

CPCM: Contextual Point Cloud Modeling for Weakly-supervised Point Cloud Semantic Segmentation

Lizhao Liu^{1,2} Zhuangwei Zhuang^{1,2} Shangxin Huang¹ Xunlong Xiao¹ Tianhang Xiang¹

Cen Chen¹ Jingdong Wang³ Mingkui Tan^{1,2†}

¹South China University of Technology ²Pazhou Lab ³Baidu Inc.

{selizhaoliu, z.zhuangwei, sevtars, sexxl, sexiangtianhang}@mail.scut.edu.cn,

{chencen, mingkuitan}@scut.edu.cn, wangjingdong@baidu.com

Abstract

We study the task of weakly-supervised point cloud semantic segmentation with sparse annotations (e.g., less than 0.1% points are labeled), aiming to reduce the expensive cost of dense annotations. Unfortunately, with extremely sparse annotated points, it is very difficult to extract both contextual and object information for scene understanding such as semantic segmentation. Motivated by masked modeling (e.g., MAE) in image and video representation learning, we seek to endow the power of masked modeling to learn contextual information from sparsely-annotated points. However, directly applying MAE to 3D point clouds with sparse annotations may fail to work. First, it is non-trivial to effectively mask out the informative visual context from 3D point clouds. Second, how to fully exploit the sparse annotations for context modeling remains an open question. In this paper, we propose a simple yet effective Contextual Point Cloud Modeling (CPCM) method that consists of two parts: a region-wise masking (RegionMask) strategy and a contextual masked training (CMT) method. Specifically, RegionMask masks the point cloud continuously in geometric space to construct a meaningful masked prediction task for subsequent context learning. CMT disentangles the learning of supervised segmentation and unsupervised masked context prediction for effectively learning the very limited labeled points and mass unlabeled points, respectively. Extensive experiments on the widely-tested ScanNet V2 and S3DIS benchmarks demonstrate the superiority of CPCM over the state-of-the-art.

1. Introduction

With the growing demand for autonomous driving and robotic navigation, point cloud semantic segmentation becomes an indispensable technique for accurate 3D environ-

[†]Corresponding author.

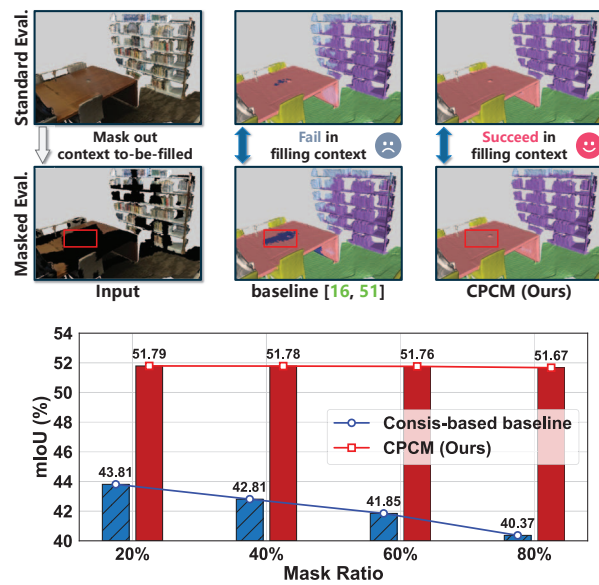


Figure 1: Effectiveness of the proposed CPCM on context comprehension ability compared to the consistency-based baseline [16, 53]. We conduct masked evaluations to inspect the model’s contextual understanding ability. The visual comparison of results from different methods (mask ratio = 40%) and the performance w.r.t. different mask ratios are shown in the top and bottom panels, respectively.

ment perception [18, 27, 51]. Recent years have witnessed great progress in fully-supervised learning in point cloud segmentation [2, 6, 10, 14, 31, 32, 39, 47, 56]. However, densely-annotating point-wise labels are time-consuming, labor-intensive as well as economic-inefficient to obtain since the number of points in point cloud data can easily reach tens of thousands of magnitude [42, 48]. It goes without saying that diving into point cloud semantic segmentation from sparse labels is crucial to reduce the annotation cost and expand the application boundary [9, 20, 22].

Very recently, to reduce the reliance on dense labels

while still delivering satisfactory point cloud semantic segmentation performance, most effort has been put into learning from the weakly-annotated labels [9, 16, 25, 42, 48, 49, 52, 53]. Among several types of weakly-annotated labels, the partial point-wise labeling scheme offers the best trade-off between annotation cost and segmentation performance [9, 22]. In the partially annotated point cloud data, the labeled part typically occupies a very small portion of points (*e.g.*, 0.1%) per scene [9]. In this case, directly applying supervised cross-entropy loss only on the limited labeled part is prone to overfitting [25, 33, 43]. As a result, the primary challenge is learning from a significant proportion of unlabeled points to improve model generalization performance, rather than utilizing only the labeled points [16, 53].

Existing methods seek to tackle the challenge by exploiting different levels of feature consistency under various data augmentations. To be specific, researchers resort to enforcing feature consistency between differently augmented or geometrically calibrated point clouds by discriminating points from different scenes with contrastive learning [8, 11, 16, 45], exploring color & geometric smoothness [48, 52], more advanced consistency loss such as JS-divergence [53] and similarity weighted loss [43]. However, given limited annotations, exploring feature consistency only would be insufficient to capture the complex structures of point clouds, making it very difficult to extract both contextual and object information for satisfactory segmentation performance. To inspect the consistency-based methods’ comprehension of scene context, we conduct a pilot study by masked evaluation: evaluate the segmentation performance given a *context-to-be-filled* point cloud. As shown in Figure 1, the performance of the consistency-based method degenerates drastically, indicating a poor understanding of the scene context, even in this simple case. Thus, comprehending the complex scene context from mass unlabeled points remains an unresolved issue.

Motivated by masked modeling (*e.g.*, MAE [7]) in image and video that learns good representations by masking random patches of the input image and reconstructing the missing information, we seek to endow the power of masked modeling for weakly-supervised point cloud segmentation. However, directly employing MAE to 3D point clouds with sparse annotations may fail to work due to the following reasons. **First**, since 3D point clouds are typically unordered and irregular, it is nontrivial to mask out the informative visual context from the 3D point clouds for subsequent context learning. **Second**, considering the limited but valuable labeled data in the weakly-annotated point cloud, how to fully exploit the labeled points in masked modeling remains an open question.

To address the above issues, we propose a simple yet effective Contextual Point Cloud Modeling (CPCM) that consists of two parts: region-wise masking (RegionMask) strat-

egy and a contextual masked training (CMT) method. To be specific, RegionMask evenly divides the geometric space into a set of cuboids and masks all points within the cuboids selected with a given mask ratio. Different from the trivial point-wise masking solution [26] that performs point-wise random masking, our RegionMask masks the point cloud continuously in the geometric space to provide a meaningful masked context prediction task. Beyond that, RegionMask is able to control the difficulty of the masked feature prediction task by adjusting a hyper-parameter region size, showing flexibility in handling different amounts of annotation. Similar to MAE [7], we expect that with a very high mask ratio (*i.e.*, 0.75), the model is able to learn more visual concepts [7], thereby mastering the contextual information. However, as shown in our experiments, directly incorporating the masked modeling objective into the consistency-based training framework impedes learning from the limited but valuable labeled points, resulting in degenerated performance. To resolve this problem, we propose a contextual masked training (CMT) method that adds an extra masked feature prediction branch into the consistency-based framework, which not only paves the way for learning labeled data but allows the model to effectively learn the complex scene context. The proposed CPCM achieves state-of-the-art performance on two widely-tested benchmarks ScanNet V2 and S3DIS. For example, on ScanNet V2 [4], CPCM outperforms SQN [9] by **5.6%** mIoU on online test set.

Our contributions are summarized as follows:

- We propose contextual point cloud modeling that incorporates masked modeling into the consistency-based training framework to effectively learn contextual information from sparsely-annotated data.
- We propose a region-wise masking strategy that masks the point cloud continuously to construct the meaningful masked prediction task and a contextual masked training method that facilitates the learning from limited labeled data and masked context prediction.
- To the best of our knowledge, we are the first to explore 3D masked modeling on weakly-supervised point cloud segmentation. Extensive experiments on widely-tested benchmarks demonstrate the superior performance of the proposed CPCM.

2. Related Work

Fully-supervised point cloud segmentation. There are mainly three kinds of fully-supervised methods proposed to encode the 3D point cloud into effective representations for semantic segmentation, including point-based [10, 13, 15, 31, 32, 41, 54], voxel-based [3, 6, 14, 17, 23, 24, 34, 35, 44] and hybrid methods [2, 47]. Early attempts [37, 39, 56]

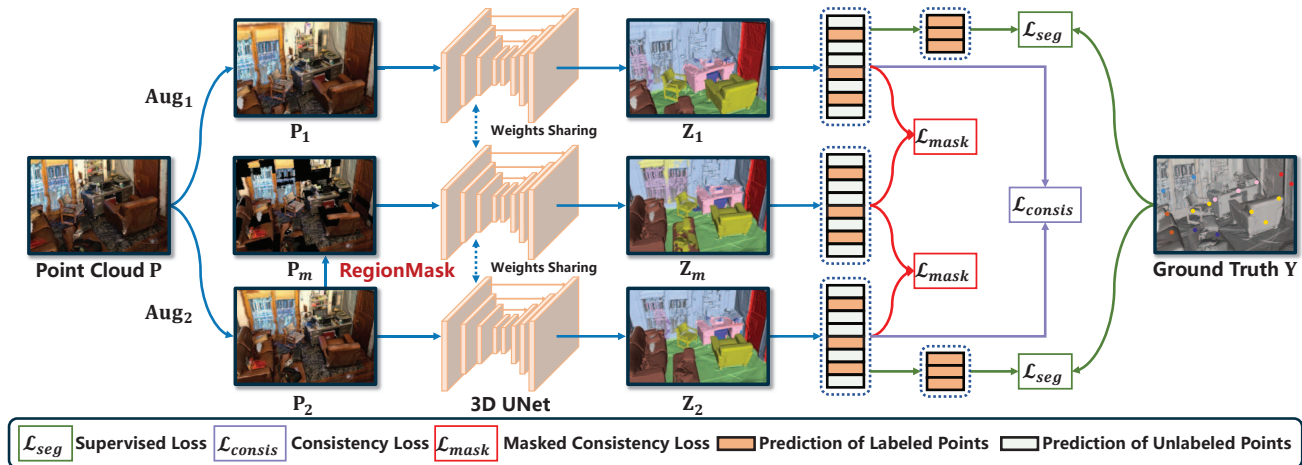


Figure 2: Overall scheme of our CPCM method. Given a point cloud P , we first apply two random augmentations and our region-wise masking to obtain the augmented point clouds P_1, P_2 and the masked point cloud P_m , respectively. Then, the features Z_1, Z_2, Z_m are extracted by a weight-sharing 3D UNet. The supervised cross-entropy loss \mathcal{L}_{seg} is computed over labeled features and a consistency loss \mathcal{L}_{consis} is computed on Z_1, Z_2 . Last, our masked consistency loss \mathcal{L}_{mask} enforces the feature consistency between Z_1, Z_m and Z_2, Z_m to help the model focus on learning contextual information.

simply employ the 2D convolution on the projected point cloud image, which is efficient but the projection process causes the loss of 3D geometric detail. The point-based methods are proposed to directly process the irregular and unordered points with order-agnostic architectures such as PointNet [31] and PointNet++ [32] that can be naturally applied to the point cloud but are less effective than 2D convolution in encoding the contextual information [10, 13, 15, 41, 54]. The voxel-based methods [3, 6, 14] combine the neighboring points into regular grids and often leverage sparse convolution [6, 17, 34, 35, 44] to handle the sparse voxelized data. The latest works combine the merits from both worlds and form hybrid methods, but also bring more complex architecture design and extra training costs [2, 47]. Overall, the fully-supervised point cloud segmentation methods have a strong dependence on densely-annotated labels, limiting their application scenarios.

Weakly-supervised point cloud segmentation. Learning from weakly annotated point cloud data has become a hot research topic [9, 25, 42, 43, 48, 52, 53], which not only reduces the annotation cost but also turns out to be a more general solution for real-life segmentation scenarios [22, 42]. For the partially labeled point cloud, the supervised cross-entropy loss is suitable to learn from the labeled points, which, however, is prone to learn an overfit segmentation model due to the very limited annotations [25, 33, 43]. Thus, existing approaches focus on learning the major unlabeled part and can be grouped into two paradigms: pseudo labeling [9, 25, 42] and consistency-based regularization [16, 48, 49, 52, 53]. The pseudo-labeling methods predict pseudo-labels of the unlabeled

points to explore them. MPRM [42] trains a segmentation model on the sub-cloud labels and uses the class activation map [55] to pseudo-label the whole sub-cloud to train the final model. OTOC [25] improves the quality of the pseudo labels with multi rounds self-training. SQN [9] leverages the geometric prior to better use limited labels. Since the pseudo label is destined to be inaccurate, consistency-based approaches learn the feature consistency across augmentations [16, 43, 48, 49, 52, 53] or calibrated views [43] to use mass unlabeled data. MIL [49] enforce scene-level feature consistency for model optimization. Moreover, point-wise consistency is also leveraged by considering the color or geometric smoothness [48, 52], feature similarity [43, 53] or using pseudo-labeling as guidance [16]. However, feature consistency across augmentations may not fully comprehend the complex structures of weakly-annotated point clouds. Instead, we propose to learn masked feature consistency to better explore the contextual information.

Masked modeling for vision. Masked modeling has been a long endeavor to learn effective representation from vision data. Early attempts reconstruct RGB features from masked images [30], which are improved by masking a very high ratio of image content to learn meaningful visual representation [7, 21, 46, 50]. Moreover, masked supervised learning improves the perception of contextual information in fully-supervised image semantic segmentation [57]. Recently, researchers apply the masked modeling approach to learn unlabeled point cloud data [19, 26, 28]. Unlike the above settings, weakly-supervised point cloud segmentation provides both labeled and unlabeled data. Moreover, applying masked modeling tailored for unsupervised / fully-

supervised learning to both labeled and unlabeled data simultaneously is rarely explored. In this paper, we propose a contextual masked training method to learn from the limited supervision and the masked feature prediction task for weakly-supervised point cloud semantic segmentation.

3. Contextual Point Cloud Modeling

Notations. Formally, a point cloud data is a collection of N points $\mathbf{P} = \{p_1, p_2, \dots, p_N\}$, where each point p_n often comprises the geometric location and RGB information, *i.e.*, $p_n = \mathbf{P}[n] = (x_n, y_n, z_n, r_n, g_n, b_n)$. We use $[\cdot]$ as the index operation that retrieves the corresponding element (can be a vector or a scalar) from a set or a matrix. To accomplish the point cloud semantic segmentation task, given a point cloud \mathbf{P} and a segmentation network $f_\theta(\cdot)$ parameterized by θ , we expect the model to produce point-wise classification features¹ $\mathbf{Z} = \text{Softmax}(f_\theta(\mathbf{P}))$, where $\mathbf{Z}[n] \in (0, 1)$, $\text{argmax}(\mathbf{Z}[n]) \in \mathcal{C}$ and $\mathcal{C} = \{0, 1, 2, \dots, C - 1\}$ is a predefined category set with C classes. Unlike the fully-supervised point cloud semantic segmentation that provides the label of every point in \mathbf{P} , only sparse annotations are available in weakly-supervised point cloud semantic segmentation. The weakly-labeled point cloud data comprises two parts, the labeled part and the unlabeled part, *i.e.*, $(\mathbf{P}, \mathbf{Y}) = \{(p_s, y_s) \mid s \in \mathcal{S}\} \cup \{(p_u, \emptyset) \mid u \in \mathcal{U}\}$, where \mathcal{S}, \mathcal{U} denote the index sets of the labeled and unlabeled points respectively and \emptyset is a special token denoting the label is unavailable. During model training, a dataset $\mathcal{D} = \{(\mathbf{P}, \mathbf{Y})\}$ includes hundreds of or thousands of point cloud & weak-label pair is provided.

3.1. Problem Definition

With the limited labeled data and a mass of unlabeled data, weakly-supervised point cloud semantic segmentation focuses on learning useful representations from a large amount of unlabeled data to improve model generalization. Existing approaches often achieve this by enforcing point-wise feature consistency across augmentations [16, 48, 49, 53]. Given a weakly-labeled point cloud data (\mathbf{P}, \mathbf{Y}) , two random augmentations² are applied $\mathbf{P}_1 = \text{Aug}_1(\mathbf{P})$ and $\mathbf{P}_2 = \text{Aug}_2(\mathbf{P})$. Based on this, point-wise classification for two point clouds is calculated by $\mathbf{Z}_1 = \text{Softmax}(f_\theta(\mathbf{P}_1))$, $\mathbf{Z}_2 = \text{Softmax}(f_\theta(\mathbf{P}_2))$. The general form for the consistency-based method is as follows:

$$\mathcal{L}_{\text{CB}} = \mathcal{L}_{\text{seg}} + \alpha \mathcal{L}_{\text{consis}}, \quad (1)$$

where \mathcal{L}_{seg} and $\mathcal{L}_{\text{consis}}$ denote supervised cross-entropy loss and the consistency loss introduced below and α is a hyper-parameter that controls optimization strength on the consistency loss. The supervised loss \mathcal{L}_{seg} is computed

¹We use the term features and logits interchangeably for convenience.

²Details on the data augmentation are put in the supplementary.

Algorithm 1 Training method for CPCPM

Require: The training dataset $\mathcal{D} = \{(\mathbf{P}, \mathbf{Y})\}$, the point cloud segmentation network $f_\theta(\cdot)$, the region size G , the mask ratio R , the weighting factor α, β , the learning rate η .

Ensure: Optimized point cloud segmentation network f_θ .

- 1: Randomly initializes the model parameters θ .
 - 2: **while** not converge **do**
 - 3: Obtain a weakly-labeled point cloud data (\mathbf{P}, \mathbf{Y}) from \mathcal{D} .
 - 4: Obtain the labeled indexes \mathcal{S} from \mathbf{Y} .
 - 5: *// perform two random augmentations*
 - 6: $\mathbf{P}_1 \leftarrow \text{Aug}_1(\mathbf{P}), \mathbf{P}_2 \leftarrow \text{Aug}_2(\mathbf{P})$.
 - 7: Compute region-wise masking flag \mathbf{M} by Eqn. (4).
 - 8: Compute region-wise masked point cloud \mathbf{P}_m by Eqn. (7).
 - 9: *// perform segmentation for augmented point clouds*
 - 10: $\mathbf{Z}_1 \leftarrow \text{Softmax}(f_\theta(\mathbf{P}_1)), \mathbf{Z}_2 \leftarrow \text{Softmax}(f_\theta(\mathbf{P}_2))$.
 - 11: *// perform segmentation for the masked point cloud*
 - 12: $\mathbf{Z}_m \leftarrow \text{Softmax}(f_\theta(\mathbf{P}_m))$.
 - 13: Compute the cross-entropy loss \mathcal{L}_{seg} by Eqn. (2).
 - 14: Compute the consistency loss $\mathcal{L}_{\text{consis}}$ by Eqn. (3).
 - 15: Compute the masked consistency loss $\mathcal{L}_{\text{mask}}$ by Eqn. (9).
 - 16: Compute the overall training objective $\mathcal{L}_{\text{CPCM}}$ by Eqn. (8).
 - 17: *// update network parameters via gradient descent*
 - 18: $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{CPCM}}$.
 - 19: **end while**
-

over limited labeled points:

$$\mathcal{L}_{\text{seg}} = \frac{1}{|\mathcal{S}|} \sum_{s \in \mathcal{S}} \text{CE}(\mathbf{Z}_1[s], \mathbf{Y}[s]) + \text{CE}(\mathbf{Z}_2[s], \mathbf{Y}[s]), \quad (2)$$

where $\text{CE}(\cdot, \cdot)$ is the cross-entropy loss. In the meanwhile, the consistency loss $\mathcal{L}_{\text{consis}}$ enforces point-wise feature consistency as follows:

$$\mathcal{L}_{\text{consis}} = \frac{1}{N} \sum_n \text{JS}(\mathbf{Z}_1[n], \mathbf{Z}_2[n]), \quad (3)$$

where $\text{JS}(\cdot, \cdot)$ minimizes the Jensen-Shannon divergence of different features. Feature consistency from different augmentations can exploit the unlabeled data but may not be informative enough to comprehend the complex structure of the point cloud data, failing to effectively explore the contextual information such as space, color and semantic continuity that is crucial for satisfactory segmentation. Attracted by the strong context modeling ability of masked modeling in image and video representation learning, we seek to endow the power of masked modeling to weakly-supervised point cloud segmentation. However, designing an effective masking strategy for 3D point cloud data and developing a compatible training scheme to fully exploit the limited labeled data for masked modeling remain open questions.

Overview. To answer the above questions, we propose Contextual Point Cloud Modeling (CPCM) to model the contextual information effectively with two steps: **First**, we propose a region-wise masking strategy that masks the point cloud in the continuous geometric space, providing meaningful missing context to be filled. **Second**, we propose a




Figure 3: Comparisons of different masking strategies. The proposed region-wise masking removes meaningful context to be filled. We set the mask ratio = 25% for visualization.

contextual masked training that facilitates the learning of limited labeled points and masked feature prediction tasks by adding an extra stream for masked feature extraction. Then, we enforce the feature consistency between masked and unmasked features to learn effective contextual representations. The overall framework and algorithm of CPCM are shown in Figure 2 and Algorithm 1, respectively.

3.2. Region-wise Point Cloud Masking

In this section, we introduce our region-wise masking scheme that provides an effective supervision signal for the model to learn contextual information. To formulate the masking strategy, we first define $\mathbf{M} \in \mathbb{R}^N$ as a zero-one vector to indicate whether a point in point cloud³ $\mathbf{P} \in \mathbb{R}^{N \times 6}$ is masked or not and denote the mask ratio as R ($0 \leq R \leq 1$), *i.e.*, the number of the masked points is $R * N$. Then, the masked point cloud \mathbf{P}_m is computed in a point-wise setting the color information to zero⁴:

$$\mathbf{P}_m[n] = [x_n, y_n, z_n, \mathbf{M}[n] \cdot r_n, \mathbf{M}[n] \cdot g_n, \mathbf{M}[n] \cdot b_n]. \quad (4)$$

To obtain a masked point cloud, a straightforward solution, termed PointMask, is to randomly sample each point (or voxel) with the given mask ratio R

$$\mathbf{M}[n] = \mathbb{1}\{q \leq R\}, \quad q \sim U[0, 1], \quad (5)$$

where $\mathbb{1}\{\cdot\}$ is the indicator function and q is a random variable drawn from the uniform distribution $U[0, 1]$. As shown in Table 4, PointMask delivers unsatisfactory improvement compared to the baseline, especially with a very high mask ratio (*i.e.*, 0.75). We attribute this failure to the following reasons: The PointMask strategy tends to decrease the resolution of the point cloud (see Figure 3b), which does not effectively mask meaningful visual words [7] to predict.

To reasonably remove some contextual information from a point cloud, we introduce Region-wise Masking (RegionMask) that evenly splits the scene into cuboids

and masks the points within the randomly selected cuboids. We first define the region size G to denote the number of cuboids. Note that a cuboid that parallels the axes in a 3D coordinate system is represented by $[(x_{\min}, y_{\min}, z_{\min}), (x_{\max}, y_{\max}, z_{\max})]$. Assuming that the minimal cuboid covering a point cloud is $[(0, 0, 0), (l, w, h)]$. We evenly partition the scene into a set of cuboid regions \mathcal{H} *i.e.*, ($|\mathcal{H}| = G^3$) as follows:

$$\begin{aligned} \mathcal{H} &= \left\{ [(x_i, y_j, z_k), (x_{i+1}, y_{j+1}, z_{k+1})] \right\}, \\ x_i &= i \cdot \frac{l}{G}, y_j = j \cdot \frac{w}{G}, z_k = k \cdot \frac{h}{G}, \\ i, j, k &\in \{0, 1, \dots, G-1\}, \end{aligned} \quad (6)$$

where x_i, y_j, z_k are the evenly split points along the x, y, z axes and $(\frac{l}{G}, \frac{w}{G}, \frac{h}{G})$ are the length, width, height of a region, respectively. Then, we randomly select $R \cdot G^3$ regions \mathcal{H}^m and compute the mask flag \mathbf{M} as follows:

$$\mathbf{M}[n] = \mathbb{1}\{(x_n, y_n, z_n) \in \mathcal{H}^m\}, \quad (7)$$

where \in denotes a point that lies within a cuboid or not. Then, the masked point cloud is computed by Eqn. (4). As shown in Figure 3c, RegionMask masks the unordered and irregular point cloud continuously, providing meaningful context-to-be-filled patterns such as partial inner-instance mask and cross-instance mask. Moreover, as shown in Section 4.3, RegionMask is able to flexibly cope with different amounts of annotation by adjusting the region size.

3.3. Contextual Masked Training Method

In this section, we introduce our contextual masked training method for learning the contextual information between the masked and unmasked data. We first consider the mask operation as a ‘‘strong augmentation’’ and incorporate it directly into the consistency-based training framework. However, as shown in Figure 4, the training cross-entropy error significantly increases and the performance drops considerably. These results indicate that the input distribution is significantly altered by the mask operation, which impedes learning from limited but valuable labeled points.

Training objective. Taking both the learning from limited labeled data and the learning of contextual information into account, we propose to add an extra branch to perform the masked features prediction task while leaving the two weakly-supervised branches untouched. To be specific, given a weakly-labeled point cloud data (\mathbf{P}, \mathbf{Y}) , we obtain two point clouds $\mathbf{P}_1, \mathbf{P}_2$ by two random augmentations and the masked version \mathbf{P}_m by the proposed RegionMask. Then, we extract their corresponding features $\mathbf{Z}_1, \mathbf{Z}_2, \mathbf{Z}_m$ with the segmentation model $\text{Softmax}(f_\theta(\cdot))$. Last, the overall training objective for our contextual masked training is as follows

$$\mathcal{L}_{\text{CPCM}} = \mathcal{L}_{\text{seg}} + \alpha \mathcal{L}_{\text{consis}} + \beta \mathcal{L}_{\text{mask}}, \quad (8)$$

³For convenience, we refer to the point cloud data as a matrix.

⁴The coordinate x, y, z is left untouched since the sparse convolution operation in 3D UNet requires it for the convolution kernel construction.

Method	Setting	mIoU (%)	ceiling	floor	wall	beam	column	window	door	chair	table	bookcase	sofa	board	clutter
MinkNet* [3]	Fully	68.2	91.7	98.7	83.8	0.0	24.7	56.8	72.1	91.5	83.5	73.3	70.8	81.3	58.4
PointNet [31]		41.1	88.8	97.3	69.8	0.1	4.0	46.3	10.8	58.9	52.6	5.9	40.3	26.4	33.2
KPCConv [40]		67.1	92.8	97.3	82.4	0.0	23.9	58.0	69.0	91.0	81.5	75.3	75.4	66.7	58.9
RandLA-Net [10]		62.4	91.2	95.7	80.1	0.0	25.2	62.3	47.4	75.8	83.2	60.8	70.8	65.2	54.0
RFCR [5]		68.7	94.2	98.3	84.3	0.0	28.5	62.4	71.2	92.0	82.6	76.1	71.1	71.6	61.3
II Model [12]	10%	46.3	91.8	97.1	73.8	0.0	5.1	42.0	19.6	66.7	67.2	19.1	47.9	30.6	41.3
MT [38]		47.9	92.2	96.8	74.1	0.0	10.4	46.2	17.7	67.0	70.7	24.4	50.2	30.7	42.2
10×Fewer [48]		48.0	90.9	97.3	74.8	0.0	8.4	49.3	27.3	69.0	71.7	16.5	53.2	23.3	42.8
SPT [52]	1%	61.8	91.5	96.9	80.6	0.0	18.2	58.1	47.2	75.8	85.7	65.3	68.9	65.0	50.2
PSD [53]		63.5	92.3	97.7	80.7	0.0	27.8	56.2	62.5	78.7	84.1	63.1	70.4	58.9	53.2
HybridCR [16]		65.3	92.5	93.9	82.6	0.0	24.2	64.4	63.2	78.3	81.7	69.0	74.4	68.2	56.5
II Model [12]	0.2%	44.3	89.1	97.0	71.5	0.0	3.6	43.2	27.4	62.1	63.1	14.7	43.7	24.0	36.7
MT [38]		44.4	88.9	96.8	70.1	0.1	3.0	44.3	28.8	63.6	63.7	15.5	43.7	23.0	35.8
10×Fewer [48]		44.5	90.1	97.1	71.9	0.0	1.9	47.2	29.3	62.9	64.0	15.9	42.2	18.9	37.5
SQN [9]	0.1%	61.4	91.7	95.6	78.7	0.0	24.2	55.9	63.1	62.9	70.5	67.8	60.7	56.1	50.6
CPCM (Ours)		66.3 (+4.9)	91.4	95.5	82.0	0.0	30.8	54.1	70.1	87.6	79.4	70.0	67.0	77.8	56.6
PSD [53]	0.03%	48.2	87.9	96.0	62.1	0.0	20.6	49.3	40.9	55.1	61.9	43.9	50.7	27.3	31.1
HybridCR [16]		51.5	85.4	91.9	65.9	0.0	18.0	51.4	34.2	63.8	78.3	52.4	59.6	29.9	39.0
MIL [49]	0.02%	51.4	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
MIL* [49]		52.1	89.2	95.5	74.8	0.2	19.2	41.1	23.1	76.3	64.7	62.6	27.8	57.8	44.8
CPCM (Ours)		62.3 (+10.2)	92.6	95.6	79.4	0.0	17.8	49.3	59.4	85.7	75.6	69.1	60.7	68.2	55.8

Table 1: Comparisons with state-of-the-art methods on S3DIS area5 test set. * denotes results based on our reimplementation.

where β is a hyper-parameter to control the optimization strength of contextual masked learning and \mathcal{L}_{mask} is our masked consistency loss introduced below.

Masked consistency loss. We seek to learn contextual information through masked and unmasked features. To this end, we propose to minimize the distribution gap between masked and unmasked features. In this way, the model shall learn to leverage the unmasked part in the masked point cloud *i.e.*, the surrounding context, thereby improving segmentation performance. Specifically, with the features $\mathbf{Z}_1, \mathbf{Z}_2, \mathbf{Z}_m$ respectively extracted from the two randomly augmented and the masked point clouds, we introduce our masked consistency loss as follows:

$$\mathcal{L}_{mask} = \frac{1}{N} \sum_n JS(\mathbf{Z}_1[n], \mathbf{Z}_m[n]) + JS(\mathbf{Z}_2[n], \mathbf{Z}_m[n]), \quad (9)$$

where the unmasked features $\mathbf{Z}_1, \mathbf{Z}_2$ are considered as the “ground truth” and we detach the gradients of $\mathbf{Z}_1, \mathbf{Z}_2$ during masked consistency loss calculation.

4. Experiments

Datasets. We consider two benchmark datasets ScanNet V2 [4] and S3DIS [1]. ScanNet V2 has 20 semantic classes and the number of training / validation / testing scans is 1,201 / 312 / 100 respectively. We evaluate our model on both val and online test set following [9, 16, 49]. S3DIS, a large-scale point cloud dataset, contains 6 areas with 271 rooms and 13 semantic categories. We adopt the widely-used area5 test set [48, 53] for evaluation, where the number of training and testing scans is 204 and 68, respectively.

Implementation details. We implement our method using MinkowskiEngine [3], a sparse convolution library based on PyTorch [29], as done in previous works [9, 49]. As for the model architecture, we adopt the 34-layer Sparse Residual U-Net [36] following previous works [8, 45].

For evaluation, we use the class-wise Intersection over Union (IoU) and mean IoU (mIoU) metrics. For optimization, we employ the SGD optimizer with $\text{lr} = 1e^{-2}$, weight decay = $1e^{-3}$, the polynomial learning rate scheduler with decay rate = 0.9 and set the batch size to 2 and 4 for ScanNet V2 and S3DIS, respectively. During training, the voxel size is set to 2cm and 5cm for ScanNet V2 and S3DIS, respectively. All models are trained for 180 epochs. We choose JS-divergence as our consistency loss [53]. We refer to the annotation ratio $< 0.1\%$ (including 20 points on ScanNet V2) as the extreme-limited annotations and $\geq 0.1\%$ as the limited annotations. As for the region size G and mask ratio R in RegionMask, we set the mask ratio $R = 0.75$ and set $G = 8$ and $= 4$ for the extreme-limited and limited annotations, respectively. As for (α, β) in \mathcal{L}_{CPCM} , we set $(\alpha, \beta) = (5, 10)$ and $= (1, 5)$ for the extreme-limited and limited annotations, respectively.⁵ All experiments are conducted on 2 and 1 TITAN 3090 GPU(s) for ScanNet V2 and S3DIS, respectively. Our source code is publicly available at <https://github.com/lizhaoliu-Lec/CPCM>.

4.1. Comparison with State-of-the-arts

Quantitative results on S3DIS. We provide the quantitative results on S3DIS in Table 1. For fair comparisons, our approach is evaluated under the same settings used by prior works *i.e.*, the annotation ratio being 0.2%, 0.1%, and 0.02%. The proposed CPCM consistently outperforms the previous state-of-the-art across different annotation ratios, often by a large margin. To be specific, CPCM outperforms SQN by 4.9% under the 0.1% setting and beats MIL by 10.2% under the extreme-limited annotation setting 0.02%. Notably, our CPCM trained by 0.1% label is able to sur-

⁵Analysis on hyper-parameters α, β are put in the supplementary.

Method	Setting	Val	Test
PointNet++ [32]	Fully	N/A	33.9
KPConv [40]		N/A	68.4
MinkNet [3]		72.9	73.6
MPRM [42]	Scene	21.9	N/A
WYPR [33]		29.6	24.0
MIL [49]		26.2	N/A
MPRM [42]	Subcloud	43.2	41.1
MIL [49]		47.4	45.8
SPT [52]	1%	N/A	51.1
PSD [53]		N/A	54.7
HybridCR [16]		56.9	56.8
SQN [9]	0.1%	58.4	56.9
CPCM (Ours)		63.8 (+5.4)	62.5 (+5.6)
WYPR [33]	20 pts	51.5	N/A
OTOC [†] [25]		55.1	N/A
MIL [49]		57.8	54.4
CPCM (Ours)		62.7 (+4.9)	62.8 (+8.4)

Table 2: Comparisons with state-of-the-art methods on ScanNet V2. † indicates results reproduced by MIL [49].

pass the HybridCR trained by 1% label. By diving into per-class mIoU, we observe that our CPCM performs well in relatively small instance categories in a scene such as “chair”, “table”, and “sofa” that tend to be misclassified, which cannot be accomplished without effectively understanding the scene context. Moreover, with 0.1% annotations only, CPCM achieves competitive performance to the fully supervised MinkNet (66.3 vs. 68.2), closing the gap between fully and weakly supervised methods.

Quantitative results on ScanNet V2. We evaluate our approach under 0.1% and 20 points (pts) settings on ScanNet V2 and the quantitative results are shown in Table 2. Although the amount of annotation is very limited, the proposed CPCM provides substantial improvements over prior SoTAs. Specifically, on the validation set, CPCM leads SQN by 5.4% under the 0.1% setting and MIL by 4.9% under the 20 pts setting. Moreover, on the private test set, CPCM still leads SQN and MIL by 5.6% and 8.4% respectively, showing the strong generalization ability of CPCM.

4.2. Ablation Analysis on CPCM

Comparisons to baselines. Since our implementation is based on the fully-supervised MinkNet and the weakly-supervised consis-based method, we directly compare them to investigate the effectiveness of CPCM. The results are shown in Table 3. MinkNet performs decently with 0.1% annotation ratio but suffers from extreme-limited annotation 0.01%. The consis-based method delivers noticeable improvements on both datasets for all settings, showing that it is a strong baseline. Unsurprisingly, the proposed CPCM completely beats the MinkNet and the consis-based baseline, often by a large margin. Notably, when it comes to the extreme-limited 0.01% setting, CPCM boosts the performance of MinkNet by 14.6% and 11.6% on ScanNet V2 and S3DIS, respectively. These results demonstrate the advantage of CPCM that effectively comprehends the scene

Method	\mathcal{L}_{consis}	\mathcal{L}_{mask}	ScanNet V2		S3DIS	
			0.01%	0.1%	0.01%	0.1%
MinkNet	✗	✗	37.6	60.3	47.7	62.9
Consis-based	✓	✗	44.2 (+6.6)	61.8 (+1.5)	52.9 (+5.2)	64.9 (+2.0)
CPCM (Ours)	✓	✓	52.2 (+14.6)	63.8 (+3.5)	59.3 (+11.6)	66.3 (+3.4)

Table 3: Comparisons with two strong baselines: *fully-supervised* method MinkNet trained on weakly-annotated labels and the *weakly-supervised* consis-based method.

context over the strong consis-based baseline.

Masking Strategy	ScanNet V2 (0.01%)		S3DIS (0.01%)	
	0.15	0.75	0.15	0.75
Consis-based	44.2		52.9	
PointMask	42.3 (-1.9)	48.2 (+4.0)	52.3 (-0.6)	55.1 (+2.2)
RegionMask (Ours)	46.5 (+2.3)	52.2 (+8.0)	55.8 (+2.9)	59.3 (+6.4)

Table 4: Ablation studies on different masking strategies. The contextual masked training modeling scheme is employed. Otherwise, all masking strategies show degenerated performance compared to the consis-based baseline.

Region masking. Since random point masking is a common solution in masked vision modeling and has recently been applied to unsupervised point cloud data learning [26]. We investigate the behavior of PointMask under both low and high mask ratios and the results are put in Table 4. On one hand, when the mask ratio is low (0.15), PointMask performs even slightly worse than the consis-based baseline while the proposed RegionMask boosts the performance by 2.3% and 2.9% on the ScanNet V2 and S3DIS, respectively. On the other hand, when the mask ratio is high (0.75), RegionMask considerably improves the performance while PointMask brings only a relatively marginal boost. We conclude that RegionMask is able to mask more meaningful visual words than PointMask under both low and high mask ratios, paving the path of promising masked vision modeling for weakly-supervised point cloud segmentation.





Figure 4: Evolution of training cross-entropy (CE) error and test mIoU w.r.t. training epochs on S3DIS (0.01%).

Contextual masked training. We investigate the effectiveness of the proposed contextual masked training (CMT) by removing the masking stream, resulting in a consistency-based framework with a “masking augmentation”. As shown in Figure 4, the training cross-entropy error drastically increases without CMT, which indicates simply in-

corporating “masking augmentation” hampers the learning of limited but valuable labeled data. With CMT, the segmentation model shows low training cross-entropy error as well as high test mIoU. Moreover, we also put the quantitative results in Table 5 and observe a noticeable performance drop when discarding CMT. Then, with CMT, CPCM achieves substantial improvements over the consis-based baseline. These results verify that CPCM facilitates the learning of valuable annotation but also rich context information, achieving substantial improvements.

RM	CMT	ScanNet V2		S3DIS	
		0.01%	0.1%	0.01%	0.1%
✗	✗	44.2	61.8	52.9	64.9
✓	✗	41.6 (-2.6)	58.6 (-3.2)	51.1 (-1.8)	63.6 (-1.3)
✓	✓	52.2 (+6.7)	63.8 (+2.0)	59.3 (+6.4)	66.3 (+1.4)

Table 5: Ablation studies on our contextual masked training scheme. RM and CMT are short for RegionMask strategy and contextual masked training, respectively.



(a) Effect of region size.

(b) Effect of mask ratio.

Figure 5: Further analysis on the proposed CPCM. (a) We investigate the effect of region size on S3DIS under 0.01% and 0.1% settings. (b) We investigate the effect of mask ratio on S3DIS and ScanNet V2 under the 0.01% setting.

4.3. Further Analysis on CPCM

Region size. As the region size increases, the task of contextual information comprehension becomes easier since the masked region to predict becomes smaller. Therefore, we are able to control the difficulty of the context comprehension task by varying the region size. With less annotation, we may set the masked features prediction task easier. In Figure 5a, the optimal region size becomes smaller when the annotation ratio goes up *i.e.*, 8 for 0.01% and 4 for 0.1%, which verifies the flexibility of the proposed RegionMask strategy for handling different annotation ratios.

Mask ratio. More meaningful visual context will be covered as the mask ratio grows. As shown in Figure 5b, the segmentation performance is constantly boosted by a larger mask ratio up to 0.75, showing the strong potential of our CPCM to effectively explore the scene context. The optimal mask ratio is 0.75 and exceeds which the masked context prediction task becomes too hard to achieve the best result.

⁶More qualitative results can be found in the supplementary.

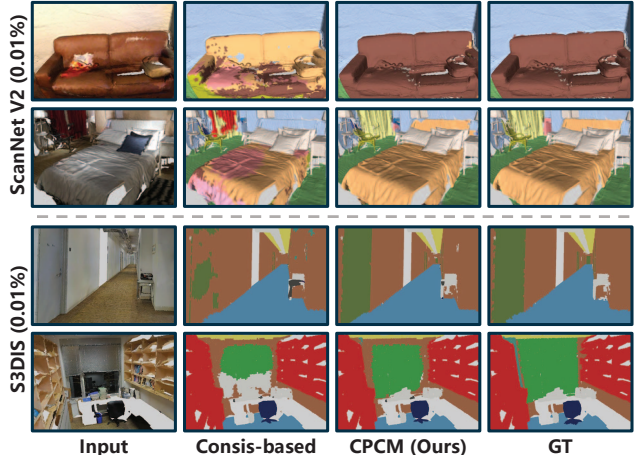


Figure 6: Qualitative comparison between the consis-based method and our CPCM on the ScanNet V2 and S3DIS.⁶

Qualitative results. To intuitively understand our CPCM’s ability to effectively comprehend contextual information, we provide visual comparison results in Figure 6. We first observe that CPCM shows advantages in understanding semantic categories with diverse appearances (sofa, row 1) and covering geometrically large objects (curtain and bed, row 2). Moreover, we recognize that CPCM does an excellent job at distinguishing both geometric and appearance similar categories (door and wall, row 3) and objects with complex structures (window, row 4).

5. Conclusion

In this work, we study the learning of contextual information in the weakly-supervised point-cloud segmentation task which is not well-explored by existing methods. To this end, we proposed CPCM to model the contextual relationship among mass unlabeled points by enforcing the masked feature consistency. We first introduce a region-wise masking strategy to effectively and flexibly mask the point cloud to produce context-to-be-filled data for subsequent learning. Then, we proposed a contextual masked training method to help the model capture contextual information from both limited labeled data and the masked features prediction task. Extensive experiments on the weakly-supervised point cloud segmentation benchmarks show the superior performance of our method. In the future, we will further explore the masked modeling scheme in the weakly-supervised point cloud detection and instance segmentation.

Acknowledgements. This work was partially supported by Key-Area Research and Development Program of Guangdong Province 2019B010155001, National Natural Science Foundation of China (NSFC) (62072190), National Natural Science Foundation of China (NSFC) 61836003 (key project), Program for Guangdong Introducing Innovative and Entrepreneurial Teams 2017ZT07X183.

References

- [1] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *CVPR*, pages 1534–1543, 2016. 6
- [2] Ran Cheng, Ryan Razani, Ehsan Taghavi, Enxu Li, and Bingbing Liu. 2-s3net: Attentive feature fusion with adaptive feature selection for sparse semantic segmentation network. In *CVPR*, pages 12547–12556, 2021. 1, 2, 3
- [3] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, pages 3075–3084, 2019. 2, 3, 6, 7
- [4] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 2, 6
- [5] Jingyu Gong, Jiachen Xu, Xin Tan, Haichuan Song, Yanyun Qu, Yuan Xie, and Lizhuang Ma. Omni-supervised point cloud segmentation via gradual receptive field component reasoning. In *CVPR*, pages 11673–11682, 2021. 6
- [6] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, pages 9224–9232, 2018. 1, 2, 3
- [7] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 2, 3, 5
- [8] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *CVPR*, pages 15587–15597, 2021. 2, 6
- [9] Qingyong Hu, Bo Yang, Guangchi Fang, Yulan Guo, Aleš Leonardis, Niki Trigoni, and Andrew Markham. Sqn: Weakly-supervised semantic segmentation of large-scale 3d point clouds. In *ECCV*, pages 600–619, 2022. 1, 2, 3, 6, 7
- [10] 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. In *CVPR*, pages 11108–11117, 2020. 1, 2, 3, 6
- [11] Li Jiang, Shaoshuai Shi, Zhuotao Tian, Xin Lai, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Guided point contrastive learning for semi-supervised point cloud semantic segmentation. In *ICCV*, pages 6423–6432, 2021. 2
- [12] Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016. 6
- [13] Loïc Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *CVPR*, pages 4558–4567, 2018. 2, 3
- [14] Truc Le and Ye Duan. Pointgrid: A deep network for 3d shape understanding. In *CVPR*, pages 9204–9214, 2018. 1, 2, 3
- [15] Jiaxin Li, Ben M Chen, and Gim Hee Lee. So-net: Self-organizing network for point cloud analysis. In *CVPR*, pages 9397–9406, 2018. 2, 3
- [16] Mengtian Li, Yuan Xie, Yunhang Shen, Bo Ke, Ruizhi Qiao, Bo Ren, Shaohui Lin, and Lizhuang Ma. Hybridcr: Weakly-supervised 3d point cloud semantic segmentation via hybrid contrastive regularization. In *CVPR*, pages 14930–14939, 2022. 1, 2, 3, 4, 6, 7
- [17] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. *NeurIPS*, 31, 2018. 2, 3
- [18] Ying Li, Lingfei Ma, Zilong Zhong, Fei Liu, Michael A Chapman, Dongpu Cao, and Jonathan Li. Deep learning for lidar point clouds in autonomous driving: A review. *TNNLS*, 32(8):3412–3432, 2020. 1
- [19] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. In *ECCV*, pages 657–675, 2022. 3
- [20] Lizhao Liu, Junyi Cao, Minqian Liu, Yong Guo, Qi Chen, and Mingkui Tan. Dynamic extension nets for few-shot semantic segmentation. In *ACM MM*, pages 1441–1449, 2020. 1
- [21] Lizhao Liu, Shangxin Huang, Zhuangwei Zhuang, Ran Yang, Mingkui Tan, and Yaowei Wang. Das: Densely-anchored sampling for deep metric learning. In *ECCV*, pages 399–417, 2022. 3
- [22] Minghua Liu, Yin Zhou, Charles R Qi, Boqing Gong, Hao Su, and Dragomir Anguelov. Less: Label-efficient semantic segmentation for lidar point clouds. In *ECCV*, pages 70–89, 2022. 1, 2, 3
- [23] Yongcheng Liu, Bin Fan, Gaofeng Meng, Jiwen Lu, Shiming Xiang, and Chunhong Pan. Densepoint: Learning densely contextual representation for efficient point cloud processing. In *ICCV*, pages 5239–5248, 2019. 2
- [24] Yongcheng Liu, Bin Fan, Shiming Xiang, and Chunhong Pan. Relation-shape convolutional neural network for point cloud analysis. In *CVPR*, pages 8895–8904, 2019. 2
- [25] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. One thing one click: A self-training approach for weakly supervised 3d semantic segmentation. In *CVPR*, pages 1726–1736, 2021. 2, 3, 7
- [26] Chen Min, Dawei Zhao, Liang Xiao, Yiming Nie, and Bin Dai. Voxel-mae: Masked autoencoders for pre-training large-scale point clouds. *arXiv preprint arXiv:2206.09900*, 2022. 2, 3, 7
- [27] Anh Nguyen and Bac Le. 3d point cloud segmentation: A survey. In *RAM*, pages 225–230, 2013. 1
- [28] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *ECCV*, pages 604–621, 2022. 3
- [29] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *NeurIPS*, pages 8024–8035, 2019. 6
- [30] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, pages 2536–2544, 2016. 3

- [31] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017. 1, 2, 3, 6
- [32] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 30, 2017. 1, 2, 3, 7
- [33] Zhongzheng Ren, Ishan Misra, Alexander G Schwing, and Rohit Girdhar. 3d spatial recognition without spatially labeled 3d. In *CVPR*, pages 13204–13213, 2021. 2, 3, 7
- [34] Dario Rethage, Johanna Wald, Jurgen Sturm, Nassir Navab, and Federico Tombari. Fully-convolutional point networks for large-scale point clouds. In *ECCV*, pages 596–611, 2018. 2, 3
- [35] Gernot Riegler, Ali Osman Ulusoy, and Andreas Geiger. Octnet: Learning deep 3d representations at high resolutions. In *CVPR*, pages 3577–3586, 2017. 2, 3
- [36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, pages 234–241, 2015. 6
- [37] Hang Su, Subhransu Maji, Evangelos Kalogerakis, and Erik Learned-Miller. Multi-view convolutional neural networks for 3d shape recognition. In *ICCV*, pages 945–953, 2015. 2
- [38] A Tarvainen and H Valpola. Weight-averaged consistency targets improve semi-supervised deep learning results. corr abs/1703.01780. *arXiv preprint arXiv:1703.01780*, 1(5), 2017. 6
- [39] Maxim Tatarchenko, Jaesik Park, Vladlen Koltun, and Qian-Yi Zhou. Tangent convolutions for dense prediction in 3d. In *CVPR*, pages 3887–3896, 2018. 1, 2
- [40] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *ICCV*, pages 6411–6420, 2019. 6, 7
- [41] Xu Wang, Jingming He, and Lin Ma. Exploiting local and global structure for point cloud semantic segmentation with contextual point representations. *NeurIPS*, 32, 2019. 2, 3
- [42] Jiacheng Wei, Guosheng Lin, Kim-Hui Yap, Tzu-Yi Hung, and Lihua Xie. Multi-path region mining for weakly supervised 3d semantic segmentation on point clouds. In *CVPR*, pages 4384–4393, 2020. 1, 2, 3, 7
- [43] Jiacheng Wei, Guosheng Lin, Kim-Hui Yap, Fayao Liu, and Tzu-Yi Hung. Dense supervision propagation for weakly supervised semantic segmentation on 3d point clouds. *arXiv preprint arXiv:2107.11267*, 2021. 2, 3
- [44] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *CVPR*, pages 9621–9630, 2019. 2, 3
- [45] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *ECCV*, pages 574–591, 2020. 2, 6
- [46] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, pages 9653–9663, 2022. 3
- [47] Jianyun Xu, Ruixiang Zhang, Jian Dou, Yushi Zhu, Jie Sun, and Shiliang Pu. Rpvnet: A deep and efficient range-point-voxel fusion network for lidar point cloud segmentation. In *ICCV*, pages 16024–16033, 2021. 1, 2, 3
- [48] Xun Xu and Gim Hee Lee. Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In *CVPR*, pages 13706–13715, 2020. 1, 2, 3, 4, 6
- [49] Cheng-Kun Yang, Ji-Jia Wu, Kai-Syun Chen, Yung-Yu Chuang, and Yen-Yu Lin. An ml-derived transformer for weakly supervised point cloud segmentation. In *CVPR*, pages 11830–11839, 2022. 2, 3, 4, 6, 7
- [50] Kun Yi, Yixiao Ge, Xiaotong Li, Shusheng Yang, Dian Li, Jianping Wu, Ying Shan, and Xiaohu Qie. Masked image modeling with denoising contrast. *arXiv preprint arXiv:2205.09616*, 2022. 3
- [51] Dimitris Zermas, Izzat Izzat, and Nikolaos Papanikolopoulos. Fast segmentation of 3d point clouds: A paradigm on lidar data for autonomous vehicle applications. In *ICRA*, pages 5067–5073, 2017. 1
- [52] Yachao Zhang, Zonghao Li, Yuan Xie, Yanyun Qu, Cuihua Li, and Tao Mei. Weakly supervised semantic segmentation for large-scale point cloud. In *AAAI*, pages 3421–3429, 2021. 2, 3, 6, 7
- [53] Yachao Zhang, Yanyun Qu, Yuan Xie, Zonghao Li, Shanshan Zheng, and Cuihua Li. Perturbed self-distillation: Weakly supervised large-scale point cloud semantic segmentation. In *ICCV*, pages 15520–15528, 2021. 1, 2, 3, 4, 6, 7
- [54] Hengshuang Zhao, Li Jiang, Chi-Wing Fu, and Jiaya Jia. Pointweb: Enhancing local neighborhood features for point cloud processing. In *CVPR*, pages 5565–5573, 2019. 2, 3
- [55] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *CVPR*, pages 2921–2929, 2016. 3
- [56] Zhuangwei Zhuang, Rong Li, Kui Jia, Qicheng Wang, Yuanqing Li, and Mingkui Tan. Perception-aware multi-sensor fusion for 3d lidar semantic segmentation. In *ICCV*, pages 16280–16290, 2021. 1, 2
- [57] Hasib Zunair and A Ben Hamza. Masked supervised learning for semantic segmentation. In *BMVC*, 2022. 3