

ReDAL: Region-based and Diversity-aware Active Learning for Point Cloud Semantic Segmentation

Tsung-Han Wu¹ Yueh-Cheng Liu^{1†} Yu-Kai Huang^{1†}
 Hsin-Ying Lee¹ Hung-Ting Su¹ Ping-Chia Huang¹ Winston H. Hsu^{1,2}

¹National Taiwan University

²Mobile Drive Technology

Abstract

Despite the success of deep learning on supervised point cloud semantic segmentation, obtaining large-scale point-by-point manual annotations is still a significant challenge. To reduce the huge annotation burden, we propose a Region-based and Diversity-aware Active Learning (ReDAL), a general framework for many deep learning approaches, aiming to automatically select only informative and diverse sub-scene regions for label acquisition. Observing that only a small portion of annotated regions are sufficient for 3D scene understanding with deep learning, we use softmax entropy, color discontinuity, and structural complexity to measure the information of sub-scene regions. A diversity-aware selection algorithm is also developed to avoid redundant annotations resulting from selecting informative but similar regions in a querying batch. Extensive experiments show that our method highly outperforms previous active learning strategies, and we achieve the performance of 90% fully supervised learning, while less than 15% and 5% annotations are required on S3DIS and SemanticKITTI datasets, respectively.

1. Introduction

Point cloud semantic segmentation is crucial for various emerging applications such as indoor robotics and autonomous driving. Many supervised approaches [19, 20, 30, 27, 6, 26] along with several large-scale datasets [1, 7, 10, 5] are recently provided and have made huge progress.

Although recent deep learning methods have achieved great success with the aid of massive datasets, obtaining a large-scale point-by-point labeled dataset is still costly and challenging. Specifically, the statistics show that there would be more than 100,000 points in a room-sized point cloud scene [1, 5]. Furthermore, the annotation process of 3D point-wise data is much more complicated than that of

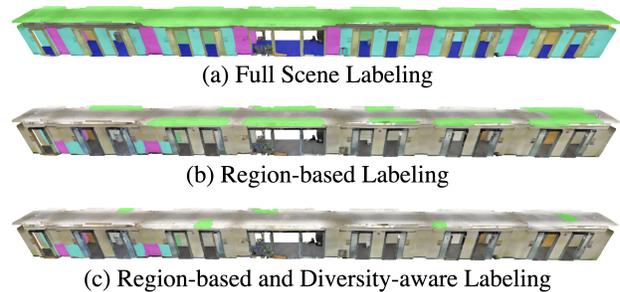


Figure 1. **Human labeling efforts (colored areas) of different learning strategies.** (a) In supervised training or traditional deep active learning, all points in a single point cloud are required to be labeled, which is labor-intensive. (b) Since few regions contribute to the model improvement, our region-based active learning strategy selects only a small portion of informative regions for label acquisition. Compared with case (a), our approach greatly reduces the cost of semantic labeling of walls and floors. (c) Moreover, considering the redundant labeling where repeating visually similar regions in the same querying batch, we develop a diversity-aware selection algorithm to further reduce redundant labeling (e.g., ceiling colored in green in (b) and (c)) effort by penalizing visually similar regions.

2D data. Unlike simply selecting closed polygons to form a semantic annotation in a 2D image [22], in 3D point-by-point labeling, annotators are asked to perform multiple 2D annotations from different viewpoints during the annotation process [10] or to label on 3D space with brushes through multiple zooming in and out and switching the brush size [5]. Therefore, such numerous points and the complicated annotation process significantly increase the time and cost of manual point-by-point labeling.

To alleviate the huge burden of manual point-by-point labeling in large-scale point cloud datasets, some previous works have tried to reduce the total number of labeled point cloud scans [14] or lower the annotation density within a

[†]Co-second authors contribute equally.

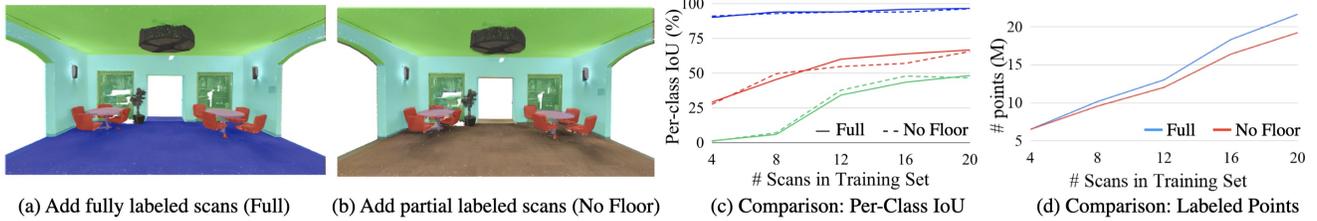


Figure 2. **Not all annotated regions contribute to the model’s improvement.** This toy experiment compares the performance contribution of fully labeled (a) and partially (b, w/o floor) labeled scans on S3DIS [1] dataset. Specifically, the training dataset contains only 4 fully-labeled point cloud scans at the beginning. Another 4 fully or partially labeled scans are then added into the dataset at each following iteration. As shown in (c), compared to using all labels (solid line), removing floor labels (dash line) leads to similar performance on all classes including floor (blue), chairs (red), and bookcases (green). Additionally, (d) demonstrate that 12% of point annotation (21.7M fully labeled points versus 19.1M partially labeled points at 20 scans) is saved by simply removing the floor labels. Therefore, this shows that not all annotated regions contribute to the model’s improvement, and we can save the annotation costs by selecting key regions to annotate while maintaining the original performance.

single point cloud scan [34]. However, they neglect that regions in a point cloud scan may not contribute to the performance equally. As can be observed from Figure 2, for a deep learning model, only 4 labeled point cloud scans are needed to reach over 0.9 IoU on large uniform objects, such as floors. However, 20 labeled scans are required to achieve 0.5 IoU on small items or objects with complex shapes and colors, like chairs and bookcases. Therefore, we argue that an effective point selection is essential for lowering annotation costs while preserving model performance.

In this work, we propose a novel *Region-based and Diversity-aware Active Learning* (ReDAL) framework general for many deep learning network architectures. By actively selecting data from a huge unlabeled dataset for label acquisition, only a small portion of informative and diverse sub-scene regions is required to be labeled.

To find out the most informative regions for label acquisition, we utilize the combination of the three terms, *softmax entropy*, *color discontinuity*, and *structural complexity*, to calculate the information score of each region. Softmax entropy is a widely used approach to measure model uncertainty, and areas with large color differences or complex structures in a point cloud provide more information because semantic labels are usually not smooth in these areas. As shown in the comparison of Figure 1 (a, b), the region-based active selection strategy significantly reduces the annotation effort of original full scene labeling.

Furthermore, to avoid redundant annotation resulting from multiple individually informative but duplicated data in a query batch, which is a common problem in deep active learning, we develop a novel diversity-aware selection algorithm considering both region information and diversity. In our proposed method, we first extract all regions’ features, then measure the similarity between regions in the feature space, and finally, use a greedy algorithm to penalize multiple similar regions appearing in the same querying batch.

As can be observed from the comparison of Figure 1 (b, c), our region-based and diversity-aware selection strategy can avoid querying labels for similar regions and further reduce the effort of manual labeling.

Experimental results demonstrate that our proposed method significantly outperforms existing deep active learning approaches on both indoor and outdoor datasets with various network architectures. On S3DIS [1] and SemanticKITTI [5] datasets, our proposed method can achieve the performance of 90% fully supervised learning, while less than 15%, 5% annotations are required. Our ablation studies also verify the effectiveness of each component in our proposed method.

To sum up, our contributions are highlighted as follows,

- We pave a new path for 3D deep active learning that utilizes region segmentation as the basic query unit.
- We design a novel diversity-aware active selection approach to avoid redundant annotations effectively.
- Experimental results show that our method can highly reduce human annotation effort on different state-of-the-art deep learning networks and datasets, and outperforms existing deep active learning methods.

2. Related Work

2.1. Point Cloud Semantic Segmentation with less labeled data

In the past decade, many supervised point cloud semantic segmentation approaches have been proposed [13, 19, 20, 30, 27, 6, 3, 15, 26]. However, despite the continuous development of supervised learning algorithms and the simplicity of collecting 3D point cloud data in large scenes, the cost of obtaining manual point-by-point marking is still high. As a result, many researchers began to study how to achieve similar performance with less labeled data.

Some have tried to apply transfer learning to this task. Wu *et al.* [33] developed an unsupervised domain adaptation approach to make the model perform well in real-world scenarios given only synthetic training sets. However, their method can only be applied to a single network architecture [32] instead of a general framework.

Some others applied weakly supervised learning to reduce the cost of labeling. [34] utilized gradient approximation along with spatial and color smoothness constraints for training with few labeled scattered points. However, this operation does not save much cost, since annotators still have to switch viewpoints or zoom in and out throughout a scene when labeling scattered points. Besides, [31] designed a multi-path region mining module to help the classification model learn local cues and to generate pseudo point-wise labels at subcloud-level, but their performance is still far from the current state-of-the-art method compared with the fully-supervised result.

Still some others leveraged active learning to alleviate the annotation burden. [16] designed an active learning approach to reduce the workload of CRF-based semantic labeling. However, their method can not be applied to current large-scale datasets for two reasons. First, the algorithm highly relies on the result of over-segmentation pre-processing and the algorithm cannot perfectly cut out small blocks with high purity in the increasingly complex scenes of the current real-world datasets. Second, the computation of pair-wise CRF is extremely high and thus not suitable for large-scale datasets. In addition to the above practice, [14] proposed segment entropy to measure the informativeness of single point cloud scan in a deep active learning pipeline.

To the best of our knowledge, we are the first to design a region-based active learning framework general for many deep learning models. Furthermore, our idea of reducing redundant annotation through diversity-aware selection is totally different from those previous works.

2.2. Deep Active Learning

Sufficient labeled training data is vital for supervised deep learning models, but the cost of manual annotation is often high. Active Learning [24] aims to reduce labeling cost by selecting the most valuable data for label acquisition. [28] proposed the first active learning framework on deep learning where a batch of items, rather than a single sample in traditional active learning, is queried in each active selection for acceleration.

Several past deep active learning practices are based on model uncertainty. [28] is the first work that applied least confidence, smallest margin [21] and maximum entropy [25] to deep active learning. [29] introduced semi-supervision to active learning, which assigned pseudo-labels for instances with the highest certainty. [8, 9] combined bayesian active learning with deep learning, which

estimated model uncertainty by MC-Dropout.

In addition to model uncertainty, many recent deep active learning works take in-batch data diversity into account. [23, 12, 2] stated that neglecting data correlation would cause similar items to appear in the same querying batch, which further leads to inefficient training. [23] converted batch selection into a core-set construction problem to ensure diversity in labeled data; [12, 2] tried to consider model uncertainty and data diversity at the same time. Empirically, uncertainty and diversity are two key indicators of active learning. [11] is a hybrid method that enjoys the benefit of both by dynamically choosing the best query strategy in each active selection step.

To the best of our knowledge, we design the first 3D deep active learning framework combining uncertainty, diversity and point cloud domain knowledge.

3. Method

In this section, we describe our region-based and diversity-aware active learning pipeline in detail. Initially, we have a 3D point cloud dataset D , which can be divided into two parts. One is a subset D_L containing randomly selected point cloud scans with complete annotations, and the other is a large unlabeled set D_U without any annotation.

In traditional deep active learning, the network is trained on the current labeled set D_L under supervision initially. Then, select a batch of data for label acquisition from the unlabeled set D_U according to a certain strategy. Finally, move the newly labeled data from D_U to D_L ; then, go back to step one to re-train or fine-tune the network and repeat the loop until the budget of the annotation is exhausted.

3.1. Overview

We use a sub-scene region as the fundamental query unit in our proposed ReDAL method. In traditional deep active learning, the smallest unit for label querying is a sample, which is a whole point cloud scan in our task. However, based on the prior experiment shown in Figure 2, we know that some labeled regions may contribute little to the model improvement. Therefore, we change the fundamental unit of label querying from a point cloud scan to a sub-scene region in a scan.

Instead of using model uncertainty as the only criterion to determine the selection common in 2D active learning, we leverage the domain knowledge from 3D computer vision and include two informative cues, *color discontinuity* and *structural complexity*, in the selection indicators. Moreover, to avoid redundant labeling caused by multiple duplicate regions in a querying batch, we design a simple yet effective diversity-aware selection strategy to mitigate the problem and improve the performance.

Our region-based and diversity-aware active learning can be divided into 4 steps: (1) Train on current labeled dataset

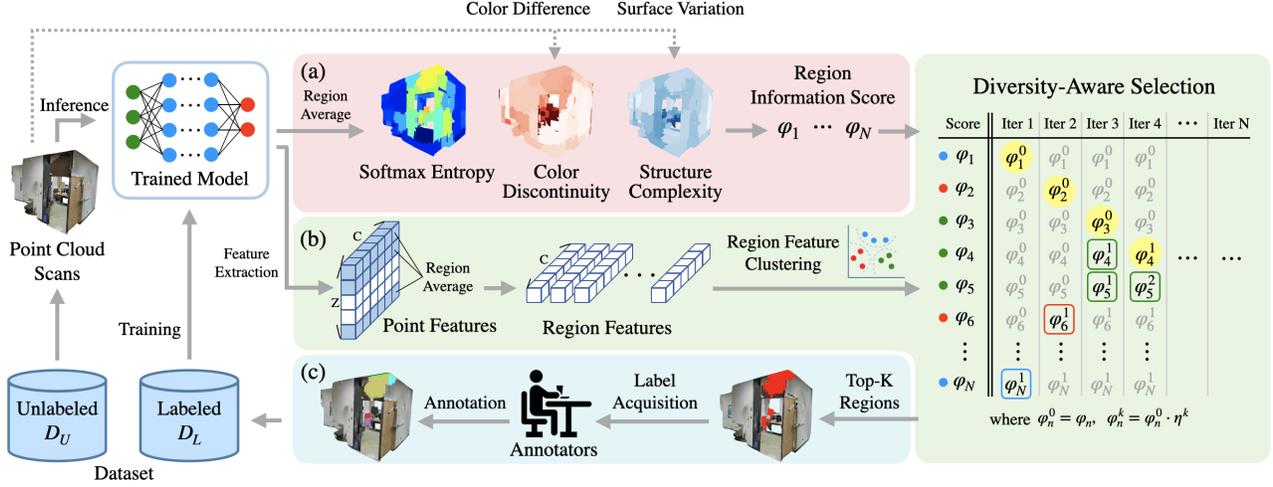


Figure 3. **Region-based and Diversity-Aware Active Learning Pipeline.** In the proposed framework, a point cloud semantic segmentation model is first trained in supervision with labeled dataset D_L . The model then produces softmax entropy and features of all regions from the unlabeled dataset D_U . (a) Softmax entropy along with color discontinuity and structure complexity calculated from the unlabeled regions serves as selection indicators (Sec. 3.2), and (b) generates scores which are then adjusted by penalizing regions belonging to the same clusters grouped by the extracted features (Sec. 3.3). (c) The top-ranked regions are labeled by annotators and added to the labeled dataset D_L for the next phase (Sec. 3.4).

D_L in a supervised manner. (2) Calculate the region information score φ for each region with three indicators: softmax entropy, structure complexity and color discontinuity as shown in Figure 3 (a) (Sec. 3.2). (3) Perform diversity-aware selection by measuring the similarity between all regions and using a greedy algorithm to penalize similar regions appearing in a querying batch as shown in Figure 3 (b) (Sec. 3.3). (4) Select top-K regions for label acquisition, and move them from the unlabeled dataset D_U into the current labeled dataset D_L as shown in Figure 3 (c) (Sec. 3.4).

3.2. Region Information Estimation

We divide a large-scale point cloud scan into some sub-scene regions as the fundamental label querying units using VCCS [17] algorithm, an unsupervised over-segmentation method that groups similar points into a region. The original purpose of this algorithm was to cut a point cloud into multiple small regions with high segmentation purity to reduce the computational burden of the probability statistical model. Different from the original purpose of requiring high purity, our method merely utilizes the algorithm to divide a scan into sub-scenes of median size for better annotation and learning. An ideal sub-scene consists of several but not complicated semantic meanings, while preserving geometric structures of point cloud.

In each active selection step, we calculate the information score of a region from three aspects: (1) softmax entropy, (2) color discontinuity, and (3) structural complexity, which is described in detail as follows.

3.2.1 Softmax Entropy

Softmax entropy is a widely used approach to measure the uncertainty in active learning [28, 29]. We first obtain the softmax probability of all point cloud scans in the unlabeled set D_U with the model trained in the previous active learning phase. Then, given the softmax probability P of a point cloud scan, we calculate the region entropy H_n for the n -th region R_n by averaging the entropy of points belonging to the region R_n as shown in Eq. 1.

$$H_n = \frac{1}{|R_n|} \sum_{i \in R_n} -P_i \log P_i \quad (1)$$

3.2.2 Color Discontinuity

In 3D computer vision, the color difference is also an important clue since areas with large color differences are more likely to indicate semantic discontinuity. Therefore, it is also included as an indicator for measuring regional information. For all points in a given point cloud with color intensity value I , we compute the 1-norm color difference between a point i and its k -nearest neighbor d_i ($|d_i| = k$). Then we produce the region color discontinuity score C_n for the n -th region R_n by averaging the values of points belonging to the region R_n as shown in Eq. 2.

$$C_n = \frac{1}{k \cdot |R_n|} \sum_{i \in R_n} \sum_{j \in d_i} \|I_i - I_j\|_1 \quad (2)$$

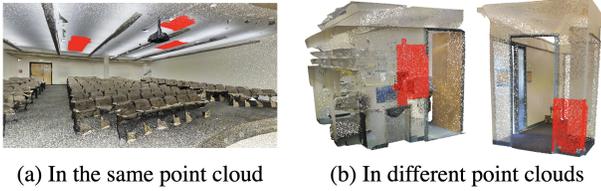


Figure 4. **Our method is able to find out visually similar regions not only in the same point cloud (a) but also in different point clouds (b).** The areas colored in red are the ceiling in an auditorium (a) and walls next to the door (b). These regions may cause redundant labeling effort if appearing in the same querying batch, and thus they are filtered by our diversity-aware selection (Sec. 3.3.1).

3.2.3 Structural Complexity

We also include structure complexity as an indicator, since complex surface regions, boundary places, or corners in a point cloud are more likely to indicate semantic discontinuity. For all points in a given point cloud, we first compute the surface variation σ based on [4, 18]. Then, we calculate the region structure complexity score S_n for the n -th region R_n by averaging surface variation of points belonging to the region R_n as shown in Eq. 3.

$$S_n = \frac{1}{|R_n|} \sum_{i \in R_n} \sigma_i \quad (3)$$

After calculating the softmax entropy, color discontinuity, and structural complexity of each region, we combine these terms linearly to form the region information score φ_n of the n -th region R_n as shown in Eq. 4.

$$\varphi_n = \alpha H_n + \beta C_n + \gamma S_n \quad (4)$$

Finally, we rank all regions in descending order based on the region information scores and produce a sorted information list $\varphi = (\varphi_1, \varphi_2, \dots, \varphi_N)$. The above process is illustrated in Figure 3 (a).

3.3. Diversity-aware Selection

With the sorted region information list φ , a naive way is to select the top-ranked regions for label acquisition directly. Nevertheless, this strategy results in multiple visually similar regions being in the same batch as shown in Figure 4. These regions, though informative individually, provide less diverse information for the model.

To avoid visually similar regions appearing in a querying batch, we design a diversity-aware selection algorithm divided into two parts: (1) region similarity measurement and (2) similar region penalization.

3.3.1 Region Similarity Measurement

We measure the similarity among regions in the feature space rather than directly on point cloud data because the scale, shape, and color of each region are totally different.

Given a point cloud scan with Z points, we record the output before the final classification layer as the point features with shape $Z \times C$. Then, we produce the region features by calculating the average of the point features of the points belonging to the same region. Finally, we gather all point cloud regions and use k -means algorithm to cluster these region features. The above process can be seen in the middle of Figure 3 (b). After clustering regions, we regard regions belonging to the same cluster as similar regions. An example is shown in Figure 4.

3.3.2 Similar Region Penalization

To select diverse regions, a greedy algorithm takes the sorted list of information scores $\varphi \in \mathbb{R}^N$ as input and re-scores all regions by penalizing regions with lower scores that belong to the same clusters containing regions with higher scores.

The table in the right of Figure 3 (b) offers an example where the algorithm loops through all regions one by one. The scores of regions ranked below while belonging to the same cluster as the current region are multiplied by a decay rate η . To be specific, the red, green, and blue dots in the left of the table denote the cluster indices of regions, and φ_n^k denotes the score φ_n of the region R_n with k -time penalization. Yellow circles under φ_n^k indicate the current region to be compared in each iteration, and rounded rectangles mark regions belonging to the same cluster as the current region. In the first iteration, φ_N is penalized as R_N and R_0 belong to the same cluster denoted by blue dots. R_N 's score is then replaced by φ_N^1 to mark the first decay. In the third iteration, φ_4 and φ_5 are both penalized as R_3 , R_4 and R_5 belong to the same cluster denoted by green dots. Their scores are thus substituted by φ_4^1 and φ_5^1 . The same logic applies to the other iterations. Then we obtain the adjusted scores φ_N^* for label acquisition.

Note that in our implementation shown in Algorithm 1, we penalize the corresponding importance weight W , which are initialized to 1 for all M clusters, instead of directly penalizing the score for efficiency. Precisely, in each iteration, we adjust the score of the pilot by multiplying the importance weight of its cluster. Then, the importance weight of the cluster is multiplied by the decay rate η .

3.4. Region Label Acquisition

After getting the final scores φ^* by considering region diversity, we select regions into a querying batch in decreasing order according to φ^* for label acquisition until the budget of this round is exhausted. Note that in each label ac-

Algorithm 1: Similar Region Penalization

Input: Original sorted information score $\varphi \in \mathbb{R}^N$
and corresponding M -cluster region labels
 $L \in \mathbb{R}^N$; cluster importance weight
 $W \in \mathbb{R}^M$ and decay rate η

Output: Final Region information score $\varphi^* \in \mathbb{R}^N$

Init: $W_m \leftarrow 1 \quad \forall 1 \leq m \leq M$;

for $i \leftarrow 1$ **to** N **do**

$\varphi_i^* \leftarrow \varphi_i \cdot W_{L_i}$;
 $W_{L_i} \leftarrow W_{L_i} \cdot \eta$;

end

return φ^*

quisition step, we set the budget as a fixed number of total points instead of a fixed number of regions for fair comparison since each region contains different number of points.

For experiments, after selecting the querying batch data, we regard the ground truth region annotation as the labeled data obtained from human annotators. Then, these regions are moved from unlabeled set D_U to labeled set D_L . Note that different from the 100% fully labeled initial training point cloud scans, since we regard a region as the basic labeling unit, many point cloud scans with only a small portion of labeled region are appended to the labeling data set D_L in every active selection step as shown in Figure 3 (c).

Finishing the active selection step containing region information estimation, diversity-aware selection and region label acquisition, we repeat the active learning loop to finetune the network on the updated labeled dataset D_L .

4. Experiments

4.1. Experimental Settings

In order to verify the effectiveness and universality of our proposed active selection strategy, we conduct experiments on two different large-scale datasets and two different network architectures. The implementation details are explained in the supplementary material due to limited space.

Datasets. We use S3DIS [1] and SemanticKITTI [5] as representatives of indoor and outdoor scenes, respectively. S3DIS is a commonly used indoor scene segmentation dataset. The dataset can be divided into 6 large areas, with a total of 271 rooms. Each room has a corresponding dense point cloud with color and position information. We evaluate the performance of all label acquisition strategies on the Area5 validation set and perform active learning on the remaining datasets. SemanticKITTI is a large-scale autonomous driving dataset with 43552 point cloud scans from 22 sequences. Each point cloud scan is captured by LiDAR sensors with only position information. We evalu-

ate the performance of all label acquisition strategies on the official validation split (seq 08) and perform active learning on the whole official training split (seq 00~07 and 09~10).

Network Architectures. To verify an active strategy on various deep learning networks, we use MinkowskiNet [6], based on sparse convolution, and SPVCNN [26], based on point-voxel CNN thanks to the great performance on large-scale point cloud datasets with high inference speed.

Active Learning Protocol. For all experiments, we first randomly select a small portion ($x_{init}\%$) of fully labeled point cloud scans from the whole training data as the initial labeled data set D_L and treat the rest as the unlabeled set D_U . Then, we perform K rounds of the following actions: (1) Train the deep learning model on D_L in a supervised manner. (2) Select a small portion ($x_{active}\%$) of data from D_U for label acquisition according to different active selection strategies. (3) Add the newly labeled data into D_L and finetune the deep learning model.

We choose $x_{init} = 3\%$, $K = 7$, and $x_{active} = 2\%$ for S3DIS dataset, and $x_{init} = 1\%$, $K = 5$, and $x_{active} = 1\%$ for SemanticKITTI dataset. To ensure the reliability of the experimental results, we perform the experiments three times and record the average value for each setting.

4.2. Comparison among different active selection strategies.

We compare our proposed method with 7 other active selection strategies, including random point cloud scans selection (**RAND**), softmax confidence (**CONF**) [28], softmax margin (**MAR**) [28], softmax entropy (**ENT**) [28], MC-dropout (**MCDR**) [8, 9], core-set approach, (**Core-Set**) [23] and segment-entropy (**SEGENT**) [14]. The implementation details are explained in the supplementary material.

The experimental results can be seen in Figure 5. In each subplot, the x-axis means the percentage of labeled points and the y-axis indicates the mIoU achieved by the network. Our proposed ReDAL significantly surpasses other existing active learning strategies under any combination.

In addition, we observe that random selection (**RAND**) outperforms any other active learning methods except ours on four experiments. For uncertainty-based methods, such as ENT and MCDR, since the model uncertainty value is dominated by background area, the performance is not as expected. The same, for pure diversity approach, such as CSET, since the global feature is dominated by the background area, simply clustering global feature cannot produce diverse label acquisition. The experimental results further verify our proposal of changing the fundamental querying units from a scan to a region is a better choice.

On the S3DIS [1] dataset, our proposed active selection strategy can achieve more than 55 % mIoU with 15 % la-

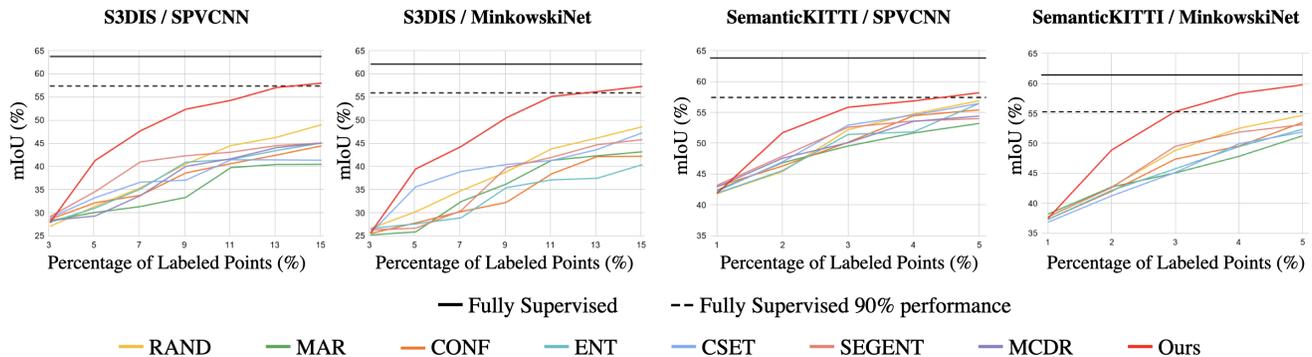


Figure 5. **Experimental results of different active learning strategies on 2 datasets and 2 network architectures.** We compare our region-based and diversity-aware active selection strategy with other existing baselines. It is obvious that our proposed method outperforms any existing active selection approaches under any combinations. Furthermore, our method is able to reach 90% fully supervised result with only 15%, 5% labeled points on S3DIS [1] and SemanticKITTI [5] dataset respectively.

Method	avg	road	person	bicycle	bicyclist
RAND	54.7	90.2	52.0	9.5	47.7
Full	61.4	93.5	65.0	20.3	78.4
ReDAL	59.8	91.5	63.4	29.5	84.1

Table 1. **Results of IoU performance (%) on SemanticKITTI [5].** Under only 5 % of annotated points, our proposed ReDAL outperforms random selection and is on par with full supervision (Full).

Method	road	person	bicycle	bicyclist
RAND	206	0.42	0.15	0.10
Full	205	0.35	0.17	0.13
ReDAL	168	1.20	0.25	0.21

Table 2. **Labeled Class Distribution Ratio (%)**. With limited annotation budgets, our active method ReDAL queries more labels on small objects like a person but less on large uniform areas like roads. The selection strategy can mitigate the label imbalance problem and improve the performance on more complicated object scenes without hurting much on large areas as shown in Table 1.

belonging points, while others cannot reach 50 % mIoU under the same condition. The main reason for such a large performance gap is that these room-sized point cloud data in the dataset is very different. Compared with other active selection methods querying a batch of point cloud scans, our region-based label acquisition make the model be trained on more diverse labeled data.

As for the SemanticKITTI [5], we find that with merely less than 5% of labeled data, our active learning strategy can achieve 90% of the result of fully supervised methods. With the network architecture of MinkowskiNet, our active selection strategy can even reach 95 % fully supervised result with only 4 % labeled points.

Furthermore, in Table 1, the performance of some small or complicated class objects like bicycles and bicyclists is even better than fully supervised one. Table 2 shows the reason that our selected algorithm focuses more on those small or complicated objects. In other words, our ReDAL does not waste annotation budgets on easy cases like uniform surfaces, which again realizes the observation and motivation in the introduction. Besides, our ReDAL selection strategy makes it more friendly to real-world applications, such as autonomous driving, since it puts more emphasis on important and valuable semantics.

4.3. Ablation Studies

We verify the effectiveness of all components in our proposed active selection strategy on S3DIS dataset [1]. The results are shown in Figure 8.

First, changing the labeling units from scans to regions contributes the most to the improvement as shown from the comparison of the purple line (ENT), the yellow line (ENT+Region), and the light blue line (RAND+Region). By applying region-based selection, the mIoU performance improves over 10 % under both networks.

Furthermore, our diversity-aware selection also plays a key role in the active selection process shown from the comparison of the yellow line (ENT+Region) and the green line (ENT+Region+Div). Without the aid of this component, the performance of region-based entropy is even lower than random region selection under SPVCNN network architecture as shown from the comparison of the yellow line (ENT+Region) and the light blue line (RAND+Region).

As for adding extra information of color discontinuity and structure complexity, it contributes little to SPVCNN but is helpful to MinkowskiNet when the percentage of labeled points is larger than 9 % as shown in the comparison of the green line (w/o color and structure) and the red line (w/ color and structure).

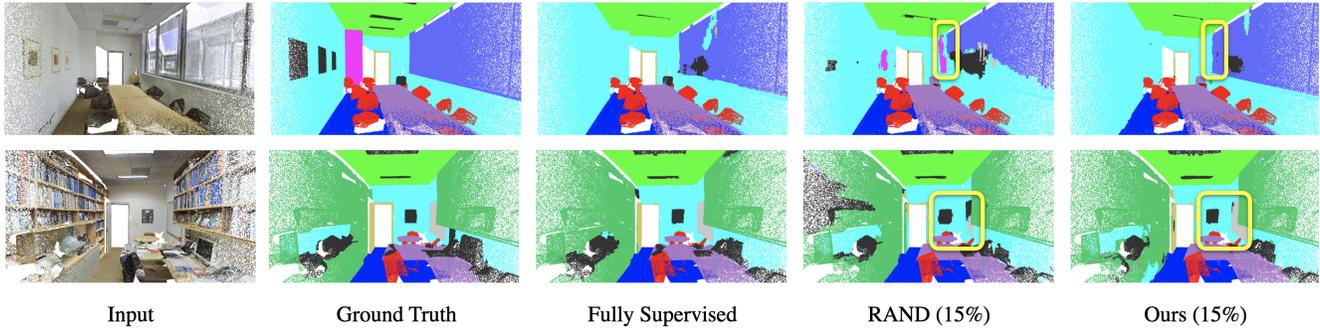


Figure 6. **Visualization for the inference result on S3DIS dataset with SPVCNN network architecture.** We show some inference examples on S3DIS Area 5 validation set. With our active learning strategy, the model can produce sharp boundaries (shown on the yellow bounding box in the first row) and recognize small objects, such as boards and chairs (shown on the yellow bounding box in the second row) with only 15 % labeled points.

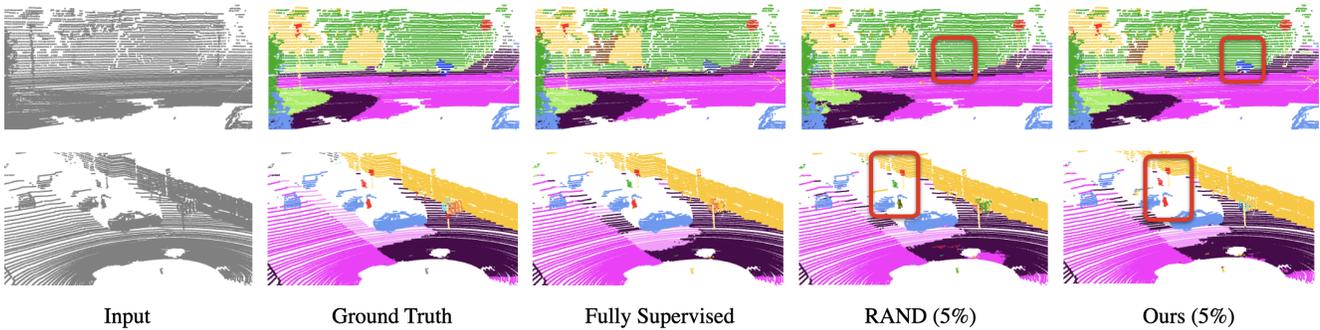


Figure 7. **Visualization for the inference result on SemanticKITTI dataset with MinkowskiNet network architecture.** We show some inference examples on the SemanticKITTI sequence 08 validation set. With our active learning strategy, the model can correctly recognize small vehicles (shown on the red bounding box in the first row) and identify people on the side walk (shown on the red bounding box in the second row) with merely 5 % labeled points.

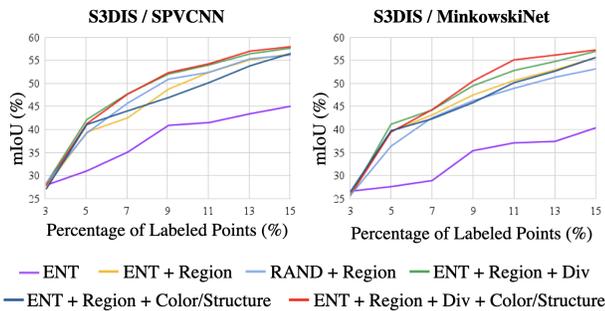


Figure 8. **Ablation Study.** The best combinations are altering labeling units from scans to regions (+Region), applying diversity-aware selection (+Div), and additional region information (+Color/Structure). Best viewed in color. (Sec. 4.3)

Note that the performance of “ENT+Region” (the yellow line) is similar to “ENT+Region+Color/Structure” (the dark blue line). The reason is that without our diversity module, the selected query batch is still full of duplicated regions. The result also validates the importance of our diversity-aware greedy algorithm.

5. Conclusion

We propose ReDAL, a region-based and diversity-aware active learning framework, for point cloud semantic segmentation. The active selection strategy considers region information and diversity, concentrating the labeling effort on the most informative and distinctive regions rather than full scenes. This approach can be applied to many deep learning network architectures and datasets, substantially reducing the annotation cost, and greatly outperforms existing active learning strategies.

Acknowledgement

This work was supported in part by the Ministry of Science and Technology, Taiwan, under Grant MOST 110-2634-F-002-026, Mobile Drive Technology (FIH Mobile Limited), and Industrial Technology Research Institute (ITRI). We benefit from NVIDIA DGX-1 AI Supercomputer and are grateful to the National Center for High-performance Computing.

References

- [1] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1534–1543, 2016.
- [2] Jordan T Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *ICLR*, 2020.
- [3] Matan Atzmon, Haggai Maron, and Yaron Lipman. Point convolutional neural networks by extension operators. *ACM Trans. Graph.*, 37(4), July 2018.
- [4] Dena Bazazian, Josep R Casas, and Javier Ruiz-Hidalgo. Fast and robust edge extraction in unorganized point clouds. In *2015 international conference on digital image computing: techniques and applications (DICTA)*, pages 1–8. IEEE, 2015.
- [5] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9297–9307, 2019.
- [6] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3075–3084, 2019.
- [7] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017.
- [8] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059, 2016.
- [9] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep bayesian active learning with image data. In *International Conference on Machine Learning*, pages 1183–1192, 2017.
- [10] Timo Hackel, Nikolay Savinov, Lubor Ladicky, Jan D Wegner, Konrad Schindler, and Marc Pollefeys. Semantic3d. net: A new large-scale point cloud classification benchmark. *arXiv preprint arXiv:1704.03847*, 2017.
- [11] Wei-Ning Hsu and Hsuan-Tien Lin. Active learning by learning. In *Twenty-Ninth AAAI conference on artificial intelligence*. Citeseer, 2015.
- [12] Andreas Kirsch, Joost van Amersfoort, and Yarin Gal. Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. In *Advances in Neural Information Processing Systems*, pages 7026–7037, 2019.
- [13] Felix Järemo Lawin, Martin Danelljan, Patrik Tosteberg, Goutam Bhat, Fahad Shahbaz Khan, and Michael Felsberg. Deep projective 3d semantic segmentation. In *International Conference on Computer Analysis of Images and Patterns*, pages 95–107. Springer, 2017.
- [14] Y Lin, G Vosselman, Y Cao, and MY Yang. Efficient training of semantic point cloud segmentation via active learning. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 2:243–250, 2020.
- [15] Zhijian Liu, Haotian Tang, Yujun Lin, and Song Han. Point-voxel cnn for efficient 3d deep learning. In *Advances in Neural Information Processing Systems*, pages 965–975, 2019.
- [16] Huan Luo, Cheng Wang, Chenglu Wen, Ziyi Chen, Dawei Zai, Yongtao Yu, and Jonathan Li. 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.
- [17] Jeremie Papon, Alexey Abramov, Markus Schoeler, and Florentin Worgotter. Voxel cloud connectivity segmentation-supervoxels for point clouds. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2027–2034, 2013.
- [18] Mark Pauly, Richard Keiser, and Markus Gross. Multi-scale feature extraction on point-sampled surfaces. In *Computer graphics forum*, volume 22, pages 281–289. Wiley Online Library, 2003.
- [19] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.
- [20] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30:5099–5108, 2017.
- [21] Nicholas Roy and Andrew McCallum. Toward optimal active learning through monte carlo estimation of error reduction. *ICML, Williamstown*, pages 441–448, 2001.
- [22] Bryan C Russell, Antonio Torralba, Kevin P Murphy, and William T Freeman. Labelme: a database and web-based tool for image annotation. *International journal of computer vision*, 77(1-3):157–173, 2008.
- [23] Ozan Sener and Silvio Savarese. Active learning for convolutional neural networks: A core-set approach. In *International Conference on Learning Representations*, 2018.
- [24] Burr Settles. Active learning literature survey. Technical report, University of Wisconsin-Madison Department of Computer Sciences, 2009.
- [25] Claude E Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [26] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3d architectures with sparse point-voxel convolution. In *European Conference on Computer Vision*, pages 685–702. Springer, 2020.
- [27] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6411–6420, 2019.

- [28] D. Wang and Y. Shang. A new active labeling method for deep learning. In *2014 International Joint Conference on Neural Networks (IJCNN)*, pages 112–119, 2014.
- [29] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. Cost-effective active learning for deep image classification. *IEEE Transactions on Circuits and Systems for Video Technology*, 27(12):2591–2600, 2016.
- [30] Lei Wang, Yuchun Huang, Yaolin Hou, Shenman Zhang, and Jie Shan. Graph attention convolution for point cloud semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 10296–10305, 2019.
- [31] Jiacheng Wei, Guosheng Lin, Kim-Hui Yap, Tzu-Yi Hung, and Lihua Xie. Multi-path region mining for weakly supervised 3d semantic segmentation on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4384–4393, 2020.
- [32] Bichen Wu, Alvin Wan, Xiangyu Yue, and Kurt Keutzer. Squeezeseg: Convolutional neural nets with recurrent crf for real-time road-object segmentation from 3d lidar point cloud. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 1887–1893. IEEE, 2018.
- [33] Bichen Wu, Xuanyu Zhou, Sicheng Zhao, Xiangyu Yue, and Kurt Keutzer. 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)*, pages 4376–4382. IEEE, 2019.
- [34] Xun Xu and Gim Hee Lee. Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13706–13715, 2020.