

RI-MAE: Rotation-Invariant Masked AutoEncoders for Self-Supervised Point Cloud Representation Learning

Kunming Su¹, Qiuxia Wu^{1*}, Panpan Cai¹, Xiaogang Zhu²,
Xuequan Lu³, Zhiyong Wang⁴, Kun Hu⁴

¹South China University of Technology

²The University of Adelaide

³The University of Western Australia

⁴The University of Sydney

kunmingsu07@gmail.com, {qxwu, caipp}@scut.edu.cn, xiaogang.zhu@adelaide.edu.au,
bruce.lu@uwa.edu.au, {zhiyong.wang, kun.hu}@sydney.edu.au

Abstract

Masked point modeling methods have recently achieved great success in self-supervised learning for point cloud data. However, these methods are sensitive to rotations and often exhibit sharp performance drops when encountering rotational variations. In this paper, we propose a novel Rotation-Invariant Masked AutoEncoders (RI-MAE) to address two major challenges: 1) achieving rotation-invariant latent representations, and 2) facilitating self-supervised reconstruction in a rotation-invariant manner. For the first challenge, we introduce RI-Transformer, which features disentangled geometry content, rotation-invariant relative orientation and position embedding mechanisms for constructing rotation-invariant point cloud latent space. For the second challenge, a novel dual-branch student-teacher architecture is devised. It enables the self-supervised learning via the reconstruction of masked patches within the learned rotation-invariant latent space. Each branch is based on an RI-Transformer, and they are connected with an additional RI-Transformer predictor. The teacher encodes all point patches, while the student solely encodes unmasked ones. Finally, the predictor predicts the latent features of the masked patches using the output latent embeddings from the student, supervised by the outputs from the teacher. Extensive experiments demonstrate that our method is robust to rotations, achieving the state-of-the-art performance on various downstream tasks.

Code — <https://github.com/kunmingsu07/RI-MAE>.

Technical Appendices — <https://arxiv.org/abs/2409.00353>.

1 Introduction

Pioneered by PointNet (Qi et al. 2017a), deep neural networks have achieved great success in point cloud representation learning, and shown impressive results on various downstream tasks such as point cloud classification and segmentation (Qi et al. 2017b; Wang et al. 2019; Li et al. 2018). However, such approaches are supervised and necessitate data annotated with labor-intensive efforts for training models. In contrast, self-supervised learning methods learn to derive representation from unlabeled data by designing pretext

*Corresponding author




Figure 1: Illustration of the consistent relative orientation between point patches. A point cloud can be partitioned into several point patches, each of which can be regarded as an aligned patch rotated from a canonical pose with a specific rotation. While the overall pose of the point cloud can undergo various rotations, the relative rotation ΔR between any two patches remains constant and thus is rotation-invariant.

tasks, achieving notable success in the fields of natural language processing (NLP) and computer vision (CV). Inspired by this success, many self-supervised point cloud representation learning approaches have been proposed (Yang et al. 2018; Liu et al. 2019; Afham et al. 2022; Chen et al. 2021; Poursaeed et al. 2020; Rao, Lu, and Zhou 2020; Sanghi 2020; Tsai et al. 2022; Sun et al. 2022). They learn to acquire point cloud representations through a variety of tasks, such as diverse forms of point cloud reconstruction and contrastive learning tasks.

Recently, building on the success of masked language modeling (MLM) and masked image modeling (MIM), as exemplified by BERT (Devlin et al. 2018) and Masked AutoEncoder (MAE) (He et al. 2021), several studies have introduced transformer-based masked point modeling (MPM) approaches (Yu et al. 2022; Pang et al. 2022; Zhang et al. 2022; Dong et al. 2023; Zhang et al. 2023; Qi et al. 2024; Chen et al. 2024a; Qi et al. 2023). They divide the input point cloud into point patches and apply a substantial mask-

ing ratio to randomly mask these patches. Subsequently, a transformer-based architecture is utilized to encode the unmasked patches and to reconstruct the masked ones, serving as a self-supervised task in the pretraining phase. Thanks to the capabilities of the transformer model and the strategically designed pretext tasks, the encoder is capable of capturing high-level point cloud features from partially unmasked data. This process leads to remarkable enhancements in performance on downstream tasks through fine-tuning.

Nonetheless, existing MPM methods are sensitive to rotations. They overlook the effects of point cloud orientation, assuming that the input point clouds are pre-aligned. This assumption does not hold in real-world scenarios, where point clouds typically involve arbitrary rotations due to varying capture settings. As a result, the performance on downstream tasks significantly drops when faced with rotated point cloud data. To address this essential problem, this study aims to devise a rotation-invariant MPM approach for point cloud data. We begin by identifying and analyzing two primary challenges to integrate rotation invariance into MPM:

1. *How to achieve rotation-invariant latent representations?* Existing methods attain rotation invariance by transforming the input point cloud into handcrafted features (Chen et al. 2019; Li et al. 2021b), designing rotation-invariant convolutions (Chen and Cong 2022; Zhang, Hua, and Yeung 2022), or establishing global reference frames (Zhao et al. 2022; Li et al. 2021a). Nevertheless, directly integrating these solutions into transformer-based MPM poses significant challenges, as they demand specialized network designs that are inherently incompatible with the transformer framework. Despite the introduction of a rotation-invariant transformer in (Yu, Zhang, and Cai 2023), its complex architecture hinders the development of a universally applicable MPM framework. Hence, instead of making extensive modifications to the Transformer, we aim to achieve rotation invariance with minimal modifications, preserving the advantages of transformer-based MPM architectures.
2. *How to reconstruct masked patches in a rotation-invariant manner?* Conventional self-supervised reconstruction strategy for MPM is directly guided by the the raw input point cloud, which contains orientation patterns and is not applicable for constructing a rotation-invariant MPM approach. For two point clouds that are identical in content but differ in orientation, their learned latent representations should be identical. However, a conventional reconstruction process is expected to yield two distinct results when supervision is applied to revert them back to their original input orientations. Such ambiguous supervision for reconstruction causes discrepancies in losses and less desired reconstructions. As such, our goal is to devise effective reconstruction mechanisms that is not within the raw input space.

Based on these insights, we propose a novel Rotation Invariant Masked AutoEncoders for self-supervised learning on point clouds, namely RI-MAE. Addressing the first challenge, we propose a rotation-invariant transformer, RI-

Transformer, for point clouds, as the backbone of RI-MAE. Specifically, a content-orientation disentanglement mechanism is introduced to align point patches to their canonical poses for rotation-invariant representations. To assist this representation learning process, a rotation-invariant orientation embedding (RI-OE) and a rotation-invariant position embedding (RI-PE) are proposed. RI-OE leverages the rotation-invariant nature of the relative orientation between two patches as illustrated in Fig. 1. The orientation of a point patch is defined by a rotation matrix relative to its canonical pose. RI-PE adopts the center coordinates and rotation matrix of a point patch to address the issue of conventional rotation-sensitive point cloud position embeddings. With these compatible mechanisms, our RI-Transformer inherently supports rotation invariance, enhancing the usage of general transformers for the learning on point clouds.

For the second challenge, we propose a novel dual-branch student-teacher architecture for self-supervised reconstruction of masked point patches within a learned rotation-invariant latent space shared by the two branches. Specifically, the dual-branch consists of two RI-Transformers, each corresponding to one branch, and they are connected with an additional RI-Transformer predictor. The teacher branch processes and encodes all point patches, while the student branch solely addresses the visible set. Subsequently, the predictor is employed to predict the features of the masked point patches using the student latent embeddings of visible point patches, along with the supervision from the teacher branch’s outputs.

Our main contributions are summarized as follows:

1. We propose RI-Transformer with disentangled geometry content, and rotation-invariant patch orientation and position embedding mechanisms to learn rotation-invariant latent representations for point clouds.
2. We propose a dual-branch architecture, RI-MAE, based on the RI-Transformer backbone, enabling self-supervised masked patch reconstruction within a learned rotation-invariant latent space.
3. Comprehensive experiments demonstrate the state-of-the-art performance on multiple downstream tasks.

2 Related Work

2.1 Rotation-Invariant Point Cloud Analysis

Deep neural networks have become a cornerstone for learning point cloud representations, giving rise to numerous methods (Qi et al. 2017a,b; Wang et al. 2019; Li et al. 2018). Despite their impressive success across a variety of downstream tasks, these methods exhibit a notable sensitivity to point cloud rotations. In response to the challenges posed by rotation perturbations, a number of rotation-robust methods have been proposed, based on techniques such as spherical convolutional neural networks and their variants (Estevés et al. 2018; You et al. 2021), regular icosahedral lattice (Rao, Lu, and Zhou 2019), rotation conditions (Qiu et al. 2022), tensor field networks (Poulenard and Guibas 2021) and adversarial training (Wang, Yang, and Tao 2022). These methods have marked encouraging improvements in handling rotated point clouds. Nevertheless, these methods continue to

face challenges in scenarios when confronted with data subjected to arbitrary rotations since they fail to obtain identical representations from point clouds in different poses.

To learn consistent features across a range of potential rotations, many methods (Chen et al. 2019; Li et al. 2021b; Gu et al. 2021) transform the input point cloud into rotation-invariant, handcrafted features, such as distances and angles. These features, derived from point coordinates and normal vectors, remain independent regarding the global orientation of point clouds. Drawing inspiration from 2D image representation learning, several methods (Chen and Cong 2022; Zhang, Hua, and Yeung 2022) introduced rotation-invariant convolutions to extract point cloud features. However, by solely focusing on distances and angles, these methods may lose other vital information. To address this, some methods align point clouds to canonical poses, thereby learning rotation-invariant representations through the establishment of global reference frames (Li et al. 2021a; Zhang et al. 2020; Lou et al. 2023). Unlike these existing approaches generally requiring complex feature engineering and network designs, our method eliminates the need for intricate feature or network designs with a simplified process, and allows for training without the necessity of annotating data.

2.2 Self-Supervised Learning for Point Clouds

Self-supervised learning for point clouds has been explored extensively, leading to many promising approaches. These approaches train models with carefully designed pretext tasks instead of using labeled data. Common pretext tasks include point cloud reconstruction efforts such as input reconstruction (Yang et al. 2018), local-to-global reconstruction (Liu et al. 2019), restoration of anomalous parts (Chen et al. 2021), and completion of occlusions (Wang et al. 2020). Additionally, various contrastive learning approaches have been explored (Xie et al. 2020; Rao, Lu, and Zhou 2020; Sanghi 2020; Afham et al. 2022), which learn point cloud representations by differentiating between pre-defined positive and negative samples. Other pretext tasks for self-supervised learning include orientation estimation (Poursead et al. 2020), mixing and disentangling techniques (Sun et al. 2022), and pose disentanglement strategies (Tsai et al. 2022).

Recently, inspired by the great achievement of BERT (Devlin et al. 2018) and MAE (He et al. 2021), many Transformer-based MPM methods for point cloud have been proposed and significantly outperform other methods in downstream tasks. Point-BERT (Yu et al. 2022) introduced the BERT pretraining scheme to point clouds, which trains models with a task to predict discrete tokens generated by dVAE (Rolfe 2016). Point-MAE (Pang et al. 2022) proposed a self-supervised pretext task to reconstruct masked input point patches based on the known tokens. Point-M2AE (Zhang et al. 2022) employed a pyramid encoder and decoder design to conduct masked point cloud reconstruction in a hierarchical manner. ACT (Dong et al. 2023), I2P-MAE (Zhang et al. 2023), ReCon (Qi et al. 2023) and ReCon++ (Qi et al. 2024) improved the 3D masked modeling by introducing pre-trained model from other modalities, such as images and natural languages. Point-GPT (Chen et al. 2024a)

arranged point patches in an ordered sequence and extended the concept of generative pre-training transformer (GPT) to point clouds. Despite achieving remarkable results, these MPM methods are sensitive to rotation, constraining their scalability and generalizability.

3 Method

3.1 Revisiting Masked Point Modeling

We first revisit the masked point modeling (MPM) framework for point clouds. Existing methods, such as Point-MAE, segment the input point cloud \mathbf{X} into G point patches $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^G$ using Farthest Point Sampling (FPS) and K-Nearest Neighborhood (KNN) algorithms. Subsequently, a masking ratio α is applied, partitioning the point patches into a visible set $\mathbf{P}_v = \{\mathbf{p}_i^v\}_{i=1}^V$ and a masked set $\mathbf{P}_m = \{\mathbf{p}_i^m\}_{i=1}^M$, where $V = (1 - \alpha) \times G$ and $M = \alpha \times G$ are the numbers of visible and masked patches, respectively. Following this, a mini-PointNet (Qi et al. 2017a) is utilized to tokenize the visible point patches as point tokens $\mathbf{T}_v = \{\mathbf{t}_i^v\}_{i=1}^V$ and the coordinates of patch centers are projected as position embeddings $\{\mathbf{pos}_i\}_{i=1}^G$ via a learnable MLP (multi-layer perceptron). To this end, an encoding stage is introduced, during which only the visible tokens \mathbf{T}_v are fed into a transformer-based encoder:

$$\mathbf{T}_e = \text{Encoder}(\mathbf{T}_v). \quad (1)$$

In the decoding stage, a full set of tokens consisting of encoded visible tokens \mathbf{T}_e and learnable mask tokens $\mathbf{T}_m = \{\mathbf{t}_i^m\}_{i=1}^M$ are fed into a transformer-based decoder. Every masked token shares a learned masked embedding vector \mathbf{t}^m , matching the dimension of \mathbf{t}_i^v . The decoder’s output for these masked tokens is denoted as $\mathbf{H}^m = \{\mathbf{h}_i^m\}_{i=1}^M$. Finally, a prediction head that consists of fully connected (FC) layers, formulates the point coordinates of the masked point patches \mathbf{P}_m based on \mathbf{H}^m as follows:

$$\hat{\mathbf{P}}_m = \text{FC}(\mathbf{H}^m) = \text{FC}(\text{Decoder}(\mathbf{T}_e, \mathbf{T}_m)), \quad (2)$$

where $\hat{\mathbf{P}}_m$ can be viewed as the estimation of \mathbf{P}_m . For simplicity, the position embeddings, which are integrated with the respective tokens in both the encoder and the decoder to indicate location context, are not depicted in Eq. (1) - (2).

We will next explain our proposed method, which is illustrated in Fig. 2. We first present a rotation-invariant transformer backbone, RI-Transformer for point clouds; and then a rotation-invariant MPM approach, RI-MAE, is introduced.

3.2 Rotation-Invariant Transformer

We introduce a rotation-invariant transformer, namely RI-Transformer, as an effective representation learning backbone for point clouds. The overall framework is shown in Fig. 2. Our RI-Transformer aims to decompose the point patches into 1) geometric content tokens and 2) orientation embeddings; and to provide location information with 3) rotation-invariant position embeddings.

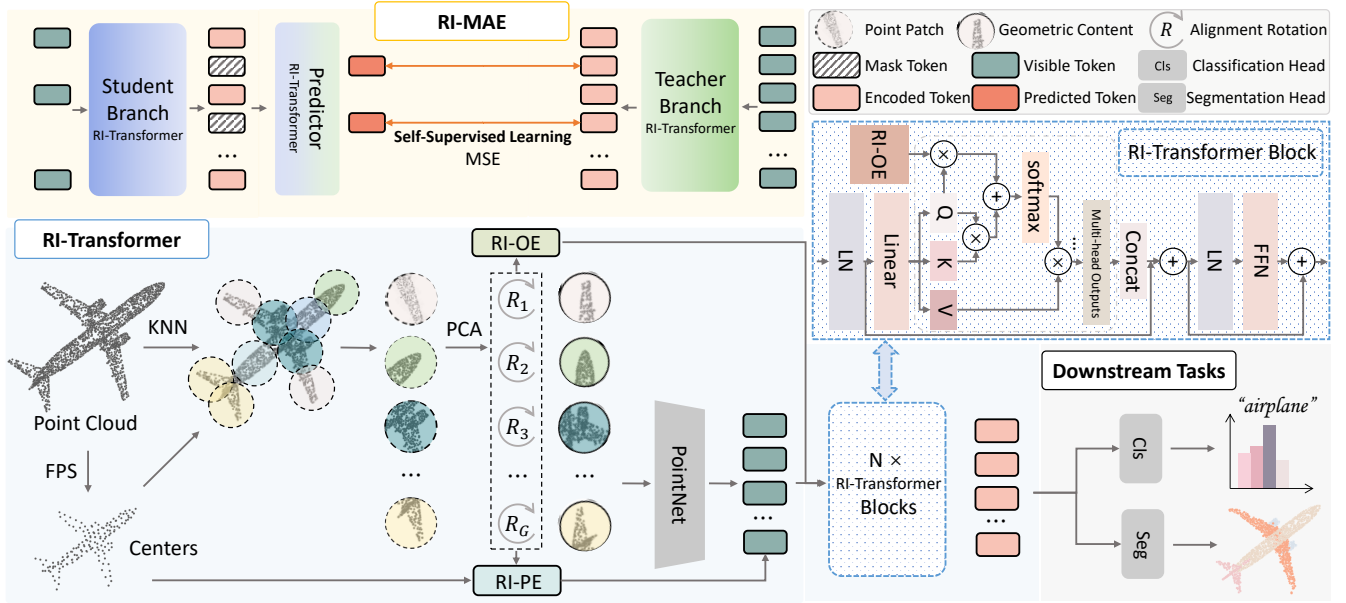


Figure 2: Overview of the proposed RI-MAE architecture. The input point cloud is first divided into point patches via FPS and KNN, and PCA is utilized to align the patches and obtain rotation matrices relative to canonical poses. Then geometry content tokens, RI-OEs, and RI-PEs are formulated as RI-Transformer’s inputs. Finally, task heads are employed for downstream tasks. A dual-branch student-teacher scheme is devised to conduct the self-supervised learning to pretrain the RI-Transformer.

Content-Oriented Disentanglement & Content Tokenization This disentanglement mechanism allows our framework to learn and utilize the geometric content and orientation knowledge in a flexible manner. They can be further formulated in a transformer architecture, enhancing the pipeline’s capabilities beyond what conventional MPM methods offer.

For point cloud \mathbf{X} , each of its point patches $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^G$ can be viewed as having been rotated from a canonical pose with a specific rotation, i.e.,

$$\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^G = \{\bar{\mathbf{p}}_i \mathbf{R}_i\}_{i=1}^G, \quad (3)$$

where $\bar{\mathbf{p}}_i$ is the point patch \mathbf{p}_i in the canonical pose, and \mathbf{R}_i is the rotation that rotates $\bar{\mathbf{p}}_i$ into \mathbf{p}_i . As a result, a point patch can be decomposed into two components: the canonical content information $\bar{\mathbf{p}}_i$, which is rotation-invariant, and the orientation information \mathbf{R}_i .

To obtain $\bar{\mathbf{p}}_i$ and \mathbf{R}_i , we conduct principal component analysis (PCA) to \mathbf{p}_i , which is widely utilized for transforming point cloud into canonical pose (Li et al. 2021a; Xie et al. 2023; Yu, Zhang, and Cai 2023; Kim, Park, and Han 2020; Xiao et al. 2020; Yu et al. 2020). Next, we utilize a mini-PointNet and an MLP to encode $\bar{\mathbf{p}}_i$ and \mathbf{R}_i into the rotation-invariant token $\bar{\mathbf{t}}_i$ and the orientation embedding \mathbf{ori}_i , respectively. Formally, we have:

$$\bar{\mathbf{p}}_i, \mathbf{R}_i = \text{PCA}(\mathbf{p}_i), \quad (4)$$

$$\bar{\mathbf{t}}_i = \text{PointNet}(\bar{\mathbf{p}}_i), \quad (5)$$

$$\mathbf{ori}_i = \text{MLP}(\mathbf{R}_i). \quad (6)$$

Rotation-Invariant Orientation Embedding (RI-OE) Importantly, the rotation \mathbf{R}_i , which aligns \mathbf{p}_i to its canon-

ical form $\bar{\mathbf{p}}_i$, varies when the point set \mathbf{P} undergoes a rotation characterized by \mathbf{R} . To elaborate, we have the following relationship:

$$\bar{\mathbf{p}}_i, \mathbf{R}'_i = \text{PCA}(\mathbf{p}_i \mathbf{R}), \mathbf{R}'_i = \mathbf{R}_i \mathbf{R}, \quad (7)$$

where \mathbf{R}'_i is the rotation that rotates the canonical $\bar{\mathbf{p}}_i$ into the rotated point patch $\mathbf{p}_i \mathbf{R}$. Hence, the orientation embedding \mathbf{ori}_i is affected as well, since the MLP itself does not inherently ensure rotation invariance.

Fully discarding orientation information and solely retaining geometric content intuitively maintain a rotation-invariant approach. However, this also results in the loss of relative orientation knowledge between point patches within a point cloud, and undermines the quality of point cloud representations. Fortunately, the relative orientation between two patches remains invariant to rotation \mathbf{R} and can be conveniently expressed using our decomposed orientation component, as shown in Fig. 1.

In detail, the relative orientation between patches \mathbf{p}_i and \mathbf{p}_j can be formulated as $\mathbf{R}_{ij} = \mathbf{R}_j \mathbf{R}_i^T$, where $i, j \in \{1, 2, \dots, G\}$. Upon applying the rotation \mathbf{R} to \mathbf{X} and \mathbf{P} : $\mathbf{X}' = \mathbf{X} \mathbf{R}$ and $\mathbf{P}' = \mathbf{P} \mathbf{R} = \{\mathbf{p}_i \mathbf{R}\}_{i=1}^G = \{\bar{\mathbf{p}}_i \mathbf{R}'_i\}_{i=1}^G$ can be obtained. Subsequently, regarding \mathbf{p}'_i and \mathbf{p}'_j , we have the relative orientation \mathbf{R}'_{ij} between them:

$$\begin{aligned} \mathbf{R}'_{ij} &= \mathbf{R}'_j \mathbf{R}'_i{}^T = (\mathbf{R}_j \mathbf{R})(\mathbf{R}_i \mathbf{R})^T \\ &= \mathbf{R}_j \mathbf{R} \mathbf{R}^T \mathbf{R}_i^T = \mathbf{R}_j \mathbf{R}_i^T = \mathbf{R}_{ij}. \end{aligned} \quad (8)$$

Therefore, we embed this relative patch orientation as the rotation-invariant information for characterizing the relations between patches. For this purpose, inspired by (Wu

et al. 2021), we devise a rotation-invariant orientation embedding (RI-OE) mechanism for self-attentions. Given the input token $\bar{\mathbf{T}} = \{\bar{\mathbf{t}}_i\}_{i=1}^G$, its self-attention with RI-OE is formulated as:

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{softmax}\left(\frac{\mathbf{Q}\mathbf{K}^T + \mathbf{B}}{\sqrt{d_k}}\right)\mathbf{V}, \quad (9)$$

where $\mathbf{Q} = \bar{\mathbf{T}}\mathbf{W}^Q$, $\mathbf{K} = \bar{\mathbf{T}}\mathbf{W}^K$, $\mathbf{V} = \bar{\mathbf{T}}\mathbf{W}^V$ are the query, key and value matrices; \mathbf{W}^Q , \mathbf{W}^K and \mathbf{W}^V are learnable matrices; and $\mathbf{B} = \{b_{ij}\} \in \mathbb{R}^{G \times G}$ is RI-OE. We formulate b_{ij} by establishing bilinear relations that integrate the features of the i^{th} patch token with the relative orientation information pertaining to the j^{th} patch:

$$b_{ij} = \bar{\mathbf{t}}_i \mathbf{W}^Q \mathbf{r}_{ij}^T, \quad (10)$$

$$\mathbf{r}_{ij} = \text{MLP}(\mathbf{R}_{ij}). \quad (11)$$

Note that for computational efficiency, we adopt the query weight matrix for the bilinear operator empirically. Instead of predicting the relative rotation with learnable module as done in existing method (Deng et al. 2023), our RI-OE directly employs rotation matrices to delineate the relative orientation, thereby avoiding intricate network designs.

Rotation-Invariant Position Embedding (RI-PE) Position embeddings offer critical location information for each point patch in MPM methods. Existing studies typically embed the coordinates of patch centers with MLPs. However, this practice is sensitive to rotations, as the coordinates of centres significantly alter with the rotation of point clouds. To address this, we introduce the RI-PE mechanism.

Specifically, for a point patch \mathbf{p}_i with center coordinate \mathbf{c}_i , we utilize its rotation matrix \mathbf{R}_i . The position embedding ripos_i for \mathbf{p}_i can be formulated as:

$$\text{ripos}_i = \text{MLP}(\mathbf{c}_i \mathbf{R}_i^T), \quad (12)$$

where $\mathbf{c}_i \mathbf{R}_i^T$ represents the location of the origin in the local reference frame of \mathbf{p}_i . The rotation invariance of ripos_i is discussed in the technical appendix.

To this end, given the proposed mechanisms including RI-OE and RI-PE, our novel transformer architecture, RI-Transformer, can be constructed with the property of rotation invariance. Compare with existing rotation-invariant methods, our RI-Transformer effectively derives two embeddings rooted in content-orientation disentanglement to preserve the intrinsic relative pose and local positional information. This helps avoid complex feature and network design as in existing methods, and maintain the capability to adopt any general transformer architectures seamlessly.

3.3 Rotation-Invariant MAE

Existing MPM methods utilize raw point clouds to guide the reconstruction of masked patches. However, this approach is not suitable for building a rotation-invariant MPM. Specifically, for a given point cloud in different directions, a rotation-invariant MPM should yield identical point cloud latent representations; however, the coordinates of the masked patches are different, necessitating that the decoder reconstructs distinct patches from the same latent representation. This conflict results in unstable loss, making

the model difficult for optimization. We provide theoretical proof for this in the technical appendix.

Therefore, we introduce a rotation-invariant MAE architecture, RI-MAE, for rotation-invariant self-supervised learning. RI-MAE is with a dual-branch student-teacher architecture, as illustrated in Fig. 2. It enables the reconstruction of masked patches within a learned rotation-invariant latent space instead of the rotation-sensitive input space, avoiding rotation interference. Overall, the two branches involves two RI-Transformers correspondingly and are connected with an additional RI-Transformer predictor. The teacher branch processes and encodes all point patches $\mathbf{P} = \{\mathbf{p}_i\}_{i=1}^G$. In the student branch, only the visible set $\mathbf{P}_v = \{\mathbf{p}_i^v\}_{i=1}^V$ is fed for encoding. Subsequently, the RI-Transformer predictor, with a single Transformer layer, is employed to predict the features of the masked point patches $\bar{\mathbf{Z}}^t = \{\bar{\mathbf{z}}_i^t\}_{i=1}^M$. It uses the output embeddings of the visible point patches from the student branch, along with essential mask tokens, for the prediction. The supervision is based on the corresponding outputs from the teacher branch. The prediction is denoted as $\bar{\mathbf{Z}}^s = \{\bar{\mathbf{z}}_i^s\}_{i=1}^M$. Given that the reconstruction process is carried out within the learned rotation-invariant latent space, our dual-branch model guarantees the rotation invariance of the reconstruction. This attribute ensures the model’s ability to pre-train effectively, even in scenarios where the orientation of the point cloud is unknown.

Note that the teacher and student branches are initialized randomly. The student updates its parameters through back-propagation, while the teacher operates as momentum-based, with its parameters being updated as the exponential moving average (EMA) of the student’s parameters. Following the training of the dual-branch architecture, only the student encoder is utilized for downstream tasks.

Loss Function. Given the target masked point patch representations $\mathbf{Z}^t = \{\bar{\mathbf{z}}_i^t\}_{i=1}^M$ from the teacher branch and the predicted representations $\bar{\mathbf{Z}}^s = \{\bar{\mathbf{z}}_i^s\}_{i=1}^M$ outputted from the predictor, a loss function is defined based on the mean squared error (MSE):

$$\mathcal{L} = \text{MSE}(\bar{\mathbf{Z}}^t, \bar{\mathbf{Z}}^s) = \frac{1}{M} \sum_{i=1}^M \|\bar{\mathbf{z}}_i^t - \bar{\mathbf{z}}_i^s\|_2^2. \quad (13)$$

4 Experiments

4.1 Implementation Details

Following existing MPM methods, we utilized FPS and KNN to divide an input point cloud into G point patches with $K = 32$ points for each patch. For point cloud classification, we set $G = 64$, while for the segmentation tasks, $G = 256$. The RI-Transformer encoders in RI-MAE contain 12 transformer layers, while the predictor has only one transformer block. For each transformer block, we set the internal dimension to 384 and the number of heads to 6.

For training details, we utilized an AdamW optimizer with a cosine learning rate decay, applying a decay factor 0.05, and incorporated a 10-epoch warm-up phase. The learning rate was set at 0.0005 for both pretraining and classification, while for segmentation, it was adjusted to 0.0002.

Methods	z/z	$SO3/SO3$	$z/SO3$
RaRI-Conv (Chen and Cong 2022)	83.3	83.3	83.3
CRIN \ddagger (Lou et al. 2023)	84.7	84.7	84.7
Yu et al. (Yu, Zhang, and Cai 2023)	86.6	86.3	<u>86.6</u>
LocoTrans (Chen et al. 2024b)	85.0	84.5	85.0
Transformer (Yu et al. 2022)	79.9	58.7	20.0
Point-BERT (Yu et al. 2022)	87.4	81.6	21.5
Point-MAE \dagger (Pang et al. 2022)	88.3	85.4	23.8
Point-M2AE \dagger (Zhang et al. 2022)	91.2	<u>86.7</u>	26.7
ACT \dagger (Dong et al. 2023)	93.3	83.6	22.7
PointGPT-S \dagger (Chen et al. 2024a)	91.7	85.4	18.8
RI-Transformer	88.3	88.3	88.3
RI-MAE	91.6	91.6	91.6
RI-MAE \dagger	<u>91.9</u>	91.9	91.9

Table 1: Real-world classification accuracy (%) on ScanObjectNN. \dagger indicates the method uses 2048 points as input. \ddagger denotes additional normal vectors were fed.

Methods	10-way	
	10-shot	20-shot
CrossPoint (Afham et al. 2022)	27.0 \pm 2.9	26.8 \pm 2.7
Transformer (Yu et al. 2022)	29.0 \pm 5.6	31.8 \pm 3.3
Point-BERT (Yu et al. 2022)	<u>36.5 \pm 3.2</u>	36.4 \pm 2.2
Point-MAE (Pang et al. 2022)	36.4 \pm 2.4	<u>37.0 \pm 1.9</u>
Point-M2AE (Zhang et al. 2022)	35.9 \pm 4.2	<u>36.7 \pm 3.1</u>
ACT (Dong et al. 2023)	33.8 \pm 6.7	36.2 \pm 4.7
PointGPT-S (Chen et al. 2024a)	32.6 \pm 3.7	32.0 \pm 2.7
RI-Transformer	78.9 \pm 6.4	86.8 \pm 5.2
RI-MAE	90.2 \pm 5.5	93.7 \pm 3.5

Table 2: Few-shot object classification results (%) on ModelNet40 in the $z/SO3$ scenario.

We pretrained RI-MAE on ShapeNet (Chang et al. 2015), which consists of more than 50,000 3D models, covering 55 categories. Following ClusterNet (Chen et al. 2019), we conducted experiments in three scenarios to evaluate the rotation invariance: 1) z/z for training and testing with azimuthal rotations; 2) $z/SO3$ for training with azimuthal rotations and testing with arbitrary rotations; and 3) $SO3/SO3$, entailing both training and testing under arbitrary rotations. All experiments were conducted utilizing two RTX 2080Ti GPUs, with the PyTorch framework version 1.7.

4.2 Downstream Tasks

Real-World Classification on ScanObjectNN. We evaluate the proposed method on a classification task with a real-world dataset - ScanObjectNN (Uy et al. 2019), which contains 2,902 point clouds of 15 categories collected in real world. There are three variants of ScanObjectNN: *OBJ_BG*, *OBJ_ONLY*, and *PB_T50_RS*. We evaluate our approach on the *OBJ_BG* variant by incorporating additional MLP classification heads, with 1,024 points sampled, and compare our

Methods	mIoU $_I$	mIoU $_C$
RIFrame (Li et al. 2021b)	82.5	79.4
LGR-Net (Zhao et al. 2022)	82.8	80.1
Luo et al. (Luo et al. 2022)	-	81.0
RaRI-Conv (Chen and Cong 2022)	83.8	-
Xie et al. (Xie et al. 2023)	83.9	<u>81.5</u>
CRINet (Lou et al. 2023)	-	80.5
Yu et al. (Yu, Zhang, and Cai 2023)	-	80.3
LocoTrans (Chen et al. 2024b)	<u>84.0</u>	80.1
Transformer (Yu et al. 2022)	29.1	33.4
Point-BERT (Yu et al. 2022)	29.2	33.8
Point-MAE (Pang et al. 2022)	31.5	33.8
Point-M2AE (Zhang et al. 2022)	35.4	38.7
ACT (Dong et al. 2023)	33.3	35.7
PointGPT-S (Chen et al. 2024a)	32.3	34.3
RI-Transformer	84.0	81.6
RI-MAE	84.3	82.1

Table 3: Part segmentation on ShapeNetPart with mean IoU for all instances mIoU $_I$ (%) and mean IoU for all categories mIoU $_C$ (%), in the $z/SO3$ scenarios.

method with state-of-the-art rotation-invariant and MPM approaches. It is worth emphasizing that, for the MPM methods, experiments in rotation scenarios were conducted utilizing the open-source code. The results are detailed in Table 1. Specifically, we trained an RI-Transformer without utilizing the proposed dual-branch self-supervised learning strategy as a baseline. The baseline RI-Transformer exhibits superior performance compared to the state-of-the-art rotation-invariant methods, notably outperforming the second-best approach (Yu, Zhang, and Cai 2023) by a margin of +1.7%. Regrading RI-MAE, with the dual-branch learning strategy, it delivers improvements of +2.6% over the RI-Transformer baseline. These results indicate the superiority of the proposed mechanisms. The results of the remaining two variants are included in the technical appendix.

Few-shot Learning on ModelNet40. ModelNet40 is a synthetic shape dataset that contains 12,311 CAD models across 40 categories, with 1,024 points sampled from each model (Wu et al. 2014). We conducted few-shot learning on ModelNet40, following practices from existing MPM methods. The performance is assessed across two settings: *10-way 10-shot* and *10-way 20-shot*, where *K-way N-shot* indicates that *N* objects are sampled for training from each of *K* randomly selected classes, followed by testing on 20 unseen instances randomly chosen from each of these *K* classes. All experiments were conducted in the $z/SO3$ scenario to evaluate rotation robustness. We report both the average accuracy and the standard deviation from 10 independent runs in Table 2. The results show that our RI-MAE significantly outperforms existing methods, showing more than 50% improvements. We also compare performance in other train/test scenarios and report the results of *5-way 10-shot* and *5-way 20-shot* settings in the technical appendix.

Methods	z/z	$SO3/SO3$	$z/SO3$
GCConv (Zhang et al. 2020)	-	55.8	<u>55.7</u>
LGR-Net (Zhao et al. 2022)	43.4	43.4	43.4
Xie et al. (Xie et al. 2023)	44.6	44.6	44.6
Yu et al. (Yu, Zhang, and Cai 2023)	51.2	51.2	51.2
LocoTrans (Chen et al. 2024b)	-	56.0	54.2
Transformer (Yu et al. 2022)	60.0	<u>57.1</u>	4.9
Point-MAE (Pang et al. 2022)	<u>60.8</u>	<u>56.6</u>	5.9
ACT (Dong et al. 2023)	61.2	56.1	5.4
RI-Transformer	59.9	59.9	59.9
RI-MAE	60.3	60.3	60.3

Table 4: Semantic segmentation accuracy with mIoU (%) on S3DIS in three train/test scenarios on Area 5.




Figure 3: Visualization of part segmentation results on ShapeNet in the $z/SO3$ scenario. The leftmost column is the ground truth and the rest columns are the testing results of RI-MAE under arbitrary rotations.

Part Segmentation on ShapeNet. The evaluation on part segmentation is with ShapeNet, which consists of 16,881 objects across 16 categories. In line with existing MPM methods, we sample 2,048 points from each object as input. A simple segmentation head is incorporated as in Point-MAE (Pang et al. 2022). The segmentation results in the $z/SO3$ scenario are reported in Table 3. It is evident that existing MPM methods are sensitive to rotations, resulting in significantly unsatisfactory accuracy levels. In contrast, our RI-MAE demonstrates robustness to rotations and achieves the best performance, with a mean intersection over union (mIoU) score of 84.3% $mIoU_I$ and 82.1% $mIoU_C$. In addition, the part segmentation results of RI-MAE in the $z/SO3$ scenario are visualized in Fig. 3, indicating consistent predictions regardless of rotation. We also provide the results in the z/z scenario in the technical appendix.

Semantic Segmentation on S3DIS. We expanded our study to include challenging semantic segmentation on the large-scale 3D scenes dataset, S3DIS (Armeni et al. 2016), which contains point clouds scanned from 271 rooms across

Row	RI-OE	RI-PE	Dual-Branch	AE	Acc.
#1			✓		82.6
#2	✓		✓		88.6
#3		✓	✓		88.5
#4	✓	✓		✓	86.7
#5	✓	✓	✓		91.6

Table 5: Ablation study on RI-OE, RI-PE and dual-branch.

six indoor areas, with all points annotated into 13 semantic categories. We sampled 4,096 points per room, utilizing only xyz coordinates as input, without incorporating rgb information. Areas 1 - 4 and 6 are for training and Area 5 is for testing. As detailed in Table 4, it is evident that our method outperforms other methods when tested with rotated data. It is noteworthy that our RI-Transformer significantly surpasses Yu et al. (Yu, Zhang, and Cai 2023), a heavily modified transformer, proving to be an effective backbone.

4.3 Ablation Study

We conducted ablation studies to evaluate the effectiveness of the proposed mechanisms: 1) rotation-invariant orientation embedding, 2) rotation-invariant position embedding and 3) the dual-branch architecture. We report the results on the OBJ_BG subset of the ScanObjectNN dataset, specifically under the $z/SO3$ scenario.

Effectiveness of RI-OE and RI-PE. We first analyze the contribution of the proposed RI-OE and RI-PE. As shown in Table 5, comparing the row #1, #2, #3 and #5, both our RI-OE and RI-PE deliver significant improvements and the best performance is achieved when the two mechanisms are employed. By incorporating the relative orientation and location information offered by RI-OE and RI-PE, our RI-MAE gains a deeper comprehension of the intricate structure within the input point cloud, facilitating the acquisition of a higher quality representation of the point cloud.

Effectiveness of the Dual-Branch. We also conducted an ablation study on the dual-branch architecture, where we replaced the dual-branch structure with an autoencoder (AE) architecture for pre-training. As evident in row #4 and #5 in Table 5, our dual-branch architecture significantly surpasses the AE structure. As deliberated, when endeavoring AE to establish a rotation-invariant MPM framework, the uncertain reconstruction targets can lead to unstable loss and suboptimal solution. In contrast, our dual-branch architecture performs reconstructions within the learned rotation-invariant latent space, thus eliminating the influence of rotations and guaranteeing the efficacy of the pre-training.

5 Conclusion

In this paper, we present RI-MAE, a rotation-invariant masked point modeling approach, for self-supervised point cloud representation learning. RI-Transformer backbone is proposed to derive rotation-invariant latent space and a dual-branch design enables rotation-invariant masked point patches reconstruction. Extensive experiments demonstrate the state-of-the-art performance of our method.

References

- Afham, M.; Dissanayake, I.; Dissanayake, D.; Dharmasiri, A.; Thilakarathna, K.; and Rodrigo, R. 2022. CrossPoint: Self-supervised Cross-modal Contrastive Learning for 3D Point Cloud Understanding. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9902–9912.
- Armeni, I.; Sener, O.; Zamir, A.; Jiang, H.; Brilakis, I. K.; Fischer, M.; and Savarese, S. 2016. 3D Semantic Parsing of Large-Scale Indoor Spaces. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1534–1543.
- Chang, A. X.; Funkhouser, T. A.; Guibas, L. J.; Hanrahan, P.; Huang, Q.-X.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; Xiao, J.; Yi, L.; and Yu, F. 2015. ShapeNet: An Information-Rich 3D Model Repository. *ArXiv*, abs/1512.03012.
- Chen, C.; Li, G.; Xu, R.; Chen, T.; Wang, M.; and Lin, L. 2019. ClusterNet: Deep Hierarchical Cluster Network with Rigorously Rotation-invariant Representation for Point Cloud Analysis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4994–5002.
- Chen, G.; Wang, M.; Yang, Y.; Yu, K.; Yuan, L.; and Yue, Y. 2024a. PointGPT: Auto-regressively Generative Pre-training from Point Clouds. *Advances in Neural Information Processing Systems*, 36.
- Chen, R.; and Cong, Y. 2022. The Devil is in the Pose: Ambiguity-free 3D Rotation-invariant Learning via Pose-aware Convolution. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7472–7481.
- Chen, Y.; Duan, L.; Zhao, S.; Ding, C.; and Tao, D. 2024b. Local-consistent Transformation Learning for Rotation-invariant Point Cloud Analysis. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5418–5427.
- Chen, Y.; Liu, J.; Ni, B.; Wang, H.; Yang, J.; Liu, N.; Li, T.; and Tian, Q. 2021. Shape Self-correction for Unsupervised Point Cloud Understanding. In *IEEE/CVF International Conference on Computer Vision*, 8362–8371.
- Deng, J.; Wang, C.; Lu, J.; He, J.; Zhang, T.; Yu, J.; and Zhang, Z. 2023. SE-ORNet: Self-Ensembling Orientation-Aware Network for Unsupervised Point Cloud Shape Correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5364–5373.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *ArXiv*, abs/1810.04805.
- Dong, R.; Qi, Z.; Zhang, L.; Zhang, J.; Sun, J.; Ge, Z.; Yi, L.; and Ma, K. 2023. Autoencoders as Cross-Modal Teachers: Can Pretrained 2D Image Transformers Help 3D Representation Learning? In *International Conference on Learning Representations*.
- Esteves, C.; Allen-Blanchette, C.; Makadia, A.; and Daniilidis, K. 2018. Learning SO(3) Equivariant Representations with Spherical CNNs. In *European Conference on Computer Vision*, 52–68.
- Gu, R.; Wu, Q.; Xu, H.; Ng, W. W.; and Wang, Z. 2021. Learning Efficient Rotation Representation for Point Cloud via Local-global Aggregation. In *IEEE International Conference on Multimedia and Expo*, 1–6. IEEE.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Doll'ar, P.; and Girshick, R. B. 2021. Masked Autoencoders Are Scalable Vision Learners. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15979–15988.
- Kim, S.; Park, J.; and Han, B. 2020. Rotation-invariant Local-to-global Representation Learning for 3D Point Cloud. *Advances in Neural Information Processing Systems*, 33: 8174–8185.
- Li, F.; Fujiwara, K.; Okura, F.; and Matsushita, Y. 2021a. A Closer Look at Rotation-invariant Deep Point Cloud Analysis. In *IEEE/CVF International Conference on Computer Vision*, 16198–16207.
- Li, X.; Li, R.; Chen, G.; Fu, C.-W.; Cohen-Or, D.; and Heng, P.-A. 2021b. A Rotation-invariant Framework for Deep Point Cloud Analysis. *IEEE Transactions on Visualization and Computer Graphics*, 28(12): 4503–4514.
- Li, Y.; Bu, R.; Sun, M.; Wu, W.; Di, X.; and Chen, B. 2018. PointCNN: Convolution on X-transformed Points. *Advances in Neural Information Processing Systems*, 31.
- Liu, X.; Han, Z.; Wen, X.; Liu, Y.-S.; and Zwicker, M. 2019. L2G Auto-encoder: Understanding Point Clouds by Local-to-Global Reconstruction with Hierarchical Self-Attention. In *ACM International Conference on Multimedia*.
- Lou, Y.; Ye, Z.; You, Y.; Jiang, N.; Lu, J.; Wang, W.; Ma, L.; and Lu, C. 2023. CRIN: Rotation-Invariant Point Cloud Analysis and Rotation Estimation via Centrifugal Reference Frame. *ArXiv*, abs/2303.03101.
- Luo, S.; Li, J.; Guan, J.; Su, Y.; Cheng, C.; Peng, J.; and Ma, J. 2022. Equivariant Point Cloud Analysis via Learning Orientations for Message Passing. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18910–18919.
- Pang, Y.; Wang, W.; Tay, F. E.; Liu, W.; Tian, Y.; and Yuan, L. 2022. Masked Autoencoders for Point Cloud Self-supervised Learning. In *European Conference on Computer Vision*, 604–621. Springer.
- Poulenard, A.; and Guibas, L. J. 2021. A Functional Approach to Rotation Equivariant Non-linearities for Tensor Field Networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13174–13183.
- Poursaeed, O.; Jiang, T.; Qiao, Q.; Xu, N.; and Kim, V. G. 2020. Self-Supervised Learning of Point Clouds via Orientation Estimation. In *International Conference on 3D Vision*, 1018–1028.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. *Advances in Neural Information Processing Systems*, 30.
- Qi, Z.; Dong, R.; Fan, G.; Ge, Z.; Zhang, X.; Ma, K.; and Yi, L. 2023. Contrast with Reconstruct: Contrastive 3D Representation Learning Guided by Generative Pretraining. In *International Conference on Machine Learning*.

- Qi, Z.; Dong, R.; Zhang, S.; Geng, H.; Han, C.; Ge, Z.; Yi, L.; and Ma, K. 2024. ShapeLLM: Universal 3D Object Understanding for Embodied Interaction. *ArXiv*, abs/2402.17766.
- Qiu, Z.; Li, Y.; Wang, Y.; Pan, Y.; Yao, T.; and Mei, T. 2022. SPE-Net: Boosting Point Cloud Analysis via Rotation Robustness Enhancement. In *European Conference on Computer Vision*, 593–609. Springer.
- Rao, Y.; Lu, J.; and Zhou, J. 2019. Spherical Fractal Convolutional Neural Networks for Point Cloud Recognition. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 452–460.
- Rao, Y.; Lu, J.; and Zhou, J. 2020. Global-local Bidirectional Reasoning for Unsupervised Representation Learning of 3D Point Clouds. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5376–5385.
- Rolfe, J. T. 2016. Discrete Variational Autoencoders. *ArXiv*, abs/1609.02200.
- Sanghi, A. 2020. Info3D: Representation Learning on 3D Objects Using Mutual Information Maximization and Contrastive Learning. In *European Conference on Computer Vision*, 626–642. Springer.
- Sun, C.; Zheng, Z.; Wang, X.; Xu, M.; and Yang, Y. 2022. Self-supervised Point Cloud Representation Learning via Separating Mixed Shapes. *IEEE Transactions on Multimedia*.
- Tsai, M.-S.; Chiang, P.-Z.; Tsai, Y.-H.; and Chiu, W.-C. 2022. Self-Supervised Feature Learning from Partial Point Clouds via Pose Disentanglement. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 1031–1038.
- Uy, M. A.; Pham, Q.-H.; Hua, B.-S.; Nguyen, D. T.; and Yeung, S.-K. 2019. Revisiting Point Cloud Classification: A New Benchmark Dataset and Classification Model on Real-World Data. In *IEEE/CVF International Conference on Computer Vision*, 1588–1597.
- Wang, H.; Liu, Q.; Yue, X.; Lasenby, J.; and Kusner, M. J. 2020. Unsupervised Point Cloud Pre-training via Occlusion Completion. In *IEEE/CVF International Conference on Computer Vision*, 9762–9772.
- Wang, R.; Yang, Y.; and Tao, D. 2022. ART-Point: Improving Rotation Robustness of Point Cloud Classifiers via Adversarial Rotation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14371–14380.
- Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2019. Dynamic Graph CNN for Learning on Point Clouds. *ACM Transactions on Graphics*, 38(5): 1–12.
- Wu, K.; Peng, H.; Chen, M.; Fu, J.; and Chao, H. 2021. Rethinking and Improving Relative Position Encoding for Vision Transformer. In *IEEE/CVF International Conference on Computer Vision*, 10013–10021.
- Wu, Z.; Song, S.; Khosla, A.; Yu, F.; Zhang, L.; Tang, X.; and Xiao, J. 2014. 3D ShapeNets: A Deep Representation for Volumetric Shapes. In *IEEE Conference on Computer Vision and Pattern Recognition*, 1912–1920.
- Xiao, Z.; Lin, H.; Li, R.; Geng, L.; Chao, H.; and Ding, S. 2020. Endowing Deep 3D Models with Rotation Invariance Based on Principal Component Analysis. In *IEEE International Conference on Multimedia and Expo*, 1–6. IEEE.
- Xie, L.; Yang, Y.; Wang, W.; Lin, B.; Cai, D.; and He, X. 2023. General Rotation Invariance Learning for Point Clouds via Weight-feature Alignment. *ArXiv*, abs/2302.09907.
- Xie, S.; Gu, J.; Guo, D.; Qi, C. R.; Guibas, L.; and Litany, O. 2020. PointContrast: Unsupervised Pre-training for 3D Point Cloud Understanding. In *European Conference on Computer Vision*, 574–591. Springer.
- Yang, Y.; Feng, C.; Shen, Y.; and Tian, D. 2018. FoldingNet: Point Cloud Auto-encoder via Deep Grid Deformation. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 206–215.
- You, Y.; Lou, Y.; Shi, R.; Liu, Q.; Tai, Y.-W.; Ma, L.; Wang, W.; and Lu, C. 2021. PRIN/SPRIN: On Extracting Point-Wise Rotation Invariant Features. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44: 9489–9502.
- Yu, J.; Zhang, C.; and Cai, W. 2023. Rethinking Rotation Invariance with Point Cloud Registration. In *AAAI Conference on Artificial Intelligence*, volume 37, 3313–3321.
- Yu, R.; Wei, X.; Tombari, F.; and Sun, J. 2020. Deep Positional and Relational Feature Learning for Rotation-invariant Point Cloud Analysis. In *European Conference on Computer Vision*, 217–233. Springer.
- Yu, X.; Tang, L.; Rao, Y.; Huang, T.; Zhou, J.; and Lu, J. 2022. Point-BERT: Pre-training 3D Point Cloud Transformers with Masked Point Modeling. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 19313–19322.
- Zhang, R.; Guo, Z.; Gao, P.; Fang, R.; Zhao, B.; Wang, D.; Qiao, Y.; and Li, H. 2022. Point-M2AE: Multi-scale Masked Autoencoders for Hierarchical Point Cloud Pre-training. *Advances in Neural Information Processing Systems*, 35: 27061–27074.
- Zhang, R.; Wang, L.; Qiao, Y.; Gao, P.; and Li, H. 2023. Learning 3D Representations from 2D Pre-trained Models via Image-to-point Masked Autoencoders. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21769–21780.
- Zhang, Z.; Hua, B.-S.; Chen, W.; Tian, Y.; and Yeung, S.-K. 2020. Global Context Aware Convolutions for 3D Point Cloud Understanding. In *International Conference on 3D Vision*, 210–219.
- Zhang, Z.; Hua, B.-S.; and Yeung, S.-K. 2022. RICov++: Effective Rotation Invariant Convolutions for 3D Point Clouds Deep Learning. *International Journal of Computer Vision*, 130(5): 1228–1243.
- Zhao, C.; Yang, J.; Xiong, X.; Zhu, A.; Cao, Z.; and Li, X. 2022. Rotation Invariant Point Cloud Analysis: Where Local Geometry Meets Global Topology. *Pattern Recognition*, 127: 108626.