



HAL
open science

You Never Get a Second Chance To Make a Good First Impression: Seeding Active Learning for 3D Semantic Segmentation

Nermin Samet, Oriane Siméoni, Gilles Puy, Georgy Ponimatkin, Renaud Marlet, Vincent Lepetit

► **To cite this version:**

Nermin Samet, Oriane Siméoni, Gilles Puy, Georgy Ponimatkin, Renaud Marlet, et al.. You Never Get a Second Chance To Make a Good First Impression: Seeding Active Learning for 3D Semantic Segmentation. 2023 IEEE/CVF International Conference on Computer Vision (ICCV), Oct 2023, Paris, France. 10.1109/ICCV51070.2023.01691 . hal-04496540

HAL Id: hal-04496540

<https://hal.science/hal-04496540>

Submitted on 8 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

You Never Get a Second Chance To Make a Good First Impression: Seeding Active Learning for 3D Semantic Segmentation

Nermin Samet¹, Oriane Siméoni², Gilles Puy², Georgy Ponimatkin¹, Renaud Marlet^{1,2}, Vincent Lepetit¹
¹LIGM, Ecole des Ponts, Univ Gustave Eiffel, CNRS, Marne-la-Vallée, France ²Valeo.ai, Paris, France

Abstract

We propose *SeedAL*, a method to seed active learning for efficient annotation of 3D point clouds for semantic segmentation. Active Learning (AL) iteratively selects relevant data fractions to annotate within a given budget, but requires a first fraction of the dataset (a ‘seed’) to be already annotated to estimate the benefit of annotating other data fractions. We first show that the choice of the seed can significantly affect the performance of many AL methods. We then propose a method for automatically constructing a seed that will ensure good performance for AL. Assuming that images of the point clouds are available, which is common, our method relies on powerful unsupervised image features to measure the diversity of the point clouds. It selects the point clouds for the seed by optimizing the diversity under an annotation budget, which can be done by solving a linear optimization problem. Our experiments demonstrate the effectiveness of our approach compared to random seeding and existing methods on both the S3DIS and SemanticKitti datasets. Code is available at <https://github.com/nerminsamet/seedal>.

1. Introduction

We are interested in the efficient annotation of sparse 3D point clouds (as captured indoors by depth cameras or outdoors by automotive lidars) for semantic segmentation.

Modern AI systems require training on large annotated datasets to reach a high performance on such a task. As annotating is costly (more than capturing the data itself), several approaches have been proposed to achieve a more frugal learning, such as semi-supervision [71], weak supervision [88], few-shot [52] and zero-shot learning [56], self-supervision [37, 1] and, as studied here, active learning [57].

Active learning (AL) methods iteratively select relevant fractions of a dataset to be annotated within a given budget so that, after a few iterations, the model learned on the annotated fractions reaches a performance close to the performance of the model learned on the fully-annotated dataset,

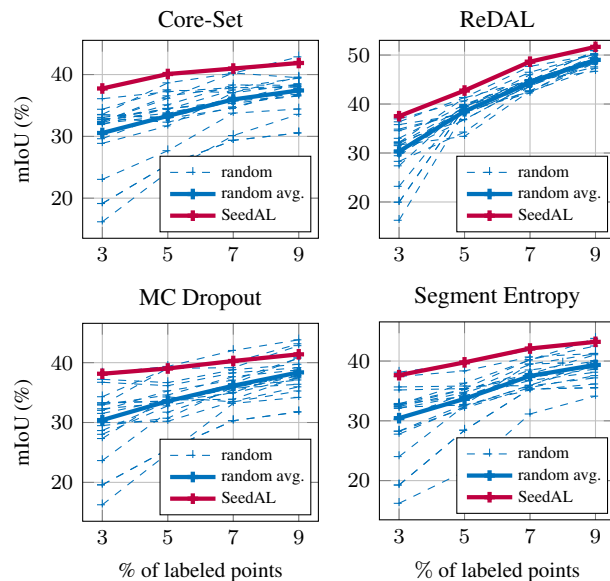


Figure 1: *Impact of active learning seed on performance.* We show the variability of results obtained with 20 different random seeds (blue dashed lines), within an initial annotation budget of 3% of the dataset, when using various active learning methods for 3D semantic segmentation of S3DIS. We compare it to the result obtained with our seed selection strategy (solid red line), named SeedAL, which performs better or on par with the best (lucky) random seeds among 20, and “protects” from very bad (unlucky) random seeds.

although at a much lower annotation cost. This AL selection typically targets the most uncertain [38] or most diverse [62] data, that are assumed to have the most positive impact when annotated and used for training.

The criteria used in AL methods for selecting data to annotate generally require a preliminary fraction of the dataset to be already annotated. Assuming this fraction is representative enough of the rest of the dataset, it can be used to estimate the benefit of annotating other data fractions. But such a preliminary fraction is not available at the first AL iteration, when confronted to a completely unannotated dataset.

There is thus a cold start problem, which is to determine a first fraction of the dataset to annotate — the seed.

Initial set. Nevertheless, most AL publications pay little attention to the choice of this seed before iterating their AL method. They just pick a random fraction of the dataset. Although results are sometimes presented averaged over a few runs, it can be insufficient given the high variance level. As shown in [Figure 1](#), it is particularly true in our 3D semantic segmentation context. Besides, nothing prevents in practice from drawing an “unlucky” seed, leading to significant underperformance or annotation overheads.

To address this issue, our method, named SeedAL, automatically selects extremely good AL seeds, and thus also protects from drawing unlucky seeds, as illustrated on [Figure 1](#). Only few approaches have recently been proposed to construct such good seeds to get AL on track right away [[84](#), [54](#), [40](#), [29](#), [48](#), [81](#), [13](#), [50](#)], and they mostly regard text or image classification. To the best of our knowledge, no such AL seeding approach exists for 3D point clouds.

Self-supervised features, for text or images [[20](#), [31](#), [12](#)], are the key to most of these approaches. It thus seems natural to use self-supervised 3D features in our case. But existing pre-trained 3D backbones are not as versatile as their 2D counterpart, and provide lower quality features (see [Sect. 3.3](#)). Our approach to seed AL for 3D point cloud is to leverage high-quality pretrained features [[12](#)] for images that are views of the scanned scenes.

As we rely on image features like [[54](#), [40](#), [29](#), [81](#), [13](#)], it seems natural as well to reuse these AL seeding methods for 3D scenes. But a direct transposition to 3D is not sensible because, in our case, the initial budget for the seed is not given in terms of number of images but in the size of scenes (number of 3D points): selecting a scene just because it contains a single image of interest would be suboptimal.

Our approach integrates scenes and views to make a better use of the initial annotation budget. While 2D approaches use feature clustering to guide seed selection, we show it is suboptimal in our context (2D views of 3D scenes) and propose a better formulation based on binary linear programming. As our algorithm has a higher complexity than clustering, we also develop heuristics to scale to large datasets by first extracting a small-enough pool of good candidates to select from.

Last, we observe that usual pretrained features, generally created from object-centric datasets such as ImageNet, fail to convey the diversity of complex multi-object scenes, by giving too much importance to big or repeated objects. We address this issue by analyzing images at patch level.

Contributions. To the best of our knowledge, we are the first to address AL seeding for 3D point clouds:

- We study the sensitivity of AL seeding for the 3D semantic segmentation of point clouds and the relevance of 2D images features to select 3D scenes.

- Leveraging these studies, we propose a general approach to seed any AL method for point clouds originating from scenes with available image views.
- We present experimental evidence of the effectiveness of our AL seeding method, which consistently improves over random seeds and 2D-inspired baselines, requiring less iterations (thus less annotations) and/or reaching higher segmentation performance.

2. Related Work

We discuss in this section different approaches to reduce the annotation work on point clouds, then focus on Active Learning methods, and finally existing works on seed construction for active learning—even though none of them consider point cloud annotations.

Towards less supervision for 3D point clouds. The best approaches for 3D semantic segmentation of dense indoor point clouds [[55](#), [17](#), [87](#), [70](#)] or sparse outdoor point clouds [[79](#), [32](#), [89](#), [68](#)] are trained under full supervision, i.e., they require manual annotations of all points in all scenes, which is a notoriously costly task [[5](#), [26](#)]. Several approaches are currently explored to mitigate these costs. For example, one can leverage self-supervision [[76](#), [86](#), [35](#), [42](#), [51](#), [61](#)] to pre-train a neural network using many unlabeled data with an annotation-free pretext task and then fine-tune this network using few annotations for the task of interest. One can also rely on domain adaptation techniques [[83](#), [59](#), [60](#), [77](#), [80](#)] to exploit existing annotated datasets on a source domain and avoid annotating a new dataset for a target domain. Weak supervision avoids the burden of complete point cloud annotation by, e.g., annotating a few points [[45](#), [78](#), [16](#), [85](#)], annotating few regions in a scene [[44](#), [74](#)], or using scene level labels [[58](#), [74](#)]. Alternatively, active learning techniques, which we describe next, allow for a smart selection of the data to be annotated.

Active Learning for 3D point clouds. To minimize annotation costs, active learning strategies [[9](#), [25](#), [62](#)] aim to select the most relevant data to annotate for the considered model and task. Most selection strategies rely on diversity criteria [[25](#), [62](#), [3](#)] or on a measure of the uncertainty of the model [[9](#), [72](#)]. In computer vision, active learning has been mostly developed for classification tasks of 2D images [[63](#)]. It has been adapted to the semantic segmentation of 3D point clouds combined with markov random fields [[46](#)] and more recently deep learning for object detection [[21](#), [47](#)] and semantic segmentation [[43](#), [75](#), [34](#)]. Diversity of selected scenes has been ensured by selecting core sets [[62](#), [75](#)] in the feature space—in that space each scene is represented by aggregating the features of its corresponding points. The model uncertainty can be measured over each point class-scores [[72](#), [43](#)]*—e.g.*, using the entropy or the margin between the two highest scores—

and averaged per scene in order to select the most confusing to the model. In [43], the authors improve results by computing the uncertainty scores at the level of pre-computed segments. Requiring several forward passes, ensembling techniques [6, 22, 23] have been used to evaluate the model’s confidence for 3D scenes [21, 75]. Moving away from scene-based AL, [65, 75, 34] propose region-based strategies—which mix difficulty and diversity—that allow them to further reduce cost by selecting only local set of points to be annotated. In this work, instead of proposing a new AL strategy, we show that a smart selection of the first set of data with SeedAL—usually over-looked and randomly selected—can boost all methods drastically.

Cold start problem. The need for an initially annotated fraction of the data to bootstrap an AL method is similar to what has been identified as a cold start problem in collaborative filtering and recommendation systems, where new users start off with an empty profile [49].

In active learning, the term “cold start problem” has in fact been used with two close but different meanings. On the one hand [33, 39, 82], it refers to the strong bias induced on the first AL iterations by a too small initial annotated set [19]. In particular, uncertainty-based methods are not effective when trained with too little data [30]. Note that, in this setting, a seed has nevertheless to be given as input, although if it is possible to set its size automatically to limit this bias, by deciding when to stop annotating the first data fraction constituting the seed [24]. On the other hand [84, 13], the cold start problem also refers to the lack of a priori information to select a relevant seed in the first place, which is what we are addressing here.

Seeding active learning. Despite the high variance of performance with random seed selection, it appears it is difficult for AL approaches to create good seeds [13]. In fact, only few recent approaches propose to automatically construct AL seeds [84, 54, 40, 29, 48, 81, 13, 50]. To the best of our knowledge, none applies to 3D point clouds, and there is only one concerning semantic segmentation [50].

All approaches, except [50], only address text or image classification, which is a much coarser-grained task than semantic segmentation. The general idea of most approaches for image classification [54, 40, 29, 81, 13] is to use a self-supervised pretrained network, clustering the image features to estimate diversity groups and picking images close to cluster centers to get representatives, possibly focusing on hard-to-contrast and diverse data [13] by exploiting contrastive self-supervised features [14]. Alternatively, [48] creates seeds using a core-set approach based on the Wasserstein distance between feature distributions of candidate data for the seed and the unannotated dataset. While this form of feature clustering or distance minimization between distributions also makes sense for 3D data, seman-

tic segmentation requires extracting much more detailed information from each single datum to account for complex scenes (whether indoors or outdoors) that capture sets of objects, as opposed to object-centered pictures targeted by image classification in other approaches.

The only AL seeding approach we know for 3D, and in fact for semantic segmentation too, applies to dense 3D medical images (CT-scans) [50], which significantly differ from sparse 3D point clouds captured by depth cameras or lidar scanners. The seeding strategy in [50] is based on a heuristic pseudo-labeling defined by a hand-parameterized, rule-based segmentation. It consists in thresholding CT data within a typical window of values for abdominal soft tissues, then extracting largest connected components and selecting major organs of the abdomen as foreground. This pseudo-labeling is specific to that peculiar kind of medical images and cannot be directly transposed to point clouds, although heuristic hand-designed pre-segments have also been used on 3D point clouds for label-efficient semantic segmentation [44].

3. Preliminary Study

In this section, after introducing our notations and formally describing the problem, we highlight the sensitivity of various AL methods to the seed and then motivate the use of DINO [12] features for the selection of the seed. We use the results of this preliminary study to construct our method, SeedAL, which is fully described in the next section.

3.1. Problem setup and notations

Our method works in the following generic setting. We assume that a scene i is captured by a depth sensor, providing a point cloud P^i , and by one or multiple cameras, providing several views $V_1^i, \dots, V_{N_i}^i$ of the scene. For example, RGB-D cameras capture conveniently both the depth and the views simultaneously. Such a multi-modal setting is common, whether indoors [2, 18] or outdoors [5, 10].

Our dataset is made of a set of point clouds, with their respective views, and our goal is to select a good seed S , i.e., the initial set of the point clouds which will be fully annotated for the first active learning cycle. As mentioned earlier, our approach for selecting S is based on a pre-trained self-supervised image backbone. We denote this backbone $\phi(\cdot)$ and use it to extract D -dimensional ℓ_2 -normalized features from each view: $\phi(V_n^s) \in \mathbb{R}^D$.

3.2. The high variance induced by random seeds

Most AL publications use as AL seeds random fraction of the dataset — which is one of the many issues for reproducibility [40, 36]. A seed is typically picked within a maximum size, given as an initial annotation budget. For AL method evaluation, because of a high variance level due

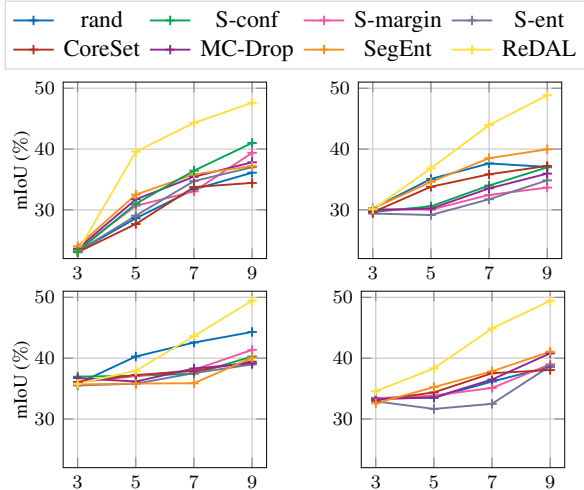


Figure 2: *Influence of the choice of the seed.* We visualize results obtained with several AL methods (introduced later in section 5) when using 4 different random seed. With the exception of ReDAL, which is consistently better, the ranking of the methods varies with the seed.

to the random AL seed, a common practice is to average results over few runs, randomizing the AL seed and the training iterations [24, 41, 11, 40, 54, 66, 36]. Even so, it can be not enough given the variance level. As shown in Figure 1, it is particularly true in our problem. Besides, nothing prevents from drawing an “unlucky” seed, leading to significant underperformance or annotation overheads.

We first consider the common case of random seeds, i.e., picking random scenes until the annotation budget is reached. In Figure 1, we illustrate the effect of drawing 20 such random seeds: it shows significant variations of performance, both at initial stage and after some iterations. Furthermore, as shown in Figure 2, the performance hierarchy of scene-based AL methods differs depending on the seed.

A natural way to deal with both intra-method and inter-method variations is to consider the average and variance of the method performance over a large-enough number of random seeds. It allows evaluating the sensitivity of an AL method to the seed and it allows a better comparison of different AL methods, although still with a high variance. However, it does not provide any practical hint how to choose and cold-start these methods. Besides, it is sub-optimal as it does not make the most of best seeds.

3.3. Image features to characterize 3D scenes

At first glance, it seems more natural and appropriate to directly use 3D features, as opposed to rely on extra 2D images. However, 3D backbones are not as versatile as 2D ones because of the huge domain gaps in 3D data, compared to picture data. 3D data require specific architectures and tuning due to the extreme variations of 3D

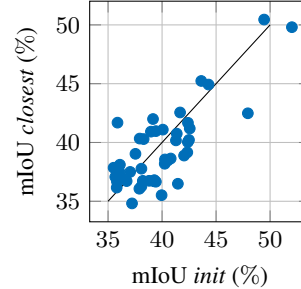


Figure 3: *Correlation between results of AL methods when using a random set and the closest in the DINO feature space [12].* Each point corresponds to a cycle of an AL method. Pearson correlation coefficient is 0.78. Results are produced on the S3DIS dataset [2].

point distribution, depending on the type of scenes (indoors vs outdoors, static vs dynamic) and on the kind of sensor used (photogrammetry, variants of depth cameras, lidars with widely differing acquisition patterns and resolutions). Besides, current self-supervised 3D features have a lower quality than their 2D counterpart in the sense that self-supervised 3D features can be improved by distilling, in 3D, high-quality 2D image features of views of the same scenes [61]. In fact, the quality of self-trained 2D image features [28, 27, 12, 4], is now such that, even if self-trained on different datasets, it can be used for object discovery in new data [67, 73, 53]. The best results have actually been achieved with DINO [12]; [81] also exploit DINO in their AL method. There is no such result with pretrained 3D features yet. Last, 2D backbones are pretrained efficiently and on very large datasets, which boosts quality and favors generalizability, compared to the computation burden and practical size of pretraining data in 3D.

Here, rather than relying on a distillation from 2D to 3D as in [61], it is simpler and more efficient to directly use 2D image features to select corresponding 3D fractions of the dataset. Moreover, it makes our approach agnostic to the 3D backbone, which avoids having to handle the domain gaps of 3D data. Last, we can leverage very large image datasets, such as ImageNet, as it has been shown to transfer well to tasks on other image datasets [12]; pretraining specifically on the images of the target dataset is not needed, and could even be detrimental if the dataset is not large enough.

3.4. Relevance of DINO features

To evaluate the effectiveness of DINO features, we first select a subset of the random seeds generated previously. Next, we create the most similar seed (in terms of DINO features) to each selected random seed, and run active learning methods with all of the seeds. Finally, we calculate the correlation between the random seeds and their newly generated most similar correspondences. More specifically, for

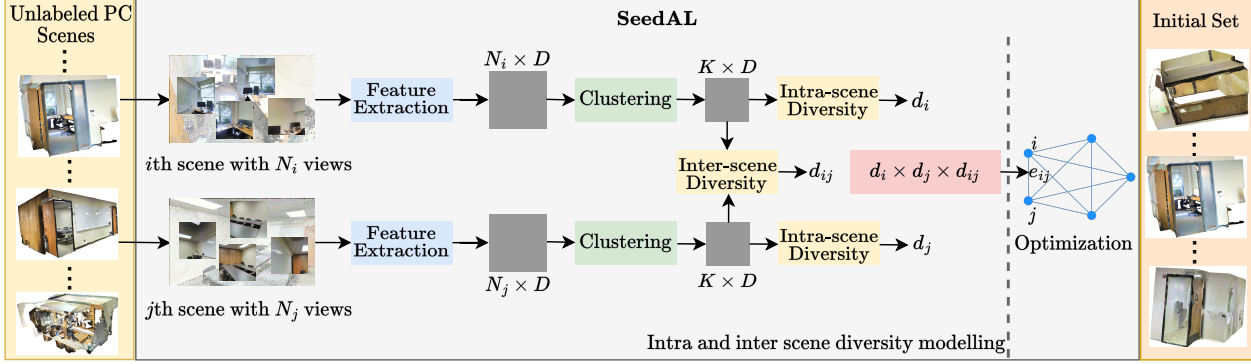


Figure 4: Overall processing pipeline of SeedAL.

each selected seed, denoted as S , we replace all scenes in each seed by identifying alternative scenes with the minimum distance in DINO space, and create a new seed denoted as \hat{S} . Formally, for each scene i , we compute the average feature $\Phi^i = \sum_{n=1}^{N_i} \phi(V_{N_i}^i) / N_i$ of all views. Then, we create \hat{S} by selecting, for each scene $i \in S$, the scene $\arg \min_{j \notin S} \|\Phi^j - \Phi^i\|_2$. Then, we apply several active learning (AL) methods to the newly created initial sets, and compare the results obtained from S and \hat{S} . Our analysis reveals a strong positive correlation between the two sets of results with a Pearson correlation coefficient of 0.78. The results indicate the applicability of leveraging DINO features for selecting the initial seed.

4. SeedAL: Method

Our approach to seed selection posits that a good set should have two key characteristics: Firstly, each scene within the set should exhibit a sufficient level of internal diversity. Secondly, there should be a notable degree of diversity between the scenes themselves. To identify the scenes that should be in the seed, we represent the unlabeled dataset as a fully connected graph G , in which each node represents a scene. The edge weights, denoted by e_{ij} , between two scenes i and j , encapsulate both intra and inter-scene diversity measures. These measures are formally introduced below. By looking for a sub-graph that optimizes the sum of edge values while satisfying a budget constraint, we can find a seed with good diversity. We propose a linear optimization framework for the selection of this sub-graph and explain its details in the following paragraphs. We illustrate our method in Figure 4.

4.1. Diversity measures

Intra-scene diversity measure. Given that a scene may contain near-duplicate views, e.g., captured from very similar viewpoints, it is important to eliminate redundancy to obtain representative features that accurately capture the characteristics of the scene. To this end, for each scene i ,

we cluster the set of view-features $\phi(V_1^i), \dots, \phi(V_{N_i}^i)$ in K clusters using k-means. The resulting set of representative features, denoted by $\Phi_1^i, \dots, \Phi_K^i$, contains K distinctive features of the scene. To calculate the diversity within a scene i , we calculate the pairwise dissimilarities between the cluster centers $\Phi_1^i, \dots, \Phi_K^i$: $(1 - \Phi_k^i \cdot \Phi_{k'}^i)$ for all $k > k' \in \{1, \dots, K\}$, where “ \cdot ” denotes the scalar product in \mathbb{R}^D . Finally, we average all these pairwise dissimilarities to obtain the internal diversity d_i of scene i :

$$d_i = \frac{2}{K(K-1)} \sum_{k=1}^K \sum_{k'=k+1}^K (1 - \Phi_k^i \cdot \Phi_{k'}^i). \quad (1)$$

Inter-scene diversity measure. Given two different scenes i and j , we compute their inter-scene diversity d_{ij} by relying on their respective cluster centers: $\Phi_1^i, \dots, \Phi_K^i$ and $\Phi_1^j, \dots, \Phi_K^j$. Specifically, we average the pairwise dissimilarities between all pairs of cluster centers:

$$d_{ij} = \frac{1}{K^2} \sum_{k=1}^K \sum_{k'=1}^K (1 - \Phi_k^i \cdot \Phi_{k'}^j). \quad (2)$$

Combined intra and inter-scene diversity measure. To combine the intra and inter-scene diversities, we simply compute $e_{ij} = d_{ij} d_i d_j$. Hence, scenes i and j in the graph G are highly connected if they are both highly diverse internally *and* mutually diverse.

4.2. Linear optimization for seed selection

We cast our seed selection as the identification of a fully connected subgraph of G such that the sum of edge weights is maximized under the constraint of a budget b for the initial annotations. Here, we define this budget in terms of the number of points to be annotated. The connected subgraph of G could be identified by solving the following problem:

$$\arg \max_{\{x_i\}_i} \sum_{i=1}^M \sum_{j=i+1}^M e_{ij} x_i x_j \quad \text{s. t.} \quad \sum_{i=1}^M c_i x_i \leq b, \quad (3)$$

where the x_i are boolean variables, indicating whether scene i is selected or not. c_i the annotation cost for scene i , which in practice we take as the number of points in scene i . The objective term $\sum_{i=1}^M \sum_{j=i+1}^M e_{ij} x_i x_j$, where M is the number of scenes, permits us to select scenes with the highest diversity possible while the constraint $\sum_{i=1}^M c_i x_i \leq b$ ensures that we do not exceed our annotation budget b .

However, this is a quadratic problem and M is large in practice. We can turn this problem into the linear problem

$$\arg \max_{\substack{\{y_{ij}\}_{i,j>i} \\ \{x_i\}_i}} \sum_{i=1}^M \sum_{j=i+1}^M e_{ij} y_{ij} \text{ s. t. } \begin{cases} \sum_{i=1}^M c_i x_i \leq b \\ y_{ij} \leq x_i \\ y_{ij} \leq x_j \end{cases} \quad (4)$$

by introducing the boolean variables y_{ij} . Because it is an optimization problem and thanks to the additional constraints, $y_{ij} = 1 \Leftrightarrow x_i = x_j = 1$. The advantage of this formulation is that it is much more efficient to solve when the number of scenes M is large.

In practice, we observed that it can still be computationally challenging for very large M . To address this issue, we first extract the top L edges in G with the highest weights e_{ij} and collect all scenes on both ends on these edges. Then, we rebuild a smaller graph using only these scenes and apply the optimization to this reduced graph. We use the solver from SCIP [7] to find the optimal solution.

5. Experiments

We present in this section the experimental protocol that we use to evaluate our method SeedAL and discuss results. We also provide an ablation study.

5.1. Experimental setup

Active learning strategies. We test our method with diverse active learning strategies that exploit either the model uncertainty or enforce diversity, that are either scene-based or region-based. In particular, we compare to a random selection of scenes (rand), the diversity-based method CoreSet [62] (CoreSet) and the uncertainty-based methods Softmax Confidence [72] (S-conf), Softmax Margin [72] (S-margin), Softmax Entropy [72] (S-ent), MC-Dropout [23] (MC-Drop) and Segment Entropy [43] (SegEnt). We also consider the region-based method ReDAL [75] (ReDAL).

Seeding baselines. We compare SeedAL to two clustering-based baselines inspired from [84]. For this purpose, for each scene i , we first compute the average feature $\Phi^i = \sum_{n=1}^{N_i} \phi(V_{N_i}^i) / N_i$ of all views. The first simple baseline, KMcentroid, starts with a clustering of the features Φ^1, \dots, Φ^M into K clusters using k-means. Then, for each cluster, we search the scene i whose feature Φ^i is the closest to the cluster centroid, we select these K scenes, and repeat

this search process until we reach the annotation budget. The most costly samples are then removed iteratively until the annotation budget is satisfied. We also consider a variant of KMcentroid, called KMfurthest, inspired by the core-set [62] selection. It starts as in KMcentroid with a selection of the K scenes closest to each cluster centroid. Then, we continue spending the annotation budget by selecting, iteratively, the scene j whose feature Φ^j is the farthest away from all cluster centroids, to favor a selection of diverse scenes. We also compare to the natural baseline *random*.

Datasets. Following [75], we focus on the semantic segmentation task and evaluate SeedAL on two datasets representative of indoor and outdoor scenes, namely S3DIS [2] and SemanticKITTI [5]. S3DIS is composed of 271 scenes extracted from 6 major indoor areas; for each scene is provided a dense point cloud with color. As common practice, we provide AL results on the Area 5 validation set, and use the rest for training. We also evaluate on SemanticKITTI, a large-scale autonomous driving dataset composed of 22 driving sequences, totaling 43,552 point clouds, each accompanied with images. Following the official protocol, we evaluate on the validation split (seq 08) and train the models on the entire official training split (seq 00-07 and 09-10).

Network architecture and evaluation metric. Following previous works [75], we consider SPVCNN [69], that is based on point-voxel CNN and achieves good results both on indoor and outdoor scenes. We evaluate the semantic segmentation task in all experiments with the mIoU metric.

Technical details. In all our experiments, we set K as the number of the classes for each dataset: 13 for S3DIS, 19 for SemanticKITTI., L is set as 100 for S3DIS and 1000 for SemanticKITTI. On S3DIS, the RGB images depicting each scene constitute our views. We use the class token extracted at the last layer of the DINO-pretrained ViT-B/8 as our view features. On SemanticKITTI, each scene is depicted by a single image frame. To account for multi-object images in outdoor scenes, we consider that each 8×8 patches of a frame constitute one view and we take the corresponding patch feature at the last layer of the ViT-B/8 as the view feature. Also, as there is a significant number of redundant and highly similar scenes in SemanticKITTI, we use a greedy algorithm to sparsify the sequences. The details of this sparsification are in the supplementary material.

5.2. Indoor evaluation

We evaluate here our method on S3DIS. We present in Figure 5 the results obtained with our AL seeding strategy, SeedAL, and with the considered baselines. First, we observe that using SeedAL boosts significantly all 8 AL methods evaluated over randomly selected seeds (dashed blue curves). SeedAL is consistently at the top of the graphs while random seeds yields, overall, unstable results.

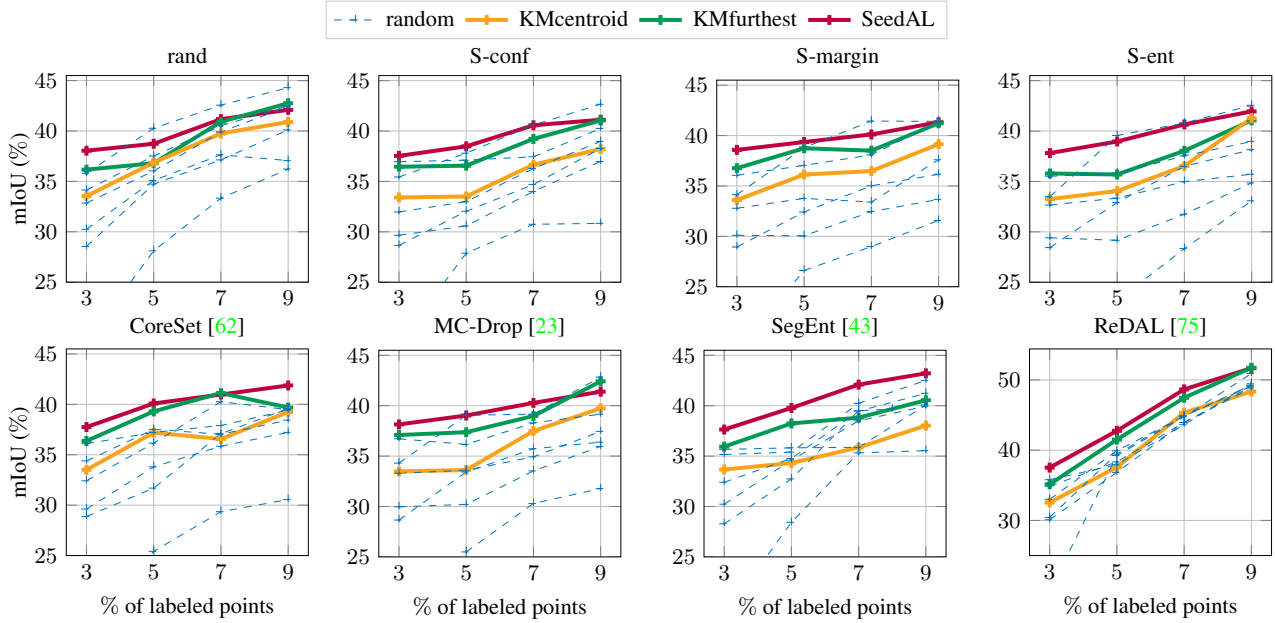


Figure 5: AL results on S3DIS. Comparison of active learning results when using SeedAL vs several baselines.

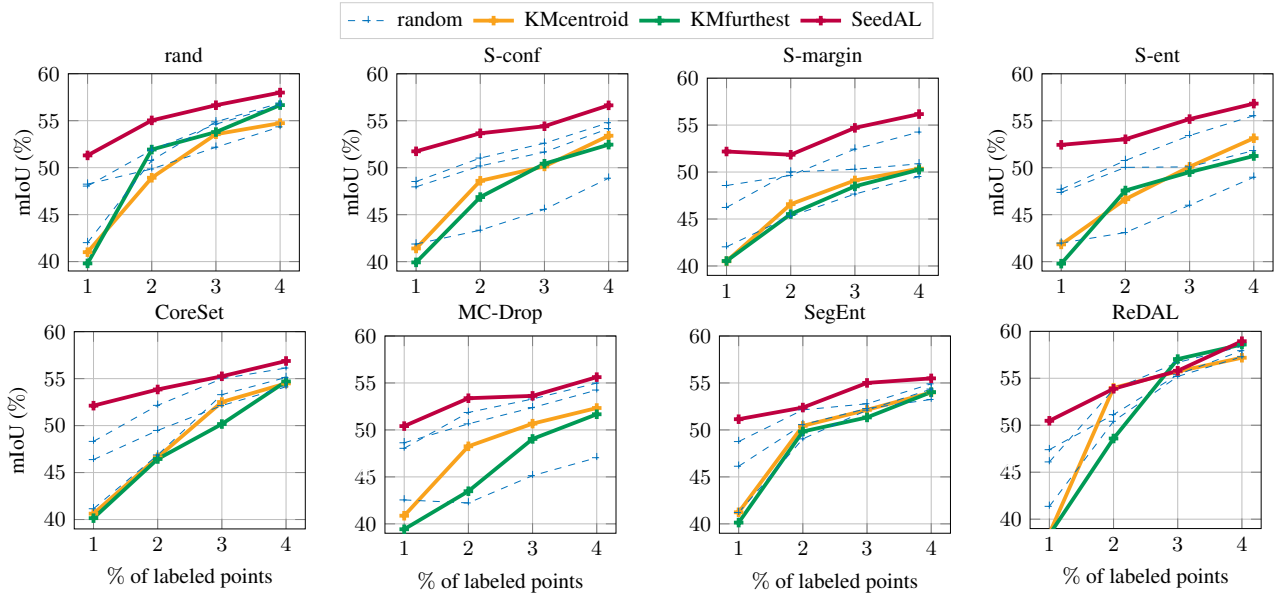


Figure 6: AL results on SemanticKITTI. Comparison of active learning results when using SeedAL vs several baselines.

Regarding our baselines, KMfurthest achieves consistently better performances than KMcentroid over all AL strategies. We explain this result by the fact that KMfurthest contains more diverse scenes than KMcentroid. We observe that for certain methods, KMfurthest achieves similar levels of performance as SeedAL. However, it is important to note that SeedAL is more robust and it consistently obtains strong performance in both early and late active learning iterations.

Making a good first impression. We highlight here the

boost obtained with SeedAL at the *very first cycle* (when no AL has been applied yet). We observe that with SeedAL, results are +1.9 pt over the best random seed, +7.8 pt v.s. the average of random seeds and +18.8 pt when comparing to the most unlucky seed. We surpass also our two baselines KMcentroid and KMfurthest by 4.9 pt and 2.4 pt, respectively. With only 3% of data well-selected thanks to SeedAL, we obtain better results than with 4 out of 5 random seeds with 5%, or even 7% (when employing S-conf, S-margin, S-ent, CoreSet and MC-Drop), of the data. These

results prove that our seeding selection is a much better strategy than betting on the quality of the random seed.

We also compare our method with SSDR-AL [64]. Differently than other AL methods, SSDR-AL’s initialization picks superpoints to label, rather than full scenes. As our current formulation selects scenes rather than regions, it cannot directly be applied in this case. The only way we can compare quantitatively with [64] is on S3DIS regarding the proportion of points to label in order to reach the same level as 90% of full supervision: we boost ReDAL to outperform SSDR-AL, only requiring 9% labeled points instead of 11.7%, as shown in Table 1.

Method (init.)	SSDR (rand)	ReDAL (rand)	ReDAL (ours-DINO)	ReDAL (ours-MoCo)
% labeled pts (↓)	11.7 %	13%	11%	9%

Table 1: Comparing SSDR’s initialization with SeedAL in terms of % of pts to label to get to 90% of full supervision.

5.3. Outdoor evaluation

We now focus on the challenging SemanticKitti dataset, which depicts highly redundant urban images. We present results in Figure 6. We also show results with the baseline rand, which is known to surpass all AL methods except the region-based ReDAL.

Boosting active learning. First, we observe that SeedAL significantly surpasses all seeding baselines, either random or diversity-based. Second, it achieves better results compared to any randomly sampled seed with all AL methods and at all cycles but one (with ReDAL). SeedAL yields an improvements of 3.3 pt and 9.7 pt in the first cycle over the best and worst random seed, respectively. We also observe that the results obtained with KMcentroid and KMfurthest are rather disappointing: the AL results are falling below the worst random seed at the first and second cycles. Surprisingly, ReDAL reaches better results with KMfurthest when using 3% of the dataset. However, our seeding strategy appears to be *agnostic* to the AL method and boost all of them consistently. Overall, the best results are again obtained with our seeding strategy and ReDAL.

SeedAL is especially effective for large-scale datasets, enabling scene-based AL methods to achieve almost 90% of the fully supervised performance using just 4% labeled points. This is comparable to the region-based method ReDAL which achieves 93% of the fully supervised performance. SeedAL bridges the gap between complex and costly methods and simpler AL methods on the large-scale SemanticKITTI dataset.

Improving random selection. rand is known to be a hard-to-beat baseline on this dataset [75]. We observe that thanks to SeedAL, the performance of *rand* goes as high as 58.0%, v.s. 58.9% for the best considered method (ReDAL). The

baseline rand is not itself an active learning strategy — given that scenes are randomly sampled — but these results show that a good seeding strategy impacts greatly all results, even those of the random baseline.

We also compare our method with LiDAL [34] on SemanticKITTI. LiDAL is an AL method specifically developed for 3D LiDAR semantic segmentation. As it is shown in Table 2, LiDAL benefits more from SeedAL than random selection. It gets 90% of full supervision’s mIoU (58.0%) with 2.1% labels, vs 2.6% with a random seed.

Seeding method	Init (1%)	2%	3%	4%	5%
rand [34]	48.8	57.1	58.7	59.3	59.5
SeedAL (ours)	52.6	57.8	59.3	60.3	60.6

Table 2: LiDAL [34] with random or SeedAL’s seeds.

5.4. Ablation experiments

We study the relevance of the different components used to construct SeedAL. We conduct experiments on S3DIS with all the active learning methods described in Sec. 5.1. Due to space constraint, we report CoreSet and ReDAL results and include the others in the supplementary material.

Intra-scene diversity. We highlight the relevance of selecting scenes with high intra-scene diversity instead of high intra-scene similarity. We compute a diversity and a similarity score per scene. Then we select the top scenes according to each metric until exhausting our annotation budget. The diversity and similarity scores are respectively computed by averaging $(1 - \phi(V_{N_i}^i) \cdot \phi(V_{N_j}^i))$ and $(\phi(V_{N_i}^i) \cdot \phi(V_{N_j}^i))$ between all views of a scene. Results in Figure 7 (a) show that using the intra-scene diversity yields the best results.

Clustering views. We show that computing the dissimilarity score from the cluster centroids as in (1) is a better strategy than computing it as in the previous paragraph. Performances obtained with both strategies are presented in Figure 7 (b). Clustering removes redundant features and yields a better estimate of the scene diversity, hence selection of a significantly better seed as demonstrated by the results for 3% of labeled points.

Inter-scene diversity. We justify why we select scenes with large inter-scene diversity rather than scenes with large inter-scene similarity. In a first experiment, we replace the edge weight e_{ij} in problem (4) by d_{ij} , the inter-scene dissimilarity of (2), and, in a second experiment, by the inter-scene similarity, obtained by averaging $(\Phi_k^i \cdot \Phi_{k'}^j)$ over all pairs (k, k') . Figure 7 (c) shows that selecting scenes with high inter-scene dissimilarity yields the best seeds.

Combined intra and inter-scene diversity. Finally, we show the interest of combining the intra and inter-scene diversity, i.e., using SeedAL, in Figure 7 (d). In this figure, we compare the performance of seeds selected by solving

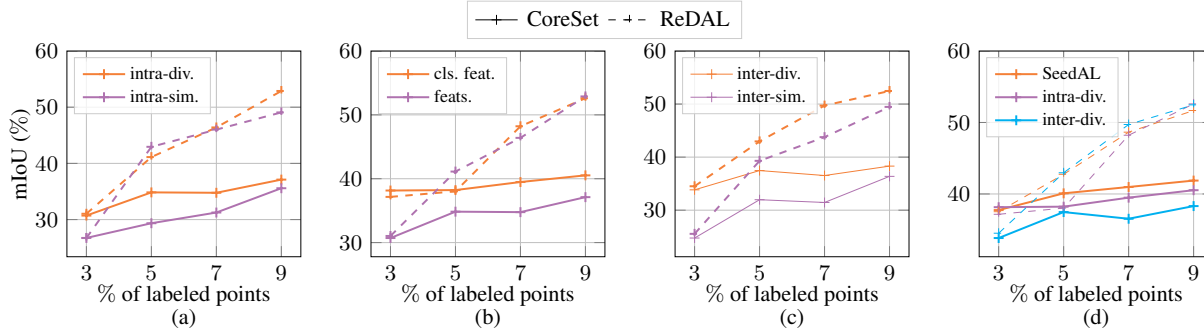


Figure 7: *Ablation study.* We evaluate here results obtained with different seeding strategies. (a) Seeds made of scenes with high intra-diversity (intra-div.) or high intra-similarity (intra-sim.). (b) Seeds selected with two different intra-diversity metrics: view features (feats.) or computed after clustering the view features (cls. feat.). (c) Seeds made of scenes with high inter-diversity (inter-div.) or high inter-similarity (inter-sim.). (d) Seeds selected with SeedAL, considering only inter-diversity (inter-div.) or intra-diversity (intra-div.).

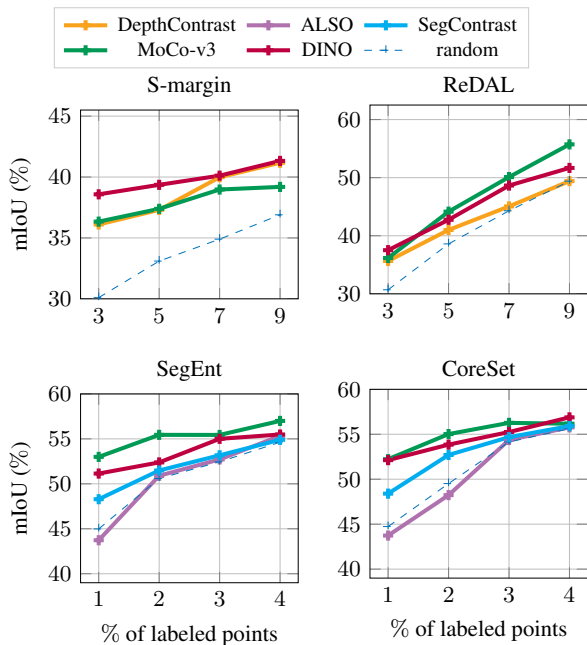


Figure 8: SeedAL results on S3DIS (first row) and SemanticKITTI (second row) using features from DepthContrast, SegContrast, ALSO, MoCo-v3, DINO. ‘random’ is an average over the random seeds.

equation (4) with $e_{ij} = d_i d_j d_{ij}$ (SeedAL), or $e_{ij} = d_{ij}$ (inter-div). We also show the performance of seeds made of scenes with top d_i (intra-div) intra-diversity.

Next we show the effect of different 2D and 3D self-supervised features to construct SeedAL. Due to space constraint, we report results from two different methods for each dataset and include the rest in supplementary material.

SeedAL with different 2D features. We experimented with self-supervised 2D features from MoCo-v3 [15] as an alternative to DINO. The results are reported in Figure 8. On SemanticKITTI, SeedAL with MoCo-v3 performs on par or

slightly better than SeedAL with DINO, but is slightly inferior to SeedAL with DINO on S3DIS, except for ReDAL. DINO and MoCo-v3 are pretrained on ImageNet, which is very different from S3DIS and SemanticKITTI; it shows the foremost generalizability of these features.

SeedAL with pure 3D features. We also experimented with self-supervised 3D features: DepthContrast [86] pretrained on ScanNet for indoor RGBD (no model is available for outdoors), as well as SegContrast [51] and ALSO [8] pretrained on SemanticKITTI. To adapt SeedAL to work with 3D features, we simply replace the set of 2D image features of a scene with 3D point features obtained from the point cloud of the same scene, which we thus average to obtain whole-scene features. Results are illustrated in Figure 8: DepthContrast is only slightly below 2D features, SegContrast is not as good but still better than rand, while ALSO underperforms, likely because it is not contrastive. Note that there is no pretrained 3D network that works both for indoors and outdoors, contrary to 2D-based models.

6. Conclusion

In this work, we have shown the influence of seeds on the performance of AL methods for point clouds, and proposed a method for efficiently selecting a seed that yields to good performance. Our approach works out of the box: it is agnostic to the AL method and to the 3D backbone used for the downstream task, and it does not require training on the dataset. We hope it will help scaling up point cloud semantic segmentation for practical applications.

Acknowledgments

We thank Samet Hicsonmez for insightful discussions. This work was granted access to HPC resources of IDRIS under GENCI allocations 2023-AD011013267R1, 2023-AD011013262R1 and 2023-AD011013413R1.

References

- [1] Saleh Albelwi. Survey on Self-Supervised Learning: Auxiliary Pretext Tasks and Contrastive Learning Methods in Imaging. *Entropy*, 24(4), 2022. 1
- [2] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. 3, 4, 6
- [3] Jordan T. Ash, Chicheng Zhang, Akshay Krishnamurthy, John Langford, and Alekh Agarwal. Deep batch active learning by diverse, uncertain gradient lower bounds. In *International Conference on Learning Representations*, 2020. 2
- [4] Adrien Bardes, Jean Ponce, and Yann LeCun. VICReg: Variance-Invariance-Covariance Regularization for Self-Supervised Learning. In *International Conference for Learning Representations (ICLR)*, 2022. 4
- [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 *International Conference on Computer Vision (ICCV)*, pages 9297–9307, 2019. 2, 3, 6
- [6] William H. Beluch, Tim Genewein, Andreas Nürnberger, and Jan M. Köhler. The Power of Ensembles for Active Learning in Image Classification. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. 3
- [7] Ksenia Bestuzheva, Mathieu Besançon, Wei-Kun Chen, Antonia Chmiela, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Oliver Gaul, Gerald Gamrath, Ambros Gleixner, Leona Gottwald, Christoph Graczyk, Katrin Halbig, Alexander Hoen, Christopher Hojny, Rolf van der Hulst, Thorsten Koch, Marco Lübbecke, Stephen J. Maher, Frederic Matter, Erik Mühmer, Benjamin Müller, Marc E. Pfetsch, Daniel Rehfeldt, Steffan Schlein, Franziska Schlösser, Felipe Serrano, Yuji Shinano, Boro Sofranac, Mark Turner, Stefan Vigerske, Fabian Wegscheider, Philipp Wellner, Dieter Weninger, and Jakob Witzig. The SCIP Optimization Suite 8.0. ZIB-Report 21-41, Zuse Institute Berlin, 2021. 6
- [8] Alexandre Boulch, Corentin Sautier, Björn Michele, Gilles Puy, and Renaud Marlet. ALSO: Automotive lidar self-supervision by occupancy estimation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. 9
- [9] Clemens-Alexander Brust, Christoph Kading, and Joachim Denzler. Active Learning for Deep Object Detection. In *International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP)*, 2019. 2
- [10] Holger Caesar, Varun Bankiti, Alex H. Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuScenes: A Multimodal Dataset for Autonomous Driving. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 3
- [11] Razvan Caramalau, Binod Bhattarai, and Tae-Kyun Kim. Visual Transformer for Task-Aware Active Learning, 2021. arXiv preprint arXiv:2106.03801. 4
- [12] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *International Conference on Computer Vision (ICCV)*, 2021. 2, 3, 4
- [13] Liangyu Chen, Yutong Bai, Siyu Huang, Yongyi Lu, Bihan Wen, Alan L. Yuille, and Zongwei Zhou. Making Your First Choice: To Address Cold Start Problem in Vision Active Learning. In *NeurIPS Workshop on Human in the Loop Learning*, 2022. 2, 3
- [14] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved Baselines with Momentum Contrastive Learning, 2020. 3
- [15] Xinlei Chen*, Saining Xie*, and Kaiming He. An empirical study of training self-supervised vision transformers. In *International Conference on Computer Vision (ICCV)*, 2021. 9
- [16] Mingmei Cheng, Le Hui, Jin Xie, and Jian Yang. SSPC-Net: Semi-Supervised Semantic 3D Point Cloud Segmentation Network. In *AAAI*, 2021. 2
- [17] Jaesung Choe, Chunghyun Park, Francois Rameau, Jaesik Park, and In So Kweon. PointMixer: MLP-Mixer For Point Cloud Understanding. In *European Conference on Computer Vision*, 2022. 2
- [18] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-Annotated 3D Reconstructions of Indoor Scenes. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 3
- [19] Sanjoy Dasgupta. Two Faces of Active Learning. *TCS*, 412(19):1767–1781, Apr. 2011. 3
- [20] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding. In *NAACLHLT*, 2019. 2
- [21] Di Feng, Xiao Wei, Lars Rosenbaum, Atsuto Maki, and Klaus Dietmayer. Deep Active Learning for Efficient Training of a LiDAR 3D Object Detector. *2019 IEEE Intelligent Vehicles Symposium (IV)*, pages 667–674, 2019. 2, 3
- [22] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep Bayesian Active Learning with Image Data. In *International Conference on Machine Learning*, pages 1183–1192, 2017. 3
- [23] Yarin Gal, Riashat Islam, and Zoubin Ghahramani. Deep Bayesian Active Learning with Image Data. In *International Conference on Machine Learning*, pages 1183–1192, 2017. 3, 6
- [24] Mingfei Gao, Zizhao Zhang, Guo Yu, Sercan Ö Arık, Larry S. Davis, and Tomas Pfister. Consistency-Based Semi-Supervised Active Learning: Towards Minimizing Labeling Cost. In *European Conference on Computer Vision*, 2020. 3, 4
- [25] Yonatan Geifman and Ran El-Yaniv. Deep Active Learning over the Long Tail. In *arXiv Preprint*, 2017. 2
- [26] Kyle Genova, Xiaoqi Yin, Abhijit Kundu, Caroline Pantofaru, Forrester Cole, Avneesh Sud, Brian Brewington, Brian Shucker, and Thomas Funkhouser. Learning 3D Semantic Segmentation with Only 2D Image Supervision. In *International Conference on 3D Vision*, pages 361–372, 2021. 2

- [27] Spyros Gidaris, Andrei Bursuc, Gilles Puy, Nikos Komodakis, Matthieu Cord, and Patrick Pérez. OBoW: Online Bag-of-Visual-Words Generation for Self-Supervised Learning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 4
- [28] Jean-Bastien Grill, Florian Strub, Florent Althé, Corentin Tallec, Pierre H. Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Daniel Guo, Mohammad Gheshlaghi Azar, Bilal Piot, Koray Kavukcuoglu, Rémi Munos, and Michal Valko. Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning. In *Advances in Neural Information Processing Systems*, 2020. 4
- [29] Guy Hacohen, Avihu Dekel, and Daphna Weinshall. Active Learning on a Budget: Opposite Strategies Suit High and Low Budgets. In *International Conference on Machine Learning*, 2022. 2, 3
- [30] Guy Hacohen, Avihu Dekel, and Daphna Weinshall. Active Learning on a Budget: Opposite Strategies Suit High and Low Budgets. In *International Conference on Machine Learning*, 2022. 3
- [31] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum Contrast for Unsupervised Visual Representation Learning. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 2
- [32] Yuenan Hou, Xinge Zhu, Yuexin Ma, Chen Change Loy, and Yikang Li. Point-to-Voxel Knowledge Distillation for LiDAR Semantic Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. 2
- [33] Neil Houlsby, Jose Miguel Hernandez-Lobato, and Zoubin Ghahramani. Cold-Start Active Learning with Robust Ordinal Matrix Factorization. In *International Conference on Machine Learning*, 2014. 3
- [34] Zeyu Hu, Xuyang Bai, Runze Zhang, Xin Wang, Guangyuan Sun, Hongbo Fu, and Chiew-Lan Tai. LiDAL: Inter-Frame Uncertainty Based Active Learning for 3D LiDAR Semantic Segmentation. In *European Conference on Computer Vision*, pages 248–265, 2022. 2, 3, 8
- [35] Siyuan Huang, Yichen Xie, Song-Chun Zhu, and Yixin Zhu. Spatio-Temporal Self-Supervised Representation Learning for 3D Point Clouds. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [36] Yilin Ji, Daniel Kaestner, Oliver Wirth, and Christian Wressnegger. Randomness Is the Root of All Evil: More Reliable Evaluation of Deep Active Learning. In *IEEE Winter Conference on Applications of Computer Vision*, 2023. 3, 4
- [37] Longlong Jing and Yingli Tian. Self-Supervised Visual Feature Learning with Deep Neural Networks: A Survey. *IEEE Transactions on Robotics and Automation*, 43(11):4037–4058, 2021. 1
- [38] Ajay J. Joshi, Fatih Porikli, and Nikolaos Papanikolopoulos. Multi-Class Active Learning for Image Classification. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. 1
- [39] Ksenia Konyushkova, Sznitman Raphael, and Pascal Fua. Learning Active Learning from Data. In *International Conference on Neural Information Processing*, pages 4228–4238, 2017. 3
- [40] Adrian Lang, Christoph Mayer, and Radu Timofte. Best Practices in Pool-Based Active Learning for Image Classification, 2022. OpenReview. 2, 3, 4
- [41] Xingjian Li, Pengkun Yang, Tianyang Wang, Xueying Zhan, Min Xu, Dejing Dou, and Chengzhong Xu. Deep Active Learning with Noise Stability, 2021. arXiv preprint arXiv:2205.13340. 4
- [42] Hanxue Liang, Chenhan Jiang, Dapeng Feng, Xin Chen, Hang Xu, Xiaodan Liang, Wei Zhang, Zhenguang Li, and Luc Van Gool. Exploring Geometry-Aware Contrast and Clustering Harmonization for Self-Supervised 3D Object Detection. In *International Conference on Computer Vision (ICCV)*, 2021. 2
- [43] Y. Lin, G. Vosselman, Y. Cao, and M. Y. Yang. Efficient Training of Semantic Point Cloud Segmentation via Active Learning. *SCIENCE*, 2:243–250, 2020. 2, 3, 6
- [44] Minghua Liu, Yin Zhou, Charles R. Qi, Boqing Gong, Hao Su, and Dragomir Anguelov. LESS: Label-Efficient Semantic Segmentation for LiDAR Point Clouds. In *European Conference on Computer Vision*, 2022. 2, 3
- [45] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. One Thing One Click: A Self-Training Approach for Weakly Supervised 3D Semantic Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1726–1736, June 2021. 2
- [46] 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. *SCIENCE*, PP:1–14, May 2018. 2
- [47] Yadan Luo, Zhuoxiao Chen, Zijian Wang, Xin Yu, Zi Huang, and Mahsa Baktashmotlagh. Exploring Active 3D Object Detection from a Generalization Perspective. In *International Conference for Learning Representations (ICLR)*, 2023. 2
- [48] Rafid Mahmood, Sanja Fidler, and Marc T. Law. Low-Budget Active Learning via Wasserstein Distance: An Integer Programming Approach. In *International Conference on Learning Representations*, 2022. 2, 3
- [49] David Maltz and Kate Ehrlich. Pointing the Way: Active Collaborative Filtering. In *CHI*, 1995. 3
- [50] Vishwesh Nath, Dong Yang, Holger R. Roth, and Daguang Xu. Warm Start Active Learning with Proxy Labels and Selection via Semi-Supervised Fine-Tuning. In *Conference on Medical Image Computing and Computer Assisted Intervention*, 2022. 2, 3
- [51] Lucas Nunes, Rodrigo Marcuzzi, Xieyuanli Chen, Jens Behley, and Cyrill Stachniss. SegContrast: 3D Point Cloud Feature Representation Learning through Self-Supervised Segment Discrimination. *IEEE Robotics and Automation Letters*, 7(2):2116–2123, 2022. 2, 9
- [52] Archit Parnami and Minwoo Lee. Learning from Few Examples: A Summary of Approaches to Few-Shot Learning, 2022. arXiv preprint arXiv:2203.04291. 1
- [53] Georgy Ponimatkin, Nermin Samet, Yang Xiao, Yuming Du, Renaud Marlet, and Vincent Lepetit. A Simple and Powerful Global Optimization for Unsupervised Video Object Segmentation. In *IEEE Winter Conference on Applications of Computer Vision*, 2023. 4

- [54] Kossar Pourahmadi, Parsa Nooralinejad, and Hamed Pirsiavash. A Simple Baseline for Low-Budget Active Learning, 2021. arXiv preprint arXiv:2110.12033. [2](#), [3](#), [4](#)
- [55] Guocheng Qian, Yuchen Li, Houwen Peng, Jinjie Mai, Hasan Abed Al Kader Hammoud, Mohamed Elhoseiny, and Bernard Ghanem. PointNeXt: Revisiting PointNet++ with Improved Training and Scaling Strategies. In *Advances in Neural Information Processing Systems*, 2022. [2](#)
- [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models from Natural Language Supervision. In *International Conference on Machine Learning*, 2021. [1](#)
- [57] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Brij B. Gupta, Xiaojiang Chen, and Xin Wang. A Survey of Deep Active Learning. *ACM Comput. Surv.*, 54(9), Oct. 2021. [1](#)
- [58] Zhongzheng Ren, Ishan Misra, Alexander G. Schwing, and Rohit Girdhar. 3D Spatial Recognition Without Spatially Labeled 3D. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13204–13213, June 2021. [2](#)
- [59] Cristiano Saltori, Fabio Galasso, Giuseppe Fiameni, Nicu Sebe, Elisa Ricci, and Fabio Poiesi. CoSMix: Compositional Semantic Mix for Domain Adaptation in 3D LiDAR Segmentation. In *European Conference on Computer Vision*, 2022. [2](#)
- [60] Cristiano Saltori, Stéphane Lathuilière, Nicu Sebe, Elisa Ricci, and Fabio Galasso. SF-UDA3D: Source-Free Unsupervised Domain Adaptation for LiDAR-Based 3D Object Detection. In *3DV*, 2020. [2](#)
- [61] Corentin Sautier, Gilles Puy, Spyros Gidaris, Alexandre Boulch, Andrei Bursuc, and Renaud Marlet. Image-To-LiDAR Self-Supervised Distillation for Autonomous Driving Data. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [2](#), [4](#)
- [62] Ozan Sener and Silvio Savarese. Active Learning for Convolutional Neural Networks: A Core-Set Approach. In *International Conference for Learning Representations (ICLR)*, 2018. [1](#), [2](#), [6](#)
- [63] Burr Settles. Active Learning Literature Survey. Technical Report, University of Wisconsin-Madison Department of Computer Sciences, 2009. [2](#)
- [64] Feifei Shao, Yawei Luo, Ping Liu, Jie Chen, Yi Yang, Yulei Lu, and Jun Xiao. Active learning for point cloud semantic segmentation via spatial-structural diversity reasoning. In *Proceedings of the 30th ACM International Conference on Multimedia*, 2022. [8](#)
- [65] Xian Shi, Xun Xu, Ke Chen, Lile Cai, Chuan-Sheng Foo, and Kui Jia. Label-Efficient Point Cloud Semantic Segmentation: An Active Learning Approach. In *arXiv Preprint*, 2021. [3](#)
- [66] Changjian Shui, Fan Zhou, Christian Gagné, and Boyu Wang. Deep Active Learning: Unified and Principled Method for Query and Training. In *International Conference on Artificial Intelligence and Statistics*, 2020. [4](#)
- [67] Oriane Siméoni, Gilles Puy, Huy V. Vo, Simon Roburin, Spyros Gidaris, Andrei Bursuc, Patrick Pérez, Renaud Marlet, and Jean Ponce. Localizing Objects with Self-Supervised Transformers and No Labels. In *British Machine Vision Conference*, 2021. [4](#)
- [68] 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, 2020. [2](#)
- [69] 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, 2020. [6](#)
- [70] 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 *International Conference on Computer Vision (ICCV)*, 2019. [2](#)
- [71] Jesper E. Van Engelen and Holger H. Hoos. A Survey on Semi-Supervised Learning. *Machine Learning*, 109(2):373–440, 2020. [1](#)
- [72] Dan Wang and Yi Shang. A New Active Labeling Method for Deep Learning. In *International Joint Conference on Neural Networks*, pages 112–119, 2014. [2](#), [6](#)
- [73] Yangtao Wang, Xi Shen, Shell Xu Hu, Yuan Yuan, James L. Crowley, and Dominique Vaufreydaz. Self-Supervised Transformers for Unsupervised Object Discovery Using Normalized Cut. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. [4](#)
- [74] 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 *Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020. [2](#)
- [75] Tsung-Han Wu, Yueh-Cheng Liu, Yu-Kai Huang, Hsin-Ying Lee, Hung-Ting Su, Ping-Chia Huang, and Winston H. Hsu. ReDAL: Region-Based and Diversity-Aware Active Learning for Point Cloud Semantic Segmentation. In *International Conference on Computer Vision (ICCV)*, pages 15510–15519, 2021. [2](#), [3](#), [6](#), [8](#)
- [76] Saining Xie, Jiatao Gu, Demi Guo, Charles R. Qi, Leonidas J. Guibas, and Or Litany. PointContrast: Unsupervised Pre-Training for 3D Point Cloud Understanding. In *European Conference on Computer Vision*, pages 574–591, 2020. [2](#)
- [77] Qiangeng Xu, Yin Zhou, Weiyue Wang, Charles R. Qi, and Dragomir Anguelov. SPG: Unsupervised Domain Adaptation for 3D Object Detection via Semantic Point Generation. In *International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [78] Xun Xu and Gim Hee Lee. Weakly Supervised Semantic Point Cloud Segmentation: Towards 10x Fewer Labels. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#)
- [79] Xu Yan, Jiantao Gao, Chaoda Zheng, Chao Zheng, Ruimao Zhang, Shuguang Cui, and Zhen Li. 2DPASS: 2D Priors Assisted Semantic Segmentation on LiDAR Point Clouds. In *European Conference on Computer Vision*, 2022. [2](#)

- [80] Jihan Yang, Shaoshuai Shi, Zhe Wang, Hongsheng Li, and Xiaojuan Qi. ST3D: Self-Training For Unsupervised Domain Adaptation on 3D Object Detection. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [2](#)
- [81] Ofer Yehuda, Avihu Dekel, Guy Hacohen, and Daphna Weinshall. Active Learning through a Covering Lens. In *Advances in Neural Information Processing Systems*, 2022. [2](#), [3](#), [4](#)
- [82] Ofer Yehuda, Avihu Dekel, Guy Hacohen, and Daphna Weinshall. Active learning through a covering lens. In *Advances in Neural Information Processing Systems*, 2022. [3](#)
- [83] Li Yi, Boqing Gong, and Thomas Funkhouser. Complete & Label: A Domain Adaptation Approach to Semantic Segmentation of LiDAR Point Clouds. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [2](#)
- [84] Michelle Yuan, Hsuan-Tien Lin, and Jordan Boyd-Graber. Cold-Start Active Learning through Self-Supervised Language Modeling. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7935–7948, Nov. 2020. [2](#), [3](#), [6](#)
- [85] Yachao Zhang, Zonghao Li, Yuan Xie, Yanyun Qu, Cuihua Li, and Tao Mei. Weakly Supervised Semantic Segmentation for Large-Scale Point Cloud. In *American Association for Artificial Intelligence Conference*, 2021. [2](#)
- [86] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-Supervised Pretraining of 3D Features on Any Point-Cloud. In *International Conference on Computer Vision (ICCV)*, pages 10252–10263, 2021. [2](#), [9](#)
- [87] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip H. S. Torr, and Vladlen Koltun. Point Transformer. In *International Conference on Computer Vision (ICCV)*, 2021. [2](#)
- [88] Zhi-Hua Zhou. A Brief Introduction to Weakly Supervised Learning. *SCIENCE*, 5(1):44–53, 2017. [1](#)
- [89] Xinge Zhu, Hui Zhou, Tai Wang, Fangzhou Hong, Yuexin Ma, Wei Li, Hongsheng Li, and Dahua Lin. Cylindrical and Asymmetrical 3D Convolution Networks for LiDAR Segmentation. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. [2](#)