

Deep Positional and Relational Feature Learning for Rotation-Invariant Point Cloud Analysis

Ruixuan Yu^{1,2}, Xin Wei¹, Federico Tombari^{2,3}, and Jian Sun¹

¹ Xi'an Jiaotong University, China

² Technical University of Munich, Germany

³ Google

{yuruixuan123, wxmath}@stu.xjtu.edu.cn,
tombari@google.com, jiansun@xjtu.edu.cn

Abstract. In this paper we propose a rotation-invariant deep network for point clouds analysis. Point-based deep networks are commonly designed to recognize roughly aligned 3D shapes based on point coordinates, but suffer from performance drops with shape rotations. Some geometric features, e.g., distances and angles of points as inputs of network, are rotation-invariant but lose positional information of points. In this work, we propose a novel deep network for point clouds by incorporating positional information of points as inputs while yielding rotation-invariance. The network is hierarchical and relies on two modules: a positional feature embedding block and a relational feature embedding block. Both modules and the whole network are proven to be rotation-invariant when processing point clouds as input. Experiments show state-of-the-art classification and segmentation performances on benchmark datasets, and ablation studies demonstrate effectiveness of the network design.

Keywords: Rotation-invariance, point cloud, deep feature learning.

1 Introduction

Point clouds are widely employed as a popular 3D representation for objects and scenes. They are generated by most current acquisition techniques and 3D sensors, and used within well-studied application fields such as autonomous driving, archaeology, robotics, augmented reality, to name a few. Among these applications, shape recognition and segmentation are two fundamental tasks focusing on automatically recognizing and segmenting 3D objects or object parts [1–3].

The majority of 3D object recognition approaches are currently based on deep learning [1, 2, 4–7]. Most of the point-based methods take positional information, such as point coordinates or normal vectors on aligned 3D objects, as inputs for the network, then learn deep features suitable for the task at hand. These methods now achieve state-of-the-art performance for 3D recognition. PointNet [8] firstly designed a network to process point clouds by taking point coordinates as inputs. Following works such as [1, 2, 9–11] developed various convolution operations on point clouds which brought performance improvements. These

advances justify that designing networks based on the positional information of 3D points on aligned object shapes is an effective way for shape recognition.

Nevertheless, in many scenarios such as, e.g., 6D pose estimation for robotics/AR [12, 13] and CAD models [14] in industrial environments, or analysis of molecules [15, 16] where small scale objects are uncontrolled, a major limitation of above point-based methods is that they tend to be rotation-sensitive, and their performance drops dramatically when tested on shapes under arbitrary rotations in the 3D rotation group $SO(3)$. A remedy for this is to develop robustness against rotations by augmenting training dataset with arbitrary rotations, e.g., the $SO(3)/SO(3)$ mode in Spherical CNN [17]. However, it is not efficient to use data augmentation for learning rotation-invariance, since the enlarged dataset requires a large computational cost. Also, the developed robustness to rotations of the network will be up to the seen augmentations during training.

To deal with rotation-sensitivity,

an alternative way is to represent shapes with geometric features such as distances and angles. Spherical CNN [17] and α^3 SCNN [19] are representative works that take distances and angles as spherical functions. They learn rotation-equivariant features by defining convolutions on spherical functions, and then aggregate features by global average pooling for rotation-invariance. Note that these representations are not rigorously rotation-invariant due to discretization on sphere.

ClusterNet [18] and RConvNet [20] also take rotation-invariant geometric features as inputs and design specific convolutions on local point clouds. However, directly transforming point coordinates to rotation-invariant geometric features may lose positional information of point cloud, which is essential to recognize 3D shapes when they are (roughly) aligned. For example, in Fig. 1, three local surfaces within the circles are flat surfaces belonging to different parts of two shapes. If we represent the center points of these local surfaces by geometric features (e.g., in ClusterNet [18]) agnostic to their positional information, these features may not distinguish different parts of the same shape or different shapes.

To achieve rotation-invariant shape recognition with high accuracy, we propose a novel deep architecture by learning positional and relational features of point cloud. It takes both point coordinates and geometric features as network inputs, and achieves rotation-invariance by designing a *Positional Feature Embedding block (PFE-block)* and a *Relational Feature Embedding block (RFE-block)*, both of which are rotation-invariant. For both shape classification and segmentation, invariant features on point cloud guarantee that classification label of shape and segmentation label of each point are invariant to shape rotation.




Fig. 1. The three circles highlight examples of local planar surfaces. Their center points are similar with geometric features [18], but distinguishable by point coordinates of selected pose as features in our approach. Features of each center point are shown in rows of tables.

It seems to be contradictory to learn rotation-invariant features with point coordinates as input. We design the PFE-block as composition of a pose expander, a pose selector and a positional feature extractor. The PFE-block is proven to be able to produce invariant features for points agnostic to shape rotations. The pose expander maps the shape into a rotation-invariant pose space, then the pose selector selects a unique pose from the space, whose point coordinates are more discriminative than geometric features as shown in Fig. 1. With this selected pose, we extract its positional features by a positional feature extractor. The RFE-block further enhances the deep feature representation of the point cloud with relational convolutions, where the weights are learned based on the relations of neighboring points. This block is also rotation-invariant.

As a summary, we propose two novel rotation-invariant network blocks and a deep hierarchical network jointly learning positional and relational features as shown in Fig. 2. Our network is one of the very few works that achieve rigorous rotation-invariance. For both point cloud shape classification and segmentation, we achieve state-of-the-art performances on commonly used benchmark datasets.

2 Related Work

2.1 Point-based deep learning

Point cloud is a basic representation for 3D shape, and point coordinates are taken as raw features by most point-based networks to carry out tasks on aligned shapes. PointNet [8] is the first effort that takes point coordinates as raw features and embeds positional information of points to deep features followed by max-pooling to be a global shape descriptor. Afterwards, a series of works attempt to improve feature embedding by designing novel deep networks on point clouds. PointNet++ [1] builds hierarchical architecture with PointNet as local encoder. More works such as SpiderCNN [9], PointCNN [2], KCNet [21], PointConv [22], RS-CNN [23] learn point cloud features with various local convolutions starting from point coordinates or additional geometric features. Though these works have achieved state-of-the-art performance for various shape analysis tasks, they are sensitive to shape rotation. In our work, we build a novel network with building blocks being able to achieve rotation-invariance, but with positional raw features as inputs. Our network can explore and encode positional information of 3D shapes while precluding the disadvantage of rotation-sensitiveness.

2.2 Rotation-robust representations

There are several ways to improve rotation-robustness. Data augmentation by randomly rotating shapes in training set is widely applied in training 3D networks as in [8, 9]. Furthermore, transformer network [24] was generalized and utilized in 3D recognition [8, 25, 26]. The transformer implicitly learns alignments among shapes either in 3D space or feature space which may reduce the impact of shape rotations. As shown in [18], these techniques commonly improved robustness but can not achieve rigorous invariance to rotations.




Fig. 2. PR-invNet consists of geometric feature extraction, positional feature embedding block, relational feature embedding block. PR-invNet is designed for classification and segmentation. Shape coarsening is not utilized in segmentation.

Rotation-invariant geometric features can be taken as network inputs. Traditionally, these hand-crafted features are widely utilized to represent 3D shapes in [27–31], and further introduced as inputs of deep networks [18, 20, 32]. Most of them utilize distances and angles [20, 32] and sine / cosin values of the angles [18] as raw point-wise features. They encode relations among points and achieve invariance to arbitrary shape rotations. However, the geometric features may lose discriminative clues that are contained in point positions / orientations.

Another way to achieve rotation-robustness is to learn rotation-equivariant features. [33] built filters on spherical harmonics and learned locally equivariant features to 3D rotations and translations. Group convolution [34] is a generalized convolution to achieve rotation-equivariant feature learning based on rotation-group. This idea is extended to 3D domain [35] with various discretized rotation-groups. Spherical CNN [17] and a^3 SCNN [19] defined group convolution on sphere which can be taken as an approximation to infinite group $SO(3)$. They also require the input of the networks to be rotation-equivariant if regarding features as functions over shapes. Commonly, these methods lead to rotation-robustness by global aggregation such as average-pooling or max-pooling, but rigorous rotation-invariance is not ensured due to the discretized rotation groups.

Instead of requiring rotation-invariant geometric features as inputs, our network can directly learn positional features from point coordinates as inputs while achieving rigorous rotation-invariance based on PFE-block. It embeds the input shape into a rotation-invariant pose space, from which we derive a discriminative pose by pose selector to be sent to positional feature extractor, and this operation is rotation-invariant and justified to be effective in experiments.

3 Method

In this section, we introduce our rotation-invariant network, dubbed *PR-InvNet*. It aims at deep positional and relational feature learning, based on, respectively, the PFE-block and RFE-block. As in Fig. 2, it consists of several stages. The initial feature embedding stage is composed of geometric feature extraction and




Fig. 3. Illustration of geometric feature extraction for point o which is based on distances and angles. See Sect. 3.1 for details.

RFE-block. The output features of the PFE and RFE blocks are max-pooled over point cloud and concatenated, then fed to a MLP for shape classification. For point cloud segmentation, we switch off shape coarsening operation, and the concatenated features of PFE-blocks and RFE-blocks are further concatenated with globally max-pooled features before being fed to MLP for point label prediction. Each stage is now detailed in the following.

3.1 Geometric feature extraction

PR-invNet takes geometric features similar to [18] as additional raw input features complementary to point coordinates⁴. As shown in Fig. 3, c and o are the centers of the global shape and a local patch respectively. To compute geometric features of point o , we first consider its neighboring points. The feature of each of its neighboring point q is computed as $[|\vec{co}|, |\vec{oq}|, \cos(\alpha), \cos(\beta)]$, with $[\cdot]$ being the concatenation operation. Then we concatenate features of all neighboring points together with the distance of $|\vec{co}|$ as the geometric feature of point o , i.e.,

$$g_o = [|\vec{co}|, \{|\vec{co}|, |\vec{oq}|, \cos(\alpha), \cos(\beta)\}_{q \in N(o)}], \quad (1)$$

where $N(o)$ denotes set of neighboring points of o . For simplicity, we denote it as $\{g_i\}_{i=1}^N$ with N as number of points. It can be noted that g_i is rotation-invariant, please refer to supplemental material for proof. In our work, for each given point, the above geometric features are extracted and concatenated over a multi-scale neighborhoods (three neighborhoods determined by Euclidean distance with 8, 16, 32 points), and 8 neighboring points are uniformly sampled at each scale. Then the geometric feature of each point is in length of $3 \times 4 \times 8 + 1 = 97$.

3.2 Positional feature embedding block

The PFE-block aims to perform rotation-invariant feature embedding by incorporating positional information encoded in point coordinates. Since point coordinates are sensitive to shape rotations, one solution is to take all rotated versions (i.e., poses) of a shape into consideration, and fuse extracted features from them by global pooling. However, this is inefficient since there are infinite

⁴ Even without geometric features, our network would be still rotation-invariant.

number of rotations in $SO(3)$ to ensure rigorous rotation-invariance. We propose an idea that first maps input shape to a rotation-invariant pose space by *pose expander*, then selects a representative pose by *pose selector*. Point coordinates of selected pose are utilized to extract positional features by *feature extractor*.

Pose expander. The pose expander aims to map a shape to produce a rotation-invariant pose space. Given a shape with point cloud $P \in \mathcal{R}^{N \times 3}$, it is obvious that the pose space $\{\mathcal{T}_g P\}_{g \in SO(3)}$ containing all possible rotated shapes in $SO(3)$ is rotation-invariant. We aim to derive a compact pose space of input shape guaranteeing rotation-invariance. We achieve this by first normalizing the pose of each shape via Principal Components Analysis (PCA), then applying a discretized rotation-group in $SO(3)$ to the normalized shape to derive a pose space.

We first apply PCA on the point coordinates of a shape to normalize its pose using a coordinate system composed of the three eigen-vectors. We indicate the normalized shape as \hat{P} . Note that pose normalization by PCA is not injective. For example, eigendecomposition may result in different signs of eigen-vectors, therefore results in different normalized shapes. In Fig. 4, a shape with two rotated versions may produce two different normalized shapes, which are symmetric.




Fig. 4. Two rotated versions of one shape may result in two symmetric and different poses by PCA.

To solve the ambiguity of PCA-normalized shape, we enumerate all possible signs of eigen-vectors, and construct a PCA-normalized pose space with eight normalized shapes based on a symmetric group G_S

$$G_S = \{g_s | g_s = \begin{bmatrix} x & 0 & 0 \\ 0 & y & 0 \\ 0 & 0 & z \end{bmatrix}, x, y, z \in \{-1, 1\}, s = 1, 2, \dots, 8\}, \quad (2)$$

Then the PCA-normalized pose space is obtained as

$$H_S(P) = \{\hat{P}g_s | g_s \in G_S\}, \quad (3)$$

where $H_S(P)$ contains eight possible PCA-normalized shapes corresponding to different possible signs of eigen-vectors, and this pose space is denoted as *sym-space*. Theorem 1 proves that $H_S(P)$ is rotation-invariant for input shape P .

Though sym-space is rotation-invariant, PCA may not always well align the poses of different shapes even from the same category. We further expand the pose space from $H_S(P)$ to increase the possibility of containing aligned poses for different shapes. By discretizing $SO(3)$ using limited representative rotations, we enlarge sym-space $H_S(P)$ to *rot-space* $H_{SR}(P)$ using rotation group G_R as

$$H_{SR}(P) = \{\hat{P}g_s g_r | g_s \in G_S, g_r \in G_R\}, \quad (4)$$

where G_R can be any discretized subgroup of $SO(3)$. $H_{SR}(P)$ is also rotation-invariant for P as proved by Theorem 1. Compared with $H_S(P)$, rot-space

$H_{SR}(P)$ contains more poses with higher chances to include the same or similar poses of different shapes. In our implementation, we use rotation group A_5 , i.e., alternating group of a regular dodecahedron (or a regular icosahedron) [36] to construct the rot-space. By deleting duplicate shapes, we finally map the shape with rotation set $A_5 \times Z_2$ into rot-space with 120 poses.

Pose selector. Given a shape P , we have derived a rotation-invariant pose space $H_{SR}(P)$, i.e., the rot-space. One naive way is to extract features from different poses of the shape in pose space followed by orderless operations, e.g., average- or max-pooling, to aggregate features as a rotation-invariant representation for shape P . However, it is computationally intensive to extract features from multiple poses of shape. We design a simple but effective method by selecting the most representative shape pose using a pose selector, and learn positional features from this single selected shape instead than from all shapes in pose space.

The pose selector is designed as a *multi-head neural network* $\Psi(\cdot)$ over shape poses in $H_{SR}(P)$. We score each pose by $v_s = \Psi(\tilde{P}_s)$ for pose $\tilde{P}_s \in H_{SR}(P)$, $s = 1, \dots, |H_{SR}(P)|$. $v_s \in \mathcal{R}^{N_h}$ is a vector generated by N_h heads of $\Psi(\cdot)$. In our implementation, $\Psi(\cdot)$ is designed as a simple network with a point-wise fully connected (FC) layer, a max-pooling layer over all points, followed by an additional FC layer and softmax with N_h output scores for a point cloud. Similar to multi-head attention [37], the multiple elements in vector v_s reflect the responses of \tilde{P}_s to different modes of the neural network. For a shape pose \tilde{P}_s , its final score is set as the largest response in vector v_s , i.e., $\bar{v}_s = \max\{v_s\}$, and \bar{v}_s is a scalar. Then the selected representative shape \tilde{P} from pose space $H_{SR}(P)$ for shape P is the shape pose with the largest score.

Theorem 1 proves that the sym-space, rot-space and the selected pose are rotation-invariant to P (please see supplementary material for proof).

Theorem 1. Denoting $P \in \mathcal{R}^{N \times 3}$ as point cloud of a shape, its sym-space $H_S(P)$ in Eqn. (3) and rot-space $H_{SR}(P)$ in Eqn. (4) are rotation-invariant for P , i.e., for rotated shape $Q = PR$ of P with any rotation matrix $R \in \mathcal{R}^{3 \times 3}$, we have $H_S(P) = H_S(Q)$ and $H_{SR}(P) = H_{SR}(Q)$. Assuming the pose selector is injective, the selected pose of P is rotation-invariant, i.e., $\tilde{P} = \tilde{Q}$.

Feature extractor. With the selected shape pose $\tilde{P} \in \mathcal{R}^{N \times 3}$ of input shape P , we can extract its point-wise features based on the point coordinates like the traditional point-based networks, e.g., [1, 2, 9]. Specifically, given point-wise coordinates \tilde{p}_i , geometric feature g_i (extracted from input shape P), and already extracted point features f_i , we concatenate $[g_i, f_i, \tilde{p}_i]$ as point-wise input features for our feature extractor. Motivated by the structure of PointNet++ [1], for each point i , its point feature is extracted as

$$f'_i = \Gamma(\{[g_j, f_j, \tilde{p}_j]\}_{j \in N(i)}) \quad (5)$$

where $N(i)$ is the local neighborhood of point i with K points including point i . $\Gamma(\cdot)$ is the feature extractor, which is designed as a shared MLP for points,




Fig. 5. Overview of the PFE-block. It maps a shape to be rotation-invariant pose space, i.e., rot-space, then selects a discriminative pose via pose selector, followed by feature extractor. This block is rotation-invariant.

followed by max-pooling over neighborhood. Fig. 5 shows the pipeline of the PFE-block. Since the selected pose is rotation-invariant to input shape by Theorem 1, its point coordinates \tilde{p}_i are also rotation-invariant. Hence, features in Eqn. (5) are rotation-invariant because all inputs of Γ are rotation-invariant.

3.3 Relational feature embedding block

The relational feature embedding block (RFE-block) is specifically designed to explore and learn point features by modeling interactions of neighboring points. For each point i , we aggregate the features of its local neighborhood with K points including point i by convolution. Assuming the neighboring points are indexed by k sorted by increasing Euclidean distance to point i , and corresponding point features are $\{f_k \in \mathcal{R}^d\}_{k=1}^K$, the updated feature $f'_i \in \mathcal{R}^{d'}$ for point i is

$$f'_{il} = \sum_{k \in N(i)} \sum_{j=1}^d f_{kj} W_{kjl}, \text{ for } l = 1, \dots, d' \quad (6)$$

where f_{kj} is the j -th feature channel of the k -th neighboring point, f'_{il} is the l -th channel of the updated feature f'_i . W_{kjl} is defined as

$$W_{kjl} = \Phi([C_k^p, C_k^g, C_k^f, g_k])_j \Theta_{kjl}, \quad (7)$$

where $C_k^p, C_k^g, C_k^f \in \mathcal{R}^K$ are K -d vectors denoting correlations of centered point coordinates, geometric features, point features respectively between k -th point and all K points in local neighborhood, followed by l_2 -normalization. For example, for centered point coordinates, i.e., coordinates subtracted by that of center pixel i , $C_k^p = \text{Normalize}([p_k^\top p_1, p_k^\top p_2, \dots, p_k^\top p_K])$, where we use p to represent the centered point coordinates. C_k^g, C_k^f are similarly defined as inner products of corresponding features. $\Phi(\cdot)$ is modeled as MLP, embedding relationships of neighboring points to convolution weight with learnable parameters Θ_{kjl} . It is easy to verify that weights in Eqn. (7) are rotation-invariant because the inputs of Φ are rotation-invariant defined by inner products.

3.4 Network architecture

PR-invNet is designed by concatenating rotation-invariant PFE-block and RFE-block. As shown in Fig. 2, given a shape, we first initialize the point-wise features by initial feature embedding, followed by three stages concatenating PFE-block and RFE-block to extract positional and relational features. For classification, we use an additional shape coarsening operation at each stage to construct a hierarchy of coarsened points, implemented by farthest point sampling with sampling rate of $1/2$. The intermediate features generated by PFE-blocks and RFE-blocks of different stages are concatenated and aggregated by a MLP set as two layers of FC+BN+ReLU+Dropout, followed by a FC to output confidence for shape categorization or segmentation. The number of hidden units of the first two FCs are 512, 256, and dropout probability is 0.5. The number of neighboring points in Eqns. (5-6) is 16. For an input shape of PR-invNet, we conduct PCA to normalize shape pose only one time, and all PFE-blocks in different stages derive rot-space by Eqn.(4) starting from point cloud of this normalized pose or its coarsened point clouds. More details of pose selector Ψ , feature extractor Γ , convolution filters Φ , Θ in RFE-block are introduced in supplementary material.

4 Experimental results

In this section we evaluate and compare with state-of-the-art methods for rotated shape classification and segmentation on the following 3D datasets.

ModelNet40 [6] consists of 12,311 shapes from 40 categories, with 9,843 training and 2,468 test objects for shape classification.

ShapenetCore55 [38] contains two subsets: ‘normal’, ‘perturbed’, with aligned and randomly rotated shapes in $SO(3)$ respectively. 51,190 models are categorized into 55 classes with training, validation, test sets in ratios of 70%, 10%, 20%. We uniformly sample points from the original mesh data for both subsets.

ShapeNet part [39] contains 16,880 shapes from 16 category of shapes, with 14,006 / 2,874 shapes for training / test, annotated with 50 parts.

Each shape is represented by 1024 points. To train our PR-invNet, we take Adam optimizer with initial learning rate, epoch number as 0.001, 250, and the learning rate is exponentially decayed with decay rate and decay step as 0.7, 200000. The batch size for shape classification and segmentation are 16 and 8 respectively. It takes 286.6ms, 672.1ms in one iteration for shape classification and segmentation respectively on a NVIDIA 1080 Ti GPU.

4.1 3D shape classification

We mainly compare with point-based methods and rotation-robust methods.

ModelNet40. As done in Spherical CNN [17], we compare in the following three modes. (1) Both training and test on data augmented by azimuthal rotation (z/z); (2) Training with azimuthal rotation and test with arbitrary rotation

Table 1. Shape classification accuracy on ModelNet40 (in %).

	Method	Input size	z/z	SO(3)/SO(3)	z/SO(3)
Rotation-sensitive	VoxNet [40]	30^3	83.0	87.3	-
	SubVolSup [7]	30^3	88.5	82.7	36.6
	SubVolSup MO [7]	30^3	89.5	85.0	45.5
	MVCNN 12x [4]	12×224^2	89.5	77.6	70.1
	MVCNN 80x [4]	80×224^2	90.2	86.0	81.5
	PointNet [8]	1024×3	87.0	80.3	12.8
	PointNet++ [1]	1024×3	89.3	85.0	28.6
	PointCNN [2]	1024×3	91.3	84.5	41.2
	DGCNN [10]	1024×3	92.2	81.1	20.6
Rotation-robust	Spherical CNN [17]	$2 \times 64 \times 64$	88.9	76.9	86.9
	a^3 SCNN [19]	$2 \times 165 \times 65$	89.6	87.9	88.7
	RICNN [20]	1024×3	86.5	86.4	86.4
	ClusterNet [18]	1024×3	87.1	87.1	87.1
	RRI-PointNet++	1024×3	79.4	79.4	79.4
	Proposed	1024×3	89.2	89.2	89.2

(z/SO(3)); (3) Both training and test with arbitrary rotation augmented data (SO(3)/SO(3)). The results are presented in Table 1. The traditional point-based methods, including PointNet [8], PointNet++ [1], PointCNN [2], DGCNN [10], perform best in z/z mode, but decline sharply when testing on data augmented with arbitrary rotations (z/SO(3)). Though using arbitrarily augmented training data in SO(3)/SO(3) mode, their performances still drop with large gaps. For these rotation-robust methods including Spherical CNN [17], a^3 SCNN [19], RICNN [20], ClusterNet [18] and RRI⁵-PointNet++, they are more robust in different modes. Considering that Spherical CNN and a^3 SCNN rely on discretized angles in the sphere, they are not rigorously rotation-invariant and performance drops in modes (2)-(3). ClusterNet, RICNN and RRI-PointNet++ are rotation-invariant relying on input geometric features. Our PR-invNet is also rotation-invariant, and it achieves highest performance in modes (2-3) and second highest performance in mode (1) than all the rotation-robust methods, demonstrating its superiority for arbitrarily rotated shape classification.

ShapeNetCore55. We compare with point-based methods including PointNet [8], PointNet++ [1], DGCNN [10], RS-CNN [23], SpiderCNN [9], and also rotation-robust Spherical CNN [17] and rotation-invariant RICNN [20], RRI-PointNet++. The results are presented in Table 2. As shown in Table 2, we achieve best performance in both aligned and perturbed datasets, compared with Spherical CNN, RICNN and RRI-PointNet++. Among all methods, we achieve highest accuracy on the perturbed dataset.

4.2 3D shape segmentation

⁵ RRI is an essential part of ClusterNet, and the codes are provided by the authors.

Table 2. Shape classification accuracy on ShapeNetCore55 (in %). All the results of compared methods are achieved by running their code.


	Method	Input size	Aligned	Perturbed
Rotation-sensitive	PointNet [8]	1024×3	83.4	74.6
	PointNet++ [1]	1024×3	84.2	70.9
	DGCNN [10]	1024×3	86.3	74.1
	RS-CNN [23]	1024×3	85.9	73.4
	SpiderCNN [9]	1024×3	79.8	64.4
Rotation-robust	Spherical CNN [17]	2×64^2	76.2	73.8
	RIConvNet [20]	1024×3	78.5	76.9
	RRI-PointNet++	1024×3	70.8	67.9
	Proposed	1024×3	78.9	77.6

Table 3. Shape segmentation mIoU on ShapeNet part segmentation dataset (in %).

	Method	Input	z/z	$SO(3)/SO(3)$	$z/SO(3)$
Rotation-sensitive	PointNet [8]	xyz	76.2	74.4	37.8
	PointNet++ [1]	xyz+normal	80.7	76.7	48.2
	PointCNN [2]	xyz	81.5	71.4	34.7
	DGCNN [10]	xyz	78.8	73.3	37.4
	SpiderCNN [9]	xyz+normal	81.8	72.3	42.9
Rotation-robust	RIConvNet [20]	xyz	75.6	75.5	75.3
	Proposed	xyz	79.4	79.4	79.4

We next apply PR-invNet to point cloud segmentation to predict the part label of each point. We use the same experimental settings as RIConvNet [20], and mean per-class IoU (mIoU, %) are presented in Table 3. Compared with all rotation-sensitive

point-based methods and rotation-robust methods, our PR-invNet outperforms them significantly in both $SO(3)/SO(3)$ and $z/SO(3)$ modes, i.e., more than 2.7% and 4.1% mIoU improvements. We also report the per-class IoU for above settings in $z/SO(3)$, $SO(3)/SO(3)$ modes respectively in the supplemental material, and our PR-invNet achieves better performance than the rotation-robust RIConvNet on all categories. In Fig. 6, we show the segmentation results in $z/SO(3)$ mode of several objects as well as the corresponding ground truth labels, and our predictions labels are reasonable and close to ground truth.

**Fig. 6.** Shape segmentation results.




Fig. 8. Rotation groups.




Fig. 9. Head numbers.




Fig. 10. Noise levels.

Table 4. Classification accuracy on ModelNet40 (%).

Method	Geo-fea	PFE-block	RFE-block	Acc.
PR-invNet-noGeo	×	✓	✓	88.2
PR-invNet-noPFE	✓	×	✓	87.8
PR-invNet-noRFE	✓	✓	×	88.0
PR-invNet	✓	✓	✓	89.2




Fig. 7. Stage numbers.

4.3 Ablation study

In this section, we conduct ablation study on ModelNet40 to justify the effects of our network design, and also test the robustness of PR-invNet to noises.

Effect of proposed blocks. To evaluate the effectiveness of geometric feature extraction, PFE-block, and RFE-block, we conduct experiments that utilize networks without above blocks respectively, i.e., PR-invNet-noGeo, PR-invNet-noPFE, PR-invNet-noRFE, on ModelNet40, and present the experiment results in Table 4. Compared with PR-invNet, the networks of PR-invNet-noGeo, PR-invNet-noPFE, PR-invNet-noRFE achieve lower performance, demonstrating the need for each component. Note that all networks of above variants are also rigorously rotation-invariant, and they also achieve better performance than most of the rotation-sensitive and rotation-robust methods whose results are presented in Table 1 in the $SO(3)/SO(3)$ and $z/SO(3)$ modes.

Effect of stage number. To evaluate the effect of number of stages in PR-invNet, we compare performance of PR-invNet- i with $i = 0, 1, 2, \dots, 5$, and i denotes number of stages excluding initial feature embedding stage. As shown in Fig. 7, network with 3 stages achieve the highest accuracy. Deeper architectures have larger capacity but marginally decreased performances. This phenomenon was also observed in other graph CNNs [41], and training of deeper networks on point clouds deserves to be more investigated in future work.

Design of global shape feature. As shown in Fig. 2, learned features from all PFE-blocks and RFE-blocks are aggregated by MLP for classification. We evaluate PR-invNet-PFE and PR-invNet-RFE that only utilize features from PFE-blocks and RFE-blocks respectively, achieving 88.5%, 87.4% classification

accuracies on ModelNet40. Result using only features of the last RFE-block is 88.3%. This demonstrates that concatenating all features of PFE-block and RFE-block is more effective.

Effect of PCA-normalization and rot-space. In PR-invNet, we take PCA to normalize shape as network input. We also conduct experiments with networks without PCA-normalization and only with PCA (without using sym-space and rot-space), i.e., PR-invNet-noPCA, PR-invNet-PCA. Their classification accuracies in ModelNet40 with z/z mode are 90.0% and 89.1% respectively, but drop to 44.3%, 74.8% in z/SO(3) mode, which are not rotation-invariant. In SO(3)/SO(3), our pose selector over rot-space upon PCA(89.2%) outperforms pure PCA for alignment(88.3%). Removing both PCA and pose selector achieves 87.3%, showing our pose selection is more effective than PCA.

Effect of rotation group G_R . In pose expander of PFE-block, we use the rotation group G_R to construct rot-space, which is taken as alternating group A_5 of a regular dodecahedron (or icosahedron). Other alternating (symmetry) group can also be used, such as identity group (I_R), or A_4 , S_4 [36] respectively corresponding to regular tetrahedron, cube (or octahedron). The classification accuracies using these groups are in Fig. 8. Compared with other groups, A_5 has the maximum number of elements, resulting in largest rot-space and achieves the highest accuracy. Operating on this largest space, the pose selector has more probability to select a better pose for shape recognition.

Effect of pose selector. In PFE-block, we use a pose selector to select the representative shape pose in rot-space. However, we can also extract features from all shapes from rot-space and then perform max-pooling or average-pooling over them to aggregate features. We compare our pose selector with these aggregation methods and present the classification accuracy and computational cost in Table 5. Since it is infeasible to extract features from all shapes in rot-space containing 120 poses, we compare these variants on the sym-space containing eight shape poses, i.e., rot-space with identity rotation group. We experiment with batch size 12, and as shown in Table 5, our pose selection method achieves higher performance while needing less GPU memory and computational time.

Effect of head number in Ψ . For pose selector Ψ , we design it as a multi-head neural network, which has 200 heads. We conduct experiments to demonstrate the effect of head numbers. As shown in Fig. 9, the Ψ with head number as 200 achieves higher performance. Note that Ψ with multi-heads all achieve better accuracies than that with one head.

Table 5. Cla. Acc. on ModelNet40.

Method	Acc. (%)	Memory (GB)	Time (ms)
Average-pooling	88.1	10.2	496.5
Max-pooling	88.3	10.2	512.3
Proposed	88.5	4.4	251.0

Table 6. Cla. Acc. on ModelNet40 (%).

Method	C_k^p	C_k^g	C_k^f	g_k	Acc.
PR-invNet-P	×	✓	✓	✓	88.1
PR-invNet-G	✓	×	✓	✓	88.8
PR-invNet-F	✓	✓	×	✓	89.9
PR-invNet-g	✓	✓	✓	×	88.5
PR-invNet	✓	✓	✓	✓	89.2

Design of RFE-block. In RFE-block, we design convolution filters based on three terms, i.e., the relations of point position, geometric feature, and point feature in Eqn. (7). In Table 6, we show our networks using only one of these terms, resulting in networks of PR-invNet-P, PR-invNet-G, PR-invNet-F respectively. We also present result of PR-invNet-g only using geometric feature in Eqn. (7). Their classification accuracies are all inferior compared with full version. We also conduct experiment that without normalization in Eqn. (7), achieving classification accuracy as 88.6%.

Robustness to noise. We train PR-invNet on ModelNet40 training dataset and test it on test data with various levels of noise under the $z/\text{SO}(3)$ mode. We add Gaussian noise with different standard deviation (Std) on each point (coordinates within a unit ball) independently. The overall classification accuracies are in Fig. 10. PR-invNet keeps robustness under noise with Std of 0.01.




Fig. 11. Illustration of selected shapes corresponding to different heads in multi-head score network. Shapes in a box highlighted by same color correspond to same head.

Visualization for pose selector. In PFE-block, we design multi-head score network Ψ to select a pose from rot-space $H_{SR}(P)$ based on scoring each shape pose by its maximum score belonging to multiple heads. Therefore the selected pose from rot-space corresponds to a certain head. Here we illustrate selected shape poses organized by corresponding heads highlighted in different colors in Fig. 11. We observe that shapes corresponding to each head have similar poses. The multi-head score network may enable to learn to align the shapes to clusters of poses using heads in multi-head score network.

5 Conclusions

In this work, we focus on rotation-invariant deep network design on point clouds by proposing two rotation-invariant network blocks. The constructed PR-invNet was extensively justified to be effective for rotated 3D shape classification and segmentation. In future work, we are interested to deeply investigate pose alignment capability of PFE-block, and further improve the network architecture.

Acknowledgement. This work was supported by NSFC (11971373, 11690011, U1811461, 61721002) and National Key R&D Program 2018AAA0102201.

References

1. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: *Advances in Neural Information Processing Systems*. (2017) 5099–5108
2. Li, Y., Bu, R., Sun, M., Wu, W., Di, X., Chen, B.: Pointcnn: Convolution on x-transformed points. In: *Advances in Neural Information Processing Systems*. (2018) 820–830
3. Pham, Q.H., Nguyen, T., Hua, B.S., Roig, G., Yeung, S.K.: Jsis3d: joint semantic-instance segmentation of 3d point clouds with multi-task pointwise networks and multi-value conditional random fields. In: *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. (2019) 8827–8836
4. Su, H., Maji, S., Kalogerakis, E., Learned-Miller, E.: Multi-view convolutional neural networks for 3d shape recognition. In: *IEEE International Conference on Computer Vision (ICCV)*. (2015) 945–953
5. Esteves, C., Xu, Y., Allen-Blanchette, C., Daniilidis, K.: Equivariant multi-view networks. In: *IEEE International Conference on Computer Vision (ICCV)*. (2019) 1568–1577
6. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: *CVPR*. (2015) 1912–1920
7. Qi, C.R., Su, H., Nießner, M., Dai, A., Yan, M., Guibas, L.J.: Volumetric and multi-view cnns for object classification on 3d data. In: *CVPR*. (2016) 5648–5656
8. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: *CVPR*. (2017) 652–660
9. Xu, Y., Fan, T., Xu, M., Zeng, L., Qiao, Y.: Spidercnn: Deep learning on point sets with parameterized convolutional filters. In: *European Conference on Computer Vision (ECCV)*. (2018) 87–102
10. 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) (2019) 1–12
11. Rao, Y., Lu, J., Zhou, J.: Spherical fractal convolutional neural networks for point cloud recognition. In: *CVPR*. (2019) 452–460
12. a, Y.Z., a b, G.J., b, Y.L.: A novel finger and hand pose estimation technique for real-time hand gesture recognition. *Pattern Recognition* **49** (2016) 102–114
13. Xiang, Y., Schmidt, T., Narayanan, V., Fox, D.: Posecnn: A convolutional neural network for 6d object pose estimation in cluttered scenes. *arXiv:1711.00199* (2017)
14. Harald, W., Didier, S.: Tracking of industrial objects by using cad models. *Journal of Virtual Reality and Broadcasting* **4**(1)
15. Bero, S.A., Muda, A.K., Choo, Y., Muda, N.A., Pratama, S.F.: Rotation analysis of moment invariant for 2d and 3d shape representation for molecular structure of ats drugs. In: *In 4th World Congress on Information and Communication Technologies (WICT)*. (2014) 308–313
16. Berenger, F., Voet, A., Lee, X.Y., Zhang, K.Y.J.: A rotation-translation invariant molecular descriptor of partial charges and its use in ligand-based virtual screening. *Journal of Cheminformatics* **6**(1) (2014) 23–23
17. Esteves, C., Allen-Blanchette, C., Makadia, A., Daniilidis, K.: Learning so(3) equivariant representations with spherical cnns. In: *European Conference on Computer Vision (ECCV)*. (2018) 52–68
18. Chen, C., Li, G., Xu, R., Chen, T., Wang, M., Lin, L.: Clusternet: Deep hierarchical cluster network with rigorously rotation-invariant representation for point cloud analysis. In: *CVPR*. (2019) 4994–5002

19. Liu, M., Yao, F., Choi, C., Sinha, A., Ramani, K.: Deep learning 3d shapes using alt-az anisotropic 2-sphere convolution. In: ICLR. (2019)
20. Zhang, Z., Hua, B.S., Rosen, D.W., Yeung, S.K.: Rotation invariant convolutions for 3d point clouds deep learning. In: International Conference on 3D Vision (3DV). (2019) 204–213
21. Shen, Y., Feng, C., Yang, Y., Tian, D.: Mining point cloud local structures by kernel correlation and graph pooling. In: CVPR. (2018) 4548–4557
22. Wu, W., Qi, Z., Fuxin, L.: Pointconv: Deep convolutional networks on 3d point clouds. In: CVPR. (2019) 9621–9630
23. Liu, Y., Fan, B., Xiang, S., Pan, C.: Relation-shape convolutional neural network for point cloud analysis. In: CVPR. (2019) 8895–8904
24. Jaderberg, M., Simonyan, K., Zisserman, A., et al.: Spatial transformer networks. In: Advances in Neural Information Processing Systems. (2015) 2017–2025
25. Bas, A., Huber, P., Smith, W.A., Awais, M., Kittler, J.: 3d morphable models as spatial transformer networks. In: IEEE International Conference on Computer Vision Workshops. (2017) 904–912
26. Mukhaimar, A., Tennakoon, R., Lai, C.Y., Hoseinnezhad, R., Bab-Hadiashar, A.: Pl-net3d: Robust 3d object class recognition using geometric models. *IEEE Access* **7** (2019) 163757–163766
27. Stein, F., Medioni, G.: Structural indexing: Efficient 3-d object recognition. *IEEE Transactions on Pattern Analysis & Machine Intelligence* (2) (1992) 125–145
28. Sun, Y., Abidi, M.A.: Surface matching by 3d point’s fingerprint. In: IEEE International Conference on Computer Vision (ICCV). (2001) 263–269
29. Zhong, Y.: Intrinsic shape signatures: A shape descriptor for 3d object recognition. In: International Conference on Computer Vision Workshops. (2009) 689–696
30. Rusu, R.B., Blodow, N., Beetz, M.: Fast point feature histograms (fpfh) for 3d registration. In: IEEE International Conference on Robotics and Automation. (2009) 3212–3217
31. Tombari, F., Salti, S., Di Stefano, L.: Unique signatures of histograms for local surface description. In: European Conference on Computer Vision (ECCV). (2010) 356–369
32. Deng, H., Birdal, T., Ilic, S.: Ppf-foldnet: Unsupervised learning of rotation invariant 3d local descriptors. In: European Conference on Computer Vision (ECCV). (2018) 602–618
33. Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., Riley, P.: Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv:1802.08219* (2018)
34. Cohen, T., Welling, M.: Group equivariant convolutional networks. In: International Conference on Machine Learning (ICML). (2016) 2990–2999
35. Worrall, D., Brostow, G.: Cubenet: Equivariance to 3d rotation and translation. In: European Conference on Computer Vision (ECCV). (2018) 567–584
36. Zimmermann, B.P.: On finite groups acting on spheres and finite subgroups of orthogonal groups. *arXiv:1108.2602* (2011)
37. 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. (2017) 5998–6008
38. Savva, M., Yu, F., Su, H., Aono, M., Chen, B., Cohen-Or, D., Deng, W., Su, H., Bai, S., Bai, X., et al.: Shrec16 track: large-scale 3d shape retrieval from shapenet core55. In: Eurographics Workshop on 3D Object Retrieval. (2016)

39. Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., et al.: Shapenet: An information-rich 3d model repository. arXiv preprint arXiv:1512.03012 (2015)
40. Maturana, D., Scherer, S.: Voxnet: A 3d convolutional neural network for real-time object recognition. In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). (2015) 922–928
41. Li, G., Muller, M., Thabet, A., Ghanem, B.: Deepgcns: Can gcns go as deep as cnns? In: CVPR. (2019) 9267–9276