








LiDAL: Inter-frame Uncertainty Based Active Learning for 3D LiDAR Semantic Segmentation

Zeyu Hu¹^{*}, Xuyang Bai¹, Runze Zhang², Xin Wang², Guangyuan Sun², Hongbo Fu³, and Chiew-Lan Tai¹

¹ Hong Kong University of Science and Technology
{zhuam,xbaiad,taicl}@cse.ust.hk

² Lightspeed & Quantum Studios, Tencent
{ryanrzzhang,alexinwang,gerrysun}@tencent.com

³ City University of Hong Kong
hongbofu@cityu.edu.hk

Abstract. We propose LiDAL, a novel active learning method for 3D LiDAR semantic segmentation by exploiting inter-frame uncertainty among LiDAR frames. Our core idea is that a well-trained model should generate robust results irrespective of viewpoints for scene scanning and thus the inconsistencies in model predictions across frames provide a very reliable measure of uncertainty for active sample selection. To implement this uncertainty measure, we introduce new inter-frame divergence and entropy formulations, which serve as the metrics for active selection. Moreover, we demonstrate additional performance gains by predicting and incorporating pseudo-labels, which are also selected using the proposed inter-frame uncertainty measure. Experimental results validate the effectiveness of LiDAL: we achieve 95% of the performance of fully supervised learning with less than 5% of annotations on the SemanticKITTI and nuScenes datasets, outperforming state-of-the-art active learning methods. Code release: <https://github.com/hzykent/LiDAL>

Keywords: Active Learning, 3D LiDAR Semantic Segmentation

1 Introduction

Light detection and ranging (LiDAR) sensors capture more precise and farther-away distance measurements than conventional visual cameras, and have become a necessity for an accurate perception system of outdoor scenes. These sensors generate rich 3D geometry of real-world scenes as 3D point clouds to facilitate a thorough scene understanding, in which 3D LiDAR semantic segmentation serves as a cornerstone. The semantic segmentation task is to parse a scene and assign an object class label to each point in 3D point clouds, thus providing point-wise perception information for numerous downstream applications like robotics [41] and autonomous vehicles [18].

^{*} intern at Tencent Lightspeed & Quantum Studios




Fig. 1: **Illustration of inter-frame uncertainty.** While in one frame (Left) an object is correctly predicted as “vehicle” (highlighted in a red box), in the subsequent frame (Right) a large part of this object is mistakenly predicted as “fence” when scanned from a different viewpoint.

Thanks to the large-scale LiDAR datasets [1,3] made publicly available in recent years, the state of the art in 3D LiDAR semantic segmentation has been significantly pushed forward [40,11,57]. However, the requirement of fully labeled point clouds for existing segmentation methods has become a major obstacle to scaling up the perception system or extending it to new scenarios. Typically, since a LiDAR sensor may perceive millions of points per second, exhaustively labeling all points is extremely laborious and time-consuming. It poses demands on developing label-efficient approaches for 3D LiDAR semantic segmentation.

Active learning provides a promising solution to reduce the costs associated with labeling. Its core idea is to design a learning algorithm that can interactively query a user to label new data samples according to a certain policy, leading to models trained with only a fraction of the data while yielding similar performances. Inspired by 2D counterparts [42,7,27,19,35], some previous works have explored active learning for 3D LiDAR semantic segmentation [26,21,49,34]. However, these methods almost exclusively operate on single LiDAR frames. Such a strategy is surprising since, unlike most 2D datasets, in which images are captured as independent samples, 3D LiDAR datasets are generally scanned as continuous point cloud frames. As a consequence, inter-frame constraints naturally embedded in the LiDAR scene sequences are largely ignored. We believe such constraints are particularly interesting for examining the quality of network predictions; i.e., the same object in a LiDAR scene should receive the same label when scanned from different viewpoints (see Fig. 1).

In this work, we propose to exploit inter-frame constraints in a novel view-consistency-based uncertainty measure. More specifically, we propose new inter-frame divergence and entropy formulations based on the variance of predicted score functions across continuous LiDAR frames. For a given (unlabeled) object (e.g., the one in a red box in Fig. 1), if its predicted labels differ across frames, we assume faulty network predictions and then strive to obtain user-specified labels for the most uncertain regions. In addition to the main active learning formulation, we also explore further improvements to the labeling efficiency with

self-training by utilizing the proposed uncertainty measure in an inverse way. During each active learning iteration, we augment the user-specified labels with pseudo-labels generated from the most certain regions across frames to further boost performance without extra annotations or much computational cost.

To summarize, our contributions are threefold:

1. We propose a novel active learning strategy for 3D LiDAR semantic segmentation by estimating model uncertainty based on the inconsistency of predictions across frames.
2. We explore self-training in the proposed active learning framework and show that further gains can be realized by including pseudo-labels.
3. Through extensive experiments, we show that the proposed active learning strategy and self-training technique significantly improve labeling efficiency over baselines, and establish the state of the art in active learning for 3D LiDAR semantic segmentation.

2 Related Work

Compared to fully supervised methods [40,13,57,37,5,12], label-efficient 3D semantic segmentation is a relatively open research problem. Previous explorations can be roughly divided into five categories: transfer learning, unsupervised and self-supervised learning, weakly-supervised learning, active learning, and self-training. LiDAL falls into both the active learning and self-training categories.

Transfer Learning. Taking advantage of existing fully labeled datasets, transfer learning has been introduced to 3D semantic segmentation for reducing the annotation costs. Various domain adaptation approaches have been developed to make them perform well in novel scenarios given only labeled data from other domains [22,54,17] or even synthetic training sets [48]. They achieve fairly decent results but still require fully labeled data from a source domain and fail to generalize to new scenarios that are highly different from the source.

Unsupervised and Self-supervised Learning. Leveraging the colossal amount of unlabeled data, pre-trained models can be fine-tuned on a small set of labeled data to alleviate the over-dependence on labels and thus achieve satisfactory performances [8,39,36,51,14]. Pseudo tasks used for pre-training include reconstructing space [30], contrast learning [50,24,9], ball cover prediction [33], instance discrimination [56], and point completion [44], etc. Compared to other label-efficient counterparts, these methods require more labeled data and most of them only apply to object-level point clouds.

Weakly-supervised Learning. Instead of point-by-point labeling in fully supervised learning, weak labels take various forms like scene-level or sub-cloud-level labels [28,46], 2D supervision [43], fewer point labels [53,4,55,10,47], seg-level labels [25,38], and box-level labels [23], etc. These methods can reduce the number of labeled samples, but either require intricate labeling processes or produce much more inferior results than the fully-supervised counterparts.

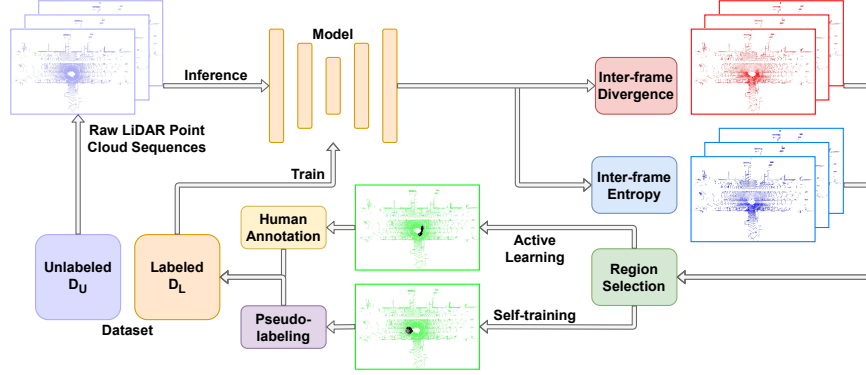


Fig. 2: **Pipeline of LiDAL.** In each round of active learning, we first train a 3D LiDAR semantic segmentation network in supervision with labeled dataset D_L . Second, we use the trained network to compute an inter-frame divergence score and an inter-frame entropy score for all regions from the unlabeled dataset D_U . We then select a batch of regions based on these scores for active learning and self-training, and finally request their respective labels from the human annotation and the pseudo-labeling. The process is repeated until the labeling budget is exhausted or all training data is labeled.

Active Learning. During network training, active learning methods iteratively select the most valuable data for label acquisition. The very few existing methods have explored uncertainty measurements like segment entropy [20], color discontinuity [49], and structural complexity [49,34]. The existing methods take LiDAR data as separated frames and only consider intra-frame information. Inspired by a 2D work operating on multi-view images [35], we take advantage of the inter-frame constraints for active learning in this work. Different from this 2D work, due to the distinct natures of 2D images and 3D point clouds, we design novel uncertainty formulations and selection strategies. Moreover, we propose a joint active learning and self-training framework to further exploit the inter-frame constraints.

Self-training. Building on the principle of knowledge distillation, previous methods generate pseudo labels to expand sparse labels [52,25] or to facilitate the network training using only scene-level supervision [28]. In this work, we develop a pseudo-labeling method applied in conjunction with our active learning framework to achieve even greater gains in efficiency.

3 Method

3.1 Overview

The goal of LiDAL is to train a well-performing 3D LiDAR semantic segmentation model with a constrained annotation budget. Specifically, we assume the

availability of data $D = \{D_L, D_U\}$. The data consists of sequences of LiDAR frames, each provided with an ego-pose of the scanning device. Initially, D_L is a small set of LiDAR frames randomly selected from D with each frame having its label annotation, and D_U is a large unlabeled set without any annotations. Following previous works [35,49], we use a sub-scene region as a fundamental query unit to focus on the most informative parts of the LiDAR frames.

As illustrated in Fig. 2, our LiDAL method consists of four main steps: 1. Train the network to convergence using the currently labeled dataset D_L . 2. Calculate the model uncertainty scores for each region of D_U with two indicators: inter-frame divergence and inter-frame entropy (Section 3.2). 3. Select regions based on the uncertainty measures for active learning and self-training (Section 3.3). 4. Obtain labels from human annotation and pseudo-labeling. These steps will be repeated until the labeling budget is exhausted or all training data is labeled.

3.2 Uncertainty Scoring

At each iteration, after the first step of training the network on D_L , our active learning method LiDAL then aims at predicting which samples from D_U are the most informative to the network at the current state. To this end, we introduce two novel uncertainty scoring metrics named *inter-frame divergence* and *inter-frame entropy*. Fig. 3 provides an overview of the scoring process.

Inter-frame Divergence. In a nutshell, the proposed inter-frame divergence score aims at estimating which objects are consistently predicted the same way, irrespective of the scanning viewpoints.

For each frame, we first calculate its point-wise class probability maps using the current trained segmentation network. To attain robust probability predictions, we perform data augmentations with random translation, rotation, scaling, and jittering. The probability P for a point p in frame F_i to belong to class c is given by:

$$P_i^p(c) = \frac{1}{D} \sum_{d=1}^D P_{i,d}^p(c), \quad (1)$$

where D is the number of augmented inference runs of the segmentation network, and $P_{i,d}^p(c)$ is the softmax probability of point p belonging to class c in the augmented inference run d .

Next, using the provided ego-pose, we register each frame in the world coordinate system. For each point in a given frame, we find its corresponding points in the neighboring frames and assign to it their associated probability distributions. Implementation details can be found in **Supplementary Section A**. Each point p in frame F_i is now associated with a set of probability distributions Ω_i^p , each coming from a neighboring frame:

$$\Omega_i^p = \{P_j^p, j | F_j^p \text{ corresponds to } F_i^p\}, \quad (2)$$

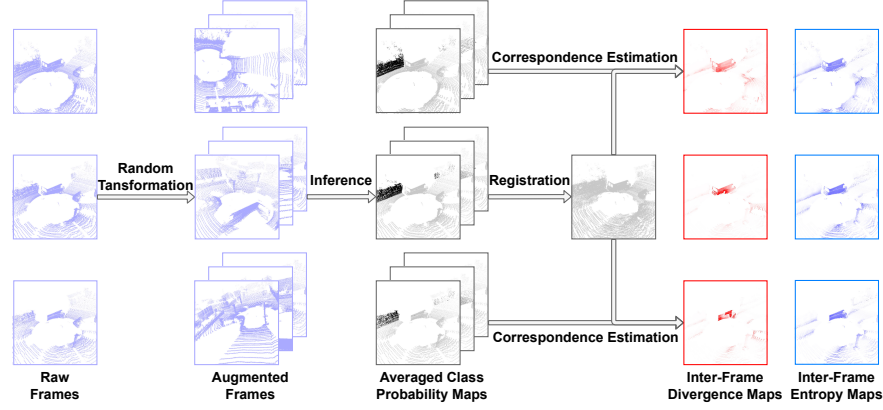


Fig. 3: **Illustration of uncertainty scoring.** For each unlabeled LiDAR frame in the dataset, we first obtain its averaged class probability predictions from augmented inference runs. Next, we register each frame with its provided ego-pose. For each point in the frame, we then find its corresponding points in neighboring frames and assign to it their associated class probability predictions. With the aggregated multiple class probability predictions per point, we compute the inter-frame divergence and entropy scores. We assign these two scores to each unlabeled region by averaging the scores of all the points contained in it.

where F_j^p denotes the point p in frame F_j , and F_j represents one of the neighboring frames of F_i .

In our setting, when estimating the point correspondences between neighboring frames, we assume that the objects in the scene are static and thus the points in the same registered position represent the same object. The moving objects are not specially treated for two reasons. First, they contribute only a small portion of the dataset. Second, when estimating correspondences after registration, the prediction disagreements introduced by the 3D motions can be seen as inter-frame inconsistency and help the system select these informative regions.

The inter-frame divergence corresponds to the average pairwise KL divergence between the class distribution at any given point and the class distributions assigned to that point from the neighboring frames. It effectively captures the degree of agreement between the prediction in the current frame with the predictions coming from the neighboring frames. Specifically, we define the inter-frame divergence score FD for a point p in frame F_i as follows:

$$FD_i^p = \frac{1}{|\Omega_i^p|} \sum_{P_j^p \in \Omega_i^p} D_{KL}(P_i^p || P_j^p), \quad (3)$$

where $D_{KL}(P_i^p || P_j^p)$ is the KL Divergence between distributions P_i^p and P_j^p .

Inter-frame Entropy. After measuring how inconsistent the predictions are across frames, we then define the inter-frame entropy score, which indicates the amount of uncertainty for the network to process a certain point. For a point p in frame F_i , with the aggregated probability distributions Ω_i^p , the mean distribution M_i^p can be calculated as:

$$M_i^p = \frac{1}{|\Omega_i^p|} \sum_{P_j^p \in \Omega_i^p} P_j^p, \quad (4)$$

which can be seen as the marginalization of the prediction probabilities over the scanning viewpoints.

The inter-frame entropy score FE is defined as the entropy of the mean class probability distribution M_i^p :

$$FE_i^p = - \sum_c M_i^p(c) \log(M_i^p(c)). \quad (5)$$

A high inter-frame entropy score implies that on average, the prediction of the current network for this point is significantly uncertain. Since the mean class probability distribution is the average result from both the augmented inference runs and the aggregation of corresponding points, the inter-frame entropy score estimates both the intra-frame uncertainty under random affine transformations and inter-frame uncertainty under viewpoint changes.

3.3 Region Selection

To select the most informative parts of the unlabeled dataset, we opt for using sub-scene regions as the fundamental label querying units, following previous works [35,49]. Our implementation uses the constrained K-means clustering [2] algorithm for region division. An ideal sub-scene region consists of one or several object classes and is lightweight to label for the annotator.

For each region r , the two scores FD_i^r and FE_i^r are computed as the average of the inter-frame divergence and inter-frame entropy scores of all the points contained in r :

$$FD_i^r = \frac{1}{|r|} \sum_{p \in r} FD_i^p, \quad (6)$$

$$FE_i^r = \frac{1}{|r|} \sum_{p \in r} FE_i^p, \quad (7)$$

where $|r|$ is the number of points contained in region r .

Active Learning. We now discuss our active learning strategy utilizing the proposed inter-frame uncertainty scores. Our strategy to select the next region for labeling consists of two steps. First, we look for the region r from frame F_i that has the highest inter-frame divergence score in D_U :

$$(i, r) = \arg \max_{(j, s) \in D_U} FD_j^s, \quad (8)$$




Fig. 4: **Examples of uncertainty scores.** Red and blue indicate inter-frame divergence and entropy, respectively. The darker the color, the higher the value. Due to the sparsity and varying-density property of LiDAR point clouds, neural networks tend to generate class distributions that are more uniform for farther away sparse points. As highlighted by the dotted red boxes, this property results in misleadingly high values for far away sparse points in terms of inter-frame entropy but affects less on the inter-frame divergence scores.

where (j, s) refers to region s from frame F_j .

Since the inter-frame divergence indicates that for each region how inconsistent the predictions are across frames, the scores are similar for all the regions that are in correspondence. To determine which one of the regions in correspondence contains the largest amount of beneficial information to improve the network, we retrieve the set of regions representing the same part of the outdoor scene and denote this set as S . We then look for the region from S with the highest inter-frame entropy score:

$$(k, t) = \arg \max_{(j, s) \in S} \{FE_j^s | (j, s) \text{ and } (i, r) \text{ overlap}\}, \quad (9)$$

where (k, t) refers to the selected region t from frame F_k . Implementation details can be found in **Supplementary Section A**.

The selected region is added to D_L and all regions in set S are then removed from D_U to avoid label redundancy. The process is repeated until reaching the labeling budget. The active learning algorithm is summarized in Alg. 1.

One possible alternative strategy is to first find the region with the highest inter-frame entropy score and then select the one with the highest inter-frame divergence score in the corresponding set. A similar strategy is implemented in a previous work operating on multi-view images [35]. However, unlike 2D images with dense and uniformly sampled pixels, 3D LiDAR frames have sparse and varying-density points. Specifically, the farther from the scanning viewpoint, the sparser the LiDAR points. As shown in Fig. 4, due to the sparsity, neural networks tend to predict uniform class distributions for peripheral points. This property will result in misleadingly high inter-frame entropy scores for points far away from the scanning viewpoint (Equation 5), while the inter-frame divergence scores remain stable (Equation 3). Considering the robustness of the system, we

opt for the proposed strategy instead of the possible alternative. A quantitative comparison can be found in Section 4.4.

Self-training. To further exploit the inter-frame constraints embedded in the LiDAR sequences, we leverage the fact that our measure of viewpoint inconsistency can also help us identify reliable regions with high-quality pseudo-labels, which can be directly injected into the training set.

On the contrary with respect to active learning, which selects samples with the most uncertain predictions, self-training aims at acquiring confident and accurate pseudo-labels. To this end, we conduct a reversed process of the proposed active learning strategy. Specifically, we first look for the region r from frame F_i that has the lowest inter-frame divergence score in D_U :

$$(i, r) = \arg \min_{(j, s) \in D_U} FD_j^s, \quad (10)$$

We then look for the region from the corresponding set S that has the lowest inter-frame entropy score:

$$(k, t) = \arg \min_{(j, s) \in S} \{FE_j^s | (j, s) \text{ and } (i, r) \text{ overlap}\}, \quad (11)$$

The pseudo-label of the selected region is retrieved from the network predictions and all regions in set S will not be used for further pseudo-labeling. The process is repeated until reaching the target number of pseudo-labels. In order to prevent label drifting [6], we reset the pseudo label set at each iteration and only select regions that are not already selected in the previous iteration. The self-training algorithm is summarized in Alg. 2.

Algorithm 1 Active Learning

Input:

Data set D , labeled set D_L ,
annotation budget B , metric M

Init:

Added samples $A \leftarrow \{\}$
Unlabeled set $D_U \leftarrow D \setminus D_L$

repeat

$(i, r) \leftarrow \arg \max_{(j, s) \in D_U} MFD_j^s$
Retrieve $S \triangleright$ Corresponding set
 $(k, t) \leftarrow \arg \max_{(j, s) \in S} MFE_j^s$
 $A \leftarrow A \cup (k, t)$
 $D_L \leftarrow D_L \cup (k, t)$
 $D_U \leftarrow D_U \setminus S$

until $|A| = B$ or $|D_U| = 0$

return A

Algorithm 2 Self-training

Input:

Data set D , labeled set D_L ,
previous pseudo set P ,
target number T , metric M

Init:

New pseudo set $P' \leftarrow \{\}$
 $D'_U \leftarrow (D \setminus D_L) \setminus P \triangleright$ No re-labeling

repeat

$(i, r) \leftarrow \arg \min_{(j, s) \in D'_U} MFD_j^s$
Retrieve $S \triangleright$ Corresponding set
 $(k, t) \leftarrow \arg \min_{(j, s) \in S} MFE_j^s$
 $P' \leftarrow P' \cup (k, t)$
 $D'_U \leftarrow D'_U \setminus S$

until $|P'| = T$ or $|D'_U| = 0$

return P'

4 Experiments

To demonstrate the effectiveness of our proposed method, we now present various experiments conducted on two large-scale 3D LiDAR semantic segmentation datasets, i.e., SemanticKITTI [1] and nuScenes [3]. We first introduce the datasets and evaluation metrics in Section 4.1, and then present the experimental settings in Section 4.2. We report the results on the SemanticKITTI and nuScenes datasets in Section 4.3, and the ablation studies in Section 4.4.

4.1 Datasets and Metrics

SemanticKITTI [1]. SemanticKITTI is a large-scale driving-scene dataset derived from the KITTI Vision Odometry Benchmark and was collected in Germany with the Velodyne-HDLE64 LiDAR. The dataset consists of 22 sequences containing 43,552 point cloud scans. We perform all our experiments using the official training (seq 00-07 and 09-10) and validation (seq 08) split. 19 classes are used for segmentation.

nuScenes [3]. nuScenes was collected in Boston and Singapore with 32-beam LiDAR sensors. It contains 1,000 scenes of 20s duration annotated with 2Hz frequency. Following the official train/val splits, we perform all label acquisition strategies on the 700 training sequences (28k scans) and evaluate them on 150 validation sequences (6k scans). 16 classes are used for segmentation.

Metrics. For evaluation, we report mean class intersection over union (mIoU) results for both the SemanticKITTI and nuScenes datasets following the official guidance.

4.2 Experimental Settings

Network Architectures. To verify the effectiveness of the proposed active learning strategy on various network architectures, we adopt MinkowskiNet [5] based on sparse convolution, and SPVCNN [37] based on point-voxel CNN, as our backbone networks for their great performance and high efficiency. We make the same choices for the network architectures as ReDAL [49], a recent state-of-the-art active learning method of 3D semantic segmentation, for better comparison.

Baseline Active Learning Methods. We select eight baseline methods for comparison, including random frame selection (RAND_{fr}), random region selection (RAND_{re}), segment-entropy (SEGENT) [20], softmax margin (MAR) [16,29,45], softmax confidence (CONF) [32,45], softmax entropy (ENT) [15,45], core-set selection (CSET) [31], and ReDAL [49]. The implementation details for all the methods are explained in **Supplementary Section B**.

Learning Protocol. Following the same protocol as ReDAL [49], the model is initialized by training on $x_{\text{init}}\%$ of randomly selected LiDAR frames with full annotations. The active learning process consists of K rounds of the following




Fig. 5: Mean intersection over union scores on SemanticKITTI Val [1]. Detailed results can be found in **Supplementary Section C**.

actions: 1. Finetune the model on the current labeled set D_L . 2. Select $x_{active}\%$ of data from the current unlabeled set D_U for human annotation according to different active selection strategies. 3. Update D_L and D_U .

The labeling budget is measured by the percentage of labeled points. For both SemanticKITTI and nuScenes datasets, we use $x_{init} = 1\%$, $K = 4$, and $x_{active} = 1\%$. For self-training, the target number of pseudo-labels in terms of percentage of labeled points $T = 1\%$. To ensure the reliability of the results, all the experiments are performed three times and the average results are reported. More training details can be found in **Supplementary Section A**.

4.3 Results and Analysis

In this section, we present the performance of our approach compared to the baseline methods on the SemanticKITTI and nuScenes datasets. Fig. 5 and Fig. 6 show the comparative results. In each subplot, the x-axis represents the percentage of labeled points and the y-axis indicates the mIoU score achieved by the respective networks, which are trained with data selected through different active learning strategies.

Since most of the baseline methods are not designed for LiDAR point clouds, we re-implement these methods for LiDAR data based on their official codes. For ReDAL [49], in its published paper, it is evaluated on the SemanticKITTI dataset but not on the nuScenes dataset. For the SemanticKITTI dataset, we find that the reported scores of its initial networks (trained with 1% of randomly selected frames) are way lower than our implementations (41.9 vs 48.8 for SPVCNN and 37.5 vs 47.3 for MinkowskiNet). We retrained its networks and got better results using its official code but with a finer training schedule (details can be found in **Supplementary Section A**). Both the retrained results and the reported results are presented in Fig. 5. For nuScenes, we adapt its official code and report the results.

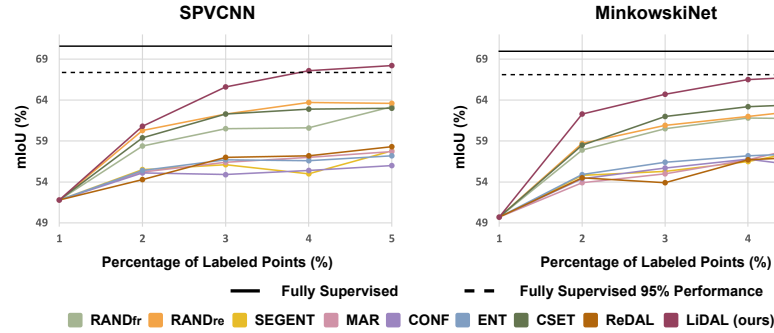


Fig. 6: **Mean intersection over union scores on nuScenes Val [3].** Detailed results can be found in **Supplementary Section C**.

SemanticKITTI. As shown in Fig. 5, our proposed LiDAL significantly surpasses the existing active learning strategies using the same percentages of labeled data. Specifically, our method is able to reach nearly 95% of the fully supervised performance with only 5% of labeled data for the SPVCNN network and even achieve about 98% of the fully supervised result for the MinkowskiNet network.

In addition, we notice that many active learning strategies perform worse than the random baseline and some even bring negative effects on the network performance (e.g., the mIoU scores may drop after adding the data samples selected by SEGENT). For uncertainty-based methods, such as SEGENT and CONF, since the model uncertainty values are calculated only within each frame and are biased by the peripheral points due to the scanning property of LiDAR, their performances are degraded. Even for pure diversity-based approaches, such as CSET, since the LiDAR datasets are captured as continuous frames and have plenty of redundant information among neighboring frames, simply clustering features may fail to produce diverse label acquisition.

Moreover, we observe that the performance gap between $RAND_{re}$ and $RAND_{fr}$ is trivial. It showcases that if not combined with effective uncertainty measures, region-based training brings little benefit to the network performance.

nuScenes. We also evaluate our algorithm on the novel nuScenes dataset and report the results in Fig. 6. As shown in the figure, our method outperforms all the competitors in terms of mIoU under all the experimental settings. Specifically, for both SPVCNN and MinkowskiNet, our method achieves more than 95% of the fully supervised performances with only 5% of labeled data.

Compared to the results of SemanticKITTI, similar phenomena can be witnessed that many active learning strategies perform worse than the random baseline. However, the negative effects are alleviated and the mIoU scores consistently increase by adding data samples selected by most strategies. A possible explanation is that, since the nuScenes dataset contains 1,000 scenes of tens of

Table 1: **Ablation study: building components.** FD: inter-frame divergence score; NMS: non-maximum suppression, i.e., select the region with the highest score in the corresponding set; FE: inter-frame entropy score.

FD	Frame-level	Region-level	NMS	FE	Pseudo	mIoU(%)
✓	✓					51.8
✓		✓				52.5
✓		✓	✓			55.5
✓		✓		✓		56.4
✓		✓		✓	✓	57.1

frames while the SemanticKITTI dataset contains only 22 scenes of thousands of frames, the network is less likely to be biased by the data selected from nuScenes than that from SemanticKITTI.

4.4 Ablation Study

In this section, we conduct a number of controlled experiments that demonstrate the effectiveness of the building modules in LiDAL, and also examine some specific decisions in our LiDAL design. All the experiments are conducted on the SemanticKITTI validation set evaluating the performance of the SPVCNN network trained on the data selected in the first active learning round, keeping all the hyper-parameters the same. More ablation studies can be found in **Supplementary Section D**.

Building Components. In Table 1, we evaluate the effectiveness of each component of our method. **1. Effect of region-level labeling.** “FD + Frame-level” represents the baseline, which is to select frames with the highest average inter-frame divergence scores for training. By changing from “FD + Frame-level” to “FD + Region-level” (selecting regions), we can improve the performance by 0.7%. This improvement is brought by focusing on the most informative parts of the scenes. **2. Effect of active selection strategy.** “FD + Region-level + NMS” refers to selecting only the region with the highest inter-frame divergence score in the corresponding set. By avoiding the label redundancy, we can gain about 3% of improvement. “FD + Region-level + FE” refers to the proposed selection strategy described in Section 3.3. From the proposed inter-frame entropy measure, we further improve about 0.9%. **3. Effect of pseudo-labels.** “FD + Region-level + FE + Pseudo” denotes the complete strategy of LiDAL. The introduction of pseudo-labels brings around 0.7% of performance improvement.

Region Selection Strategies. In Section 3.3, we discuss two possible region selection strategies for both active learning and self-training. We advocate the proposed one that first finds corresponding sets using the inter-frame divergence

Table 2: **Ablation study:** **(Left)** Region selection strategy; **(Right)** Target number of pseudo-label.

Strategy	mIoU(%)	Target Number(%)	mIoU(%)
FE + FD	55.7	0.0	56.4
FD + FE	57.1	0.5	56.8
		1.0	57.1
		2.0	56.4

scores and then selects regions with the inter-frame entropy scores. To justify our choice, we implement both strategies and report the results in Table 2 (Left). “FD + FE” refers to the proposed strategy and “FE + FD” refers to the possible alternative strategy that first finds corresponding sets using the entropy scores and then selects regions with the divergence scores. As shown in the table, the proposed strategy significantly outperforms the possible alternative strategy. It may be caused by the misleadingly high entropy values of peripheral points, as illustrated in Fig. 4.

Target Number of Pseudo-labels. In Section 3.3, we explore self-training in the proposed active learning framework and show that further gains can be realized by including pseudo-labels in Table 1. To investigate the impact of pseudo-labels, we inject different numbers of pseudo-labels into the training set and report the results in Table 2 (Right). We observe that with the increasing number of pseudo-labels, the gain of network performance first increases and then decreases. We speculate that adding pseudo-labels with a reasonable number will improve the network performance but superfluous pseudo-labels may bring unhelpful training biases and label noises. A further study on pseudo-labels can be found in **Supplementary Section D**.

5 Conclusion

In this paper, we have presented a novel active learning strategy for 3D LiDAR semantic segmentation, named LiDAL. Aiming at exploiting the inter-frame constraints embedded in LiDAR sequences, we propose two uncertainty measures estimating the inconsistencies of network predictions among frames. We design a unified framework of both active learning and self-training by utilizing the proposed measures. Extensive experiments show that LiDAL achieves state-of-the-art results on the challenging SemanticKITTI and nuScenes datasets, significantly improving over strong baselines. For future works, one straightforward direction is to explore the potential of inter-frame constraints for RGB-D sequences of indoor scenes. Moreover, we believe that future works with special treatments for moving objects will further improve the performance.

Acknowledgements. This work is supported by Hong Kong RGC GRF 16206722 and a grant from City University of Hong Kong (Project No. 7005729).

References

1. Behley, J., Garbade, M., Milioto, A., Quenzel, J., Behnke, S., Stachniss, C., Gall, J.: Semantickitti: A dataset for semantic scene understanding of lidar sequences. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9297–9307 (2019)
2. Bradley, P.S., Bennett, K.P., Demiriz, A.: Constrained k-means clustering. Microsoft Research, Redmond **20**(0), 0 (2000)
3. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
4. Cheng, M., Hui, L., Xie, J., Yang, J.: Sspc-net: Semi-supervised semantic 3d point cloud segmentation network. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 1140–1147 (2021)
5. Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3075–3084 (2019)
6. Feng, Q., He, K., Wen, H., Keskin, C., Ye, Y.: Active learning with pseudo-labels for multi-view 3d pose estimation. arXiv preprint arXiv:2112.13709 (2021)
7. Górriz, M., Giró Nieto, X., Carlier, A., Faure, E.: Cost-effective active learning for melanoma segmentation. In: ML4H: Machine Learning for Health NIPS, Workshop at NIPS 2017. pp. 1–5 (2017)
8. Hassani, K., Haley, M.: Unsupervised multi-task feature learning on point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8160–8171 (2019)
9. Hou, J., Graham, B., Nießner, M., Xie, S.: Exploring data-efficient 3d scene understanding with contrastive scene contexts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15587–15597 (2021)
10. Hu, Q., Yang, B., Fang, G., Guo, Y., Leonardis, A., Trigoni, N., Markham, A.: Sqn: Weakly-supervised semantic segmentation of large-scale 3d point clouds with 1000x fewer labels. arXiv preprint arXiv:2104.04891 (2021)
11. Hu, Q., Yang, B., Xie, L., Rosa, S., Guo, Y., Wang, Z., Trigoni, N., Markham, A.: Randla-net: Efficient semantic segmentation of large-scale point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11108–11117 (2020)
12. Hu, Z., Bai, X., Shang, J., Zhang, R., Dong, J., Wang, X., Sun, G., Fu, H., Tai, C.L.: Vmnet: Voxel-mesh network for geodesic-aware 3d semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15488–15498 (2021)
13. Hu, Z., Zhen, M., Bai, X., Fu, H., Tai, C.L.: Jsenet: Joint semantic segmentation and edge detection network for 3d point clouds. In: European Conference on Computer Vision. pp. 222–239. Springer (2020)
14. Huang, S., Xie, Y., Zhu, S.C., Zhu, Y.: Spatio-temporal self-supervised representation learning for 3d point clouds. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6535–6545 (2021)
15. Hwa, R.: Sample selection for statistical parsing. Computational linguistics **30**(3), 253–276 (2004)
16. Joshi, A.J., Porikli, F., Papanikolopoulos, N.: Multi-class active learning for image classification. In: 2009 IEEE conference on computer vision and pattern recognition. pp. 2372–2379. IEEE (2009)

17. Langer, F., Milioto, A., Haag, A., Behley, J., Stachniss, C.: Domain transfer for semantic segmentation of lidar data using deep neural networks. In: 2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 8263–8270. IEEE (2020)
18. Li, B., Zhang, T., Xia, T.: Vehicle detection from 3d lidar using fully convolutional network. arXiv preprint arXiv:1608.07916 (2016)
19. Li, H., Yin, Z.: Attention, suggestion and annotation: a deep active learning framework for biomedical image segmentation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 3–13. Springer (2020)
20. Lin, Y., Vosselman, G., Cao, Y., Yang, M.: Efficient training of semantic point cloud segmentation via active learning. *ISPRS Annals of Photogrammetry, Remote Sensing & Spatial Information Sciences* **5**(2) (2020)
21. Lin, Y., Vosselman, G., Cao, Y., Yang, M.Y.: Active and incremental learning for semantic als point cloud segmentation. *ISPRS journal of photogrammetry and remote sensing* **169**, 73–92 (2020)
22. Liu, W., Luo, Z., Cai, Y., Yu, Y., Ke, Y., Junior, J.M., Gonçalves, W.N., Li, J.: Adversarial unsupervised domain adaptation for 3d semantic segmentation with multi-modal learning. *ISPRS Journal of Photogrammetry and Remote Sensing* **176**, 211–221 (2021)
23. Liu, Y., Hu, Q., Lei, Y., Xu, K., Li, J., Guo, Y.: Box2seg: Learning semantics of 3d point clouds with box-level supervision. arXiv preprint arXiv:2201.02963 (2022)
24. Liu, Y., Yi, L., Zhang, S., Fan, Q., Funkhouser, T., Dong, H.: P4contrast: Contrastive learning with pairs of point-pixel pairs for rgb-d scene understanding. arXiv e-prints pp. arXiv–2012 (2020)
25. Liu, Z., Qi, X., Fu, C.W.: 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. pp. 1726–1736 (2021)
26. Luo, H., Wang, C., Wen, C., Chen, Z., Zai, D., Yu, Y., Li, J.: Semantic labeling of mobile lidar point clouds via active learning and higher order mrf. *IEEE Transactions on Geoscience and Remote Sensing* **56**(7), 3631–3644 (2018)
27. Mackowiak, R., Lenz, P., Ghorri, O., Diego, F., Lange, O., Rother, C.: Cereals-cost-effective region-based active learning for semantic segmentation. In: BMVC (2018)
28. Ren, Z., Misra, I., Schwing, A.G., Girdhar, R.: 3d spatial recognition without spatially labeled 3d. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13204–13213 (2021)
29. Roth, D., Small, K.: Margin-based active learning for structured output spaces. In: European Conference on Machine Learning. pp. 413–424. Springer (2006)
30. Sauder, J., Sievers, B.: Self-supervised deep learning on point clouds by reconstructing space. *Advances in Neural Information Processing Systems* **32** (2019)
31. Sener, O., Savarese, S.: Active learning for convolutional neural networks: A core-set approach. arXiv preprint arXiv:1708.00489 (2017)
32. Settles, B., Craven, M.: An analysis of active learning strategies for sequence labeling tasks. In: proceedings of the 2008 conference on empirical methods in natural language processing. pp. 1070–1079 (2008)
33. Sharma, C., Kaul, M.: Self-supervised few-shot learning on point clouds. *Advances in Neural Information Processing Systems* **33**, 7212–7221 (2020)
34. Shi, X., Xu, X., Chen, K., Cai, L., Foo, C.S., Jia, K.: Label-efficient point cloud semantic segmentation: An active learning approach. arXiv preprint arXiv:2101.06931 (2021)

35. Siddiqui, Y., Valentin, J., Nießner, M.: Viewal: Active learning with viewpoint entropy for semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9433–9443 (2020)
36. Sun, W., Tagliasacchi, A., Deng, B., Sabour, S., Yazdani, S., Hinton, G., Yi, K.M.: Canonical capsules: Unsupervised capsules in canonical pose. arXiv preprint arXiv:2012.04718 (2020)
37. Tang, H., Liu, Z., Zhao, S., Lin, Y., Lin, J., Wang, H., Han, S.: Searching efficient 3d architectures with sparse point-voxel convolution. In: European conference on computer vision. pp. 685–702. Springer (2020)
38. Tao, A., Duan, Y., Wei, Y., Lu, J., Zhou, J.: Seggroup: Seg-level supervision for 3d instance and semantic segmentation. arXiv preprint arXiv:2012.10217 (2020)
39. Thabet, A.K., Alwassel, H., Ghanem, B.: Mortonnet: Self-supervised learning of local features in 3d point clouds (2019)
40. Thomas, H., Qi, C.R., Deschaud, J.E., Marcotegui, B., Goulette, F., Guibas, L.J.: Kpconv: Flexible and deformable convolution for point clouds. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6411–6420 (2019)
41. Thrun, S., Montemerlo, M., Dahlkamp, H., Stavens, D., Aron, A., Diebel, J., Fong, P., Gale, J., Halpenny, M., Hoffmann, G., et al.: Stanley: The robot that won the darpa grand challenge. *Journal of field Robotics* **23**(9), 661–692 (2006)
42. Vezhnevets, A., Buhmann, J.M., Ferrari, V.: Active learning for semantic segmentation with expected change. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3162–3169. IEEE (2012)
43. Wang, H., Rong, X., Yang, L., Feng, J., Xiao, J., Tian, Y.: Weakly supervised semantic segmentation in 3d graph-structured point clouds of wild scenes. arXiv preprint arXiv:2004.12498 (2020)
44. Wang, H., Liu, Q., Yue, X., Lasenby, J., Kusner, M.J.: Unsupervised point cloud pre-training via occlusion completion. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9782–9792 (2021)
45. Wang, K., Zhang, D., Li, Y., Zhang, R., Lin, L.: Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology* **27**(12), 2591–2600 (2016)
46. Wei, J., Lin, G., Yap, K.H., Hung, T.Y., Xie, L.: Multi-path region mining for weakly supervised 3d semantic segmentation on point clouds. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 4384–4393 (2020)
47. Wei, J., Lin, G., Yap, K.H., Liu, F., Hung, T.Y.: Dense supervision propagation for weakly supervised semantic segmentation on 3d point clouds. arXiv preprint arXiv:2107.11267 (2021)
48. Wu, B., Zhou, X., Zhao, S., Yue, X., Keutzer, K.: Squeezesegv2: Improved model structure and unsupervised domain adaptation for road-object segmentation from a lidar point cloud. In: 2019 International Conference on Robotics and Automation (ICRA). pp. 4376–4382. IEEE (2019)
49. Wu, T.H., Liu, Y.C., Huang, Y.K., Lee, H.Y., Su, H.T., Huang, P.C., Hsu, W.H.: Redal: Region-based and diversity-aware active learning for point cloud semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15510–15519 (2021)
50. Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In: European conference on computer vision. pp. 574–591. Springer (2020)

51. Xu, C., Yang, S., Zhai, B., Wu, B., Yue, X., Zhan, W., Vajda, P., Keutzer, K., Tomizuka, M.: Image2point: 3d point-cloud understanding with pretrained 2d convnets. arXiv preprint arXiv:2106.04180 (2021)
52. Xu, K., Yao, Y., Murasaki, K., Ando, S., Sagata, A.: Semantic segmentation of sparsely annotated 3d point clouds by pseudo-labelling. In: 2019 International Conference on 3D Vision (3DV). pp. 463–471. IEEE (2019)
53. Xu, X., Lee, G.H.: Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13706–13715 (2020)
54. Yi, L., Gong, B., Funkhouser, T.: Complete & label: A domain adaptation approach to semantic segmentation of lidar point clouds. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15363–15373 (2021)
55. Zhang, Y., Li, Z., Xie, Y., Qu, Y., Li, C., Mei, T.: Weakly supervised semantic segmentation for large-scale point cloud. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 35, pp. 3421–3429 (2021)
56. Zhang, Z., Girdhar, R., Joulin, A., Misra, I.: Self-supervised pretraining of 3d features on any point-cloud. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10252–10263 (2021)
57. Zhu, X., Zhou, H., Wang, T., Hong, F., Ma, Y., Li, W., Li, H., Lin, D.: Cylindrical and asymmetrical 3d convolution networks for lidar segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9939–9948 (2021)