

PointDC: Unsupervised Semantic Segmentation of 3D Point Clouds via Cross-modal Distillation and Super-Voxel Clustering

Zisheng Chen^{1*}, Hongbin Xu^{1,2*}, Weitao Chen², Zhipeng Zhou³, Haihong Xiao¹, Baigui Sun²,
Xuansong Xie², Wenxiong kang^{1,4†}

¹South China University of Technology

²Alibaba Group, ³Chinese Academy of Science, ⁴Pazhou Laboratory

halveschen@163.com hongbinxu1013, hillskyxm@gmail.com auwxkang@scut.edu.cn

Abstract

Semantic segmentation of point clouds usually requires exhausting efforts of human annotations, hence it attracts wide attention to the challenging topic of learning from unlabeled or weaker forms of annotations. In this paper, we take the first attempt for fully unsupervised semantic segmentation of point clouds, which aims to delineate semantically meaningful objects without any form of annotations. Previous works of unsupervised pipeline on 2D images fails in this task of point clouds, due to: 1) Clustering Ambiguity caused by limited magnitude of data and imbalanced class distribution; 2) Irregularity Ambiguity caused by the irregular sparsity of point cloud. Therefore, we propose a novel framework, PointDC, which is comprised of two steps that handle the aforementioned problems respectively: Cross-Modal Distillation (CMD) and Super-Voxel Clustering (SVC). In the first stage of CMD, multi-view visual features are back-projected to the 3D space and aggregated to a unified point feature to distill the training of the point representation. In the second stage of SVC, the point features are aggregated to super-voxels and then fed to the iterative clustering process for excavating semantic classes. PointDC¹ yields a significant improvement over the prior state-of-the-art unsupervised methods, on both the ScanNet-v2 (+18.4 mIoU) and S3DIS (+11.5 mIoU) semantic segmentation benchmarks.

1. Introduction

Semantic segmentation of 3D point cloud is a crucial problem that assigns each individual point to a known ontology. The segmentation models can delineate objects at

*These authors contributes equally.

†Corresponding author.

¹The code is released at: <https://github.com/SCUT-BIP-Lab/PointDC>.

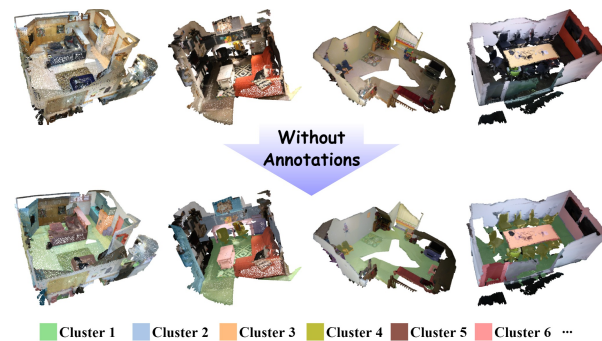


Figure 1. From unannotated point clouds, we would like a segmentation system to discover the semantic concepts automatically without any supervision.

a fine granularity of object boundary, which is helpful for multiple downstream applications, such as robotic navigation, autonomous vehicles, and scene parsing. Despite the immense progress of fully-supervised schemes in 3D semantic segmentation, the success crucially relies on large-scale datasets and annotations. Unfortunately, it requires exhausting efforts to conduct semantic-level per-point annotations (e.g. ≈ 22.3 minutes per indoor scene for annotation [7]).

To reduce the efforts on the tedious process of semantic annotations, several works were proposed to create 3D semantic segmentation systems that can be trained from weaker forms of annotations, including projected 2D images [32], subcloud-level [34], segment-level [28], and point-level annotations [23, 35, 14]. However, few works attempt to handle the great challenge of 3D semantic segmentation without any form of human annotations or motion cues. In this paper, we aim to build an unsupervised 3D semantic segmentation framework that can automatically excavate meaningful semantic features from point clouds, as shown in Fig. 1.

Some efforts of unsupervised semantic segmentation

have been witnessed in the research field of 2D visual images. Independent Information Clustering (IIC) [16] and PiCIE [5] try to excavate semantically meaningful features through self-supervised learning based on transformation invariance and equivariance, meantime conducting a clustering process to optimize the compactness of the learned semantic clusters. STEGO [12] further distills from a self-supervisedly pretrained Transformer to form compact semantic clusters of the corpora. Whereas the existing line of unsupervised clustering pipeline could not be simply migrated to 3D point clouds for the following reasons:

1) **Clustering Ambiguity:** For a 2D unsupervised system, the major premise is that the pictures are meaningful images collected by human (rather than meaningless case like an empty image with a single color). The nature prior during the collection of human enables the effectiveness of large-scale unsupervised clustering on 2D images as long as the dataset is large enough. However, for a 3D unsupervised system, on the one hand, the huge cost for data collection limits the magnitude and diversity of the dataset; on the other hand, the imbalanced occupancy of 3D space aggravates the long-tail distribution effect among different classes when clustering among points, resulting in the ignorance of classes with fewer points.

2) **Irregularity Ambiguity:** Without a regular grid-like structure, point clouds might have variations in the density of local areas. During the calculation of clusters, the points from dense areas inherently weigh more than the points from sparse areas. As a result, the original K-means clustering in [16, 5] might overly focus on the dense regions and ignores the sparse regions.

In this paper, we take the first attempt for unsupervised 3D semantic segmentation and introduce **PointDC** (**P**oint cloud **C**ross-modal **D**istillation and **S**uper-Voxel **C**lustering), which is capable of discovering and segmenting semantic objects from point clouds without any human annotations. Directing to handle the aforementioned problems of Clustering Ambiguity and Irregularity Ambiguity, we respectively adopt Cross-Modal Distillation (CMD) and Super-Voxel Clustering (SVC) in our PointDC framework.

1) To handle the former problem, CMD could integrate the multi-view visual clues to distill the corresponding point feature in the point cloud. As cognitive scientists [8, 27] argue, humans are proficient at mapping the visual concepts learned from 2D images to understand the 3D world. We obtain multi-view images by observing 3D point clouds from different viewpoints first and feed them to a self-supervisedly pretrained visual model such as DINO [4] to extract unsupervised visual features. By back-projecting the multi-view visual features into the corresponding points in 3D space, we can aggregate the features from different views to formulate a unified multi-view representation, and distill the learning of point representation. The involvement

of multi-view visual cues can effectively diminish the ambiguity during clustering among points, and provide a coarse understanding of 3D semantic features.

2) To handle the latter problem, instead of clustering on the original point space, SVC rasterizes the 3D space into super-voxels and assigns each point to the corresponding super-voxel. The features of points in the same voxel are aggregated together via Super-Voxel Pooling for a unified permutation-invariant representation. During each iteration phase of clustering process, we first assign the point features to the super-voxels and then cluster among these voxels. Afterward, the feature of each super-voxel is assigned back to the occupying points, and assume these points in the local super-voxel share a common semantic feature.

The overall pipeline of our PointDC framework includes 2 steps: 1) CMD is utilized to distill the learning of point representation first; 2) Then SVC is applied iteratively to optimize the clustered semantic representation on point clouds. For evaluation, we conduct extensive experiments on the challenging ScanNet-v2 [7] and S3DIS [2]. Compared with state-of-the-art on existing unsupervised methods for point cloud semantic segmentation, our PointDC achieves an improvement on both the ScanNet-v2 (+18.4 mIOU) and S3DIS (+11.5 mIOU).

In summary, our contribution is threefold.

- We take the first attempt for unsupervised 3D semantic segmentation without any kinds of human annotations.
- We propose PointDC, a novel framework for unsupervised 3D semantic segmentation. It is comprised of 2 steps: 1) Cross-Modal Distillation that distills multi-view visual features to the point-based representations; 2) Super-Voxel Clustering that regularizes the point features with voxelized representation through super-voxel pooling, and iteratively clusters to optimize the semantic features on point clouds.
- The proposed method achieves superior improvement compared with existing unsupervised methods for 3D semantic segmentation on various challenging datasets, demonstrating its effectiveness.

2. Related Work

3D Semantic Segmentation 3D semantic segmentation approaches can be divided into 2 categories: point-based methods [25, 26, 33, 21, 19, 18] and voxel-based methods [11, 6]. In point-based methods [33, 21, 19, 18], the information of points is fused from their neighboring areas computed from K-NN or spherical search for effective 3D representations. In voxel-based methods, the points in the 3D space are converted to voxels with 3D-grid structure. In these voxelized representations, standard convolution operations can be applied to extract features from 3D information. Due to the sparsity of point clouds, sparse convolution

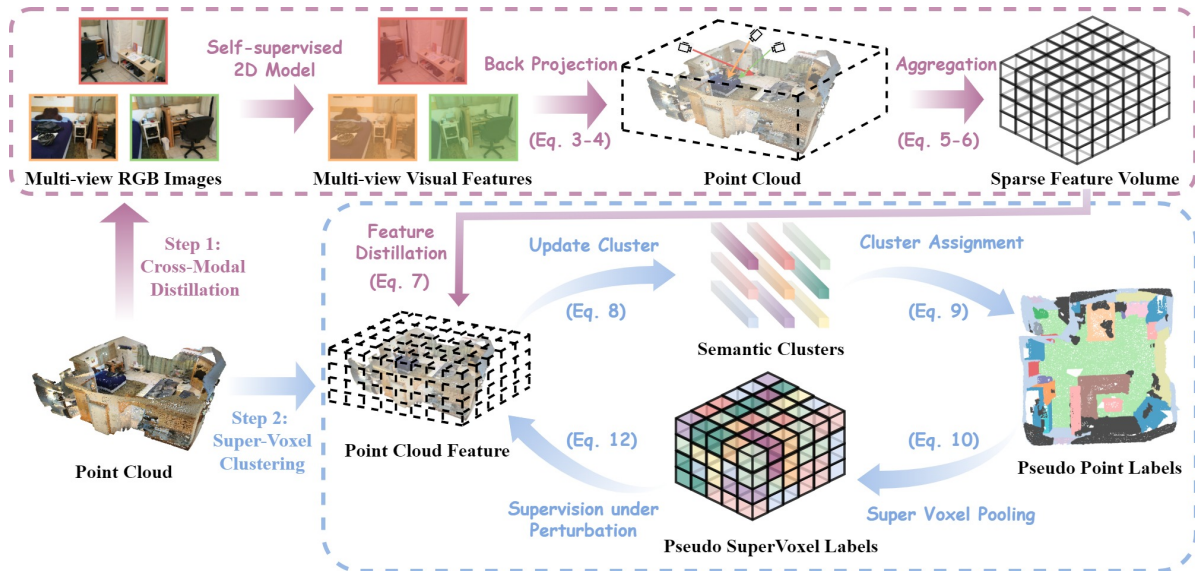


Figure 2. Overview of PointDC framework. The training contains 2 steps: Cross-Modal Distillation and Super-Voxel Clustering.

[11, 6] is adopted to process the voxelized representation of point clouds. Recently, the Transformer structure [40] is also used to handle point clouds, as a novel alternative to the classic convolutional structure. However, most of the previous works are designed for fully-supervised schemes or utilizing weaker forms of annotations [35]. In this work, we focus on handling the 3D semantic segmentation problem without using any human annotations.

2D Unsupervised Learning There has been a number of recent progresses in unsupervised 2D semantic segmentation. DeepCluster [3] obtains latent semantic representations by clustering in the low-dimensional feature space. IIC [16] proposes invariant-information clustering in pixel-level representation to learn by clustering with a self-supervised manner. PiCIE [5] leverages the inductive bias of photometric invariance and geometric equivariance, bringing two sets of images with different transformations close to each other in the feature space. MaskContrast [30] uses an unsupervised saliency detection model to obtain a binary encoding of the silhouette of salient objects in the image, and further adopts contrastive learning to pull the distances of pixel-wise features inside the mask and push the ones between different masks. STEGO [12] extracts features from pretrained models and proposes a novel contrastive loss that encourages features to form compact clusters while preserving the relationships across the corpora.

3D Unsupervised Learning 3D unsupervised learning can be roughly divided into two categories: generation-based methods[31, 20] and contrastive-learning-based methods[36, 13, 15]. Generation-based methods let the model to complete the input point clouds which are occluded or down-sampled. Contrastive-learning-based methods define the augmentation (rotation, color jittering, differ-

ent views) of the given point clouds as positive samples and others as negative samples.

Cross-Modal Learning Cross-Modal learning utilizes data from different modalities[1, 37, 22]. [37] directly converts a 2D pretrained model into a point cloud model using filter dilation. [1] defines images as strong positive samples, then enforces both intra-modal and cross-modal global feature correspondence in the invariant space.

3. Method

In this work, we first attempt in learning unsupervised 3D semantic segmentation models only given uncurated and unlabeled datasets of point clouds. This section begins by introducing the problem statement (Sec. 3.1). We formulate this task as point-level clustering among the whole point cloud dataset to discover semantic meaningful clusters, and introduce the preliminary of learning by clustering pipelines in Sec. 3.2. To handle the problem of Clustering Ambiguity and Irregularity Ambiguity in existing clustering pipelines, we propose Cross-Modal Distillation (CMD) (Sec. 3.3) and Super-Voxel Clustering (SVC) (Sec. 3.4). Finally, the overall training process of our PointDC framework is introduced in Sec. 3.5. The overview of the PointDC is shown in Fig. 2.

3.1. Problem Statement

Given the unlabeled dataset \mathcal{U} from some domain \mathcal{D} which has M point clouds in total. The i -th point cloud $P_i \in \mathbb{R}^{N \times 6}$ represents a scene with N points of 6 dimensions including the XYZ coordinates and RGB intensities of points. On this dataset \mathcal{U} , we aim to discover a set of virtual semantically meaningful classes \mathcal{C} and learn a semantic

feature extractor f_θ parameterized by θ . When provided an unseen point cloud from domain \mathcal{D} during evaluation, f_θ should be able to assign every point a label from the discovered classes \mathcal{C} .

3.2. Preliminary

We begin with preliminaries of prior works that learn an end-to-end neural network for clustering unlabeled data [16, 5, 12, 24, 38]. The key point in these works is that clustering data into classes requires strong feature representation, meantime the learning of feature representations also needs precise class labels. To handle this chicken-and-egg problem, the simplest solution is the one defined by DeepCluster [29]. Following the procedure of an E-M algorithm, we can alternate between clustering using currently extracted feature representation, and using the clustered results as pseudo-labels to supervise the training of the feature extractor. We can still follow this simple strategy for point cloud semantic segmentation task, by alternating the clustering process among instance-wise feature to point-wise feature representation.

Concretely, suppose that we have a set of unlabeled point clouds P_i , where i represents the index in the dataset. The extracted feature tensor is denoted as $f_\theta(P_i) \in \mathbb{R}^{N \times D}$, where D is the feature dimension. Denote $f_\theta(P_i)[j]$ as the feature vector on the j -th point of point cloud P_i and $\mu \in \mathbb{R}^{C \times D}$ as the randomly initialized cluster centroids. The baseline of learning by clustering can be summarized as follows:

1. Optimizing the following object function and use K-Means to cluster the current feature among all points in the dataset:

$$\min_{y, \mu} \sum_{i, j} \|f_\theta(P_i)[j] - \mu[y_{ij}]\| \quad (1)$$

where y_{ij} is the assigned label of the j -th point of i -th point cloud.

2. Use the clustered labels as pseudo-labels to train the segmentation network:

$$\min_{\theta, \omega} \sum_{i, j} L_{CE}(g_\omega(f_\theta(P_i)[j]), y_{ij}) \quad (2)$$

where L_{CE} is the cross-entropy function and g_ω is the segmentation head parameterized by ω .

3.3. Cross-Modal Distillation

As discussed in Sec. 1, the Clustering Ambiguity problem in point cloud is caused by the extremely imbalanced occupancy of different classes in 3D space and the lack of large-scale point cloud datasets with diversity and magnitude comparable with image datasets. Instead of directly



Figure 3. Visualization of the clustering results among multi-view feature maps extracted by DINO [4]. It demonstrates that the multi-view features are semantically correlated.

clustering on the point cloud dataset, we propose Cross-Modal Distillation (CMD) as an initialization step before clustering. In intuition, the multi-view visual features are semantically correlated coarsely, as shown in Fig. 3. Hence, the distillation from multi-view visual modality to 3D point cloud modality can provide a reliable initialization for clustering.

In CMD, the point cloud P_i is converted to images I_{iv} by observing P_i on different viewpoints, where v is the index of viewpoints. Given a self-supervised pretrained 2D neural network h , we can obtain the visual feature map $h(I_{iv})$ of each image I_{iv} . Suppose that the intrinsic and extrinsic matrix of camera k is respectively K_v and T_v . By projecting the 3D point cloud P_i to each of the multi-view images, we can calculate the cross-modal correspondence between the 3D points and 2D pixels.

Given the j -th point of P_i , we can calculate its projection \hat{p}_{ijv} on view v via:

$$z_{ijv} \hat{p}_{ijv} = K_v T_v P_i[j] \quad (3)$$

where z_{ijv} is the depth value on $\hat{p}_{ijv} = [u_{ijv}, v_{ijv}, 1]^T$. u_{ijv} and v_{ijv} are the image pixel coordinates on width and height.

Then, we can filter the invalid projections which are outside the imaging areas with $u_{ijv} > W$, $u_{ijv} < 0$, $v_{ijv} > H$, and $v_{ijv} < 0$. Denote that the remaining projections as \tilde{p}_{ijv} , and its depth value as \tilde{z}_{ijv} . Since there might exist multiple projections to the same pixel, we filter the occluded points by finding the projection with minimum depth value:

$$j^* = \arg \min_j \tilde{z}_{ijv} \quad (4)$$

where j^* represents the index of the corresponding point on point cloud P_i .

Afterward, we can warp the pixel-wise feature to the corresponding point indexed by j^* :

$$H_{iv} = h(I_{iv})[\tilde{p}_{ij^*v}] \quad (5)$$

where $H_{iv} \in \mathbb{R}^{N \times V \times D}$ is the point-wise feature projected from visual features on multiple views.

Suppose that the super-voxel segmentation function is $S_{vox}(\cdot)$. The feature $F_i \in \mathbb{R}^{M \times D}$ on the super-voxel $S_{vox}(P_i) \in \mathbb{R}^{M \times 3}$ converted from original point cloud P_i can be calculated through:

$$F_i[k] = \frac{1}{\|\mathcal{N}(k)\|} \sum_{j \in \mathcal{N}(k)} \max_v H_{iv}[j] \quad (6)$$

where k is the index of super-voxel and $\mathcal{N}(k)$ represents the indices of points belonging to the k -th super-voxel. The multi-view features are firstly aggregated via Max-Pooling among different views and then the feature on each super-voxel is aggregated via Avg-Pooling among the occupying points.

Finally, in the stage of CMD, the point feature extractor is then supervised by the visual feature aggregated from multi-view images. CMD optimizes the following function while traversing the whole dataset:

$$\min_{\theta} \sum_i \left\| \left(\frac{1}{\|\mathcal{N}(k)\|} \sum_{j \in \mathcal{N}(k)} (f_{\theta}(P_i)[j]) \right) - F_i \right\|_2^2 \quad (7)$$

where we first aggregate the point feature $f_{\theta}(P_i)$ in each super-voxel via AVG-Pooling and distill it with F_i .

3.4. Super-Voxel Clustering

As discussed in Sec. 1, the Irregularity Ambiguity problem comes from the irregular structure of point cloud. In a scene, the dense points might disturb the distance-based clustering metric, leading to the ignorance of sparse points. To handle this issue, we propose Super-Voxel Clustering (SVC), an iterative learning by clustering pipeline that alternates point-based clustering with super-voxel-based clustering.

In SVC, given the point cloud $P_i \in \mathbb{R}^{N \times 3}$. Denote that the feature extractor $f_{\theta}(\cdot)$ is parameterized by θ . The number of clusters is C , the output feature dimension of $f_{\theta}(\cdot)$ is D , and the number of super-voxels is M . The learning process can be summarized as follows:

1. Extract point-wise feature from point cloud P_i and update the super-voxelized features via K-Means:

$$y, \mu^* = \arg \min_{y, \mu} \sum_{i, k} \|S_{vox}(f_{\theta}(P_i))[k] - \mu[y_{ik}]\|_2^2 \quad (8)$$

where i is the index of point cloud in the dataset and k is the index of super-voxel. The super-voxel aggregation function $S_{vox}(\cdot)$ aggregates the feature of points belonging to same super-voxel: $S_{vox}(f_{\theta}(P_i)) = \frac{1}{\|\mathcal{N}(k)\|} \sum_{j \in \mathcal{N}(k)} f_{\theta}(P_i)[j]$. $y \in \mathbb{R}^{M \times C}$ is the assigned label on super-voxels, and $\mu^* \in \mathbb{R}^{C \times D}$ is the clustered centroids.

2. To assign the label to each point of the point cloud, we use the distance to the clustered centroids μ^* to create

soft assignment of probabilities towards each class.

$$\hat{y}_{ij} = -\log \left(\frac{e^{-\cos(f_{\theta}(P_i)[j], \mu^*[\hat{y}_{ij}])}}{\sum_l e^{-\cos(f_{\theta}(P_i)[j], \mu_l^*)}} \right) \quad (9)$$

where \hat{y}_{ij} is the assigned point-wise pseudo labels on the j -th point of point cloud P_i . $\cos(\cdot, \cdot)$ is the cosine distance.

3. Then, we adopt Super-Voxel Pooling to filter the assigned point-wise labels in each voxel of the super-voxel.

$$\tilde{y}_{ik} = \frac{1}{\|\mathcal{N}(k)\|} \sum_{j \in \mathcal{N}(k)} \hat{y}_{ij} \quad (10)$$

where \tilde{y}_{ik} is the soft label filtered by the prior of super-voxels and k is the index of super-voxel.

4. Finally, we can train the point-wise representation with \tilde{y}_{ik} .

$$\theta^* = \arg \min_{\theta} \left\| \sum_k \sum_{j \in \mathcal{N}(k)} L_{CE}(f_{\theta}(P_i)[j], \phi(\tilde{y}_{ik})) \right\|_2^2 \quad (11)$$

where L_{CE} is the cross-entropy function, and $\phi(\cdot) = \text{onehot}(\arg \max(\cdot))$ converts the input an ont-hot label.

5. Repeat previous 1-4 operations iteratively.

Furthermore, the inductive bias of invariance and equivariance to specific transformations are used to promote the pseudo-label training in Eq. 11. Following the self-supervised prior, the extracted features should maintain the invariance towards transformations like color jittering or gaussian noise, and the extracted feature should be warped or rotated accordingly to the geometric transformation on the original point cloud to keep the equivariance. Denote that $\pi_{inv}(\cdot)$ and $\pi_{equ}(\cdot)$ are respectively the transformation for invariance and equivariance. Then the inductive bias could be appended to promote Eq. 11:

$$\theta^* = \arg \min_{\theta} \left\| \sum_k \sum_{j \in \mathcal{N}(k)} L_{CE}(f_{\theta}(\pi_{equ}(\pi_{inv}(P_i)))[j], \phi(\pi_{equ}(\tilde{y}_{ik}))) \right\|_2^2 \quad (12)$$

3.5. Overall Framework

As presented in Fig. 2, our PointDC framework includes two training stages: CMD and SVC. In the first stage of CMD, we can obtain multi-view images by observing the point cloud from different viewpoints. A self-supervisedly pretrained 2D model is used to extract the feature maps from multi-view images. Then we can back-project (Eq. 5) each pixel of the multi-view images to its corresponding point of the point cloud by calculating Eq. 3 and 4. Since one point might have multiple projections on different images, we aggregate the cross-view feature via the Global Max-Pooling in Eq. 6. Following the assumption that each super-voxel

Methods	Unsupervised		Linear Probe	
	mIoU	Acc	mIoU	Acc
[3] DeepCluster	3.88	19.75	4.04	23.55
[16] IIC	3.98	20.47	4.00	23.06
[5] PiCIE	4.10	22.81	4.34	27.04
[36] PC-HC	4.63	21.75	4.72	43.09
[36] PC-NCE	3.93	21.24	4.29	44.21
[31] OcCo	3.17	19.97	3.37	21.65
[1] CrossPoint	3.81	20.52	3.94	22.92
[13] CSC	4.64	18.24	5.31	28.5
[15] STRL	4.13	19.38	4.25	29.70
PointDC	25.74	63.69	28.78	71.62

Table 1. Comparison of unsupervised segmentation on the ScanNet-v2 validation set. PointDC significantly outperforms prior art in both unsupervised clustering and linear probe metrics.

Method	mIoU
[3] DeepCluster	3.7
[16] IIC	3.7
[5] PiCIE	3.9
[36] PC-HC	3.9
[36] PC-NCE	3.8
[13] CSC	4.5
[15] STRL	4.1
PointDC	22.9

Table 2. Comparison of unsupervised segmentation on the ScanNet-v2 test set. The results of the online benchmark are reported.

contains similar semantics, we further aggregate the features of points belonging to the same super-voxel via the Global Avg-Pooling in Eq. 6. The feature extractor is then distilled by the multi-view cues in Eq. 7. In the second stage of SVC, we perform K-Means clustering on the super-voxels aggregated from point-wise features in local regions in Eq. 8. Then we assign the label to each point of the point cloud in a non-parametric manner based on the distance to clusters in Eq. 9. The label for each super-voxel is then aggregated via Super-Voxel Pooling among points located in the same voxel in Eq. 10. Finally, the pseudo label of super-voxel is used to train the feature extractor under random perturbation of transformations in Eq. 12.

4. Experiments

4.1. Experiment Details

Datasets and Metric: We conduct experiments on 2 point cloud benchmarks, ScanNet-v2[7] and S3DIS[2]. ScanNet-v2[7] contains 1613 3D scans from 707 unique indoor scenes, all annotated with 20 classes. Following the official setting, we use 1201 scenes and 312 scenes as training set and validation set, respectively. The remaining 100 scenes are used as test set. S3DIS[2] contains 271 indoor scenes with 13 classes. We follow the official train/validation split, training on Areas 1,2,3,4,6 and then testing on Area 5. As our method requires image data and camera intrinsic and

Method	mIoU	Acc
[3] DeepCluster	5.46	19.75
[16] IIC	5.33	21.47
[5] PiCIE	5.90	25.05
[36] PC-HC	9.27	26.87
[36] PC-NCE	8.86	23.32
[13] CSC	11.09	34.83
[15] STRL	10.21	37.40
PointDC	22.59	54.10

Table 3. Comparison of unsupervised segmentation on the S3DIS validation set (Area 5).

extrinsic parameters, we use ScanNet-v2 2D data and 2D-3D-S[2]. 2D-3D-S contains multi-view images corresponding to the scenes in S3DIS, the corresponding depth maps as well as the internal and external camera parameters. We utilize the intersection-over-union as evaluation metric of the 3D semantic segmentation results, and report the mean result (mIOU) over all categories for comparison with other approaches. Moreover, we also utilize the accuracy over all categories in the results.

Implementation Details: In the stage of CMD, we experiment with pretrained image segmentation model STEGO[12] noted as h in Sec. 3.3. For the backbone of point feature extractor f_θ , we use the Sparse 3D-UNet[10] from [11]. The points are voxelized following the procedure of [17] and then fed to f_θ for feature extraction. For the super-voxels segmentation [9], we use the mesh segment results [7] for ScanNet-v2. For the super-voxel partition in S3DIS, we utilize the geometric partition results described in [19].

Unsupervised Clustering and Linear Probe: Since we are agnostic to the ground truth label, the clustered results might have a random permutation of order compared with the ground truth. To evaluate the quality of an unsupervised method, we follow the 2 protocols used in previous works [30, 5]: unsupervised clustering and linear probe. For unsupervised clustering, we do not have access to the ground truth labels, but we can use a Hungarian matching algorithm to align our unlabeled clusters and the ground truth labels for evaluation. This measures how consistent the predicted semantic segments are with the ground truth annotations and diminish the aforementioned permutations of the predicted class labels. For linear probe, we train a linear projection from the features to the class labels with cross-entropy loss. This measures the feature quality of the learned feature representation.

4.2. 3D Unsupervised Semantic Segmentation

Evaluation on ScanNet-v2: We conduct the unsupervised clustering and linear probe test on the ScanNet-v2 validation set and report the results in Tab. 1. Previous methods of learning by clustering including DeepCluster [3], IIC [16],

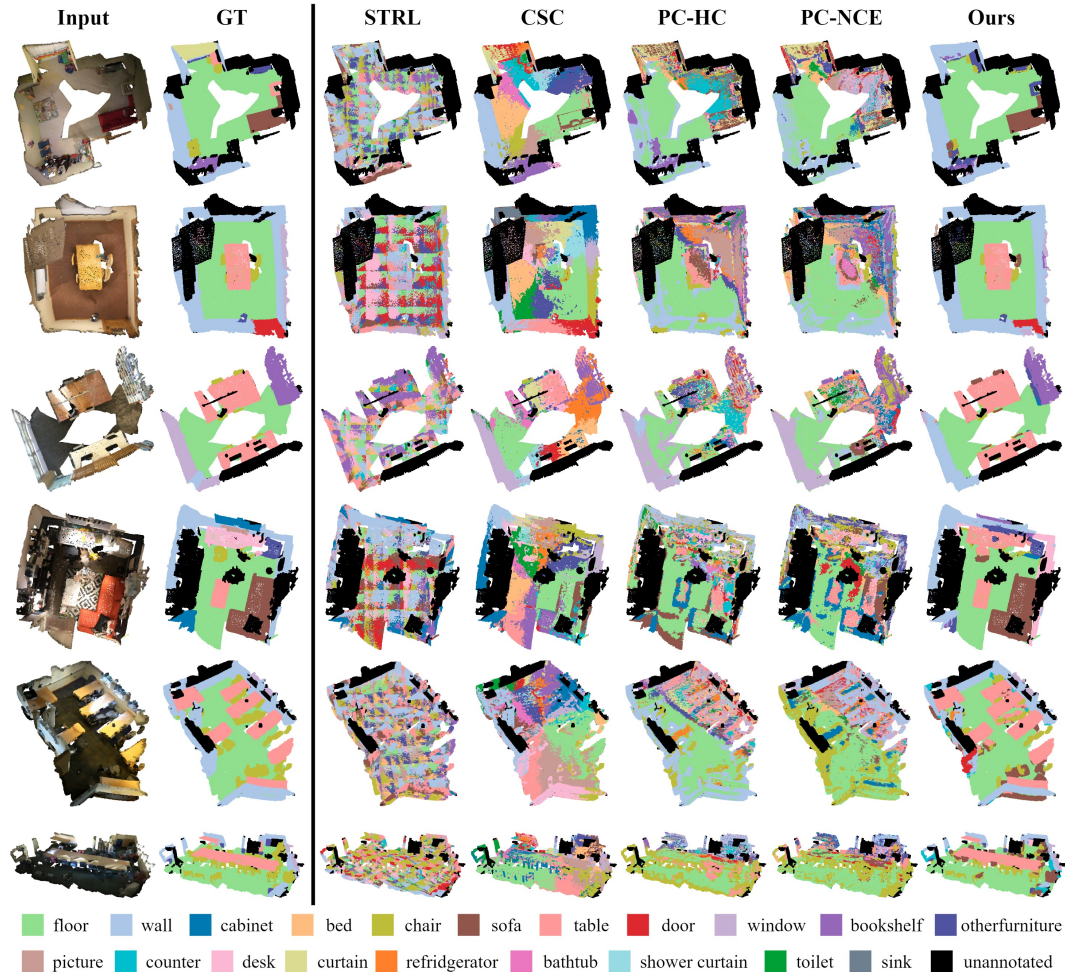


Figure 4. Qualitative comparison of unsupervised segmentation on ScanNet-v2 validation set. Each of the aligned ground truth labels and clusters is assigned a color. For better understanding, we show some the color and name matches in the bottom.

Random Clustering	Multi-View Clustering	CMD (Eq. 7)	Basic-SVC (Eq. 8 and Eq. 12)	Non-parametric Classifier (Eq. 9)	Super-Voxel Pooling (Eq. 10)	mIoU	Acc
✓						3.86	13.18
✓	✓					13.58	34.00
✓	✓	✓				20.29	42.12
✓	✓	✓	✓			23.64	61.53
✓	✓	✓	✓	✓		24.85	61.79
✓	✓	✓	✓	✓	✓	25.74	63.69

Table 4. Ablation experiments of PointDC on the ScanNet-v2 validation set.

PiCIE [5] are compared in the table. Furthermore, we also compare with previous state-of-the-art unsupervised pre-training methods for point clouds including: PointContrast [36], OcCo [31], CrossPoint [1], CSC [13] and STRL [15]. In the table, PC-HC and PC-NCE respectively represent the PointContrast model trained with Hardest-Contrastive loss and PointInfoNCE loss respectively. Without using any kinds of human annotations, our method outperforms all other methods as shown in the table. In particular, PointDC improves by **+21.10** unsupervised mIoU, **+40.88** unsupervised accuracy, **+23.47** linear probe mIoU, and **27.41** lin-

ear probe compared with the next best baseline. In Tab. 2, we present the evaluation results of unsupervised methods on the ScanNet-v2 test set. Similarly, we can also find a large improvement of **+18.4** unsupervised mIoU compared with the next best baseline. Moreover, the qualitative comparisons on the ScanNet-v2 validation set are provided in Fig. 4. In this table, we also compare our method with a concurrent work GrowSP [39] which has good clustering performance, and the results prove that the proposed method achieve better performance. As the figure reveals, our method is able to precisely locate semantically mean-




Figure 5. Visualization of PointDC’s segmentation results under different iterations during training.

ingful objects compared with other unsupervised methods.

Evaluation on S3DIS: We also evaluate the proposed method on S3DIS dataset to validate the effectiveness of the proposed method, and show the results in Tab. 3. The unsupervised semantic segmentation methods of DeepCluster [3], IIC [16], PiCIE [5] and the unsupervised pre-training point cloud methods of PointContrast [36], CSC [13], STRL [15] are used for comparison in the table. As shown in the table, PointDC achieves an improvement of **+11.50** unsupervised mIoU and **+16.70** unsupervised Accuracy.

4.3. Ablation Experiment

Ablation Study of PointDC Framework: To validate the effectiveness of the proposed PointDC framework, we conduct ablation experiments of each related component of the framework, and present the results in Tab. 4. In the table, ‘Random Clustering’ means the baseline of learning by clustering on 3D point clouds discussed in Sec. 3.2. ‘Multi-view Clustering’ means that a 2D unsupervised learning by clustering framework [12] is conducted on the multi-view images of point clouds, and the 2D segmentation results are fused to construct 3D segmentation results. It is compared to demonstrate the difference of results existing framework clustered on 2D images and our PointDC framework clustered on 3D point clouds. ‘CMD’ represents the model trained with Cross-Modal Distillation (Eq. 7), and ‘Basic SVC’ includes a simplified version of Super-Voxel-Clustering, only utilizing Eq. 8 and Eq. 12 in the training. ‘Non-parametric Classifier’ means that we adopt assign the label to each point based on the distance to the cluster centroids instead of a learnable classifier in previous clustering pipelines [5], as shown in Eq. 9. ‘Super-Voxel Pooling’ means that we apply AVG-Pooling towards the point-wise pseudo labels in super-voxels to filter the noise caused by the irregularity of points, as shown in Eq. 10. From Tab. 4, it can be found that each component of PointDC can improve the performance of unsupervised 3D segmentation effectively.

Ablation Study of Iteration: Since our PointDC framework is a learning-by-clustering framework, the model should converge as the clusters optimize in different iterations. Hence, we conduct experiments of PointDC under different iterations to validate the results. As shown in Fig. 5, we visualize the segmentation results of the proposed method under different iterations. From the figure, we can

Method	mIoU	Acc
[31] OcCo	3.17	19.97
[1] CrossPoint	3.81	20.52
[15] STRL	4.13	19.38
PointDC*	11.50	39.65
PointDC	14.06	45.36

Table 5. Comparison of unsupervised segmentation methods with the same backbone of DGCNN on ScanNet-v2 validation set. * denotes that only CMD is used.

find that the model becomes better and better along with the training iterations and the update of clustering results.

4.4. Free-Model

To validate the effectiveness of our method on different models, we additionally employ DGCNN [33] for 3D unsupervised semantic segmentation. With the same backbone of DGCNN, unsupervised methods of OcCo [31], Cross-Point [1], and STRL [15] are used for comparison in Tab. 5. It demonstrates that our method outperforms previous best results with **+9.93** unsupervised mIoU and **+24.84** unsupervised accuracy.

5. Conclusion

We take the first attempt at the challenging topic of unsupervised semantic segmentation of 3D point clouds without any human annotations, and introduce a novel framework, PointDC. It contains two steps: Cross-Modal Distillation (CMD) and Super-Voxel Clustering (SVC). In the first stage of CMD, the multi-view features of the point cloud are back-projected to the 3D space and aggregated together in the super-voxels to distill the training of point representation. In the next stage of SVC, the point representations are aggregated to super-voxels and then fed to the iterative clustering process for learning semantically meaningful representations. As the evaluation results on different point cloud benchmarks, our method achieves superior performance on both the ScanNet-v2 (+18.4 mIoU) and S3DIS (+11.5 mIoU).

6. Acknowledgement

This work was supported by the National Natural Science Foundation of China (No.61976095) and the Nat-

ural Science Foundation of Guangdong Province, China (No.2022A1515010114). This work was also supported by Alibaba Group through Alibaba Research Intern Program.

References

- [1] Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga Rodrigo. Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9902–9912, 2022.
- [2] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017.
- [3] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 132–149, 2018.
- [4] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021.
- [5] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16794–16804, 2021.
- [6] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3075–3084, 2019.
- [7] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017.
- [8] Judy S DeLoache, Mark S Strauss, and Jane Maynard. Picture perception in infancy. *Infant behavior and development*, 2:77–89, 1979.
- [9] Pedro F Felzenszwalb and Daniel P Huttenlocher. Efficient graph-based image segmentation. *International journal of computer vision*, 59:167–181, 2004.
- [10] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. *arXiv: Computer Vision and Pattern Recognition*, 2017.
- [11] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9224–9232, 2018.
- [12] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snavely, and William T Freeman. Unsupervised semantic segmentation by distilling feature correspondences. *arXiv preprint arXiv:2203.08414*, 2022.
- [13] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15587–15597, 2021.
- [14] Qingyong Hu, Bo Yang, Guangchi Fang, Yulan Guo, Aleš Leonardis, Niki Trigoni, and Andrew Markham. Sqn: Weakly-supervised semantic segmentation of large-scale 3d point clouds. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVII*, pages 600–619. Springer, 2022.
- [15] Siyuan Huang, Yichen Xie, Song-Chun Zhu, and Yixin Zhu. Spatio-temporal self-supervised representation learning for 3d point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6535–6545, 2021.
- [16] Xu Ji, Joao F Henriques, and Andrea Vedaldi. Invariant information clustering for unsupervised image classification and segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9865–9874, 2019.
- [17] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and Pattern recognition*, pages 4867–4876, 2020.
- [18] Loic Landrieu and Mohamed Boussaha. Point cloud over-segmentation with graph-structured deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7440–7449, 2019.
- [19] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4558–4567, 2018.
- [20] Ruihui Li, Xianzhi Li, Chi-Wing Fu, Daniel Cohen-Or, and Pheng-Ann Heng. Pu-gan: A point cloud upsampling adversarial network. *International Conference on Computer Vision*, 2019.
- [21] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *Advances in neural information processing systems*, 31, 2018.
- [22] Yueh-Cheng Liu, Yu-Kai Huang, Hung-Yueh Chiang, Hung-Ting Su, Zhe-Yu Liu, Chin-Tang Chen, Ching-Yu Tseng, and Winston H Hsu. Learning from 2d: Contrastive pixel-to-point knowledge transfer for 3d pretraining. *arXiv preprint arXiv:2104.04687*, 2021.
- [23] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. One thing one click: A self-training approach for weakly supervised 3d semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1726–1736, 2021.
- [24] Guofeng Mei, Litao Yu, Qiang Wu, Jian Zhang, and Mohammed Bennamoun. Unsupervised learning on 3d point clouds by clustering and contrasting. *arXiv preprint arXiv:2202.02543*, 2022.
- [25] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification

- and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [26] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [27] Susan Ann Rose. Infants’ transfer of response between two-dimensional and three-dimensional stimuli. *Child Development*, pages 1086–1091, 1977.
- [28] An Tao, Yueqi Duan, Yi Wei, Jiwen Lu, and Jie Zhou. Seg-group: Seg-level supervision for 3d instance and semantic segmentation. *IEEE Transactions on Image Processing*, 31:4952–4965, 2022.
- [29] Kai Tian, Shuigeng Zhou, and Jihong Guan. Deepcluster: A general clustering framework based on deep learning. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2017, Skopje, Macedonia, September 18–22, 2017, Proceedings, Part II 17*, pages 809–825. Springer, 2017.
- [30] Wouter Van Gansbeke, Simon Vandenhende, Stamatios Georgoulis, and Luc Van Gool. Unsupervised semantic segmentation by contrasting object mask proposals. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10052–10062, 2021.
- [31] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9782–9792, 2021.
- [32] Haiyan Wang, Xuejian Rong, Liang Yang, Jinglun Feng, Jizhong Xiao, and Yingli Tian. Weakly supervised semantic segmentation in 3d graph-structured point clouds of wild scenes. *arXiv preprint arXiv:2004.12498*, 2020.
- [33] 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)*, 38(5):1–12, 2019.
- [34] Jiacheng Wei, Guosheng Lin, Kim-Hui Yap, Tzu-Yi Hung, and Lihua Xie. Multi-path region mining for weakly supervised 3d semantic segmentation on point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4384–4393, 2020.
- [35] Yushuang Wu, Zizheng Yan, Shengcai Cai, Guanbin Li, Yizhou Yu, Xiaoguang Han, and Shuguang Cui. Pointmatch: a consistency training framework for weakly supervised semantic segmentation of 3d point clouds. *arXiv preprint arXiv:2202.10705*, 2022.
- [36] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 574–591. Springer, 2020.
- [37] Chenfeng Xu, Shijia Yang, Tomer Galanti, Bichen Wu, Xiangyu Yue, Bohan Zhai, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with 2d image pretrained models. *arXiv preprint arXiv:2106.04180*, 2021.
- [38] Ling Zhang and Zhigang Zhu. Unsupervised feature learning for point cloud by contrasting and clustering with graph convolutional neural network. *arXiv preprint arXiv:1904.12359*, 2019.
- [39] Zihui Zhang, Bo Yang, Bing Wang, and Bo Li. Growsp: Unsupervised semantic segmentation of 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17619–17629, 2023.
- [40] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021.