

AdaCo: Overcoming Visual Foundation Model Noise in 3D Semantic Segmentation via Adaptive Label Correction

Pufan Zou^{1,2*}, Shijia Zhao^{1,2*}, Weijie Huang^{1,2}, Qiming Xia^{1,2},
Chenglu Wen^{1,2†}, Wei Li^{3†}, Cheng Wang^{1,2}

¹Fujian Key Laboratory of Sensing and Computing for Smart Cities, Xiamen University, China,

²Key Laboratory of Multimedia Trusted Perception and Efficient Computing,
Ministry of Education of China, Xiamen University, China,

³Inceptio

{zoupufan, zhaoshijia4567, weijie, xiaqiming}@stu.xmu.edu.cn,
clwen@xmu.edu.cn, liweimcc@gmail.com, cwang@xmu.edu.cn

Abstract

Recently, Visual Foundation Models (VFMs) have shown a remarkable generalization performance in 3D perception tasks. However, their effectiveness in large-scale outdoor datasets remains constrained by the scarcity of accurate supervision signals, the extensive noise caused by variable outdoor conditions, and the abundance of unknown objects. In this work, we propose a novel label-free learning method, **Adaptive Label Correction (AdaCo)**, for 3D semantic segmentation. AdaCo first introduces the Cross-modal Label Generation Module (**CLGM**), providing cross-modal supervision with the formidable interpretive capabilities of the VFMs. Subsequently, AdaCo incorporates the Adaptive Noise Corrector (**ANC**), updating and adjusting the noisy samples within this supervision iteratively during training. Moreover, we develop an Adaptive Robust Loss (**ARL**) function to modulate each sample’s sensitivity to noisy supervision, preventing potential underfitting issues associated with robust loss. Our proposed AdaCo can effectively mitigate the performance limitations of label-free learning networks in 3D semantic segmentation tasks. Extensive experiments on two outdoor benchmark datasets highlight the superior performance of our method.

Introduction

As a core task of 3D perception, 3D semantic segmentation intends to identify different objects and parts, which has made significant progress in several fields such as robotics, autonomous driving, and mixed reality. Existing advanced methods (Hu et al. 2020; Zhu et al. 2021) inevitably rely on precise point-wise semantic annotations to train a well-performance segmenter. However, large-scale annotation for outdoor scenes is time-consuming and labor-intensive.

In the past few years, many methods adopt weakly supervised learning (Hu et al. 2022; Liu, Qi, and Fu 2021; Huang et al. 2024), semi-supervised learning (Li et al. 2021), and other methods (Zhang et al. 2023) to further reduce the need

*These authors contributed equally.

†Corresponding author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

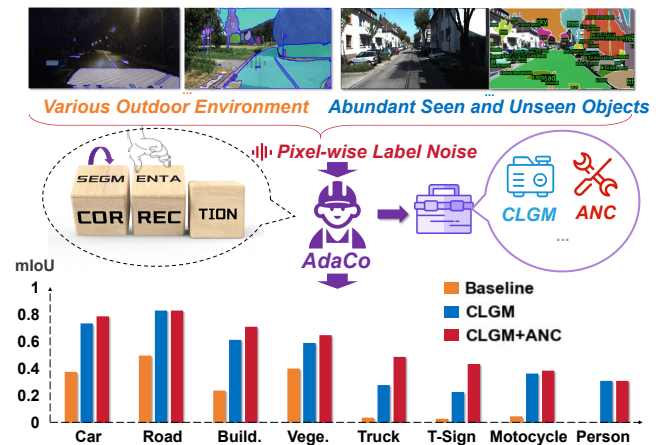


Figure 1: In outdoor scenes, variable environmental conditions and diverse object categories introduce significant noise when using VFMs for 3D perception tasks. Adaco effectively improves the label quality of the baseline method, MaskCLIP, in major categories in road scenes. "Build.", "Vege.", "T-Sign" means Building, Vegetation, Traffic-Sign.

for semantic annotation. However, these methods still rely on a certain amount of high-quality 3D annotations.

With the powerful understanding and generalization ability of the Visual Foundation Models (VFMs) in zero-shot 2D semantic segmentation, especially SAM (Kirillov et al. 2023) for segment masks and CLIP (Radford et al. 2021) for semantic labels, recent works (Chen et al. 2023b; Peng et al. 2023) show promising advances in label-free 3D semantic segmentation. Technically, these works generate 3D semantic labels by unprojecting 2D semantic features, extracted by VFMs, onto the corresponding point cloud.

However, these VFMs-assisted label-free methods also suffer some challenges. (1) *Inaccurate pseudo labels*. Current approaches follow the fashion of first extracting semantics from 2D images using VFMs and subsequently embedding them in 2D/3D features. Though refinements based on label consistency and spatial-temporal cues are well ex-

plored, the 3D pseudo labels remain inaccurate with great room to improve, particularly for challenging labels such as ‘truck’ and ‘person’ as indicated by the yellow bar in the bottom chart of Fig. 1. (2) *Unexpected label noise*. As image-to-point calibration errors cannot be fully eliminated, transferring 2D labels from VFMs into 3D may cause problems like edge blurring. Besides, the influence of illumination, occlusion, and other factors on the image may also introduce noise in the labels. (3) *Unrobust training process*. If the pseudo labels incorporate significant noise, the 3D segmenter is prone to fitting the noise during training, leading to an effectively unusable model in practical applications.

To address the above challenges, we propose a novel label-free 3D semantic segmentation method, AdaCo, which aims to refurbish noisy labels generated under the guidance of VFMs while alleviating the impact of error samples on the network. Firstly, we propose a Cross-modal Label Generation Module (CLGM) to generate 3D point annotations. Specifically, unlike embedding 2D labels into 3D features to yield pseudo labels in previous works, the CLGM lifts VFMs-produced descriptions into 3D space for voting. Without information loss in embedding, the CLGM generates more decent 3D pseudo labels than our baseline, shown in Fig. 1. Secondly, to alleviate the noise issue in the pseudo labels caused by the calibration errors and imaging variation, we propose an Adaptive Noise Corrector (ANC) to iteratively correct the mislabeled and fill the unlabeled point-wise semantic annotation to clean samples for the network training. Lastly, we further combine ANC with an Adaptive Robust Loss (ARL) to reduce the sensitivity of the 3D segmenter to noisy samples during the training process.

Extensive experiments on widely used datasets SemanticKITTI (Behley et al. 2019) and nuScenes (Caesar et al. 2019) demonstrate the superior performance of our AdaCo over the current baselines. The main contributions of this work are as follows:

- We propose a label-free method, Adaptive Label Correction (**AdaCo**), for 3D semantic segmentation. AdaCo outperforms state-of-the-art label-free 3D segmenters by a large margin.
- We propose a Cross-modal Label Generation Module (**CLGM**), leverages VFMs to generate initial noisy 3D pseudo labels from 2D images. Subsequently, we propose the Adaptive Noise Corrector (**ANC**) to refurbish these noisy labels iteratively.
- We propose Adaptive Robust Loss (**ARL**) to alleviate the network’s tendency to fit the residual noisy labels, enforcing the penalties for mislabeled labels by the composite loss function.

Related Work

Cross-Modality Supervision for LiDAR Semantic Segmentation. 3D semantic segmentation tasks typically require high-quality point-wise annotations, which can be labor-intensive and costly. To this end, 2D3DNet (Genova et al. 2021) utilizes 2D pseudo labels derived from a pre-trained 2D segmenter to train a 3D segmenter, and apply

a credible sparsity filter to deal with noisy 2D pseudo labels. 2D VFMs such as CLIP (Radford et al. 2021) and SAM (Kirillov et al. 2023), leveraging their powerful image understanding and generalization capabilities, are also employed to produce higher-quality 2D pseudo labels for label-free 3D semantic segmentation. OpenScene (Peng et al. 2023) and CLIP2Scene (Chen et al. 2023b) utilize the CLIP model to generate embeddings between images, text, and points, and design semantic consistency regularization to select positive and negative point samples, training the 3D network using contrastive loss. Chen et al. (Chen et al. 2023a) use SAM to refine the CLIP-predicted 2D pseudo labels as 3D cross-modal noisy labels and mitigate the impact of incorrect labels through latent space consistency optimization. However, their performance remains limited by the noise.

Noisy Label Learning for Deep Neural Networks.

Noisy label learning enables the robust training of deep neural networks in noisy label settings. To alleviate the impact of noisy labels, some works (Sukhbaatar et al. 2014; Bekker and Goldberger 2016; Lee et al. 2019) change the architecture to model the noise transition matrix of a noisy dataset and output the estimated-based label transition probability to improve the network’s generalization ability. Some other works (Xia et al. 2020; Wei et al. 2021) use regularization techniques to regulate the network overfitting to noisy labels during training. Still, the label noise rate restricts the model’s accuracy. (Zheng et al. 2020) propose the LRT method and test the confidence level of the noisy classifier to determine whether a label is correct. (Ye et al. 2021) introduce the label corrector to refurbish the noisy ground truth in the PCSS. However, these methods remain limited for complex open scenes, and the number of network warm-up rounds is manually preset. Our proposed ANC and ARL self-adapt to the learning situation of different samples, continuously refurbish the supervision signals and learn clean samples while mitigating noise label fitting.

Methodology

As shown in Fig. 2, our AdaCo comprises three modules: (1) CLGM utilizes VFMS for image semantic understanding and labeling 3D point clouds; (2) ANC refurbishes noisy labels using historical predictions and density-based clustering; (3) ARL applies constraints to the network to prevent overfitting to noisy labels. Both ANC and ARL adaptively adjusted according to sample learning.

Cross-modal Label Generation Module

Many works (Chen et al. 2023b; Peng et al. 2023; Chen et al. 2023a) use VFMs like CLIP or SAM for enhancing the feature extraction of 2D semantics. Motivated by their success, we introduce them to point-wise semantic annotations in 3D scenes. In this case, we propose the Cross-modal Label Generation Module (CLGM) to provide the point-wise pseudo label, serving as the exclusive supervision in network training while identifying more abundant semantic categories.

Pixel-wise Annotation Generation. As shown in Fig. 3, we first propose the 2D Pseudo Label Generation Engine

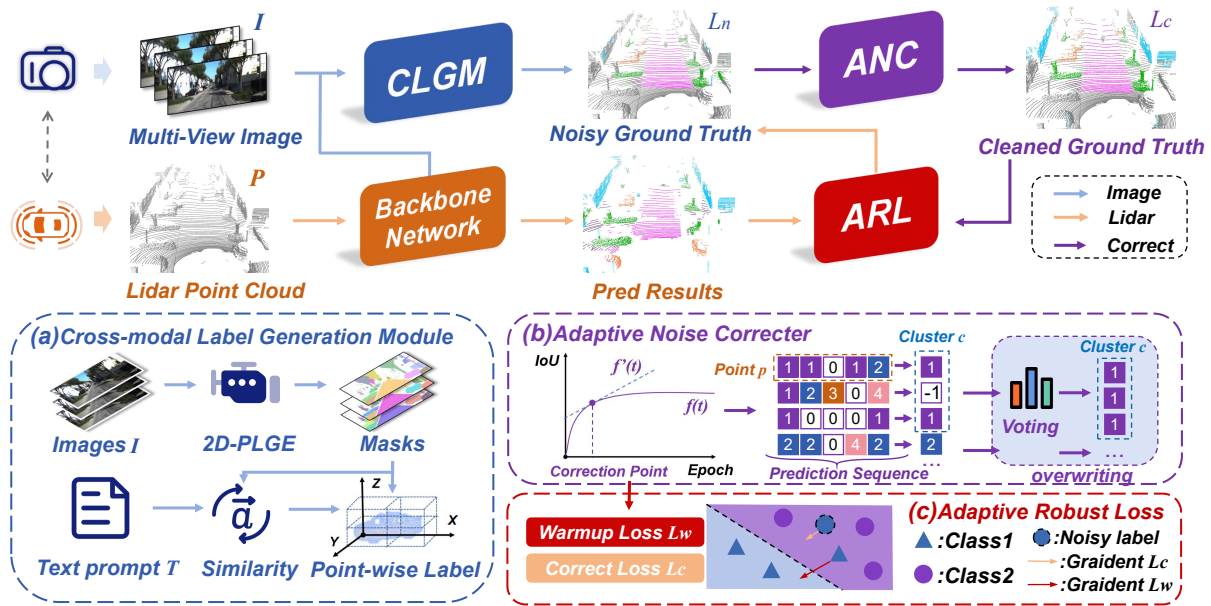


Figure 2: The illustration of our label-free 3D semantic segmentation method AdaCo. (a) The 3D noisy pseudo labels are obtained by transferring the VFMs-generate pixel-wise annotation to each point based on CLGM. (b) These noisy pseudo labels are then adaptively refurbished by the proposed ANC. Furthermore, (c) ARL is employed to adaptively adjust the loss method based on the training IoU curve and enforce different penalties for noisy and clean labels.

(2D-PLGE) to generate pixel-wise pseudo label. Specifically, take as input multi-view images $\mathcal{I} \in \mathbb{R}^{K \times 3 \times H \times W}$, we obtain a set of class-agnostic binary masks $\mathcal{M}^K = \{m_1^K, m_2^K, \dots, m_n^K | m_n^K \in \{0, 1\}^{W \times H}\}$ with the assistance of the FastSAM (Zhao et al. 2023), where m_n^K denoting the n -th mask in the K -th view image, W and H are the width and height of the image, respectively. For each mask m_n^K , we fed it along with its corresponding image patch into the SSA-Engine (Chen, Yang, and Zhang 2023) to generate the semantic description $\mathcal{D}^K = \{D_1^K, D_2^K, \dots, D_n^K\}$. Then, we utilize the category vocabulary as a textual prompt T to calculate the semantic similarity between T and the generated descriptions. Ultimately, we select the semantic class with the highest similarity to the text prompt as the final annotation y_{2D}^k for each image pixel.

Point-wise Label Unprojection and Voxel Refinement.

By utilizing the existing pixel-to-point calibration relationships and matrix computations, we derive point-wise annotations y_{3D} by back-projecting pixel-wise annotations y_{2D}^k from different viewpoints onto the point cloud. Leveraging the continuity of camera capture and point cloud acquisition across frames, we voxelize the point cloud and align the adjacent frame’s point cloud to the current frame’s coordinate system using the calibration matrix. We then vote the semantic annotation of the back projection point by point, to initially refine the quality of the pseudo label.

Adaptive Noise Corrector

Issues like lens distortion and sampling rate which lead to blurred object edges and the occlusion problem also intro-

duce extensive noise. To address this, we designed an Adaptive Noise Corrector (ANC) that utilizes the network’s early learning with clean samples and point clustering constraints.

Adaptive Correction Timing Search.

Inspired by (Liu et al. 2022), we first harness the early learning capabilities of deep neural networks to extract knowledge from clean data samples. This process will continuously update the noisy labels by the network’s historical predictions on unknown samples and dynamically adjust warm-up iterations according to each sample’s learning curve. As the green curve in Fig.4 shows, monitoring training mIoU and early learning mIoU on the training set, we note a rapid early increase in both, followed by a slowdown or decline at the same time as the model memorizes noise. Therefore, this turning point can be regarded as the pivotal moment at which each sample achieves saturation in learning from the clean data. For each iteration t , we let each sample s from the sample set S record its training mIoU, and then use the least squares method to fit an exponential function f to the mIoU as:

$$f_s(t) = a_s \cdot \left(1 - \exp\left(-\frac{1}{c_s} \cdot t^{b_s}\right) \right), \quad \forall t \in T, \forall s \in S \quad (1)$$

where a , b , and c are the fitting parameters. Subsequently, we calculate the derivative of the fitting function f' with respect to t from 1 to the current iteration t_c . We then set a threshold r . When the change in the derivative exceeds this threshold r , we consider it the moment t_c to correct the annotation.

Reliable Labels Prediction. After a correction round, t_c , we analyze the network’s historical predictions for each in-

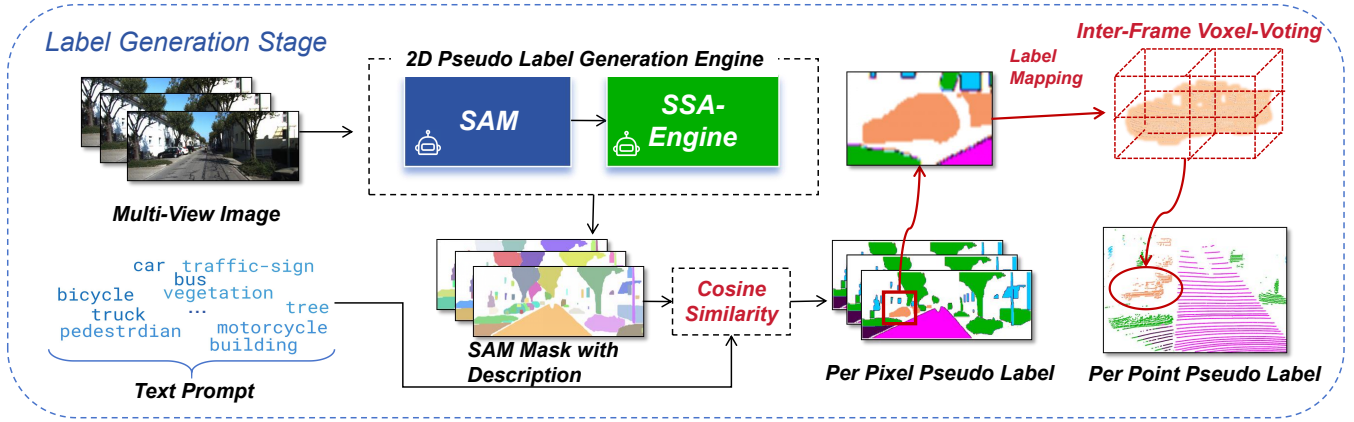


Figure 3: CLGM pipeline. We utilize the composition of SAM and SSA-Engine as our 2D-PLGE to segment masks with semantic description, and then calculate the semantic similarity between text prompt and description word by word, and map the class corresponding to the highest semantic similarity to each pixel as its label. Finally, we transfer the pixel-wise label to the point after refined by inter-frame voxel-voting.

stance, using the mode of these predictions to refurbish the noise labels. Additionally, following PNAL (Ye et al. 2021), we consider samples with consistent historical predictions and their labels as reliable. The point x_s in sample s is defined as a reliable sample if its confidence in historical prediction satisfies $F(x_s; q) \geq \gamma$ ($0 \leq \gamma \leq 1$). The predictive confidence is defined as:

$$F(x_s; q) = -\frac{\text{entropy}(P(k|x_s; q))}{\log \frac{1}{K}}, \quad (2)$$

where q is the length of history predictions. We hope that the length of the historical prediction record q can be self-adapted according to the sample, that is, the number of training rounds before the correction round t_c occurs, but q will not exceed the upper bound t_m to save memory used during training. The label history $H_{x_s}(q)$ of the point x_s denotes the predicted labels of the previous q rounds. The predicted labels of the points at time t as \hat{y}_t , and $H_{x_s}(q) = \{\hat{y}_{t1}, \dots, \hat{y}_{tq}\}$. For the point x_s in sample s , we denote its output probability of the label k as $P(k|x_s; q)$, where $k \in \{1, \dots, K\}$. We can calculate the probability as:

$$P(k|x_s; q) = \frac{\sum_{\hat{y} \in H_{x_s}(q)} [\hat{y} = k]}{q}, \quad (3)$$

$$q = t_c \quad \text{if } t_c < t_m, \quad t_m \quad \text{otherwise}, \quad (4)$$

where $[\cdot]$ is the Iverson bracket. We define the set of reliable points in sample s as X_{reliable}^s . For each point in the reliable set, we calculated its reliable labels as:

$$y_{x_s}^* = \text{argmax} P(k|x_s; q). \quad (5)$$

Clustering Label Propagation. We employ the DBSCAN (Ester et al. 1996) to provide instance information within the sample, addressing the issue of edge blur caused by direct back-projection. Initially, we remove ground points from the point cloud scene and apply the DBSCAN clustering algorithm to the remaining non-ground points to form




Figure 4: The learning curve of different mIoU in the SemanticKITTI *train* set. The training mIoU is calculated with the noisy ground truth, while the early learning IoU curve is calculated with the correct ground truth.

m clusters $C = \{c_1, \dots, c_m\}$. For each cluster, we replace the labels of the intersection points with the set of reliable points X_{reliable}^s . The new labels L_{ref} are determined by a voting mechanism within the intersection area. We have counted the frequency of each label within the intersection area and constructed a label frequency vector $\text{Fre}_s = \{\text{Fre}_s^0, \dots, \text{Fre}_s^K\}$ based on this information, where Fre_s^k ($0 \leq k \leq K$) represents the frequency of class k in this intersection area. The final winner label L_{ref} is randomly selected under the guidance of this vector, where $L_{ref} = \text{Random} \left\{ k | \text{Fre}_s^k \geq \frac{\max(\text{Fre}_s)}{\omega} \right\}$. Unlike the method used in the (Ye et al. 2021), after updating a portion of the labels, we enhance the network’s memory for clean samples by applying the complete label information from the entire sample set during gradient calculation. The Algorithm 1 shows the process of the Adaptive Noise Corrector.

Adaptive Robust Loss

The robust loss function is crucial for training with noisy labels. However, according to (Ma et al. 2020; Zhang and Sabuncu 2018; Wang et al. 2019), it may hinder optimal con-

vergence and early learning, potentially causing underfitting. On the other hand, cross-entropy loss offers rapid noise-free label early learning ability. Therefore, we propose optimizing training with two adjustable loss functions, guided by sample learning. In the initial phase of training, before the correction round t_c , our loss function primarily comprises cross-entropy loss on the 3D network and mean squared error loss between feature Z_{2D} from 2D knowledge branch and Z_C from camera branch following (Cen et al. 2023). The ARL in the warmup stage is defined below:

$$L_{\text{warmup}} = L_{\text{CE}} + L_{\text{MSE}}(Z_{2D}, Z_C) + L_{\text{Lovász}}(y, \hat{y}), \quad (6)$$

where \hat{y} is the prediction of 3D network, $L_{\text{Lovász}}$ denotes the Lovasz loss function (Berman, Triki, and Blaschko 2018).

Once the correction operation is initiated, we introduce additional terms to the loss function, including normalized cross-entropy loss and mean absolute error loss following (Ma et al. 2020), while diminishing the dominant weight of the preceding loss functions. The normalized cross-entropy loss and mean absolute error loss are as follows:

$$L_{\text{NCE}} = \frac{\sum_{k=1}^K q(k|x) \log p(k|x)}{\sum_{j=1}^K \sum_{k=1}^K q(y=j|x) \log p(k|x)}, \quad (7)$$

where the x denotes a sample in the dataset, K is the number of class. $q(y=k|x)$ denotes the distribution over different labels in class k for sample x , and the $\sum_k^K q(y=k|x) = 1$. The $p(k|x)$ denotes the probability output of the network as $p(k|x) = \frac{e^{z_k}}{\sum_{j=1}^K e^{z_j}}$, z_k represents the logits output of the network in class k .

$$L_{\text{MAE}} = \frac{1}{N} \sum_{i=1}^N |y_i - \hat{y}_i|, \quad (8)$$

where N presents the number of points. The additional loss term is calculated as:

$$L_{\text{corr}} = \lambda L_{\text{NCE}} + \beta L_{\text{MAE}}, \quad (9)$$

where the λ we set it to 100, and β we set it to 1 as (Ma et al. 2020) suggests. The overall adaptive robust loss L_{ARL} for our training model is calculated as:

$$L_{\text{ARL}} = \begin{cases} L_{\text{warmup}} & t < t_c \\ L_{\text{corr}} + L_{\text{warmup}} + \sigma L_{\text{CE}} & t \geq t_c, \end{cases} \quad (10)$$

where the σ we set to reduce the overfitting of the cross-entropy function to noise while ensuring gradient sensitivity.

Experiments

Aiming to evaluate the robust performance of our proposed AdaCo in the environment of noisy annotations, we conduct experiments on two large-scale outdoor LiDAR scene datasets, SemanticKITTI (Behley et al. 2019) and nuScenes (Caesar et al. 2019), and compare our method with the previous 3D label-free semantic segmentation methods. At the same time, we conducted a complete ablation experiment to verify the feasibility of each module. We will elaborate on the complete process of our experiments in terms of dataset, training details, ablation study, etc.

In our experiments, we utilize the CMDFusion (Cen et al. 2023) backbone network as our baseline, which is guided by pixel-wise pseudo labels generated from multi-view images in each dataset using MaskCLIP (Zhou, Loy, and Dai 2022).

Algorithm 1: Adaptive Noise Corrector

Require: LiDAR points X_s of sample s , the noisy label L_s , K classes, The clusters C ;
Ensure: The refurbished label L_{ref} ;
1: $H_x(q) = [], I_s = []$
2: **for** epoch $t = 0, 1, 2, \dots, T$ **do**
3: **for** sample $s \in S$ **do**
4: $\hat{y}_s = \text{Model}(X_s)$
5: $I_s.append(\text{mIoU}(\hat{y}_s, L_s))$
6: $f(t) = \text{fitFunc}(I_s)$
7: **if** $\frac{|f'(1) - f'(t)|}{f'(1)} > r$ **then**
8: $X_{\text{reliable}}^s = \{x_s | F(x_s; q) > \text{threshold}\}$
9: $y_{\text{reliable}}^s = \text{argmax}(\frac{\sum_{\hat{y} \in H_x(q)} [\hat{y}=k]}{|H_x(q)|})$
10: $C_{i^*} = \{x_s | \exists x_s \in X_{\text{reliable}}^s \cap x_s \in X_{C_i}, C_i \in C\}$
11: $\text{Fre}_{C_{i^*}}^k = \sum [y_{\text{reliable}}^s = k] [x_s \in X_{C_i}], (k \in [1, 2, \dots, K])$
12: $L_{\text{ref}} = \text{Random} \left\{ k | \text{Fre}_s^k \geq \frac{\max(\text{Fre}_s)}{\omega} \right\}$
13: ...
14: **end if**
15: $H_x(q).append(\hat{y}_s)$
16: **end for**
17: **end for**

Datasets and Metrics. SemanticKITTI (Behley et al. 2019) has 22 sequences with 20 semantic classes. We use 19,130 samples for training from sequences 00-10, and 4,071 samples for validation from sequences 08. The dataset provides 2 images of the front-view camera for each sample.

The nuScenes (Caesar et al. 2019) dataset is widely used in 3D perception tasks (Xia et al. 2023, 2024). It contains 28,130 samples for training, 6,019 samples for validation, 6,008 samples for testing, and 17 semantic classes. Each sample has LiDAR data of 32 beams and RGB images captured by 6 cameras, covering the whole 360 angle.

We present the results with the official evaluation metrics: mean Intersection over Union (mIoU) and mean Accuracy (mAcc). All experiments are conducted on point clouds and not on images in both datasets.

Implementation Details. Our backbone network settings follow the CMDFusion (Cen et al. 2023). We set voxel size to [0.05m,0.05m,0.05m] for SemanticKITTI and [0.1m,0.1m,0.1m] for nuScenes. The 2D-PLGE is kept frozen for inference. To obtain shapes of instances in the image that are as complete and non-overlapping as possible, we set the segmentation threshold of FastSAM (Zhao et al. 2023) to 0.7 and the mask confidence value to 0.4. For the voxel size of inter-frame voting, we adopted the same settings as those in (Cen et al. 2023). In the CLGM, we utilize dataset categories as text prompts and apply word2vec (Church 2017) to assess the word-by-word similarity be-

Settings	Methods	Publication	mIoU(%)	
			nuScenes	SemanticKITTI
Label-Free	MaskCLIP (Zhou, Loy, and Dai 2022)	ECCV 2022	12.8	-
	MaskCLIP++ (Zhou, Loy, and Dai 2022)	ECCV 2022	15.3	-
	CLIP2Scene (Chen et al. 2023b)	CVPR 2023	20.8	-
	OpenScene (Peng et al. 2023)	CVPR 2023	14.6	-
	TowardsVFMs (Chen et al. 2023a)	NIPS 2023	26.8	-
	Hierarchical (Kang et al. 2024)	CVPR 2024	23.0	-
Unsupervised	GrowSp (Zhang et al. 2023)	CVPR 2023	-	13.2
	U3DS ³ (Liu et al. 2024)	CVPR 2024	-	14.2
Label-Free	Baseline (MaskCLIP [†]) (Zhou, Loy, and Dai 2022)	-	16.5	8.1
	Ours	-	31.2	25.7

Table 1: Comparison results on official SemanticKITTI single scan validation set and the nuScenes val set with current state-of-the-art label-free and unsupervised methods for point cloud semantic segmentation tasks. † represents the reproduced result.

tween the 2D-PLGE-generated mask descriptions and the prompts. In the ANC, we set the derivative threshold r for label correction to 0.9, and our refinement is conducted at the sample level. Each sample records its mIoU for each round of training to fit a curve, whose parameters satisfy $0 \leq a \leq 1$, $b \geq 0$, $c \geq 0$. We record the network’s prediction results for each round of the sample to form a historical prediction sequence, with the maximum sequence length set to 5 due to memory constraints and they only correct once during the training phase to avoid the label quality decrease. The parameter ω we choose the winner label in the cluster is set to be 3. In the ARL, we set the parameters λ and β to be 100 and 1 as (Ma et al. 2020) suggests, emphasizing more active learning and less passive learning. The weight σ for minimizing cross-entropy loss is set to -0.99, and we have conducted various experiments to determine the optimal parameter settings. We train the models using the SGD optimizer with a learning rate of 0.24 for 40 epochs and adopt the cosine annealing warm restart strategy for learning rate decay. All experiments are trained on 4 RTX 3090 GPUs.

Comparison Results

Results on SemanticKITTI. SemanticKITTI has two images of the front-view camera per sample and we fed them into MaskCLIP to generate the 2D pseudo labels, then we used the pixel-to-point calibration matrix to transfer the pixel-wise pseudo labels to the point and training on CMD-Fusion (Cen et al. 2023) as our baseline. As shown in Tab. 1, AdaCo achieves a 17.6 mIoU improvement over the baseline. Furthermore, compared to the unsupervised methods, our method respectively achieves 12.5 and 11.5 higher mIoU. Fig. 5 presents a qualitative comparison between the ground truth, noisy ground truth, baseline, and our AdaCo.

Results on NuScenes. In this experiment, we compared our AdaCo method with existing 3D label-free semantic segmentation methods that use 2D VFMs to guide 3D networks. As shown in Tab. 1, our method outperforms the state-of-the-art method by 4.4 in mIoU, and the qualitative results in Fig. 5 demonstrate the effectiveness of our approach.

CLGM	ANC	ARL	mIoU(%)	mAcc(%)
			8.1	35.1
✓			19.5	58.2
✓	✓		25.3	73.2
✓		✓	21.4	65.9
✓	✓	✓	25.7	74.9

Table 2: Ablation study on SemanticKITTI *validation* set.

Ablation Study

To validate the effectiveness of each component in AdaCo, we conducted ablation experiments on the SemanticKITTI validation set and reported the results in Tab. 2. In addition, we have also verified the impact of hyperparameters for each component in AdaCo and reported the results in Tab. 3.

Effect of CLGM. As the Tab. 2 indicates, our CLGM demonstrates better semantic comprehension capabilities than the baseline. In Tab. 3, the results presented in the *Adjacent Frame* row illustrate the impact of the length of adjacent frames. As the length increased, the quality of the labels improved. Due to memory constraints, we did not use longer adjacent frames to optimize the labels.

Effect of ANC. As shown in Tab. 2, our ANC strategy significantly enhances the quality of CLGM-generate pseudo labels. In Tab. 3, we adjusted the maximum historical prediction record values for various lengths in the *History Seq.* row, suggesting a shorter length would result in a lack of reference for ANC to correct samples. Due to memory constraints and performance considerations, the optimal length was determined to be 5. Additionally, the results in *Adaptive Correction* row indicate that adaptive correction only corrected once per sample performed better than correct at each time the derivative of the fitting function started to decrease.

Effect of ARL. As depicted in Tab. 2, ARL effectively mitigates the interference of noisy pseudo labels on the model, demonstrating a more robust segmentation performance. Moreover, we refine the ARL by adjusting the

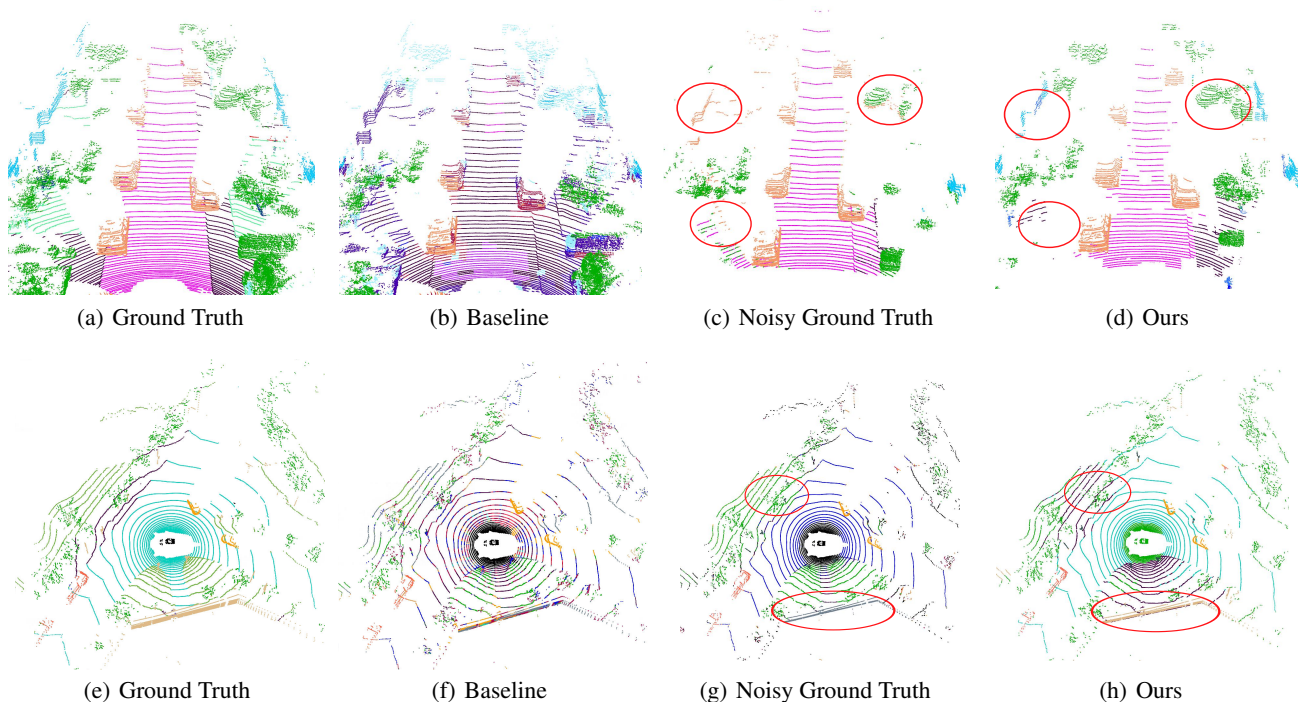


Figure 5: Qualitative results in SemanticKITTI (a)-(d) and nuScenes (e)-(h), the noisy ground truth is generated from CLGM. As shown by the red circle in the figure, our Adaco remembered clean samples and mitigated the blurring of object edges.

Ablation Target	Settings	mIoU (%)	mAcc(%)
Adjacent Frame	1, w/o train	25.4	74.6
	2, w/o train	25.9	75.0
History Seq.	3	18.5	56.2
	4	21.6	65.5
	5	25.3	73.2
Adaptive Correction	Correct Each Down	18.9	57.3
	Correct Once	24.0	70.4
Robust Loss Weight	$\lambda/\beta = 10, \sigma = 0$	18.0	54.7
	$\lambda/\beta = 10, \sigma = -0.9$	18.1	55.3
	$\lambda/\beta = 100, \sigma = -0.9$	18.4	56.1
	$\lambda/\beta = 10, \sigma = -0.99$	21.4	65.9
	$\lambda/\beta = 100, \sigma = -0.99$	21.6	66.3
	Full	25.7	74.9

Table 3: The impact of the module hyperparameters on SemanticKITTI validation set.

weights of the various active, passive, and unnormalized cross-entropy losses during the correction phase. As shown in the *Robust Loss Weight* row in Tab. 3, the model performance tended to improve when decreasing the contribution of the unnormalized loss to the total loss. Besides, the model’s resilience to noise is also strengthened after incorporating the additional robust noise loss terms. Fig. 4 shows that ARL has better early learning capabilities in the initial stage of the model and is less prone to fitting noisy samples.

Conclusion

We propose a novel label-free point cloud semantic segmentation approach, AdaCo, which leverages VFMs to provide cross-modal 3D semantic pseudo labels and adaptively adjusts the noisy labels within the samples. Furthermore, it robustly learns from noisy information. Firstly, we employ the CLGM to extract semantics from images and transfer them to point clouds to provide point-wise pseudo labels. Next, we propose ANC and ARL to mitigate the impact of noisy labels provided by the CLGM on network learning. ANC refurbishes the supervision signals by integrating the instance shapes within the scene with historical predictions. ARL employs a combination of normalized cross-entropy loss and robust loss to enhance the penalty for incorrect annotations on network learning. **Limitations:** AdaCo relies on pixel-to-point calibration for semantic label transfer and is limited by the semantic understanding of VFMs in driving contexts, potentially failing with small image targets and diverse categories. Future work may explore decoupling 2D image calibration from 3D point clouds for cross-dataset supervision or offering a more precise semantic capture of road images.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No.62171393), and the Fundamental Research Funds for the Central Universities (No.20720220064).

References

- Behley, J.; Garbade, M.; Milioto, A.; Quenzel, J.; Behnke, S.; Stachniss, C.; and Gall, J. 2019. SemanticKITTI: A Dataset for Semantic Scene Understanding of LiDAR Sequences. In *Proc. of the IEEE/CVF International Conf. on Computer Vision (ICCV)*.
- Bekker, A. J.; and Goldberger, J. 2016. Training deep neural networks based on unreliable labels. In *2016 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2682–2686. IEEE.
- Berman, M.; Triki, A. R.; and Blaschko, M. B. 2018. The lovász-softmax loss: A tractable surrogate for the optimization of the intersection-over-union measure in neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4413–4421.
- Caesar, H.; Bankiti, V.; Lang, A. H.; Vora, S.; Liong, V. E.; Xu, Q.; Krishnan, A.; Pan, Y.; Baldan, G.; and Beijbom, O. 2019. nuScenes: A multimodal dataset for autonomous driving. *arXiv preprint arXiv:1903.11027*.
- Cen, J.; Zhang, S.; Pei, Y.; Li, K.; Zheng, H.; Luo, M.; Zhang, Y.; and Chen, Q. 2023. CMDFusion: Bidirectional Fusion Network With Cross-Modality Knowledge Distillation for LiDAR Semantic Segmentation. *IEEE Robotics and Automation Letters*, 9(1): 771–778.
- Chen, J.; Yang, Z.; and Zhang, L. 2023. Semantic Segment Anything. <https://github.com/fudan-zvg/Semantic-Segment-Anything>. Accessed: 2024-08-10.
- Chen, R.; Liu, Y.; Kong, L.; Chen, N.; Xinge, Z.; Ma, Y.; Liu, T.; and Wang, W. 2023a. Towards label-free scene understanding by vision foundation models. In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Chen, R.; Liu, Y.; Kong, L.; Zhu, X.; Ma, Y.; Li, Y.; Hou, Y.; Qiao, Y.; and Wang, W. 2023b. Clip2scene: Towards label-efficient 3d scene understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7020–7030.
- Church, K. W. 2017. Word2Vec. *Natural Language Engineering*, 23(1): 155–162.
- Ester, M.; Kriegel, H.-P.; Sander, J.; Xu, X.; et al. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *kdd*, volume 96, 226–231.
- Genova, K.; Yin, X.; Kundu, A.; Pantofaru, C.; Cole, F.; Sud, A.; Brewington, B.; Shucker, B.; and Funkhouser, T. 2021. Learning 3d semantic segmentation with only 2d image supervision. In *2021 International Conference on 3D Vision (3DV)*, 361–372. IEEE.
- 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 *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVII*, 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/CVF conference on computer vision and pattern recognition*, 11108–11117.
- Huang, W.; Zou, P.; Xia, Y.; Wen, C.; Zang, Y.; Wang, C.; and Zhou, G. 2024. OPOCA: One Point One Class Annotation for LiDAR Point Cloud Semantic Segmentation. *IEEE Transactions on Geoscience and Remote Sensing*, 62: 1–10.
- Kang, X.; Chu, L.; Li, J.; Chen, X.; and Lu, Y. 2024. Hierarchical Intra-modal Correlation Learning for Label-free 3D Semantic Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 28244–28253.
- 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.
- Lee, K.; Yun, S.; Lee, K.; Lee, H.; Li, B.; and Shin, J. 2019. Robust inference via generative classifiers for handling noisy labels. In *International conference on machine learning*, 3763–3772. PMLR.
- Li, H.; Sun, Z.; Wu, Y.; and Song, Y. 2021. Semi-supervised point cloud segmentation using self-training with label confidence prediction. *Neurocomputing*, 227–237.
- Liu, J.; Yu, Z.; Breckon, T. P.; and Shum, H. P. 2024. U3DS3: Unsupervised 3D Semantic Scene Segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3759–3768.
- Liu, S.; Liu, K.; Zhu, W.; Shen, Y.; and Fernandez-Granda, C. 2022. Adaptive early-learning correction for segmentation from noisy annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2606–2616.
- 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.
- Ma, X.; Huang, H.; Wang, Y.; Romano, S.; Erfani, S.; and Bailey, J. 2020. Normalized loss functions for deep learning with noisy labels. In *International conference on machine learning*, 6543–6553. PMLR.
- Peng, S.; Genova, K.; Jiang, C.; Tagliasacchi, A.; Pollefeys, M.; Funkhouser, T.; et al. 2023. Openscene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 815–824.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, 8748–8763. PMLR.
- Sukhbaatar, S.; Bruna, J.; Paluri, M.; Bourdev, L.; and Fergus, R. 2014. Training convolutional networks with noisy labels. *arXiv preprint arXiv:1406.2080*.
- Wang, Y.; Ma, X.; Chen, Z.; Luo, Y.; Yi, J.; and Bailey, J. 2019. Symmetric cross entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF international conference on computer vision*, 322–330.

Wei, H.; Tao, L.; Xie, R.; and An, B. 2021. Open-set label noise can improve robustness against inherent label noise. *Advances in Neural Information Processing Systems*, 34: 7978–7992.

Xia, Q.; Deng, J.; Wen, C.; Wu, H.; Shi, S.; Li, X.; and Wang, C. 2023. Coin: Contrastive instance feature mining for outdoor 3d object detection with very limited annotations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 6254–6263.

Xia, Q.; Ye, W.; Wu, H.; Zhao, S.; Xing, L.; Huang, X.; Deng, J.; Li, X.; Wen, C.; and Wang, C. 2024. HINTED: Hard Instance Enhanced Detector with Mixed-Density Feature Fusion for Sparsely-Supervised 3D Object Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15321–15330.

Xia, X.; Liu, T.; Han, B.; Gong, C.; Wang, N.; Ge, Z.; and Chang, Y. 2020. Robust early-learning: Hindering the memorization of noisy labels. In *International conference on learning representations*.

Ye, S.; Chen, D.; Han, S.; and Liao, J. 2021. Learning with noisy labels for robust point cloud segmentation. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6443–6452.

Zhang, Z.; and Sabuncu, M. 2018. Generalized cross entropy loss for training deep neural networks with noisy labels. *Advances in neural information processing systems*, 31.

Zhang, Z.; Yang, B.; Wang, B.; and Li, B. 2023. Growsp: Unsupervised semantic segmentation of 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17619–17629.

Zhao, X.; Ding, W.; An, Y.; Du, Y.; Yu, T.; Li, M.; Tang, M.; and Wang, J. 2023. Fast segment anything. *arXiv preprint arXiv:2306.12156*.

Zheng, S.; Wu, P.; Goswami, A.; Goswami, M.; Metaxas, D.; and Chen, C. 2020. Error-bounded correction of noisy labels. In *International Conference on Machine Learning*, 11447–11457. PMLR.

Zhou, C.; Loy, C. C.; and Dai, B. 2022. Extract free dense labels from clip. In *European Conference on Computer Vision*, 696–712. Springer.

Zhu, X.; Zhou, H.; Wang, T.; Hong, F.; Li, W.; Ma, Y.; Li, H.; Yang, R.; and Lin, D. 2021. Cylindrical and asymmetrical 3d convolution networks for lidar-based perception. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(10): 6807–6822.