

Weakly Supervised Semantic Segmentation for Large-Scale Point Cloud

Yachao Zhang¹, Zonghao Li¹, Yuan Xie^{2*}, Yanyun Qu^{1*}, Cuihua Li¹, Tao Mei³

¹ School of Informatics, Xiamen University, Fujian, China

² School of Computer Science and Technology, East China Normal University, Shanghai, China

³ JD AI Research, Beijing, China

{yachaozhang,zonghaoli}@stu.xmu.edu.cn, yxie@cs.ecnu.edu.cn, {yyqu,chli}@xmu.edu.cn, tmei@jd.com

Abstract

Existing methods for large-scale point cloud semantic segmentation require expensive, tedious and error-prone manual point-wise annotations. Intuitively, weakly supervised training is a direct solution to reduce the cost of labeling. However, for weakly supervised large-scale point cloud semantic segmentation, too few annotations will inevitably lead to ineffective learning of network. We propose an effective weakly supervised method containing two components to solve the above problem. Firstly, we construct a pretext task, *i.e.*, point cloud colorization, with a self-supervised learning to transfer the learned prior knowledge from a large amount of unlabeled point cloud to a weakly supervised network. In this way, the representation capability of the weakly supervised network can be improved by the guidance from a heterogeneous task. Besides, to generate pseudo label for unlabeled data, a sparse label propagation mechanism is proposed with the help of generated class prototypes, which is used to measure the classification confidence of unlabeled point. Our method is evaluated on large-scale point cloud datasets with different scenarios including indoor and outdoor. The experimental results show the large gain against existing weakly supervised methods and comparable results to fully supervised methods.

Introduction

3D scene understanding is required in numerous applications, in particular robotics, autonomous driving and virtual reality. Large-scale point cloud semantic segmentation as a fundamental task attracts more and more attention. The success of deep neural networks in point cloud semantic segmentation is attribute to their ability to scale up with more well-labeled training data (Qi et al. 2017a,b; Zhang, Hua, and Yeung 2019; Li, Chen, and Hee Lee 2018; Wang et al. 2019; Wu, Qi, and Fuxin 2019; Yang et al. 2019; Hu et al. 2020).

Fully supervised point cloud semantic segmentation methods need expensive, tedious and error-prone manual point-wise annotations. One direct solution is to achieve effective segmentation via weakly supervised learning with annotating partial points or semantic category which is still in its infancy. Xu and Lee proposed a weakly supervised ap-




Figure 1: Visualize the results of semantic segmentation, the misclassification points are signed in red. From left to right, the columns show the results with one labeled point, 1% labeled points, 10% labeled points and the ground truth.

proach (Xu and Lee 2020) whose results are close to the previous fully supervised performance with $10\times$ fewer labeled points. This approach only deal with an instance (ShapeNet-Part) or a block ($1m \times 1m$ on S3DIS) of point cloud in a small scale.

For large-scale point clouds, the network cannot effectively learn feature representation with a few labeled points. We consider the two strategies to solve this problem: 1) What prior knowledge can be transferred to the segmentation network for improving feature representation in weakly supervised learning? 2) How are the labels of given points propagated to unlabeled points effectively?

Firstly, transfer learning provides a promising way and is relatively successful in 2D vision tasks, such as image classification, segmentation, and so on. But point cloud tasks lack a well-annotated and category-extensive dataset like ImageNet (Deng et al. 2009) for pre-training. We consider to use self-supervised learning for knowledge transfer. However, how to use enormous amounts of unlabeled point cloud to generate labels by itself and learn a semantically related representation for subsequent tasks is very challenging. Secondly, propagating pseudo labels to unlabeled points is a common method for weakly supervised and semi-supervised task to learn effectively. Usually, this requires constructing a fully-connected graph with all points. However, for large-scale point clouds ($\sim 10^6$ points), the fully-connected graph

*Corresponding Author

is unachievable due to large GPU memory consumption.

To address the above difficulties, we present an effective weakly supervised method for large-scale point cloud semantic segmentation. Firstly, we notice that points of the same semantic class have similar color distribution, and point cloud with color is essentially free. We choose point cloud colorization as a pretext task for transfer learning. Specifically, we use color space transformation to construct a self-supervised network for color space completion. We further introduce a local perceptual regularization term to enhance the local representation which is consistent with the goal of segmentation task. As a result, it can learn a prior-based initialization distribution. After that, we use the pre-trained knowledge to fine-tune weakly semantic segmentation network. Our learning scheme allows to transfer the knowledge from the self-supervised task to the weakly supervised task and improves the effectiveness of feature representation. Moreover, we propose a sparse label propagation method induced by class prototype which can gradually propagate pseudo labels to unlabeled points. It is computational friendly that can be adapted to large-scale tasks.

Our contributions can be summarized as follows:

- A weakly supervised segmentation method is proposed for large-scale point cloud which only needs very small number of labeled points.
- We adopt the heterogeneous transfer learning method and construct a self-supervised pretext task by point cloud colorization. It can learn a prior distribution and be generalized well to segmentation tasks.
- We propose an efficient sparse label propagation method which can propagate labels to unlabeled points. It expands the supervision information and has low computational complexity.
- Extensive experimental results demonstrate that our method achieves the comparable to or even exceeds fully supervised competitors.

Related Work

Semantic segmentation for large-scale point cloud. PointNet and PointNet++ (Qi et al. 2017a,b) are pioneering approaches for point clouds. While recent works have shown promising results on small point clouds, most of them cannot directly scale up to large-scale point clouds ($\sim 10^6$ points) due to high computational and memory costs (Hu et al. 2020).

SPG (Landrieu and Simonovsky 2018) processes the large-scale point clouds through a graph convolution-based method. FCPN (Rethage et al. 2018) preprocesses large-scale point clouds into voxels. However, both the graph partitioning and voxelization are computationally expensive. Recently, RandLA-Net (Hu et al. 2020) provides an efficient and lightweight neural architecture for large-scale point cloud semantic segmentation. The state-of-the-art methods mentioned above are fully supervised that require a large number of point clouds with dense point-wise annotation. Obviously, this point-wise annotation is labor-intensive and time-consuming.

Self-supervised learning on point cloud. Self-supervised learning springs up in computer vision recently. It aims to learn good representations from unlabeled visual data, reducing or even eliminating the need for the costly collection of manual labels (Newell and Deng 2020). Recent self-supervised learning makes great successes in feature representation which achieves comparable or outperforms results produced by supervised pre-training (Bachman, Hjelm, and Buchwalter 2019; He et al. 2020).

Unlike 2D images, self-supervised learning in point cloud is rarely used. Previous works on unsupervised 3D representation learning (Achlioptas et al. 2018; Gadelha, Wang, and Maji 2018; Hassani and Haley 2019; Li, Chen, and Hee Lee 2018; Sauder and Sievers 2019; Yang et al. 2018) mainly focus on representing an instance (ShapeNet (Chang et al. 2015)). It is difficult to directly apply the feature representation learned on an instance to real-world large-scale point cloud tasks due to the large domain gaps (Xie et al. 2020).

PointContrast (Xie et al. 2020) concerns on contrastive embeddings and proposed a pre-training task for 3D point cloud understanding. The core of this method is that different views of point cloud should be mapped to similar embeddings for matched points. It achieves promising results on fully-supervised downstream tasks. Xu and Lee (Xu and Lee 2020) proposed a Siamese self-supervision branch by augmenting the training sample with a random in-plane rotation and flipping, and then made the original and augmented point-wise predictions be consistent. This method (Xu and Lee 2020) treats self-supervision as a branch of multi-task and only the current training samples are used. Therefore, it cannot make full use of other massive point clouds to learn representations with strong generalization.

Weakly Supervised Point Cloud Semantic Segmentation. There are few researches on weakly supervised point cloud semantic segmentation. Following the weakly supervised manner called incomplete supervision in (Zhou 2018), the recent work (Xu and Lee 2020) utilizes $10\times$ fewer labeled points to achieve comparable performance to fully supervised method in small scale part segmentation and small blocks ($\sim 10^3$ points) semantic segmentation. MPRM (Wei et al. 2020) introduces a multi-path region mining module to generate pseudo point-level label from a classification network. The classification network is trained by the subcloud-level weakly labels. For the scene-level as an input, the performance will degrade. However, up to now, none of them can be generalized to large-scale problem well.

Proposed Method

Overview

Our method aims to exploit the knowledge transfer and label propagation to solve the problem of unstable and poor representation produced by the network under weakly supervised large-scale point cloud segmentation. We propose an effective weakly supervised large-scale method and depict the overview framework in Figure 2.

We take point cloud colorization as a self-supervised pretext task to learn a prior-based initialization distribution. A local perceptual regularization is proposed to learn the con-

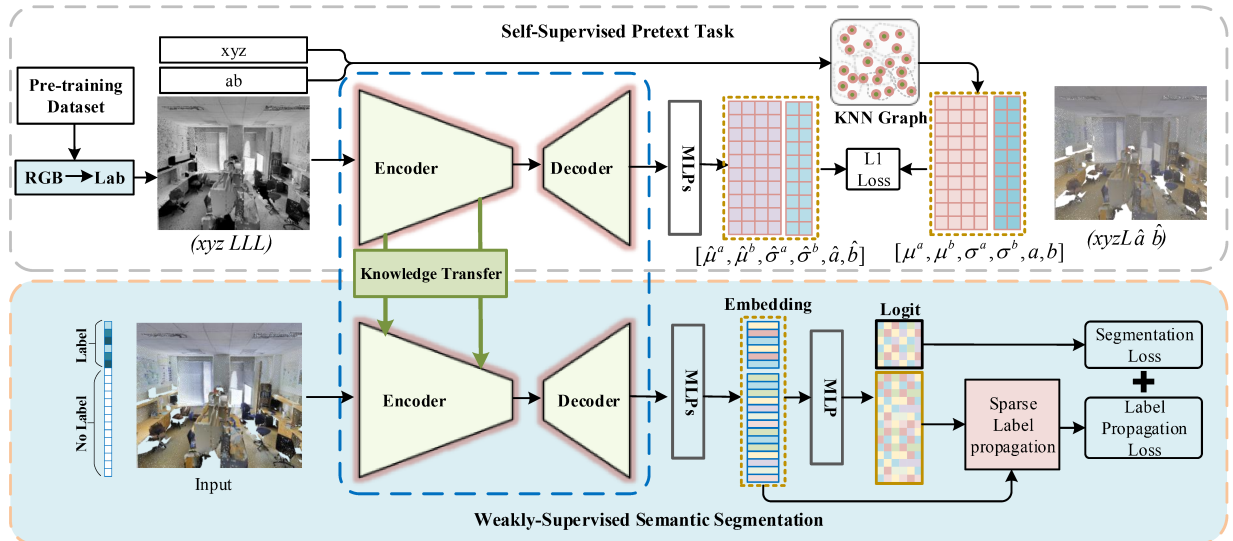


Figure 2: The framework of our method consists of three parts: i) Self-supervised pretext task learns a prior knowledge. ii) The prior knowledge is used to fine-tune the weakly-supervised semantic segmentation network. iii) Sparse label propagation generates pseudo label for unlabeled data to improve the effectiveness of weakly-supervised task.

textual information. Then, we use pre-trained parameters of encoder to initialize the weakly-supervised network for improving the effectiveness of feature presentation.

Furthermore, we leverage labeled points to directly supervise the network and fine-tune network parameters. We also introduce a non-parametric label propagation method for weakly supervised semantic segmentation. Some unlabeled points are assigned pseudo labels through the similarity between class prototypes (Li et al. 2020; Qiu, Yao, and Mei 2017) and embeddings of unlabeled points. Therefore, more supervised information is introduced to improve effectiveness of training. Considering of computational and memory efficiency for large-scale point cloud, we choose RandLA-Net (Hu et al. 2020) as the backbone, which is an efficient and lightweight neural architecture for large-scale point clouds semantic segmentation. In the following, we describe the self-supervised pretext task and the sparse label propagation method.

Self-Supervised Pretext Task

Colorization provides a powerful supervisory signal unlike training from scratch in 2D vision task. Training data are easy to collect, so any point cloud with color can be used as training data. Due to the progress of point cloud acquisition equipment, we have access to enormous amounts of unlabeled point cloud data with color information. We investigate and implement self-supervised learning on point cloud colorization which is treated as a pretext task. Point colorization aims to guide the self-supervised model to learn the feature representation.

It is recognized that the **Lab** color space favor for perceptual distance (Zhang, Isola, and Efros 2016; Larsson, Maire, and Shakhnarovich 2017). Thus, we perform point cloud colorization by **a**, **b** completion in this colorspace. Given

the lightness channel **L**, the network predicts the **a** and **b** color channels and the local Gaussian distribution for each point. Notably, the value in channel **L** is replicated to three times of each point to keep the same dimension as the input of segmentation task. Therefore, the input point cloud $X^s = [x_1, x_2, \dots, x_{N^s}] \in \mathcal{R}^{N^s \times 6}$ consists of N^s 3D points with the xyz coordinates and three **L**. N^s is the number of points in one point cloud.

Moreover, we implement RandLA-Net on the self-supervised task by modifying the final output layer. That is, the output of the network is a 6-dimension vector which contains the predicted \hat{a} , \hat{b} and the corresponding local mean and variance.

Loss of Self-Supervised Task The loss is inherited from standard regression problem, which minimizes \mathcal{L}_1 error between the prediction and ground truth. Given a point cloud with coordinates and triplicate lightness values, the self-supervised pretext task learns a mapping $\hat{Y} = \mathcal{F}(X^s; \Theta)$, $\hat{Y} = \{\hat{a}, \hat{b}, \hat{\mu}^a, \hat{\sigma}^a, \hat{\mu}^b, \hat{\sigma}^b\}$, where \hat{a} , \hat{b} , $\hat{\mu}$ and $\hat{\sigma}$ denote predicted **a**, **b** and corresponding local mean and variance, respectively. The loss of self-supervised task can be formulated as:

$$\mathcal{L}_{ab} = \frac{1}{2N^s} \sum_{i=1}^{N^s} (\|a_i - \hat{a}_i\|_1 + \|b_i - \hat{b}_i\|_1). \quad (1)$$

In addition, to learn the local color distribution of every point, we introduce a local perceptual regularization term. If the network can predict the color distribution (mean and variance) of the neighbors, it can embed the local information which is consistent with the segmentation task using local features for weakly supervised semantic segmentation. Given a point x_i as the centroid, the local neighbor $\mathcal{N}(x_i)$ is

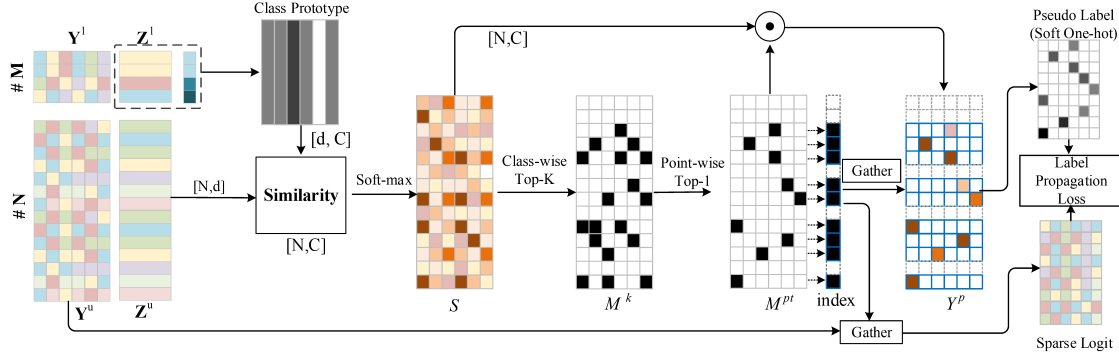


Figure 3: The framework of sparse label propagation. \odot is Hadamard product and **Gather** denotes the operator of getting elements by index.

calculated by KNN according to the Euclidean distance. The ground truth μ_i^a and σ_i^a of a channel can be obtained by:

$$\mu_i^a = \frac{1}{K} \sum_{j=1}^K a_j, \forall x_j \in \mathcal{N}(x_i), \quad (2)$$

$$\sigma_i^a = \sqrt{\frac{1}{K} \sum_{j=1}^K (a_j - \mu_i^a)^2 + \varepsilon}, \forall x_j \in \mathcal{N}(x_i), \quad (3)$$

where ε is a very small constant. μ_i^b, σ_i^b can be obtained in the same way. We formulate the local perceptual regularization term as:

$$\mathcal{L}_{local} = \frac{1}{4N^s} \sum_{i=1}^{N^s} (\|\mu_i^a - \hat{\mu}_i^a\|_1 + \|\sigma_i^a - \hat{\sigma}_i^a\|_1 + \|\mu_i^b - \hat{\mu}_i^b\|_1 + \|\sigma_i^b - \hat{\sigma}_i^b\|_1). \quad (4)$$

The total loss \mathcal{L}_p of self-supervised pretext task can be expressed as:

$$\mathcal{L}_p = \mathcal{L}_{ab} + \mathcal{L}_{local}. \quad (5)$$

Discussion *Why is the knowledge learned from the pretext task beneficial for semantic segmentation?*

- Pretext task learns similar feature distributions as the semantic segmentation task. Objects in the same category usually have similar color distribution, for example the vegetation is typically green, and the road is black. The surface color texture of the scene provides ample cues for many categories.
- Pretext task embeds the local feature representation. We introduce a local perceptual regularization term to constrain the local color distribution to be consistent with the original distribution. Thus, it allows the network to embed more local information. Therefore, it may enhance the embedding of local features for semantic segmentation task.

Sparse Label Propagation

Segmentation performance degrades significantly with few labeled points. The main reason is that supervisory information provided by few labeled points can not be propagated well to unlabeled points. Therefore, we use the labeled

points to assign pseudo labels for unlabeled points, and further provide additional supervised information to improve the representation of the weakly supervised network.

In order to achieve this goal, the following items require to be taken into account: 1) Computational complexity is not high and memory recourse is not large. Large-scale point clouds usually contain $\sim 10^6$ points, if using all the points as nodes to construct a fully-connected graph, it will consume a lot of memory and computing resources. 2) The anchor points should be sparse. Some ambiguous points should not be given labels to train the network. 3) The propagated label should be soft. The propagated labels should be related to their similarity, and the higher of the similarity, the more similar the label should be.

We design a sparse label propagation method. The overall framework is shown in Figure 3. It consists of three parts: class prototype generation, class assignment matrix construction, sparse pseudo-label generation.

Class prototype generation. In the last two layers of the network, we output embedding $\mathbf{Z} = [z_1^l, z_2^l, \dots, z_M^l; z_1^u, z_2^u, \dots, z_N^u] \in \mathcal{R}^{(M+N) \times d}$ and the corresponding prediction $\mathbf{Y} = [y_1^l, y_2^l, \dots, y_M^l; y_1^u, y_2^u, \dots, y_N^u] \in \mathcal{R}^{(M+N) \times C}$ which comes from M labeled points and N unlabeled points. We use $\mathbf{Z}^l, \mathbf{Z}^u$ to represent the embedding of labeled and unlabeled points, respectively. Firstly we generate C prototypes to represent the C classes according to the labeled points. Specifically, we simply take the mean of the labeled point embeddings \mathbf{Z}^l for each class. For class c , the prototype ρ_c is given by:

$$\rho_c = \frac{1}{|\mathcal{I}_c|} \sum_{z_i^l \in \mathcal{I}_c} z_i^l, \quad (6)$$

where \mathcal{I}_c denotes the embedding sets of class c , $c = \{1, 2, \dots, C\}$ for labeled points.

Class assignment matrix construction. We leverage embedding \mathbf{Z}^u of unlabeled points and the class prototypes to construct a similarity matrix $\mathcal{W} \in \mathcal{R}^{N \times C}$ by:

$$\mathcal{W} = \exp\left(-\frac{\|z_i^u - \rho_c\|^2}{\sigma}\right), \quad (7)$$

where σ is a hyper-parameter. Each column of \mathcal{W} represents

the similarity between the unlabeled point and the class prototypes. We use *soft-max* to convert the similarity into a class assignment probability S as:

$$S = \mathcal{P}(i \mapsto c) = \frac{\exp(\mathcal{W})}{\sum_{c=1}^C \exp(\mathcal{W})}. \quad (8)$$

Sparse pseudo label generation. There are some points with low similarity to each category. These points are not suitable to provide supervisory information to train the network. Specifically, for each class, according to the class assignment matrix S , we select the *top-K* unlabeled points and get the mask $M^k \in \{0, 1\}^{N \times C}$, where $m_{ic}^k = 1$ means the embedding of *i*-th point is the first *K* points similar to the class *c*. N is the number of unlabeled points. This is a label expansion method with a balanced number of categories. It can alleviate the category imbalance to a certain extent.

As an unlabeled point may belong to multiple categories, we choose the most similar category and generate a binary mask. According to M^k , we get the point mask $M^{pt} \in \{0, 1\}^N$. $m_i^{pt} = 1$ denotes that the *i*-th point is assigned a pseudo labels. We can get the sparse pseudo label $Y^p \in \mathcal{R}^{N \times C}$ by:

$$Y^p = M^{pt} \odot S, \quad (9)$$

where \odot denotes Hadamard product. Y^p is the form of soft one-hot. The label propagation loss \mathcal{L}_{sp} can be formulated as:

$$\mathcal{L}_{sp} = -\frac{1}{\|M^{pt}\|_1} \sum_{i=1}^N m_i^{pt} \sum_{c=1}^C y_{ic}^p \log y_{ic}^u, \quad (10)$$

where y_{ic}^u is the probability that the unlabeled point *i* is classified as category *c*, and $\|\cdot\|_1$ denotes *L1* norm.

Compared with the traditional fully-connected graph label propagation method, our method is computation efficient. The complexity of our method is $\mathcal{O}(NCd)$, while the method of fully-connected graph is $\mathcal{O}((N+M)^2d)$. C is the number of category and d denotes the dimension of \mathbf{Z}^u . The magnitude of C is $\sim 10^1$, which is much smaller than $N(\sim 10^6)$.

Loss of Weakly Supervised Task

The loss of weakly supervised semantic segmentation includes two terms: segmentation loss and label propagation loss.

$$\mathcal{L}_{total} = \mathcal{L}_{seg} + \lambda \mathcal{L}_{sp}, \quad (11)$$

We utilize a softmax cross-entropy loss on the labeled points. For labeled points, the segmentation loss is formulated as:

$$\mathcal{L}_{seg} = -\frac{1}{M} \sum_{i=1}^M \sum_{c=1}^C y_{ic} \log \frac{\exp(y_{ic}^l)}{\sum_{c=1}^C \exp(y_{ic}^l)}, \quad (12)$$

where y_{ic} is the ground truth of labeled point *i*. M denotes the number of labeled points.

In the early stage of training, the embedding is unreliable. Class prototypes generated by embedding can not represent classes well. Thus, the label propagation loss should not be introduced to optimize network parameters. As the embedding gets better, the weight of the pseudo loss should

increase. Therefore, we introduce a non-linear parameter λ to balance the two losses. In this work, λ is formulated as:

$$\lambda = \begin{cases} 0, & epoch < 30 \\ e^{\frac{epoch}{max_epoch} - 1}, & otherwise \end{cases} \quad (13)$$

where *epoch* and *max_epoch* denote the current training epoch and the total epochs, respectively.

Experiments and Analysis

In this section, we firstly introduce the experimental settings. Then, we evaluate the weakly semantic segmentation performance on indoor and outdoor large-scale point clouds, respectively. Furthermore, we perform ablation study to evaluate the importance of the main components.

Experiment Settings

Dataset Setting for Self-supervised Task Previous works on unsupervised 3D representation learning (Achlioptas et al. 2018; Gadelha, Wang, and Maji 2018; Hassani and Haley 2019; Li, Chen, and Hee Lee 2018; Sauder and Sievers 2019; Yang et al. 2018) mainly focused on ShapeNet (Chang et al. 2015) which is a dataset of single-object CAD point cloud models. However, the real-world large-scale point cloud usually contains multiple models of different categories. Pre-training on ShapeNet have poor scalability because of a large domain gap. We choose ScanNet (Dai et al. 2017) as the pre-training dataset which is a big real-world point cloud dataset containing 2.5 million views in more than 1500 indoor scans.

To train self-supervised pretext task, we convert the **RGB** space to the **Lab** space and split 6 channels corresponding to the coordinate (x, y, z) and color (L, a, b) into the given channels (x, y, z, L) and the prediction channels (a, b) .

Dataset Setting for Weakly Supervised Segmentation

In order to evaluate the performance of our network on weakly semantic segmentation tasks, we experiment on the indoor scene dataset S3DIS (Armeni et al. 2016) and ScanNetv2 (Dai et al. 2017) and outdoor dataset Semantic3D (Hackel et al. 2017). Note that our method is pre-trained on ScanNet dataset. However, the pre-training task does not use the semantic labels.

Implementation Details Here weakly supervised settings are studied. i) 1 point label (*1pt*), we assume there is only one point within each category labeled with ground-truth for each scene. ii) (*x%*) denote *x* percentage points with ground-truth. We set $x = \{1, 10\}$ in our experiments. iii) The cost of labeling 1% points manually is much higher than *1pt*. We consider a cheap way: super-point (*SPT*), which annotates a local area instead of one point. In all settings, the annotated points are randomly selected.

Our pre-training and weakly semantic segmentation network are built on the efficient backbone of RandLA-Net. The two networks can be trained built upon TensorFlow with a single NVIDIA Titan RTX. We use Adam optimizer with default parameters. The number of neighboring points *K* is

Setting	Method	Area5	6-fold
Fully	PointNet [’17]	41	47.6
	DGCNN [’19]	-	56.1
	RSNet [’18]	56.5	-
	PointCNN [’18]	57.3	65.4
	ShellNet [’19]	-	66.8
	RandLA-Net [’20]	62.8*	70.0
1pt (0.2%)	II Model [’16]	44.3	-
	MT [’17]	44.4	-
	Xu [’20]	44.5	-
0.2%	Ours	56.4	-
1pt (0.03%)	Baseline	40.4	-
	Ours	45.8	-
1%	Ours	61.8	65.9
10%	II Model [’16]	46.3	-
	MT [’17]	47.9	-
	Xu [’20]	48.0	-
	Ours	64.0	68.1

Table 1: Comparisons of performance on S3DIS (Armeni et al. 2016) (mIoU %). Area5 denotes evaluation on Area 5. 6-fold is cross evaluation of 6 areas. The superscript * indicates the result evaluated by the official code.

16. Both the self-supervised pretext task and weakly semantic segmentation network are trained 80 epochs. We calculate the mean Intersect over Union (mIoU) to evaluate the performance.

Evaluation on S3DIS

We conduct the comparison experiments with state-of-the-art weakly supervised (Weakly) and fully supervised (Fully) on the indoor dataset S3DIS. The former contains: II Model (Laine and Aila 2016), MT (Tavainen and Valpola 2017) and Xu (Xu and Lee 2020). The latter includes: PointNet (Qi et al. 2017a), DGCNN (Wang et al. 2019), RSNet (Huang, Wang, and Neumann 2018), PointCNN (Li et al. 2018), ShellNet (Zhang, Hua, and Yeung 2019), RandLA-Net (Hu et al. 2020).

The results are presented in Table 1. Overall, we observe the consistent improvement in the performance of segmentation with more labeled points. Our 1pt denotes only one labeled point for each category in the entire rooms instead of small blocks (e.g., 1×1 meter). Therefore, our 1pt represents that total labeled points are about 0.03%. While the total labeled points of (Xu and Lee 2020) are about 0.2%. Ours still achieves better performance against Xu (Xu and Lee 2020) at 1pt setting. Under 1% setting, we achieve comparable performance against the fully supervised method, RandLA-Net, and even exceed the previously fully supervised methods by a large margin. At 10% setting, ours is overwhelmingly superior to the previous weakly supervised methods and even achieves 1.2% improvement against RandLA-Net (Hu et al. 2020) on Area 5. These results show the effectiveness of our method.

The qualitative results are shown in Figure 1 on S3DIS dataset under 1pt, 1% and 10% settings. It can be seen that

Setting	Method	mIoU
Fully	PointNet++ [’17]	33.9
	SPLATNet [’18]	39.3
	PointCNN [’18]	45.8
	KPConv [’19]	68.4
Weakly	MPRM (subcloud) [’20]	41.1
	Ours (<i>SPT</i>)	49.0
	Ours (1%)	51.1
	Ours (10%)	52.0

Table 2: Quantitative results on ScanNetv2 (mIoU %).

at 1% setting, our method can correctly classify except for challenging boundaries.

Evaluation on ScanNetv2

We compare our method with five fully supervised methods: PointNet++ (Qi et al. 2017b), SPLATNet (Su et al. 2018), PointCNN (Li et al. 2018), KPConv (Thomas et al. 2019), and a weakly supervised method: MPRM (Wei et al. 2020). Table 2 shows the comparison results on test set. From *SPT* to 10% setting, ours achieves significant gain and make a breakthrough margin compared with fully supervised method PointCNN. At *SPT* setting, ours gains 7.9% against MPRM which is trained with subcloud-level semantic labels (subcloud). We cannot directly and quantitatively evaluate the annotated labor of *SPT* and MPRM. Compared to MPRM, ours needs to draw an extra little region at random. But it reduces the labor of dividing into subclouds and repeatedly annotating semantic category for each subcloud.

Evaluation on Semantic3D

For outdoor point cloud semantic segmentation, we compare ours against five state-of-the-art fully supervised methods: SnapNet (Boulch, Le Saux, and Audebert 2017), SEGCloud (Tchapmi et al. 2017), ShellNet (Zhang, Hua, and Yeung 2019), KPConv (Thomas et al. 2019) and RandLA-Net (Hu et al. 2020). The results of online testing are summarized in Table 3.

Clearly, the more labeled points, the better the semantic segmentation performance is. With 1% labeled points, we achieve 72.6% mIoU and outperforms fully supervised ShellNet (Zhang, Hua, and Yeung 2019) by the gain of over 3.3%. With 10% labeled points, our method is close to the current state-of-the-art RandLA-Net. Therefore, our method is also effective in outdoor dataset.

Ablation Study

We conducted ablation studies to evaluate the importance of the main components: self-supervised pre-training task and sparse label propagation. *Scratch* denotes random initialization. *Pre-training* represents that using the parameters of self-supervised pretext task for initialization. $\lambda = 0$ represent without sparse label propagation. $\lambda = 1$ and ‘nonlinear’ denote a constant 1 and a nonlinear way weighting the sparse label propagation, respectively.

Self-supervised pre-training or training from scratch.

Setting	Method	mIoU	OA
Fully	SnapNet [’17]	59.1	88.6
	SEGCloud [’17]	61.3	88.1
	ShellNet [’19]	69.3	93.2
	KPConv [’19]	74.6	92.9
	RandLA-Net [’20]	77.4	94.8
Weakly	Ours (1%)	72.6	93.7
	Ours (10%)	73.3	94.0

Table 3: Quantitative results on Semantic3D (reduced-8) (Hackel et al. 2017). ‘OA’ denotes overall accuracy.

Setting		1pt	SPT	1%	10%
Scratch	$\lambda = 0$	40.4	55.8	58.6	60.6
Pre-training	$\lambda = 0$	45.4	57.5	60.1	61.3
	$\lambda = 1$	-	-	-	-
	nonlinear	45.8	58.2	61.8	64.0


Table 4: Comparisons of different components on S3DIS Area 5 (Armeni et al. 2016) (mIoU %).

From the comparison of the first and second rows of Table 4, the pre-training makes the performance better, while training from scratch drops in performance for 5.0%, 1.7%, 1.5%, 0.7%. We hypothesize that our pre-training is effective because it can transfer knowledge learned from supervised tasks to subsequent tasks. This encourages the subsequent task to learn a more robust feature distribution. We can still find that the fewer annotation points, the greater the effect of self-supervision. We infer that fewer labeled points cannot learn a good feature representation. Therefore, it needs more additional knowledge to promote feature representation learning.

Sparse label propagation. We further analyse the effects of sparse label propagation. Comparing the second and fourth rows of Table 4, our *nonlinear* obtains 0.4%, 0.7%, 1.7%, and 2.7% improvement over $\lambda = 0$, respectively. The more labeled points, the better performance of the semantic segmentation achieves. Because more labeled points can learn better representations and produce accurate class prototypes. When λ is a constant 1, the network training collapses. Because the initial training (the first 30 epochs) cannot learn a good feature representation. It is unreasonable to use these features for label propagation. In the process of feature learning gradually getting better, a larger weight is gradually introduced to the loss term.

Effectiveness and efficiency. We conduct a pilot study to understand the effectiveness and show the mIoU curve of our method and baseline on validation set. From Figure 4, our method converges faster and has a higher stable value which demonstrates our method is more effective for weakly supervised large-scale point cloud segmentation.

Like RandLA-Net (Hu et al. 2020), we feed the entire scene to the network. We list the training time of each epoch (RandLA-Net: 267s and Ours: 280s) and total test time (RandLA-Net: 115s and Ours: 116s) on Area-5 of S3DIS dataset. As the sparse label propagation is introduced in the



(a) On S3DIS

(b) On Semantic3D

Figure 4: mIoU curve of our method and baseline on the validation set.



Figure 5: Visualize the results of point cloud colorization.

training phase, training time of per-epoch increases by 13s. But their test time are almost the same.

We also visualize the results of colorization in Figure 5 on S3DIS. Although our point cloud colorization focuses on the representation learning, it can be seen that the colorization is close to real-world color.

Conclusion

We present a weakly supervised semantic segmentation method for large-scale point cloud. With the help of a self-supervised pretext task and sparse label propagation, our method significantly outperforms the weakly supervised and almost reaches the accuracy of fully-supervised methods on three challenging large-scale point cloud dataset. The experimental results show self-supervised knowledge transfer is an effective way to improve weak supervised performance and our method make a breakthrough margin against the compared weakly supervised method. Furthermore, the more the labeled points, the better performance our method achieves, which is consistent with our expectation. In the future, we expect future work to explore more effective self-supervised knowledge representation for weakly or semi-supervised point cloud tasks.

Acknowledgments

This work is supported by the National Natural Science Foundation of China under Grant 61876161, Grant 61772524, the National Key Research and Development Program of China No.2020AAA0108301, and Natural Science Foundation of Shanghai No.20ZR1417700.

References

- Achlioptas, P.; Diamanti, O.; Mitliagkas, I.; and Guibas, L. 2018. Learning representations and generative models for 3d point clouds. In *ICML*, 40–49.
- 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.
- Bachman, P.; Hjelm, R. D.; and Buchwalter, W. 2019. Learning representations by maximizing mutual information across views. In *NeurIPS*, 15535–15545.
- Boulch, A.; Le Saux, B.; and Audebert, N. 2017. Unstructured point cloud semantic labeling using deep segmentation networks. *3DOR* 2: 7.
- Chang, A. X.; Funkhouser, T.; Guibas, L.; Hanrahan, P.; Huang, Q.; Li, Z.; Savarese, S.; Savva, M.; Song, S.; Su, H.; et al. 2015. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*.
- 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.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. Imagenet: A large-scale hierarchical image database. In *CVPR*, 248–255.
- Gadelha, M.; Wang, R.; and Maji, S. 2018. Multiresolution tree networks for 3d point cloud processing. In *ECCV*, 103–118.
- Hackel, T.; Savinov, N.; Ladicky, L.; Wegner, J. D.; Schindler, K.; and Pollefeys, M. 2017. SEMANTIC3D.NET: A new large-scale point cloud classification benchmark. In *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, volume IV-1-W1, 91–98.
- Hassani, K.; and Haley, M. 2019. Unsupervised multi-task feature learning on point clouds. In *CVPR*, 8160–8171.
- He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 9729–9738.
- 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, Q.; Wang, W.; and Neumann, U. 2018. Recurrent slice networks for 3d segmentation of point clouds. In *CVPR*, 2626–2635.
- Laine, S.; and Aila, T. 2016. Temporal ensembling for semi-supervised learning. In *ICLR*.
- Landrieu, L.; and Simonovsky, M. 2018. Large-scale point cloud semantic segmentation with superpoint graphs. In *CVPR*, 4558–4567.
- Larsson, G.; Maire, M.; and Shakhnarovich, G. 2017. Colorization as a proxy task for visual understanding. In *CVPR*, 6874–6883.
- Li, J.; Chen, B. M.; and Hee Lee, G. 2018. So-net: Self-organizing network for point cloud analysis. In *CVPR*, 9397–9406.
- Li, X.; Yu, L.; Fu, C.-W.; Cohen-Or, D.; and Heng, P.-A. 2020. Unsupervised Detection of Distinctive Regions on 3D Shapes. *ACM Transactions on Graphics* 39(5): 1–14.
- Li, Y.; Bu, R.; Sun, M.; Wu, W.; Di, X.; and Chen, B. 2018. Pointcnn: Convolution on x-transformed points. In *NeurIPS*, 820–830.
- Newell, A.; and Deng, J. 2020. How useful is self-supervised pretraining for visual tasks? In *CVPR*, 7345–7354.
- 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*, 5099–5108.
- Qiu, Z.; Yao, T.; and Mei, T. 2017. Learning spatio-temporal representation with pseudo-3d residual networks. In *ICCV*, 5533–5541.
- Rethage, D.; Wald, J.; Sturm, J.; Navab, N.; and Tombari, F. 2018. Fully-convolutional point networks for large-scale point clouds. In *ECCV*, 596–611.
- Sauder, J.; and Sievers, B. 2019. Self-supervised deep learning on point clouds by reconstructing space. In *NeurIPS*, 12962–12972.
- Su, H.; Jampani, V.; Sun, D.; Maji, S.; Kalogerakis, E.; Yang, M.-H.; and Kautz, J. 2018. Splatnet: Sparse lattice networks for point cloud processing. In *CVPR*, 2530–2539.
- Tarvainen, A.; and Valpola, H. 2017. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 1195–1204.
- Tchapmi, L.; Choy, C.; Armeni, I.; Gwak, J.; and Savarese, S. 2017. Segcloud: Semantic segmentation of 3d point clouds. In *International Conference on 3D Vision (3DV)*, 537–547.
- Thomas, H.; Qi, C. R.; Deschard, J.-E.; Marcotegui, B.; Goulette, F.; and Guibas, L. J. 2019. Kpconv: Flexible and deformable convolution for point clouds. In *ICCV*, 6411–6420.
- Wang, Y.; Sun, Y.; Liu, Z.; Sarma, S. E.; Bronstein, M. M.; and Solomon, J. M. 2019. Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics* 38(5): 1–12.
- Wei, J.; Lin, G.; Yap, K.-H.; Hung, T.-Y.; and Xie, L. 2020. Multi-path region mining for weakly supervised 3D semantic segmentation on point clouds. In *CVPR*, 4384–4393.
- Wu, W.; Qi, Z.; and Fuxin, L. 2019. Pointconv: Deep convolutional networks on 3d point clouds. In *CVPR*, 9621–9630.
- Xie, S.; Gu, J.; Guo, D.; Qi, C. R.; Guibas, L. J.; and Litany, O. 2020. PointContrast: Unsupervised pre-training for 3D point cloud understanding. In *ECCV*.

Xu, X.; and Lee, G. H. 2020. Weakly supervised semantic point cloud segmentation: towards 10x fewer labels. In *CVPR*, 13706–13715.

Yang, J.; Zhang, Q.; Ni, B.; Li, L.; Liu, J.; Zhou, M.; and Tian, Q. 2019. Modeling point clouds with self-attention and gumbel subset sampling. In *CVPR*, 3323–3332.

Yang, Y.; Feng, C.; Shen, Y.; and Tian, D. 2018. Foldingnet: Point cloud auto-encoder via deep grid deformation. In *CVPR*, 206–215.

Zhang, R.; Isola, P.; and Efros, A. A. 2016. Colorful image colorization. In *ECCV*, 649–666.

Zhang, Z.; Hua, B.-S.; and Yeung, S.-K. 2019. Shellnet: Efficient point cloud convolutional neural networks using concentric shells statistics. In *ICCV*, 1607–1616.

Zhou, Z.-H. 2018. A brief introduction to weakly supervised learning. *National Science Review* 5(1): 44–53.