

# P-SLCR: Unsupervised Point Cloud Semantic Segmentation via Prototypes Structure Learning and Consistent Reasoning

Lixin Zhan<sup>1, 2</sup>, Jiang Jie<sup>1, 2\*</sup>, Tianjian Zhou<sup>1, 2</sup>, Yukun Du<sup>1</sup>, Yan Zheng<sup>1</sup>, Xuehu Duan<sup>1, 2</sup>

<sup>1</sup>College of Systems Engineering, National University of Defense Technology

<sup>2</sup>Laboratory for Big Data and Decision, National University of Defense Technology

zhanlixin98@outlook.com, jiejiang@nudt.edu.cn, zhoutianjian2000@nudt.edu.cn, duyukun-nudt@outlook.com, zhengyan24@nudt.edu.cn, duanxuehu24@nudt.edu.cn

## Abstract

Current semantic segmentation approaches for point cloud scenes heavily rely on manual labeling, while research on unsupervised semantic segmentation methods specifically for raw point clouds is still in its early stages. Unsupervised point cloud learning poses significant challenges due to the absence of annotation information and the lack of pre-training. The development of effective strategies is crucial in this context. In this paper, we propose a novel prototype library-driven unsupervised point cloud semantic segmentation strategy that utilizes Structure Learning and Consistent Reasoning (P-SLCR). First, we propose a Consistent Structure Learning to establish structural feature learning between consistent points and the library of consistent prototypes by selecting high-quality features. Second, we propose a Semantic Relation Consistent Reasoning that constructs a prototype inter-relation matrix between consistent and ambiguous prototype libraries separately. This process ensures the preservation of semantic consistency by imposing constraints on consistent and ambiguous prototype libraries through the prototype inter-relation matrix. Finally, our method was extensively evaluated on the S3DIS, SemanticKITTI, and ScanNet datasets, achieving the best performance compared to unsupervised methods. Specifically, the mIoU of 47.1% is achieved for Area-5 of the S3DIS dataset, surpassing the classical fully supervised method PointNet by 2.5%.

**Code** — <https://github.com/lixinzhan98/P-SLCR>

## Introduction

Semantic segmentation of point clouds is a crucial task in computer 3D vision, representing fundamental research with growing significance across various applications in recent years. Since the inception of PointNet (Qi et al. 2017a), there has been a surge in neural network-based semantic segmentation research for point clouds (Hu et al. 2020; Zhao et al. 2021a; Zhan, Li, and Min 2023; Wu et al. 2024), leading to significant enhancements in the accuracy of point segmentation (Huang et al. 2024; Li et al. 2023). However, these methodologies heavily rely on manual annotation, with unstructured 3D data annotation typically requiring greater human and material resources than annotating 2D images. Cer-

\*Corresponding author.

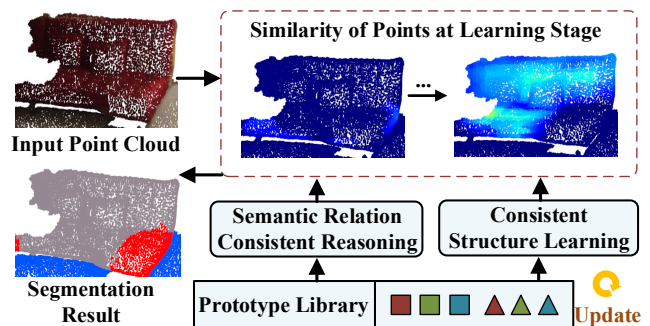


Figure 1: P-SLCR architecture for structural learning and consistent reasoning through a learnable prototype-based library. The consistent points are continuously expanding and increasingly resembling the prototype.

tain approaches alleviate this issue by utilizing a limited set of point cloud labels (Liu et al. 2022b; Zhang et al. 2021a; Shi et al. 2022; Unal, Dai, and Van Gool 2022; Hu et al. 2022; Zhan et al. 2025a; Li et al. 2022; Zhan et al. 2025b) during model training. Despite achieving promising outcomes, these methods still depend on manual annotations for data alignment. Extending to new scenarios is challenging due to the need for manual labeling, a highly time-consuming process. While unsupervised methods for 2D images (Ji, Henriques, and Vedaldi 2019; Cho et al. 2021) are booming, unsupervised understanding of 3D data is just beginning. Some researchers have also achieved unsupervised learning of 3D data through domain adaptive strategies (Zhao et al. 2021b; Saltori et al. 2023), but our study does not rely on any transfer learning.

Our study aims to delineate 3D features without the manual annotation. While numerous methods have emerged to tackle unsupervised point cloud semantic segmentation tasks, representative approaches include PointDC (Chen et al. 2023), GrowSP (Zhang et al. 2023), U3DS<sup>3</sup> (Liu et al. 2024) and LogoSP (Zhang et al. 2025). GrowSP posits that individual 3D points often lack practical significance, while a cluster or set of 3D points holds practical relevance. GrowSP (Zhang et al. 2023) and U3DS<sup>3</sup> (Liu et al. 2024) render the unsupervised task (Chen et al. 2023; Wu et al. 2023) of large-scale point cloud analysis feasible through

the application of over-segmentation for clustering, employing superpoint segmentation and superpoint region expansion. However, the pseudo-labels generated by the clustering algorithm are not completely trustworthy, and the direct use of the full set of pseudo-labels to supervise network learning may not be conducive to distinguishing salient features across categories. In this way, the prototype features are not representative and do not fully exploit the structural information of the point cloud.

To alleviate the above problems, we propose an unsupervised method for semantic segmentation of point clouds via Prototypes Consistent Structure Learning and Reasoning strategy. As shown in Fig. 1, we built two prototype libraries, the consistent prototype library, and the ambiguous prototype library. These libraries are updated by the EMA (Tarvainen and Valpola 2017) algorithm based on the local clustering centers of midpoint batches. Consistent features typically demonstrate higher accuracy than ambiguous features, facilitating the utilization of confidence in selecting high-quality features and establishing a structural similarity learning connection between consistent features and prototypes. The consistent prototype prioritizes learning robust features within each category and can guide the learning of ambiguous prototypes. Semantic features of both consistent and ambiguous prototypes are maintained consistently through constraints on the consistent prototype. The learning process is dynamic, with the potential for all ambiguous points to be gradually included in the consistent set through continuous learning, thus completing the feature space division and achieving precise semantic segmentation. Our contribution is as follows:

- We propose an innovative unsupervised framework for semantic segmentation of point clouds, emphasizing consistent structure learning and consistent reasoning guided by a dynamic prototype library.
- We propose a consistent structure learning that utilizes plausibility to select high-quality features and establish consistent structure learning between consistent point features and a library of consistent prototypes.
- We propose a semantic relation consistent reasoning designed to assist in training consistent prototypes, guiding the learning process for ambiguous prototypes. This constraint can be applied to uphold semantic feature consistency within both the consistent and ambiguous prototype libraries.

## Related Work

### Fully Supervised Semantic Segmentation

Fully supervised point cloud semantic segmentation has achieved great success in recent years. It can be categorized into view-based methods (Tatarchenko et al. 2018; Milioto et al. 2019; Kundu et al. 2020), voxel-based methods (Meng et al. 2019; Zhu et al. 2021), and point-based methods. The first two categories convert sparse data into regular data, inevitably losing structural information. Point-based methods, stemming from PointNet (Qi et al. 2017a) and PointNet++ (Qi et al. 2017b), advocate for the employment of

permutation invariant functions to extract point cloud features directly. Follow-up work RandLA-Net (Hu et al. 2020) balances efficiency and segmentation accuracy in handling large-scale semantic segmentation of point clouds. The advent of Transformer (Vaswani 2017) has spurred rapid advancements in Transformer-based approaches (Zhao et al. 2021a; Wu et al. 2022, 2024) within point cloud semantic segmentation, yielding promising outcomes. However, these methods rely on a large amount of annotated information, which is time-consuming and expensive.

### Weakly Supervised Semantic Segmentation

To reduce the burden of complete annotation, researchers have shifted their focus toward training semantic segmentation models with a limited number of annotations. Some of these methods use fewer 3D labels (Liu et al. 2022b; Shi et al. 2022; Unal, Dai, and Van Gool 2022; Hu et al. 2022; Li et al. 2022), and others use a small number of fully labeled scenes (Chibane et al. 2022; Liu et al. 2022a; Ren et al. 2021). To enrich annotation information, one may leverage consistency perturbation (Zhang et al. 2021b) or employ pseudo-label generation (Cheng et al. 2021). The consistency perturbation relies on ensuring that perturbed samples maintain consistency with the predictions made for the original sample. The methodology for generating pseudo-labels entails utilizing predicted values from unlabeled points to formulate labels that are integrated into the model for training. Some approaches employ self-supervised pre-training (Hou et al. 2021; Zhang, Bai, and Li 2022), and supervised signals can originate from diverse modalities, including 2D images (Robert, Vallet, and Landrieu 2022; Shin et al. 2022) or language models (Zhang et al. 2022; Rozenberszki, Litany, and Dai 2022). Nevertheless, these methods still need manual annotation for data alignment, preventing the network model from autonomously discerning semantic classes.

### Unsupervised Semantic Segmentation

Several works have proposed methods for 2D unsupervised semantic segmentation (Ji, Henriques, and Vedaldi 2019; Cho et al. 2021). Nevertheless, the most current methods are inapplicable to point clouds because of the distinctions between images and point clouds. Research has initiated exploration into unsupervised semantic segmentation of 3D point clouds, yielding promising results. GrowSP (Zhang et al. 2023) represents a pioneering effort in unsupervised semantic segmentation of point clouds. The approach involves initially segmenting superpoints using VCCS (Papon et al. 2013), synthesizing superpoint features, and subsequently clustering the superpoints while assigning pseudo-labels. During the later stages of training, larger superpoints can be synthesized. U3DS<sup>3</sup> (Liu et al. 2024) establishes a generalized unsupervised segmentation method by means of spatial clustering, iterative training by clustering cosine distances between the center of mass and points, and expanding the training data by color changes. However, none of them offer effective coaching strategies. Pseudo-labels generated by clustering algorithms are not entirely consistent, and violently using all of them to supervise network learning may

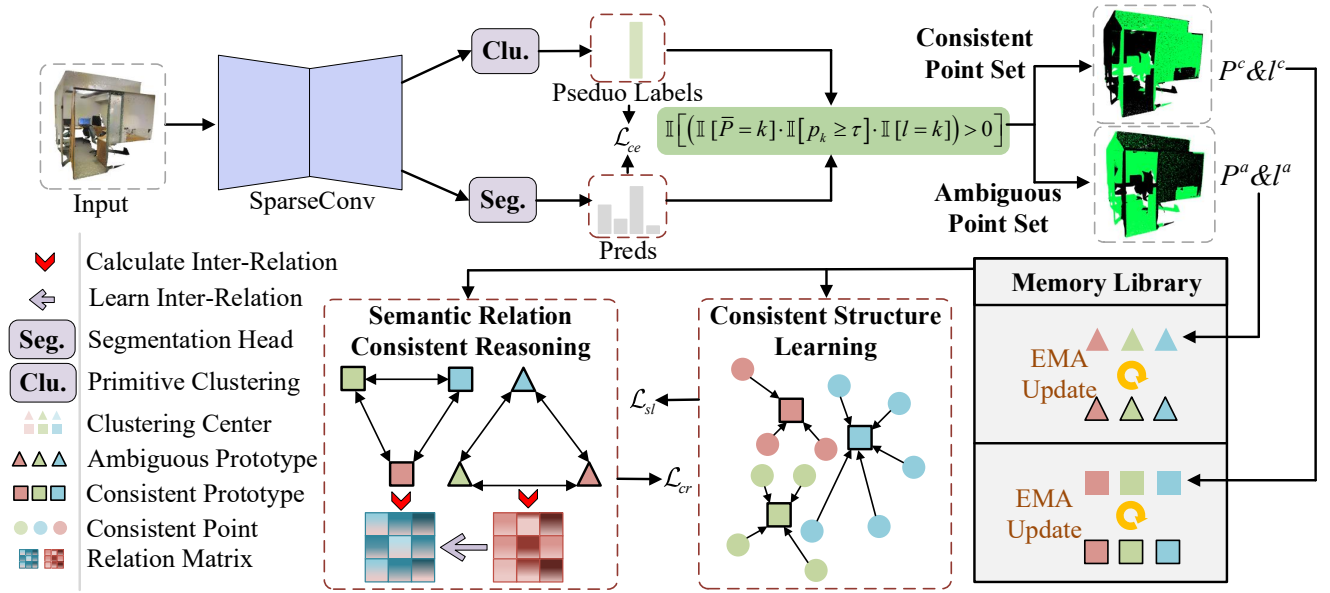


Figure 2: An overview of our P-SLCR with features extracted by SparseConv, pseudo-labels and predictions obtained by clustering and segmenting the header. The proposed reliability classifies the scene into consistent and ambiguous points and establishes an updatable prototype library to construct consistent structure learning and semantic relation consistent reasoning based on the prototype library.

be detrimental to distinguishing salient features of a category. The semantic prototype is not universal and does not fully exploit the structural information of the point cloud.

## P-SLCR

### Overview

Fig. 2 illustrates the overall framework of our proposed **P-SLCR** model, which performs unsupervised point cloud semantic segmentation through *prototype library-driven structure learning* and *semantic relation-consistent reasoning*. Given an input point cloud, a feature extractor  $G$  encodes geometric and color cues into point-wise embeddings, which are aggregated into semantically homogeneous superpoints. P-SLCR maintains a *dual prototype library* that dynamically separates reliable (consistent) and uncertain (ambiguous) representations according to feature reliability. The consistent prototype library captures stable semantics across scenes, while the ambiguous one serves as a buffer to model uncertain regions and guide their refinement. A structure-consistent reasoning module enforces semantic and relational alignment between the two prototype sets, encouraging globally coherent segmentation. Through iterative optimization, refined pseudo-labels generated from the prototype reasoning loop are fed back to supervise  $G$ , forming a closed learning cycle that progressively enhances both feature quality and segmentation consistency.

### Separation of Reliable Points

Without loss of generality, for points in a point cloud scene, can obtain the features  $f$  obtained after the network  $G$ , the pseudo-labels  $l$ , and the initial prototypes  $\mu$ , which are usu-

ally the clustering centers of the categories. It is reasonable to assume (Appendix for details.) that if the network  $G$  is adequately robust and the prototype  $\mu$  is also robust. Theoretically, the local clustering centers predicted by the network  $G$  for each batch are expected to be infinitesimally close to the corresponding prototype  $\mu$ , reflecting local consistency during training.

In reality,  $G$  is not a highly efficient learner, and  $\mu$  lacks the required robustness. This approach can be executed as a training strategy for the model.

However, it is not desirable to directly and crudely calculate the clustering center in a single batch based on pseudo-labels. The pseudo-label  $l$  assigned by the clustering may lack credibility and contain numerous errors. Therefore, we introduce the concept of consistent points, which mainly filters out points with high confidence. Formally, the consistent points should be satisfied that the clustering assigned pseudo-label  $l$  agrees with the network  $G$  prediction  $\bar{p}$ , where  $\bar{p} = \text{argmax}(\mathbf{p})$ . If  $p_c$  exceeds the threshold  $\tau$ , the consistent set is defined as illustrated in Formula 1:

$$R = \mathbf{1} \left[ \sum_{k=1}^K (\mathbf{1}[\bar{p} = k] \cdot \mathbf{1}[p_k \geq \tau] \cdot \mathbf{1}[l = k]) \right], \quad (1)$$

where  $\mathbf{1}$  is the indicator function,  $K$  denotes the total number of categories, and  $\tau$  is the confidence threshold. The generated  $R$  is a binary mask.

The set of consistent point clouds  $P^c$  and their respective pseudo-labels  $l^c$  are filtered by the mask  $R$ . Likewise, the ambiguous point cloud set  $P^a$  and its associated pseudo-label  $l^a$  are determined by the complement of the mask

$(1 - R)$ , as shown in Formulas 2 and 3:

$$P^c = R \cdot P, \quad P^a = (1 - R) \cdot P, \quad (2)$$

$$l^c = R \cdot l, \quad l^a = (1 - R) \cdot l. \quad (3)$$

## Library of Prototypes

To aid in learning prototype structures, we established two prototype memory banks, including a consistent prototype library and an ambiguous prototype library. The Exponential Moving Average (EMA) is employed to calculate the clustering centers when merging various batches, and then get the consistent and ambiguous prototypes driven by receiving multiple batches of samples. Initially, the consistent and ambiguous prototype memory banks are the clustering centers computed during clustering, i.e.,  $\mu^c = \mu^a = \mu$ . The specific update calculation for the prototype libraries are shown in Formulas 4 and 5:

$$\mu_k^c \leftarrow \alpha \mu_k^c + (1 - \alpha) \bar{\mu}_k^c, \quad (4)$$

$$\mu_k^a \leftarrow \alpha \mu_k^a + (1 - \alpha) \bar{\mu}_k^a, \quad (5)$$

where  $\bar{\mu}_k^c$  and  $\bar{\mu}_k^a$  represent the category  $k$  clustering center feature representation in a single batch of samples.  $\alpha$  is the EMA parameter, with a default setting of 0.99.

Enhancing the robust features in the prototype library involves calculating category centers for each batch and updating the library consistently. For a single batch, the consistent category clustering center  $\bar{\mu}_k^c$  and the ambiguous category clustering center  $\bar{\mu}_k^a$  can be defined as Formulas 6 and 7:

$$\bar{\mu}_k^{c(i)} = \frac{1}{|\Omega_k^{c(i)}|} \sum_{\mathbf{p}_j^c \in \Omega_k^{c(i)}} G(\mathbf{p}_j^c), \quad (6)$$

$$\bar{\mu}_k^{a(i)} = \frac{1}{|\Omega_k^{a(i)}|} \sum_{\mathbf{p}_j^a \in \Omega_k^{a(i)}} G(\mathbf{p}_j^a), \quad (7)$$

where  $\Omega_k^{c(i)}$  and  $\Omega_k^{a(i)}$  are the set of points pseudo-labeled  $k$  of  $i$  batch for consistent and ambiguous points respectively. This can be expressed as  $\Omega_k^{c(i)} = \{\mathbf{p}_j^c \mid l_j^c = k\}$  and  $\Omega_k^{a(i)} = \{\mathbf{p}_j^a \mid l_j^a = k\}$ . Where  $i = \{1, \dots, BN_L\}$ ,  $|\cdot|$  denotes the number of sets, and  $BN_L$  denotes the number of batch.

## Consistent Structure Learning

To leverage the potential of valid points, a structural error matrix is formulated by prototypes from a consistent prototype library with the consistent point feature of the Clustering centre for corresponding categories within a given batch. To learn consistent structures, the exact definition is depicted in Formula 8:

$$M_k^c = \|\mu_k^c - G(\mathbf{p}_j^c)\|_2, \quad (8)$$

where  $M_k^c$  represents the structural error between the prototype of category  $k$  in the consistent prototype library and the consistent feature within a single batch, with  $\|\cdot\|_2$  denoting the Euclidean distance.

Theoretically, minimizing the structural error facilitates the learning of consistent structures. Hence, the dependable structural loss can be represented by Formula 9:

$$\mathcal{L}_{sl} = \sum_{k=1}^K M_k^c. \quad (9)$$

Structure learning of consistent points is necessary, and the visualisation results in Fig. 7 validate our strategy. Consistent structure learning is equivalent to a reduction in the distance of a specific category from the prototype and a brightening of the colour.

## Semantic Relation Consistent Reasoning

To constrain consistent and ambiguous points to have a consistent semantic representation. An artificial assumption is made that prototypical features acquired at consistent points are more precise than those obtained at ambiguous points.

Therefore, we introduce consistent and ambiguous prototypes from the consistent reasoning constraints prototype library. To maintain consistency in the similarity between consistent and ambiguous prototypes, distinct consistent and ambiguous similarity matrices are computed separately, as demonstrated in Formula 10:

$$e_{ij}^c = \mu_i^c \cdot \mu_j^{cT}, \quad e_{ij}^a = \mu_i^a \cdot \mu_j^{aT}, \quad i, j = \{1, 2, \dots, K\}. \quad (10)$$

The similarity matrix can be normalized using Formula 11:

$$\bar{e}_{ij}^c = e_{ij}^c / \sum_{j=1}^K e_{ij}^c, \quad \bar{e}_{ij}^a = \log(e_{ij}^a / \sum_{j=1}^K e_{ij}^a). \quad (11)$$

The relationship between consistent and ambiguous prototypes is often in error. For consistency preservation, it is possible to keep the information entropy between them relatively close. Hence, the consistent reasoning loss function is defined by Formula 12:

$$\mathcal{L}_{cr} = \frac{1}{K^2} \sum_{i=1}^K \sum_{j=1}^K \bar{e}_{ij}^c \log(\bar{e}_{ij}^c / \bar{e}_{ij}^a). \quad (12)$$

Consistent reasoning extends the range of consistent points and enhances the robustness of the consistent prototype. In the visualisation results in Fig. 7, as the reasoning continues, it turns out that the more ambiguous points that should have belonged to the category are classified as consistent points.

## Overall Objective Function

In summary, the overall loss function can be formulated as Formula 13:

$$\mathcal{L}_{total} = \mathcal{L}_{ce} + \lambda_1 \mathcal{L}_{sl} + \lambda_2 \mathcal{L}_{cr}, \quad (13)$$

where  $\lambda_1$  and  $\lambda_2$  are hyperparameters.

	Method	OA(%)	mAcc(%)	mIoU(%)
Supervised Methods	PointNet	77.5	59.1	44.6
	PointNet++	77.5	62.6	50.1
	SparseConv	88.4	69.2	60.8
semi-supervised (10%)	Jiang et al.	69.1	-	57.7
weakly supervised (1pt)	MT	-	-	44.4
	Zhang et al.	-	-	48.2
Unsupervised Methods	KMeans	22.1	21.2	9.4
	DBSCAN	17.5	19.8	9.2
	IIC-S-PFH	31.2	16.3	9.1
	PICIE-S-PFH	48.4	40.4	25.2
	GrowSP	78.4	<b>57.2</b>	44.5
	U3DS <sup>3</sup>	75.5	55.8	42.8
	PointDC	54.1	-	22.6
	P-SLCR (ours)	<b>79.7</b>	<b>57.2</b>	<b>47.1</b>

Table 1: Quantitative results for semantic segmentation on the S3DIS dataset. Only 12 categories evaluated on the Area-5.

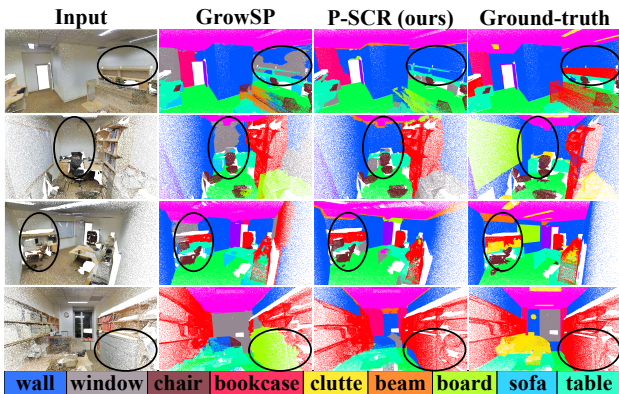


Figure 3: The qualitative results on S3DIS demonstrate the segmentation performance in comparison to GrowSP. Each color represents a semantic class. We highlight the differences with black circles.

## Implementation Details

The input point cloud generally comprises 6D raw features with normalized coordinates XYZ and color RGB. SemanticKITTI (Behley et al. 2019) lacks color information. A straightforward and efficient neural network framework is implemented based on the SparseConv (Graham, Engelcke, and Van Der Maaten 2018) architecture, without any pre-training, utilizing the KMeans algorithm for clustering hyperpoint features. Color enhancements (Liu et al. 2023) were applied, involving color shifting, contrast enhancement, and dithering techniques. The number of semantic primitives,  $C$ , must be defined for training purposes.  $\lambda_1$  and  $\lambda_2$  are 0 for the first half of the training phase and set to 1 for the second half. During testing, semantic categories of these primitives are clustered using the KMeans (Lloyd 1982) algorithm. Given the mismatch between the test labels and predicted classes, the Hungarian algorithm aligns the predicted classes with the true labels.

	Method	OA(%)	mAcc(%)	mIoU(%)
Supervised Methods	PointNet	75.9	67.1	49.4
	PointNet++	77.1	74.1	55.1
	SparseConv	89.4	78.1	69.2
Unsupervised Methods	GrowSP	76.0	59.4	44.6
	P-SLCR (ours)	<b>79.0</b>	<b>62.2</b>	<b>47.5</b>

Table 2: Quantitative results for semantic segmentation on the S3DIS 6-Fold. Only 12 categories evaluated.

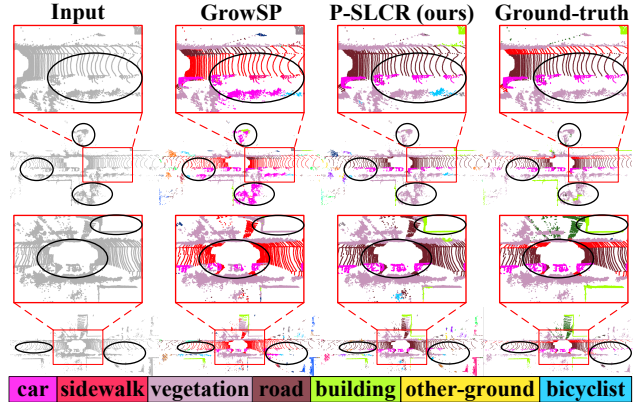


Figure 4: The qualitative results on SemanticKITTI demonstrate the segmentation performance in comparison to GrowSP. We zoomed in on local regions to show details and highlighted differences with black circles.

## Experiments

The effectiveness of our method was evaluated by testing it on three comprehensive point cloud datasets, encompassing both indoor (Armeni et al. 2016; Dai et al. 2017) and outdoor (Behley et al. 2019) scenarios. The overall accuracy (oAcc), the mean accuracy (mAcc), and the mean Intersection over Union (mIoU), which are commonly used in point cloud semantic segmentation methods, are mainly used as evaluation metrics. The comparison methods include PointNet (Qi et al. 2017a), PointNet++ (Qi et al. 2017b), SparseConv (Graham, Engelcke, and Van Der Maaten 2018), Zhang et al. (Zhang et al. 2021b), MT (Tarvainen and Valpola 2017), Jiang et al. (Jiang et al. 2021), KMeans (Lloyd 1982), DBSCAN (Ester et al. 1996), IIC-S-PFH (Ji, Henriques, and Vedaldi 2019), PICIE-S-PFH (Cho et al. 2021), PointDC (Chen et al. 2023), GrowSP (Zhang et al. 2023) and U3DS<sup>3</sup> (Liu et al. 2024).

### Evaluation on S3DIS

S3DIS (Armeni et al. 2016) is a large-scale indoor point cloud dataset divided into 6 areas, labeling the point cloud into 13 classes with one clutter class. For comparison with previous methods, we remove clutter classes during testing. It is also compared with fully supervised methods, semi-supervised methods, weakly supervised methods and the latest unsupervised methods.

The results for S3DIS Area-5 are shown in Table 1, our method (P-SLCR) achieves the best performance for all metrics in the unsupervised method comparison. Compared to

	Method	OA(%)	mAcc(%)	mIoU(%)
	KMeans	8.2	8.1	2.5
	DBSCAN	17.8	7.5	6.8
Unsupervised Methods	IIC-S-PFH	23.4	9.0	4.6
	PICIE-S-PFH	42.7	11.5	6.8
	GrowSP	38.3	19.7	13.2
	U3DS <sup>3</sup>	34.8	<b>23.1</b>	14.2
	P-SLCR (ours)	<b>55.9</b>	21.1	<b>15.3</b>

Table 3: Quantitative results for semantic segmentation on the SemanticKITTI dataset. All 19 categories evaluated on the validation set.

	Method	mIoU(%)
	PointNet	14.6
Supervised Methods	PointNet++	20.1
	SparseConv	53.2
Unsupervised Methods	GrowSP	14.3
	P-SLCR (ours)	<b>15.9</b>

Table 4: Quantitative results for semantic segmentation on the SemanticKITTI dataset. All 19 categories evaluated on the online test set.

the next best unsupervised method, GrowSP, P-SLCR outperforms with a 1.3% higher OA and a 2.6% higher mIoU. P-SLCR remains competitive even when compared with weakly supervised methods. In addition, it is worth noting that P-SLCR outperforms the classical fully supervised network PointNet for the first time and by 2.5% in mIoU. The results for S3DIS 6-Fold are shown in Table 2, P-SLCR is substantially ahead of GrowSP in all experimental metrics.

Fig. 3 presents qualitative results obtained by the state-of-the-art unsupervised method GrowSP and our proposed P-SLCR. GrowSP tends to misclassify walls and bookcases as windows or boards, whereas P-SLCR produces more accurate and semantically consistent segmentation results.

### Evaluation on SemanticKITTI

SemanticKITTI (Behley et al. 2019) is a large outdoor driving dataset collected using laser LIDAR. In total, there are 22 sequences and 19 semantic categories. Comparisons are made primarily with unsupervised methods that have migrated from other fields, as well as with more recent unsupervised point cloud methods.

As shown in Table 3, our method (P-SLCR) obtains the optimal OA and mIoU. It is 0.9% ahead of the U3DS<sup>3</sup> (Liu et al. 2024), in terms of mIoU. Notably, it is substantially ahead of other unsupervised methods in OA, about 20% ahead of the state-of-the-art unsupervised point cloud methods. The results for SemanticKITTI online test are shown in Table 4, P-SLCR outperforms GrowSP by 1.6% on mIoU.

Fig. 4 presents the visual outcomes obtained through the latest unsupervised method, GrowSP (Zhang et al. 2023). GrowSP tends to misclassify extensive road segments as sidewalk, a discrepancy effectively addressed by P-SLCR. Moreover, GrowSP struggles with distinguishing between cars and vegetation, leading to semantic confusion. It fre-

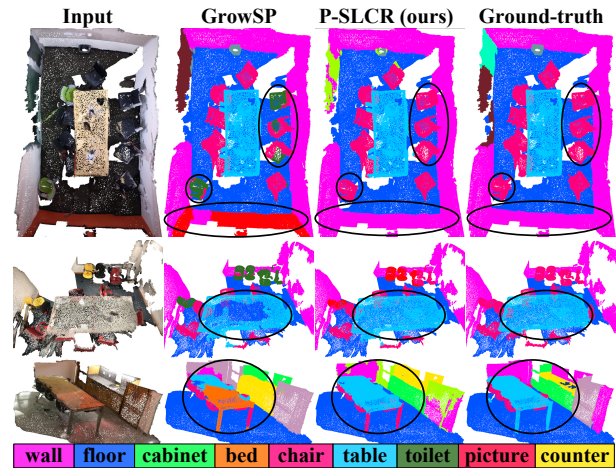


Figure 5: The qualitative results on ScanNet demonstrate the segmentation performance in comparison to GrowSP. Each color represents a semantic class. We highlight the differences with black circles.

	Method	OA(%)	mAcc(%)	mIoU(%)
	KMeans	10.2	10.4	3.4
	DBSCAN	15.3	10.1	6.1
Unsupervised Methods	IIC-S-PFH	18.9	6.3	3.0
	PICIE-S-PFH	23.6	15.1	7.4
	GrowSP	57.3	44.2	25.4
	U3DS <sup>3</sup>	60.1	46.8	27.3
	P-SLCR (ours)	<b>61.4</b>	<b>49.0</b>	<b>29.0</b>

Table 5: Quantitative results for semantic segmentation on the ScanNet dataset. All 20 categories evaluated on the validation set.

quently separates cars and vegetation erroneously, a challenge absent in P-SLCR, underscoring the robustness of P-SLCR.

### Evaluation on ScanNet

ScanNet (Dai et al. 2017) is an indoor dataset in RGB-D captured from real scenes, with 41 categories in the training set and 20 categories in the validation and test sets. The main comparison is with state-of-the-art unsupervised point cloud semantic segmentation methods.

As shown in Table 5, our method (P-SLCR) achieved the best performance in all metrics, with 1.3%, 2.2%, and 1.7% improvement in OA, mAcc, and mIoU, respectively, over the next best method (Liu et al. 2024).

Fig. 5 displays that P-SLCR accurately segments the categories of chairs, tables, and walls. GrowSP often incorrectly divides objects, especially smaller categories like chairs and tables, resulting in a single class being split into multiple categories.

### Ablation Study

A series of ablation studies were conducted on the S3DIS (Armeni et al. 2016) dataset Area-5 to assess the efficacy

	OA(%)	mAcc(%)	mIoU(%)
(1) Remove Reliable Structure Learning	79.36	55.63	45.65
(2) Remove Consistent Reasoning	79.08	54.09	42.37
(3) Remove Color Enhancement	79.28	57.50	45.21
(4) $\tau = 0.6$ for Degree of Reliability	79.19	57.08	43.80
(5) $\tau = 0.7$ for Degree of Reliability	79.72	57.23	<b>47.08</b>
(6) $\tau = 0.8$ for Degree of Reliability	78.88	56.76	44.50
(7) Number of Semantic Primitive = 200	78.81	<b>58.45</b>	44.25
(8) Number of Semantic Primitive = 300	79.72	57.23	<b>47.08</b>
(9) Number of Semantic Primitive = 400	78.01	58.36	45.27
(10) $\mathcal{L}_{cr}$ Change to using Cross-Entropy	<b>80.27</b>	56.63	44.72

Table 6: The OA, mAcc and mIoU scores of all ablated networks on Area-5 of S3DIS based on our full P-SLCR.




Figure 6: Similarity measure matrix visualization results for consistent and ambiguous prototypes. For presentation purposes, we clustered the number of prototypes into categories of visualization results for ScanNet validation sets.

of our proposed components and hyperparameter selection. Four groups of ablation studies were carried out, as illustrated in Table 6.

**(1)-(3) Contribution of Components.** The contributions of each component to the complete method are illustrated by individually removing them. The degree of reliability and semantic primitive category is set to 0.7 and 300, respectively, and other settings are consistent with the full method. Removing consistent reasoning has the greatest impact on the model, with mIoU plummeting by 4.71%.

**(4)-(6) Degree of Reliability.** The selection of dependable points is influenced by the sensitivity to different levels of reliability. Typically, with increasing experience, the reliability threshold is set above 0.5. In this study, we have considered three specific values: 0.6, 0.7, and 0.8. A notable advantage is observed at a threshold of 0.7, leading us to adopt this value as the degree of reliability in our comprehensive methodology.

**(7)-(9) Number of Semantic Primitive.** We investigated




Figure 7: The visualization shows the distance from the prototype for different categories of consistent points. The visualization results for the validation set based on the ScanNet.

the sensitivity to the number of distinct semantic primitive categories, as indicated by the number of clustering centers. Within this framework, we examined three values for comparison: 200, 300, and 400. Notably, a substantial advantage was observed at 300, leading to the adoption of this value as the number of semantic primitive categories in our complete.

**(10)  $\mathcal{L}_{cr}$  Change to using Cross-Entropy.** Replacing the constraint with cross-entropy loss caused a sharp drop in mIoU. It shows that excessive constraints on the similarity distribution between consistent and ambiguous and prototypes are detrimental to semantic recognition.

**Analysis:** Fig. 6 illustrates the similarity measure matrices of both consistent and ambiguous prototypes. At epoch 10, the similarity range is minimal, suggesting little distinction among prototype categories. Introducing the P-SLCR method led to heightened discrepancies among prototypes by epoch 830, revealing prominent features. Notably, the similarity to the prototype itself significantly surpassed the similarity to other categories. Comparison with the unsupervised point cloud semantic segmentation methods demonstrates the state-of-the-art of our approach (P-SLCR). However, we found that the accuracy of unsupervised methods is still low, especially when there is no color information. During training, the model seems to increasingly prioritize approaching areas with higher point density, potentially leading to the neglect of certain correct edges (Fig. 7). (See the Appendix for more discussion)

## Conclusion

We propose a new architecture for prototypes consistent structure learning and consistent reasoning model (P-SLCR), including a library of learnable prototypes. The pseudo-labels generated during clustering, the training phase selects consistent points by reliability for structure learning and constructs semantic consistent reasoning between consistent and ambiguous prototypes. We validated P-SLCR against the latest models using the S3DIS, ScanNet, and SemanticKITTI datasets. The results confirm that our unsupervised approach surpasses the fully supervised method, introducing a novel framework for 3D unsupervised learning. Our future work will continue to explore the potential of P-SLCR for other 3D unsupervised tasks.

## References

- Armeni, I.; Sener, O.; Zamir, A. R.; Jiang, H.; Brilakis, I.; Fischer, M.; and Savarese, S. 2016. 3d semantic parsing of large-scale indoor spaces. In *CVPR*, 1534–1543.
- Behley, J.; Garbade, M.; Milioto, A.; Quenzel, J.; Behnke, S.; Stachniss, C.; and Gall, J. 2019. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *ICCV*, 9297–9307.
- Chen, Z.; Xu, H.; Chen, W.; Zhou, Z.; Xiao, H.; Sun, B.; Xie, X.; et al. 2023. PointDC: Unsupervised Semantic Segmentation of 3D Point Clouds via Cross-modal Distillation and Super-Voxel Clustering. In *ICCV*, 14290–14299.
- Cheng, M.; Hui, L.; Xie, J.; and Yang, J. 2021. Spsc-net: Semi-supervised semantic 3d point cloud segmentation network. In *AAAI*, volume 35, 1140–1147.
- Chibane, J.; Engelmann, F.; Anh Tran, T.; and Pons-Moll, G. 2022. Box2mask: Weakly supervised 3d semantic instance segmentation using bounding boxes. In *ECCV*, 681–699. Springer.
- Cho, J. H.; Mall, U.; Bala, K.; and Hariharan, B. 2021. Picie: Unsupervised semantic segmentation using invariance and equivariance in clustering. In *CVPR*, 16794–16804.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 5828–5839.
- Ester, M.; Kriegel, H.-P.; Sander, J.; Xu, X.; et al. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *KDD*, volume 96, 226–231.
- Graham, B.; Engelcke, M.; and Van Der Maaten, L. 2018. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, 9224–9232.
- Hou, J.; Graham, B.; Nießner, M.; and Xie, S. 2021. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *CVPR*, 15587–15597.
- Hu, Q.; Yang, B.; Fang, G.; Guo, Y.; Leonardis, A.; Trigoni, N.; and Markham, A. 2022. Sqn: Weakly-supervised semantic segmentation of large-scale 3d point clouds. In *ECCV*, 600–619.
- Hu, Q.; Yang, B.; Xie, L.; Rosa, S.; Guo, Y.; Wang, Z.; Trigoni, N.; and Markham, A. 2020. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *CVPR*, 11108–11117.
- Huang, W.; Zou, P.; Xia, Y.; Wen, C.; Zang, Y.; Wang, C.; and Zhou, G. 2024. OPOCA: One Point One Class Annotation for LiDAR Point Cloud Semantic Segmentation. *IEEE Trans. Geosci. Remote Sens.*
- Ji, X.; Henriques, J. F.; and Vedaldi, A. 2019. Invariant information clustering for unsupervised image classification and segmentation. In *ICCV*, 9865–9874.
- Jiang, L.; Shi, S.; Tian, Z.; Lai, X.; Liu, S.; Fu, C.-W.; and Jia, J. 2021. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *ICCV*, 6423–6432.
- Kundu, A.; Yin, X.; Fathi, A.; Ross, D.; Brewington, B.; Funkhouser, T.; and Pantofaru, C. 2020. Virtual multi-view fusion for 3d semantic segmentation. In *ECCV*, 518–535. Springer.
- Li, M.; Xie, Y.; Shen, Y.; Ke, B.; Qiao, R.; Ren, B.; Lin, S.; and Ma, L. 2022. Hybridcr: Weakly-supervised 3d point cloud semantic segmentation via hybrid contrastive regularization. In *CVPR*, 14930–14939.
- Li, W.; Zhan, L.; Min, W.; Zou, Y.; Huang, Z.; and Wen, C. 2023. Semantic Segmentation of Point Cloud With Novel Neural Radiation Field Convolution. *IEEE Geosci. Remote Sens.*
- Liu, J.; Yu, Z.; Breckon, T. P.; and Shum, H. P. 2024. U3DS3: Unsupervised 3D Semantic Scene Segmentation. In *IEEE Winter Applications Computer Vision*, 3759–3768.
- Liu, K.; Zhao, Y.; Nie, Q.; Gao, Z.; and Chen, B. M. 2022a. Weakly supervised 3d scene segmentation with region-level boundary awareness and instance discrimination. In *ECCV*, 37–55. Springer.
- Liu, L.; Zhuang, Z.; Huang, S.; Xiao, X.; Xiang, T.; Chen, C.; Wang, J.; and Tan, M. 2023. Cpcm: Contextual point cloud modeling for weakly-supervised point cloud semantic segmentation. In *CVPR*, 18413–18422.
- Liu, M.; Zhou, Y.; Qi, C. R.; Gong, B.; Su, H.; and Angelov, D. 2022b. Less: Label-efficient semantic segmentation for lidar point clouds. In *ECCV*, 70–89. Springer.
- Lloyd, S. 1982. Least squares quantization in PCM. *IEEE Trans. Inform. Theory*, 28(2): 129–137.
- Meng, H.-Y.; Gao, L.; Lai, Y.-K.; and Manocha, D. 2019. Vv-net: Voxel vae net with group convolutions for point cloud segmentation. In *ICCV*, 8500–8508.
- Milioto, A.; Vizzo, I.; Behley, J.; and Stachniss, C. 2019. Rangenet++: Fast and accurate lidar semantic segmentation. In *IROS*, 4213–4220.
- Papon, J.; Abramov, A.; Schoeler, M.; and Worgotter, F. 2013. Voxel cloud connectivity segmentation-supervoxels for point clouds. In *CVPR*, 2027–2034.
- Qi, C. R.; Su, H.; Mo, K.; and Guibas, L. J. 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 652–660.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, volume 30.
- Ren, Z.; Misra, I.; Schwing, A. G.; and Girdhar, R. 2021. 3d spatial recognition without spatially labeled 3d. In *CVPR*, 13204–13213.
- Robert, D.; Vallet, B.; and Landrieu, L. 2022. Learning multi-view aggregation in the wild for large-scale 3d semantic segmentation. In *CVPR*, 5575–5584.
- Rozenberszki, D.; Litany, O.; and Dai, A. 2022. Language-grounded indoor 3d semantic segmentation in the wild. In *ECCV*, 125–141. Springer.
- Saltori, C.; Galasso, F.; Fiameni, G.; Sebe, N.; Poiesi, F.; and Ricci, E. 2023. Compositional semantic mix for domain adaptation in point cloud segmentation. *IEEE TPAMI*, 45(12): 14234–14247.

- Shi, H.; Wei, J.; Li, R.; Liu, F.; and Lin, G. 2022. Weakly supervised segmentation on outdoor 4D point clouds with temporal matching and spatial graph propagation. In *CVPR*, 11840–11849.
- Shin, I.; Tsai, Y.-H.; Zhuang, B.; Schulter, S.; Liu, B.; Garg, S.; Kweon, I. S.; and Yoon, K.-J. 2022. Mm-tta: multi-modal test-time adaptation for 3d semantic segmentation. In *CVPR*, 16928–16937.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *NeurIPS*, 30.
- Tatarchenko, M.; Park, J.; Koltun, V.; and Zhou, Q.-Y. 2018. Tangent convolutions for dense prediction in 3d. In *CVPR*, 3887–3896.
- Unal, O.; Dai, D.; and Van Gool, L. 2022. Scribble-supervised lidar semantic segmentation. In *CVPR*, 2697–2707.
- Vaswani, A. 2017. Attention is all you need. *NeurIPS*.
- Wu, X.; Jiang, L.; Wang, P.-S.; Liu, Z.; Liu, X.; Qiao, Y.; Ouyang, W.; He, T.; and Zhao, H. 2024. Point Transformer V3: Simpler Faster Stronger. In *CVPR*, 4840–4851.
- Wu, X.; Lao, Y.; Jiang, L.; Liu, X.; and Zhao, H. 2022. Point transformer v2: Grouped vector attention and partition-based pooling. *NeurIPS*, 35: 33330–33342.
- Wu, X.; Wen, X.; Liu, X.; and Zhao, H. 2023. Masked scene contrast: A scalable framework for unsupervised 3d representation learning. In *CVPR*, 9415–9424.
- Zhan, L.; Du, Y.; Jiang, J.; Wei, Y.; Zhou, T.; and Yang, Z. 2025a. BGC-Net: Bilateral Graph Convolutional Network for Weakly-supervised Semantic Segmentation of Large-scale Point Clouds. *IEEE Trans. Circuit Syst. Video Technol.*
- Zhan, L.; Li, W.; Jiang, J.; Zhou, T.; Wen, C.; and Wang, C. 2025b. QPCR: Weakly Supervised Semantic Segmentation of Large-Scale Point Cloud via Gradual Query Points Component Reasoning. *IEEE Trans. Geosci. Remote Sens.*, 63: 3512190.
- Zhan, L.; Li, W.; and Min, W. 2023. FA-ResNet: Feature affine residual network for large-scale point cloud segmentation. *Int. J. Appl. Earth Obs. Geoinf.*, 118: 103259.
- Zhang, R.; Guo, Z.; Zhang, W.; Li, K.; Miao, X.; Cui, B.; Qiao, Y.; Gao, P.; and Li, H. 2022. Pointclip: Point cloud understanding by clip. In *CVPR*, 8552–8562.
- Zhang, Y.; Li, Z.; Xie, Y.; Qu, Y.; Li, C.; and Mei, T. 2021a. Weakly supervised semantic segmentation for large-scale point cloud. In *AAAI*, volume 35, 3421–3429.
- Zhang, Y.; Qu, Y.; Xie, Y.; Li, Z.; Zheng, S.; and Li, C. 2021b. Perturbed self-distillation: Weakly supervised large-scale point cloud semantic segmentation. In *ICCV*, 15520–15528.
- Zhang, Z.; Bai, M.; and Li, E. L. 2022. Self-supervised pre-training for large-scale point clouds. *NeurIPS*, 35: 37806–37821.
- Zhang, Z.; Dai, W.; Wen, H.; and Yang, B. 2025. LogoSP: Local-global Grouping of Superpoints for Unsupervised Semantic Segmentation of 3D Point Clouds. In *CVPR*, 1374–1384.
- Zhang, Z.; Yang, B.; Wang, B.; and Li, B. 2023. Growsp: Unsupervised semantic segmentation of 3d point clouds. In *CVPR*, 17619–17629.
- Zhao, H.; Jiang, L.; Jia, J.; Torr, P. H.; and Koltun, V. 2021a. Point transformer. In *ICCV*, 16259–16268.
- Zhao, S.; Wang, Y.; Li, B.; Wu, B.; Gao, Y.; Xu, P.; Darrell, T.; and Keutzer, K. 2021b. epointda: An end-to-end simulation-to-real domain adaptation framework for lidar point cloud segmentation. In *AAAI*, volume 35, 3500–3509.
- Zhu, X.; Zhou, H.; Wang, T.; Hong, F.; Ma, Y.; Li, W.; Li, H.; and Lin, D. 2021. Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In *CVPR*, 9939–9948.