

# BaSAL: Size-Balanced Warm Start Active Learning for LiDAR Semantic Segmentation

Jiarong Wei, Yancong Lin\* and Holger Caesar

**Abstract**—Active learning strives to reduce the need for costly data annotation, by repeatedly querying an annotator to label the most informative samples from a pool of unlabeled data, and then training a model from these samples. We identify two problems with existing active learning methods for LiDAR semantic segmentation. First, they overlook the severe class imbalance inherent in LiDAR semantic segmentation datasets. Second, to bootstrap the active learning loop when there is no labeled data available, they train their initial model from randomly selected data samples, leading to low performance. This situation is referred to as the cold start problem. To address these problems we propose BaSAL, a size-balanced warm start active learning model, based on the observation that each object class has a characteristic size. By sampling object clusters according to their size, we can thus create a size-balanced dataset that is also more class-balanced. Furthermore, in contrast to existing information measures like entropy or CoreSet, size-based sampling does not require a pretrained model, thus addressing the cold start problem effectively. Results show that we are able to improve the performance of the initial model by a large margin. Combining warm start and size-balanced sampling with established information measures, our approach achieves comparable performance to training on the entire SemanticKITTI dataset, despite using only 5% of the annotations, outperforming existing active learning methods. We also match the existing state-of-the-art in active learning on nuScenes. Our code is available at: <https://github.com/Tony-WJR/BaSAL>.

## I. INTRODUCTION

**A**UTONOMOUS vehicles are often equipped with LiDAR, a time-of-flight sensor that accurately scans the environment and excels at 3D perception tasks. Semantic segmentation is a popular task for LiDAR-based perception systems whose goal is to assign a predefined class label for each point in the LiDAR scan. Existing methods on this task are mostly learning-based and evaluated on large-scale benchmark datasets like SemanticKITTI [1] and nuScenes [2], [3]. However, training such models requires a large amount of data that is expensive to label. Active learning is a machine learning technique that reduces the demand for annotation by interactively selecting informative samples and querying an oracle for annotation [4].

Research on active learning for LiDAR (semantic) segmentation [5], [6] focuses on novel information measures that quantify the importance of unlabeled data, such that a limited annotation budget is spent on the most informative samples. We identify two problems that have not been explicitly addressed in existing active learning methods for LiDAR segmentation: class imbalance and the cold start problem.

\* Corresponding author.

All authors are with the Intelligent Vehicles Group, Delft University of Technology, the Netherlands.

The performance on the LiDAR segmentation task is commonly measured using the mean Intersection over Union (mIoU) metric, which assigns the same weight to all classes. However, since there is a strong class imbalance, models are trained on fewer points of rare classes, thus yielding a lower performance in general. Consequently, points belonging to rare classes have a relatively larger impact on the mIoU metric. In contrast, the active learning budget is measured by the number of points, which treats points of all classes the same. This reveals a fundamental mismatch between the performance metric and the active learning budget. By reducing class imbalance in the sampled data, we can reduce this mismatch.

In our initial approach to address class imbalance, we tried to extract more samples of rare classes using the prediction result of pretrained networks. We utilized a model trained on a small subset of labels to classify the unlabeled pool, specifically targeting samples predicted as rare classes. Our preliminary studies, however, revealed this strategy to be ineffective due to the model's poor performance and inaccurate predictions on rare classes. To find a more effective reflection of classes, we observed that objects with identical labels typically share similar sizes. Consequently, by creating a size-balanced dataset, we inherently foster a more class-balanced dataset.

The cold start problem [7] in active learning is about which data to select first for label acquisition and model training when there is no labeled data to start with. Existing methods initialize the active learning loop by annotating randomly sampled data from a pool of unlabeled data [5], [6] for training. Then they select the most informative samples using their proposed information measure based on the outputs of the trained model. However, the performance of the initial model by random sampling is often lacking, which negatively affects subsequent active learning iterations. In contrast, we discovered that sampling objects based on their size does not require a pretrained model. It thus can be easily employed in the very first iteration of active learning, also making the selected data more class-balanced. If the initial model trained by this sampling strategy is verified to perform better than the traditional random sampling, it can be a warm start solution to the cold start problem.

Following the above intuitions, we propose BaSAL, a size-balanced warm start active learning strategy that addresses class imbalance and the cold start problem. To address class imbalance, we introduce Average Point Information (API), a metric assessing the informativeness of each object cluster based on its size. We use the API metric to create partitions

of clusters that have approximately the same amount of information and allocate the labeling budget equally to all partitions. To initiate model training, we randomly draw samples from each partition for label acquisition. For subsequent iterations of active learning, we adopt the well-established information measures, entropy [8] and CoreSet [9], to rank unlabeled data and select top-ranking clusters from each partition for further label acquisition. Experiments on SemanticKITTI show that by labeling only 5% points, we are able to achieve a comparable performance of fully supervised learning. We also outperform existing active learning methods by a large margin on SemanticKITTI. On nuScenes, our model performs competitively with the state-of-the-art LiDAL [6] method and reaches 95% performance of fully supervised learning.

In summary, our contribution can be outlined as follows:

- Based on the insight that it is non-trivial to sample objects according to their class, we instead propose a method that addresses class imbalance by sampling object clusters according to their size.
- We initialize the active learning loop by equally distributing the labeling budget to all size-based partitions and then randomly drawing samples from each partition, thus addressing the cold start problem.
- Our experiments demonstrate that BaSAL outperforms the state-of-the-art in active learning on SemanticKITTI and matches it on nuScenes.

## II. RELATED WORK

In this section we discuss the literature on active learning, class imbalance and the cold start problem in active learning. We focus on pool-based active learning that assumes a large pool of unlabeled data available [4].

### A. Active Learning

There are numerous strategies to select new samples for label acquisition, including uncertainty-based methods, diversity-based methods or combinations of multiple approaches. Uncertainty-based methods select hard examples by measuring model disagreement [10]–[12], entropy [8], [13], predicted loss [14], discriminator scores [15], or geometric distances to the decision boundary [16]. However, uncertainty-based methods tend to draw similar samples without taking diversity into account. Other works [9], [17]–[19] propose a diversity-based strategy that finds a subset of samples that best represents the entire pool. Follow-up works [20]–[22] combine model uncertainty and diversity.

Recent research on active learning for LiDAR segmentation takes similar strategies, but exploits task-specific prior knowledge. LiDAL [6] considers inconsistencies in model predictions across frames as the uncertainty measure for informative sample selection. ReDAL [5] divides a point cloud into regions and then selects diverse regions according to multiple cues, including softmax entropy, color discontinuity, and structural complexity. [23] proposes a diversity-based method for 3D object detection by enforcing both spatial diversity and temporal diversity. [24] uses detection

and prediction entropy as the information measures for the prediction and planning tasks. Similarly, we consider both uncertainty and diversity as information measures.

### B. Class Imbalance in Active Learning

Class imbalance is a common problem for datasets collected in the wild. [25] summarizes two kinds of techniques to cope with the class imbalance problem: Density-sensitive active learning and skew-specialized active learning [26]. The former assigns an informativeness score to each sample by imposing an assumption on the input space. Examples include information density [27], pre-clustering [28], and alternate density-sensitive heuristics [29]. The latter incorporates a bias towards underrepresented classes, thus resulting in a more balanced sampling [30]–[32]. [33] addresses class imbalance in active learning for visual tasks by introducing a sample balancing step that prioritizes minority classes. Most methods require a pretrained model. In contrast, we tackle class imbalance in LiDAR segmentation without pretraining by creating size-balanced partitions, as size is a characteristic trait of class.

### C. Cold Start in Active Learning

Cold start in active learning refers to the problem of which data to label first given an unlabeled pool of data [7]. The labeled data is then used to train an initial model which serves as the starting point for active learning. [34], [35] take advantage of pretrained models and cluster unlabeled samples in the embedding space. Those samples closer to the cluster centers are selected for annotation, as they better represent categories than random samples. [36] focuses primarily on selecting examples that are hard to learn via self-supervised contrastive learning, based on the assumption that if a model can not separate a sample from others, this sample is expected to exhibit typical characteristics, such as common visual patterns in vision that are shared by others. [37] proposes TypiClust that aims to find out typical examples which better represent the entire dataset. The typicality of a sample is measured by the inverse of the average Euclidean distance to its K-nearest neighbors in the feature space. [38] proposes ProbCover, a solution for cold start that maximizes the probability of covering the unlabeled set in the embedding space. [7] uses pseudo labels to pretrain models and then rank unlabeled data by uncertainty, after which the top-ranking samples are used for label acquisition. We handle the cold start problem by a simple size-based sampling.

## III. METHOD

Fig. 1 shows an overview of BaSAL, which consists of a preprocessing step and the active learning loop. We first preprocess the input point clouds to create multiple partitions (Section III-A). During the active learning loop, we use size-based sampling (Section III-B) in both the model initialization phase (warm start) and subsequent active learning iterations of the loop to ensure a size-balanced training set. After warm starting the model, we adopt the

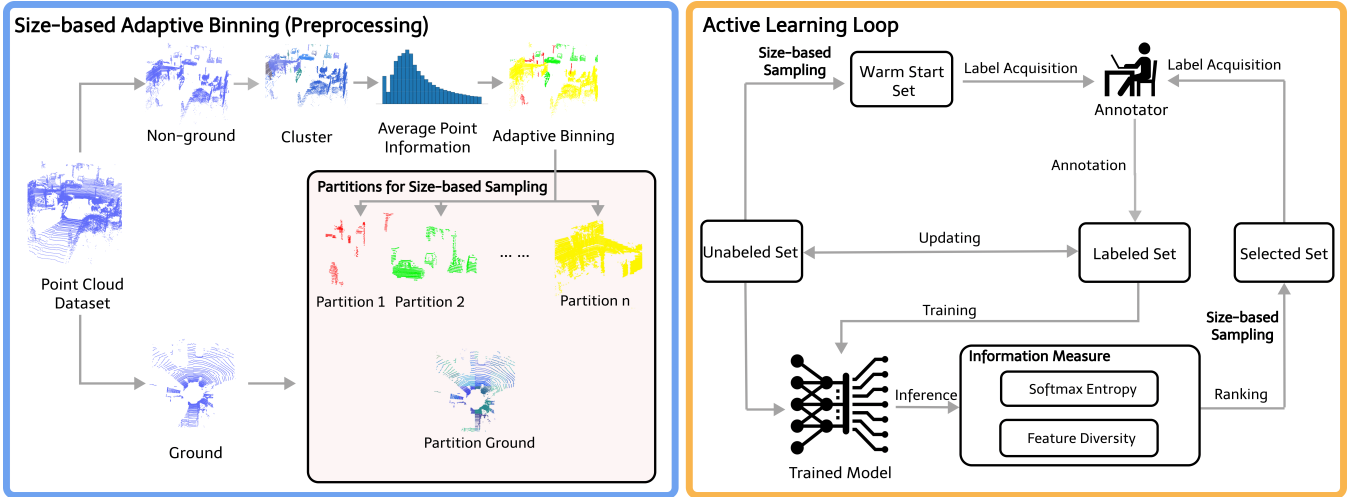


Fig. 1: **Overview of BaSAL.** Our framework consists of a preprocessing step and the active learning loop. Size-balanced sampling is used to determine both the warm start set for the initialization of the active learning loop and the selected set for subsequent active learning iterations of the active learning loop.

established information measures (Section III-C) to rank and select the most informative clusters from each partition, until the annotation budget is reached. Then we add them to the labeled set and retrain the model.

#### A. Size-based Adaptive Binning (Preprocessing)

Given a dataset  $\mathcal{D}$ , we first apply ground plane detection and obtain the ground points which are further split into a number of grids  $\mathcal{G}$ . The remaining non-ground points are split into a collection of class-agnostic clusters  $\mathcal{C}$  using HDBSCAN [39]. Based on the assumption that objects of similar sizes tend to share the same semantic labels, we group the class-agnostic clusters  $\mathcal{C}$  by size, i.e. the sum of the dimensions of the 3D bounding box that encloses a cluster. We round all sizes to integers. To balance the information among sizes, we calculate the Average Point Information (API) per size, defined by Eq. (1), where  $C_s$  is the number of clusters and  $P_s$  is the number of accumulated points from all clusters, for a particular size  $s$ .  $\frac{P_s}{C_s}$  stands for the average number of points per cluster and the additional logarithm penalizes larger clusters, as each additional point brings less information gain. We adaptively bin all sizes by API and obtain  $B$  partitions such that the accumulated API inside each partition is approximately the same. We show an example in Fig. 2, where  $B$  is set to 3.

$$API_s = \left( \frac{\log P_s}{C_s} \right)^{-1} \quad (1)$$

#### B. Size-based Sampling

We employ size-based sampling in two places in our active learning pipeline. In the initial active learning iteration of the loop, most works [5], [6] use random sampling. Instead, we allocate the budget equally to all partitions (including the ground partition), and then randomly sample clusters from each partition. This achieves the desired warm start. In subsequent iterations, we rank clusters within each size-based




Fig. 2: **Adaptive binning over sizes.** We split all size groups into 3 partitions, as indicated by colors. The accumulated API from each partition is approximately the same. We balance object classes by creating a size-based partition.

partition using standard information measures (Section III-C) and sample the top-ranking clusters for annotation and further training. The budget remains approximately the same over partitions.

#### C. Information Measures

To rank clusters, we combine two well-established information measures: Softmax entropy [8] and feature diversity (CoreSet) [9]. Entropy aims to select the data that the model is most uncertain of. CoreSet prioritizes unlabeled data that is far from the labeled data in the feature space to ensure diversity. After ranking, we choose the top-ranking clusters for training in the next active learning iteration.

1) *Softmax Entropy*: We calculate the entropy for each unlabeled cluster  $j$  via Eq. (2), where  $N$  is the number of points in cluster  $j$  and  $p_{n,c}$  is the predicted probability of point  $n$  and class  $c$ .

$$E_j = -\frac{1}{N} \sum_{n=1}^N \sum_{c=1}^C p_{n,c} \log(p_{n,c}) \quad (2)$$




Fig. 3: **Experiment results of different active learning strategies on SemanticKITTI [1], nuScenes [2] using SPVCNN [40], Minkowski [41] network.** We compare BaSAL with other existing works. The solid line is the performance of the fully supervised model. The dashed line indicates 95% performance of the fully supervised model. Our model outperforms all existing active learning approaches on SemanticKITTI and gets on par performance with the state-of-the-art active learning method LiDAL [6] on nuScenes.

2) *Feature Diversity*: We use the feature from the last layer in the encoder-decoder network before classification and calculate the mean feature vector  $\mathbf{f}$  for each cluster. The feature diversity  $D_j$  is defined as the summed Euclidean distance between an unlabeled cluster  $j$  and all labeled clusters from the same partition, as in Eq. (3), where  $i$  is the index of a labeled cluster.

$$D_j = \sum_{i=1}^{N_l} \|\mathbf{f}_i - \mathbf{f}_j\|_2 \quad (3)$$

3) *Combination*: To create the final ranking of unlabeled clusters, we sort all unlabeled clusters by entropy and diversity per partition, as defined in Eq. (4), where  $r_{E_j}$  and  $r_{D_j}$  denote the ranks for the  $j$ -th cluster in terms of entropy and diversity.  $\mathcal{U}_b$  contains the indices of all unlabeled clusters in partition  $b$ , and  $R_b$  is the overall ranking of these clusters.

$$R_b = \operatorname{argsort}_{j \in \mathcal{U}_b} \left( \frac{1}{r_{E_j}} + \frac{1}{r_{D_j}} \right) \quad (4)$$

#### IV. EXPERIMENTS

We conduct extensive experiments on SemanticKITTI [42] and nuScenes [2] datasets, and compare with baselines [5], [6] on LiDAR semantic segmentation task.

##### A. Datasets and Evaluation Metric

On SemanticKITTI, we take sequences 00 – 07 and seq 09 – 10 for training and report the performance on validation sequence 08. On nuScenes, we train our model on the official training split which contains 700 scenes, and report the performance on the validation split with 150 scenes. We compare the mean Intersection over Union (mIoU) metric over 19 classes in SemanticKITTI and 16 classes in nuScenes.

##### B. Experimental Settings

1) *Network Architectures*: Following ReDAL [5] and LiDAL [6], we test our model with two backbone architectures: SPVCNN [40] and MinkowskiNet [41].

2) *Baselines*: We compare with random point selection (RAND), softmax confidence (CONF) [43], softmax margin (MAR) [43], softmax entropy (ENT) [43], MC-dropout (MCDR) [44], CoreSet selection (CSET) [9], segment-entropy (SEGMENT) [45], ReDAL [5] and LiDAL [6]. Experimental results of the baselines come from LiDAL [6].

3) *Active Learning Protocol*: BaSAL training has two stages: warm start initialization and subsequent active learning iterations. We warm start the model by size-based sampling. The annotation budget in this phase is  $x_{init}\%$ . After warm start, we conduct  $K$  active learning iterations. During each iteration, we interactively select top-ranking clusters from each partition for label acquisition such that the total annotation is no greater than  $x_{active}\%$ . Then we load the previous checkpoint and fine-tune the model with all labeled data.

The labeling budget is measured using the percentage of the labeled points with respect to the entire dataset. For both SemanticKITTI and nuScenes, we set  $x_{init} = 1$ ,  $K = 4$ , and  $x_{active} = 1$ . The annotation budget is uniformly distributed over all partitions (including the ground partition) in each iteration. We report the average over three runs.

##### C. Implementation Details

1) *Network Training*: All experiments are conducted with a single A40 GPU. For SemanticKITTI and nuScenes, we set the training batch size to 10/30 and validation batch size to 20/60. We first train our network using the warm start data for 100/200 epochs. Then we finetune the model for 30/150 epochs for each active learning iteration. For both datasets, we train the networks with the cross-entropy loss and the Adam optimizer with an initial learning rate of  $1e-3$ .

2) *Size-based Adaptive Binning*: Given a point cloud dataset, we first use PatchWork++ [46] with the default parameters to separate the ground points and the non-ground points. The ground points are then divided into grids of size  $10m \times 10m$ , named Partition Ground. For non-ground points, we use HDBSCAN clustering [39] algorithm to

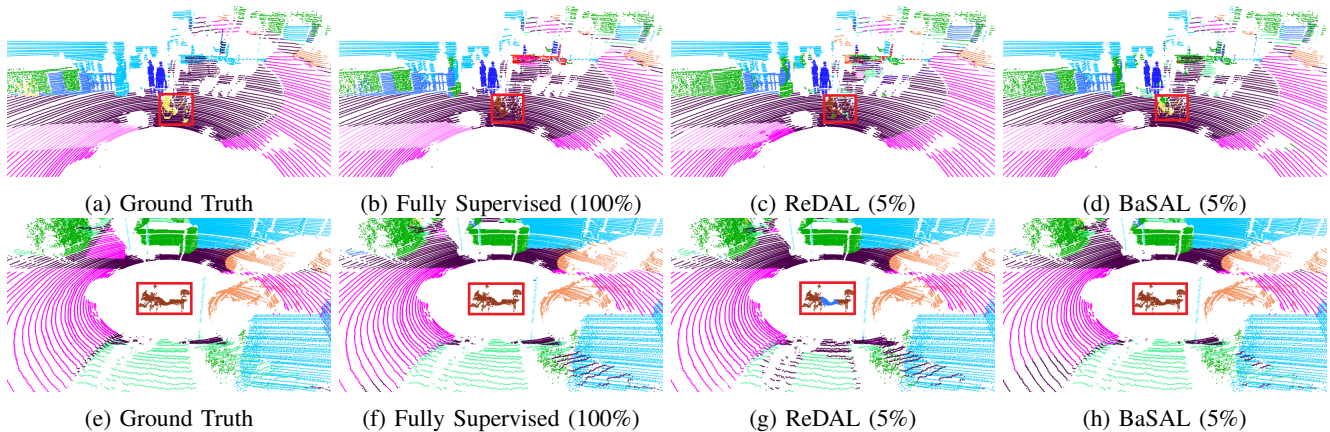


Fig. 4: **Qualitative comparison on SemanticKITTI [1] using the Minkowski [41] backbone.** We visualize semantic segmentation results on two examples. Our model successfully detects the bicycle in (d), as indicated by the red box. In comparison, other models misclassify the bicycle as sidewalks in (b) and (c). In the second example, our model better detects the motorcycle in (h), while ReDAL [5] over-segments it. We improve the performance on underrepresented classes.

split them into clusters. We set  $min\_cluster$ ,  $min\_samples$  and  $cluster\_selection\_epsilon$  to be 20, 10 and 0.5 for SemanticKITTI and 20, 1 and 0.5 for nuScenes. We then calculate the size of all the non-ground clusters. The size of a cluster is the rounded sum of the length, width, and height of its bounding box. We filter those clusters with the size larger than 25 meters. As mentioned in Section III-A, we calculate the Average Point Information (API) metric and adaptively group the sizes to 3 partitions.

3) *Size-based Sampling*: The size-based sampling method is used in both the warm start step and subsequent active learning iterations. For warm start (1% budget), we uniformly sample clusters from Partition 1, Partition 2, Partition 3, and Partition Ground until the accumulated number of points in each Partition reaches 0.25% of the entire dataset. For each subsequent active learning iteration, we increase the point budget by 1%, which is also uniformly allocated to all partitions (0.25% each).

#### D. Main Results

Fig. 3 compares BaSAL with the baselines. The x-axis represents the percentage of labeled points with respect to the entire training set, and the y-axis represents performance (mIoU). On SemanticKITTI, our model consistently outperforms all baselines in all settings. When the annotation budget increases to 5%, BaSAL (with Minkowski backbone) is able to match the performance of fully supervised training on the entire dataset, demonstrating the benefit of active learning in reducing annotation effort. BaSAL with the SPVCNN backbone achieves 98% the performance of the fully supervised model. On nuScenes, BaSAL surpasses all the baselines given 1%, 2%, and 3% labeled data. LiDAL [6] is able to match the performance of BaSAL when the annotation budget is up to 5%. Both models reach 95% performance of fully supervised learning. We also compare our warm start with the random sampling cold start at 1%. Notably, the advantage of our model over other baselines is more than

5% mIoU on both datasets, validating the advantage of our warm start in tackling the cold start problem.

#### E. Ablation Study

1) *Method Ablation*: We numerically evaluate the contribution of each component in our design on SemanticKITTI [1] in Tab. III. RS stands for random sampling from all unlabeled clusters. SS represents size-based sampling. EN and FD indicate the information measures: entropy and feature diversity, respectively. We first ablate our size-based partition. Without relying on any learning-based information measures, our model gains a 3.3% improvement, by simply replacing the random sampling with our size-based partition in the cold start setting. Notably, this is the most significant increase in our ablation study. Warm starting the model further results in a 0.3% improvement. Next, we evaluate the effect of information measures. Adding softmax entropy as an information measure improves the result by 0.7%. Taking feature diversity as an extra measure marginally improves the result by 0.2%.

2) *Number of Partitions*: We also experiment with different numbers of partitions when warm starting the model. We set the number of partitions  $B$  to 3, 6, 12, and 25, and achieve the best performance when  $B$  is set to 3 and 6. When  $B$  gets larger to 12 and 25, the performance drops. In general, we find the number of partitions has a minor impact on the performance when  $B$  is between 3 and 6. We set  $B$  to 3 in our implementation.

#### F. Analysis of Class Imbalance

1) *Analysis on Less Frequent Classes*: We show our results on less frequent classes in Tab. II. We outperform ReDAL and fully supervised learning by a substantial margin on four classes: bicycle, motorcycle, person and bicyclist. For example, the advantage over baselines is more than 20% on the bicycle class and approximately 10% on the bicyclist class. It is worth noting that these two classes only occupy

TABLE I: **Class distribution on SemanticKITTI [1].** We calculate the class distribution over the entire dataset and over partitions of different size. There are 4 partitions (including the ground partition). Classes are highly imbalanced over the dataset and most classes have a characteristic size (highlighted in bold text).

	car	bicycle	motorcycle	truck	other vehicle	person	bicyclist	motorcyclist	road	parking	sidewalk	other ground	building	fence	vegetation	trunk	terrain	pole	traffic sign
Entire dataset	5.7	0.01	0.05	0.3	0.3	0.04	0.02	0.005	27.3	2.0	19.7	0.5	10.2	3.9	18.6	0.8	10.3	0.3	0.07
Partition1 (small)	26.7	23.4	<b>77.3</b>	3.6	17.7	<b>53.9</b>	<b>68.2</b>	<b>76.3</b>	0.1	0.1	0.2	1.2	6.5	4.4	9.3	<b>48.7</b>	1.1	<b>42.9</b>	<b>50.5</b>
Partition2 (medium)	<b>63.9</b>	14.3	10.6	29.6	<b>56.5</b>	23.2	24.3	18.1	0.1	0.2	0.2	1.6	13.1	9.0	18.0	33.2	1.7	19.6	23.6
Partition3 (large)	6.5	<b>40.4</b>	6.9	<b>62.9</b>	23.7	17.8	5.6	0.8	0.1	0.1	0.4	4.4	<b>71.1</b>	<b>49.8</b>	<b>50.7</b>	14.7	3.4	31.6	25.7
Ground Partition	2.8	21.9	5.2	3.9	2.1	5.1	1.9	4.9	<b>99.7</b>	<b>99.6</b>	<b>99.3</b>	<b>92.7</b>	9.3	36.8	22.1	3.5	<b>93.9</b>	5.8	0.2

TABLE II: **Performance on less frequent classes.** We measure mIoU (%) on SemanticKITTI [1] using the Minkowski [41] backbone. We also calculate the proportion of points per class with respect to the full dataset. Our model boosts the performance on rare classes substantially, verifying its advantage in tackling class imbalance. FS stands for fully supervised learning on the full dataset.

	bicycle	motorcycle	person	bicyclist	motorcyclist
%	0.01	0.05	0.04	0.02	0.005
RAND	9.5	45.0	52.0	47.8	0.0
ReDAL	29.6	58.6	63.4	84.1	<b>0.5</b>
FS	20.4	63.9	65.0	78.5	0.4
Ours	<b>43.5</b>	<b>70.6</b>	<b>70.4</b>	<b>88.4</b>	0.2

TABLE III: **Ablation study.** We conduct our ablation on SemanticKITTI [1]. We train all models from scratch in iteration 1 and fine-tune them in iteration 2-5. Our size-based partition contributes the most to the final performance. RS: random sampling from all unlabeled data; SS: size-based sampling; EN: softmax entropy as the information measure; FD: feature diversity as the information measure [9].

Iteration 1	Iteration 2-5				mIoU(%)
	Info measure		Data Sampling		
Initialization	EN	FD	RS	SS	
Cold start			✓		56.8
				✓	60.1
Warm start	✓			✓	60.4
	✓	✓		✓	61.1
				✓	61.3

0.01% and 0.02% of the entire dataset. We attribute the remarkable improvement to our size-based partition which better handles class-balance. We show qualitative results on two samples in Fig. 4. In the top row, our model is able to recognize the bicycle, labeled by the red box in (d), while other models fail. In the second row, our model better segments the motorcycle than ReDAL [5], verifying our advantage in segmenting underrepresented classes.

2) *Analysis on Size-based Partition:* We also analyze the distribution of all classes on SemanticKITTI in Tab. I. We first enumerate the proportion of each class and then calculate the distribution of each class after size-based partition. All the numbers are vertically normalized over partitions per class. One key observation is that each object class has a size bias, which aligns with our assumption that size is an informative cue to distinguish classes. For example, the less frequent classes, *person*, *motorcyclist* and *motorcycle*, come mostly from *partition 1*. Compared to random sampling from the entire dataset, sampling from *partition 1* is likely

to increase the chance of obtaining an underrepresented object class, thus reducing the class imbalance considerably. However, we also notice that the majority of the *bicycle* class falls in *partition 3*. We suspect that this class may not be well segmented by the HDBSCAN [39] clustering algorithm.

## V. DISCUSSION

Active learning in large-scale autonomous driving datasets aims to reduce the heavy labeling effort. However, recent works quantify labeling effort differently. [47] counts the number of frames and bounding boxes. ReDAL [5], LiDAL [6], Just Label What You Need [24] count the number of point labels. Some other works measure the number of clicks of an annotator, such as OneThingOneClick [48] and LESS [49]. LESS achieves approximately 100% performance using only 0.1% of the annotations. A crucial difference between our method and LESS is that they require the annotator to inspect the entire dataset, whereas in our case the annotator only inspects 5% of the data.

## VI. CONCLUSION

We present BaSAL, a size-based warm start strategy for active learning on LiDAR semantic segmentation. BaSAL reduces the class imbalance in large-scale autonomous driving datasets by sampling from a size-based partition. Our intuition is that each object class has a characteristic size. Thus we are able to indirectly address the class imbalance and the cold start problem by grouping class-agnostic clusters according to size. We fuse the size-based partition into state-of-the-art active learning models and improve the performance considerably. Particularly, on SemanticKITTI we achieve the same performance as fully supervised learning on the entire dataset, while using only 5% of the annotation. Meanwhile, we boost the performance on less frequent classes significantly. Future work can explore better ways to create object clusters and develop novel cost functions to accurately quantify human annotation effort. On-device active learning is also an interesting direction without an excessive demand for computational resources.

## REFERENCES

- [1] J. Behley, M. Garbade, A. Milioto, J. Quenzel, S. Behnke, C. Stachniss, and J. Gall, "SemanticKITTI: A dataset for semantic scene understanding of lidar sequences," in *ICCV*, 2019, pp. 9297–9307. **1, 4, 5, 6**
- [2] H. Caesar, V. Bankiti, A. H. Lang, S. Vora, V. E. Liang, Q. Xu, A. Krishnan, Y. Pan, G. Baldan, and O. Beijbom, "nuScenes: A multimodal dataset for autonomous driving," in *CVPR*, 2020, pp. 11 621–11 631. **1, 4**

- [3] W. K. Fong, R. Mohan, J. V. Hurtado, L. Zhou, H. Caesar, O. Beijbom, and A. Valada, "Panoptic nusences: A large-scale benchmark for lidar panoptic segmentation and tracking," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 3795–3802, 2022. 1
- [4] B. Settles, "Active learning literature survey," *University of Wisconsin-Madison Department of Computer Sciences*, 2009. 1, 2
- [5] T.-H. Wu, Y.-C. Liu, Y.-K. Huang, H.-Y. Lee, H.-T. Su, P.-C. Huang, and W. H. Hsu, "Redal: Region-based and diversity-aware active learning for point cloud semantic segmentation," in *ICCV*, 2021, pp. 15 510–15 519. 1, 2, 3, 4, 5, 6
- [6] Z. Hu, X. Bai, R. Zhang, X. Wang, G. Sun, H. Fu, and C.-L. Tai, "LiDAL: Inter-frame uncertainty based active learning for 3d lidar semantic segmentation," in *ECCV*, 2022, pp. 248–265. 1, 2, 3, 4, 5, 6
- [7] V. Nath, D. Yang, H. R. Roth, and D. Xu, "Warm start active learning with proxy labels and selection via semi-supervised fine-tuning," in *MICCAI 2022*, 2022. 1, 2
- [8] A. Holub, P. Perona, and M. C. Burl, "Entropy-based active learning for object recognition," in *CVPR Workshop*. IEEE, 2008, pp. 1–8. 2, 3
- [9] O. Sener and S. Savarese, "Active learning for convolutional neural networks: A core-set approach," *arXiv preprint arXiv:1708.00489*, 2017. 2, 3, 4, 6
- [10] W. H. Beluch, T. Genewein, A. Nürnberger, and J. M. Köhler, "The power of ensembles for active learning in image classification," in *CVPR*, 2018, pp. 9368–9377. 2
- [11] B. Lakshminarayanan, A. Pritzel, and C. Blundell, "Simple and scalable predictive uncertainty estimation using deep ensembles," *NeurIPS*, vol. 30, 2017. 2
- [12] K. S. Tan, H. Caesar, and O. O. Beijbom, "Cross-modality active learning for object detection," U.S. Patent 11521010, Jan. 2024. 2
- [13] Y. Gal, R. Islam, and Z. Ghahramani, "Deep bayesian active learning with image data," in *International conference on machine learning*. PMLR, 2017, pp. 1183–1192. 2
- [14] D. Yoo and I. S. Kweon, "Learning loss for active learning," in *CVPR*, 2019, pp. 93–102. 2
- [15] S. Sinha, S. Ebrahimi, and T. Darrell, "Variational adversarial active learning," in *ICCV*, 2019, pp. 5972–5981. 2
- [16] S. Tong and D. Koller, "Support vector machine active learning with applications to text classification," *Journal of machine learning research*, vol. 2, no. Nov, pp. 45–66, 2001. 2
- [17] H. T. Nguyen and A. Smeulders, "Active learning using pre-clustering," in *Proceedings of the twenty-first international conference on Machine learning*, 2004, p. 79. 2
- [18] Y. Guo, "Active instance sampling via matrix partition," *NeurIPS*, vol. 23, 2010. 2
- [19] D. Gudovskiy, A. Hodgkinson, T. Yamaguchi, and S. Tsukizawa, "Deep active learning for biased datasets via fisher kernel self-supervision," in *CVPR*, 2020, pp. 9041–9049. 2
- [20] A. Kirsch, J. Van Amersfoort, and Y. Gal, "Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning," *NeurIPS*, vol. 32, 2019. 2
- [21] J. T. Ash, C. Zhang, A. Krishnamurthy, J. Langford, and A. Agarwal, "Deep batch active learning by diverse, uncertain gradient lower bounds," *arXiv preprint arXiv:1906.03671*, 2019. 2
- [22] A. Ghita, B. Antoniuussen, W. Zimmer, R. Greer, C. Creß, A. Møgel-mose, M. M. Trivedi, and A. C. Knoll, "ActiveAnno3D—an active learning framework for multi-modal 3d object detection," *arXiv preprint arXiv:2402.03235*, 2024. 2
- [23] H. Liang, C. Jiang, D. Feng, X. Chen, H. Xu, X. Liang, W. Zhang, Z. Li, and L. Van Gool, "Exploring geometry-aware contrast and clustering harmonization for self-supervised 3d object detection," in *ICCV*, 2021, pp. 3293–3302. 2
- [24] S. Segal, N. Kumar, S. Casas, W. Zeng, M. Ren, J. Wang, and R. Urtasun, "Just label what you need: fine-grained active selection for perception and prediction through partially labeled scenes," *arXiv preprint arXiv:2104.03956*, 2021. 2, 6
- [25] J. Attenberg and Ş. Ertekin, "Class imbalance and active learning," *Imbalanced Learning: Foundations, Algorithms, and Applications*, pp. 101–149, 2013. 2
- [26] K. Tomanek and U. Hahn, "Reducing class imbalance during active learning for named entity annotation," in *Proceedings of the fifth international conference on Knowledge capture*, 2009, p. 105–112. 2
- [27] B. Settles and M. Craven, "An analysis of active learning strategies for sequence labeling tasks," in *proceedings of the 2008 conference on empirical methods in natural language processing*, 2008, p. 1070–1079. 2
- [28] H. T. Nguyen and A. Smeulders, "Active learning using pre-clustering," in *Proceedings of the twenty-first international conference on Machine learning*, 2004, p. 79. 2
- [29] P. Donmez and J. G. Carbonell, "Paired-sampling in density-sensitive active learning," *Carnegie Mellon University*, 2008. 2
- [30] K. Tomanek and U. Hahn, "Reducing class imbalance during active learning for named entity annotation," in *Proceedings of the fifth international conference on Knowledge capture*, 2009, pp. 105–112. 2
- [31] S. Ertekin, J. Huang, L. Bottou, and L. Giles, "Learning on the border: active learning in imbalanced data classification," in *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, 2007, pp. 127–136. 2
- [32] M. Bloodgood and K. Vijay-Shanker, "Taking into account the differences between actively and passively acquired data: The case of active learning with support vector machines for imbalanced datasets," *arXiv preprint arXiv:1409.4835*, 2014.
- [33] U. Aggarwal, A. Popescu, and C. Hudelot, "Active learning for imbalanced datasets," in *WACV*, 2020, pp. 1428–1437. 2
- [34] M. Yuan, H.-T. Lin, and J. Boyd-Graber, "Cold-start active learning through self-supervised language modeling," *arXiv preprint arXiv:2010.09535*, 2020. 2
- [35] K. Pourahmadi, P. Nooralinejad, and H. Pirsiavash, "A simple baseline for low-budget active learning," *arXiv preprint arXiv:2110.12033*, 2021. 2
- [36] L. Chen, Y. Bai, S. Huang, Y. Lu, B. Wen, A. Yuille, and Z. Zhou, "Making your first choice: to address cold start problem in medical active learning," in *Medical Imaging with Deep Learning*. PMLR, 2024, pp. 496–525. 2
- [37] G. Hacoheh, A. Dekel, and D. Weinshall, "Active learning on a budget: Opposite strategies suit high and low budgets," *arXiv preprint arXiv:2202.02794*, 2022. 2
- [38] O. Yehuda, A. Dekel, G. Hacoheh, and D. Weinshall, "Active learning through a covering lens," *NeurIPS*, vol. 35, pp. 22 354–22 367, 2022. 2
- [39] L. McInnes, J. Healy, and S. Astels, "HDBScan: Hierarchical density based clustering," *The Journal of Open Source Software*, vol. 2, no. 11, p. 205, 2017. 3, 4, 6
- [40] H. Tang, Z. Liu, S. Zhao, Y. Lin, J. Lin, H. Wang, and S. Han, "Searching efficient 3d architectures with sparse point-voxel convolution," in *European conference on computer vision*, 2020, pp. 685–702. 4
- [41] C. Choy, J. Gwak, and S. Savarese, "4d spatio-temporal convnets: Minkowski convolutional neural networks," in *CVPR*, 2019, pp. 3075–3084. 4, 5, 6
- [42] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for autonomous driving? the kitti vision benchmark suite," in *2012 IEEE conference on computer vision and pattern recognition*. IEEE, 2012, pp. 3354–3361. 4
- [43] D. Wang and Y. Shang, "A new active labeling method for deep learning," in *2014 International joint conference on neural networks (IJCNN)*. IEEE, 2014, pp. 112–119. 4
- [44] Y. Gal and Z. Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059. 4
- [45] Y. Lin, G. Vosselman, Y. Cao, and M. Yang, "Efficient training of semantic point cloud segmentation via active learning," *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, vol. 2, pp. 243–250, 2020. 4
- [46] S. Lee, H. Lim, and H. Myung, "Patchwork++: Fast and robust ground segmentation solving partial under-segmentation using 3D point cloud," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst.*, 2022, pp. 13 276–13 283. 4
- [47] Z. Liang, X. Xu, S. Deng, L. Cai, T. Jiang, and K. Jia, "Exploring diversity-based active learning for 3d object detection in autonomous driving," *arXiv preprint arXiv:2205.07708*, 2022. 6
- [48] Z. Liu, X. Qi, and C.-W. Fu, "One thing one click: A self-training approach for weakly supervised 3d semantic segmentation," in *CVPR*, 2021, pp. 1726–1736. 6
- [49] M. Liu, Y. Zhou, C. R. Qi, B. Gong, H. Su, and D. Anguelov, "Less: Label-efficient semantic segmentation for lidar point clouds," in *ECCV*, 2022, pp. 70–89. 6