

Self-Supervised Feature Learning from Partial Point Clouds via Pose Disentanglement

Meng-Shiun Tsai^{*1} Pei-Ze Chiang^{*1} Yi-Hsuan Tsai² Wei-Chen Chiu¹
¹National Yang Ming Chiao Tung University, Taiwan ²NEC Labs America

Abstract

Self-supervised learning on point clouds has gained a lot of attention recently, since it addresses the label-efficiency and domain-gap problems on point cloud tasks. In this paper, we propose a novel self-supervised framework to learn informative representations from partial point clouds. We leverage partial point clouds scanned by LiDAR that contain both content and pose attributes, and we show that disentangling such two factors from partial point clouds enhances feature representation learning. To this end, our framework consists of three main parts: 1) a completion network to capture holistic semantics of point clouds; 2) a pose regression network to understand the viewing angle where partial data is scanned from; 3) a partial reconstruction network to encourage the model to learn content and pose features. To demonstrate the robustness of the learnt feature representations, we conduct several downstream tasks including classification, part segmentation, and registration, with comparisons against state-of-the-art methods. Our method not only outperforms existing self-supervised methods, but also shows a better generalizability across synthetic and real-world datasets.

1. Introduction

Point clouds provide one of the most intuitive representations for 3D object models, and they are extensively adopted as the data format for the recognition tasks in different application scenarios, such as autonomous vehicles, robotics, and architectonics. Recently, with the rapid development of deep learning techniques, various network architectures are proposed for learning effective features of the point cloud data, e.g., PointNet [32], PointNet++ [33], and DGCNN [41]. However, the success of learning the recognition models based on these feature backbones typically relies on the large-scale supervised dataset, in which it could be quite expensive to manually collect the ground truth labels and the learnt feature representations of point cloud data are less generalizable across different tasks.

* Both authors contributed equally to the paper

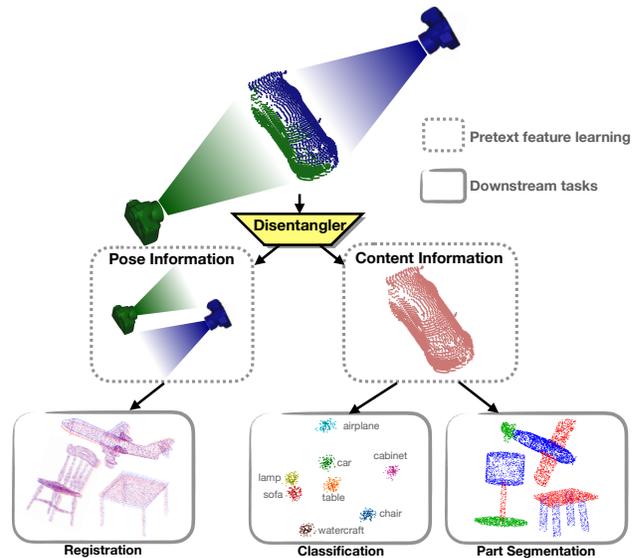


Figure 1. We propose to learn the feature representation of 3D point cloud data in a self-supervised manner via point cloud completion. Our model learns to disentangle the feature into the content and pose parts, where the former is beneficial for the downstream tasks of classification and part segmentation, while the latter helps to improve the task of point cloud registration.

These potential issues motivate various research development in learning effective feature representations from point clouds via unsupervised or self-supervised manner, such as solving jigsaw puzzles [2, 35], contrastive learning [34, 43], and local structure modeling [18, 19]. In this paper, we aim to seek for a different solution that is also intuitive and suitable for the point cloud data. Inspired by the success of feature learning in the image [30] and the natural language [38] domains, which have been developed for years, we adopt an analogous strategy of image inpainting [30], but in a 3D point cloud version via *learning to complete* the partial point cloud of a 3D object. As such, similar benefits for feature learning as shown in image inpainting [30] can be achieved by point cloud completion, which not only needs the semantic understanding of the point cloud data, but also learns better holistic feature representations via reconstructing the

plausible missing parts.

However, simply completing the partial point cloud, as in the case of image domain, may not suffice the need for learning effective feature representations, as the point cloud data has different characteristics from images. That is, given a 3D object point cloud, there could be multiple angles to view this data, such that each angle produces a different partial point cloud (in the top of Figure 1). Therefore, although the final completed point cloud is the same for all the partial point clouds obtained from the same object, the completion network may require to learn a different feature representation for each view-angle during the completion process. To take view-angle into consideration, instead of using existing point cloud completion frameworks [26, 48], which only produce one feature representation that *entangles* both the view-angle information and the content cue for the point cloud data, we propose to *disentangle* these two factors and learn more effective feature representations in a self-supervised manner (see Figure 1).

Specifically, we utilize two encoders to extract the content and pose (i.e., view-angle) features individually from each partial point cloud input. Then, in addition to performing completion using the content feature, the pose feature should be able to predict the pose of the input data and guide the reconstruction in specific view-angle. Thus, we introduce another two modules: 1) a pose regression network to predict the view-angle of the partial point cloud using the pose feature; and 2) a partial reconstruction network to recover a specific partial point cloud, through the combination of a content feature from a view-angle i and a pose feature from another view-angle j . As a result, no matter which view-angle the content feature is extracted from, the partial reconstruction should be mainly guided by the pose feature. Therefore, our framework encourages the content and pose features to be more compact on themselves while being more disentangled to each other, which better facilitates the feature representation learning process.

We conduct extensive experiments to show the effectiveness of our self-supervised framework: 1) our learned content feature is able to provide favorable performance against state-of-the-art methods on the downstream tasks of classification and part segmentation, and 2) our learned pose features contribute to the pose-relevant downstream task of point cloud registration. The main contributions of our work are summarized as follows:

- We propose a self-supervised framework based on point cloud completion for learning the feature representations from the partial point cloud data.
- We develop a pipeline to disentangle point cloud feature representations into the content and pose factors, which enables the model to learn effective features.
- We show that the learnt content and pose features improve several downstream tasks (i.e., classification, part

segmentation, and registration), while showing the generalizability across synthetic and real-world datasets.

2. Related Work

Supervised Learning for Point Cloud Data. Due to the popularity of 3D scanning technologies and the advancement of deep learning techniques, learning to extract features from the point cloud data for recognition tasks is one of the active research topics. One early attempt via deep networks is PointNet [32] which uses the max-pooling operation to aggregate the point-wise features into the global one, thus achieving the permutation invariant property. While PointNet neglects the local structure between neighboring points, its successor, PointNet++ [33], adopts the hierarchical neural network to progressively extract the features of a point cloud from multiple resolutions. Furthermore, DGCNN [41] leverages the graph neural network [23] to process the nearby points and their edges via the k-nearest neighbor algorithm in each of the intermediate feature spaces. Other methods that also try to capture the relationships among the local regions in point clouds are developed in [24, 27, 44].

However, all the aforementioned works are based on the complete point cloud with supervised information, which is costly to collect in terms of time and expense. Even when we are able to use the synthetic dataset where the supervision is easier to obtain, the feature learnt from the synthetic data could be less generalizable to the real-world data due to the domain gap [3]. In this work, we thereby focus on the self-supervised learning scheme to learn effective representations that can generalize better across datasets.

Unsupervised / Self-Supervised Feature Learning for Point Cloud Data. Recently, various works have been proposed to explore the unsupervised and self-supervised schemes for learning feature representations of 3D point clouds, where the training objectives or the labels are produced from the data itself. An intuitive objective is derived via performing the self-reconstruction [1, 12], in which the semantic and compact feature representations are learnt in the bottleneck between the encoding and decoding processes. For instance, the encoder of 3DCapsuleNet [49] is designed to process 3D point clouds and aggregate the final latent capsules based on the dynamic routing, followed by the decoder to reconstruct the original point cloud composed of multiple point patches. [45] proposes an autoencoder architecture where the rich point-wise features in multiple stages are progressively produced through the seed generation module. Eckart *et al.* [13] further extend the traditional autoencoder paradigm to form the bottleneck layer being modeled by a discrete generative model. In addition to the self-reconstruction, MAP-VAE [18] proposes the multi-angle analysis to introduce the local self-reconstruction, leading a better modeling on the local geometry of point clouds.

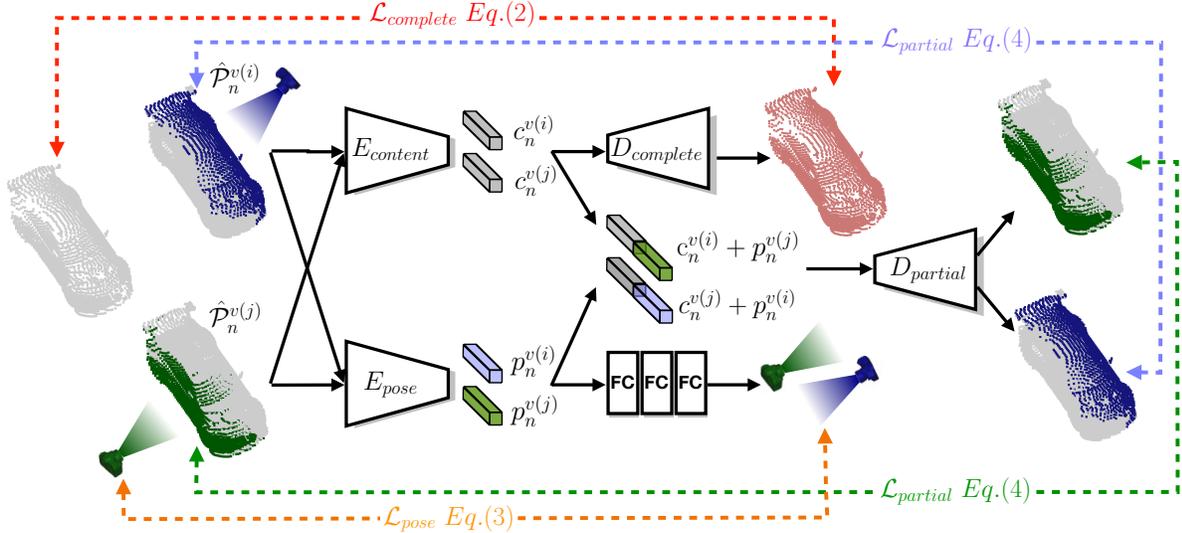


Figure 2. Our proposed framework aims to learn two distinct features (c_n, p_n) via disentanglement that consists of three main branches: 1) a completion branch containing a content encoder $E_{content}$ and a decoder $D_{complete}$ that learns the content features c_n , through completing partial point clouds of the same 3D object. 2) a pose regression branch with a pose encoder E_{pose} that learns the pose features p_n , through predicting view-angles where a partial point cloud is scanned from. 3) a partial reconstruction branch $D_{partial}$, which leverages combined content features with exchanged pose features to reconstruct the partial inputs. Note that “+” denotes the concatenation operation and dashed lines in colors are the loss functions we employ for three branches.

Another effective technique is stemmed from contrastive learning. For instance, PointContrast [43] proposes to use a pretext task, where the mapping between two point clouds taken by viewing from different perspectives should be consistent even after performing the random geometric transformations on them. [20] relies on the similar consistency but extends to include the spatial augmentations on the data, and adopts the self-ensembling mechanism to drive the learning. Other methods are inspired by the techniques originally applied on 2D images. For instance, Sauder *et al.* [35] adopts the jigsaw puzzle idea on 3D point cloud, in which the point cloud is decomposed into multiple parts followed by random permutation, and then the network is trained by putting these parts back to their original positions. Similarly, Allegro *et al.* [2] further integrate multi-task learning into the same framework as [35]. In addition, [9] destroys shape parts and learns the feature representation via the process of distinguishing the destroyed parts and restoring them.

Different from the above methods, our approach is inspired by the idea of 2D image inpainting [30] but in a 3D version, in which we propose to self-learn the completion process from partial point clouds as the pretext task. Furthermore, beyond point cloud completion, we extend our framework to consider the pose information, so that features can be disentangled during the completion process to learn effective feature representations. We note that, a recent work [39] also adopts the idea of point cloud completion to drive feature learning but does not fully leverage the pose information, in which this work can be treated as an ablated variant of our model. Later in experiments we provide the

ablation study to verify the benefit of our model design with respect to this work.

Feature Disentanglement. Disentanglement aims to decompose the feature representation into multiple parts for better explaining the factors behind data variation. Numerous methods are developed in 2D images, such as the conditional GAN [28], auxiliary classifier GAN [29], InfoGAN [8] and conditional VAE [36], all of which are used to disentangle the latent space of GAN [17] and VAE [22]. Moreover, introducing disentanglement into representation learning benefits various computer vision applications. For instance, [25, 51] disentangle the pose information from person images to obtain the purified feature of person, which improves the performance in the task of person re-identification. Zhou *et al.* [50] factors out the shape and pose features from the 3D mesh data to boost the performance in the tasks of pose transfer and shape retrieval. In our proposed framework, we disentangle the content and pose information of the 3D point cloud representation, in which the content feature is beneficial for the downstream tasks of classification and part segmentation, while the pose feature boosts the task of point cloud registration.

3. Proposed Method

As motivated in the introduction, the objective of our proposed framework is to perform self-supervised learning via partial point cloud completion and learn the 3D point cloud feature representations, which is decomposed into the content and pose parts. The overall architecture of our

proposed framework is illustrated in Figure 2, where there are three main branches: completion branch, pose regression branch, and partial reconstruction branch. We now detail our proposed framework in the following subsections.

3.1. Completion Branch

The purpose of this branch is to perform 3D partial point cloud completion, which is proposed in this paper as an analogous strategy to 2D image inpainting [30] for self-learning feature representations. To this end, given a dataset composed of N complete point clouds of 3D objects $\mathbb{P} = \{\mathcal{P}_n\}_{n=1}^N$, for each of the complete point cloud \mathcal{P}_n , we are able to generate many of its corresponding partial point clouds $\{\hat{\mathcal{P}}_n^{v(k)}\}_{k=1}^K$ as being scanned from K different pre-defined viewpoints $\{v(k)\}_{k=1}^K$. Then, for an input partial point cloud $\hat{\mathcal{P}}_n^{v(k)}$, we use the content encoder $E_{content}$ to extract its content feature $c_n^{v(k)} = E_{content}(\hat{\mathcal{P}}_n^{v(k)})$, followed by the completion decoder $D_{complete}$ to reconstruct the corresponding complete point cloud \mathcal{P}_n from $c_n^{v(k)}$.

Here, we adopt the standard point cloud feature extractor (e.g., PointNet [32] or DGCNN [41]) as the content encoder $E_{content}$. For the completion decoder $D_{complete}$, we utilize a morphing-based decoder in MSN [26] to predict \mathcal{P}_n . To measure the completion quality, we use Earth Mover’s Distance (EMD) [14] loss (which is typically approximated via the auction algorithm [6] in practical implementation). Therefore, given two point clouds $\mathcal{P}_i, \mathcal{P}_j \in \mathbb{R}^3$ with an equal size $|\mathcal{P}_i| = |\mathcal{P}_j|$, the EMD loss function is:

$$d_{EMD}(\mathcal{P}_i, \mathcal{P}_j) = \min_{\phi(\mathcal{P}_i) \rightarrow \mathcal{P}_j} \sum_{x \in \mathcal{P}_i} \|x - \phi(x)\|_2, \quad (1)$$

where ϕ is an optimal bijection function that allows each point in \mathcal{P}_i to find a nearest corresponding point in \mathcal{P}_j .

Moreover, in addition to the EMD loss that aims at guiding the entire generated point cloud to maximally cover the ground truth one, we also adopt the expansion loss proposed in [26] to better handle local regions in the point cloud. We provide details of this expansion loss in the supplementary material. As a result, our loss function in the completion branch for a point cloud \mathcal{P}_n in a view-angle $v(k)$ can be written as:

$$\mathcal{L}_{complete} = d_{EMD}(\mathcal{P}_n, D_{complete}(c_n^{v(k)})) + \lambda_{ex} * \mathcal{L}_{expan}(D_{complete}(c_n^{v(k)})). \quad (2)$$

where we follow [26] to set $\lambda_{ex} = 0.1$.

3.2. Pose Regression Branch

Although the self-supervised objective in Section 3.1 enables our completion model to learn feature representations as an initial step, it is still insufficient to fully exploit the rich information contained in partial point cloud captured under

various view-angles. Thus, in the pose regression branch, we aim to learn the pose feature that can predict the pose of the partial data and assist in the feature learning process. We first use the standard point cloud feature extractor (e.g., PointNet [32] or DGCNN [41]) as the pose encoder E_{pose} to extract the pose feature $p_n^{v(k)} = E_{pose}(\hat{\mathcal{P}}_n^{v(k)})$, which represents the feature for the view-angle $v(k)$ of the point cloud \mathcal{P}_n . Then, three fully-connected layers are constructed as the pose regressor for predicting the view-angle.

Specifically, our pose regression branch is learnt to predict the camera position where the partial data is scanned from. We define our camera position in the spherical coordinate system (γ, θ, ϕ) , where we only rotate the data along the polar angle θ and the azimuthal angle ϕ . Note that we fix the radial distance γ to have a consistent point cloud density and use the degree as unit of θ and ϕ . Although we pre-define the poses and can consider pose prediction as a classification problem, we find that retaining it as the regression task serves better for feature learning (please refer to our supplementary material). To optimize the pose regression branch, we use the mean square error (MSE). Given a partial point cloud $\hat{\mathcal{P}}_n^{v(k)}$, our pose regression branch should predict the camera position $\hat{v}(k)$. The loss function thus can be written as:

$$\mathcal{L}_{pose} = \|v(k) - \hat{v}(k)\|_2. \quad (3)$$

3.3. Partial Reconstruction Branch

Based on Section 3.1 and 3.2, we have constructed two self-supervised objectives, individually for the point cloud completion and the pose regression tasks. However, there is still a lack of how to make connections between content and pose features, such that the learnt feature representation in each branch is more compact and meaningful. To achieve it, we propose to disentangle these two factors via the partial reconstruction branch.

Given two partial point clouds $\hat{\mathcal{P}}_n^{v(i)}$ and $\hat{\mathcal{P}}_n^{v(j)}$ derived from the same complete point cloud \mathcal{P}_n but with different viewpoints (i.e., $i \neq j$), we hypothesize that the content features from two views should be similar to each other (i.e., $c_n^{v(i)} \sim c_n^{v(j)}$), as they are from the same 3D object and are used to obtain the same completion output. As such, if we combine one of the content features with a specific pose feature, this combined feature should follow the pose information to reconstruct the partial point cloud, regardless of which content feature is selected.

To realize this objective, we use a partial decoder $D_{partial}$, which takes the concatenated content and pose features as the input, but with different viewpoints, e.g., concatenation of $c_n^{v(i)}$ and $p_n^{v(j)}$, to reconstruct a partial point cloud $\hat{\mathcal{P}}_n^{v(j)}$. The loss of this partial point cloud reconstruc-

tion is based on the EMD loss (1):

$$\mathcal{L}_{partial} = d_{EMD}(\hat{\mathcal{P}}_n^{v(i)}, D_{partial}(c_n^{v(j)}, p_n^{v(i)})) + d_{EMD}(\hat{\mathcal{P}}_n^{v(j)}, D_{partial}(c_n^{v(i)}, p_n^{v(j)})), \quad (4)$$

where we consider two terms by exchanging the content and pose features in two viewpoints.

3.4. Model Pre-training as Pretext

Overall Objective. Without the use of any annotations, we combine all the aforementioned loss functions from the three branches as our final self-supervised objectives:

$$\mathcal{L}_{all} = \lambda_c \mathcal{L}_{complete} + \lambda_{pa} \mathcal{L}_{partial} + \lambda_{po} \mathcal{L}_{pose}, \quad (5)$$

where the hyperparameters λ control the balance between loss functions, for all the experiments shown later, we set $\lambda_c = \lambda_{pa} = 0.5$ to equally balance two reconstruction-based objectives and set $\lambda_{po} = 0.01$.

Dataset, Data Generation, and Implementation Details.

We follow the standard self-supervised setting in [15, 35] and use the ShapeNetCore.v1 dataset [7] for the pretext task in model pre-training, where our learnt content and pose encoders (i.e., $E_{content}$ and E_{pose}) will be used to extract the content and pose features from the point cloud data respectively to perform various downstream tasks. In the pretext task, we follow the similar experimental settings as in MSN [26] for point cloud completion, where we choose a total of 35,827 3D models of 8 classes (i.e., table, chair, car, airplane, sofa, lamp, watercraft, and cabinet) from the ShapeNet. We use Blender software [10] to render the partial point clouds for each of the 3D models as being scanned from 26 pre-defined viewpoints (details are in the supplementary material). The 3D point coordinates in the point cloud are normalized into a unit sphere (i.e., within $[-1, 1]$). For each point cloud, we uniformly sample 1024 points. The pre-training of our proposed model on the ShapeNet runs for 50 epochs with the Adam optimizer. The learning rate is initialized to 0.001 and decreased by 90% every 20 epochs, and the batch size is set to 32.

Model Architecture. In our experiments, we use the standard point cloud feature extractors (e.g., PointNet [32] or DGCNN [41]) as our encoders, and adopt the morphing-based decoder proposed by MSN [26] for our decoders. In particular, as the decoder generates the point cloud via composing multiple local patches, we utilize 16 local patches for the completion decoder $D_{complete}$, while the partial decoder $D_{partial}$ only uses a single local patch for partial point cloud reconstruction. For both the content encoder $E_{content}$ and pose encoder E_{pose} , their architectures are identical to each other but without sharing weights. For the pose regressor, we use three fully-connected layers with BatchNorm and ReLU operations between every adjacent layer. More details of

the architecture are provided in the supplementary materials. Source code and models will be released to the public.

4. Experimental Results and Analysis

Following the setting in self-supervised point cloud methods [15, 35], we conduct extensive experiments to show the effectiveness and robustness of the point cloud features learnt by our proposed pretext task pre-trained on ShapeNet [7]. Specifically, during training the classifiers of downstream tasks, the pretrained content and pose encoders via our pretext task remain fixed, in which we use the content features for point cloud classification and part segmentation, and the pose features for the registration. In the following, we first present the analysis of our pretext task via point cloud completion then provide results of each downstream task.

4.1. Pretext Analysis

Features Distribution. We assess the quality of our learnt content and pose features by visualizing the feature embedding using t-SNE in Figure 3. We extract both content and pose features from the fixed $E_{content}$ and E_{pose} respectively, which is pre-trained on ShapeNet using the DGCNN backbone at 50 epoch. In Figure 3.a and 3.c, we plot feature distributions with respect to the class label in different colors (note that class labels are not used in training). It shows that data with the same class is grouped together for the content feature, while the pose feature is more invariant to the class label. Similar observations are also found with respect to the pose label in Figure 3.b and 3.d, where the content feature is more invariant to the pose label. This evidence verifies our proposed feature disentanglement process for learning the distinct characteristics of the content and pose features.

Relevance b/w Pretext Loss and Downstream Accuracy.

We plot the training loss with respect to downstream classification accuracy in Figure 4 to verify the effectiveness of our objective function. Both Figure 4.a and Figure 4.b show the gradual accuracy gain on the downstream tasks as the training loss decreases, which indicates that the optimization of our model is effective in feature learning.

4.2. Downstream Task 1: Classification

Dataset. ModelNet40 [42] and ScanObjectNN [3] benchmarks are adopted for the downstream task of classification, and make comparisons against several state-of-the-art supervised and unsupervised methods. ModelNet40 is composed of 12,311 CAD models from 40 categories, in which it is split into 9,843 and 2,468 models for training and testing respectively. For the ScanObjectNN dataset, which is a real-world point cloud dataset based on scanned indoor scene data, we use its two variants, ‘‘OBJ_ONLY’’ and ‘‘PB.T50_RS’’, that are widely adopted for evaluating pre-trained features in the

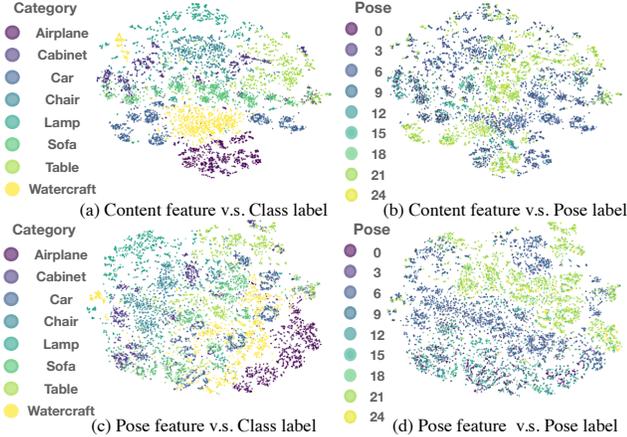


Figure 3. t-SNE visualization of content and pose features labeled with class and pose labels respectively on the ShapeNet dataset. The class label indicates the 8 categories data that we used for pre-training, while the pose label indicates the 26 viewpoints that we use for rendering partial data. Note that the pose labels with closer value denote they have similar viewpoints.

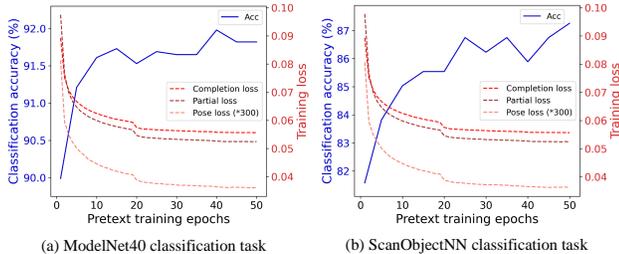


Figure 4. Pretext training loss versus downstream classification accuracy. Note that our pretext model is trained on the ShapeNet dataset, while two classification tasks are conducted on the ModelNet40 (a) and ScanObjectNN (b) benchmarks respectively.

classification downstream tasks [2, 34]. “OBJ_ONLY” contains 2,890 models with 15 categories, and “PB_T50_RS” is stemmed from “OBJ_ONLY” but being augmented with various perturbations (e.g., translation, rotation, scaling, and cropping with background, resulting in total 14,298 models), and hence is considered to be the hardest case in the ScanObjectNN. We follow the default setting to the train/test split for “OBJ_ONLY” and “PB_T50_RS” by 2,309/581 and 11,416/2,882 respectively.

Experimental Setting. We follow the standard experimental procedure as in [18, 35], where a linear Support Vector Machine (SVM) [11] is trained on the 3D point cloud features extracted from each of the methods, to evaluate the effectiveness of the feature representations in classification (note that our model uses the content feature here).

Our proposed model is pre-trained on ShapeNet, which uses the same pre-trained dataset as the unsupervised learning methods [15, 19, 21, 35, 46, 49] presented in our experiments. Please note that for both our proposed model and the unsupervised learning methods, the training set of Mod-

Table 1. **Classification accuracy (%) on the ModelNet40 dataset.** We compare our proposed model with various state-of-the-art self-supervised methods (denoted as “SSL”). In the bottom two groups, we compares methods either using the PointNet-based or DGCNN-based backbone. The first two supervised methods are served as references as they fully train the entire network. (†: reproduced by training the pretext model on the ShapeNet dataset.)

Methods	SSL	Accuracy
PointNet [32]	✗	89.2
DGCNN [41]	✗	92.9
GraphTER† [16]	✓	87.8
FoldingNet [46]	✓	88.4
PointCapsNet [49]	✓	88.9
Multi-Tasks [19]	✓	89.1
Yang <i>et al.</i> [45]	✓	90.9
PointGLR† [34]	✓	91.7
Sauder <i>et al.</i> [35]	✓	87.3
ACD [15] (PointNet++)	✓	89.8
Chen <i>et al.</i> [9]	✓	89.9
PN_ours	✓	90.1
Jing <i>et al.</i> [21]	✓	89.8
Sauder <i>et al.</i> [35]	✓	90.6
STRL [20]	✓	90.9
ParAE [13]	✓	91.6
DGCNN_ours	✓	92.0

Table 2. **Classification accuracy (%) on the ScanObjectNN dataset.** We compare our proposed model with state-of-the-art self-supervised methods (denoted as “SSL”), where supervised methods are served as references as they fully train the entire network. (†: reproduced by training the pretext model on ShapeNet.)

Methods	SSL	OBJ_ONLY	PB_T50_RS
Pointnet [32]	✗	79.2	68.2
DGCNN [41]	✗	86.2	78.1
GraphTER† [16]	✓	72.8	60.3
PointGLR† [34]	✓	85.2	73.4
PN_ours	✓	84.0	70.6
DGCNN_ours	✓	87.3	74.8

elNet40 (or ScanObjectNN) is only used to train the linear SVM classifier. For fair comparisons with GraphTER [16] and PointGLR [34], we reproduce their performance by training their pretext model on ShapeNet, where the same reproduction of [16, 34] is also applied later on other downstream tasks of part segmentation and point cloud registration.

Results. Table 1 and Table 2 show the results on the ModelNet40 and ScanObjectNN datasets respectively. Please note that, for all tables in this paper, we denote our full model of using PointNet [32] and DGCNN [41] as the backbone of encoders as “PN_ours” and “DGCNN_ours”.

In Table 1, we show that our proposed model performs favorably against other unsupervised methods, thus verifying the robustness of our self-supervised model in learning more holistic and effective feature representations. Furthermore,

Table 3. **Part segmentation results on the ShapeNet-Part dataset.** We compare our proposed model with state-of-the-art self-supervised methods (denoted as “SSL”), where supervised methods are served as references as they fully train the entire network. (ψ : pretext model is trained on ModelNet40; *: reported by MAP [18]; \dagger : reproduced by training the pretext model on the ShapeNet dataset.)

Methods	SSL	mIoU	Aero	Bag	Cap	Car	Chair	Ear ph.	Guitar	Knife	Lamp	Laptop	Motor	Mug	Pistol	Rocket	Skate	Table
PointNet [32]	\times	83.7	83.4	78.7	82.5	74.9	89.6	73.0	91.5	85.9	80.8	95.3	65.2	93.0	81.2	57.9	72.8	80.6
DGCNN [41]	\times	85.2	84.0	83.4	86.7	77.8	90.6	74.7	91.2	87.5	82.8	95.7	66.3	94.9	81.1	63.5	74.5	82.6
Latent-GAN* [1]	\checkmark	57.0	54.1	48.7	62.6	43.2	68.4	58.3	74.3	68.4	53.4	82.6	18.6	75.1	54.7	37.2	46.7	66.4
MAP ψ [18]	\checkmark	68.0	62.7	67.1	73.0	58.5	77.1	67.3	84.8	77.1	60.9	90.8	35.8	87.7	64.2	45.0	60.4	74.8
GraphTER \dagger [16]	\checkmark	82.3	81.7	76.0	83.2	74.9	84.8	64.6	90.9	87.1	82.5	95.6	56.3	93.2	81.6	56.2	71.0	67.8
PointGLR \dagger [34]	\checkmark	84.6	81.5	84.7	81.7	75.4	89.8	76.1	89.8	84.9	83.3	95.2	65.4	92.8	79.9	55.7	73.7	83.5
PN_ours	\checkmark	83.8	81.6	73.2	83.5	74.2	89.0	70.9	89.9	85.5	80.6	95.0	66.4	92.6	82.0	53.7	72.7	82.9
DGCNN_ours	\checkmark	85.1	82.3	83.5	84.5	77.3	89.8	76.3	91.0	87.3	84.2	95.5	67.8	92.5	82.8	52.1	73.9	83.5

the competitive performance (or with a small gap) of our proposed model compared to the supervised baselines shows the practical potential of our 3D point cloud feature representations. In Table 2, we show that our proposed model even outperforms all the supervised and self-supervised learning baselines in “OBJ_ONLY”, while providing comparable performance in “PB_T50_RS” with respect to the supervised methods. In particular, such results verify the generalizability of the feature representation learnt by our proposed method across synthetic and real-world datasets.

4.3. Downstream Task 2: Part Segmentation

Dataset. In addition to having the downstream task of classification for showcasing the capacity of our content features in modeling the holistic point cloud, here we experiment on another task, part segmentation, to investigate how the fine-grained information of local 3D points is maintained in our learnt content features. We adopt the ShapeNet-Part dataset [47] for performing this evaluation, which is commonly used on the part segmentation task. ShapeNet-Part dataset consists of 16,872 models with 16 categories, with being split into 13,998 and 2,874 for training and testing respectively. Depending on the object category, 3D points are annotated by 2 to 6 part labels, where there are in total 50 distinct labels for the whole dataset.

Experimental Setting. We follow the same architecture for the part classifier designed for part segmentation task in the PointNet [32] framework. Similar to the classification task, our proposed self-supervised model is first pre-trained on ShapeNet, and the model is fixed as a feature extractor while training a part segmentation classifier. Following [32], we extract point-wise features of the final convolutional layer (before max-pooling) in the content encoder $E_{content}$ as local features, while using the content features as global features for learning the segmentation part classifier. The evaluation metric is based on the average of intersection-over-union (i.e., mIoU).

Results. Table 3 summarizes the evaluation results and the comparison of our proposed model against several meth-

ods on the ShapeNet-Part dataset. Overall, our model provides better performance than the self-supervised methods, in which it demonstrates that the content features learnt via disentangling from partial point clouds with different viewpoints are able to benefit the extraction of more discriminative point-wise features. Moreover, our proposed model is competitive in comparison to the supervised methods that train the entire network on ShapeNet-Part, which shows the robustness of our learnt 3D point cloud features that are used to only train the part classifier.

4.4. Downstream Task 3: Point Cloud Registration

Dataset. Other than studying the content features in classification and part segmentation, here we adopt a pose-relevant downstream task, i.e., registration, to evaluate the pose features learnt by our proposed model. We conduct experiments for point cloud registration based on ModelNet40. We follow the same procedure of data preparation as in DCP [40], in which the whole ModelNet40 dataset is randomly split into training and testing sets regardless of object categories. For each point cloud \mathcal{P} , a randomly-drawn rigid transformation \mathcal{T} is applied on \mathcal{P} to obtain the transformed point cloud $\mathcal{T}(\mathcal{P})$, where $\{\mathcal{P}, \mathcal{T}(\mathcal{P})\}$ together with \mathcal{T} becomes an input-output pair for learning the registration model.

Experimental Setting. We adopt the same architecture for the registration model and follow the training procedure proposed by DCP [40] which uses DGCNN [41] as their feature extractor. Our proposed model is first pre-trained on ShapeNet, and then we replace the feature extractor in DCP with our pre-trained pose encoder E_{pose} to train the registration task. Note that E_{pose} is fixed during the training process to verify the effectiveness of our learnt features. For evaluation, we focus on estimating the orientation/rotation between point clouds in this experiments and adopt the root mean squared error (RMSE) as the evaluation metric.

Results. As shown in Table 4, the pose features learnt from our model pre-trained on the ShapeNet dataset outperform other self-supervised methods on RMSE(R). Moreover, our proposed model is competitive against the state-of-the-

Table 4. **Registration results on the ModelNet40 dataset.** We compare our proposed model with supervised methods reported by DCP [40] and self-supervised methods reproduced by using the official implementations. (†: reproduced by training the pretext model on the ShapeNet dataset.)

Methods	SSL	RMSE(R) ↓
PointNetLK [4]	✗	15.095
FGR [5]	✗	9.363
DCP-v2 [40] (PN)	✗	7.061
DCP-v2 [40] (DGCNN)	✗	1.143
PointGLR† [34]	✓	3.900
GraphTER† [16]	✓	2.927
PN_ours	✓	6.719
DGCNN_ours	✓	1.557

art supervised methods that fully train the entire network on ModelNet40, which verifies the robustness of our pose features on the point cloud registration task.

4.5. Ablation Study

Comparisons between Content and Pose Features. In the previous experimental results, we adopt the content feature for the tasks of classification and part segmentation, while the pose feature is utilized for the registration. Here we provide results of using different features for these downstream tasks. Table 5 provides the results of applying either the content or pose feature on all the downstream tasks.

We observe that, in comparison to the pose feature, the content feature works best for the classification and part segmentation tasks, where the former needs more semantic and holistic understanding of the point clouds while the latter needs fine-grained modeling on the local geometric structures. This verifies that the content feature learnt by our proposed method is able to compactly model both the global and local structural information of the input point cloud. When it comes to the downstream task of registration, the pose feature instead contributes better than the content feature. Such results are also aligned with the findings from several prior works [4, 31, 37], where the camera pose estimation task could provide benefits for the registration problem.

Table 5. **Content v.s. pose features.** Comparisons between content and pose features in the classification, part segmentation and registration tasks. (MN40: ModelNet40; SON: we use “OBJ_ONLY” on ScanObjectNN; Regis.: we use RMSE(R) as evaluation metric.)

Backbone	Feature	Classification		Part Seg.	Regis.
		MN40	SON		
PN	content	90.11	83.99	83.83	6.875
	pose	88.98	78.49	83.54	6.719
DGCNN	content	91.98	87.26	85.05	2.920
	pose	90.36	82.79	84.46	1.557

Comparisons among Different Model Designs. We study different model designs here. In addition to our full model, we experiment other designs: 1) the framework of only using the completion branch as proposed by the work of [39] (denoted as “Comp-only”); 2) the framework of only using the pose regression branch (denoted as “PR-only”); 3) the joint learning framework (denoted as “JL”), where the completion and pose regression branches share one encoder trained in a multi-tasking manner. More architecture details of these model variants are provided in the supplementary material. The quantitative comparison among these designs is provided in Table 6.

As we expected, Comp-only can learn more semantic information and obtain better results than PR-only in the classification and part segmentation tasks, while PR-only learns more camera related information and obtains better results than Comp-only in the registration task. Furthermore, our full model can extract more robust and effective features than the other designs, especially in the real-world data classification task. Interestingly, JL obtains worse results than PR-only in the classification task. We hypothesize that, as completion and pose regression have quite different characteristics, having one encoder to jointly learn both tasks may result in worse feature learning, thus verifying again the contribution of our proposed disentanglement method.

Table 6. **Comparisons with different model designs.** Note that “Comp-only” only uses the completion branch, “PR-only” only adopts the pose regression branch, and “JL” indicates the joint learning manner with the shared encoder. (MN40: ModelNet40; SON: we use “OBJ_ONLY” on ScanObjectNN.)

Backbone	Methods	Classification		Part Seg.	Regis.
		MN40	SON		
PN	Comp-only	89.71	83.13	83.72	7.858
	PR-only	89.42	79.00	83.60	6.865
	JL	89.34	78.49	83.77	6.843
	Ours full	90.11	83.99	83.83	6.719
DGCNN	Comp-only	91.37	86.06	84.93	3.231
	PR-only	90.64	82.62	84.47	1.644
	JL	90.52	81.93	84.49	1.741
	Ours full	91.98	87.26	85.05	1.557

5. Conclusions

In this paper, we propose a self-supervised framework for learning feature representations of 3D point clouds. Based on the objectives composed of completion, reconstruction, and pose regression for the partial point cloud data, our model learns to disentangle the content and pose factors. Our learnt content and pose feature representations of 3D point clouds experimentally demonstrate the superior performance in comparison to other self-supervised methods in various downstream tasks such as classification, part segmentation, and registration.

References

- [1] Panos Achlioptas, Olga Diamanti, Ioannis Mitliagkas, and Leonidas Guibas. Learning representations and generative models for 3d point clouds. In *International Conference on Machine Learning (ICML)*, 2018. 2, 7
- [2] Antonio Alliegro, Davide Boscaini, and Tatiana Tommasi. Joint supervised and self-supervised learning for 3d real-world challenges. *ArXiv:2004.07392*, 2020. 1, 3, 6
- [3] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 2, 5
- [4] Yasuhiro Aoki, Hunter Goforth, Rangaprasad Arun Srivatsan, and Simon Lucey. Pointnetk: Robust & efficient point cloud registration using pointnet. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 8
- [5] Raouf Benjemaa and Francis Schmitt. Fast global registration of 3d sampled surfaces using a multi-z-buffer technique. *Image and Vision Computing*, 1999. 8
- [6] Dimitri P Bertsekas. Auction algorithms for network flow problems: A tutorial introduction. *Computational optimization and applications*, 1992. 4
- [7] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *ArXiv:1512.03012*, 2015. 5
- [8] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2016. 3
- [9] Ye Chen, Jinxian Liu, Bingbing Ni, Hang Wang, Jiancheng Yang, Ning Liu, Teng Li, and Qi Tian. Shape self-correction for unsupervised point cloud understanding. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 3, 6
- [10] Blender Online Community. *Blender - a 3D modelling and rendering package*. Blender Foundation, 2018. 5
- [11] Corinna Cortes and Vladimir Vapnik. Support-vector networks. *Machine learning*, 1995. 6
- [12] Haowen Deng, Tolga Birdal, and Slobodan Ilic. Ppf-foldnet: Unsupervised learning of rotation invariant 3d local descriptors. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [13] Benjamin Eckart, Wentao Yuan, Chao Liu, and Jan Kautz. Self-supervised learning on 3d point clouds by learning discrete generative models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 6
- [14] Haoqiang Fan, Hao Su, and Leonidas J Guibas. A point set generation network for 3d object reconstruction from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 4
- [15] Matheus Gadelha, Aruni RoyChowdhury, Gopal Sharma, Evangelos Kalogerakis, Liangliang Cao, Erik Learned-Miller, Rui Wang, and Subhransu Maji. Label-efficient learning on point clouds using approximate convex decompositions. In *European Conference on Computer Vision (ECCV)*, 2020. 5, 6
- [16] Xiang Gao, Wei Hu, and Guo-Jun Qi. Graphner: Unsupervised learning of graph transformation equivariant representations via auto-encoding node-wise transformations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 6, 7, 8
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2014. 3
- [18] Zhizhong Han, Xiyang Wang, Yu-Shen Liu, and Matthias Zwicker. Multi-angle point cloud-vae: unsupervised feature learning for 3d point clouds from multiple angles by joint self-reconstruction and half-to-half prediction. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 2, 6, 7
- [19] Kaveh Hassani and Mike Haley. Unsupervised multi-task feature learning on point clouds. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 1, 6
- [20] Siyuan Huang, Yichen Xie, Song-Chun Zhu, and Yixin Zhu. Spatio-temporal self-supervised representation learning for 3d point clouds. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 3, 6
- [21] Longlong Jing, Yucheng Chen, Ling Zhang, Mingyi He, and Yingli Tian. Self-supervised feature learning by cross-modality and cross-view correspondences. *ArXiv:2004.05749*, 2020. 6
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. In *International Conference on Learning Representations (ICLR)*, 2014. 3
- [23] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2016. 2
- [24] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. 2
- [25] Yu-Jhe Li, Ci-Siang Lin, Yan-Bo Lin, and Yu-Chiang Frank Wang. Cross-dataset person re-identification via unsupervised pose disentanglement and adaptation. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 3
- [26] Minghua Liu, Lu Sheng, Sheng Yang, Jing Shao, and Shi-Min Hu. Morphing and sampling network for dense point cloud completion. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2020. 2, 4, 5
- [27] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2
- [28] Mehdi Mirza and Simon Osindero. Conditional generative adversarial nets. *ArXiv:1411.1784*, 2014. 3
- [29] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. In *International Conference on Machine Learning (ICML)*, 2017. 3

- [30] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 1, 3, 4
- [31] Trong Phuc Truong, Masahiro Yamaguchi, Shohei Mori, Vincent Nozick, and Hideo Saito. Registration of rgb and thermal point clouds generated by structure from motion. In *IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, 2017. 8
- [32] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 1, 2, 4, 5, 6, 7
- [33] 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 (NeurIPS)*, 2017. 1, 2
- [34] Yongming Rao, Jiwen Lu, and Jie Zhou. Global-local bidirectional reasoning for unsupervised representation learning of 3d point clouds. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 1, 6, 7, 8
- [35] Jonathan Sauder and Bjarne Sievers. Self-supervised deep learning on point clouds by reconstructing space. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. 1, 3, 5, 6
- [36] Kihyuk Sohn, Honglak Lee, and Xinchen Yan. Learning structured output representation using deep conditional generative models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2015. 3
- [37] Chi-Yi Tsai and Chih-Hung Huang. Indoor scene point cloud registration algorithm based on rgb-d camera calibration. *Sensors*, 2017. 8
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. 1
- [39] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 3, 8
- [40] Yue Wang and Justin M Solomon. Deep closest point: Learning representations for point cloud registration. In *IEEE International Conference on Computer Vision (ICCV)*, 2019. 7, 8
- [41] 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, 4, 5, 6, 7
- [42] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015. 5
- [43] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas J Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European Conference on Computer Vision (ECCV)*, 2020. 1, 3
- [44] Yifan Xu, Tianqi Fan, Mingye Xu, Long Zeng, and Yu Qiao. Spidercnn: Deep learning on point sets with parameterized convolutional filters. In *European Conference on Computer Vision (ECCV)*, 2018. 2
- [45] Juyoung Yang, Pyunghwan Ahn, Doyeon Kim, Haeil Lee, and Junmo Kim. Progressive seed generation auto-encoder for unsupervised point cloud learning. In *IEEE International Conference on Computer Vision (ICCV)*, 2021. 2, 6
- [46] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 6
- [47] 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. *SIGGRAPH Asia*, 2016. 7
- [48] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *International Conference on 3D Vision (3DV)*, 2018. 2
- [49] Yongheng Zhao, Tolga Birdal, Haowen Deng, and Federico Tombari. 3d point capsule networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 6
- [50] Keyang Zhou, Bharat Lal Bhatnagar, and Gerard Pons-Moll. Unsupervised shape and pose disentanglement for 3d meshes. In *European Conference on Computer Vision (ECCV)*, 2020. 3
- [51] Yang Zou, Xiaodong Yang, Zhiding Yu, BVK Kumar, and Jan Kautz. Joint disentangling and adaptation for cross-domain person re-identification. In *European Conference on Computer Vision (ECCV)*, 2020. 3