

Point Cloud Semantic Segmentation with Sparse and Inhomogeneous Annotations

Zhiyi Pan^{1,2}, Nan Zhang¹, Wei Gao^{1*}, Shan Liu³, Ge Li¹

¹Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology, Shenzhen Graduate School, Peking University

²Peng Cheng Laboratory

³Media Lab, Tencent

{panzhiyi, zhangnan}@stu.pku.edu.cn, gaowei262@pku.edu.cn, shanl@tencent.com, geli@ece.pku.edu.cn

Abstract

Utilizing uniformly distributed sparse annotations, weakly supervised learning alleviates the heavy reliance on fine-grained annotations in point cloud semantic segmentation tasks. However, few works discuss the inhomogeneity of sparse annotations, albeit it is common in real-world scenarios. Therefore, this work introduces the probability density function into the gradient sampling approximation method to qualitatively analyze the impact of annotation sparsity and inhomogeneity under weakly supervised learning. Based on our analysis, we propose an **Adaptive Annotation Distribution Network** (AADNet) capable of robust learning on arbitrarily distributed sparse annotations. Specifically, we propose a label-aware point cloud downsampling strategy to increase the proportion of annotations involved in the training stage. Furthermore, we design the multiplicative dynamic entropy as the gradient calibration function to mitigate the gradient bias caused by non-uniformly distributed sparse annotations and explicitly reduce the epistemic uncertainty. Without any prior restrictions and additional information, our proposed method achieves comprehensive performance improvements at multiple label rates and different annotation distributions.

Code — <https://github.com/panzhiyi/AADNet>

1 Introduction

Point cloud has become one of the most popular 3D representations due to its flexible ability to represent 3D objects and the wide deployment of capture devices. Point cloud semantic segmentation plays a crucial role in various applications (Hu et al. 2023; Cui et al. 2021) by exploiting the abundant 3D geometric information from point clouds. However, point cloud semantic segmentation suffers from heavy reliance on fine-grained annotated scenes (Xu and Lee 2020). The increase in data dimensionality brings an order of magnitude increase in structural complexity and processing difficulty, further exacerbating the annotation burden. Therefore, weakly supervised point cloud semantic segmentation is a promising topic for deployment and applications.

Existing weakly supervised methods rely on simulating sparse annotations through uniformly distributed sam-




Figure 1: Visual comparison about full annotations, uniform sparse annotations, and non-uniform sparse annotations. Based on the data presented in regarding the number of labeled categories per scene, the distribution disparity between non-uniform sparse annotation and full annotation is pronounced.

pling on a fully labeled dataset. We refer to such annotations as *uniformly distributed sparse annotations*. The simulated dataset maintains distribution consistency between sparse and full annotations through uniformly random sampling (Xu and Lee 2020) or “One-Thing-One-Click” (Liu, Qi, and Fu 2021). However, it places strict constraints on the annotation distribution.

Non-uniformly distributed sparse annotations are inevitable in real-world scenarios, which can be illustrated by the following processes in weakly supervised learning:

- The projection relationship between the 3D point cloud and the 2D annotation plane causes the non-uniformly distributed points on the annotation plane, directly leading to inhomogeneous annotations. Moreover, to improve annotation efficiency, annotators are more likely to annotate categories that are easier to label, which exacerbates the inhomogeneity.
- Supervoxels are typically used to expand sparse annotation information, during which uniformly distributed sparse annotated points are transformed into non-uniformly distributed sparse annotated supervoxels.
- Self-training in weakly supervised learning treats high-confidence predictions as pseudo-labels for subsequent network learning. Since high-confidence points are usu-

*Wei Gao is the corresponding author.

ally concentrated in local regions with significant patterns, the pseudo-labels are also unevenly distributed.

As illustrated in Fig. 1, there is a significant distribution gap between non-uniformly distributed sparse annotations and full annotations. Without the uniform sampling prior, the inhomogeneous annotations impede weakly supervised learning. The current approaches fail to consider the essential contribution of sparse labeling inhomogeneity to point cloud segmentation. Therefore, our work extends the annotation requirements from uniformly distributed sparse annotations to arbitrarily distributed sparse annotations.

To investigate the impact of annotations, we introduce the probability density function into the gradient sampling approximation analysis (Xu and Lee 2020). According to the Central Limit Theorem, the gradient discrepancy between weak supervision and full supervision follows a normal distribution, where the sparsity and inhomogeneity of labeling affect its mean and variance, respectively. Based on this analysis, we propose an **Adaptive Annotation Distribution Network** (AADNet) to learning from arbitrarily distributed annotations, which is composed of label-aware downsampling strategy and multiplicative dynamic entropy.

Downsampling is essential for point cloud scenes due to computational limitations. However, random downsampling (Hu et al. 2020) directly borrowed from the full supervision would miss valuable annotations entailed for training. Conversely, sampling only labeled points would severely sacrifice the structural information from point cloud resampling. Therefore, we propose the label-aware point cloud downsampling strategy (LaDS) to increase the proportion of labeled points involved in the training stage and retain rich structured information to a great extent.

We propose the multiplicative dynamic entropy with asynchronous training (MDE-AT) to accommodate the annotation inhomogeneity. We prove that an ideal calibration function for correcting the gradient bias should be inversely proportional to the probability density. However, the probability density of the annotation distribution is agnostic. Therefore, we utilize multiplicative dynamic entropy as an alternative calibration function. Additionally, to improve the epistemic certainty (Kendall and Gal 2017), we employ asynchronous training that iteratively imposes entropy loss and partial cross-entropy loss with the calibration function.

The contributions of this work can be summarized in the following three aspects:

- To the best of our knowledge, we are the first to discuss the inhomogeneity of sparse annotations in weakly supervised learning. The quantitative impact of non-uniformly distributed sparse annotations on network learning is revealed through gradient sampling approximation analysis.
- Based on the analysis, we propose an Adaptive Annotation Distribution Network that can handle sparsely annotated data with arbitrary distributions.
- The effectiveness of AADNet is comprehensively validated on multiple annotation settings and datasets.

2 Related Work

For weakly supervised point cloud semantic segmentation, researchers have utilized well-established techniques in 2D vision, such as class activation maps (Zhou et al. 2016), perturbation consistency (Pan et al. 2021), label propagation (Iscen et al. 2019), and contrastive learning (Chen et al. 2020). These techniques leverage rich hypothetical priors, processing paradigms, or additional supervision to enable weakly supervised point cloud semantic segmentation. Rather than relying on a single technique, most methods build their effective weak supervision frameworks by integrating multiple techniques.

Class Activation Maps. In weakly supervised point cloud semantic segmentation, Class Activation Maps (CAMs) provide category localization information from both scene-level and sub-scene-level annotations. To further improve the semantic localization capability of CAMs, MPRM (Wei et al. 2020) and Song *et al.* (Song et al. 2022) act CAMs and geometric projection to obtain point-level pseudo-labels. MILTrans (Yang et al. 2022) utilizes an adaptive global weighted pooling mechanism on CAMs to mitigate the negative effects of irrelevant classes and noisy points.

Perturbation Consistency. Consistency between the corresponding features of the original and perturbed point clouds provides additional supervision signals. Xu *et al.* (Xu and Lee 2020) and PSD (Zhang et al. 2021b) apply perturbations such as rotation and mirror flipping to the point cloud scene, while MILTrans (Yang et al. 2022) maintains consistency on the pair point clouds before and after downsampling. DAT (Wu et al. 2022) and HybridCR (Li et al. 2022) dynamically construct point cloud replicas using adaptive gradient and embedding network learning, respectively. RPSC (Lan et al. 2023) and PointMatch (Wu et al. 2023) leverage pseudo-labels to convey consistent information.

Label Propagation. The label propagation generates high-quality pseudo-labels in self-training. For instance, OTOC (Liu, Qi, and Fu 2021) and OTOC++ (Liu, Qi, and Fu 2023) propose RelationNets to accurately measure the similarity between 3D graph nodes and propagate labeling information. Zhang *et al.* (Zhang et al. 2021a) develops a sparse label propagation algorithm guided by the category prototype (Snell, Swersky, and Zemel 2017).

Contrastive Learning. Contrastive learning selects anchor points and imposes restrictions on corresponding positive point pairs and negative point pairs. HybridCR (Li et al. 2022) constructs contrastive learning on transformed point cloud pairs, local geometry pairs, and category prototype pairs, while MILTrans (Yang et al. 2022) performs contrastive loss between category pairs in the scene. Besides, pre-training methods with contrastive learning (Xie et al. 2020; Hou et al. 2021) can also achieve semantic segmentation tasks under few annotations.

Others. SQN (Hu et al. 2022) collects a group of hierarchical representations within the locally labeled neighborhoods through interpolation. LessIsMore (Pan et al. 2024c) proposes a sparse labeling recommendation framework for weakly supervised learning. REAL (Kweon, Kim, and Yoon 2024) introduces 2D information from SAM (Kirillov et al. 2023) to weak supervision. DGNet (Pan et al. 2024a) explic-

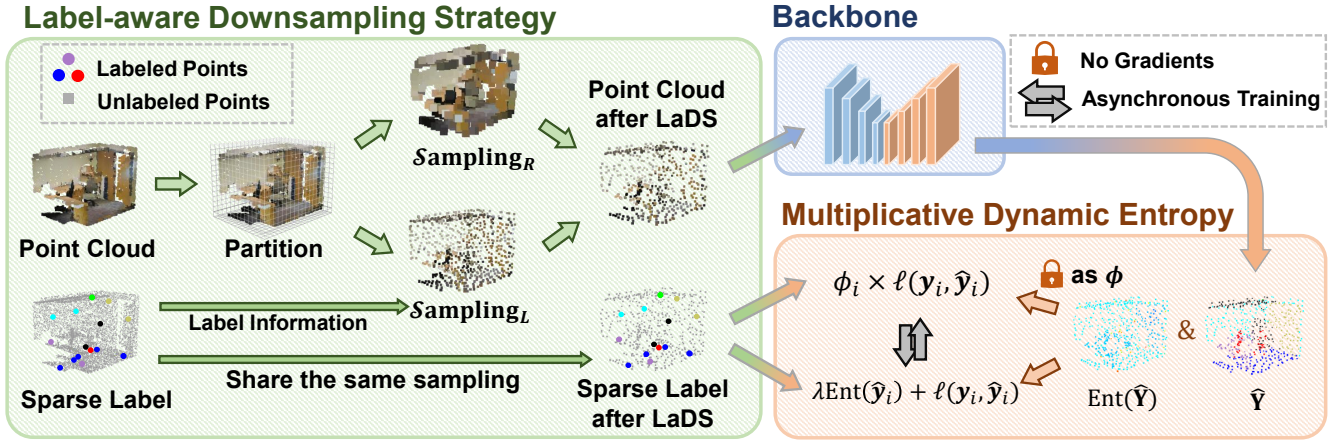


Figure 2: The framework of AADNet consists of two modules: (1) Label-aware downsampling strategy (LaDS) to boost the training annotation rate and maintain structural diversity. (2) Multiplicative dynamic entropy with asynchronous training (MDE-AT) to correct for gradient bias due to sparse labeling distributions as well as to explicitly improve epistemic certainty.

itly aligns the feature space and prior distribution to provide additional supervised signals.

The above methods are verified under the uniformly distributed sparse annotations. However, uniformly distributed sampling is elusive to achieve. Consequently, we propose a point cloud semantic segmentation method with adaptive annotation distribution for arbitrary sparse annotation scenarios.

3 Method

In Sec. 3.1, we give a brief overview of weakly supervised point cloud semantic segmentation. The gradient sampling approximation analysis with probability density function is introduced to investigate the annotation distribution in Sec. 3.2. Based on the analysis, we propose the Adaptive Annotation Distribution Network, which is composed of a label-aware downsampling strategy (Sec. 3.3) and multiplicative dynamic entropy (Sec. 3.4). The framework of AADNet is illustrated in Fig. 2.

3.1 Preliminaries

A point cloud scene in the training set for semantic segmentation can be denoted as $\mathbf{P} = \{\mathbf{X}, \mathbf{Y}\}$, where \mathbf{X} and \mathbf{Y} represent point set and its corresponding label, respectively. In detail, $\mathbf{X} = [\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N]$ and $\mathbf{x}_i \in \mathbb{R}^{1 \times F}$, in which N denotes the number of points and F is the dimension number of initial feature on each point \mathbf{x}_i . Without losing generality, $\mathbf{Y} = [\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_M]$ and $\mathbf{y}_i \in \{0, 1\}^{1 \times C}$, where M and C indicate the number of labeled points and categories, respectively. When $N \gg M$, the task is termed as a weakly supervised point cloud semantic segmentation task, and the label rate is $\frac{M}{N}$. The partial cross-entropy loss, which is commonly used in weak supervision, is defined as:

$$\mathcal{L}_p = \frac{1}{M} \sum_{i=1}^M \ell(\mathbf{y}_i, \hat{\mathbf{y}}_i) = \frac{1}{M} \sum_{i=1}^M \mathbf{y}_i \log(\hat{\mathbf{y}}_i^\top), \quad (1)$$

where $\hat{\mathbf{y}}_i \in [0, 1]^{1 \times C}$ denotes the prediction probability at point \mathbf{x}_i .

3.2 Gradient Sampling Approximation Analysis

For two networks with the same structure and initialization parameters, it is assumed that the closer the gradients are, the more likely they will converge to the same prediction results for the same input (Xu and Lee 2020). The gradients of cross-entropy $\nabla \mathcal{L}$ and partial cross-entropy $\nabla \mathcal{L}_p$ are formulated as:

$$\begin{aligned} \nabla \mathcal{L} &= \frac{1}{N} \sum_{i=1}^N g(\hat{\mathbf{y}}_i), \\ \nabla \mathcal{L}_p &= \frac{1}{M} \sum_{i=1}^M g(\hat{\mathbf{y}}_i), \end{aligned} \quad (2)$$

where $g(\hat{\mathbf{y}}_i) = \nabla \ell(\mathbf{y}_i, \hat{\mathbf{y}}_i)$, denotes the loss gradient at \mathbf{x}_i . Weakly supervised learning can be viewed as fully supervised learning with M times sampling of gradients, and the average gradient value of M times sampling is equal to $\nabla \mathcal{L}_p$. Assuming that $g(\hat{\mathbf{y}}_i)$ satisfies the independent identical distribution, the uniformly distributed sampled $s(g(\hat{\mathbf{y}}_i))$ follows the distribution $\mathcal{D}_u(E_u, V_u)$. The expectation $E_u = \mathbb{E}_{\mathbf{x}_i \sim \mathcal{D}}[p_i g(\hat{\mathbf{y}}_i)] = \sum_i \frac{1}{N} g(\hat{\mathbf{y}}_i) = \nabla \mathcal{L}$, where \mathcal{D} is the distribution of points and $p_i = \frac{1}{N}$ denotes the probability density function of uniformly distributed sampling. According to the Central Limit Theorem, the following convergent distribution can be obtained under uniformly distributed sampling:

$$(\nabla \mathcal{L}_p - \nabla \mathcal{L}) \sim \mathcal{N}(0, \frac{V_u}{M}). \quad (3)$$

Eq. 3 illustrates that the gradient difference between full annotations and uniformly sparse annotations obeys a normal distribution with 0 expectation and $\frac{V_u}{M}$ variance, which indicates that the average gradient of uniform sparse annotations is more likely to be similar to that of full annotations when the label rate becomes larger.

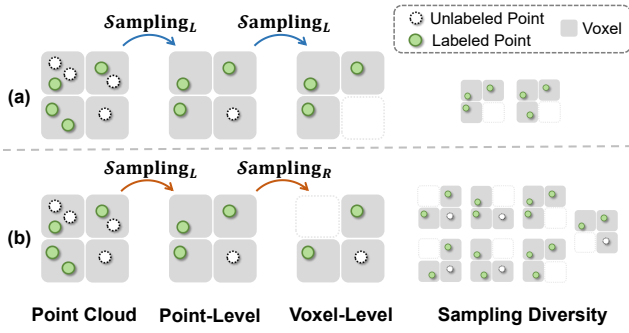


Figure 3: Compared with (a) voxel-level label-first downsampling, (b) voxel-level random downsampling retains diverse scene structural information.

Based on the previous assumption, the prediction results of the model supervised by uniform sparse annotations at an appropriate label rate are consistent with those of the fully supervised semantic segmentation model. However, this corollary is not universal for the non-uniformly distributed sparse annotations. We assume that non-uniformly sampled gradient $s_n(g(\hat{\mathbf{y}}_i))$ follows the distribution $\mathcal{D}_n(E_n, V_n)$. The expectation $E_n = \mathbb{E}_{\mathbf{x}_i \sim \mathcal{D}}[p'_i g(\hat{\mathbf{y}}_i)] = \sum_i p'_i g(\hat{\mathbf{y}}_i) = \nabla \mathcal{L} + \Delta$, where p'_i denotes the probability density function of the non-uniformly sampling distribution and the bias $\Delta = \sum_i (p'_i - p_i) g(\hat{\mathbf{y}}_i)$. Similarly, according to the Central Limit Theorem, we obtain the following convergent distribution under non-uniformly distributed sampling:

$$(\nabla \mathcal{L}_p - \nabla \mathcal{L}) \sim \mathcal{N}(\Delta, \frac{V_n}{M}). \quad (4)$$

Eq. 4 demonstrates that in the case of non-uniformly distributed sparse annotations, the average gradient still differs by Δ even at large label rates. We conclude that **the performance of weak supervision is not only affected by the annotation sparsity but also the annotation inhomogeneity**. Without any priors, uniformly distributed sparse labeling is the optimal sparse labeling.

3.3 Label-aware Downsampling Strategy

The downsampling strategy determines the upper bound of the receptive field and the diversity of the structural information of the scene (Hu et al. 2020). However, the random downsampling strategy used by full supervision is not applicable to sparse labeling scenarios. The random downsampling treats labeled and unlabeled points equally, significantly weakening the semantic supervision information of the sampled point cloud. An intuitive alternative is preferentially selecting all labeled points. Although it can greatly preserve labeled points after downsampling, it compromises the diversity of structural information when sampling the point cloud repeatedly. Therefore, a label-aware downsampling strategy is proposed to preserve labeled points and abundant structural topology.

The label-aware point cloud downsampling strategy is shown in Fig. 2. First, we partition the entire point cloud scene by voxels, *i.e.*, $\mathcal{V} = \{\mathbf{v}_j\}$ and $\mathbf{v}_{j_1} \cap \mathbf{v}_{j_2} = \emptyset$ when

$j_1 \neq j_2$. Based on the voxel partition, the downsampling on the point cloud decomposes into point-level label-first downsampling Sampling_L and voxel-level random downsampling Sampling_R .

- 1) **Point level label-first downsampling Sampling_L** . For the voxel that contains labeled points, one labeled point is randomly selected among the labeled points within the voxel, and for the voxel that does not contain any labeled point, one point is randomly selected within the voxel. The loss of structural diversity with point-level label-first sampling is almost negligible since the sampling is over a small range of voxels.
- 2) **Voxel-level random downsampling Sampling_R** . At the voxel level, we retain the conventional random downsampling strategy. In contrast to voxel-level label-first downsampling, voxel-level random downsampling sacrifices a small percentage of the labeled points but retains diverse scene structural information when sampling repeatedly in different training epochs. A comparison of sampling strategies with voxel-level label-first downsampling and voxel-level random downsampling is visualized in Fig. 3.

The points \mathbf{P}' sampled by the label-aware point cloud downsampling strategy can be defined as:

$$\mathbf{P}' = \text{Sampling}_R \left(\left\{ \text{Sampling}_L(\mathbf{v}_j, 1) \mid \mathbf{v}_j \in \mathcal{V} \right\}, K \right), \quad (5)$$

where K is a predefined number of sampling. $\text{Sampling}(\mathbf{v}, t)$ denotes sample t times from the set \mathbf{v} .

It is worth noting that the label-aware downsampling strategy targets annotation sparsity and imposes no restrictions on annotation inhomogeneity.

3.4 Multiplicative Dynamic Entropy

Based on the methodological analysis in Sec. 3.2, we impose a gradient calibration function ϕ on \mathcal{L}_p to counteract the gradient expectation bias caused by the non-uniformly distributed sparse annotations. Ideally, using multiplicative calibration function $\phi_i = \frac{p_i}{p'_i}$, the expectation of the corrected gradient satisfies:

$$E_n = \mathbb{E}_{\mathbf{x}_i \sim \mathcal{D}}[p'_i \phi_i g(\hat{\mathbf{y}}_i)] = \mathbb{E}_{\mathbf{x}_i \sim \mathcal{D}}[p_i g(\hat{\mathbf{y}}_i)] = \nabla \mathcal{L}. \quad (6)$$

As shown in Fig. 4, the gradient bias Δ is corrected by ϕ . Compared to the ideal additive calibration function $\psi_i = (p'_i - p_i) g(\hat{\mathbf{y}}_i)$, the multiplicative calibration function ϕ is gradient-independent and easier to estimate.

Nevertheless, the probability density function p'_i of non-uniformly distributed sampling in the ideal ϕ is agnostic. Taking the negative correlation between labeling difficulty and labeling probability p'_i as a guideline, we utilize the entropy function as the dynamic calibration function, *i.e.*, $\phi_i = \text{Ent}(\hat{\mathbf{y}}_i) = \hat{\mathbf{y}}_i \log(\hat{\mathbf{y}}_i^T)$. Shannon entropy as an alternative multiplicative calibration function offers two major benefits:




Figure 4: Schematic diagram of gradient calibration by MDE.

- 1) **Adaptive.** Entropy is a dynamic function related to sparsely labeled scenes in the training stage, which can adaptively correct for different point cloud scenes.
- 2) **Representativity.** The classification difficulty of points reflects the annotation distribution. Since entropy possesses the ability to reflect classification difficulty, it also represents annotation distribution.

It is worth noting that although ϕ is a function on predictions $\{\hat{y}_i\}$, it does not generate gradients during training, but only serves to weight the partial cross-entropy loss dynamically.

To make entropy more accurate and reduce the epistemic uncertainty of the network, we conversely constrain its learning as a loss term. Entropy loss and partial cross-entropy loss are organized in an asynchronous training manner. The reason for introducing asynchronous training is that synchronous training will change the optimization objective of the network and prevent partial cross-entropy loss from converging. The total loss function \mathcal{L}_{AAD}^k at the k -th training epoch can be defined as:

$$\mathcal{L}_{AAD}^k = \mathbb{1}(k \bmod 2\tau \geq \tau) \frac{1}{M} \sum_{i=1}^M \left[\phi_i \times \ell(\mathbf{y}_i, \hat{\mathbf{y}}_i) \right] + \mathbb{1}(k \bmod 2\tau < \tau) \frac{1}{M} \sum_{i=1}^M \left[\lambda \text{Ent}(\hat{\mathbf{y}}_i) + \ell(\mathbf{y}_i, \hat{\mathbf{y}}_i) \right], \quad (7)$$

where $\mathbb{1}(\cdot)$ denotes the indicator function, λ is a predefined balance weight, and τ denotes the step interval of asynchronous training. To learn efficiently, AADNet still imposes partial cross-entropy loss while training via $\text{Ent}(\hat{\mathbf{y}}_i)$.

The gradient calibration function targets annotation inhomogeneity and has no restriction on annotation sparsity. It is worth noting that even under uniformly sampled sparse annotations, due to extreme sparsity, the distribution tends to exhibit inhomogeneity, allowing the gradient calibration function to play a positive role.

4 Labeling Implementation

The quality of weakly supervised annotations depends on the sparsity and inhomogeneity of the annotations. Annotation sparsity can be easily controlled by adjusting the label rate.¹ In this section, we introduce how to control the degree

¹Compared with "One-Thing-One-Click", the uniformly random sampling allows continuous adjustment of the label rate.




Figure 5: The sparse annotations with different G .

of non-uniformity through simulation, which is essential for a comprehensive evaluation.

Consider the most uneven case of sparse annotations, where all annotations are concentrated in a highly similar local region, effectively meaning that only one point cloud cluster in the scene is labeled. As the number of labeled point cloud clusters increases (with the scale of each cluster decreasing to maintain the same label rate), the homogeneity of the sparse annotations improves. When the number of annotated point cloud clusters is further increased to M , with each cluster containing only one point, the sparse annotation becomes equivalent to uniform sampling across annotations of all points. Therefore, controlling the number of point cloud clusters can effectively simulate the annotation inhomogeneity.

In our annotation implementation, under the condition of a label rate of $\frac{M}{N}$, we first define the number of point cloud clusters G , then randomly sample G points in the scene as the cluster centers. Based on the initial features of the points, we then query $\lfloor \frac{M}{G} \rfloor$ points for each cluster center as the annotated point cloud clusters. Fig. 5 illustrates the simulated sparse annotations with different cluster numbers.

5 Experiments

5.1 Experiment Setting

Dataset. S3DIS (Armeni et al. 2016) contains 271 rooms across six areas with 13 semantic categories. We train our model on Areas 1, 2, 3, 4, and 6, and evaluate the segmentation performance in Area 5. ScanNetV2 (Dai et al. 2017) consists of 1,513 scanned scenes obtained from 707 different indoor environments and provides 21 semantic categories for each point. We utilize 1,201 scenes for training and 312 scenes for validation, according to the official split.

Annotation setting. To compare with alternatives, the performance at label rates of 0.01% on S3DIS and 20 points per scene (label rate is about 0.014%) on ScanNetV2 are reported. For comprehensive validation, we set G to 1, 10, 20, and M to regulate the inhomogeneity of sparse labeling.

Implementation. We take RandLA-Net (Hu et al. 2020) and PointNeXt-L (Qian et al. 2022) as the backbones to construct AADNet. Unless otherwise noted, AADNet is trained with default settings. To prevent the entropy loss from misleading the network during the early stage of training (Pan et al. 2024b), we set the start epoch of asynchronous training to 50. The weight $\lambda = 0.01$ and step interval $\tau = 5$ in Eq. 7.




Figure 6: Comparisons between Baseline (PointNeXt) and AADNet with different label rates and numbers of labeled clusters.

Method	M/N	S3DIS	ScanNetV2
SQN (Hu et al. 2022)	10%	38.4	41.1
EDRA (Tang et al. 2024)	10%	43.3	46.2
RandLA-Net (Hu et al. 2020)	10%	40.9	42.9
AADNet (RandLA-Net)	10%	50.5	53.0
DCL (Yao et al. 2023)	10%	52.9	55.1
CPCM (Liu et al. 2023)	10%	54.9	59.5
PointNeXt (Qian et al. 2022)	10%	54.1	56.0
AADNet (PointNeXt)	10%	56.0	61.2
SQN (Hu et al. 2022)	1%	30.7	32.6
EDRA (Tang et al. 2024)	1%	34.6	35.7
RandLA-Net (Hu et al. 2020)	1%	34.1	36.4
AADNet (RandLA-Net)	1%	37.7	41.1
DCL (Yao et al. 2023)	1%	35.6	34.7
CPCM (Liu et al. 2023)	1%	38.2	41.7
PointNeXt (Qian et al. 2022)	1%	35.9	37.0
AADNet (PointNeXt)	1%	39.0	43.4

Table 1: Quantitative comparisons with inhomogeneous annotations and $G = 1$.

Our models are trained with one NVIDIA V100 GPU on S3DIS and eight NVIDIA TESLA T4 GPUs on ScanNetV2.

5.2 Comparison Results

In Fig. 6, we compare the performance of the Baseline with AADNet under different sparsity (determined by the label rate) and inhomogeneity (determined by the number of labeled clusters) annotations. AADNet outperforms the Baseline in all 16 sparse labeling scenarios with different levels of difficulty. The superiority of AADNet is more significant at low labeling ratios since the negative effects caused by inhomogeneity are more severe at a lower label rate.

To compare with other alternatives in the non-uniform distribution sparse scenario, we reproduce four open-sourced methods. We choose the most inhomogeneous sparse labeling setting, *i.e.*, $G = 1$, to demonstrate the ability to handle inhomogeneous annotations. Tab. 1 illustrates that non-uniform annotation poses a significant hindrance to

Method	S3DIS		ScanNetV2	
	M/N	mIoU	M/N	mIoU
SQN (Hu et al. 2022)	0.01%	45.3	-	-
CPCM (Liu et al. 2023)	0.01%	49.3	0.01%	52.2
EDRA (Tang et al. 2024)	0.02%	48.4	20pts	57.0
RPSC (Lan et al. 2023)	0.04%	64.0	50pts	58.7
PointMatch (Wu et al. 2023)	0.01%	59.9	20pts	62.4
MILTrans (Yang et al. 2022)	0.02%	51.4	20pts	54.4
RandLA-Net (Hu et al. 2020)	0.01%	38.6	20pts	50.7
AADNet (RandLA-Net)	0.01%	55.6	20pts	55.1
PointNeXt (Qian et al. 2022)	0.01%	44.7	20pts	54.6
AADNet (PointNeXt)	0.01%	60.8	20pts	62.5

Table 2: Quantitative comparisons under uniformly distributed sparse annotations.

network learning. The comparison results demonstrate the effectiveness of AADNet with inhomogeneous annotations.

Even though AADNet targets arbitrarily distributed labeling, it outperforms previous methods under uniformly distributed labeling in Tab. 2.

In addition, we present a qualitative comparison of the method with the baseline in Fig. 7. The baseline has a significant performance degradation on categories with variable forms, such as the bookcase with miscellaneous books and clutter on tables, and categories that are easily confused, such as wall, column, and board. Conversely, AADNet shows more complete predictions for head categories and more accurate predictions for long-tail categories.

5.3 Ablation Study

We perform extensive ablation experiments for the proposed method on S3DIS Area 5 to verify the effectiveness of each module. Baseline denotes the point cloud semantic segmentation network PointNeXt-L (Qian et al. 2022) trained with the official default settings. The cluster number G of inhomogeneous annotations in the ablation study is set to 1.

Ablation on LaDS. Tab. 3 demonstrates the segmentation performance with or without the label-first sampling. It can




Figure 7: Qualitative comparisons on S3DIS under various label settings. The label settings are 10% non-uniform (first row) and 0.01% uniform sparse annotations (second row).

Sampling Strategy		Homo. ($G = M$)		Inhomo. ($G = 1$)	
Point-level	Voxel-level	0.01%	0.1%	1%	10%
Sampling _R	Sampling _R	44.7	64.3	35.9	54.1
Sampling _R	Sampling _L	45.6	65.0	36.4	54.0
Sampling _L	Sampling _L	58.1	66.6	37.0	53.8
Sampling _L	Sampling _R	58.4	67.0	37.8	54.8

Table 3: Ablation study for LaDS on S3DIS Area 5.

be observed that point-level label-first downsampling and voxel-level label-first downsampling promote the segmentation performance, respectively. However, voxel-level label-first downsampling may destroy the sampling diversity of re-sampling and therefore ignores part of the geometric topology. Hence, using both voxel-level and point-level label-first downsampling is suboptimal. The optimal performance is achieved by only using point-level label-first sampling (*i.e.*, LaDS), which improves the label rate while avoiding compromising sampling diversity.

Ablation on MDE-AT. We validate several feasible gradient calibration functions in Tab. 4. FixedW sets the inverse of the local density of points as the gradient calibration function. ADE denotes the Additive Dynamic Entropy. MDE-ST and MDE-AT denote the multiplicative dynamic entropy with simultaneous training and asynchronous training, respectively. By observing the metrics, we have the following inferences: Fixed weights cannot adjust the weights dynamically for different point clouds. The additive calibration

Setting	Homo. ($G = M$)		Inhomo. ($G = 1$)	
	0.01%	0.1%	1%	10%
Baseline	44.7	64.3	35.9	54.1
+ FixedW	43.3	60.5	35.6	55.1
+ ADE	48.2	64.7	36.3	55.0
+ MDE-ST	-	-	-	-
+ MDE-AT	49.7	65.0	36.8	56.0

Table 4: Ablation study for MDE-AT on S3DIS Area 5. '-' indicates that partial cross-entropy loss does not converge.

Method	Labeling	Supervoxels	Pseudo-labels
Baseline	44.7%	52.1%	53.0%
AADNet	60.8% (+16.1%)	61.7% (+9.7%)	62.5% (+9.5%)

Table 5: Validation on three inhomogeneously labeled scenarios at 0.01% label rate.

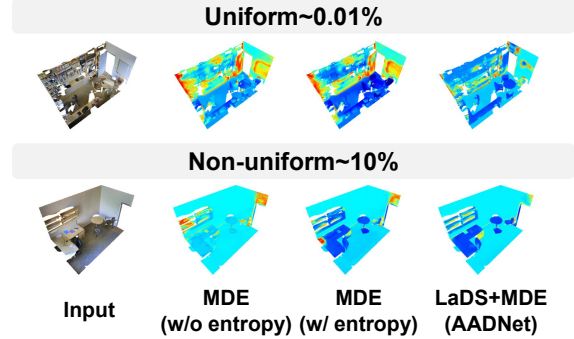


Figure 8: Shannon entropy visualization on predictions.

function is difficult to learn compared with the multiplicative one. Simultaneous training changes the learning objective of the partial cross-entropy loss function, leading to the partial cross-entropy loss function failing to converge.

5.4 Further Analysis

Beyond simulation. The simulated inhomogeneous annotations validate the effectiveness of labeling scenarios. In Tab. 5, we further report the performance under the other scenarios mentioned in the Introduction. With supervoxel or pseudo-label strategies, AADNet still has significant improvements.

Epistemic uncertainty. Fig. 8 visualize the entropy of predicted class probability. High-entropy areas are located in objects with variable forms and objects that are easily confused. By explicitly learning the entropy of the labeled points with asynchronous training, the entropy of these objects is significantly reduced, resulting in a reliable and robust semantic segmentation network. Finally, less epistemic uncertainty is achieved by combining LaDS and MDE-AT.

6 Conclusion

In this paper, we provide benchmarks for semantic segmentation with various distribution sparse annotations. By introducing the probability density function into the gradient sampling approximation analysis, we reveal the impact of sparse annotation distributions in the manual labeling process. Based on gradient analysis, we propose an adaptive annotation distribution network, consisting of the label-aware downsampling strategy and the multiplicative dynamic entropy for asynchronous training. AADNet achieves robust learning in different labeling scenarios. We expect that our work can provide novel research insights for the weakly supervised point cloud semantic segmentation community.

Acknowledgements

This work was supported in part by National Science and Technology Major Project (2024ZD01NL00101), and in part by Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology (2024B1212010006).

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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1534–1543.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PMLR.
- Cui, Y.; Chen, R.; Chu, W.; Chen, L.; Tian, D.; Li, Y.; and Cao, D. 2021. Deep learning for image and point cloud fusion in autonomous driving: A review. *IEEE Transactions on Intelligent Transportation Systems*, 23(2): 722–739.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5828–5839.
- Hou, J.; Graham, B.; Nießner, M.; and Xie, S. 2021. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 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 *European conference on computer vision*, 600–619. Springer.
- 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 11108–11117.
- Hu, Y.; Yang, J.; Chen, L.; Li, K.; Sima, C.; Zhu, X.; Chai, S.; Du, S.; Lin, T.; Wang, W.; et al. 2023. Planning-oriented autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17853–17862.
- Iscen, A.; Tolas, G.; Avrithis, Y.; and Chum, O. 2019. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5070–5079.
- Kendall, A.; and Gal, Y. 2017. What uncertainties do we need in bayesian deep learning for computer vision? *Advances in neural information processing systems*, 30.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Kweon, H.; Kim, J.; and Yoon, K.-J. 2024. Weakly Supervised Point Cloud Semantic Segmentation via Artificial Oracle. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3721–3731.
- Lan, Y.; Zhang, Y.; Qu, Y.; Wang, C.; Li, C.; Cai, J.; Xie, Y.; and Wu, Z. 2023. Weakly supervised 3D segmentation via receptive-driven pseudo label consistency and structural consistency. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 1222–1230.
- 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 14930–14939.
- 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18413–18422.
- Liu, Z.; Qi, X.; and Fu, C.-W. 2021. One thing one click: A self-training approach for weakly supervised 3d semantic segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1726–1736.
- Liu, Z.; Qi, X.; and Fu, C.-W. 2023. One Thing One Click++: Self-Training for Weakly Supervised 3D Scene Understanding. *arXiv preprint arXiv:2303.14727*.
- Pan, Z.; Gao, W.; Liu, S.; and Li, G. 2024a. Distribution Guidance Network for Weakly Supervised Point Cloud Semantic Segmentation. *Advances in neural information processing systems*.
- Pan, Z.; Jiang, P.; Wang, Y.; Tu, C.; and Cohn, A. G. 2021. Scribble-supervised semantic segmentation by uncertainty reduction on neural representation and self-supervision on neural eigenspace. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 7416–7425.
- Pan, Z.; Sun, H.; Jiang, P.; Li, G.; Tu, C.; and Ling, H. 2024b. CC4S: Encouraging Certainty and Consistency in Scribble-Supervised Semantic Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Pan, Z.; Zhang, N.; Gao, W.; Liu, S.; and Li, G. 2024c. Less is more: label recommendation for weakly supervised point cloud semantic segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4397–4405.
- Qian, G.; Li, Y.; Peng, H.; Mai, J.; Hammoud, H.; Elhoseiny, M.; and Ghanem, B. 2022. Pointnext: Revisiting pointnet++ with improved training and scaling strategies. *Advances in neural information processing systems*, 35: 23192–23204.
- Snell, J.; Swersky, K.; and Zemel, R. 2017. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30.
- Song, Y.; Sun, Z.; Li, Q.; Wu, Y.; Sun, Y.; and Luo, S. 2022. Learning indoor point cloud semantic segmentation from image-level labels. *The Visual Computer*, 38(9-10): 3253–3265.

- Tang, L.; Chen, Z.; Zhao, S.; Wang, C.; and Tao, D. 2024. All points matter: entropy-regularized distribution alignment for weakly-supervised 3D segmentation. *Advances in Neural Information Processing Systems*, 36.
- 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4384–4393.
- Wu, Y.; Yan, Z.; Cai, S.; Li, G.; Han, X.; and Cui, S. 2023. PointMatch: A consistency training framework for weakly supervised semantic segmentation of 3D point clouds. *Computers & Graphics*, 116: 427–436.
- Wu, Z.; Wu, Y.; Lin, G.; Cai, J.; and Qian, C. 2022. Dual adaptive transformations for weakly supervised point cloud segmentation. In *European conference on computer vision*, 78–96. Springer.
- Xie, S.; Gu, J.; Guo, D.; Qi, C. R.; Guibas, L.; and Litany, O. 2020. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European conference on computer vision*, 574–591. Springer.
- Xu, X.; and Lee, G. H. 2020. Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 13706–13715.
- Yang, C.-K.; Wu, J.-J.; Chen, K.-S.; Chuang, Y.-Y.; and Lin, Y.-Y. 2022. An mil-derived transformer for weakly supervised point cloud segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 11830–11839.
- Yao, B.; Xiao, H.; Zhuang, J.; and Peng, C. 2023. Weakly Supervised Learning for Point Cloud Semantic Segmentation with Dual Teacher. *IEEE Robotics and Automation Letters*.
- Zhang, Y.; Li, Z.; Xie, Y.; Qu, Y.; Li, C.; and Mei, T. 2021a. Weakly supervised semantic segmentation for large-scale point cloud. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15520–15528.
- Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2921–2929.