

# Local-consistent Transformation Learning for Rotation-invariant Point Cloud Analysis

Yiyang Chen<sup>1</sup> Lunhao Duan<sup>2</sup> Shanshan Zhao<sup>3†</sup> Changxing Ding<sup>1,4†</sup> Dacheng Tao<sup>5</sup>  
<sup>1</sup>South China University of Technology <sup>2</sup>Wuhan University <sup>3</sup>The University of Sydney  
<sup>4</sup>Pazhou Lab <sup>5</sup>Nanyang Technological University

eeyychen@mail.scut.edu.cn, lhduan@whu.edu.cn

{sshan.zhao00, dacheng.tao}@gmail.com, chxding@scut.edu.cn

## Abstract

Rotation invariance is an important requirement for point shape analysis. To achieve this, current state-of-the-art methods attempt to construct the local rotation-invariant representation through learning or defining the local reference frame (LRF). Although efficient, these LRF-based methods suffer from perturbation of local geometric relations, resulting in suboptimal local rotation invariance. To alleviate this issue, we propose a Local-consistent Transformation (LocoTrans) learning strategy. Specifically, we first construct the local-consistent reference frame (LCRF) by considering the symmetry of the two axes in LRF. In comparison with previous LRFs, our LCRF is able to preserve local geometric relationships better through performing local-consistent transformation. However, as the consistency only exists in local regions, the relative pose information is still lost in the intermediate layers of the network. We mitigate such a relative pose issue by developing a relative pose recovery (RPR) module. RPR aims to restore the relative pose between adjacent transformed patches. Equipped with LCRF and RPR, our LocoTrans is capable of learning local-consistent transformation and preserving local geometry, which benefits rotation invariance learning. Competitive performance under arbitrary rotations on both shape classification and part segmentation tasks and ablations can demonstrate the effectiveness of our method. Code will be available publicly at <https://github.com/wdttt/LocoTrans>.

## 1. Introduction

Point cloud analysis [1, 6, 10, 20, 37, 42, 45] is a fundamental task in 3D vision. Existing works [21, 29] have achieved excellent performance in classification and part segmentation for pre-aligned point cloud data. However, these works

†Corresponding authors




Figure 1. Illustration of LRFs and the points transformation. For two reference points marked by Blue and Cyan, we use the LRF in OrientedMP [16] and our proposed LCRF to transform the local patches centered in them, respectively. By observing (a) and (b), we can find that our LCRF is able to yield local-consistent transformation and thus preserve the local shapes better. In comparison, OrientedMP only generates consistent orientation in one axis (marked by Red) but fails in the other (marked by Green).

tend to suffer severe performance drops when processing 3D objects with arbitrary rotations, as they do not learn the rotation-invariant feature representations. The lack of the capability for rotation invariance limits the application of the point cloud analysis model in practical situations. Currently, many attempts [2, 12, 16, 34, 41] have been made to improve the performance under arbitrary rotations.

To achieve rotation invariance, previous works [2, 34, 43] propose to exploit local relative geometric relationships that remain unchanged under rotations, such as distance and angle, for feature learning, instead of the 3D coordinates. However, such a strategy does not fully use 3D geometric information. Recently, some methods [4, 16, 41] aim to learn rotation-invariant representations by building local reference frames (LRFs). These LRF-based methods con-

struct three orthonormal basis vectors for each point as the reference frame and then project the original local points into the frame. The orientations of LRF can be defined from geometric information, such as local barycenter, global center, and normal vector, or learned from input point cloud data. For instance, OrientedMP [16] is a pioneering work that employs the equivariant network to extract local features for building LRF.

However, constructing local rotation-invariant representation via existing LRFs inevitably changes the original local geometry. As shown in Fig. 1(a), for two adjacent points (marked by Cyan and Blue), we transform the local regions centered by them using the LRF built by OrientedMP, and different orientations are observed. As a result, the network might be struggling to learn identical representations for them due to the loss of relative pose. Moreover, this perturbation may lead to an ambiguous representation when applying vanilla convolution directly to points that have lost their relative pose and been assigned within a local patch after a k-nearest-neighbors (KNN) operation, as observed in existing works [4, 41]. To analyze the reason causing such an issue, we further visualize the learned orientations along two axes in OrientedMP. As illustrated in Fig. 1(a), the local consistency only exists in one orientation but not in the other. Therefore, we attribute the perturbation of geometry relationships to the local inconsistency of learned LRF.

Motivated by the analysis above, we propose *Local-consistent Transformation (LocoTrans)* learning strategy, aiming to learn local-consistent reference frame and achieving better local geometry preservation. Specifically, following OrientedMP, we also exploit the equivariant network for constructing LRF, as the equivariant information can be leveraged to eliminate rotations existing in data. OrientedMP utilizes this property to construct two axes of LRF based on Gram-Schmidt orthogonalization with equivariant features. However, it suffers from a learning imbalance between the two axes due to the asymmetry of their definition. Therefore, to address this issue, we select two axes of LRF that satisfy the symmetry based on the angular bisector of the equivariant features. Then two constraints including orthogonality and consistency are employed to optimize the LRF, yielding local-consistent reference frame (LCRF). Thus, LCRF is able to achieve local-consistent transformation during the construction of local rotation-invariant representation, mitigating disturbances in local geometry. As shown in Fig. 1(b), our LCRF achieves consistent orientations in both two axes for the two neighbored points. Nevertheless, since we can only achieve consistency in a local region, in the intermediate layers of the network, the relative pose between adjacent points is still changed, which is called the relative pose issue [4, 41]. We alleviate this issue by developing a relative pose recovery (RPR) module. In our RPR, we utilize the equivariant features that contain

original global information to improve the learned rotation-invariant features.

Our LocoTrans is implemented through a fusion network comprising equivariant and invariant branches, with the invariant branch relying on the equivariant one to achieve LCRF and RPR. The key contributions of this paper can be summarized as follows:

- We propose LocoTrans to achieve rotation-invariant feature learning while keeping local geometric information.
- We develop LCRF by a new definition for axes of LRF satisfying the symmetry, which is able to learn local-consistent reference frame; We also design an RPR module to address the relative pose issue.
- We conduct experiments to evaluate whether our network is rotation-invariant to arbitrary rotations in both classification and part segmentation tasks. The results show the effectiveness of our method.

## 2. Related Works

### 2.1. Point Cloud Shape Analysis

Methods for 3D point cloud shape analysis can be categorized into three primary groups: projection-based [19, 25], voxel-based [17, 30], and point-based [20, 21, 29]. The first two categories are constrained by geometric information loss and high computational costs, respectively, which currently have been explored by some works [9, 13]. In contrast, point-based methods work on point clouds without additional operations, with PointNet [20] pioneering direct processing of point sets. Subsequent works consider local context and prioritize various aspects, such as abstracting strategies [21, 44], convolution operations [35, 46], and feature aggregation [3, 32]. However, a common drawback in many existing methods is a lack of robustness to rotations, leading to a significant performance decline when confronted with arbitrarily rotated data. In our work, we address this limitation by introducing rotation invariance to the 3D network by exploring the construction of local reference frames.

### 2.2. Rotation Invariance Methods

Recently, some works [14, 33, 39] address the challenge of rotations by converting rotated point clouds to canonical poses through principal component analysis (PCA) to obtain rotation-invariant features. Li et al. [14] comprehensively summarize ambiguities in canonical poses and mitigate their impact through blending strategies. Despite these efforts, complete elimination of ambiguities remains challenging, necessitating extensive augmentation for approximate invariance. Additionally, constructing rotation-invariant features within local patches has been widely explored, involving handcrafted features with local geometry replacing Cartesian coordinates as input. RI-CNN [43]

designs a four-dimensional feature with distance and angle while PaRI-Conv [4] utilizes the normal vector to extend it to eight dimensions. In the absence of normal vectors, SGMNet [34] connects every pair of points in a local neighborhood for a full geometrical description. LRF-based methods [4, 12, 16, 41] introduce rotation invariance through building local reference frames, which can be used to transform local coordinates. However, these methods often sacrifice the original pose information between local patches, resulting in performance limitations. To address this issue, some methods [4, 41] attempt to restore relative pose with hand-crafted features. In this paper, we first alleviate relative pose changes with local-consistent reference frame and then leverage equivariant network to capture local geometry for restoring relative pose.

### 2.3. Rotation Equivariance Methods

Currently, typical solutions to equivariance can be divided into tensor field-based and vector-based. Tensor field-based methods [8, 18, 27] constrain convolutional kernels to adhere to the spherical harmonics family to achieve equivariance but are limited by significant memory requirements. In comparison, vector-based methods [5, 11, 23, 24, 26] transform standard neural network representations from scalar to vector to achieve equivariance. Typically, activation functions are often used to introduce non-linearity into the network for better learning. However, maintaining equivariance in vector-based networks requires linear layer compositions in encoders, limiting network performance. To address this limitation, some works [5, 11] explore introducing non-linearity through a combination of linear layers. VNN [5] achieves this through vector truncation operations, while GVP [11] scales vectors based on their norms. In our paper, we leverage features output by an equivariant network to enhance the learning of local rotation-invariant representation.

## 3. Method

In this paper, our objective is to learn local rotation-invariant representation with local-consistent reference frames. In the following, we first introduce the overview of our method in Sec. 3.1, then detail our Local-consistent Transformation (LocoTrans) learning in Sec. 3.2 by outlining the components of the local-consistent reference frame (LCRF) construction and the relative pose recovery (RPR) module.

### 3.1. Overview

The overall framework of our approach is shown in Fig. 2. Our LocoTrans achieves LCRF and RPR with equivariant features. Thus, in our framework, we introduce a fusion network comprising two branches: the invariant branch and the equivariant branch. Given the significant success of DGCNN [29] in the field of point cloud shape analysis,

we adopt it as the backbone for both branches. Our fusion network takes raw point clouds as input. Initially, we feed the raw points into the equivariant branch to get equivariant features that are then used to construct LCRF and used in RPR module of the invariant branch. At the output stage, the attention mechanism, widely applied in feature fusion [36, 47] to aggregate information, is employed to fuse features from the invariant branch and the equivariant branch.

Before delving into the details of our method, we establish some notations. Given a point cloud  $P = [p_1, p_2, \dots, p_N] \in \mathbb{R}^{N \times 3}$  and an arbitrary rotation  $R \in \text{SO}(3)$ , the rotation-invariant representation  $\delta(p)$  satisfies:

$$\delta(Rp) = \delta(p), \quad (1)$$

where  $p \in P$  and  $\delta(\cdot)$  is the function to build rotation-invariant representation. The equivariant features we require in our LocoTrans are generated from an equivariant network  $h$  which satisfies:

$$h(Rp) = Rh(p) = Rv, \quad (2)$$

where  $v \in \mathbb{R}^{3 \times C}$  represents the feature output from the equivariant network for each point.

### 3.2. Local-consistent Transformation Learning

To mitigate geometry perturbation caused by local inconsistency of previous LRF, LocoTrans proposes LCRF to perform a local-consistent transformation. The RPR module is also presented to further restore the relative pose. We introduce these modules separately.

**Local-consistent Reference Frame.** Given a reference point  $p_r$  in a point cloud, its LRF  $U_r \in \mathbb{R}^{3 \times 3}$  is composed of three orthonormal basis:

$$U_r = [u_{r,1}, u_{r,2}, u_{r,1} \times u_{r,2}], \quad (3)$$

where  $u_{r,1}$  and  $u_{r,2}$  are orthogonal normalized vectors. The LRF is widely used to construct local rotation-invariant representation based on each reference point and its corresponding neighbors.

To leverage equivariant features for constructing a LRF, these features are first projected to  $\tilde{v}_r = [\tilde{v}_{r,1}, \tilde{v}_{r,2}] \in \mathbb{R}^{3 \times 2}$  with a Multi-Layer Perceptron (MLP). The resulting normalized vectors,  $\tilde{v}_{r,1}$  and  $\tilde{v}_{r,2}$ , are then employed to define  $u_{r,1}$  and  $u_{r,2}$  within the LRF through the Gram-Schmidt process, as outlined below:

$$u_{r,1} = \tilde{v}_{r,1}, \quad (4)$$

$$u_{r,2} = \frac{\tilde{v}_{r,2} - \langle \tilde{v}_{r,2}, u_{r,1} \rangle u_{r,1}}{\|\tilde{v}_{r,2} - \langle \tilde{v}_{r,2}, u_{r,1} \rangle u_{r,1}\|}, \quad (5)$$

where  $\|\cdot\|$  denotes the vector norm and  $\langle \cdot, \cdot \rangle$  represents the inner product. This process is depicted in Fig. 3(a).

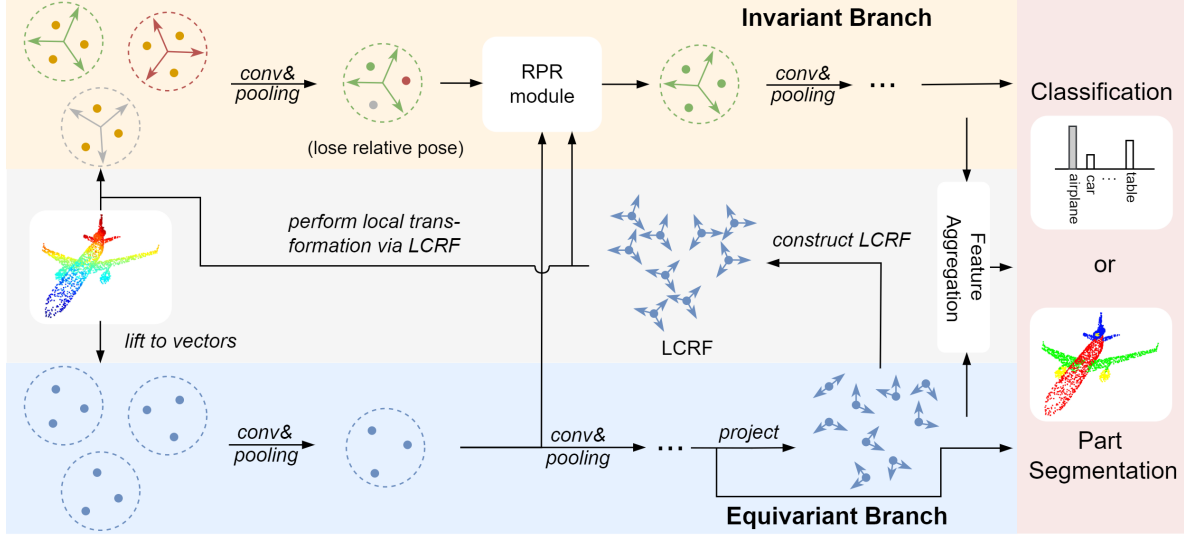


Figure 2. Overall framework of our LocoTrans. Our network consists of the invariant branch and the equivariant branch. To reduce the local geometric perturbation caused by local region transformation, we construct LCRF using the features from the equivariant branch to perform the local-consistent transformation. However, such a perturbation is difficult to completely eliminate. The RPR module is presented to alleviate this issue, restoring the relative pose between local patches. At the output level, we merge the features from two branches to aggregate information.

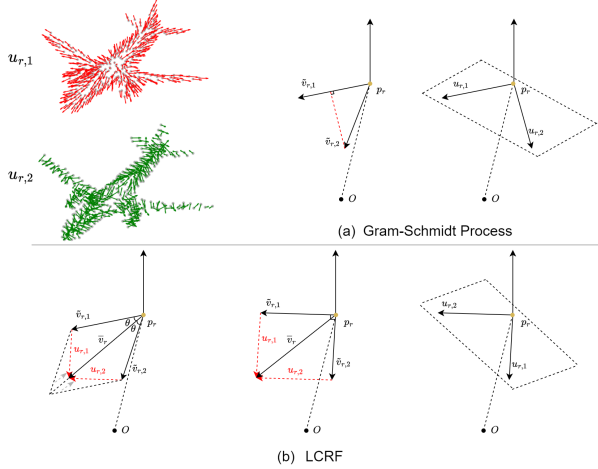


Figure 3. Comparison between the ways of generating LRF with equivariant features. (a) Gram-Schmidt process. (b) Our LCRF. The two red dashed arrows in (b) are orthogonal. We present the visualized results of  $u_{r,1}$  and  $u_{r,2}$  in LRF generated by Gram-Schmidt process.

However, as shown in Fig. 3(a),  $u_{r,1}$  maintains orientation consistency in the local area, while  $u_{r,2}$  does not. This phenomenon arises because  $u_{r,1}$  and  $u_{r,2}$ , constructed using the Gram-Schmidt process, exhibit an imbalance in learning, stemming from disparities in their definitions. Constructing local rotation-invariant representations with such inconsistent LRFs can lead to serious perturbation of local geometric relationships.

To address this issue, as shown in Fig 3(b), our LCRF introduces a vector  $\bar{v}_r$  that has the same direction as the an-

gular bisector of  $\tilde{v}_{r,1}$  and  $\tilde{v}_{r,2}$  to construct the symmetric  $u_{r,1}$  and  $u_{r,2}$ . Specifically, the length of  $\bar{v}_r$  is first defined to be equal to  $\sin \theta + \cos \theta$  to make  $u_{r,1}$  and  $u_{r,2}$  orthogonal, where  $\theta$  is half the angle between  $\tilde{v}_{r,1}$  and  $\tilde{v}_{r,2}$ . Then  $u_{r,1}$  and  $u_{r,2}$  are obtained by subtracting  $\tilde{v}_{r,1}$  and  $\tilde{v}_{r,2}$  from  $\bar{v}_r$  and normalizing to the unit vector. This process can be written as:

$$\sin \theta = \sqrt{\frac{1 - \tilde{v}_{r,1}^T \tilde{v}_{r,2}}{2}}, \cos \theta = \sqrt{\frac{1 + \tilde{v}_{r,1}^T \tilde{v}_{r,2}}{2}}, \quad (6)$$

$$\bar{v}_r = \frac{(\tilde{v}_{r,1} + \tilde{v}_{r,2})}{\|\tilde{v}_{r,1} + \tilde{v}_{r,2}\|} (\sin \theta + \cos \theta), \quad (7)$$

$$u_{r,1} = \frac{\bar{v}_r - \tilde{v}_{r,1}}{\|\bar{v}_r - \tilde{v}_{r,1}\|}, u_{r,2} = \frac{\bar{v}_r - \tilde{v}_{r,2}}{\|\bar{v}_r - \tilde{v}_{r,2}\|}. \quad (8)$$

Analysis of the orthogonality of  $u_{r,1}$  and  $u_{r,2}$  can be found in the Supplementary. Although LCRF achieves the learning balance between  $u_{r,1}$  and  $u_{r,2}$ , it still needs to be optimized for local consistency. Given a set of reference points  $P_r$  in the point cloud  $P$ , for two reference points  $p_a, p_b \in P_r$  within the local area, we construct LRFs with normalized vectors  $\tilde{v}_{a,1}, \tilde{v}_{a,2}$  and  $\tilde{v}_{b,1}, \tilde{v}_{b,2}$  from equivariant network, respectively. The LRFs constructed by the above process for  $p_a$  and  $p_b$  can be represented as  $U_a = [u_{a,1}, u_{a,2}, u_{a,1} \times u_{a,2}]$  and  $U_b = [u_{b,1}, u_{b,2}, u_{b,1} \times u_{b,2}]$ . The cosine of the angle between their axes is defined as the consistency metric between  $U_a$  and  $U_b$ . Taking  $u_{a,1}$  and

$u_{b,1}$  as an example, the computation is as follows:

$$u_{a,1}^T u_{b,1} = \frac{(\bar{v}_a - \tilde{v}_{a,1})^T (\bar{v}_b - \tilde{v}_{b,1})}{\|\bar{v}_a - \tilde{v}_{a,1}\| \|\bar{v}_b - \tilde{v}_{b,1}\|}, u_{a,1}^T u_{b,1} \in (-1, 1), \quad (9)$$

where  $\bar{v}_a$  and  $\bar{v}_b$  are obtained by Eq. 6 and Eq. 7 with  $\tilde{v}_{a,1}, \tilde{v}_{a,2}$  and  $\tilde{v}_{b,1}, \tilde{v}_{b,2}$  as input.

To learn the local-consistent LRF, we design a two-step optimization strategy to maximize the consistency. First, to better maximize the consistency of  $u_{a,1}$  and  $u_{b,1}$ , we simplify the representation of  $u_{a,1}^T u_{b,1}$ . Observing the definitions of  $u_{a,1}$  and  $u_{b,1}$ , we find that  $u_{a,1}$  and  $u_{b,1}$  can be defined with just one of the equivariant vectors when two equivariant vectors are orthogonal. Hence, for two reference points  $p_a$  and  $p_b$ , we encourage the two normalized vectors  $(\tilde{v}_1, \tilde{v}_2) \in \{(\tilde{v}_{a,1}, \tilde{v}_{a,2}), (\tilde{v}_{b,1}, \tilde{v}_{b,2})\}$  from equivariant network to be orthogonal by an orthogonality loss:

$$\mathcal{L}_{orth} = \tilde{v}_1^T \tilde{v}_2. \quad (10)$$

Once  $\tilde{v}_1$  and  $\tilde{v}_2$  are orthogonal, as shown in the center of Fig. 3(b), Eq. 9 can be rewritten as:

$$u_{a,1}^T u_{b,1} = \tilde{v}_{a,2}^T \tilde{v}_{b,2}. \quad (11)$$

In this way, maximizing the consistency between  $u_{a,1}$  and  $u_{b,1}$  is converted to maximizing the consistency between  $\tilde{v}_{a,2}$  and  $\tilde{v}_{b,2}$ . Next, we define consistency loss  $\mathcal{L}_{consist}$  to make  $\tilde{v}_1$  and  $\tilde{v}_2$  learn from each other, thereby maximizing both  $\tilde{v}_{a,1}^T \tilde{v}_{b,1}$  and  $\tilde{v}_{a,2}^T \tilde{v}_{b,2}$  as follows:

$$\mathcal{L}_{consist} = (\tilde{v}_{a,1}^T \tilde{v}_{b,1} - \tilde{v}_{a,2}^T \tilde{v}_{b,2})^2. \quad (12)$$

Maximizing the consistency between  $u_{a,2}$  and  $u_{b,2}$  can also be achieved by the above process. With our two-step optimization strategy, we can effectively enhance local consistency for both axes of LRF to alleviate the changes in relative pose.

**Relative Pose Recovery.** We can generate local rotation-invariant representation via the constructed LCRF and extract features with convolution operations. Taking DGCNN’s edge convolution [29] as an example, two consecutive convolution operations are defined as follows:

$$x_r = \max_{j \in \mathcal{N}(p_r)} \psi(U_r^T p_r, U_r^T (p_j - p_r)), \quad (13)$$

$$x'_r = \max_{j \in \mathcal{N}(x_r)} \phi(x_r, x_j - x_r), \quad (14)$$

where  $\mathcal{N}(\cdot)$  represents the neighbors around the input with KNN operation.  $x_r$  represents the feature corresponding to the reference point  $p_r$ , and  $x_j$  means the feature corresponding to  $p_j$ , the neighbor of  $p_r$ . Both  $\phi$  and  $\psi$  are MLPs. Analysis of how Eq. 13 achieves rotation invariance can be found in the Supplementary.

In Eq. 13, each reference point and its neighbors are transformed with respect to our LCRF. However, as the consistency is only maintained in local regions, the relative pose between adjacent points still changes in the intermediate layers of the network. Previous works [4, 41] have introduced a dynamic kernel  $W_j$  for each neighbor in the local patch, aiming to restore the original relative pose between neighbors and the reference point:

$$\hat{x}_j = W_j(\mathcal{F}_r^j) x_j, \quad (15)$$

where  $\mathcal{F}_r^j$  represents the original relative pose between each neighbor and the reference point. The refined edge convolution is then given by:

$$x'_r = \max_{j \in \mathcal{N}(x_r)} \phi(x_r, \hat{x}_j - x_r). \quad (16)$$

These methods obtain  $\mathcal{F}_r^j$  by hand-crafted features such as the rotation matrix and translation vector [41] and point pair feature [4]. However, these hand-crafted features may not adequately represent complex geometric relationships for constructing the relative pose. Hence, we introduce the RPR module, utilizing an equivariant network to capture essential geometric relationships for restoring the relative pose. The equivariant network can process the rotated point cloud without altering the relationship between local patches. The features extracted from it are employed to encode the correct relative pose:

$$\mathcal{F}_r^j = U_r^T (v_j - v_r), \quad (17)$$

where  $v_r$  and  $v_j$  represent the equivariant features corresponding to  $p_r$  and  $p_j$ , respectively. As shown in Eq. 2, equivariant features preserve the rotation matrix when input data is rotated, necessitating the use of the LRF to eliminate rotation.

**Training Losses.** We input the features from the invariant branch, equivariant branch, and fusion layer to three different classifiers separately for classification or segmentation. The total loss of our fusion network  $\mathcal{L}$  can be written as:

$$\mathcal{L} = \mathcal{L}_i + \mathcal{L}_e + \mathcal{L}_f + \lambda_{orth} \mathcal{L}_{orth} + \lambda_{consist} \mathcal{L}_{consist}, \quad (18)$$

where  $\mathcal{L}_i$ ,  $\mathcal{L}_e$ , and  $\mathcal{L}_f$  are the cross-entropy loss for predictions from the invariant branch, equivariant branch, and fusion layer, respectively.  $\lambda_{orth}$  and  $\lambda_{consist}$  are the corresponding hyperparameters to weight  $\mathcal{L}_{orth}$  and  $\mathcal{L}_{consist}$ .

## 4. Experiment

We evaluate the effectiveness of our method on three tasks: 3D shape classification on ModelNet40 [31] dataset (Sec. 4.2), real-world shape classification on ScanObjectNN [28] dataset (Sec. 4.3), and part segmentation on ShapeNetPart [38] dataset (Sec. 4.4). Moreover, we provide

comprehensive ablation studies and visualization of LCRF in Sec. 4.5 and Sec. 4.6. More results can be found in the Supplementary.

### 4.1. Implemented details

We use the SGD optimizer and set the initial learning rate to 0.1. For classification tasks, the learning rate will be adjusted with cosine annealing [15], and the batch size is set to 32; for the part segmentation task, the learning rate is scaled by 0.3 every 30 epochs, and the batch size is set to 16. The maximum training epoch of all tasks is set to 250.

To evaluate the performance of our method under various rotations, we follow existing training/test settings [7], i.e.,  $z/z$ ,  $z/SO(3)$ <sup>1</sup>, and  $SO(3)/SO(3)$ , to conduct experiments. Here  $z$  denotes that the input data is rotated around the vertical axis while  $SO(3)$  means rotating input arbitrarily.

### 4.2. 3D Shape Classification

**Dataset.** ModelNet40 [31] is widely used in the 3D shape analysis task, which provides 12,311 CAD models from 40 object categories. The dataset is split into 9,843 samples for training and 2,468 samples for testing. We follow prior work [19] to randomly extract 1,024 3D points from each sample for both training and testing.

**Results.** Existing methods can be categorized into three groups: rotation-sensitive, rotation-equivariant, and rotation-invariant methods. We compare our method with them in the Tab. 1. The results show that our method achieves consistent performance under three training/test settings and outperforms all methods under  $z/SO(3)$  and  $SO(3)/SO(3)$ . For rotation-sensitive methods, their performance excels in  $z/z$  but drops significantly in the  $z/SO(3)$  setting. Although in  $SO(3)/SO(3)$  setting this issue is alleviated with additional data augmentation, a notable gap still remains in comparison with rotation-invariant methods. This indicates that these methods struggle to handle unseen rotations during testing. Rotation-equivariant methods are robust to rotations, but their performance is restricted by their linear network structure. In comparison with previous rotation-invariant methods, our method can achieve rotation invariance while better preserving the local geometry information by constructing local-consistent reference frames and restoring the relative pose, leading to improved performance.

### 4.3. Real-world Shape Classification

**Dataset.** ScanObjetNN [28] is a real-world dataset containing 15,000 incomplete objects scanned from 2,902 real-world objects. To test the effectiveness of our method in the real-world scenario, we choose the OBJ\_BG subset of

<sup>1</sup>It means we apply rotation in vertical axis during training while rotating the point cloud arbitrarily for testing. The other two have similar definitions.

	Method	Input	$z/z$	$z/SO(3)$	$SO(3)/SO(3)$
Rotation-sensitive	PointNet [20] (2017)	pc	89.2	16.4	75.5
	PointNet++ [21] (2017)	pc	89.3	28.6	85.0
	PointNet++ [21] (2017)	pc+n	91.8	18.4	77.4
	DGCNN [29] (2019)	pc	92.2	20.6	81.1
Rotation-equivariant	TFN [27] (2018)	pc	88.5	85.3	87.6
	TFN-NL [18] (2021)	pc	89.7	89.7	89.7
	VN-DGCNN [5] (2021)	pc	89.5	89.5	90.2
	SVNet-DGCNN [26] (2022)	pc	90.3	90.3	90.0
Rotation-invariant	SF-CNN [22] (2019)	pc	91.4	84.8	90.1
	RI-CNN [43] (2019)	pc	86.5	86.4	86.4
	RI-GCN [12] (2020)	pc	89.5	89.5	89.5
	AECNN [41] (2020)	pc	91.0	91.0	91.0
	SGMNet [34] (2021)	pc	90.0	90.0	90.0
	Li et al. [14] (2021)	pc	90.2	90.2	90.2
	PaRI-Conv [4] (2022)	pc	91.4	91.4	91.4
	OrientedMP [16] (2022)	pc	88.4	88.4	88.9
	PaRot [40] (2023)	pc	90.9	91.0	90.8
	Ours	pc	91.6	91.6	91.5

Table 1. Classification accuracy (%) on ModelNet40 dataset. ‘pc’ and ‘n’ denote the 3D coordinates and normal vectors of input data, respectively.

ScanObjetNN which contains background noise. The subset consists of 2,890 samples in 15 categories, which are split into 2,312 samples for training and 578 samples for testing.

**Results.** As shown in Tab. 2, for the challenging real-world scenario, our method still achieves the best performance in all three settings. It is worth noting that PaRI-Conv [4], which relies on hand-crafted LRF, exhibits poorer performance than its performance on ModelNet40. The effectiveness of hand-crafted LRF is limited in real-world applications for the use of geometry information such as the global center that is sensitive to background noise. In comparison, we construct LRF with features output by equivariant branch to alleviate the impact of background noise. Although efficient, a certain absolute accuracy difference (0.5%) occurs when the training setting is different ( $z/SO(3)$  v.s.  $SO(3)/SO(3)$ ). We guess that this is because different training augmentations change the randomness of the network, which will affect learning-based LCRF. The above results demonstrate our method can also effectively tackle real-world rotations.

Method	Input	$z/z$	$z/SO(3)$	$SO(3)/SO(3)$
PointNet [20] (2017)	pc	73.3	16.7	54.7
PointNet++ [21] (2017)	pc	82.3	15.0	47.4
DGCNN [29] (2019)	pc	82.8	17.7	71.8
RI-CNN [43] (2019)	pc	-	78.4	78.1
Li et al. [14] (2021)	pc	84.3	84.3	84.3
PaRI-Conv [4] (2022)	pc	77.8	77.8	78.1
PaRot [40] (2023)	pc	-	82.1	82.6
Ours	pc	85.0	85.0	84.5

Table 2. Classification accuracy (%) on ScanObjectNN dataset.

### 4.4. Shape Part Segmentation

**Dataset.** We adopt ShapeNetPart dataset [38] to evaluate our method on the shape part segmentation task. ShapeNetPart contains 16, 881 3D samples from 16 categories, which are further subdivided into 50 parts. We refer to the data split in [20] and sample 2,048 points for each 3D object.

**Results.** In order to fairly compare with existing methods, we report our results with mean Intersection-over-Union (mIoU) over all instances and all classes together in Tab. 3. The results show our method is competitive on Class mIoU and achieves the best performance on Insta. mIoU. Furthermore, compared to our backbone DGCNN, LocoTrans significantly boosts the performance by 42.7% in z/SO(3) setting. We also give visualized segmentation results of 3D objects in Fig. 4, and the figure indicates our network is robust to arbitrary rotations in segmentation.

Method	Input	z/SO(3)		SO(3)/SO(3)	
		Class mIoU	Insta. mIoU	Class mIoU	Insta. mIoU
PointNet [20] (2017)	pc	37.8	-	74.4	-
PointNet++ [21] (2017)	pc	48.2	-	76.7	-
DGCNN [29] (2019)	pc	37.4	-	73.3	-
RI-CNN [43] (2019)	pc	75.3	-	75.5	-
AECNN [41] (2020)	pc	80.2	-	80.2	-
RI-GCN [12] (2020)	pc	77.2	-	77.3	-
VN-DGCNN [5] (2021)	pc	-	81.4	-	81.4
Li et al. [14] (2021)	pc	74.1	81.7	74.1	81.7
PaRI-Conv [4] (2022)	pc	-	83.8	-	83.8
PaRot [40] (2023)	pc	79.2	-	79.5	-
Ours	pc	80.1	84.0	80.0	83.8

Table 3. Part segmentation results on ShapeNetPart dataset. ‘Class mIoU’ stands for mIoU (%) over all classes while ‘Insta. mIoU’ denotes mIoU over all instances.

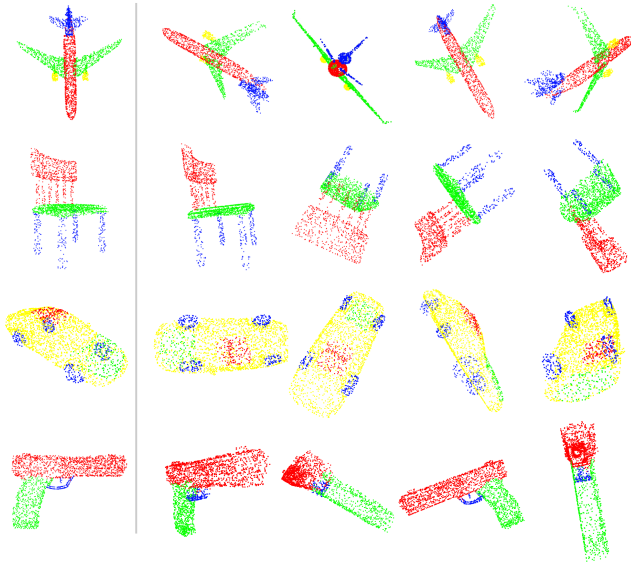


Figure 4. Visualization of part segmentation results on ShapeNetPart dataset under z/SO(3) setting. From top to bottom, we present the segmentation results for four categories: airplane, chair, car, and pistol. The leftmost column is the ground truth and the right four columns are the predictions of our network under arbitrary rotations.

#### 4.5. Ablation Studies

In this section, we conduct experiments to demonstrate the effectiveness of each component in LocoTrans. Further ablation studies are also designed on the LCRF and the RPR module for deep analysis. If no otherwise specified, all experiments are performed on ModelNet40 dataset under the z/SO(3) setting and reported with accuracy (%).

Row	Baseline	Aggregation	LCRF	RPR		z/SO(3)
				coordinate	feature	
#1	✓					89.1
#2		✓				90.4
#3		✓	✓			90.7
#4		✓		✓		90.8
#5		✓			✓	91.3
#6		✓	✓		✓	91.6

Table 4. Ablation study of our method on ModelNet40 dataset in z/SO(3) setting. ‘coordinate’ denotes the coordinates of point cloud while ‘feature’ represents the equivariant features.

**Effects of Different Components.** We carry out a series of experiments to validate our method and report the performance in Tab. 4. We take the output of the invariant branch of our network as the Baseline in Row #1 without using LCRF and RPR module. Aggregation in Row #2 refers to fusing the output of the invariant branch and the equivariant branch, which contributes to the improvement of classification accuracy. In Row #3 we replace the LRF built through the Gram-Schmidt process with our LCRF. Compared to Row #2, our LCRF improves the performance by alleviating the perturbation of local geometry with local-consistent transformation. We also test the performance of the RPR module with the original coordinates and the equivariant features as input. The results in Row #4 and Row #5 show both variations can enhance the performance of the network for addressing the loss of relative pose. Notably, the comparison between Row #4 and Row #5 highlights that features extracted by the equivariant branch outperform original coordinates in representing relative pose. Finally, our network performs optimally when all modules are in use.

**Effects of LCRF.** We also conduct experiments to explore the role of different LRFs in local rotation-invariant representation construction in Tab. 5. In Row #1, we hand-craft LRF with the global center and local barycenter, while leveraging features from the equivariant network to construct LRF based on the Gram-Schmidt process in Row #2. Due to the lack of local consistency, generating local rotation-invariant representation with either hand-crafted LRF or learning-based LRF built by Gram-Schmidt process results in worse performance. Our LCRF brings improvement by achieving the learning balance between two vectors from an equivariant network. Moreover, comparing the results on ModelNet40 and ScanObjectNN, LCRF brings more significant improvement in the latter containing background noise because data in the former is clean and easier to classify.

Row	LRF	z/SO(3)	
		ModelNet40	ScanObjectNN
#1	hand-crafted	91.2	82.8
#2	Gram-Schmidt process	91.3	84.2
#3	LCRF	91.6	85.0

Table 5. Ablation study on LRF construction.

**Effects of RPR Module.** We compare the effects of using




Figure 5. Visualization of  $u_r^1$  (Red) and  $u_r^2$  (Green) in LRF. (a) LRF generated with the Gram-Schmidt process. (b) LCRF. While the Gram-Schmidt process fails to maintain local consistency in  $u_r^2$ , Our LCRF learns local-consistent orientation in both axes.

different features to represent the relative pose in Tab. 6. In Row #1 we directly utilize the coordinates of the point cloud to represent the relative pose while in Row #2, Row #3, and Row #4, we use features extracted from the coordinates in different ways. In Row #2 hand-crafted feature means the APPF-w/oDirection proposed by PaRI-Conv [4] which encodes the relative Euclidean distance and angle in the local patch. Equivariant feature in Row #3 represents the feature extracted from the equivariant branch while invariant feature in Row #4 represents the feature extracted from the invariant branch. The comparison of Row #1 and Row #2 shows hand-crafted features cannot improve performance because they may not adequately represent essential geometric relationships. In comparison, equivariant branch can better capture the local geometry information to encode relative pose, leading to the improvement of classification accuracy in Row #3. The performance in Row #4 declines because the invariant feature itself suffers from the relative pose issue.

Row	Relative pose representation	z/SO(3)
#1	coordinate	91.0
#2	hand-crafted feature	91.0
#3	equivariant feature	91.6
#4	invariant feature	90.8

Table 6. Ablation study on relative pose representation.

**Computational Burden.** We provide the number of parameters, FLOPs, and accuracy of the DGCNN-based networks in Tab. 7. The two branches of both ‘Ours-base’ (no LCRF and no RPR, Row #2 in Tab. 4) and ‘Ours’ are built upon DGCNN (1.81M Params) and VN-DGCNN (2.89M Params). Tab. 7 shows that the two-branch structure indeed requires more costs. Note that, the fusion layer and projection layer in ‘Ours-base’ and ‘Ours’ cause extra parameters than the sum of DGCNN and VN-DGCNN. However, ‘Ours-base’ v.s. Li et al. [14] and ‘Ours-base’ v.s. PaRI-Conv [4] reveal that naively increasing parameters does not necessarily enhance performance (even degradation occurs). The ‘Ours-base’ v.s. ‘Ours’ shows that while our proposed LCRF and RPR introduce minor additional costs, they significantly improve performance, highlighting the effectiveness of our proposed strategies. Nonetheless,

the added equivariant branch does increase overall costs, prompting us to seek more cost-effective ways to leverage equivariant information in future research.

	Li et al. [14]	PaRI-Conv [4]	Ours-base	Ours
Params	2.91M	1.85M	6.22M	6.27M
FLOPs	3747M	1938M	7552M	7998M
Acc.(ModelNet40)	90.2	91.4	90.4	91.6
Acc.(ScanObjectNN)	84.3	77.8	82.8	85.0

Table 7. Computational burden on the classification network.

#### 4.6. Visualization of LCRF

We visualize two types of LRFs in Fig. 5. For the Gram-Schmidt process, it can always maintain local-consistent orientation in  $u_r^1$  but not in  $u_r^2$ . With these kinds of LRFs, points on similar local structures are transformed into different reference frames, which is not conducive for network to learn the local structure. In contrast, our LCRF maintains consistency across different axes and preserves the local geometry.

### 5. Conclusion

In this paper, we proposed Local-consistent Transformation (LocoTrans) learning to effectively achieve local rotation-invariant representation. LocoTrans is built upon the equivariant network and consists of two modules. Specifically, the LCRF constructs local-consistent reference frames to preserve local geometry relationships when performing transformation via LRF. However, the relative pose between adjacent points still changes. To further restore relative pose, our RPR module leverages the equivariant network to encode pose information from the original local coordinates and fuses pose information with neighbor features. Experimental results demonstrate the effectiveness of our method.

### 6. Acknowledgment

This work was supported by the National Natural Science Foundation of China under Grant 62076101, Guangdong Basic and Applied Basic Research Foundation under Grant 2023A1515010007, Guangdong Provincial Key Laboratory of Human Digital Twin under Grant 2022B1212010004, CAAI-Huawei MindSpore Open Fund, and TCL Young Scholars Program.



## References

- [1] Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga Rodrigo. Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9902–9912, 2022. [1](#)
- [2] Chao Chen, Guanbin Li, Ruijia Xu, Tianshui Chen, Meng Wang, and Liang Lin. Clusternet: Deep hierarchical cluster network with rigorously rotation-invariant representation for point cloud analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [1](#)
- [3] Jiajing Chen, Burak Kakillioglu, Huantao Ren, and Senem Velipasalar. Why discard if you can recycle?: A recycling max pooling module for 3d point cloud analysis. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 559–567, 2022. [2](#)
- [4] Ronghan Chen and Yang Cong. The devil is in the pose: Ambiguity-free 3d rotation-invariant learning via pose-aware convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7472–7481, 2022. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [5] Congyue Deng, Or Litany, Yueqi Duan, Adrien Poulencard, Andrea Tagliasacchi, and Leonidas J. Guibas. Vector neurons: A general framework for so(3)-equivariant networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12200–12209, 2021. [3](#), [6](#), [7](#)
- [6] Lunhao Duan, Shanshan Zhao, Nan Xue, Mingming Gong, Gui-Song Xia, and Dacheng Tao. Condaformer: Disassembled transformer with local structure enhancement for 3d point cloud understanding. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. [1](#)
- [7] Carlos Esteves, Christine Allen-Blanchette, Ameesh Makadia, and Kostas Daniilidis. Learning so(3) equivariant representations with spherical cnns. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018. [6](#)
- [8] Fabian Fuchs, Daniel Worrall, Volker Fischer, and Max Welling. Se(3)-transformers: 3d roto-translation equivariant attention networks. In *Advances in Neural Information Processing Systems*, pages 1970–1981. Curran Associates, Inc., 2020. [3](#)
- [9] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. *CVPR*, 2018. [2](#)
- [10] Zhao Jin, Munawar Hayat, Yuwei Yang, Yulan Guo, and Yinjie Lei. Context-aware alignment and mutual masking for 3D-language pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10984–10994, 2023. [1](#)
- [11] Bowen Jing, Stephan Eismann, Patricia Suriana, Raphael John Lamarre Townshend, and Ron Dror. Learning from protein structure with geometric vector perceptrons. In *International Conference on Learning Representations*, 2021. [3](#)
- [12] SEOHYUN KIM, JaeYoo Park, and Bohyung Han. Rotation-invariant local-to-global representation learning for 3d point cloud. In *Advances in Neural Information Processing Systems*, pages 8174–8185. Curran Associates, Inc., 2020. [1](#), [3](#), [6](#), [7](#)
- [13] Roman Klokov and Victor 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 (ICCV)*, 2017. [2](#)
- [14] Feiran Li, Kent Fujiwara, Fumio Okura, and Yasuyuki Matsushita. A closer look at rotation-invariant deep point cloud analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16218–16227, 2021. [2](#), [6](#), [7](#), [8](#)
- [15] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [6](#)
- [16] Shitong Luo, Jiahan Li, Jiaqi Guan, Yufeng Su, Chaoran Cheng, Jian Peng, and Jianzhu Ma. Equivariant point cloud analysis via learning orientations for message passing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18932–18941, 2022. [1](#), [2](#), [3](#), [6](#)
- [17] Daniel Maturana and Sebastian Scherer. Voxnet: A 3d convolutional neural network for real-time object recognition. In *2015 IEEE/RSJ international conference on intelligent robots and systems (IROS)*, pages 922–928. IEEE, 2015. [2](#)
- [18] Adrien Poulencard and Leonidas J. Guibas. A functional approach to rotation equivariant non-linearities for tensor field networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13174–13183, 2021. [3](#), [6](#)
- [19] Charles R. Qi, Hao Su, Matthias Niessner, Angela Dai, Mengyuan Yan, and Leonidas J. Guibas. Volumetric and multi-view cnns for object classification on 3d data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [2](#), [6](#)
- [20] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [1](#), [2](#), [6](#), [7](#)
- [21] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*. Curran Associates, Inc., 2017. [1](#), [2](#), [6](#), [7](#)
- [22] Yongming Rao, Jiwen Lu, and Jie Zhou. Spherical fractal convolutional neural networks for point cloud recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [6](#)
- [23] Victor Garcia Satorras, Emiel Hoogeboom, and Max Welling. E (n) equivariant graph neural networks. In *International conference on machine learning*, pages 9323–9332. PMLR, 2021. [3](#)
- [24] Kristof Schütt, Oliver Unke, and Michael Gastegger. Equivariant message passing for the prediction of tensorial prop-

- erties and molecular spectra. In *International Conference on Machine Learning*, pages 9377–9388. PMLR, 2021. 3
- [25] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, 2015. 2
- [26] Zhuo Su, Max Welling, Matti Pietikäinen, and Li Liu. Svnnet: Where so (3) equivariance meets binarization on point cloud representation. In *International Conference on 3D Vision*, 2022. 3, 6
- [27] Nathaniel Thomas, Tess Smidt, Steven Kearnes, Lusann Yang, Li Li, Kai Kohlhoff, and Patrick Riley. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds. *arXiv preprint arXiv:1802.08219*, 2018. 3, 6
- [28] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 5, 6
- [29] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 2019. 1, 2, 3, 5, 6, 7
- [30] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 2
- [31] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. 5, 6
- [32] Tiange Xiang, Chaoyi Zhang, Yang Song, Jianhui Yu, and Weidong Cai. Walk in the cloud: Learning curves for point clouds shape analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 915–924, 2021. 2
- [33] Zelin Xiao, Hongxin Lin, Renjie Li, Lishuai Geng, Hongyang Chao, and Shengyong Ding. Endowing deep 3d models with rotation invariance based on principal component analysis. In *2020 IEEE international conference on multimedia and expo (ICME)*, pages 1–6. IEEE, 2020. 2
- [34] Jianyun Xu, Xin Tang, Yushi Zhu, Jie Sun, and Shiliang Pu. Sgmnet: Learning rotation-invariant point cloud representations via sorted gram matrix. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10468–10477, 2021. 1, 3, 6
- [35] Mutian Xu, Runyu Ding, Hengshuang Zhao, and Xiaojuan Qi. Paconv: Position adaptive convolution with dynamic kernel assembling on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3173–3182, 2021. 2
- [36] Xu Yan, Jiantao Gao, Chaoda Zheng, Chao Zheng, Ruimao Zhang, Shuguang Cui, and Zhen Li. 2dpass: 2d priors assisted semantic segmentation on lidar point clouds. In *European Conference on Computer Vision*, pages 677–695. Springer, 2022. 3
- [37] Yuwei Yang, Munawar Hayat, Zhao Jin, Chao Ren, and Yinjie Lei. Geometry and uncertainty-aware 3D point cloud class-incremental semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21759–21768, 2023. 1
- [38] Li Yi, Vladimir G Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *ACM Transactions on Graphics (ToG)*, 35(6):1–12, 2016. 5, 6
- [39] Ruixuan Yu, Xin Wei, Federico Tombari, and Jian Sun. Deep positional and relational feature learning for rotation-invariant point cloud analysis. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part X 16*, pages 217–233. Springer, 2020. 2
- [40] Dingxin Zhang, Jianhui Yu, Chaoyi Zhang, and Weidong (Tom) Cai. Parot: Patch-wise rotation-invariant network via feature disentanglement and pose restoration. In *AAAI Conference on Artificial Intelligence*, 2023. 6, 7
- [41] Junming Zhang, Ming-Yuan Yu, Ram Vasudevan, and Matthew Johnson-Roberson. Learning rotation-invariant representations of point clouds using aligned edge convolutional neural networks. In *2020 International Conference on 3D Vision (3DV)*, pages 200–209. IEEE, 2020. 1, 2, 3, 5, 6, 7
- [42] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8552–8562, 2022. 1
- [43] Zhiyuan Zhang, Binh-Son Hua, David W. Rosen, and Sai-Kit Yeung. Rotation invariant convolutions for 3d point clouds deep learning. In *International Conference on 3D Vision (3DV)*, 2019. 1, 2, 6, 7
- [44] Hengshuang Zhao, Li Jiang, Chi-Wing Fu, and Jiaya Jia. Pointweb: Enhancing local neighborhood features for point cloud processing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [45] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip H.S. Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16259–16268, 2021. 1
- [46] Haoran Zhou, Yidan Feng, Mingsheng Fang, Mingqiang Wei, Jing Qin, and Tong Lu. Adaptive graph convolution for point cloud analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4965–4974, 2021. 2
- [47] Zhuangwei Zhuang, Rong Li, Kui Jia, Qicheng Wang, Yuanqing Li, and Mingkui Tan. Perception-aware multi-sensor fusion for 3d lidar semantic segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 16280–16290, 2021. 3