A-CNN: Annularly Convolutional Neural Networks on Point Clouds

Artem Komarichev   Zichun Zhong   Jing Hua

Department of Computer Science, Wayne State University

{artem.komarichev,zichunzhong,jinghua}@wayne.edu

Abstract

Analyzing the geometric and semantic properties of 3D point clouds through the deep networks is still challenging due to the irregularity and sparsity of samplings of their geometric structures. This paper presents a new method to define and compute convolution directly on 3D point clouds by the proposed annular convolution. This new convolution operator can better capture the local neighborhood geometry of each point by specifying the (regular and dilated) ring-shaped structures and directions in the computation. It can adapt to the geometric variability and scalability at the signal processing level. We apply it to the developed hierarchical neural networks for object classification, part segmentation, and semantic segmentation in large-scale scenes. The extensive experiments and comparisons demonstrate that our approach outperforms the state-of-the-art methods on a variety of standard benchmark datasets (e.g., ModelNet10, ModelNet40, ShapeNet-part, S3DIS, and ScanNet).

1. Introduction

Nowadays, the ability to understand and analyze 3D data is becoming increasingly important in computer vision and computer graphics communities. During the past few years, the researchers have applied deep learning methods to analyze 3D objects inspired by the successes of these techniques in 2D images and 1D texts. Traditional low-level handcrafted shape descriptors suffer from not being able to learn the discriminative and sufficient features from 3D shapes [1]. Recently, deep learning techniques have been applied to extract hierarchical and effective information from 3D shape features captured by low-level descriptors [20, 6]. 3D deep learning methods are widely used in shape classification, segmentation, and recognition, etc. But all these methods are still constrained by the representation power of the shape descriptors.

One of the main challenges to directly apply deep learning methods to 3D data is that 3D objects can be represented in different formats, i.e., regular / structured representation


Figure 1: The proposed annularly convolutional neural networks (A-CNN) model on point clouds to perform classification, part segmentation, and semantic segmentation tasks.

(e.g., multi-view images and volumes), and irregular / unstructured representation (e.g., point clouds and meshes). There are extensive approaches based on regular / structured representation, such as multi-view convolutional neural networks (CNNs) [32, 26, 10] and 3D volumetric / grid CNN methods and its variants [38, 26, 28, 35, 36, 16, 9]. These methods can be conveniently developed and implemented in 3D data structure, but they easily suffer from the heavy computation and large memory expense. So it is better to define the deep learning computations based on 3D shapes directly, i.e., irregular / unstructured representation, such as point cloud based methods [25, 27, 13, 30, 3, 18, 19, 33, 17, 42, 34, 8, 40]. However, defining the convolution on the irregular / unstructured representation of 3D objects is not an easy task. Very few methods on point clouds have defined an effective and efficient convolution on each point. Meanwhile, several approaches have been proposed to develop convolutional networks on 2D manifolds [21, 4, 23, 39]. Their representations (e.g., 3D surface meshes) have point positions as well as connectivities, which makes it relatively easier to define the convolution operator on them.

In this work, we present a new method to define and compute convolutions directly on 3D point clouds effectively and efficiently by the proposed *annular convolutions*. This new convolution operator can better capture local neighborhood geometry of each point by specifying the (regular and dilated) ring-shaped structures and directions in the computation. It can adapt to the geometric variabil-

ity and scalability at the signal processing level. Then, we apply it along with the developed hierarchical neural networks to object classification, part segmentation, and semantic segmentation in large-scale scene as shown in Fig. 1. The key *contributions* of our work are as follows:

- We propose a new approach to define convolutions on point cloud. The proposed *annular convolutions* can define arbitrary kernel sizes on each local ring-shaped region, and help to capture better geometric representations of 3D shapes;

- We propose a new multi-level hierarchical method based on *dilated rings*, which leads to better capturing and abstracting shape geometric details. The new dilated strategy on point clouds benefits our proposed closed-loop convolutions and poolings;

- Our proposed network models present new state-of-the-art performance on object classification, part segmentation, and semantic segmentation of large-scale scenes using a variety of standard benchmark datasets.

2. Related Work

Due to the scope of our work, we focus only on recently related deep learning methods, which are proposed on different 3D shape representations.

Volumetric Methods. One traditional way to analyze a 3D shape is to convert it into the regular volumetric occupancy grid and then apply 3D CNNs [38, 26]. The major limitation of these approaches is that 3D convolutions are more expensive in computations than 2D cases. In order to make the computation affordable, the volume grid size is usually in a low resolution. However, lower resolution means loosing some shape geometric information, especially in analyzing large-scale 3D shapes / scenes. To overcome these problems, octree-based methods [28, 35, 36] have been proposed to allow applying 3D CNNs on higher / adaptive resolution grids. PointGrid [16] is a 3D CNN that incorporates a constant number of points within each grid cell and allows it to learn better local geometric details. Similarly, Hua et al. [9] presented a 3D convolution operator based on a uniform grid kernel for semantic segmentation and object recognition on point clouds.

Point Cloud based Methods. PointNet [25] is the first attempt of applying deep learning directly on point clouds. PointNet model is invariant to the order of points, but it considers each point independently without including local region information. PointNet++ [27] is a hierarchical extension of PointNet model and learns local structures of point clouds at different scales. But [27] still considers every point in its local region independently. In our work, we address the aforementioned issues by defining the convolution operator that learns the relationship between neighboring points in a local region, which helps to better capture the local geometric properties of the 3D object.

Klokov et al. [13] proposed a new deep learning architecture called Kd-networks, which uses kd-tree structure to construct a computational graph on point clouds. KC-Net [30] improves PointNet model by considering the local neighborhood information. It defines a set of learnable point-set kernels for local neighboring points and presents a pooling method based on a nearest-neighbor graph. PCNN [3] is another method to apply convolutional neural networks to point clouds by defining extension and restriction operators, and mapping point cloud functions to volumetric functions. SO-Net [17] is a permutation invariant network that utilizes spatial distribution of point clouds by building a self-organizing map. There are also some spectral convolution methods on point clouds, such as SyncSpecCNN [42] and spectral graph convolution [34]. Point2Sequence [19] learns the correlation of different areas in a local region by using attention mechanism, but it does not propose a convolution on point clouds. PointCNN [18] is a different method that proposes to transform neighboring points to the canonical order and then apply convolution.

Recently, there are several approaches proposed to process and analyze large-scale point clouds from indoor and outdoor environments. Engelmann et al. [8] extended PointNet model to exploit the large-scale spatial context. Ye et al. [40] proposed a pointwise pyramid pooling to aggregate features at local neighborhoods as well as two-directional hierarchical recurrent neural networks (RNNs) to learn spatial contexts. However, these methods do not define convolutions on large-scale point clouds to learn geometric features in the local neighborhoods. TangentConv [33] is another method that defines the convolution on point clouds by projecting the neighboring points on tangent planes and applying 2D convolutions on them. The orientation of the tangent image is estimated according to the local point / shape curvature, but as we know the curvature computation on the local region of the point clouds is not stable and not robust (see the discussion in Sec. 3.4), which makes it orientation-dependent. Instead, our method proposes an annular convolution, which is invariant to the orientations of local patches. Also, ours does not require additional input features while theirs needs such features (e.g., depth, height, etc.).

Mesh based Methods. Besides point cloud based methods, several approaches have been proposed to develop convolutional networks on 3D meshes for shape analysis. Geodesic CNN [21] is an extension of the Euclidean CNNs to non-Euclidean domains and is based on a local geodesic system of polar coordinates to extract local patches. Anisotropic CNN [4] is another generalization of Euclidean CNNs to non-Euclidean domains, where classical convolutions are replaced by projections over a set of oriented anisotropic diffusion kernels. Mixture Model Networks (MoNet) [23] generalizes deep learning methods to non-Euclidean domains (graphs and manifolds) by com-

bining previous methods, e.g., classical Euclidean CNN, Geodesic CNN, and Anisotropic CNN. MoNet proposes a new type of kernel in parametric construction. Directionally Convolutional Networks (DCN) [39] applies convolution operation on the triangular mesh of 3D shapes to address part segmentation problem by combining local and global features. Lastly, Surface Networks [14] propose upgrades to Graph Neural Networks to leverage extrinsic differential geometry properties of 3D surfaces for increasing their modeling power.

3. Method

In this work, we propose a new end-to-end framework named as annularly convolutional neural networks (A-CNN) that leverages the neighborhood information to better capture local geometric features of 3D point clouds. In this section, we introduce main technique components of the A-CNN model on point clouds that include: regular and dilated rings, constraint-based k-nearest neighbors (k-NN) search, ordering neighbors, annular convolution, and pooling on rings.

3.1. Regular and Dilated Rings on Point Clouds

To extract local spatial context of the 3D shape, Point-Net++ [27] proposes multi-scale architecture. The major limitation of this approach is that multiple scaled regions may have overlaps (i.e., same neighboring points could be duplicately included in different scaled regions), which reduces the performance of the computational architecture. Overlapped points at different scales lead to redundant information at the local region, which limits a network to learn more discriminative features.

In order to address the above issue, our proposed framework is aimed to leverage a neighborhood at different scales more wisely. We propose two ring-based schemes, i.e., *regular rings* and *dilated rings*. Comparing to multi-scale strategy, the ring-based structure does not have overlaps (no duplicated neighboring points) at the query point's neighborhood, so that each ring contains its own unique points, as illustrated in Sec. 1 of Supplementary Material.

The difference between regular rings and dilated rings is that dilated rings have empty space between rings. The idea of proposed dilated rings is inspired by dilated convolutions on image processing [43], which benefits from aggregating multi-scale contextual information. Although each ring may define the same number of computation / operation parameters (e.g., number of neighboring points), the coverage area of each ring is different (i.e., dilated rings will have larger coverage than the regular rings) as depicted in Fig. 2. Regular rings can be considered as dilated rings with the dilation factor equal to 0.

The proposed regular rings and dilated rings will contribute to neighboring point search, convolution, and pool-


(a) Regular Rings   (b) Dilated Rings

Figure 2: The comparison of the regular and dilated ring-shaped structures (such as with two rings). We can see that comparing two sectors (e.g., black solid points) in the regular and dilated rings, the dilated rings cover larger space by using the same number of neighbors as in regular rings. Moreover, each ring contains unique neighboring points comparing to the other ring.

ing in the follow-up processes. First, for k-NN algorithm, we constrain search areas in the local ring-shaped neighborhood to ensure no overlap. Second, the convolutions defined on rings cover larger areas with the same kernel sizes without increasing the number of convolution parameters. Third, the regular / dilated ring architectures will help to aggregate more discriminative features after applying max-pooling at each ring of the local region. We will discuss them in more detail in the following subsections.

To justify the aforementioned statements, we will compare multi-scale approach with our proposed multi-ring scheme on object classification task in the ablation study (Sec. 5.4). The results show that ring-based structure captures better local geometric features than previous multi-scale method, since it achieves higher accuracy.

3.2. Constraint-based K-NN Search

In the original PointNet++ model, the ball query algorithm returns the first K neighbors found inside a search ball specified by a radius R and query point $\mathbf{q}_i$, so that it cannot guarantee that the closest points will always be found. However, our proposed k-NN search algorithm guarantees returning closest points inside the searching area by using the Euclidean metric. Each ring is defined by two parameters: the inner radius R_{inner} and the outer radius R_{outer} (in Fig. 2); therefore, the *constraint-based k-NN search* ensures that the closest and unique points will be found in each ring.

3.3. Ordering Neighbors

In order to learn relationships between neighboring points in a local regions, we need first to order points in a clockwise / counterclockwise manner and then apply annular convolutions. Our proposed ordering operator consists of two main steps: projection and ordering. The importance of the projection before ordering is that the dot product has its restriction in ordering points. By projecting points on

a tangent plane at a query point $\mathbf{q}_i$, we effectively order neighbors in clockwise / counterclockwise direction by taking use of cross product and dot product together. The detailed explanations of normal estimation, orthogonal projection, and ordering are given in the following subsections.

3.3.1 Normal Estimation on Point Clouds

Normal is an important geometric property of a 3D shape. We use it as a tool for projecting and ordering neighboring points at a local domain. The simplest normal estimation method approximates the normal $\mathbf{n}_i$ at the given point $\mathbf{q}_i$ by calculating the normal of the local tangent plane $\mathcal{T}_i$ at that point, which becomes a least-square plane fitting estimation problem [29]. To calculate normal $\mathbf{n}_i$, one needs to compute eigenvalues and eigenvectors of the covariance matrix $\mathbf{C}$ as:

$$\mathbf{C} = \frac{1}{K} \sum_{j=1}^{K} (\mathbf{x}_j - \mathbf{q}_i) \cdot (\mathbf{x}_j - \mathbf{q}_i)^T, \tag{1}$$

$$\mathbf{C} \cdot \mathbf{v}_\gamma = \lambda_\gamma \cdot \mathbf{v}_\gamma, \gamma \in \{0, 1, 2\},$$

where K is the number of neighboring points $\mathbf{x}_j$s around query point $\mathbf{q}_i$ (e.g., $K = 10$ in our experiments), λ_γ and $\mathbf{v}_\gamma$ are the γth eigenvalue and eigenvector of the covariance matrix $\mathbf{C}$, respectively. The covariance matrix $\mathbf{C}$ is symmetric and positive semi-definite. The eigenvectors $\mathbf{v}_\gamma$ form an orthogonal frame, in respect to the local tangent plane $\mathcal{T}_i$. The eigenvector $\mathbf{v}_0$ that corresponds to the smallest eigenvalue λ_0 is the estimated normal $\mathbf{n}_i$.

3.3.2 Orthogonal Projection

After extracting neighbors $\mathbf{x}_j, j \in \{1, ..., K\}$ for a query point $\mathbf{q}_i$, we calculate projections $\mathbf{p}_j$s of these points on a tangent plane $\mathcal{T}_i$ described by a unit normal $\mathbf{n}_i$ (estimated in Sec. 3.3.1) as:

$$\mathbf{p}_j = \mathbf{x}_j - ((\mathbf{x}_j - \mathbf{q}_i) \cdot \mathbf{n}_i) \cdot \mathbf{n}_i, \quad j \in \{1, ..., K\}. \tag{2}$$

Fig. 3 (a) illustrates the orthogonal projection of neighboring points on a ring.

3.3.3 Counterclockwise Ordering

Firstly, we use the geometric definition of the dot product to compute the angle between two vectors $\mathbf{c}$ (i.e., starts from the query point $\mathbf{q}_i$ and connects with a randomly starting point, such as $\mathbf{p}_1$) and $\mathbf{p}_j - \mathbf{q}_i$ (i.e., starts from the query point $\mathbf{q}_i$ and connects with other neighboring points $\mathbf{p}_j$):

$$cos(\theta_{\mathbf{p}_j}) = \frac{\mathbf{c} \cdot (\mathbf{p}_j - \mathbf{q}_i)}{||\mathbf{c}|| ||\mathbf{p}_j - \mathbf{q}_i||}. \tag{3}$$

We know that $cos(\theta_{\mathbf{p}_j})$ lies in $[-1, 1]$, which corresponds to angles between $[0°, 180°]$. In order to sort the neighboring points around the query point between


Figure 3: The illustration of the proposed annular convolution on a ring. (a) Projection: $\mathbf{q}_i$ is a query point. After applying the constraint-based k-NN search, neighboring points $\mathbf{X} = \{\mathbf{x}_j | j = 1, ..., K\}$ are extracted on a ring. Given the normal $\mathbf{n}_i$ at query point $\mathbf{q}_i$, we project the searched points on the tangent plane $\mathcal{T}_i$. (b) Counterclockwise Ordering: After projection, we randomly pick a starting point as our reference direction $\mathbf{c}$ and order points in counterclockwise. It is worth mentioning that we order original points $[\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_j, ..., \mathbf{x}_K]$ based on their projections. (c) Annular Convolution: Depending on the kernel size, we copy several original points from the beginning position and concatenate them to the end of the ordered points. Finally, we apply annular convolution with the given kernel.

$[0°, 360°)$, we must to decide which semicircle the considered point $\mathbf{p}_j$ belongs to as follows:

$$sign_{\mathbf{p}_j} = (\mathbf{c} \times (\mathbf{p}_j - \mathbf{q}_i)) \cdot \mathbf{n}_i, \tag{4}$$

where $sign_{\mathbf{p}_j} \geq 0$ is $\theta_{\mathbf{p}_j} \in [0°, 180°]$, and $sign_{\mathbf{p}_j} < 0$ is $\theta_{\mathbf{p}_j} \in (180°, 360°)$.

Then, we can recompute the cosine value of the angle as:

$$\angle_{\mathbf{p}_j} = \begin{cases} -cos(\theta_{\mathbf{p}_j}) - 2 & sign_{\mathbf{p}_j} < 0 \\ cos(\theta_{\mathbf{p}_j}) & sign_{\mathbf{p}_j} \geq 0. \end{cases} \tag{5}$$

Now the values of the angles lie in $(-3, 1]$, which maps angles between $[0°, 360°)$.

Finally, we sort neighboring points $\mathbf{x}_j$ by descending the value of $\angle_{\mathbf{p}_j}$ to obtain the counterclockwise order. Fig. 3 (b) illustrates the process of ordering in a local neighborhood. The neighboring points can be ordered in the clockwise manner, if we sort neighboring points $\mathbf{x}_j$ by ascending the value of $\angle_{\mathbf{p}_j}$.

Our experiments show in Sec. 5.4 that ordering points in the local regions is an important step in our framework and our model achieves better classification accuracy with ordered points than without ordering them.

3.4. Annular Convolution on Rings

Through the previous computation, we have the ordered neighbors represented as an array $[\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_K]$. In or-

der to develop the *annular convolution*, we need to loop the array of neighbors with respect to the size of the kernel (e.g., 1×3, 1×5, ...) on each ring. For example, if the convolutional kernel size is 1×3, we need to take the first two neighbors and concatenate them with the ending elements in the original array to construct a new circular array $[\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_K, \mathbf{x}_1, \mathbf{x}_2]$. Then, we can perform the standard convolutions on this array as shown in Fig. 3 (c).

There are some nice properties of the proposed annular convolutions as follows: (1) The annular convolution is invariant to the orientation of the local patch. That is because the neighbors are organized and ordered in a closed loop in each ring by concatenating the beginning with the end of the neighboring points' sequence. Therefore, we can order neighbors based on any random starting position, which does not negatively affect the convolution results. Compared with some previous convolutions defined on 3D shapes [4, 39, 33], they all need to compute the real principal curvature direction as the reference direction to define the local patch operator, which is not robust and cumbersome. In particular, 3D shapes have large areas of flat and spherical regions, where the curvature directions are arbitrary. (2) As we know, in reality, the normal direction flipping issues are widely existing in point clouds, especially the large-scale scene datasets. Under the annular convolution strategy, no matter the neighboring points are ordered in clockwise or counterclockwise manner, the results are the same. (3) Another advantage of annular convolution is that we can define an arbitrary kernel size, instead of just 1×1 kernels [25, 27]. Therefore, the annular convolution can provide the ability to learn the relationship between ordered points inside each ring as shown in Fig. 3 (c).

Annular convolutions can be applied on both regular and dilated rings. By applying annular convolutions with the same kernel size on different rings, we can cover and convolve larger areas by using the dilated structure, which helps us to learn larger spatial contextual information in the local regions. The importance of annular convolutions is shown in the ablation study in Sec. 5.4.

3.5. Pooling on Rings

After applying a set of annular convolutions sequentially, the resulting convolved features encode information about its closest neighbors in each ring as well as spatial remoteness from a query point. Then we aggregate the convolved features across all neighbors on each ring separately. We apply the max-pooling strategy in our framework. Our proposed ring-based scheme allows us to aggregate more discriminative features. The extracted max-pooled features contain the encoded information about neighbors and the relationship between them in the local region, unlike the pooling scheme in PointNet++ [27], where each neighbor is considered independently from its neighbors. In our pooling process, the non-overlapped regions (rings) will aggregate

different types of features in each ring, which can uniquely describe each local region (ring) around the query point. The multi-scale approach in PointNet++ does not guarantee this and might aggregate the same features at different scales, which is redundant information for a network. The (regular and dilated) ring-based scheme helps to avoid extracting duplicate information but rather promotes extracting multi-level information from different regions (rings). This provides a network with more diverse features to learn from. After aggregating features at different rings, we concatenate and feed them to another abstract layer to further learn hierarchical features.

4. A-CNN Architecture

Our proposed A-CNN model follows a design where the hierarchical structure is composed of a set of abstract layers. Each abstract layer consists of several operations performed sequentially and produces a subset of input points with newly learned features. Firstly, we subsample points by using Farthest Point Sampling (FPS) algorithm [22] to extract centroids randomly distributed on the surface of each object. Secondly, our constraint-based k-NN extracts neighbors of a centroid for each local region (i.e., regular / dilated rings) and then we order neighbors in a counterclockwise manner using projection. Finally, we apply sequentially a set of annular convolutions on the ordered points and maxpool features across neighbors to produce new feature vectors, which uniquely describe each local region.

Given the point clouds of 3D shapes, our proposed end-to-end network is able to classify and segment the objects. In the following, we discuss the classification and segmentation network architectures on 3D point clouds.

4.1. Classification Network

The classification network is illustrated at the top of Fig. 4. It consists of two major parts: encoder and classification. The encoder extracts features from each ring independently inside every layer and concatenates them at the end to process further to extract high-level features. The proposed architecture includes both *regular rings* and *dilated rings*. We end up using two rings per layer, because it gives us pretty good experimental results as shown in the Sec. 5. It can be easily extended to more than two rings per layer, if necessary.

We use regular rings in the first layer and dilated rings in the second layer in the encoder. Annular convolutions with the kernel sizes 1×3 and stride 1 are applied in the first two layers, followed by a batch normalization [12] (BN) and a rectified linear unit [24] (ReLU). Different rings of the same query point are processed in parallel. Then, the aggregated features from each ring concatenate together to propagate to the next layer. The last layer in the encoder performs convolutions with kernel sizes 1×1 followed by BN and ReLU layers, where only spatial positions of the sampled points
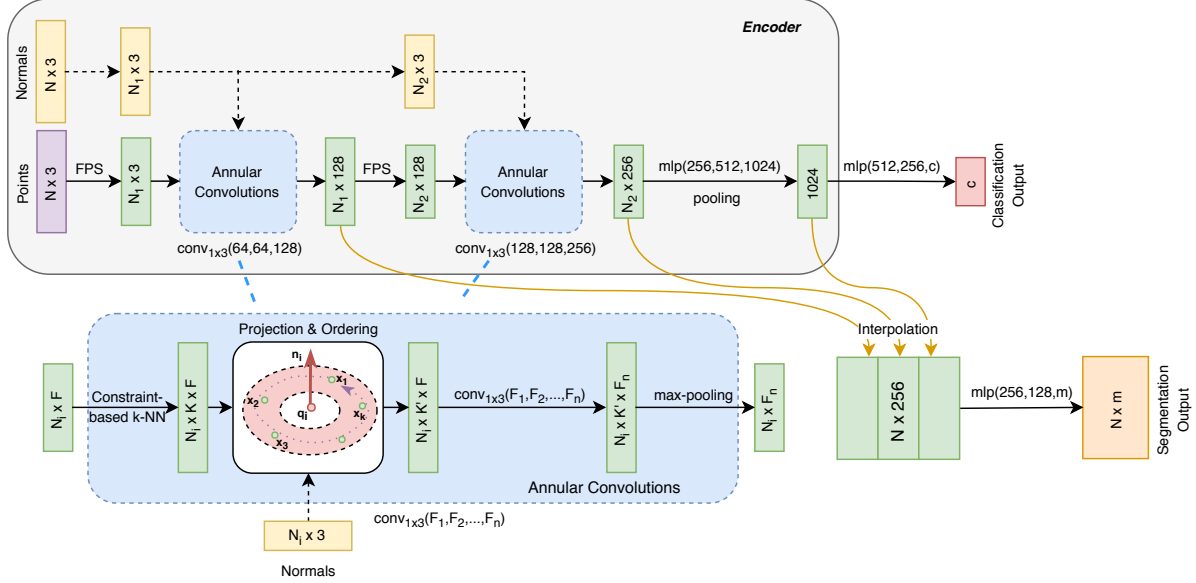
Figure 4: The architecture of A-CNN. Both classification and segmentation networks share encoder part for the feature extraction. Normals are used only to determine the order of neighboring points in the local regions (dashed arrows mean no backpropagation during training) and not used as additional features, unless it is mentioned explicitly in the experiments. N, N_1, N_2 (where $N > N_1 > N_2$) are the numbers of points as input, after the first and second layer, respectively. K and K' are the unordered and ordered points inside the local rings, respectively. c is the number of classification classes. m is the number of segmentation classes. "FPS" stands for Farthest Point Sampling algorithm. "mlp" stands for multi-layer perceptron. $conv_{1\times3}(F_1, F_2, ..., F_n)$ stands for annular convolutions with the kernel size 1×3 applied sequentially with corresponding feature map sizes $F_i, i \in 1, ..., n$.

are considered. After that aggregated high-level features are fed to the set of fully-connected layers with integrated dropout [31] and ReLU layers to calculate probability of each class. The output size of the classification network is equal to the number of classes in the dataset.

4.2. Segmentation Network

The segmentation network shares encoder part with the classification network as shown in Figure 4. In order to predict the segmentation label per point, we need to upsample the sampled points in the encoder back to the original point cloud size. As pointed out by [44], the consecutive feature propagation proposed by [27] is not the most efficient approach. Inspired from [44], we propagate features from different levels from the encoder directly to the original point cloud size, and concatenate them by allowing the network to learn the most important features from different levels as well as to learn the relationship between them.

The output of each level has different sizes due to the hierarchical feature extractions, so we have to restore hierarchical features from each level back to the original point size by using an interpolation method [27]. The interpolation method is based on the inverse squared Euclidean distance weighted average of the three nearest neighbors as:

$$f^{(l+1)}(\mathbf{x}) = \sum_{j=1}^{3} f^{(l)}(\mathbf{x}_j) \frac{w_j(\mathbf{x})}{\sum_{j=1}^{3} w_j(\mathbf{x})}, \qquad (6)$$

where $w_j(\mathbf{x}) = \frac{1}{d(\mathbf{x},\mathbf{x}_j)^2}$ is an inverse squared Euclidean

distance weight.

Then, we concatenate upsampled features from different levels and pass them through 1×1 convolution to reduce feature space and learn the relationship between features from different levels. Finally, the segmentation class distribution for each point is calculated.

5. Experiments

We evaluate our A-CNN model on various tasks such as point cloud classification, part segmentation, and large-scale scene segmentation. In the following subsections, we demonstrate more details on each task. It is noted that for the comparison experiments, best results in the tables are shown in bold font.

All models in this paper are trained on a single NVIDIA Titan Xp GPU with 12 GB GDDR5X. The training time of our model is faster than that of PointNet++ model. More details about the network configurations, training settings and timings in our experiments can be found in Sec. 2 and Tab. 2 of Supplementary Material. The source code of the framework will be made available later.

5.1. Point Cloud Classification

We evaluate our classification model on two datasets: *ModelNet10* and *ModelNet40* [38]. ModelNet is a large-scale 3D CAD model dataset. *ModelNet10* is a subset of ModelNet dataset that consists of 10 different classes with 3991 training and 908 testing objects. *ModelNet40* includes

Table 1: Classification results on *ModelNet10* and *ModelNet40* datasets. AAC is accuracy average class, OA is overall accuracy.

	ModelNet10		ModelNet40	
	AAC	OA	AAC	OA
different methods with additional input or more points				
AO-CNN [36]	-	-	-	90.5
O-CNN [35]	-	-	-	90.6
PointNet++ [27]	-	-	-	91.9
SO-Net [17]	95.5	95.7	90.8	93.4
MVCNN-MultiRes [26]	-	-	91.4	93.8
VRN Ensemble [5]	-	97.1	-	95.5
point cloud based methods with 1024 points				
PointNet [25]	-	-	86.2	89.2
Kd-Net (depth 15) [13]	93.5	94.0	88.5	91.8
Pointwise CNN [9]	-	-	81.4	86.1
KCNet [30]	-	94.4	-	91.0
PointGrid [16]	-	-	88.9	92.0
PCNN [3]	-	94.9	-	92.3
PointCNN [18]	-	-	88.1	92.2
Point2Sequence [19]	95.1	95.3	**90.4**	**92.6**
A-CNN (our)	**95.3**	**95.5**	90.3	**92.6**

40 different classes with 9843 objects for training and 2468 objects for testing. Point clouds with 10,000 points and normals are sampled from meshes, normalized into a unit sphere, and provided by [27].

For experiments on *ModelNet10* and *ModelNet40*, we sample 1024 points with normals, where normals are only used to order points in the local region. For data augmentation, we randomly scale object sizes, shift object positions, and perturb point locations. For better generalization, we apply point shuffling in order to generate different centroids for the same object at different epochs.

In Tab. 1, we compare our method with several state-of-the-art methods in the shape classification results on both *ModelNet10* and *ModelNet40* datasets. Our model achieves better accuracy among the point cloud based methods (with 1024 points), such as PointNet [25], PointNet++ [27] (5K points + normals), Kd-Net (depth 15) [13], Pointwise CNN [9], KCNet [30], PointGrid [16], PCNN [3], and PointCNN [18]. Our model is slightly better than Point2Sequence [19] on *ModelNet10* and shows comparable performance on *ModelNet40*.

Meanwhile, our model performs better than other volumetric approaches, such as O-CNN [35] and AO-CNN [36]; while we are a little worse than SO-Net [17], which uses denser input points, i.e., 5000 points with normals as the input (1024 points in our A-CNN); MVCNN-MultiRes [26], which uses multi-view 3D volumes to represent an object (i.e., 20 views of $30 \times 30 \times 30$ volume); the VRN Ensemble [5], which involves an ensemble of six models.

We also provide some feature visualization results in Sec. 3 of Supplementary Material, including global feature (e.g., t-SNE clustering) visualization and local feature (e.g., the magnitude of the gradient per point) visualization.
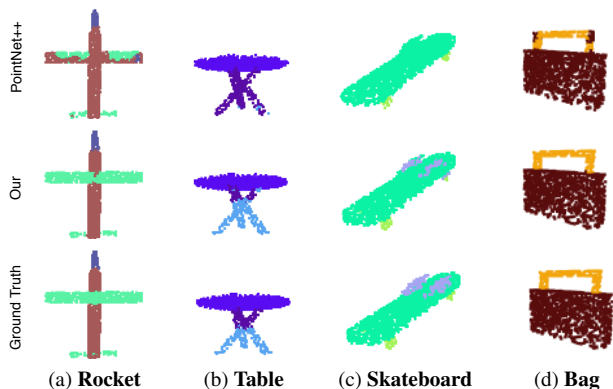

Figure 5: Qualitative results on *ShapeNet-part* dataset. We compare our results with PointNet++ [27] and ground truth.

5.2. Point Cloud Segmentation

We evaluate our segmentation model on *ShapeNet-part* [41] dataset. The dataset contains 16,881 shapes from 16 different categories with 50 label parts in total. The main challenge of this dataset is that all categories are highly imbalanced. There are 2048 points sampled for each shape from the dataset, where most shapes contain less than six parts. We follow the same training and testing splits provided in [25, 41]. For data augmentation, we perturb point locations with the point shuffling for better generalization.

We evaluate our segmentation model with two different inputs. One of the models is trained without feeding normals as additional features and the other model is trained with normals as additional features. The quantitative results are provided in Tab. 2, where mean IoU (Intersection-over-Union) is reported. The qualitative results are visualized in Fig. 5. Our approach with point locations only as input outperforms PointNet [25], Kd-Net [13], KCNet [30], and PCNN [3]; and shows slightly worse performance comparing to PointGrid [16] (volumetric method) and PointCNN [18]. Meanwhile, our model achieves the best performance with the input of point locations and normals, compared with PointNet++ [27], SyncSpecCNN [42], SO-Net [17], SGPN [37], O-CNN [35], RSNet [11], and Point2Sequence [19]. The more detailed quantitative results (e.g., per-category IoUs) and more visualization results are provided in Sec. 5 of Supplementary Material.

5.3. Semantic Segmentation in Scenes

We also evaluate our segmentation model on two large-scale indoor datasets *Stanford 3D Large-Scale Indoor Spaces (S3DIS)* [2] and *ScanNet* [7]. *S3DIS* contains 6 large-scale indoor areas with 271 rooms sampled from 3 different buildings, where each point has the semantic label that belongs to one of the 13 categories. *ScanNet* includes 1513 scanned indoor point clouds, where each voxel has been labeled with one of the 21 categories.

We employ the same training and testing strategies as

| Ceiling | Floor | Wall | Beam | Column | Window | Door | Table | Chair | Sofa | Bookcase | Board | Clutter |

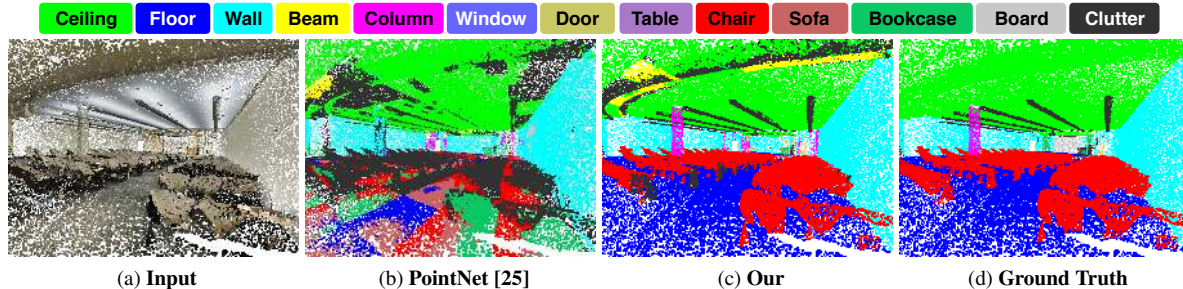| (a) **Input** | (b) **PointNet [25]** | (c) **Our** | (d) **Ground Truth** |

Figure 6: Qualitative results on *S3DIS* dataset. We compare our results with PointNet [25] and ground truth. The auditorium is a challenging room type and appears only in Area 2. Our model produces much better segmentation result, compared with the result of PointNet.

PointNet [25] on *S3DIS*, where we use 6-fold cross validation over all six areas. The evaluation results are reported in Tab. 2, and qualitative results are visualized in Fig. 6. Our model demonstrates better segmentation results compared with PointNet [25], MS+CU (2) [8], G+RCU [8], 3P-RNN [40], SPGraph [15], and TangentConv [33]. However, our model performs slightly worse than PointCNN [18] due to their non-overlapping block sampling strategy with paddings which we do not use. Meanwhile, our approach shows the best segmentation results on *ScanNet* [7] and achieves the state-of-the-art performance, compared with PointNet [25], PointNet++ [27], TangentConv [33], and PointCNN [18] according to Tab. 2.

More qualitative visualization results and data preparation details on both datasets are provided in Sec. 4 and Sec. 5, respectively, of Supplementary Material and Video.

Table 2: Segmentation results on *ShapeNet-part*, *S3DIS*, and *ScanNet*. "mean" is mean IoU (%), OA is overall accuracy.

	ShapeNet-part		*S3DIS*	*ScanNet*
	without normals	*with normals*	OA	OA
	mean	mean		
PointNet [25]	83.7	-	78.5	73.9
PointNet++ [27]	-	85.1	-	84.5
SyncSpecCNN [42]	-	84.7	-	-
O-CNN [35]	-	85.9	-	-
Kd-Net [13]	82.3	-	-	-
KCNet [30]	84.7	-	-	-
SO-Net [17]	-	84.9	-	-
SGPN [37]	-	85.8	-	-
MS+CU (2) [8]	-	-	79.2	-
G+RCU [8]	-	-	81.1	-
RSNet [11]	-	84.9	-	-
3P-RNN [40]	-	-	86.9	-
SPGraph [15]	-	-	85.5	-
TangentConv [33]	-	-	*	80.9
PCNN [3]	85.1	-	-	-
Point2Sequence [19]	-	85.2	-	-
PointGrid [16]	**86.4**	-	-	-
PointCNN [18]	86.1	-	**88.1**	85.1
A-CNN (our)	85.9	**86.1**	87.3	**85.4**

Note: * TangentConv [33] OA on *S3DIS* Area 5 is 82.5% (as reported in their paper), which is worse compared with our OA of 85.5%.

5.4. Ablation Study

The goal of our ablation study is to show the importance of the proposed technique components (in Sec. 3) in our A-CNN model. We evaluate three proposed components, such as rings without overlaps (Sec. 3.1), ordering (Sec. 3.3), and

annular convolution (Sec. 3.4) on the classification task of *ModelNet40* dataset as shown in Tab. 3. In the first experiment, we replace our proposed constraint-based k-NN on ring regions with ball query in [27], but keep ordering and annular convolutions on. In the second and third experiments, we turn off either annular convolutions or ordering, respectively; and keep the rest two components on. Our experiments show that the proposed ring-shaped scheme contributes the most to our model. It is because multi-level rings positively affect annular convolutions. Finally, A-CNN model with all three components (i.e., rings without overlaps, ordering, and annular convolutions) achieves the best results. We also discover that reducing overlap / redundancy in multi-scale scheme can improve existing methods. We evaluate the original PointNet++ [27] with and without overlap as shown in Sec. 1 of Supplementary Material.

Table 3: Ablation experiments on *ModelNet40* dataset. AAC is accuracy average class, OA is overall accuracy.

	AAC	OA
A-CNN (without rings / with overlap)	89.2	91.7
A-CNN (without annular conv.)	89.2	91.8
A-CNN (without ordering)	89.6	92.0
A-CNN (with all components)	**90.3**	**92.6**

6. Conclusion

In this work, we propose a new A-CNN framework on point clouds, which can better capture local geometric information of 3D shapes. Through extensive experiments on several benchmark datasets, our method has achieved the state-of-the-art performance on point cloud classification, part segmentation, and large-scale semantic segmentation tasks. Since our work does not solely focus on large-scale scene datasets, we will explore some new deep learning architectures to improve the current results. We will also investigate to apply the proposed framework on large-scale outdoor datasets in our future work.

References

[1] E. Ahmed, A. Saint, A. Shabayek, K. Cherenkova, R. Das, G. Gusev, D. Aouada, and B. Ottersten. Deep learning advances on different 3D data representations: A survey. *arXiv preprint arXiv:1808.01462*, 2018.

[2] I. Armeni, O. Sener, A. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. 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.

[3] M. Atzmon, H. Maron, and Y. Lipman. Point convolutional neural networks by extension operators. *ACM Transactions on Graphics*, 37(4):71:1–71:12, 2018.

[4] D. Boscaini, J. Masci, E. Rodolà, and M. Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 3189–3197, 2016.

[5] A. Brock, T. Lim, J. Ritchie, and N. Weston. Generative and discriminative voxel modeling with convolutional neural networks. *arXiv preprint arXiv:1608.04236*, 2016.

[6] S. Bu, P. Han, Z. Liu, J. Han, and H. Lin. Local deep feature learning framework for 3D shape. *Computers & Graphics*, 46:117–129, 2015.

[7] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017.

[8] F. Engelmann, T. Kontogianni, A. Hermans, and B. Leibe. Exploring spatial context for 3D semantic segmentation of point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 716–724, 2017.

[9] B.-S. Hua, M.-K. Tran, and S.-K. Yeung. Pointwise convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 984–993, 2018.

[10] H. Huang, E. Kalogerakis, S. Chaudhuri, D. Ceylan, V. G. Kim, and E. Yumer. Learning local shape descriptors from part correspondences with multiview convolutional networks. *ACM Transactions on Graphics*, 37(1):6, 2018.

[11] Q. Huang, W. Wang, and U. Neumann. Recurrent slice networks for 3D segmentation of point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2626–2635, 2018.

[12] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.

[13] R. Klokov and V. Lempitsky. Escape from cells: Deep Kd-networks for the recognition of 3D point cloud models. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 863–872, 2017.

[14] I. Kostrikov, Z. Jiang, D. Panozzo, D. Zorin, and J. Bruna. Surface networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2540–2548, 2018.

[15] L. Landrieu and M. Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4558–4567, 2018.

[16] T. Le and Y. Duan. PointGrid: A deep network for 3D shape understanding. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9204–9214, 2018.

[17] J. Li, B. M. Chen, and G. H. Lee. SO-Net: Self-organizing network for point cloud analysis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9397–9406, 2018.

[18] Y. Li, R. Bu, M. Sun, W. Wu, X. Di, and B. Chen. PointCNN: Convolution on X-transformed points. In *Advances in Neural Information Processing Systems*, pages 828–838, 2018.

[19] X. Liu, Z. Han, Y.-S. Liu, and M. Zwicker. Point2Sequence: Learning the shape representation of 3D point clouds with an attention-based sequence to sequence network. In *Association for the Advancement of Artificial Intelligence*, 2019.

[20] Z. Liu, S. Chen, S. Bu, and K. Li. High-level semantic feature for 3D shape based on deep belief networks. In *Proceedings of the IEEE International Conference on Multimedia and Expo*, pages 1–6, 2014.

[21] J. Masci, D. Boscaini, M. Bronstein, and P. Vandergheynst. Geodesic convolutional neural networks on Riemannian manifolds. In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, pages 37–45, 2015.

[22] C. Moenning and N. A. Dodgson. Fast marching farthest point sampling. Technical report, University of Cambridge, Computer Laboratory, 2003.

[23] F. Monti, D. Boscaini, J. Masci, E. Rodola, J. Svoboda, and M. Bronstein. Geometric deep learning on graphs and manifolds using mixture model CNNs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5115–1543, 2017.

[24] V. Nair and G. Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the International Conference on Machine Learning*, pages 807–814, 2010.

[25] C. Qi, H. Su, K. Mo, and L. Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660, 2017.

[26] C. Qi, H. Su, M. Nießner, A. Dai, M. Yan, and L. Guibas. Volumetric and multi-view CNNs for object classification on 3D data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5648–5656, 2016.

[27] C. Qi, L. Yi, H. Su, and L. Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, pages 5105–5114, 2017.

[28] G. Riegler, A. Ulusoy, and A. Geiger. OctNet: Learning deep 3D representations at high resolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3577–3586, 2017.

[29] R. Rusu. *Semantic 3D Object Maps for Everyday Manipulation in Human Living Environments*. PhD thesis, Computer

Science department, Technische Universitaet Muenchen, Germany, October 2009.

[30] Y. Shen, C. Feng, Y. Yang, and D. Tian. Mining point cloud local structures by kernel correlation and graph pooling. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4548–4557, 2018.

[31] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research*, 15:1929–1958, 2014.

[32] H. Su, S. Maji, E. Kalogerakis, and E. 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.

[33] M. Tatarchenko, J. Park, V. Koltun, and Q.-Y. Zhou. Tangent convolutions for dense prediction in 3D. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3887–3896, 2018.

[34] C. Wang, B. Samari, and K. Siddiqi. Local spectral graph convolution for point set feature learning. In *Proceedings of The European Conference on Computer Vision*, September 2018.

[35] P.-S. Wang, Y. Liu, Y.-X. Guo, C.-Y. Sun, and X. Tong. O-CNN: Octree-based convolutional neural networks for 3D shape analysis. *ACM Transactions on Graphics*, 36(4):72, 2017.

[36] P.-S. Wang, C.-Y. Sun, Y. Liu, and X. Tong. Adaptive O-CNN: A patch-based deep representation of 3D shapes. *ACM Transactions on Graphics*, 37(6), 2018.

[37] W. Wang, R. Yu, Q. Huang, and U. Neumann. SGPN: Similarity group proposal network for 3D point cloud instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2569–2578, 2018.

[38] Z. Wu, S. Song, A. Khosla, F. Yu, L. Zhang, X. Tang, and J. 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.

[39] H. Xu, M. Dong, and Z. Zhong. Directionally convolutional networks for 3D shape segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2698–2707, 2017.

[40] X. Ye, J. Li, H. Huang, L. Du, and X. Zhang. 3D recurrent neural networks with context fusion for point cloud semantic segmentation. In *Proceedings of The European Conference on Computer Vision*, September 2018.

[41] L. Yi, V. G. Kim, D. Ceylan, I. Shen, M. Yan, H. Su, C. Lu, Q. Huang, A. Sheffer, L. Guibas, et al. A scalable active framework for region annotation in 3D shape collections. *ACM Transactions on Graphics*, 35(6):210, 2016.

[42] L. Yi, H. Su, X. Guo, and L. Guibas. SyncSpecCNN: Synchronized spectral CNN for 3D shape segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6584–6592, 2017.

[43] F. Yu and V. Koltun. Multi-scale context aggregation by dilated convolutions. In *International Conference on Learning Representations*, 2016.

[44] L. Yu, X. Li, C.-W. Fu, D. Cohen-Or, and P.-A. Heng. PU-Net: Point cloud upsampling network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2790–2799, 2018.