

# GrowSP: Unsupervised Semantic Segmentation of 3D Point Clouds

Zihui Zhang, Bo Yang\*, Bing Wang, Bo Li  
 Shenzhen Research Institute, The Hong Kong Polytechnic University  
 vLAR Group, The Hong Kong Polytechnic University  
 zihui.zhang@connect.polyu.hk, bo.yang@polyu.edu.hk




Figure 1. Given an input point cloud with complex structures from S3DIS dataset [2], our GrowSP automatically discovers accurate semantic classes simply by progressively growing superpoints, without needing any human labels in training.

## Abstract

We study the problem of 3D semantic segmentation from raw point clouds. Unlike existing methods which primarily rely on a large amount of human annotations for training neural networks, we propose the first purely unsupervised method, called **GrowSP**, to successfully identify complex semantic classes for every point in 3D scenes, without needing any type of human labels or pretrained models. The key to our approach is to discover 3D semantic elements via progressive growing of superpoints. Our method consists of three major components, 1) the feature extractor to learn per-point features from input point clouds, 2) the superpoint constructor to progressively grow the sizes of superpoints, and 3) the semantic primitive clustering module to group superpoints into semantic elements for the final semantic segmentation. We extensively evaluate our method on multiple datasets, demonstrating superior performance over all unsupervised baselines and approaching the classic fully-supervised PointNet. We hope our work could inspire more advanced methods for unsupervised 3D semantic learning.

## 1. Introduction

Giving machines the ability to automatically discover semantic compositions of complex 3D scenes is crucial for many cutting-edge applications. In the past few years, there has been tremendous progress in fully-supervised semantic segmentation for 3D point clouds [14]. From the seminar works PointNet [40] and SparseConv [12] to a plethora of

recent neural models [21, 27, 41, 55, 60], both the accuracy and efficiency of per-point semantic estimation have been greatly improved. Unarguably, the success of these methods primarily relies on large-scale human annotations for training deep neural networks. However, manually annotating real-world 3D point clouds is extremely costly due to the unstructured data format [3, 20]. To alleviate this problem, a number of recent methods start to use fewer 3D point labels [19, 69], cheaper 2D image labels [59, 77], or active annotations [22, 63] in training. Although achieving promising results, they still need tedious human efforts to annotate or align 3D points across images for particular datasets, thus being inapplicable to novel scenes without training labels.

In this paper, we make the first step towards unsupervised 3D semantic segmentation of real-world point clouds. To tackle this problem, there could be two strategies: 1) to naïvely adapt existing unsupervised 2D semantic segmentation techniques [4, 7, 24] to 3D domain, and 2) to apply existing self-supervised 3D pretraining techniques [17, 66] to learn discriminative per-point features followed by classic clustering methods to obtain semantic categories. For unsupervised 2D semantic methods, although achieving encouraging results on color images, they can be hardly extended to 3D point clouds primarily because: a) there is no general pretrained backbone to extract high-quality features for point clouds due to the lack of representative 3D datasets akin to ImageNet [46] or COCO [29], b) they are usually designed to group pixels with similar low-level features, *e.g.* colors or edges, as a semantic class, whereas such a heuristic is normally not satisfied in 3D point clouds due to point sparsity and spatial occlusions. For self-supervised 3D pretraining methods, although the pretrained per-point

\*Corresponding author

features could be discriminative, they are lack of semantic meanings fundamentally because the commonly adopted data augmentation techniques do not explicitly capture categorical information. Section 4 clearly demonstrates that all these methods fail catastrophically on 3D point clouds.

Given a sparse point cloud composed of multiple semantic categories, we can easily observe that a relative small local point set barely contains distinctive semantic information. Nevertheless, when the size of a local point set is gradually growing, that surface patch naturally emerges as a basic element or primitive for a particular semantic class, and then it becomes much easier for us to identify the categories just by combining those basic primitives. For example, two individual 3D points sampled from a spacious room are virtually meaningless, whereas two patches might be easily identified as the *back* and/or *arm* of chairs.

Inspired by this, we introduce a simple yet effective pipeline to automatically discover per-point semantics, simply by progressively growing the size of per-point neighborhood, without needing any human labels or pretrained backbone. In particular, our architecture consists of three major components: 1) a per-point feature extractor which is flexible to adopt an existing (untrained) neural network such as the powerful SparseConv [12]; 2) a superpoint constructor which progressively creates larger and larger superpoints during training to guide semantic learning; 3) a semantic primitive clustering module which aims to group basic elements of semantic classes via an existing clustering algorithm such as K-means. The key to our pipeline is the superpoint constructor together with a progressive growing strategy in training. Basically, this component drives the feature extractor to progressively learn similar features for 3D points within a particular yet **growing superpoint**, while the features of different superpoints tend to be pushed as distinct elements of semantic classes. Our method is called **GrowSP** and Figure 1 shows qualitative results of an indoor 3D scene. Our contributions are:

- We introduce the first purely unsupervised 3D semantic segmentation pipeline for real-world point clouds, without needing any pretrained models or human labels.
- We propose a simple strategy to progressively grow superpoints during network training, allowing meaningful semantic elements to be learned gradually.
- We demonstrate promising semantic segmentation results on multiple large-scale datasets, being clearly better than baselines adapted from unsupervised 2D methods and self-supervised 3D pretraining methods. Our code is at: <https://github.com/vLAR-group/GrowSP>

## 2. Related Works

**Learning with Strong Supervision:** With the advancement of 3D scanners, acquiring point clouds becomes easier

and cheaper. In past five years, the availability of large-scale human-annotated point cloud datasets [2, 3, 10, 16, 20, 28, 52, 57] enables fully-supervised neural methods to achieve remarkable 3D semantic segmentation results. These methods generally include: 1) 2D projection-based methods [9, 25, 36, 62] which project raw point clouds onto 2D images followed by mature 2D neural architectures to learn semantics; 2) voxel-based methods [8, 12, 26, 35, 78] which usually voxelize unstructured point clouds into regular spheres, cubes, or cylinders followed by existing convolutional networks; 3) point-based methods [13, 21, 27, 34, 41, 55, 64, 76] which primarily follows the seminal PointNet [40] to directly learn per-point features using shared MLPs. The performance of these methods can be further improved by the successful self-supervised pre-training techniques in recent studies [5, 18, 23, 38, 42, 54, 58, 67, 71, 75]. Although achieving excellent accuracy on existing benchmarks, they require densely-annotated 3D data for training. This is extremely costly and prohibitive in real applications.

**Learning with Weak Supervision:** To alleviate the cost of human annotations, a number of works have started to learn 3D semantics using fewer or cheaper human labels in training. These weak labels primarily include: 1) fewer 3D point labels [19, 31, 33, 49, 56, 65, 69, 73], and 2) sub-cloud/seg-level/scene-level labels [6, 30, 43, 53, 61]. The performance of these weakly-supervised methods can also be boosted by self-supervised pre-training techniques [18, 67, 74, 75]. Apart from these weak labels, supervision signals can also come from other domains such as labeled 2D images [32, 44, 50, 68, 70, 79] or pretrained language models [15, 45, 72]. Although obtaining encouraging results, these methods still need tedious human efforts to annotate or align data points. Fundamentally, they still cannot automatically discover semantic classes.

**Unsupervised Semantic Learning:** The work [48] learns point semantics by recovering voxel positions after randomly shuffling 3D points, and Canonical Capsules [51] learns to decompose point clouds into object parts via self-canonicalization. However, both of them can only work on simple object point clouds. Technically, existing self-supervised 3D pretraining techniques [17, 66] can be used for unsupervised semantic segmentation, just by learning discriminative per-point features followed by clustering. However, as shown in Section 4.5, the pretrained point features are actually lack of semantic meanings and fail to be grouped as classes. For 2D images, a number of recent works [4, 7, 11, 24, 37] tackle the problem of unsupervised 2D semantic segmentation. However, due to the domain gap between images and point clouds, there is no existing work showing their applicability in 3D space. In fact, as demonstrated in Section 4, both the representative 2D methods IIC [24] and PICIE [7] fail catastrophically on point clouds, while our method achieves significantly better accuracy.

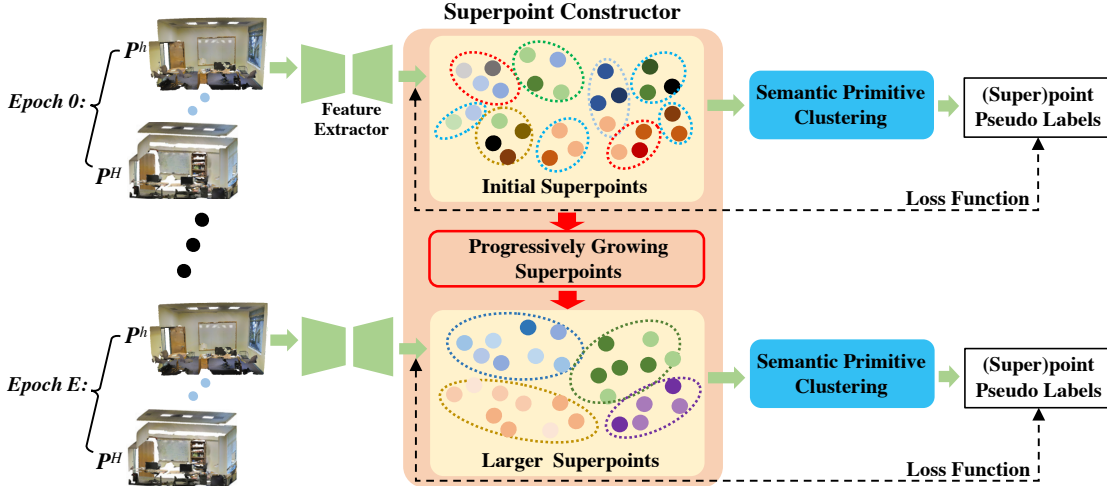


Figure 2. The general learning framework of our GrowSP. It mainly consists of three components: 1) the feature extractor which learns per-point features, 2) the superpoint constructor which progressively grows the sizes of superpoints, and 3) the semantic primitive clustering module which aims to group superpoints into semantic elements. **Note that, the superpoint constructor is no longer needed in testing.**

### 3. GrowSP

#### 3.1. Overview

Our method generally formulates the problem of unsupervised 3D semantic segmentation as joint 3D point feature learning and clustering in the absence of human labels. As shown in Figure 2, from a dataset with  $H$  point clouds  $\{\mathbf{P}^1 \dots \mathbf{P}^h \dots \mathbf{P}^H\}$ , given one single scan  $\mathbf{P}^h$  with  $N$  points as input, *i.e.*,  $\mathbf{P}^h \in \mathbb{R}^{N \times 6}$  where each point has a location  $\{x, y, z\}$  with color if available, the **feature extractor** firstly obtains per-point features  $\mathbf{F}^h \in \mathbb{R}^{N \times K}$  where the embedding length  $K$  is free to predefine, *e.g.*  $K = 128$ . We simply adopt the powerful SparseConv architecture [12] without any pretraining step as our feature extractor. Implementation details are in Appendix.

Having the input point cloud  $\mathbf{P}^h$  and its point features  $\mathbf{F}^h$  at hand which are not meaningful in the very beginning, we will then feed them into our **superpoint constructor** to progressively generate larger and larger superpoints over more and more training epochs, as detailed in Section 3.2. These superpoints will be fed into our **semantic primitive clustering module**, generating pseudo labels for all superpoints, as discussed in Section 3.3. During training, these pseudo labels will be used to optimize the feature extractor.

#### 3.2. Superpoint Constructor

This module is designed to divide each input point cloud into pieces, such that each piece as whole ideally belongs to the same category. Intuitively, compared with individual points, a single piece is more likely to have geometric meanings, thus being easier to extract high-level semantics. In order to construct high-quality superpoints and aid the network to automatically discover semantics, here we ask two key questions:

- First, what strategy should we use to construct superpoints? Naturally, if a superpoint keeps small, it can be highly homogeneous but lack of semantics. On the other hand, a larger superpoint may have better semantics but is error-prone if not constructed properly. In this regard, we propose to gradually grow the size of superpoints from small to large over more and more training epochs.
- Second, how to partition a point cloud into satisfactory pieces at the beginning, such that the network training can be bootstrapped effectively? Considering that point neural features are virtually meaningless in the early stage of network training, it is more reliable to simply leverage classic algorithms to obtain initial superpoints based on geometric features, *e.g.* surface normal or connectivity.

With these insights, we introduce the following mechanism to construct superpoints.

**Initial Superpoints:** As shown in the yellow block of Figure 2, at the beginning of network training, the initial superpoints are constructed by VCCS [39] followed by a region growing algorithm [1]. They jointly take into account the spatial/normal/normalized RGB distances between 3D points. For a specific input point cloud  $\mathbf{P}^h$ , its initial superpoints are denoted as  $\{\tilde{\mathbf{p}}_1^h \dots \tilde{\mathbf{p}}_{m^0}^h \dots \tilde{\mathbf{p}}_{M^0}^h\}$  where each superpoint  $\tilde{\mathbf{p}}_{m^0}^h$  consists of a small subset of original point cloud  $\mathbf{P}^h$ . Note that, for different point clouds, the number of their initial superpoints  $M^0$  are usually different. Implementation details are in Appendix. Figure 3 shows an example of initial superpoints for an indoor room.

**Progressively Growing Superpoints during Training:** Assuming the feature extractor is trained for epochs using Algorithm 1 which will be detailed in Section 3.4, the per-point features are expected to be more meaningful. In this regard, we turn to primarily use neural features to progres-




Figure 3. An example of initial superpoints constructed by VCCS and region growing. Each colored patch represents a superpoint.

sively construct larger superpoints for future training. As illustrated in Figure 4, each dot represents the neural embedding of a 3D point, and a red circle indicates an initial superpoint. The blue circle represents a larger superpoint absorbing one or multiple initial superpoints.




Figure 4. An illustration of progressively growing superpoints.

In particular, for a specific input point cloud  $\mathbf{P}^h$ , we have its neural features  $\mathbf{F}^h \in \mathbb{R}^{N \times K}$  and initial superpoints  $\{\tilde{\mathbf{p}}_1^h \cdots \tilde{\mathbf{p}}_{m^0}^h \cdots \tilde{\mathbf{p}}_{M^0}^h\}$ . Firstly, we will compute the mean neural features for initial superpoints, denoted as  $\{\tilde{\mathbf{f}}_1^h \cdots \tilde{\mathbf{f}}_{m^0}^h \cdots \tilde{\mathbf{f}}_{M^0}^h\}$ :

$$\tilde{\mathbf{f}}_{m^0}^h = \frac{1}{Q} \sum_{q=1}^Q \mathbf{f}_q^h, \quad \mathbf{f}_q^h \in \mathbb{R}^{1 \times K} \quad (1)$$

where  $Q$  is the total number of 3D points within an initial superpoint  $\tilde{\mathbf{p}}_{m^0}^h$ , and  $\mathbf{f}_q^h$  is the feature vector retrieved from  $\mathbf{F}^h$  for the  $q^{\text{th}}$  3D point of the superpoint.

Secondly, having these initial superpoint features, we simply use K-means to group the  $M^0$  vectors into  $M^1$  clusters, where  $M^1 < M^0$ . Each cluster represents a new and larger superpoint. In total, we get new superpoints:

$$\{\tilde{\mathbf{p}}_1^h \cdots \tilde{\mathbf{p}}_{m^1}^h \cdots \tilde{\mathbf{p}}_{M^1}^h\} \xleftarrow{\text{Kmeans}} \{\tilde{\mathbf{f}}_1^h \cdots \tilde{\mathbf{f}}_{m^0}^h \cdots \tilde{\mathbf{f}}_{M^0}^h\}$$

Note that this superpoint growing step is conducted independently on each input point cloud. The much smaller  $M^1$ , the more aggressive this growing step.

After every a certain number of training epochs, *i.e.* one round, we will compute the next level of larger superpoints by repeating above two steps. Given  $T$  levels of growing, the number of superpoints for an input point cloud will be reduced from  $M^1 \rightarrow M^2 \rightarrow M^t$  until to a small value  $M^T$ . In each epoch, all superpoints of the entire dataset will be fed into the semantic primitive clustering module.

### 3.3. Semantic Primitive Clustering

For every epoch, each input point cloud will have a number of superpoints, each of which representing a particular part of objects or stuff. As to the whole dataset, all superpoints together can be regarded as a huge set of basic semantic elements or primitives, such as chair backs, table surfaces, *etc.*. In order to discover semantics from these superpoints, two issues need to be addressed:

- First, how to effectively group these superpoints? A straightforward way is to directly cluster all superpoints into a number of object categories using an existing clustering algorithm. However, we empirically find that this is excessively aggressive, because many superpoints belonging to different categories are similar and then wrongly assigned to the same semantic group at the early training stage, and it is hard to be corrected over time. In this regard, we opt to constantly group all superpoints into a relatively large number of clusters in all training epochs.
- Second, are the neural features of superpoints discriminative enough for semantic clustering? Again, considering that the neural features of 3D points as well as superpoints are meaningless at the beginning of network training, it is more reliable to explicitly take into account point geometry features such as surface normal distributions to augment discrimination of superpoints. To this end, for each superpoint, we simply stack both its neural features and the classic PFH feature [47] for clustering.

As shown in the blue block of Figure 2, taking the first epoch as an example, given all  $H$  point clouds in the whole dataset  $\{\mathbf{P}^1 \cdots \mathbf{P}^H\}$ , we have all initial superpoints  $(\{\tilde{\mathbf{p}}_1^1 \cdots \tilde{\mathbf{p}}_{m^0}^1 \cdots\} \cdots \{\tilde{\mathbf{p}}_1^H \cdots \tilde{\mathbf{p}}_{m^0}^H \cdots\})$  and their features  $(\{\tilde{\mathbf{f}}_1^1 \cdots \tilde{\mathbf{f}}_{m^0}^1 \cdots\} \cdots \{\tilde{\mathbf{f}}_1^H \cdots \tilde{\mathbf{f}}_{m^0}^H \cdots\})$ . Each superpoint's features are geometry augmented:

$$\tilde{\mathbf{f}}_{m^0}^1 = \tilde{\mathbf{f}}_{m^0}^1 \oplus \ddot{\mathbf{f}}_{m^0}^1 \quad (2)$$

where the neural features  $\tilde{\mathbf{f}}_{m^0}^1$  are obtained by Equation 1 and concatenated with 10-dimensional PFH features  $\ddot{\mathbf{f}}_{m^0}^1$ . We simply adopt K-means to cluster all these superpoint features into  $S$  semantic primitives:

$$S \text{ primitives} \xleftarrow{\text{Kmeans}} (\{\hat{\mathbf{f}}_1^1 \cdots \hat{\mathbf{f}}_{m^0}^1 \cdots\} \cdots \{\hat{\mathbf{f}}_1^H \cdots \hat{\mathbf{f}}_{m^0}^H \cdots\})$$

**Loss Function:** Naturally, each superpoint and individual 3D points within it will be given an  $S$ -dimensional one-hot pseudo-label. For all  $S$  primitives, we use the corresponding centroids (PFH simply dropped) estimated by K-means as a classifier to classify all individual 3D points, obtaining  $S$ -dimensional logits. Lastly, the standard cross-entropy loss is applied between logits and pseudo-labels to optimize the neural extractor from scratch.

**Algorithm 1** The training pseudocode of our GrowSP. Given a dataset with  $H$  point cloud scans  $\{\mathbf{P}^1 \dots \mathbf{P}^h \dots \mathbf{P}^H\}$ .  $\hat{E}$  is a predefined number of epochs for periodically and progressively growing superpoints. The hyperparameters  $M^1$  is set as 80,  $S$  as 300,  $\hat{E}$  as 10 in all experiments.

*Epoch 0: Initial superpoints construction.*

- Apply VCCS [39] followed by region growing algorithm [1] on each point cloud scan, and obtain all initial superpoints:

$$\tilde{\mathbb{P}}_0 = (\{\tilde{\mathbf{p}}_1^1 \dots \tilde{\mathbf{p}}_{m^0}^1 \dots \tilde{\mathbf{p}}_{M^0}^1\} \dots \{\tilde{\mathbf{p}}_1^h \dots \tilde{\mathbf{p}}_{m^0}^h \dots \tilde{\mathbf{p}}_{M^0}^h\} \dots \{\tilde{\mathbf{p}}_1^H \dots \tilde{\mathbf{p}}_{m^0}^H \dots \tilde{\mathbf{p}}_{M^0}^H\});$$

*Note: the number of initial superpoints for different point cloud scans are usually different. We use the same  $M^0$  to avoid an abuse of notation.*

- Load superpoints  $\tilde{\mathbb{P}} \leftarrow \tilde{\mathbb{P}}_0$  for training;
- Predefine the number of superpoints at total  $T$  levels:  $\{M^1 \dots M^t \dots M^T\}$  where  $M^0 > M^1 > M^t > M^T$ ;

*Note: for simplicity, we choose consecutive integers in all experiments, i.e.  $M^T = M^{T-1} - 1 = \dots = M^1 - (T - 1)$ .*

- Initialize superpoint level  $t = 0$ ;

**for training epoch  $e$  in  $\{1, 2, \dots, E\}$  do**

**if  $e\% \hat{E} \neq 0$ : To optimize the neural network**

- Feed all  $H$  point clouds into the feature extractor, and obtain neural features  $\{\mathbf{F}^1 \dots \mathbf{F}^h \dots \mathbf{F}^H\}$ ;
- Obtain neural features according to Equation 1 and PFH features [47] for all superpoints  $\tilde{\mathbb{P}}$ ;
- Apply K-means to cluster all superpoints of the entire dataset into  $S$  semantic primitives, where each superpoint and individual 3D points within it will be assigned a one-hot pseudo label;
- The centroids of  $S$  semantic primitives estimated by K-means are used as a classifier to classify all individual 3D points of the dataset. Cross-entropy loss is applied between the logits and pseudo labels to optimize the whole network.

**if  $e\% \hat{E} == 0$ : To progressively grow superpoints**

- Update superpoint level  $t = t + 1$ , and get the corresponding predefined superpoint number  $M^t$ ;
- For each point cloud, obtain the latest neural features for each initial superpoint, and then apply K-means to cluster these initial superpoints into  $M^t$  new superpoints. For simplicity, we use the same value  $M^t$  for all  $H$  point clouds:

$$\tilde{\mathbb{P}}_t = (\{\tilde{\mathbf{p}}_1^1 \dots \tilde{\mathbf{p}}_{m^t}^1 \dots \tilde{\mathbf{p}}_{M^t}^1\} \dots \{\tilde{\mathbf{p}}_1^h \dots \tilde{\mathbf{p}}_{m^t}^h \dots \tilde{\mathbf{p}}_{M^t}^h\} \dots \{\tilde{\mathbf{p}}_1^H \dots \tilde{\mathbf{p}}_{m^t}^H \dots \tilde{\mathbf{p}}_{M^t}^H\});$$

- Update superpoints  $\tilde{\mathbb{P}} \leftarrow \tilde{\mathbb{P}}_t$

### 3.4. Implementation

**Training Phase:** To better illustrate our GrowSP, Algorithm 1 clearly presents all steps of our pipeline during training. Notably, our method does not need to be given the actual number of semantic classes in training, because we simply learn semantic primitives.

**Testing Phase:** Once the network is well-trained, we keep the centroids of  $S$  semantic primitives estimated by K-means on training split. In testing, these centroids are directly grouped into  $C$  semantic classes using K-means. The newly obtained centroids for the  $C$  classes are used as the final classifier. Given a test point cloud, all per-point neural features are directly classified as one of  $C$  classes, **without needing to construct superpoints anymore**. For the final evaluation metrics calculation, we follow [7] to use Hungarian algorithm to match predicted classes with ground truth labels. Implementation details are in Appendix.

## 4. Experiments

Our method is mainly evaluated on two large-scale indoor datasets and one challenging outdoor LIDAR dataset: S3DIS [2], ScanNet [10] and SemanticKITTI [3]. We also conduct cross-dataset experiments to evaluate the general-

ization ability on unseen scenes, and results are supplied in Appendix. For evaluation metrics, we report the standard mean Intersection-over-Union (mIoU), Overall Accuracy (OA), mean Accuracy (mAcc) of all classes. More quantitative and qualitative results are in Appendix.

Table 1. Quantitative results of our method and baselines on the Area-5 of S3DIS dataset [2]. Only 12 classes are evaluated. The performance standard deviations of unsupervised methods are calculated over the last five checkpoints during the final 50 epochs.

		OA(%)	mAcc(%)	mIoU(%)
Supervised Methods	PointNet [40]	77.5	59.1	44.6
	PointNet++ [41]	77.5	62.6	50.1
	SparseConv [12]	88.4	69.2	60.8
Unsupervised Methods	RandCNN	23.3±2.6	17.3±1.1	9.2±1.2
	van Kmeans	21.4±0.6	21.2±1.6	8.7±0.3
	van Kmeans-S	21.9±0.5	22.9±0.4	9.0±0.2
	van Kmeans-PFH	23.2±0.7	23.6±1.7	10.2±1.4
	van Kmeans-S-PFH	22.8±1.7	20.6±0.7	9.2±0.9
	IIC [24]	28.5±0.2	12.5±0.2	6.4±0
	IIC-S [24]	29.2±0.5	13.0±0.2	6.8±0
	IIC-PFH [24]	28.6±0.1	16.8±0.1	7.9±0.4
	IIC-S-PFH [24]	31.2±0.2	16.3±0.1	9.1±0.1
	PICIE [7]	61.6±1.5	25.8±1.6	17.9±0.9
	PICIE-S [7]	49.6±2.8	28.9±1.0	20.0±0.6
	PICIE-PFH [7]	54.0±0.8	36.8±1.7	24.4±0.6
PICIE-S-PFH [7]	48.4±0.9	40.4±1.6	25.2±1.2	
	<b>GrowSP(Ours)</b>	<b>78.4±1.5</b>	<b>57.2±1.7</b>	<b>44.5±1.1</b>

Table 2. Quantitative results of 6-fold cross validation on S3DIS dataset [2]. Only 12 classes excluding *clutter* are evaluated.

		OA(%)	mAcc(%)	mIoU(%)
Supervised Methods	PointNet [40]	75.9	67.1	49.4
	PointNet++ [41]	77.1	74.1	55.1
	SparseConv [12]	89.4	78.1	69.2
Unsupervised Methods	RandCNN	23.1	18.4	9.3
	van Kmeans	20.0	21.5	8.8
	van Kmeans-S	20.0	22.3	8.8
	van Kmeans-PFH	23.9	24.7	10.9
	van Kmeans-S-PFH	23.4	20.8	9.5
	IIC [24]	32.8	14.7	8.5
	IIC-S [24]	29.4	15.1	7.7
	IIC-PFH [24]	29.5	13.2	6.7
	IIC-S-PFH [24]	26.3	13.6	7.2
	PICIE [7]	46.4	28.1	17.8
	PICIE-S [7]	50.7	30.8	21.6
	PICIE-PFH [7]	55.0	38.8	26.6
	PICIE-S-PFH [7]	49.1	40.5	26.7
<b>GrowSP (Ours)</b>	<b>76.0</b>	<b>59.4</b>	<b>44.6</b>	

#### 4.1. Evaluation on S3DIS

The S3DIS dataset [2] consists of 6 large areas with 271 rooms. Each point belongs to one of 13 categories. We find that the *clutter* class across different rooms does not have consistent geometric patterns and semantic meanings. In the absence of human labels, to automatically discover such diverse geometries as a common category is challenging and also unreasonable. Therefore, in the final testing stage, we only group all points except *clutter* into 12 classes. The *clutter* points are excluded (masked) for metrics calculation. Note that, in training, all points including *clutter* are fed into the network, but *clutter* points are not used for clustering and loss computation. We use the standard 6-fold cross validation in our experiments.  $M^T$  is set as 20 which is slightly larger than the actual 12 semantic classes.

Since there is no existing unsupervised method for semantic segmentation on 3D point clouds, we implement the following four baselines: 1) **RandomCNN** which uses K-means to directly cluster the per-point features into 12 classes, where the features are obtained from the randomly initialized backbone as ours; 2) **vanilla K-means** which directly clusters raw 3D points ( $xyzrgb$ ) into 12 classes; 3) **IIC [24]** which is adapted from the existing unsupervised 2D method, where we change to use the same backbone as ours; 4) **PICIE [7]** which is also adapted from 2D domain with the same backbone as ours. For an extensive comparison, we also adopt the same semantic primitive clustering module on baselines, denoted as van Kmeans-S/IIC-S/PICIE-S. The PFH features are also applied to them, denoted as van Kmeans-S-PFH/IIC-S-PFH/PICIE-S-PFH. Additionally, three classic fully-supervised methods PointNet [40], PointNet++ [41], and SparseConv [12] are also included for comparison. These baselines are all carefully trained and evaluated using the same settings as ours. More details of implementation and experiments are in Appendix.

**Analysis:** As shown in Tables 1&2, our method clearly

outperforms all unsupervised baselines by large margins. Basically, the RandomCNN and vanilla K-means fail to obtain any meaningful semantic classes, due to the lack of meaningful point features. Despite using the same powerful SparseConv [12] as feature extractor, neither IIC [24] nor PICIE [7] can obtain high-quality semantic segmentation results, primarily because neither takes full use of semantics emerging from larger and larger point regions, but instead tends to simply group similar points according to low-level features. Adding semantic primitive clustering module can help PICIE get better results, but it is still significantly worse than our method. Figure 5 shows qualitative results.

#### 4.2. Evaluation on ScanNet

The ScanNet dataset [10] has 1201 rooms for training, 312 rooms for offline validation, and 100 rooms for online hidden testing. Each point belongs to one of 20 object categories or undefined background. Since there is no unsupervised result on the hidden test set and uploading baseline results to online is not allowed, we turn to compare with the above 4 unsupervised baselines on the validation split. For this dataset, the hyperparameter  $M^T$  is chosen as 30 which is slightly larger than 20 categories. During training, all raw point clouds are fed into neural networks, while the undefined points are not used for loss computation and clustering. They will also be masked out during testing.

Table 3. Quantitative results on the **validation split** of ScanNet dataset [10]. All 20 categories are evaluated.

		OA(%)	mAcc(%)	mIoU(%)
Unsupervised Methods	RandCNN	11.9±0.4	8.4±0.1	3.2±0
	van Kmeans	10.1±0.1	10.0±0.1	3.4±0
	van Kmeans-S	10.2±0.1	9.8±0.3	3.4±0.1
	van Kmeans-PFH	10.4±0.2	10.3±0.7	3.5±0.2
	van Kmeans-S-PFH	12.2±0.6	9.3±0.5	3.6±0.1
	IIC [24]	27.7±2.7	6.1±1.2	2.9±0.8
	IIC-S [24]	18.3±2.6	6.7±0.6	3.4±0.1
	IIC-PFH [24]	25.4±0.1	6.3±0	3.4±0
	IIC-S-PFH [24]	18.9±0.3	6.3±0.2	3.0±0.1
	PICIE [7]	20.4±0.5	16.5±0.3	7.6±0
	PICIE-S [7]	35.6±1.1	13.7±1.5	8.1±0.5
	PICIE-PFH [7]	23.1±1.4	14.0±0.1	8.1±0.3
	PICIE-S-PFH [7]	23.6±0.4	15.1±0.6	7.4±0.2
<b>GrowSP (Ours)</b>	<b>57.3±2.3</b>	<b>44.2±3.1</b>	<b>25.4±2.3</b>	

Table 4. Quantitative results on the **online hidden split** of ScanNet dataset [10]. All 20 categories are evaluated.

		mIoU(%)
Supervised Methods	PointNet++ [41]	33.9
	DGCNN [60]	44.6
	PointCNN [27]	45.8
	SparseConv [12]	72.5
Unsupervised Method	<b>GrowSP (Ours)</b>	26.9

**Analysis:** As shown in Table 3, all unsupervised baselines fail on this challenging dataset. IIC and PICIE are just slightly better than RandCNN, demonstrating the level of segmentation difficulty. By contrast, our GrowSP achieves




Figure 5. Qualitative results of our method and baselines. The top row is from S3DIS dataset, the row below from ScanNet. Each color represent one semantic class. Red circles highlight the differences.

very encouraging results. As shown in Table 4, our method obtains a similar mIoU score on the online benchmark. Qualitative results are presented in Figure 5.

### 4.3. Evaluation on SemanticKITTI

The SemanticKITTI dataset [3] consists of 21 sequences of 43552 outdoor LIDAR scans. It has 19130 scans for training, 4071 for validation and 20351 for online testing. Each point belongs to one of 19 semantic categories or undefined background. We set  $M^T$  as 30, and exclude(masked) undefined background points during testing, and ignore them in training. Detailed implementations and results for all categories are in appendix.

Table 5. Quantitative results on the **validation split** of SemanticKITTI dataset [3]. All 19 categories are evaluated.

		OA(%)	mAcc(%)	mIoU(%)
Supervised				
	RandCNN	25.4±3.3	6.0±0.2	3.2±0.1
	van Kmeans	8.1±0	8.2±0.1	2.4±0
	van Kmeans-S	10.3±0.3	7.7±0.1	2.6±0
Unsupervised Methods	van Kmeans-PFH	11.2±0.6	7.5±0.7	2.7±0.1
	van Kmeans-S-PFH	13.2±1.8	8.1±0.4	3.0±0.2
	IIC [24]	26.2±1.5	5.8±0.4	3.1±0.3
	IIC-S [24]	23.9±1.1	6.1±0.3	3.2±0.2
	IIC-PFH [24]	20.1±0.1	7.2±0.1	3.6±0
	IIC-S-PFH [24]	23.4±0	9.0±0	4.6±0
	PICIE [7]	22.3±0.4	14.6±0.3	5.9±0.1
	PICIE-S [7]	18.4±0.5	13.2±0.2	5.1±0.1
	PICIE-PFH [7]	46.6±0.2	10.1±0	4.7±0
	PICIE-S-PFH [7]	42.7±2.1	11.5±0.2	6.8±0.6
	<b>GrowSP (Ours)</b>	<b>38.3±1.0</b>	<b>19.7±0.6</b>	<b>13.2±0.1</b>

Table 6. Quantitative results on the **online hidden split** of SemanticKITTI dataset [3]. All 19 categories are evaluated.

		mIoU(%)
Supervised Methods		
	PointNet [40]	14.6
	PointNet++ [41]	20.1
	SparseConv [12]	53.2
Unsupervised Method	<b>GrowSP (Ours)</b>	<b>14.3</b>

**Analysis:** As shown in Tables 5&6, all unsupervised baselines fail on this outdoor dataset. Our method achieves satisfactory segmentation results on par with the fully-supervised PointNet [40]. Nevertheless, our method cannot automatically discover the minor classes such as *bike*, *cyclist*, *etc.*, which can be seen from the full category results

in appendix. We hypothesize that the failure is caused by the extreme sparsity of minor classes and the lack of discriminative raw features such as RGB colors.

### 4.4. Ablation Study

To evaluate the effectiveness of each component of our pipeline and the choices of hyperparameters, we conduct the following ablation experiments on Area-5 of S3DIS [2].

(1) **Only removing superpoint constructor.** In this setting, no superpoint will be constructed and there is no progressive growing mechanism as well. All per-point features will be directly clustered into 300 semantic primitives in training. Other settings are the same as the full method.

(2) **Only removing semantic primitive clustering.** In this setting, the semantic primitive number  $S$  is directly set as 12 (the actual number of classes). In both training and testing, all superpoints will be grouped as 12 categories.

(3) **Only removing PFH feature in semantic primitive clustering.** This aims to evaluate the benefits of explicitly including geometry features for semantic primitive clustering. All other settings are kept the same.

(4)~(6) **Sensitivity to different voxel sizes of initial superpoints.** When constructing initial superpoints for each point cloud, here we select 3 sizes of voxel grids for VCCS algorithm, *i.e.* {25cm, 50cm, 75cm}. Note that, in all experiments of our full method, we constantly set voxel size as 50cm. Intuitively, the smaller grid size, the more homogeneous initial superpoints, and the less semantics reserved.

(7)~(11) **Sensitivity to different choices of  $M^1$ .** Since  $M^0$  is determined by VCCS and we empirically find that the average number of initial superpoints across S3DIS is about 160, *i.e.*  $M^0 \approx 160$ , here we select 4 groups for  $M^1$ , *i.e.* {160, 120, 80, 40}. In our full method, we always set  $M^1$  as 80 (Algorithm 1).

(12)~(15) **Sensitivity to different choices of  $M^T$ .**  $M^T$  is the ending value of  $M$  during progressive growing, we conduct 4 ablations for  $M^T$ , *i.e.* {60, 40, 20, 12}. In our full method, we always set  $M^T$  as 20 (Algorithm 1).

(16)~(19) **Sensitivity to decreasing speeds.** We conduct 4 ablations for the decreasing speed of superpoint num-

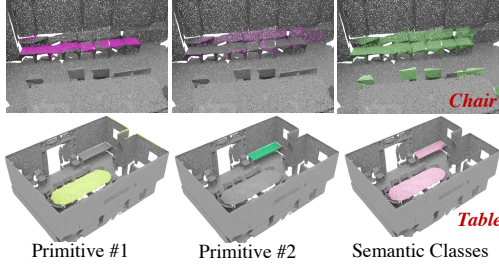


Figure 6. Qualitative results of learned semantic primitives for two classes in S3DIS: chair/table.

bers in the growing process. After every  $\{40, 13, 5, 3\}$  epochs, the number of superpoints  $M^t$  is reduced by 1 in training. In our full method, we set it as 13 (Algorithm 1).

**(20)~(23) Sensitivity to different choices of  $S$ .** The choice of semantic primitive number  $S$  also controls how aggressive our method aims to learn semantics. Here we select 4 different values for comparison, *i.e.*  $\{100, 200, 300, 500\}$ . Our full method always sets  $S$  as 300 (Algorithm 1).

**Analysis:** As shown in Table 7, we can see that: 1) Without superpoint constructor, the method nearly fails with an mIoU score of just 20%, unarguably demonstrating the key role it plays in automatically discovering complex semantic classes. 2) The simple semantic primitive clustering module is also critical, if not the same important as superpoint constructor, showing that learning semantic elements instead of classes can effectively alleviate the difficulty. 3) Explicitly adding geometry based PFH features can indeed aid the network to get better results. 4) Our method is extremely robust to different choices of all hyperparameters such as  $M^1, M^T, S$ , the voxel size for initial superpoints, and the decreasing speed. Figure 6 shows examples of learned semantic primitives. From this ablation study, we clearly see how the proposed components complement each other to achieve excellent performance without any human labels.

#### 4.5. Comparison with Self-supervised Learning

We further compare with existing self-supervised pre-training methods PointConstrast [66] (PC) and Contrastive Scene Contexts [17] (CSC) in the following two groups of experiments. Note that, PC has two versions: InfoNCE and Hardest Contrastive, denoted as PC-I/PC-H.

- **Group 1:** All PC-I/PC-H/CSC/GrowSP are well-trained on the ScanNet training set. In fact, we simply reuse the official pre-trained models of PC-I/PC-H/CSC released by authors. All these trained networks are frozen, followed by K-means to group output features into semantic categories on both ScanNet (val set) and S3DIS (Area 5).
- **Group 2:** We add a single linear classifier to each of the four pre-trained models, and only train the classifiers using full supervision until convergence (about 150 epochs for each method) on both ScanNet (training set) and S3DIS (Areas 1/2/3/4/6). After that, all four models

Table 7. The mIoU scores of all ablated networks on Area-5 of S3DIS based on our full GrowSP.

	mIoU(%)
(1) Remove Superpoint Constructor	20.3±0.4
(2) Remove Semantic Primitive Clustering	25.4±1.0
(3) Remove PFH feature	38.9±0.9
(4) 25cm voxels for initial superpoints	41.3±1.8
<b>(5) 50cm voxels for initial superpoints</b>	<b>44.5±1.1</b>
(6) 75cm voxels for initial superpoints	43.2±0.7
(7) $M^1 = 160$ for progressive growing	43.3±1.3
(8) $M^1 = 120$ for progressive growing	41.3±3.2
<b>(9) <math>M^1 = 80</math> for progressive growing</b>	<b>44.5±1.1</b>
(10) $M^1 = 40$ for progressive growing	43.1±2.0
(12) $M^T = 60$ for progressive growing	42.4±1.1
(13) $M^T = 40$ for progressive growing	43.0±0.8
<b>(14) <math>M^T = 20</math> for progressive growing</b>	<b>44.5±1.1</b>
(15) $M^T = 12$ for progressive growing	38.9±3.0
(16) Decreasing speed 40 for progressive growing	42.2±2.4
(17) Decreasing speed <b>13</b> progressive growing	<b>44.5±1.1</b>
(18) Decreasing speed 5 for progressive growing	43.2±0.9
(19) Decreasing speed 3 progressive growing	42.2±1.0
(20) $S = 100$ for semantic primitive clustering	41.4±1.2
(21) $S = 200$ for semantic primitive clustering	43.5±0.9
<b>(22) <math>S = 300</math> for semantic primitive clustering</b>	<b>44.5±1.1</b>
(23) $S = 500$ for semantic primitive clustering	43.8±1.4
<b>(24) The Full framework (GrowSP)</b>	<b>44.5±1.1</b>

are directly tested on both ScanNet (val set) and S3DIS (Area 5). Basically, this is to evaluate the quality of learned point features via linear probing.

Table 8. OA / mAcc / mIoU (%) scores of Groups 1&2.

	Group 1 (K-means)		Group 2 (Linear Probing)	
	ScanNet	S3DIS	ScanNet	S3DIS
PC-I [66]	27.6 / 10.1 / 5.1	43.8 / 18.6 / 10.4	57.1 / 19.6 / 13.3	64.3 / 32.6 / 23.1
PC-H [66]	29.5 / 12.5 / 5.8	42.8 / 17.5 / 11.3	62.6 / 18.8 / 13.3	63.4 / 36.3 / 25.9
CSC [17]	44.9 / 11.8 / 7.7	43.3 / 22.4 / 13.5	69.3 / 29.5 / 21.8	78.2 / 43.6 / 35.3
<b>Ours</b>	<b>62.9 / 44.3 / 27.7</b>	<b>56.4 / 43.1 / 28.6</b>	<b>73.5 / 42.6 / 31.6</b>	<b>80.1 / 55.4 / 44.7</b>

**Analysis:** As shown in Table 8: 1) All existing self-supervised methods completely fail to estimate semantics using K-means, while our method is significantly better. 2) Given an additional layer for linear probing, our method also shows clearly better segmentation results. We can see that the existing self-supervised pre-trained features actually do not have rich semantic category information, while our method can directly learn semantics.

## 5. Conclusion

We demonstrate that multiple 3D semantic classes can be automatically discovered using a purely unsupervised method from real-world point clouds. By leveraging a simple progressive growing strategy to create larger and larger superpoints over time, our method can successfully learn meaningful semantic elements. Extensive experiments validate the effectiveness of our approach.

**Acknowledgements:** This work was supported in part by Shenzhen Science and Technology Innovation Commission (JCYJ20210324120603011), in part by Research Grants Council of Hong Kong (25207822 & 15225522).



## References

- [1] R. Adams and L. Bischof. Seeded Region Growing. *TPAMI*, 1994. 3, 5
- [2] Iro Armeni, Sasha Sax, Amir R. Zamir, and Silvio Savarese. Joint 2D-3D-Semantic Data for Indoor Scene Understanding. *arXiv:1702.01105*, 2017. 1, 2, 5, 6, 7
- [3] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Juergen Gall. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. *ICCV*, 2019. 1, 2, 5, 7
- [4] Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. Deep Clustering for Unsupervised Learning of Visual Features. *ECCV*, 2018. 1, 2
- [5] Ye Chen, Jinxian Liu, Bingbing Ni, Hang Wang, Jiancheng Yang, Ning Liu, Teng Li, Qi Tian, Huawei Hisilicon, Huawei Car Bu, and Huawei Cloud. Shape Self-Correction for Unsupervised Point Cloud Understanding. *ICCV*, 2021. 2
- [6] Julian Chibane, Francis Engelmann, Tuan Anh Tran, and Gerard Pons-Moll. Box2Mask: Weakly Supervised 3D Semantic Instance Segmentation Using Bounding Boxes. *ECCV*, 2022. 2
- [7] Jang Hyun Cho, Utkarsh Mall, Kavita Bala, and Bharath Hariharan. PiCIE: Unsupervised Semantic Segmentation using Invariance and Equivariance in Clustering. *CVPR*, 2021. 1, 2, 5, 6, 7
- [8] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. *CVPR*, 2019. 2
- [9] Tiago Cortinhal, George Tzelepis, and Eren Erdal Aksoy. SalsaNext: Fast Semantic Segmentation of LiDAR Point Clouds for Autonomous Driving. *ISVC*, 2020. 2
- [10] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D Reconstructions of Indoor Scenes. *CVPR*, 2017. 2, 5, 6
- [11] Wouter Van Gansbeke, Simon Vandenhende, and Stamatios Georgoulis. Unsupervised Semantic Segmentation by Contrasting Object Mask Proposals. *ICCV*, 2021. 2
- [12] Benjamin Graham, Martin Engelcke, and Laurens van der Maaten. 3D Semantic Segmentation with Submanifold Sparse Convolutional Networks. *CVPR*, 2018. 1, 2, 3, 5, 6, 7
- [13] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R. Martin, and Shi-Min Hu. PCT: Point Cloud Transformer. *Computational Visual Media*, 2021. 2
- [14] Yulan Guo, Hanyun Wang, Qingyong Hu, Hao Liu, Li Liu, and Mohammed Bennis. Deep Learning for 3D Point Clouds: A Survey. *TPAMI*, 2020. 1
- [15] Huy Ha and Shuran Song. Semantic Abstraction: Open-World 3D Scene Understanding from 2D Vision-Language Models. *CoRL*, 2022. 2
- [16] Timo Hackel, N. Savinov, L. Ladicky, Jan D. Wegner, K. Schindler, and M. Pollefeys. SEMANTIC3D.NET: A New Large-scale Point Cloud Classification Benchmark. *ISPRS*, 2017. 2
- [17] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. *CVPR*, 2021. 1, 2, 8
- [18] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring Data-Efficient 3D Scene Understanding with Contrastive Scene Contexts. *CVPR*, 2021. 2
- [19] Qingyong Hu, Bo Yang, Guangchi Fang, Yulan Guo, Ales Leonardis, Niki Trigoni, and Andrew Markham. SQN: Weakly-Supervised Semantic Segmentation of Large-Scale 3D Point Clouds. *ECCV*, 2022. 1, 2
- [20] Qingyong Hu, Bo Yang, Sheikh Khalid, Wen Xiao, Niki Trigoni, and Andrew Markham. Towards Semantic Segmentation of Urban-Scale 3D Point Clouds: A Dataset, Benchmarks and Challenges. *CVPR*, 2021. 1, 2
- [21] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. RandLA-Net: Efficient Semantic Segmentation of Large-Scale Point Clouds. *CVPR*, 2020. 1, 2
- [22] Zeyu Hu, Hongbo Fu, and Chiew-lan Tai. LiDAL: Inter-frame Uncertainty Based Active Learning for 3D LiDAR Semantic Segmentation. *ECCV*, 2022. 1
- [23] Siyuan Huang, Yichen Xie, Song-Chun Zhu, and Yixin Zhu. Spatio-temporal Self-Supervised Representation Learning for 3D Point Clouds. *ICCV*, 2021. 2
- [24] Xu Ji, Andrea Vedaldi, and Joao Henriques. Invariant information clustering for unsupervised image classification and segmentation. *ICCV*, 2019. 1, 2, 5, 6, 7
- [25] Abhijit Kundu, Xiaoqi Yin, Alireza Fathi, David Ross, Brian Brewington, Thomas Funkhouser, and Caroline Pantofaru. Virtual multi-view fusion for 3D semantic segmentation. *ECCV*, 2020. 2
- [26] Huan Lei, Naveed Akhtar, and Ajmal Mian. Octree guided CNN with Spherical Kernels for 3D Point Clouds. *CVPR*, 2019. 2
- [27] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. PointCNN: Convolution On X-Transformed Points. *NIPS*, 2018. 1, 2, 6
- [28] Liqiang Lin, Yilin Liu, Yue Hu, Xingguang Yan, Ke Xie, and Hui Huang. Capturing, Reconstructing, and Simulating: the UrbanScene3D Dataset. *ECCV*, 2022. 2
- [29] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: Common Objects in Context. *ECCV*, 2014. 1
- [30] Kangcheng Liu, Yuzhi Zhao, Qiang Nie, Zhi Gao, and Ben M. Chen. Weakly Supervised 3D Scene Parsing with Region-Level Boundary Awareness and Instance Discrimination. *ECCV*, 2022. 2
- [31] Minghua Liu, Yin Zhou, Charles R Qi, Boqing Gong, Hao Su, and Dragomir Anguelov. LESS: Label-Efficient Semantic Segmentation for LiDAR Point Clouds. *ECCV*, 2022. 2
- [32] Yunze Liu, Li Yi, Shanghang Zhang, Qingnan Fan, Thomas Funkhouser, and Hao Dong. P4Contrast: Contrastive Learning with Pairs of Point-Pixel Pairs for RGB-D Scene Understanding. *arXiv:2012.13089*, 2020. 2
- [33] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. One Thing One Click: A Self-Training Approach for Weakly Supervised 3D Semantic Segmentation. *CVPR*, 2021. 2

- [34] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-Voxel CNN for Efficient 3D Deep Learning. *NeurIPS*, 2019. 2
- [35] Hsien-Yu Meng, Gao Lin, Yu-Kun Lai, and Dinesh Manocha. VV-Net: Voxel VAE Net with Group Convolutions for Point Cloud Segmentation. *ICCV*, 2019. 2
- [36] Andres Milioto, Ignacio Vizzo, Jens Behley, and Cyrill Stachniss. RangeNet++: Fast and Accurate LiDAR Semantic Segmentation. *IROS*, 2019. 2
- [37] Yassine Ouali, Céline Hudelot, and Myriam Tami. Autoregressive Unsupervised Image Segmentation. *ECCV*, 2020. 2
- [38] Yatian Pang, Wenxiao Wang, Francis E. H. Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked Autoencoders for Point Cloud Self-supervised Learning. *ECCV*, 2022. 2
- [39] Jeremie Papon, Alexey Abramov, Markus Schoeler, and Florentin Wörgötter. Voxel Cloud Connectivity Segmentation - Supervoxels for Point Clouds. *CVPR*, 2013. 3, 5
- [40] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. *CVPR*, 2017. 1, 2, 5, 6, 7
- [41] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space. *NIPS*, 2017. 1, 2, 5, 6, 7
- [42] Yongming Rao, Jiwen Lu, and Jie Zhou. Global-Local Bidirectional Reasoning for Unsupervised Representation Learning of 3D Point Clouds. *CVPR*, 2020. 2
- [43] Zhongzheng Ren, Ishan Misra, Alexander G. Schwing, and Rohit Girdhar. 3D Spatial Recognition without Spatially Labeled 3D. *CVPR*, 2021. 2
- [44] Damien Robert, Bruno Vallet, and Loic Landrieu. Learning Multi-View Aggregation In the Wild for Large-Scale 3D Semantic Segmentation. *CVPR*, 2022. 2
- [45] David Rozenberszki, Or Litany, and Angela Dai. Language-Grounded Indoor 3D Semantic Segmentation in the Wild. *ECCV*, 2022. 2
- [46] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. 1
- [47] Radu Bogdan Rusu, Nico Blodow, Zoltan Csaba Marton, and Michael Beetz. Aligning Point Cloud Views using Persistent Feature Histograms. *IROS*, 2008. 4, 5
- [48] Jonathan Sauder and Bjarne Sievers. Self-Supervised Deep Learning on Point Clouds by Reconstructing Space. *NeurIPS*, 2019. 2
- [49] Hanyu Shi, Jiacheng Wei, Ruibo Li, Fayao Liu, and Guosheng Lin. Weakly Supervised Segmentation on Outdoor 4D point clouds with Temporal Matching and Spatial Graph Propagation. *CVPR*, 2022. 2
- [50] Inkyu Shin, Yi-Hsuan Tsai, Bingbing Zhuang, Samuel Schulter, Buyu Liu, Sparsh Garg, In So Kweon, and Kuk-Jin Yoon. MM-TTA: Multi-Modal Test-Time Adaptation for 3D Semantic Segmentation. *CVPR*, 2022. 2
- [51] Weiwei Sun, Andrea Tagliasacchi, Boyang Deng, Sara Sabour, Soroosh Yazdani, Geoffrey Hinton, and Kwang Moo Yi. Canonical Capsules: Unsupervised Capsules in Canonical Pose. *NeurIPS*, 2021. 2
- [52] Weikai Tan, Nannan Qin, Lingfei Ma, Ying Li, Jing Du, Guorong Cai, Ke Yang, and Jonathan Li. Toronto-3D: A Large-scale Mobile LiDAR Dataset for Semantic Segmentation of Urban Roadways. *CVPR Workshops*, 2020. 2
- [53] An Tao, Yueqi Duan, Yi Wei, Jiwen Lu, and Jie Zhou. Seg-Group: Seg-level supervision for 3D instance and semantic segmentation. *arXiv:2012.10217*, 2020. 2
- [54] Ali Thabet, Humam Alwassel, and Bernard Ghanem. Self-supervised Learning of Local Features in 3D Point Clouds. *CVPR Workshops*, 2019. 2
- [55] Hugues Thomas, Charles R. Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J. Guibas. KPConv: Flexible and Deformable Convolution for Point Clouds. *ICCV*, 2019. 1, 2
- [56] Ozan Unal, Dengxin Dai, and Luc Van Gool. Scribble-Supervised LiDAR Semantic Segmentation. *CVPR*, 2022. 2
- [57] Nina Varney, Vijayan K. Asari, and Quinn Graehling. DALES: A Large-scale Aerial LiDAR Data Set for Semantic Segmentation. *CVPR Workshops*, 2020. 2
- [58] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matthew J. Kusner. Unsupervised Point Cloud Pre-Training via Occlusion Completion. *ICCV*, 2021. 2
- [59] Haiyan Wang, Xuejian Rong, Liang Yang, Shuihua Wang, and Yingli Tian. Towards Weakly Supervised Semantic Segmentation in 3D Graph-Structured Point Clouds of Wild Scenes. *BMVC*, 2019. 1
- [60] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic Graph CNN for Learning on Point Clouds. *TOG*, 2019. 1, 6
- [61] 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. *CVPR*, 2020. 2
- [62] Bichen Wu, Alvin Wan, Xiangyu Yue, and Kurt Keutzer. SqueezeSeg: Convolutional Neural Nets with Recurrent CRF for Real-Time Road-Object Segmentation from 3D LiDAR Point Cloud. *ICRA*, 2018. 2
- [63] Tsung-Han Wu, Yueh-Cheng Liu, Yu-Kai Huang, Hsin-Ying Lee, Hung-Ting Su, Ping-Chia Huang, and Winston H. Hsu. ReDAL: Region-based and Diversity-aware Active Learning for Point Cloud Semantic Segmentation. *ICCV*, 2021. 1
- [64] Wenxuan Wu, Zhongang Qi, and Li Fuxin. PointConv: Deep Convolutional Networks on 3D Point Clouds. *CVPR*, 2019. 2
- [65] Zhonghua Wu, Yicheng Wu, Guosheng Lin, Jianfei Cai, and Chen Qian. Dual Adaptive Transformations for Weakly Supervised Point Cloud Segmentation. *ECCV*, 2022. 2
- [66] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. *ECCV*, 2020. 1, 2, 8
- [67] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. PointContrast: Unsupervised Pre-training for 3D Point Cloud Understanding. *ECCV*, 2020. 2

- [68] 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. *ECCV*, 2022. [2](#)
- [69] Xun Xu and Gim Hee Lee. Weakly Supervised Semantic Point Cloud Segmentation:Towards 10X Fewer Labels. *CVPR*, 2020. [1](#), [2](#)
- [70] Xu Yan, Jiantao Gao, Chaoda Zheng, Chao Zheng, Ruimao Zhang, Shenghui Cui, and Zhen Li. 2DPASS: 2D Priors Assisted Semantic Segmentation on LiDAR Point Clouds. *ECCV*, 2022. [2](#)
- [71] Zhaoyuan Yin, Pichao Wang, Fan Wang, Xianzhe Xu, Hanling Zhang, Hao Li, and Rong Jin. TransFGU: A Top-down Approach to Fine-Grained Unsupervised Semantic Segmentation. *ECCV*, 2022. [2](#)
- [72] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. PointCLIP: Point Cloud Understanding by CLIP. *CVPR*, 2022. [2](#)
- [73] Yachao Zhang, Yanyun Qu, Yuan Xie, Zonghao Li, Shanshan Zheng, and Cuihua Li. Perturbed Self-Distillation: Weakly Supervised Large-Scale Point Cloud Semantic Segmentation. *ICCV*, 2021. [2](#)
- [74] Zaiwei Zhang, Min Bai, and Erran Li. Self-Supervised Pre-training for Large-Scale Point Clouds. *NeurIPS*, 2022. [2](#)
- [75] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-Supervised Pretraining of 3D Features on any Point-Cloud. *ICCV*, 2021. [2](#)
- [76] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip Torr, and Vladlen Koltun. Point Transformer. *ICCV*, 2021. [2](#)
- [77] Xiaoyu Zhu, Jeffrey Chen, Xiangrui Zeng, Junwei Liang, Chengqi Li, Sinuo Liu, Sima Behpour, and Min Xu. Weakly Supervised 3D Semantic Segmentation Using Cross-Image Consensus and Inter-Voxel Affinity Relations. *ICCV*, 2021. [1](#)
- [78] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and Asymmetrical 3D Convolution Networks for LiDAR Segmentation. *CVPR*, 2021. [2](#)
- [79] Zhuangwei Zhuang, Rong Li, Kui Jia, Qicheng Wang, Yuanqing Li, and Mingkui Tan. Perception-Aware Multi-Sensor Fusion for 3D LiDAR Semantic Segmentation. *ICCV*, 2021. [2](#)