points close to the key point get less attention than those further away, indicating that the points near the wings, where contextual differences appear, are paid more attention.

In short, our key contributions are as follows:

- We propose diffConv, a local feature aggregator which does not rely on the regular view constraint.
- We present density-dilated ball query, which adapts its view to the local point cloud density, in particular alleviating the influence of noise and outliers.

- We are the first to introduce masked attention to local point feature aggregation. In combination with Laplacian smoothing, it allows us to learn further irregularities, e.g. focusing on contextual feature differences.
- We build a hierarchical learning architecture by stacking diffConv layers. Extensive experiments demonstrate that our model performs on par with state-of-the-art methods in noise-free settings, and outperforms state-of-theart by far in the presence of noise.

2 Background and Related Work

Imitating the success of CNNs in local grid feature learning, most 3D point cloud models attempt to generalize CNNs to point clouds by fixing a regular view.

Local point feature learning with projection-based methods. Early works [21,34,44] convert the point cloud to regular 2D or 3D grids (voxels) and leverage the wellestablished standard image convolutions. Su et al. [21] projected the point cloud to multiple views and utilized 2D convolutions with max-pooling to describe the global features of points. MSNet [34] partitions the point cloud into multi-scale voxels, where 3D convolutions are applied to learn discriminative local features. These methods recast the point cloud to a regular view via the convolution-friendly grid representation. Nevertheless, they suffer from the expensive computation and inevitable loss of geometric detail due to discretization.

Point-based methods. In contrast to the expensive projection or voxelization, point-based methods [10,9,12] process the point cloud directly and efficiently. The pioneering PointNet [18] learns point features independently with multi-layer perceptrons (MLPs) and aggregates them with max-pooling. Since pointwise MLPs cannot capture local geometric structures, follow-up works obtain local point features via regular views from different types of point grouping.

Point grouping. The KNN point grouping does not guarantee balanced density [15] and adopts different region scales at different locations of the point cloud [19]. It thus comes with an inductive bias that makes it sensitive both to noise and outliers in point clouds, as well as to different local point densities, as also shown in Fig. 1. Moreover, since KNN queries a relatively small neighborhood in the point-dense regions, the features encoded in subsequent aggregation might not contain enough semantic information to be discriminative. To keep the noise and outliers isolated, some works [19,14] adopt ball query with an upper limit of k points. Balls with less than k points are filled by repeating the firstpicked point. Neighborhoods with more than k points are randomly subsampled. The upper limit k prevents the neighborhood from having too many neighbors, but does also lead to information loss. Besides, the constant ball radius and limit k are not adaptive to the uneven density of real-world point clouds.

Local point feature learning with point-based methods. PointNet++ [19] suggests a hierarchical learning scheme for local point features, including point grouping and feature aggregation. Hierarchically, PointNet++ focuses on a set of subsampled key points, which are represented by a neighborhood defined via point grouping according to pre-defined rules. Subsequently, point aggregation encodes and fuses the features of neighboring points to represent the latent shape of the point cloud structure. These point-based local feature learning methods regularize the input point cloud in the spatial or spectral domain. The point grouping in PointNet++ is done via ball query, and PointNet++ thus also inherits the corresponding inductive biases. Although the multi-scale and multi-resolution feature learning in PointNet++ can handle the uneven density to a degree, they are too expensive for application.

Graph convolutional neural networks (GCNs) [8,36,2,45] generalize standard convolutions to graph data by considering convolutions in the spectral domain. Some methods [17,25,32] adopt GCNs for point aggregation. These methods represent the point cloud as an undirected graph for spectral analysis, where neighbors are constrained to a fixed radius ball. Thus, these methods obtain a regular view via spectral analysis, effectively using a form of ball query. It follows that the resulting algorithms also come with the inductive biases associated with ball query. Additionally, GCNs entail extra computations for signal transformation to the spectral domain.

Alternative methods find other ways to define a k-size kernel for each key point, often based on KNN. In particular, DGCNN [35] groups points by their k nearest neighbors in feature space, and then aggregates the feature difference between points and their neighbors with max-pooling. In a similar way, CurveNet [38] takes guided walks and arguments the KNN neighborhood by hypothetical curves. Other works [33,27,42,39,46] learn the importance weights of the KNN features to introduce some neighborhood irregularity, but are still confined to the fixed, and sometimes spatially limited, KNN view.

Differing from the previous generalizations of CNNs, which all rely on taking a regular view of the point clouds in order to define convolutional operators, we suggest an irregular view of point cloud analysis in this paper.

3 Method

We first revisit the general formulation of local feature learning in Section 3.1. Then we propose a flexible convolution operator, namely difference graph convolution in Section 3.2. In Section 3.3 and Section 3.4, we enrich the basic diffConv with density-dilated ball query and masked attention. We present our network architectures for 3D shape classification and segmentation in supplements.

3.1 Revisit Local Point Feature Learning

Given a point cloud consisting of N points $\mathcal{P} = \{p_i | i = 1, 2, ..., N\} \in \mathbb{R}^{N \times 3}$, and a data matrix containing their feature vectors $X = [x_1, x_2, x_3, ..., x_N]^T \in \mathbb{R}^{N \times d}$ for each point p_i , local feature extraction can be formulated as:

$$g_i = \Lambda(h(p_i, p_j, x_i, x_j) | p_j \in \mathcal{N}(p_i)) \tag{1}$$

where g_i denotes the learned local feature vector, $\mathcal{N}(p_i)$ refers to the neighbors of p_i , $h(\cdot)$ is a function encoding both the coordinates p_i of the i^{th} point, and the coordinates p_j of its neighbors, as well as their feature vectors x_i and x_j , respectively. Moreover, Λ denotes an aggregation function such as MAX or AVG.

Most previous works focus on the design of $h(\cdot)$ [35,46,47,33]. Among them, one of the most widely applied methods is the edge convolution (edgeConv) appearing in DGCNN [35]. In edgeConv, the neighborhood of a point is defined by its KNN in the feature space. The feature difference $x_i - x_j$ is used to indicate the pair-wise local relationship between the point and its neighbor. The relationships are concatenated with the original point features. The concatenated pair is processed by a multi-layer perceptron l_{θ} and summarized by a MAX aggregation. The whole process can be described as

$$g_i = MAX(l_\theta(x_i - x_j || x_i)) \tag{2}$$

where || refers to concatenation, and $p_j \in \mathcal{N}(p_i)$. In this paper, DGCNN is included as a baseline model. In edgeConv, KNN brings convenience to matrixlevel operations, such as feature concatenation and linear transformation. One challenge tackled in this paper is the generalization of edgeConv to the irregular view, and in particular inclusion of feature differences into the convolution.

3.2 Difference Graph Convolution

Inspired by GCNs [8] and edgeConv [35], we propose the Difference Graph Convolution (diffConv). We present a basic version of diffConv in this subsection, with further improvements described in the following subsections. Fig. 2 gives an overview of the complete diffConv, where the querying radius is pre-computed.

By treating each point as a node, we represent the aforementioned point cloud \mathcal{P} as a directed graph $\mathcal{G} = \{X, A\}$, where A is the adjacency matrix, and X refers to point feature vectors. In \mathcal{G} , each point is connected to its neighbors. The commonly-adopted neighbor-wise feature difference in edgeConv has been proven to be very effective in capturing the local structure of point clouds [11,35,38]. When processing graphs, the Laplacian matrix efficiently measures the difference between nodes and their neighbors. Borrowing from GCNs [8], we apply Laplacian smoothing on the point feature vector by

$$S = X - \hat{A}X \tag{3}$$

where $S = [s_1, s_2, ..., s_N]^T$ contains the updated feature vectors and \hat{A} is the normalized adjacency matrix. Like edgeConv, our Difference Graph Convolution is an MLP l_{θ} on the combination of S and the original point features X

$$G = l_{\theta}(S||X) \tag{4}$$

where $G = [g_1, g_2, ..., g_N]^T$ includes the learned feature vectors. Here, X is concatenated to present the global shape information of the point cloud.



Fig. 2. Structure of our diffConv. Specifically, masking denotes the conditional statements in Eq. 9, which takes the adaptive radius r_i from density-dilated ball query. In our implementation, d_k was set to $\frac{(d+3)}{4}$. M is the number of sampled key points. The dotted line is used to avoid overlapping with the solid lines.

Relationship to edgeConv. We argue that our diffConv is a variant of edgeConv in a more flexible neighborhood. When applying diffConv on KNN grouping, it only differs in the sequence of nonlinear activation and aggregation from edge-Conv. As a demonstration, we consider a basic diffConv, with a simple binary adjacency matrix A whose entries indicate connections between points. When the neighborhood is defined by ball query (without an upper limit on the number of neighbors), the connection between p_i and p_j , A_{ij} , is given by

$$A_{ij} = \begin{cases} 1 & \text{if } \|p_i - p_j\|_2 < r \\ 0 & \text{otherwise} \end{cases}$$
(5)

where r is the priory-given ball query search radius. An arithmetic mean is leveraged to normalize the adjacency matrix. According to Eq. 3, for $p_j \in \mathcal{N}(p_i)$, the smoothed feature vector s_i calculated by

$$s_i = x_i - AVG(x_j) = AVG(x_i - x_j)$$
(6)

is the average of the neighbor-wise feature difference from Eq. 2. In line with Eq. 4, diffConv for point p_i becomes $g_i = l_{\theta} (AVG(x_i - x_j)||x_i)$. When l_{θ} is linear, diffConv can be further described by $g_i = AVG(l_{\theta}(x_i - x_j)||x_i))$. Compared with Eq. 2, in the case of linear activation, the only difference between diffConv and edgeConv is in the aggregation function. When l_{θ} is nonlinear, the feature difference is activated before aggregation, while diffConv operates in reverse order. In contrast to edgeConv, our diffConv is more flexible with respect to the definition of neighborhood - the number of neighbors is no longer fixed.



Fig. 3. Left: Ball query in a point cloud with noise. The blue circles refer to key points, green circles denote neighbors, orange circles refer to neighborhoods, and gray circles denote other points in the point cloud. **Right:** Visualizations of the kernel density on some objects from ModelNet40 [37]. Brighter color indicates higher density.

3.3 Density-dilated Ball Query

In Eq. 5, the point cloud is constructed as an undirected graph, where $A_{ij} = A_{ji}$ for point pair p_i and p_j . The directed relations between points and their neighbors are therefore neglected, as illustrated in Fig. 1b). We relax this to a density-dilated ball query where the search radius for each point is expanded according to its kernel density.

Point density, defined as the number of neighbors of each point, reflects the spatial distribution of points. Fig. 3 demonstrates a point cloud with noise. Note that under a fixed-radius ball query, contour points of the 3D object (e.g. point B) have a lower density than the points from the flat area (e.g. point A). Noise generates extreme cases (e.g. point C), which can be isolated from the object.

As illustrated in Fig. 1, in KNN, points with low density, such as noise, get a spatially larger neighborhood than others, grouping them with unrelated features far away. The adversarial effect of noise points on estimated object properties is thus negatively correlated with their density, and assigning smaller neighborhoods to low-density points therefore contributes to resisting noise.

Xiang et al. [38] claim that boundary points provide different feature differences from neighbors while points from flat areas usually have similar and unrecognizable local feature differences. To enlarge the variance of feature difference, points with a higher density (points from flat areas) thus ought to have larger local regions. The KNN grouping queries small neighborhoods for points with high density. Our ball query enlarges the neighborhood to include more contextually different neighbors, giving access to more diverse information.

Instead of counting neighbors, which is computationally expensive, we describe the spatial distribution of points using kernel density estimation [28]. For each point p_i , its kernel density d_i is estimated by a Gaussian kernel

$$d_{i} = \frac{1}{Nh} \sum_{j=1}^{N} \frac{1}{\sqrt{2\pi}} e^{-\frac{\|p_{i} - p_{j}\|_{2}^{2}}{2h^{2}}},$$
(7)

where h is a parameter controlling the bandwidth. The density d_i indicates the probability of point p_i to be located in flat areas of the object, i.e., the degree of dilation. Fig. 3 shows the kernel density of various objects, where points from flat areas have a higher density. This is in accordance with our former analysis.

Based on this density, we softly dilate the searching radius from Eq. 5 to

$$r_i = \sqrt{r^2 (1 + \hat{d}_i)},\tag{8}$$

where the square and root operations are applied to slow down the dilation speed, \hat{d}_i denotes the normalized kernel density within 0 and 1 by dividing the largest density. The dilated search radius varies from point to point, resulting in a directed graph. Since the search radius is enlarged to different degrees, the message propagation on the graph is boosted, accelerating the long-range information flow. Finally, our method is intuitively robust to noise, which is also demonstrated by our experiments in Section 4.4.

3.4 Masked Attention

The dilated ball query still treats all the neighbors inside the ball equally. We further introduce an additional, task-specific learned irregularity to the neighborhood by attention mechanism [5,30]. Different from other attention-based methods [46,33], we employ the masking mechanism in self-attention [30] to handle the no-longer-fixed neighborhood. We call this learning scheme masked attention. In a transformer built by self-attention [30] [31], there are two kinds of masks: the padding mask in the encoder is exploited to address sequences with different lengths and the sequence mask in the decoder is used to prevent ground truth leakage in sequence prediction tasks. In contrast, we employ a neighborhood mask in our masked attention to learn the local features of points. It works in the encoder while playing a different role than the padding mask. It includes three steps, the calculation of edge weight for each pair of connected points on the graph, local normalization, and balanced renormalization.

First, we revise the adjacency matrix A in Eq. 5 to

$$A_{ij} = \begin{cases} l_{\phi}(x_i || p_i) l_{\psi}(x_j || p_j)^T & \text{if } || p_i - p_j ||_2 < r_i \\ -\infty & \text{otherwise,} \end{cases}$$
(9)

where l_{ϕ} and l_{ψ} are two MLPs mapping the input to a given dimension d_k . The MLPs are employed to learn the latent relationship and alignment between p_i and p_j in the feature space. The MLPs encode both point features and coordinates, as we expect them to recognize both the semantic and geometric information. In the implementation, $-\infty$ is approximated by -10^9 , which is orders of magnitude smaller than the computed attention score.

The adjacency matrix is then normalized by softmax

$$\tilde{A}_{i,j} = \frac{e^{A_{i,j}}}{\sum_{k}^{N} e^{A_{i,k}}}.$$
(10)

Since $e^{-10^9} \approx 0$, points outside the neighborhood are automatically masked. Hence, Eq. 10 becomes the softmax normalization in neighborhoods. With the masking mechanism, we normalize the attention score locally with only matrixwise operations instead of expensive iterations. In each local region, neighbor features are aggregated dynamically and selectively.

In self-attention [30], the dot product in Eq. 9 is scaled by $\sqrt{d_k}$ before the softmax normalization, in order to prevent the product from getting extreme values as d_k increases. Here, we refine the scaling by a balanced renormalization of the attention scores \tilde{A} from Eq. 10. We apply the l_1 -norm twice on the square root of the former attention score in different dimensions by

$$\overline{A}_{ij} = \frac{\sqrt{\tilde{A}_{ij}}}{\sum_{k}^{N} \sqrt{\tilde{A}_{kj}}} \quad \text{and} \quad \hat{A}_{ij} = \frac{\overline{A}_{ij}}{\sum_{k}^{N} \overline{A}_{ik}}.$$
(11)

Here, the square root reduces the variation in attention scores, to prevent overly attention on a few points. Considering the matrix product in Eq. 3, l_1 -norm is to balance the contribution of feature channels and neighbor points respectively. The experimental results in supplemental materials illustrate the effectiveness of the proposed renormalization strategy.

Position encoding. Local spatial information is defined as the coordinate difference between points and their neighbors. Many prior works have demonstrated the importance of local spatial information in point cloud analysis [33,14,47]. We employ the Local Point-Feature Aggregation block [38] as the position encoder in diffConv. The encoded coordinates are fused with point features by

$$G = \sigma(l_{\theta}(S||X) + l_{\pi}(P)), \qquad (12)$$

where σ is a nonlinear activation function, $P = [p_1, p_2, p_3, ..., p_N]^T \in \mathbb{R}^{N \times 3}$ contains the 3D coordinates of points in \mathcal{P} , and l_{π} is the position encoder working in the dilated neighborhoods and mapping the positional difference to the same dimension as the updated point features from l_{θ} .

4 Experiments

We demonstrate our model performance in 3D shape classification, object shape part segmentation, and real-scene semantic segmentation tasks. Due to the limited space, we describe the training settings as well as the corresponding ablation studies on the hyperparameters in supplements.

4.1 Point Cloud Classification

ModelNet40 with the official split. We conducted experiments performing object shape classification on the ModelNet40 dataset [37], which contains 12,311

Table 1. 3D shape classification results on ModelNet40 dataset with official split as well as ModelNet40-C, and part segmentation results on ShapeNet Part dataset. Here, 'xyz' stands for coordinates of points; 'nor' stands for normals of points; '#points' means the number of input points in classification; 'OA' stands for the overall accuracy over instances on ModelNet40; and 'MA' stands for the mean accuracy within each category on ModelNet40; 'CER' denotes the corruption error rate on ModelNet40-C; 'mIOU' denotes the instance average Inter-over-Union on ShapeNetPart. Not all papers report MA and CER; these are recorded as '-' in the table.

Methods	Input	#points	OA(%)	MA(%)	$\operatorname{CER}(\%)$	mIOU (%)
A-CNN[10]	xyz+nor	1k	92.6	-	-	86.1
Kd-Net[9]	xyz	32k	91.8	88.5	-	82.3
SO-Net[12]	xyz+nor	5k	93.4	90.8	-	84.6
PointNet[18]	xyz	1k	89.2	-	28.3	83.7
PointNet++[19]	xyz	1k	90.7	-	23.6	85.1
DGCNN[35]	xyz	1k	92.9	90.2	25.9	85.1
CurveNet[38]	xyz	1k	93.8	-	22.7	86.6
RS-CNN[14]	xyz	1k	92.9	-	26.2	86.2
Spec-GCN [32]	xyz	1k	91.5	-	-	85.4
PointASNL[41]	xyz	1k	92.9	-	-	86.1
PCT[6]	xyz	1k	93.2	-	25.5	86.4
GDANet[40]	xyz	1k	93.4	-	25.6	86.5
Ours	xyz	1k	93.6	90.6	21.4	85.7

synthetic meshed CAD models from 40 different categories. The official split includes 9,843 models for training and 2,468 models for testing. Following most publicly available implementations [35,38,19], our validation on the official split was made by reporting the best-observed test performance over all epochs. Similar to previous works, we represented each CAD model by 1,024 points uniformly sampled from mesh faces with (x, y, z) coordinates.

We measured both the overall accuracy (OA), treating every point cloud with equal weight, and the mean accuracy (MA), where the mean was computed from the category-specific accuracies. Note that the MA up-weights the importance of small categories, giving them equal weight as the larger categories.

Our results can be found in Table 1. All results for other methods are taken from the cited papers. Note that RS-CNN, CurveNet, and GDANet adopt a voting strategy to improve the model performance. We report their classification accuracy without this trick for a fair comparison. Our model achieves the secondbest overall accuracy at 93.6%. Even compared to SO-Net, which takes 5k points with normal as input, our model has a better performance in OA.

ModelNet40 with resplits for assessing robustness. In addition to benchmarking on the official split, we also performed repeated experiments on resampled dataset splits, to obtain an evaluation that is robust to over-fitting to any specific data split. Here, we selected parameters in a more traditional fashion with

Methods	OA(%)	MA(%)	OBJ_ONLY	OBJ_BG
PointNet[18]	90.19 ± 0.58	84.55 ± 0.81	79.2	73.3
DGCNN[35]	92.89 ± 0.43	89.62 ± 0.83	86.2	82.8
CurveNet[38]	92.86 ± 0.49	89.51 ± 0.81	84.3	84.4
Ours	93.15 ± 0.34	89.86 ± 0.56	86.6	84.9

Table 2. Classification results on ModelNet40 with resplits and ScanObjectNN. OBJ_ONLY and OBJ_BG denote the overall accuracy in the corresponding variant.

the help of an independent validation set. To this end, we randomly re-split the dataset into 8,617/1,847/1,847 point clouds for training/validation/testing, respectively, and picked the model with the best validation performance for evaluation on the test. This was repeated 10 times, reporting the mean and standard deviation of the corresponding performance measures. The baseline models were trained with settings from their paper.

We compare our model performance with PointNet, DGCNN, and CurveNet, where the latter has the highest performance on the official split. According to Table 2, our method outperforms all the other models in both overall and classmean accuracy, showing a better generalization capability on the resampled data splits than other methods. Moreover, our model has the smallest standard deviation in both OA and MA, indicating a more stable performance. This supports our hypothesis that diffConv's irregular view gives discriminative local features.

ModelNet40-C with corruptions for assessing robustness. To further assess the corruption robustness, we evaluated the model trained on ModelNet40 with the official split on the ModelNet40-C dataset [22], which imposes 15 common and realistic corruptions with 5 severity levels on the original ModelNet40 point clouds. We used the corruption error rate (CER) as the evaluation metric.

The corruption performance of different models is shown in Table 1. All results for other methods are taken from [22]. Our model outperforms all existing methods with at least a 1.3% difference, showing strong robustness to corruption.

ScanObjectNN. Besides the synthetic datasets, we also validated our model on the ScanObjectNN dataset [29], with 2,902 scanned real-world objects from 15 categories, with 1,024 points each. Compared to ModelNet40, ScanObjectNN is more challenging due to the presence of noise, backgrounds, and occlusion. 80% of the objects are split into the official training set and the remaining 20% are used for testing. We conducted experiments on objects with and without background respectively (the OBJ_BG and OBJ_ONLY variant of the dataset).

Table 2 summarizes the performance of our model and the baselines. The results of PointNet and DGCNN are taken from [29], while CurveNet was retrained as it is not shown in [29]. Our model gets the highest OA on both two variants, indicating that it is practical in addressing real-world problems.



Fig. 4. Example masked attention maps for ModelNet40-C. Red dots are picked key points; green dots denote neighbors; the color scale denotes the masked attention score. Zoom in for details.

Visualizations of masked attention. Our model achieves state-of-the-art performance and shows strong robustness in 3D object classification. Fig. 4 illustrates some examples of the attention maps from the second diffConv layer in our classification network on ModelNet40. We randomly picked two key points from 8 different objects of the ModelNet40-C benchmark. As shown in the figure, the noise points are isolated and the flat-area points have a larger neighborhood than the boundary points. Besides, neighbors with larger contextual differences to the query point, e.g., the edges of the airplane and the laptop, usually get more attention. More visualizations are presented in the supplements.

3D Semantic Segmentation 4.2

We performed the scene-level understanding task on Toronto3D [24]. Compared with the indoor datasets, the Toronto3D dataset, collected by LiDAR in outdoor scenarios, contains more noise and requires more from the model's robustness. Toronto3D contains 8 different classes and 78.3M points. Following TGNet [13], we divided the scene into 5mx5m blocks, and sampled 2,048 points from each block. There are 956 samples for training and 289 for testing. The mean intersection over union over categories(mean IoU) and IoU for each category are shown in Tab. 3, where the results of other methods are taken from [24]. Our model outperforms all other methods in mean IOU, and in many categories. Specifically, our model surpasses other methods significantly in recognizing small objects, such as road marks and poles. This demonstrates the effectiveness of our method on scene-level understanding tasks.

Shape Part Segmentation 4.3

We conducted the 3D part segmentation task on the ShapeNetPart dataset [43], which consists of 16,881 CAD models from 16 categories with up to 50 different

12

 Table 3. Segmentation results on Toronto3D. Here, mean IoU denotes the category average Inter-over-Union in Toronto3D.

Methods	mean IoU	Road	Rd. mrk.	Natural	Building	Util. line	Pole	Car	Fence
PointNet++ [19]	56.55	91.44	7.59	89.80	74.00	68.60	59.53	52.97	7.54
PointNet+(MSG) [19]	53.12	90.67	0.00	86.68	75.78	56.20	60.89	44.51	10.19
DGCNN [35]	49.60	90.63	0.44	81.25	63.95	47.05	56.86	49.26	7.32
KPFCNN [26]	60.30	90.20	0.00	86.79	86.83	81.08	73.06	42.85	21.57
MS-PCNN [16]	58.01	91.22	3.50	90.48	77.30	62.30	68.54	52.63	17.12
TG-Net [13]	58.34	91.39	10.62	91.02	76.93	68.27	66.25	54.10	8.16
MS-TGNet [24]	60.96	90.89	18.78	92.18	80.62	69.36	71.22	51.05	13.59
Ours	76.73	83.31	51.06	69.04	79.55	80.48	84.41	76.19	89.83

parts. Each CAD model is represented by 2,048 uniformly-sampled points. We followed [33] and trained our model on 12,137 instances with the rest for testing. All 2,048 points in each point cloud were used.

We report the instance average Inter-over-Union (mIOU) in Table 1. The results for other methods are taken from the cited papers. Fig. 5 shows our model predictions, which are very close to the ground truth.

4.4 Performance analysis

Robustness study. According to our analysis, our method is robust to noise, especially in contrast to KNN-based models. To further verify the robustness of our model, like KCNet [20], we evaluated the model on noisy points in both classification and part segmentation tasks. A certain number of randomly picked points were replaced with uniform noise ranging from -1 to 1. In segmentation, these noise points were labeled the same as the nearest original point. Fig. 6 demonstrates the results. We compared our model with the baseline methods used in previous experiments. In classification, we also report the performance of KCNet and PointASNL, with the results taken from papers. As presented in the figure, our model shows clear robustness to noise. Even with 100 noise points, our method has an OA of 86.2% and an mIOU of 74.3%.

Computational resources We compare the corrupted error rate on ModelNet40-C, inference time (IT), floating-point operations (FLOPs) required, parameter number (Params), and memory cost against baselines in Table 4. FLOPs and



Fig. 5. Examples of ShapeNetPart segmentation. Ground truth (top); ours (bottom).

14 Lin. et al.



Fig. 6. Results of robustness study on ModelNet40 (left) and ShapeNetPart (right).

Params are estimated by PyTorch-OpCounter¹. The inference time per batch is recorded with a test batch size of 16. We conducted the experiments on a GTX 1080 Ti GPU. For a fair comparison in inference time and memory cost, each model is implemented in Python without CUDA programming. The table illustrates the efficiency of our method, which is twice as fast as CurveNet.

Table 4. Computational resource requirements.

Methods	$\operatorname{CER}(\%)$	IT (ms)	FLOPs	Params	Memory
PointNet [18]	28.3	19	0.45G	$3.47 \mathrm{M}$	$2061 \mathrm{M}$
DGCNN [35]	25.9	88	2.43G	$1.81 \mathrm{M}$	$8067 \mathrm{M}$
CurveNet [38]	22.7	190	0.66G	2.14M	4309M
Ours	21.4	78	0.28G	1.31M	$3971 \mathrm{M}$

5 Conclusion

We have reviewed contemporary local point feature learning, which heavily relies on a regular inductive bias on point structure. We question the necessity of this inductive bias, and propose diffConv, which is neither a generalization of CNNs nor spectral GCNs. Instead, equipped with Laplacian smoothing, density-dilated ball query and masked attention, diffConv analyzes point clouds through an *irregular view*. Through extensive experiments we show that diffConv achieves state-of-the-art performance in 3D object classification and scene understanding, is far more robust to noise than other methods, with a faster inference speed.

Acknowledgements

This work was supported by the DIREC project EXPLAIN-ME, the Novo Nordisk Foundation through the Center for Basic Machine Learning Research in Life Science (NNF20OC0062606), and the Pioneer Centre for AI, DNRF grant nr P1.

¹ https://github.com/Lyken17/pytorch-OpCounter

References

- Blanc, T., El Beheiry, M., Caporal, C., Masson, J.B., Hajj, B.: Genuage: visualize and analyze multidimensional single-molecule point cloud data in virtual reality. Nature Methods 17(11), 1100–1102 (2020)
- Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. Advances in neural information processing systems 29, 3844–3852 (2016)
- Dometios, A.C., Papageorgiou, X.S., Arvanitakis, A., Tzafestas, C.S., Maragos, P.: Real-time end-effector motion behavior planning approach using on-line pointcloud data towards a user adaptive assistive bath robot. In: 2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 5031–5036. IEEE (2017)
- 4. Fix, E., Hodges Jr, J.L.: Discriminatory analysis-nonparametric discrimination: Small sample performance. Tech. rep., California Univ Berkeley (1952)
- Gehring, J., Auli, M., Grangier, D., Dauphin, Y.: A convolutional encoder model for neural machine translation. In: Proceedings of the Annual Meeting of the Association for Computational Linguistics. pp. 123–135 (Jul 2017)
- 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)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 770–778 (2016)
- Kipf, T.N., Welling, M.: Semi-Supervised Classification with Graph Convolutional Networks. In: Proceedings of the 5th International Conference on Learning Representations (2017)
- Klokov, R., Lempitsky, V.: Escape from cells: Deep kd-networks for the recognition of 3d point cloud models. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 863–872 (2017)
- 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. pp. 7421–7430 (2019)
- Li, G., Muller, M., Thabet, A., Ghanem, B.: Deepgcns: Can gcns go as deep as cnns? In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9267–9276 (2019)
- Li, J., Chen, B.M., Lee, G.H.: So-net: Self-organizing network for point cloud analysis. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9397–9406 (2018)
- Li, Y., Ma, L., Zhong, Z., Cao, D., Li, J.: Tgnet: Geometric graph cnn on 3-d point cloud segmentation. IEEE Transactions on Geoscience and Remote Sensing 58(5), 3588–3600 (2019)
- 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. pp. 8895–8904 (2019)
- Liu, Z., Hu, H., Cao, Y., Zhang, Z., Tong, X.: A closer look at local aggregation operators in point cloud analysis. In: European Conference on Computer Vision. pp. 326–342. Springer (2020)
- Ma, L., Li, Y., Li, J., Tan, W., Yu, Y., Chapman, M.A.: Multi-scale point-wise convolutional neural networks for 3d object segmentation from lidar point clouds in large-scale environments. IEEE Transactions on Intelligent Transportation Systems 22(2), 821–836 (2019)

- 16 Lin. et al.
- 17. Pan, G., Wang, J., Ying, R., Liu, P.: 3dti-net: Learn inner transform invariant 3d geometry features using dynamic gcn. arXiv preprint arXiv:1812.06254 (2018)
- 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 conference on computer vision and pattern recognition. pp. 652–660 (2017)
- Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Proceedings of the 31st International Conference on Neural Information Processing Systems. p. 5105–5114 (2017)
- Shen, Y., Feng, C., Yang, Y., Tian, D.: Mining point cloud local structures by kernel correlation and graph pooling. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4548–4557 (2018)
- Su, H., Maji, S., Kalogerakis, E., Learned-Miller, E.: Multi-view convolutional neural networks for 3d shape recognition. In: Proceedings of the IEEE international conference on computer vision. pp. 945–953 (2015)
- Sun, J., Zhang, Q., Kailkhura, B., Yu, Z., Xiao, C., Mao, Z.M.: Benchmarking robustness of 3d point cloud recognition against common corruptions. arXiv preprint arXiv:2201.12296 (2022)
- Tan, M., Le, Q.: Efficientnet: Rethinking model scaling for convolutional neural networks. In: International Conference on Machine Learning. pp. 6105–6114. PMLR (2019)
- 24. Tan, W., Qin, N., Ma, L., Li, Y., Du, J., Cai, G., Yang, K., Li, J.: Toronto-3d: A large-scale mobile lidar dataset for semantic segmentation of urban roadways. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 202–203 (2020)
- Te, G., Hu, W., Zheng, A., Guo, Z.: Rgcnn: Regularized graph cnn for point cloud segmentation. In: Proceedings of the 26th ACM international conference on Multimedia. pp. 746–754 (2018)
- 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. pp. 6411–6420 (2019)
- Tian, F., Jiang, Z., Jiang, G.: Dnet: Dynamic neighborhood feature learning in point cloud. Sensors 21(7), 2327 (2021)
- 28. Turlach, B.A.: Bandwidth selection in kernel density estimation: A review. In: CORE and Institut de Statistique. Citeseer (1993)
- Uy, M.A., Pham, Q.H., Hua, B.S., Nguyen, D.T., Yeung, S.K.: Revisiting point cloud classification: A new benchmark dataset and classification model on realworld data. In: International Conference on Computer Vision (ICCV) (2019)
- 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. pp. 5998–6008 (2017)
- Veličković, P., Cucurull, G., Casanova, A., Romero, A., Liò, P., Bengio, Y.: Graph attention networks. 6th International Conference on Learning Representations (2017)
- 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)
- 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. pp. 10296–10305 (2019)
- Wang, L., Huang, Y., Shan, J., He, L.: Msnet: Multi-scale convolutional network for point cloud classification. Remote Sensing 10(4), 612 (2018)

- 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 On Graphics (tog) 38(5), 1–12 (2019)
- Wu, F., Souza, A., Zhang, T., Fifty, C., Yu, T., Weinberger, K.: Simplifying graph convolutional networks. In: International conference on machine learning. pp. 6861– 6871. PMLR (2019)
- 37. 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 conference on computer vision and pattern recognition. pp. 1912–1920 (2015)
- Xiang, T., Zhang, C., Song, Y., Yu, J., Cai, W.: Walk in the cloud: Learning curves for point clouds shape analysis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 915–924 (October 2021)
- 39. 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. pp. 3173–3182 (2021)
- Xu, M., Zhang, J., Zhou, Z., Xu, M., Qi, X., Qiao, Y.: Learning geometrydisentangled representation for complementary understanding of 3d object point cloud. arXiv preprint arXiv:2012.10921 (2020)
- Yan, X., Zheng, C., Li, Z., Wang, S., Cui, S.: 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. pp. 5589–5598 (2020)
- Yang, Y., Ma, Y., Zhang, J., Gao, X., Xu, M.: Attpnet: Attention-based deep neural network for 3d point set analysis. Sensors 20(19), 5455 (2020)
- 43. 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)
- 44. Zhang, L., Sun, J., Zheng, Q.: 3d point cloud recognition based on a multi-view convolutional neural network. Sensors **18**(11), 3681 (2018)
- Zhang, M., Cui, Z., Neumann, M., Chen, Y.: An end-to-end deep learning architecture for graph classification. In: Thirty-Second AAAI Conference on Artificial Intelligence (2018)
- Zhao, H., Jiang, L., Jia, J., Torr, P.H., Koltun, V.: Point transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 16259– 16268 (2021)
- 47. Zhou, H., Feng, Y., Fang, M., Wei, M., Qin, J., Lu, T.: Adaptive graph convolution for point cloud analysis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4965–4974 (2021)