

SnapshotNet: Self-supervised feature learning for point cloud data segmentation using minimal labeled data

Xingye Li ^{a,*}, Ling Zhang ^a, Zhigang Zhu ^{a,b}

^a The City College, The City University of New York, 160 Convent Avenue, New York, NY 10031, USA

^b The Graduate Center, The City University of New York, 365 5th Avenue, New York, NY 10016, USA



ARTICLE INFO

Communicated by Nikos Paragios

MSC:

41A05

41A10

65D05

65D17

Keywords:

Self-supervision

Point cloud

Semantic segmentation

ABSTRACT

Manually annotating complex scene point cloud datasets is both costly and error-prone. To reduce the reliance on labeled data, a new model called SnapshotNet is proposed as a self-supervised feature learning approach, which directly works on the unlabeled point cloud data of a complex 3D scene. The SnapshotNet pipeline includes three stages. In the snapshot capturing stage, snapshots, which are defined as local collections of points, are sampled from the point cloud scene. A snapshot could be a view of a local 3D scan directly captured from the real scene, or a virtual view of such from a large 3D point cloud dataset. Snapshots could also be sampled at different sampling rates or fields of view (FOVs), thus multi-FOV snapshots, to capture scale information from the scene. In the feature learning stage, a new pre-text task called multi-FOV contrasting is proposed to recognize whether two snapshots are from the same object or not, within the same FOV or across different FOVs. Snapshots go through two self-supervised learning steps: the contrastive learning step with both part contrasting and scale contrasting, followed by a snapshot clustering step to extract higher level semantic features. Then a weakly-supervised segmentation stage is implemented by first training a standard SVM classifier on the learned features with a small fraction of labeled snapshots. Then trained SVM is further used to predict labels for input snapshots and predicted labels are converted into point-wise label assignments for semantic segmentation of the entire scene using a voting procedure. The experiments are conducted on the Semantic3D dataset and the results have shown that the proposed method is capable of learning effective features from snapshots of complex scene data without any labels. Moreover, the proposed weakly-supervised method has shown advantages when comparing to the state of the art method on weakly-supervised point cloud semantic segmentation.

1. Introduction

Studies on 3D point cloud data have been gaining momentum in the field of computer vision. Deep neural networks such as PointNet (Qi et al., 2017a), DGCNN (Wang et al., 2019) have been proposed for better performances on point cloud related tasks, with the help of larger datasets such as the ModelNet (Wu et al., 2015) and Semantic3D (Hackel et al., 2017). The collective effort between deep neural networks and dedicated datasets continues to push the state of the art performance on the point cloud object classification.

On the other hand, point cloud semantic segmentation is of great interests in the applications of autonomous driving, robotics and remote sensing (Xie et al., 2020b). So far most of the deep learning driven point cloud semantic segmentation methods follow the supervised workflow, which requires densely labeled datasets, such as the 1.6 million points Oakland Dataset (Muñoz et al., 2009), the 215 million points Stanford Large-scale 3D Indoor Spaces Dataset (S3DIS) dataset (Armeni et al.,

2016) and the 4 billion points Semantic3D (Hackel et al., 2017). However, annotating large scale datasets is at a very high cost both in time and human labors. This issue is becoming more prominent in applications such as hazard assessment where drive-by and fly-by LiDAR mapping systems have been used to collect massive windstorm damage datasets in recent hurricane events (Bhargava et al., 2019; Gong et al., 2012; Gong, 2013; Hu and Gong, 2018). The fact that LiDAR is starting to be integrated into smaller mobile devices (Apple Inc., 2020), which could lead to a boom in the scale of real life complex point cloud data.

To alleviate the dependence on the labels of large datasets, unsupervised learning methods have drawn increasing attention. Among the unsupervised methods, one form known as “self-supervised learning” has been popular in the studies of image data understanding. This self-supervised approach has found success in designing “pre-text” tasks, such as jigsaw puzzle reassembly (Noroozi and Favaro, 2016), image clustering (Caron et al., 2018) and image rotation prediction (Jing et al., 2018) etc, by training deep learning models for

* Corresponding author.

E-mail address: xli020@citymail.cuny.edu (X. Li).

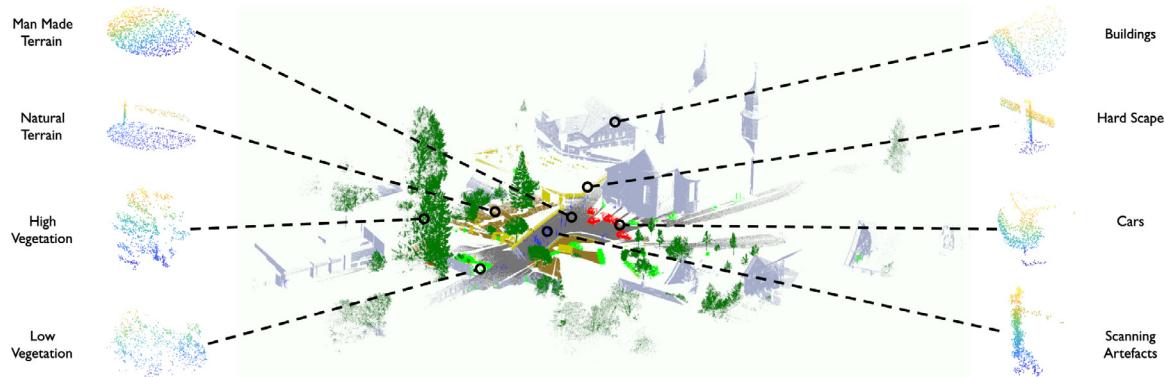


Fig. 1. Visualization of some snapshots sampled from the Semantic3D dataset. The sampling procedure makes no use of labels, therefore a snapshot may contain points from other classes. The class labels are added manually for visualization.

feature extraction without labels being involved. Based on the idea of solving pretext tasks, we previously developed the model of Contrast-ClusterNet (Zhang and Zhu, 2019), which works on unlabeled point cloud datasets by part contrasting and object clustering. While this work has shown comparable performance to its supervised counterparts on synthetic point cloud objects classification, it inherits the problem of other pretext-driven models used on image data: pretext tasks must be defined regarding the prior knowledge embedded in the data. In the context of point cloud understanding, the part contrasting and object clustering tasks assume the input data are well separated as individual objects. This assumption limits the model's power on real life scene data or where 3D data of single objects cannot be easily obtained. We addressed this limitation posed on scene point cloud by proposing a snapshot-based method (Li. and Zhu., 2021), which captures local clusters of points called snapshots from the scene as input samples to fulfill the tasks of part contrasting and object clustering. Formally, a snapshot is defined as a collection of points, sampled from a point cloud scene, without knowing their labels or even assuming they are from the same objects. (Fig. 1). It could be a real view of a local 3D scan directly captured from the real scene, or a “virtual” view of such a local 3D scan from a large 3D point cloud dataset. The effectiveness of this method is evaluated and approved by conducting classification on the captured snapshots with a single FOV in (Li. and Zhu., 2021). In this paper, we extend the idea to capture multi-FOV snapshots to further improve the performance of classification.

Besides from the limitation of making assumptions on the training data, another weakness of the Contrast-ClusterNet is that dense labels are still needed for the downstream tasks. The full supervision involved in the object classification contradicts the main goal of self-supervision, that to save labeling efforts on training data. To extend the idea of reducing labeled data usage to the downstream tasks, we seek solutions from weak supervision.

Therefore, based on the two pieces of our previous work, we further propose the SnapshotNet, which integrates multi-FOV snapshot generation, contrastive feature learning, and a weakly-supervised technique for point-wise scene segmentation using a voting mechanism. First of all, inspired by the observation that, humans are able to distinguish objects at different scales, we present a new pre-text task for contrastive learning, namely multi-FOV contrasting. When capturing a sample, we take multiple snapshots in different field-of-views (FOVs). Assuming these multi-FOV snapshots are small enough so they still represent the same object, the task of scale contrasting is to consider whether two snapshots, within one FOV or across multiple FOVs, are of the same object or not. Thus the multi-FOV contrasting includes two parts: part contrasting and scale contrasting. We will compare the performance of the scale contrasting against the part contrasting, as well as the combination of the two — the multi-FOV contrasting.

For the complete model pipeline, the captured snapshots first go through a two-step self-supervised pipeline using ContrastNet and ClusterNet consecutively for feature learning. Then a weakly-supervised

approach is implemented by training an SVM classifier on the learned features of a small portion of labeled snapshot samples (mostly cluster centers) combined with the samples automatically labeled from the clusters of the samples generated in the pipeline. Finally, the entire 3D point cloud scene is repeatedly scanned as random snapshots to go through the feature extractor and classifier. The predicted snapshot-wise label is assigned to each point of a snapshot, followed by a voting-based mechanism for the final label for each point.

This work makes the following contributions:

(1) We propose a new contrastive learning method called multi-FOV contrasting, by leveraging point cloud samples at different scales. This task devotes on predicting if two snapshots are of the same object, regardless of their sampling FOVs.

(2) We develop a three-stage approach for semantic segmentation: snapshot generation, self-supervised feature learning, and point-wise segmentation by integrating multiple weakly-supervised classified results.

(3) We study the “purity” of snapshots, and show that the self-supervised learning with impure snapshots can still effectively obtain highly useful semantic features for object classification and scene segmentation. This includes cases when some of the classes do not have well-sampled snapshots.

(4) By using the learned features and clustering to obtain larger pseudo labels with a small number of labels (thus weakly-supervised) to train a simple classifier, we design a simple voting procedure to integrating labels of randomly sample snapshots, which leads to point-wise point cloud scene segmentation performance comparable to the state of the art weakly-supervised methods.

The rest of the paper is organized as the following. Section 2 discusses related work on self-supervised learning methods on point cloud, and point cloud semantic segmentation. Section 3 describes the theory and design of the SnapshotNet for self-supervised feature learning and point cloud semantic segmentation with fewer labeled data. Section 4 details the experimental results, including the designs and evaluations of data capturing, feature learning, and segmentation. Finally Section 5 concludes the work with discussions of a few ideas for future work.

2. Background and related work

2.1. Self-supervised learning:

Self-supervised learning aims to predict for output labels that are generated from the intrinsic information of the data. This topic has been widely studied on the image data where various of pre-text tasks have been proposed, such as context prediction (Doersch et al., 2015), jigsaw puzzle reassembly (Noroozi and Favaro, 2016), image clustering (Caron et al., 2018), and image rotation prediction (Jing et al., 2018) etc, and these methods have demonstrated considerable results on ordered data such as 2D images or videos.

With the advancement in LiDAR technology, the cost for obtaining large scale point cloud data has enormously decreased. The booming in 3D point cloud data has turned the challenge from data collection to manual annotation, which is much more difficult and laborious compared to 2D data. To alleviate the use of labeled data, a number of self-supervised models have been proposed lately (Achlioptas et al., 2018; Yang et al., 2018; Zhang and Zhu, 2019; Sauder and Sievers, 2019; Xie et al., 2020a; Chen et al., 2020). In previous work of our lab, Zhang and Zhu proposed the Contrast-ClusterNet (Zhang and Zhu, 2019) with pre-text tasks of first predicting whether two segments are from the same object, leading to the ContrastNet for obtaining self-learned features, which are then used for separating the objects into different clusters using KMeans++, for training another network called ClusterNet to obtain better self-learned features. The work (Zhang and Zhu, 2019) has shown the capability of learning features in a self-supervised manner, and then using the features, an SVM classifier can be trained using labeled data for point cloud objects classification. However, this process still requires to know a set of 3D points belong to a single object (even though the label is not needed). In training the SVMs, the same amount of labeled data as in supervised models is used, therefore decreasing the benefits of leaving out annotations in self-supervised learning.

2.2. Semantic segmentation of point cloud:

With the recent works shifting focus to adapting deep learning on LiDAR point cloud data, a series of deep learning based point cloud semantic segmentation methods have been proposed. As summarized by Guo et al. (Guo et al., 2020), there are several mainstream semantic segmentation methods on point cloud data, such as the discretization-based, projection-based, and point-based methods.

The discretization-based approach is greatly inspired by the success of deep learning on 2D grid data, where the 2D data is in a regular representation, in contrast to the unordered 3D point cloud. A number of works have been proposed using the voxel-based models (Rethage et al., 2018; Hackel et al., 2017; Huang and You, 2016; Tchapmi et al., 2017), which voxelize the point cloud data to 3D grids to enable direct 3D convolutional feature extraction. Despite that this method has made significant progress on point cloud segmentation, it is very sensitive to the voxel resolution and often has strict requirements on memory and computational power.

The projection-based method, on the other hand, has shown advantages on computation efficiency. As a representation of this approach, 2D multi-views models are designed to project a 3D point cloud to 2D views from multiple directions, so that traditional convolutional networks can be applied for semantic segmentation tasks (Lawin et al., 2017; Boulch et al., 2017; Alonso et al., 2020). However, the downside of this approach is that geometrical information is often lost during the dimension reduction.

PointNet (Qi et al., 2017a) is the first deep net proposed to directly work with point cloud data without the pre-processing step of transforming the raw point cloud into voxels or 2D multi-views representations. To help catching local geometrical context, the PointNet++(Qi et al., 2017b) is developed by proposing a hierarchical network based on the PointNet. The idea of exploiting local structures of the 3D data is further explored by developing dynamic graph CNN (DGCNN) (Wang et al., 2019), which uses graphs the geometrical relations of the point cloud and operate convolutions on such graphs.

The above segmentation methods mostly rely on densely labeled data, and such datasets are proven to be costly on time and human labors. There are few works focusing on weakly-supervised scene point cloud semantic segmentation: the segmentation-aided classification (Guinard and Landrieu, 2017) is a non-parametric method, using conditional random field (CRF) to process the output of a pointwise classifier. The pseudo-labeling approach (Xu et al., 2019) trains a PointNet (Qi et al., 2017a) with a handful of labeled points and

gradually assigns pseudo-labels that are generated from the trained PointNet model to all unlabeled points, and model is also iteratively updated with more reliable pseudo-labels. Xu and Lee (Xu and Lee, 2020) proposed an incomplete supervision model with three additional training losses to constrain the model. Among them, two pieces of work (Guinard and Landrieu, 2017; Xu et al., 2019) that also worked on the outdoor datasets as ours are the baselines that our proposed SnapshotNet will be compared with.

3. Theory and design

Self-supervised learning often requires prior knowledge about the input data to ensure the intrinsic information of the data, from which the labels are derived, is consistent across all samples. This is also the case of the Contrast-ClusterNet (Zhang and Zhu, 2019), which will be used as the base model of our proposed work. As will be summarized below, it has two major modules called ContrastNet and ClusterNet. Each module is centered on a deep learning neural network DGCNN (Wang et al., 2019) capable of extracting features from the point cloud inputs. First the ContrastNet takes inputs of paired point cloud segments, which are obtained by randomly cutting the point cloud object into two halves. The job of the ContrastNet is to consider whether two segments of a pair are from the same object or not, essentially doing the task of binary classification. The trained ContrastNet is capable of extracting features at high-level due to the nature of the pre-text task of part-contrasting. The second module, the ClusterNet, is to obtain more representative and fine-grained features. Before starting training the ClusterNet, features of the raw point cloud objects are extracted by the trained ContrastNet, and these features are subsequently clustered into a much larger number of groups (than the number of object categories) using Kmeans++; in (Zhang and Zhu, 2019), experimental studies were also performed for the optimal numbers of clusters. Each object is then assigned with their cluster ID as a pseudo-label for the training of the ClusterNet.

Although there are no labels being involved in this two-step feature learning process, the nature of self-supervised learning requires some prior assumptions regarding the pre-text tasks that drive the self-supervision. In this example, such assumption is that each training sample must be an individual point cloud object to enable part contrasting. This assumption can be easily made on datasets such as ModelNet (Wu et al., 2015) and ShapeNet (Chang et al., 2015), where each sample is a synthetic CAD model of a single 3D object. However, this soon becomes a limitation on real-life point cloud datasets, such as the Okaland (Munoz et al., 2009) and Semantic3D (Hackel et al., 2017) data, where an entire point cloud is a complex scene rather than individual single objects.

To address this issue, we thus propose the SnapshotNet for the self-supervised feature learning and weakly-supervised semantic segmentation on complex scene point cloud. As illustrated in Fig. 2, our method consists of three modules: (a) snapshot capturing, (b) feature learning, and (c) segmentation. The snapshot capturing procedure, as an analogy to taking snapshots with a 3D camera, captures small areas of the entire point cloud to train the model. Then the feature learning module uses the Contrast-ClusterNet (Zhang and Zhu, 2019) as the backbone for self-supervised feature learning. Finally, in the segmentation module, a classifier is trained on few labeled data and the pseudo-labeled data for snapshot classification. A voting mechanism is followed to convert snapshot-wise predictions to point-wise predictions, achieving the goal of semantic segmentation. Each part of the pipeline will be described in details in the following subsections.

3.1. Snapshot capturing

Given a real-life point cloud dataset, the *snapshot capturing* stage applies random sampling with k-Nearest Neighbors (kNN) to obtain small collections of points as snapshots (Fig. 1). During each sampling,

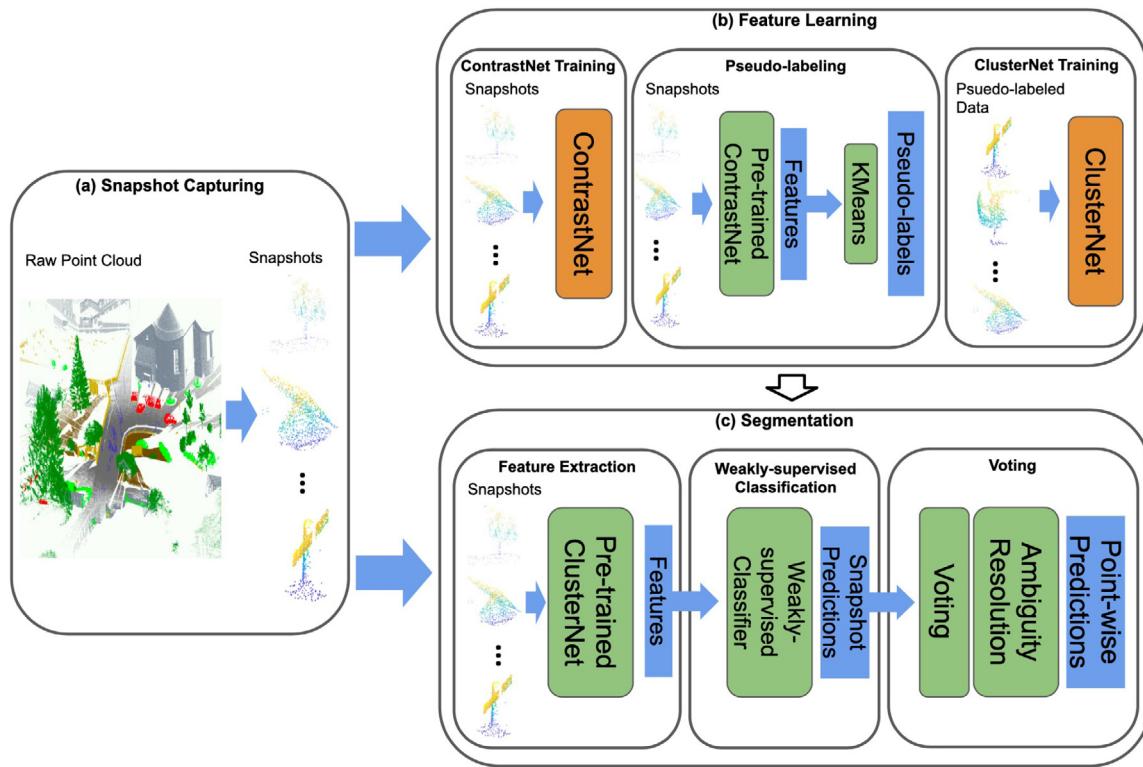


Fig. 2. The SnapshotNet pipeline: (a) Snapshot capturing from the raw point cloud scenes; (b) Feature learning by conducting contrasting tasks, snapshot clustering and cluster classification; (c) Semantic segmentation by classifying and voting on snapshots.

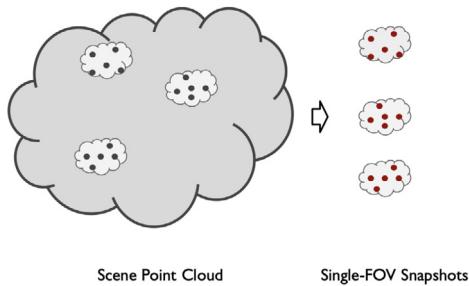


Fig. 3. Single-FOV snapshots sampling: an illustration.

an anchor point is randomly selected from the point cloud at first, and kNN gives a collection of k points nearest to the anchor point, where k defines the snapshot sampling area. This kNN strategy is a simulation of a virtual snapshot of a local 3D view, followed by point selection based on their 3D proximity to better ensure the sameness of an object. Each collection is therefore called a ‘snapshot’ of the local neighborhood in the bigger point cloud pool. Here, thus sampled snapshots share the same sampling rate with the local areas where the anchors are picked up from the scene, meaning that the area covered by each snapshot is determined by the local scene sampling rate. In other words, the snapshots have one single field of view (FOV), so they are also notated as the single-FOV snapshots (Fig. 3).

Purity of snapshots: Since the selection of an anchor point happens randomly in the point cloud, it is possible to have the anchor sitting close to the border between different semantic classes. This introduces a certain degree of noises to the snapshot by including some points from other minority classes. Compared to the object based contrasting pretext task, which in this paper is notated as ObjectNet for easy comparison with the SnapshotNet, our method further relaxes the constraint that 3D points of a sample must come from the same object.

The SnapshotNet fundamentally sees each snapshot as a collection of points that represents a small region of the bigger complex scene, where such a collection of points has a high probability of belonging to the same class. In our experiment section, we will show how the noises in snapshots will affect the performance of feature learning for later evaluation.

To quantify the noise level of sampled snapshots, we present a metric to evaluate our snapshot sampling quality, namely purity. When sampling from the Semantic3D (Hackel et al., 2017), we utilize the provided labels of the dataset to approximate the semantic label of each snapshot for the sake of snapshot classification evaluation. A label is assigned to a snapshot by voting from all points that are associated with that sample. The class label that most points agree on is chosen as the semantic label for the snapshot, the voting procedure can be parameterized as:

$$C_x = \operatorname{argmax}_i \sum_{j=1}^K I(y_j = i), \quad (1)$$

where x is the snapshot sample, y_j is the point-wise label for x ($j=1, \dots, K$), K represents the number of points in the snapshot x , and I is an indicator function for the class of each point. Thus the purity is given by

$$P(x) = \frac{\sum_{j=1}^K I(y_j = C_x)}{K} \quad (2)$$

The statistics of the voted semantic labels and the purity for each sample will be further discussed in Section 4.1 using real examples.

Multi-FOV snapshots: Inspired by zooming with a camera while taking a photo, it soon came to us that a different field of view(FOV) of a snapshot image leads to different information content. Given the same sensor size, a larger FOV might contain more objects at low details, while a smaller FOV focuses on smaller views with greater details. We adopt this observation into our design of the point cloud snapshots, to include multiple FOVs for each snapshot.

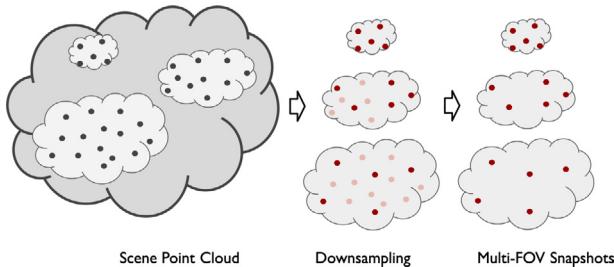


Fig. 4. Multi-FOV snapshots sampling. Three snapshots of different FOVs are pre-sampled from the scene and each one has a different size. Then the samples are downsampled to meet the same input size of the network, leading to different resolutions.

Just as the sensor size poses limitation on imaging when zooming in or out, the neural network has similar constraints when dealing with samples with different FOVs. The number of points in each snapshot is kept as the same, and this leaves only the sampling rate to be altered. Therefore, we keep the original sampling rate of the point cloud as the base, and accordingly decrease the sampling rate by grouping points at a sparser scale. Specifically, this is achieved by pre-sampling a larger number of points at the base sampling rate and then randomly drop some points to meet the input size (i.e., the number of 3D points). Fig. 4 illustrates an example of sampling snapshots in three FOVs. When sampling a snapshot at the base sampling rate, the down-sampling can be well ignored. However, when capturing a sample with an FOV two times larger than the base FOV, a pre-sampled snapshot of double amount of the points are first captured using kNN. This pre-sampled snapshot is then downsampled by a factor of two to agree on the network input size while at half the base sampling rate.

There are several intuitions behind this multi-FOV design. As described earlier, the part-contrasting exploits information at the object level by performing binary classifications on samples that share uniform sampling rate. On the other hand, the human vision is able to recognize objects at very different scales, which encourages us to further make use of the scale information. The multi-FOV snapshots are able to fill in the gap of the missing scale information, giving us an edge on contrastive learning by contrasting on various scales in addition to the part-contrasting. Secondly, the multi-FOV snapshots serve as an approach of data augmentation, to diversify the input data and indirectly making the contrasting learning more challenging to the network. Furthermore, single sampling rate is inadequate when facing a scene point cloud with objects of various scales and with different sampling resolutions. This is particularly a problem for the terrestrial scans, where the density of points rapidly changes along the distance to the scanning device. When the network is trained to take the sampling rate into account, there is the opportunity to explicitly choose an FOV that is more suitable for sampling a specific object from the scene. For instance, we would want to sample a snapshot of a small object using a small FOV to maximize the purity, and on the other hand to keep a large FOV on larger objects. This will be discussed in more details in Section 3.3.

3.2. Self-supervised learning with snapshots

After being captured from the scene, each ‘snapshot’ is viewed as a single point-cloud object and fed into the two-stage ContrastNet-ClusterNet for *feature learning*. Both networks are based on the DGCNN (Wang et al., 2019), therefore they are similar in structures.

Contrastive feature learning: The contrastive learning includes three approaches: part-contrasting, scale-contrasting, and multi-FOV contrasting. When conducting the *part-contrasting* during the training of a ContrastNet, we follow the random cutting procedure as described

in (Zhang and Zhu, 2019): two segments from the same snapshot make up a positive pair, which is labeled as 1, and on the contrary, and a negative pair consists of two segments from two different snapshots is labeled as 0 (Fig. 5(a)). The ContrastNet then learns to recognize whether the input pair is positive or negative, and the parameters are optimized by the Adam optimizer on the cross-entropy loss.

The part-contrasting considers the similarity between different parts of an object in its single-FOV snapshot, thus learning fine-grained features. The *scale-contrasting*, on the other hand, attempt to learn higher-level features for representing similarity between snapshots of an object across different scales (i.e., with different FOVs). For instance, the details of an object might get lost in a very small FOV, yet the model is still required to correctly connect this sample to its large FOV counterparts without these details. To implement this method, we similarly make up pairs from the multi-FOV snapshots: two snapshots in whichever FOVs sampled from the same anchor point form a positive pair, given a label as 1. Two snapshots from two different sampling anchors form a negative pair with a label of 0, as shown in Fig. 5(b).

The part contrasting and scale contrasting focus on very dissimilar goals, but leading to different levels of features. However, these two pre-texts are not mutually exclusive when governing the self-supervised learning. Our design of the multi-FOV snapshots provides additional room to join these two tasks when forming the training sample pairs, and we name this combination as *multi-FOV contrasting*. Now a positive pair is not limited to coming from two segments of the same single-FOV sample, we can also take two cross-FOV segments from the same sample as a positive pair, and vice versa for a negative pair (Fig. 5(c)). This formation of sample pairs is expected to push the model into learning both fine-grained and high-level abstract features.

Clustering for feature refinement: Once the ContrastNet is well trained with one of these pre-texts, we continue to adopt the idea of knowledge transfer for more refined features by learning similarities and differences of samples across different snapshots. Before starting training the ClusterNet, the learned features from the ContrastNet are used to cluster the snapshot samples into k groups with KMeans++. These cluster (group) labels are treated as pseudo-labels for the snapshots to train the ClusterNet. We use $k=300$ to cluster the snapshots of all FOVs, into new classes based on the studies in our previous work (Zhang and Zhu, 2019). Note that this number is much greater than the number of the existing semantic labels in the Semantic3D dataset (which is eight); however the large cluster number forces the ClusterNet to learn fine-grained features. The loss function defined in the work of the ClusterNet (Zhang and Zhu, 2019) is described as:

$$\min_{\theta, W} \frac{1}{N} \sum_{n=1}^N l(g_W(f_\theta(x_n)), y_n) \quad (3)$$

where the g_W is the classifier that predicts for the correct pseudo-labels y_n given the features $f_\theta(x_n)$.

3.3. Semantic segmentation with snapshots

The semantic segmentation has three major components, namely feature extraction, weakly-supervised classification, and point-wise semantic segmentation via voting ((c) in Fig. 2). The feature extraction step is a straightforward process that takes snapshots captured from the raw point cloud and extracts the deep features using the already trained ClusterNet, as described above. We then use a small fraction of the extracted features along with their labels to train an SVM classifier. This classifier serves two purposes: one is to evaluate the self-supervised features learned by the SnapshotNet in the experiments, and the second purpose is to serve as a base classifier that will further diffuse all snapshot predictions into point-wise predictions.

Classification with weak supervision and pseudo labeling: Following the self-supervised feature learning, a classifier is trained on the extracted features of labeled training data for classification. Conventionally, the training process requires as many labeled data as possible

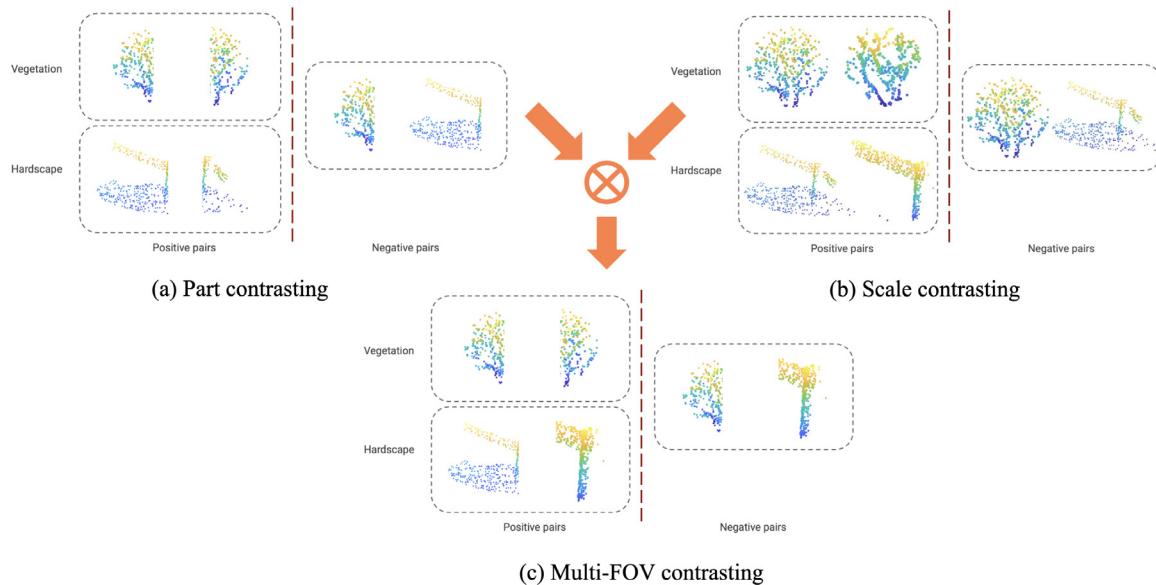


Fig. 5. Three approaches of contrastive learning by forming positive pairs or negative pairs from two samples. In the part contrasting (a), each snapshot is split into two parts and then pairs are formed by putting together of any two parts. On the left of (a), the left and right side of the tree form a positive pair because they are from the same snapshot, and the same goes for the fence, whereas on the right of (a) the tree and fence form a negative pair. The scale contrasting (b) takes two snapshots across different field-of-views as a pair. Here two snapshots of the same tree or same fence (but at different scales) form a positive pair, and a negative pair comes from different objects. The Multi-FOV contrasting (c) combines the previous two approaches by cutting each snapshot into two parts across all FOVs, and forms positive pairs where the two parts are from the same object regardless of their sampling FOVs. A negative pair has two parts from different objects.

for better performance. However, this approach in its essence is in contradiction with the objective of self-supervised learning, which aims to reduce the dependency on labeled data. Therefore, we seek solutions from weak supervision to reduce the reliance on dense labels for the downstream tasks following our self-supervised feature learning.

Here the weak supervision can be viewed from two perspectives. First is that when there are only a few labels available, we still wish to achieve comparable classification performance with the limited labels. Furthermore, the labels assigned to the snapshots are essentially coarse-grained labels because instead of point-wise labeling, each snapshot is labeled as a whole, regardless of the noises included during the sampling.

Second, a pseudo labeling technique is proposed to acquire larger training data population to feed the classifier. This technique is incorporated into the KMeans++ clustering in the feature learning module, hence named cluster-based pseudo labeling. Fig. 6 visualizes the 300 clusters of training samples using KMeans++ in the feature space against their pseudo labels. Due to the large number of clusters, each one of them facilitates only a few to a few hundreds of samples. This large collection of clusters breaks all samples into smaller groups by their similarities in the feature space, where each group hosts way less but highly alike samples. This can be seen from Fig. 6, when visualizing the clusters against their semantic labels. This property can be well used by just giving one label to each cluster and assign this label as the pseudo labels of some of the most related samples to that labeled sample in the cluster center. The selection of the nearby samples can be designed to work geometrically or statistically. In this work, a threshold is introduced to constrain the measurement of the normalized distance between each point to their cluster center. A strict threshold filters out samples far from the center to gain more accurate pseudo labeling.

Semantic segmentation by voting: Associated with our first assumption that, statistically a snapshot is able to represent a small piece of an object, we can assume that all points included in a snapshot are highly likely to belong to the same class predicted by the classifier for this snapshot. The predicted class label is assigned all points in the snapshot. Thus the point-wise segmentation problem is converted into an object classification problem. Statistically, if the snapshot capturing happens randomly, all snapshots are able to cover the whole scene after certain

number of iterations. Therefore by repeating the capture-predict-assign procedure, all points in a scene eventually get a predicted label.

The model keeps count of the points with at least one prediction to track the progress of segmentation, and a cut off threshold is set to stop the snapshot capturing. When 99.95% of the points are assigned with a prediction, the model stops taking new samples and moves to the next step of voting. Due to the randomness of the snapshot capturing, it is expected to have multiple snapshots covering the same points from the scene, which potentially assigns multiple labels to one point. To reach for a final agreement on the label, a voting procedure is designed to select the dominant label with most counts to be the final decision for each point. It is also possible that some points obtained equal numbers of different labels by the time the snapshot capturing stops. For points with such labeling conditions and particularly near the boundary of different semantic classes, it is likely two labels have the most counts at the same time and the voting would turn into a 50% chance dice rolling. To solve this problem, before the final voting we search through the whole point cloud and collect one more label by voting through kNN ($k = 5$) for those with even number of votes.

Multi-FOV snapshots for speed and accuracy: So far the Multi-FOV snapshots have been participating in the network training. Yet another important role of the multi-FOV design is to enable faster segmentation and more precise snapshot sampling leveraging our adaptive sampling technique. The adaptive sampling works to choose one of the pre-listed FOVs according to the size of objects being sampled. This process is completed in three steps: variance estimation, FOV inquiry, and snapshot down-sampling.

Variance estimation refers to the procedure of measuring how spread out the associated points are in a snapshot sample: we take the mean from each of the X , Y , Z coordinates as an imaginary center point for one pre-sampled snapshot and compute the sampling variance from the center point. Note that we use the largest FOV from the list to pre-sample a snapshot and keep the corresponding variance, with the intention to adequately differentiate the variances by maximizing the sampled area. Next is to *inquire* the most appropriate FOV for each pre-sampled snapshot based on the variance and the sampling history. During the segmentation progress, variances from all pre-sampled snapshots are kept to periodically update a KMeans for clustering, where

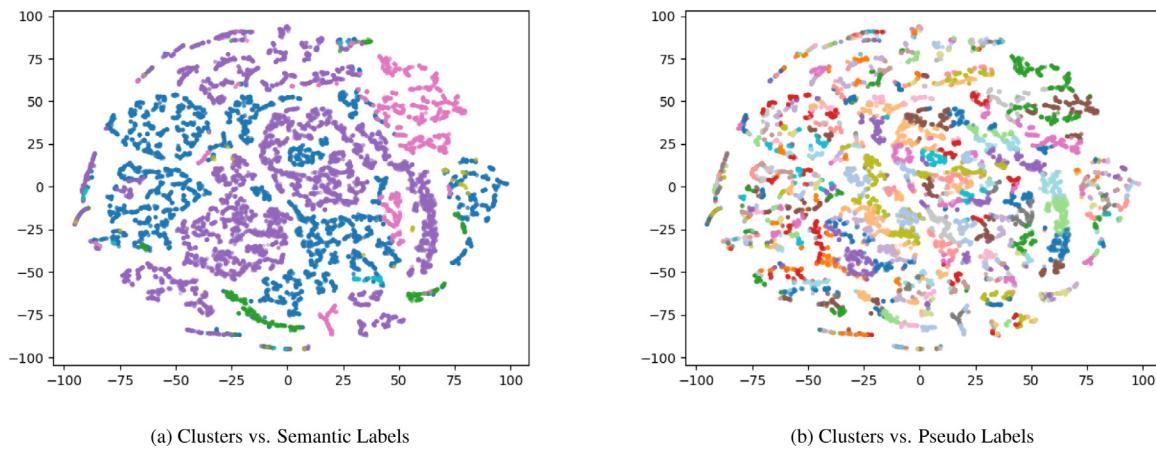


Fig. 6. Visualization of the feature embedding of the clustered snapshots from the Semantic3D. The clusters are colored by the semantic labels in 6 classes in (a), and by the pseudo labels in 300 classes in (b). The semantic labels in (a) are represented by the following colors: terrain — blue, vegetation — green, building — purple, hard scape — pink, artefacts — yellow, and cars — cyan. It is shown that samples sharing the same pseudo labels are likely to have the same semantic labels as well. Due to the large number of clusters in (b), it is not possible to distinguish all 300 colors globally, but it is possible to distinguish local clusters by their color contrasts. As a support of our claim, the average percentage of points agree with their semantic labels in each cluster among all 300 clusters is 91.67%. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

k equals to the number of the pre-listed FOVs. Before the KMeans is sufficiently trained, the model selects the smallest FOV for the next step of down-sampling. The reason for this is that the KMeans at this stage does not own enough history records to make a meaningful decision on which FOV to utilize, so the model proceeds with the most conservative option (a small FOV) to ensure a less risky and noisy down-sampling. Once the KMeans is well optimized, the model starts to inquire for cluster ID by sending in the pre-sample variance, and each cluster ID represents one of the FOVs that will be used for final sampling. In the end, the *down-sampling* follows the same principle of multi-FOV snapshots sampling, where a larger snapshot is first obtained with the assigned FOV using kNN, before points are randomly discarded to meet the network input size.

It is common to have different outdoor objects at a great range of scales. When sampling from a large regular surface, such as the ground or the building facade, the chances of including points from other objects is smaller. This observation motivates us to exploit the advantages of the adaptive sampling for a faster segmentation process. To allow this, both the down-sampled snapshot for label prediction and the pre-sampled snapshot for segmentation are kept. It can be seen from Fig. 4 that, the pre-sampled snapshots cover the same area as their corresponding down-sampled snapshots but include more points, except for the smallest FOV. Here, during the segmentation, when a prediction is acquired from the SVM, the model assigns the prediction to all points in the pre-sampled snapshot instead of the down-sampled one.

This serves as a solution to the low efficiency caused by the down-sampling operation. The discarded points are highly likely to come from the same class as their neighbor points obtained a prediction, but they will not be given a label until next time they are pre-sampled again and survived the down-sampling to go through the network, which is redundant as repetitive operations. Now that it is possible to expedite the segmentation of a large uniformed surface at a lower cost, we can leverage even smaller snapshot size to capture local structures at a higher precision. Overall, the adaptive sampling helps increasing the sampling precision for small objects while maintaining a fast segmentation speed for large surfaces.

4. Experiments and results

Extensive experiments are conducted to evaluate the effectiveness of our proposed approach for both self-supervised feature learning and weakly-supervised point cloud semantic segmentation. The implementation and experimental results are described in details in the following

sub-sections, including (1) datasets used in our experiments; (2) variations of snapshot capturing and their evaluations; (3) self-supervised feature learning by point cloud classification with fewer labeled data; and (4) evaluation of semantic segmentation by point-wise voting.

4.1. Datasets

The experiments are mostly conducted on the Semantic3D large scale point cloud classification benchmark (Hackel et al., 2017). This dataset consists of a variety of scenes across eight classes: man-made terrain, natural terrain, high vegetation, low vegetation, buildings, hard scape, scanning artefacts, and cars (Fig. 1). Considering the huge scale and high density of this dataset is beyond our computational capacity, we choose two scenes, named ‘untermaederbrunnen3’ and ‘bildstein3’ for our experiments, which consist of 27.9 million points and 7.9 million points, respectively. We use all eight labels to conduct snapshot classification for the evaluation of the self-supervised feature learning. To evaluate our weakly-supervised semantic segmentation on this dataset, we follow the experiment settings from the state of the art methods (Guinard and Landrieu, 2017; Xu et al., 2019): to combine the man made terrain and the natural terrain into a single class of terrain, and merge the high vegetation and low vegetation into vegetation.

Another dataset incorporated in the semantic segmentation experiments is the Oakland dataset (Munoz et al., 2009). In comparison to Semantic3D, the Oakland data has significantly lower but more uniform point densities as scanned from a moving platform. It has five semantic classes: scatter misc, default wire, utility pole, load bearing and facade. The class labels are drastically unbalanced, and out of the total 1.4 million points only 7000 points are default wire and 10 000 points are utility poles, while the load bearing has near 1 million points.

4.2. Snapshots capturing

Snapshot generation: In this experiment, for evaluating the self-supervised feature learning, we capture 8000 single-FOV snapshots from the Semantic3D dataset as training set and 800 samples for testing, with 1024 points in each sample. In an ideal setup where the dataset is perfectly balanced between classes, we would obtain close to 1000 training snapshots each class, but this can hardly be realized in real world scenarios (see Table 1, # of samples). The high resolution of the Semantic3D dataset (Hackel et al., 2017) poses a dilemma during the single-FOV snapshot sampling: a small sampling size is insufficient to capture details while a large sample brings burdens to

Table 1

Statistics of single-FOV snapshot sampling on the two scenes from the Semantic3D dataset (Scene 1: 'Untermaederbrunnen3'; Scene 2: 'Bildstein3'). The sampling is conducted 100 times, with 8000 samples per time. '-' indicates no snapshots being sampled.

| Statistics of Snapshot Sampling on Semantic3D dataset | | | | | | | | | | |
|---|------------|-------------------|-------------------|-------------------|------------------|-------------------|-------------------|--------------------|------------------|----------------|
| | Metrics | Man Made Terrain | Natural Terrain | High Vegetation | Low Vegetation | Buildings | Hard Scape | Scanning Artefacts | Cars | Total |
| Scene 1 | Purity (%) | 98.42 ±0.27 | 59.42 ±6.35 | 96.15 ±2.51 | 92.57 ±2.05 | 99.47 ±0.15 | 97.02 ±0.77 | 78.01 ±29.38 | 87.31 ±5.23 | 98.37 ±0.19 |
| | Samples | 2891.91 ±92.17 | 21.03 ±9.57 | 83.25 ±21.08 | 256.3 ±35.54 | 3847.68 ±92.13 | 838.01 ±61.28 | 10.34 ±6.80 | 51.48 ±16.30 | 8000 |
| Scene 2 | Purity (%) | 94.39 ±0.81 | 96.42 ±0.37 | 95.03 ±0.76 | 91.59 ±1.23 | 94.22 ±1.10 | 94.40 ±0.78 | - | 85.58 ±2.93 | 94.96 ±0.29 |
| | Samples | 1139.28 ±63.40 | 3264.49 ±81.99 | 1075.44 ±71.23 | 593.05 ±51.52 | 569.86 ±51.23 | 1240.92 ±76.18 | 0 | 116.96 ±28.20 | 8000 |

Table 2

Classification performance on the snapshots and labeled objects from the Semantic3D dataset, using the DGCNN, ObjectNet and SnapshotNet. Note that '-' means no samples from that class are obtained for testing.

| Method | Overall Accuracy (%) | Per-class Accuracy (%) | | | | | | | |
|--------------------------|----------------------|------------------------|-----------------|-----------------|----------------|-----------|------------|--------------------|------|
| | | Man Made Terrain | Natural Terrain | High Vegetation | Low Vegetation | Buildings | Hard Scape | Scanning Artefacts | Cars |
| 100% training data | | | | | | | | | |
| DGCNN | 86.5 | 98.41 | - | 100 | 85.71 | 84.38 | 47.14 | - | 20 |
| ObjectNet | 97.88 | 97 | 98 | 100 | 92 | 98 | 98 | 100 | 100 |
| SnapshotNet on objects | 98.63 | 100 | 99 | 100 | 92 | 100 | 98 | 100 | 100 |
| SnapshotNet on snapshots | 97.5 | 99.68 | - | 100 | 94.29 | 98.63 | 84.29 | - | 80 |
| 20% training data | | | | | | | | | |
| DGCNN | 84.13 | 99.68 | - | 0 | 85.71 | 79.45 | 55.71 | - | 0 |
| ObjectNet | 94.88 | 92 | 99 | 99 | 85 | 95 | 95 | 95 | 99 |
| SnapshotNet on objects | 95.63 | 98 | 95 | 100 | 89 | 91 | 95 | 100 | 97 |
| SnapshotNet on snapshots | 97.13 | 99.37 | - | 90 | 85.71 | 99.45 | 85.71 | - | 40 |
| 5% training data | | | | | | | | | |
| DGCNN | 72.88 | 54.33 | - | 0 | 62.86 | 89.86 | 92.86 | - | 0 |
| ObjectNet | 88.38 | 82 | 93 | 98 | 71 | 82 | 89 | 95 | 97 |
| SnapshotNet on objects | 90.13 | 90 | 87 | 100 | 84 | 86 | 87 | 92 | 95 |
| SnapshotNet on snapshots | 95.0 | 97.46 | - | 100 | 68.57 | 98.63 | 78.57 | - | 80 |

the computations. A compromise is made here, which takes a similar approach to the multi-FOV snapshots sampling, to take a pre-sampled snapshot with 10 times of the network input size (10240 points) and down-sample back to 1024 points per snapshot.

For the semantic segmentation, the same amount of multi-FOV snapshots are captured from both datasets at a smaller size of 512 points as the training data. The snapshot generation follows the multi-FOV sampling as described previously (Fig. 4), and we choose three FOVs for each snapshot. The original sampling rate is chosen as the base sampling rate, and the other two larger FOVs are respectively two times and ten times to the base sampling rate. Note that since each multi-FOV sample has three FOVs, the total training samples are 24 000.

Snapshot purity: Based on the proposed purity metric, we run the single-FOV snapshot capturing procedure on the Semantic3D at the base sampling rate 100 times for statistics. As shown in Table 1, the snapshot purity of a class is correlated with the numbers of sampled snapshots. From the point-wise perspective, when choosing an anchor point to find the nearest neighbors, each point in the scene has equal chance being selected as the anchor. However, class-wise speaking, when collecting points surrounding an anchor from a smaller class (i.e., a class with smaller number of points in the scene), the chance of including inter-class points is relatively higher than for a larger class, and this potentially leads to lower purity on smaller classes. Despite the fact that noises are much more likely to be included in smaller class samples, we can still see that the overall snapshot purity is above 90%. This result is in favor of our claim that, statistically each snapshot is highly capable of representing a small piece of one class from the whole point cloud. Nevertheless, we will also investigate if low purity classes can also be fairly treated.

Snapshots versus objects: To comparatively evaluate our self-supervised feature learning, we also apply the same single-FOV

sampling procedure on points grouped by the original semantic labels, instead of the whole scene of point cloud. As a result, these samples obtained exclusively from one class have 100% purity, and we refer to them as 'objects' as opposed to the 'snapshots'. Thus the Contrast-ClusterNet trained with the object-based approach is referred as ObjectNet, and is mainly for comparison purposes, even though it may have its own value if obtaining objects is a possibility.

4.3. Self-supervised feature learning

As we discussed in our previous work (Li. and Zhu., 2021), to verify the self-supervised feature learning of the SnapshotNet, we conduct experiments on both single-FOV snapshots and labeled objects derived from the scene - 'Untermaederbrunnen3' of the Semantic3D dataset. The evaluations are based on the classification accuracy on the testing samples of an SVM (with a linear kernel) trained on the extracted features of training samples. For the experiments, we train both the DGCNN and the SnapshotNet exclusively on the snapshot samples while keeping the ObjectNet trained on labeled objects. For the SnapshotNet, we also want to see if the features can be applied to object samples. Thus we use the trained model to extract features of both snapshots and objects separately to train a different classifier, and this is referred to as "SnapshotNet on snapshots" and "SnapshotNet on objects" in Table 2.

Learn with noises: Table 2 shows that, using 100% of the training data, the DGCNN has the lowest accuracy compared to the other methods. In comparison, the SnapshotNet tested on labeled objects has best performance on the total accuracy and all per-class accuracies except for low vegetation, on which is best performed by the SnapshotNet tested on snapshots. Both the SnapshotNet and the ObjectNet yield a total accuracy above 97%, and they are 10% higher than the DGCNN.

Table 3

Semantic segmentation results on Semantic3D. Three self-supervised methods are compared against the state of the art weakly-supervised methods. All 8000 labels are used in the training of the classifier.

| Method | Overall Accuracy (%) | Average F-score (%) | Per-class F-scores (%) | | | | | |
|--|----------------------|---------------------|------------------------|-------------|-------------|-------------|-------------|-------------|
| | | | Terrain | Vegetation | Building | Hardscape | Artefacts | Cars |
| Seg-aided (Guinard and Landrieu, 2017) | 83.3 | 82.3 | 98.1 | 67.0 | 98.8 | 91.5 | 51.3 | 82.3 |
| Pseudo-labeling (Xu et al., 2019) | 95.6 | 66.7 | 94.2 | 61.2 | 97.7 | 84.6 | 9.0 | 53.3 |
| Part contrasting | 96.9 | 80.2 | 97.4 | 85.5 | 98.4 | 95.1 | 32.3 | 72.5 |
| Scale contrasting | 92.1 | 74.5 | 90.6 | 75.7 | 97.5 | 87.6 | 33.4 | 62.1 |
| Multi-FOV contrasting | 97.6 | 90.2 | 98.2 | 85.6 | 98.8 | 96.6 | 74.8 | 87.0 |

This validates our claim that the proposed SnapshotNet is able to learn effective features from the raw point cloud complex scene in a self-supervised manner. It also shows that the noisy snapshots produce more powerful features by self-supervised learning than the fully-supervised DGCNN.

We see that the ObjectNet is able to achieve decent performance given that the training samples are derived from labeled data and each object sample has 100% purity. However, the snapshots used to train the SnapshotNet are often noisy. As presented in Table 1, some classes have low purity to start with, such as the Natural Terrain at a 57.47% purity, the Scanning Artefacts at 73.2% and the class of Cars at an 87.85% purity. While the snapshots from these classes are noisy, the SnapshotNet has shown high resistance over such noises in the data. (Note that there are no testing snapshots in natural terrain and scanning artefacts being sampled due to the small amount of data in the two classes) The SnapshotNet on objects performs better than the ObjectNet, which is trained on noiseless objects. It has also shown high accuracies over the aforementioned three noisy classes, which further confirms our hypothesis that powerful semantic features can be learned by predicting whether two segments are from the same snapshot and predicting the refined pseudo-labels for the snapshots, regardless of their semantic labels.

Classification with fewer labeled data: To verify the effectiveness of the proposed weakly supervised classification, we gradually reduce the amount of labels involved in the training of the SVM. The experiments are set up in the similar way by comparing the DGCNN, ObjectNet, and the SnapshotNet.

Table 2 also shows how the classification accuracy of different models vary when the percentages of the classifier training data reduce from 100% to 20% and to 5%. It can be seen that the SnapshotNet outperforms the other models on overall accuracy. The difference between the SnapshotNet (test on snapshots) and the DGCNN becomes more significant when the training data reduces, which spans from 11% to 22.12%. Comparatively, the end-to-end fully-supervised DGCNN is very sensitive to the amount of training data due to the data hungry problem, that it needs sufficient labeled data to learn representative features. The SnapshotNet however, does not rely on any labeled data for feature learning, and only need a small fraction of labeled data to train a classifier, which is a huge advantage over the fully-supervised model.

The snapshot-based SnapshotNet has also shown higher resistance on the reduction of the training data than the other training schemes: the accuracy only drops by 2.5%, to 95%, when the classifier is trained with merely 5% of the data, while the object based SnapshotNet suffers a drop of 8.5%, to 90.13%, and the ObjectNet shows a bigger decline of 9.5%, to 88.38%. We believe that the performance difference between the SnapshotNet and ObjectNet can be attributed to the use of noisy snapshots to make feature learning more robust. Compared to the high purity objects, the snapshots forces the model to distinguish whether two segments are from the same area despite they might contain points from different classes. In other words, we increase the difficulty of this pretext task by introducing noises and potentially leads to more representative features.

4.4. Semantic segmentation

Following the works of the seg-aided (Guinard and Landrieu, 2017), and pseudo-labeling (Xu et al., 2019), and for the purpose of comparison, we merge the natural terrain and man-made terrain into one class of terrain, and put together the high vegetation and low vegetation as vegetation. The experiments are carried out on these six classes cleaned from the Semantic3D dataset and the Oakland dataset, and they are organized as following:

1. Having verified that the single-FOV snapshots, despite being noisy, are able to produce meaningful features from part-contrasting, we move forward to utilize the multi-FOV snapshot to feed the other two pre-text tasks: scale contrasting, and the multi-FOV contrasting, and comparison results are given.
2. For the weakly-supervised classifier, we experiment on three different ratios of the labeled data: 100%, 20%, 10% and 5% of our total 8000 labeled training samples, which contains 8000, 800, and 400 labels respectively. These numbers of labels follow the distribution of the semantic labels in the dataset, therefore the larger classes might outnumber the small class on label numbers.
3. In addition to this, we include another test case with 30 labels per class, which is the same setup in the work of the seg-aided classification (Guinard and Landrieu, 2017) and pseudo-labeling (Xu et al., 2019). To mitigate the potential issues when using only a few labels, tests are conducted on the proposed cluster-based pseudo labeling for automatically adding more training samples.
4. To evaluate the robustness of our method over different scene point clouds, the testing results are produced from two experimental setups. One is to segment the scene point cloud from which the training samples are captured, while the other setup involves training the model on one scene but segmenting another one. In the second case, we gradually add up the fine-tuning samples from the to-be-segmented scene to find a sweet spot where minimal fine-tuning is required to achieve comparable results when performing cross-scene semantic segmentation.

Results on various contrastive learning: The results of the three contrasting approaches are listed in Table 3, where our method with the three approaches is compared with a state of the art method seg-aided classification (Guinard and Landrieu, 2017) and one comparable following study, the pseudo-labeling approach (Xu et al., 2019). By doing part contrasting, our model yields an overall accuracy (OA) at 96.9%, 1.3% higher than the pseudo-labeling method. The average F-score, however, is slightly lower than the Seg-aided classification at 80.2%. Looking at the per-class F-scores, the part contrasting produces comparable results with the seg-aided method on the classes of terrain and building, ours shows a prominent improvement over the vegetation and hardscape classes, which pushes up the per-class F-score by 18.5% and 3.6% respectively. However, the part contrasting performs noticeably worse than the seg-aided on small objects such as artefacts and cars. This seems to conform our conjecture that features learned by part contrasting are vulnerable when describing smaller items.

Table 4

Parameter studies on varying the numbers of labels involved in training. Experiments are performed on our method with the multi-FOV contrasting to compare against the state of the art methods on Semantic3D. The number of labels corresponds to 100%, 20%, 10%, and 5% of the total available labels. An additional test case using 30 labels/class is included to make a total 180 labels.

| Method | Overall Accuracy (%) | Average F-score (%) | Per-class F-scores (%) | | | | | |
|--|----------------------|---------------------|------------------------|-------------|-------------|-------------|-------------|-------------|
| | | | Terrain | Vegetation | Building | Hardscape | Artefacts | Cars |
| Seg-aided (Guinard and Landrieu, 2017) | 83.3 | 82.3 | 98.1 | 67.0 | 98.8 | 91.5 | 51.3 | 82.3 |
| Pseudo-labeling (Xu et al., 2019) | 95.6 | 66.7 | 94.2 | 61.2 | 97.7 | 84.6 | 9.0 | 53.3 |
| Ours 8000 labels | 97.6 | 90.2 | 98.2 | 85.6 | 98.8 | 96.6 | 74.8 | 87.0 |
| Ours 1600 labels | 96.9 | 84.9 | 97.3 | 84.0 | 98.9 | 95.1 | 57.5 | 76.9 |
| Ours 800 labels | 95.9 | 82.6 | 96.3 | 79.7 | 98.3 | 93.6 | 59.8 | 67.7 |
| Ours 400 labels | 94.2 | 73.3 | 94.6 | 76.0 | 97.9 | 92.6 | 31.4 | 47.8 |
| Ours 180 labels | 85.4 | 63.4 | 84.0 | 57.5 | 94.3 | 89.2 | 11.5 | 44.3 |

Table 5

Parameter studies on 3D scene segmentation with varying the numbers of labels involved in training on the Oakland outdoor dataset. Experiments are performed on our method with the multi-FOV contrasting to compare against the state of the art methods on the Oakland dataset. The number of labels corresponds to 100%, 10%, 5%, and 2.5% of the total available labels. An additional test case is included using a total of 150 labels randomly drawn from the 8000 training samples. Note that “–” suggests no training samples are available for a particular class.

| Method | Overall Accuracy (%) | Per-class F-scores (%) | | | | | | | | |
|--|----------------------|------------------------|------|---------|------|---------|------|------|---------|--------|
| | | Scatter | Misc | Default | Wire | Utility | Pole | Load | Bearing | Facade |
| Seg-aided (Guinard and Landrieu, 2017) | 96.6 | 93.7 | 46.5 | 8.7 | 99.5 | 93.9 | | | | |
| Pseudo-labeling (Xu et al., 2019) | 96.6 | 92.0 | 40.2 | 46.2 | 99.3 | 93.3 | | | | |
| Ours 8000 labels | 92.6 | 85.7 | 2.5 | – | 96.7 | 84.4 | | | | |
| Ours 800 labels | 92.2 | 84.6 | 0.2 | – | 96.7 | 83.2 | | | | |
| Ours 400 labels | 91.3 | 82.6 | 6.9 | – | 96.1 | 82.2 | | | | |
| Ours 200 labels | 89.4 | 78.6 | – | – | 95.7 | 73.8 | | | | |
| Ours 150 labels | 88.5 | 76.0 | – | – | 95.1 | 68.4 | | | | |

The scale contrasting is proposed to bring the features at a more abstract level into the play, thus to achieve better results on the small objects in the scene. There is a marginal increase on the F-score of the artefacts than the part contrasting, while at the cost of worsening on all other classes. Another small class of cars is also 10.4% worse than the part contrasting and 20.2% lower than the seg-aided method. So far the results have been suggesting that the scale contrasting might lack the capability of pushing hard for powerful versatile features, and the produced features are less descriptive at certain levels that are vital to distinguish larger objects.

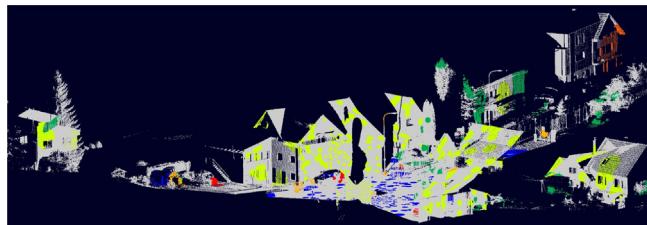
With this finding, the scale contrasting and part contrasting are combined for further experiments, seeking to strengthening the features from both perspectives. The collaborative effort of the part contrasting and scale contrasting urges the model to develop powerful features leveraging knowledge from both the object level information and higher level structural information. Our method equipped with the multi-FOV contrasting outperforms the pseudo-labeling method on the OA by 2% at 97.6%, and on F-score by 23.5%; it is 14.3% above the Seg-aided classification (Guinard and Landrieu, 2017) on OA, and a gain of the F-score is seen at 7.9%. There are some significant improvement over the small classes: the per-class F-score of the artefacts is 23.5% higher than the state of the art method at 74.8%. The class of cars has seen an increase from 82.3% to 87% using our method. The classes of vegetation and hardscape also experienced a very noticeable boost on their F-scores, and the terrain and building have a slight edge over the state of the art performance by the Seg-aided classification.

Segmentation with fewer labels on Semantic3D dataset: Table 4 shows the effect of reducing the numbers of available labels during the training of the SVM classifier on the Semantic3D. Our method with the multi-FOV contrasting is tested against the seg-aided classification and the pseudo-labeling, both using 30 labels per class. When the training data is reduced ten times to 800 labeled samples, similar performances can still be observed despite the drastic drop in the available training data. The OA of the 800 labels model is marginally greater than the pseudo-labeling approach by 0.3%, the average F-score is 15.9% higher than the pseudo-labeling approach and 0.3% over the best result from the seg-aided classification. The per-class F-score is still significantly greater than the state of the art method on the vegetation and artefacts, and is in the lead on the class of hardscape. The cars, however is

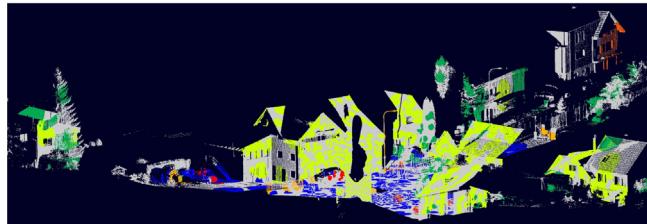
14.6% lower than the segmentation-aided classification at this level of available labels. A steady growth on this class can still be seen when the labels increase and we expect a surpassing over the seg-aided method when it is trained with more than 1600 labeled data. When further reducing the labeled data by half, the overall performance starts to drop. The per-class F-score on artefacts experiences a 28.4% decrease and for the cars it also falls by 19.9%. This suggests that a further cut down on labeled data usage by only 400 might come at a high price.

While our method has shown a superiority over the SOA methods when using as few as 800 labels, the SOA methods are only tested on 30 labels per class, making it 180 labels in total. To make a fair comparison with the two baseline methods, the training data is reduced to 30 labels per class. Apart from the artefacts holding a slight advantage over the pseudo-labeling method by 2.5%, the other classes are below the two baselines by varying degrees. Further investigating into the cause of this significant deficiency, two explanations are speculated. As commonly occurred in machine learning, underfitting of the classifier might be a big factor on the poor performance. A solution to this issue is to deploy larger amount training data. Another factor here is that, our self-supervised learning is trained on only 8000 samples from a particular statistical distribution in terms of their semantic labels, meaning that the learned features might not be as equally weighted, and the weights seem to collapse when training the classifier with uniformly distributed labels. Looking into more details, the recall rate of the artefacts is 86.0% while the precision is merely 6.2%. On the other hand, the recall of the terrain is 75.7% but the precision remains as high as 94.2%. These extremes are also seen on the vegetation and cars. This seems to meet our conjecture that the even labels bring bias into some of these classes.

Segmentation with fewer labels on Oakland dataset: Similar experiments are further conducted on the Oakland dataset. The results in Table 5 illustrates the OA and per class F-score against the two baselines. The Average F-score is not included in this test due to the lack of training samples for the disadvantaged classes such as default wire and utility pole. The Oakland dataset is extremely unbalanced between different classes and among the 8000 samples there are only 10 default wire snapshots and 6 utility pole snapshots. So when randomly selecting a smaller group of training samples from the pool, these classes



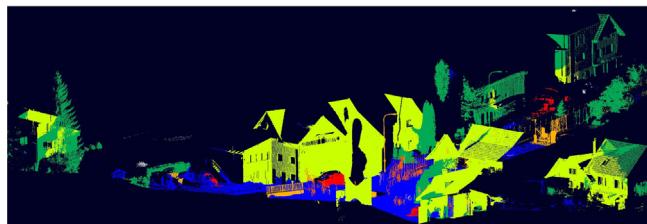
(a) Progress: 25%; Iteration: 1820; F-score: 89.5%



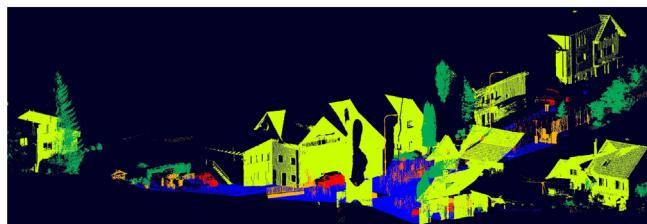
(b) Progress: 50%; Iteration: 4371; F-Score: 87.1%



(c) Progress: 75%; Iteration: 8869; F-Score: 87.6%



(d) Progress: 99.95%; Iteration: 84890; F-Score: 90.2%



(e) Ground truth

Fig. 7. Visualization of the progression of segmenting ‘Untermaederbrunnen3’. The model is governed with the pre-text of multi-FOV contrasting and the classifier is trained with 8000 labels. The bottom picture in 7(e) shows the ground truth to compare with. Colors correspond to semantic classes as the following: terrain is cyan, vegetation is green, buildings are yellow, hardscapes are orange, scanning artefacts are orange red, cars are red. The unlabeled points are in gray. We can see that snapshots gradually covers up the whole scene, during which previously falsely labeled points can be corrected by voting. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

become absent during the training, and these corresponding per-class F-scores are noted as “-”.

With the available results, we can still exploit some insights towards our model’s characteristics. The overall accuracy gradually decreases when using fewer labels, however the scatter misc and facade classes experience a steeper decline than the load bearing class does. This trend conforms the observation from the Semantic3D experiments, that

Table 6

Trade-off of the threshold section in the Cluster-based pseudo labeling on Semantic3D. 120 clusters are randomly selected from the total 300 clusters for pseudo labeling. Larger threshold value puts heavier constraints on the pseudo-labeling, leading to fewer labels but higher labeling accuracy.

| Threshold | Number of Pseudo Labels | Accuracy in Labeling |
|-----------|-------------------------|----------------------|
| 0.9 | 2631 | 98.1% |
| 0.8 | 4253 | 96.1% |
| 0.75 | 5078 | 93.7% |

the dominant classes have higher resistance on performance decay over the disadvantaged classes. Unlike the Semantic3D experiments, our model does not surpass the SOA methods with all 8000 samples. Referring to our previous work on this problem (Li. and Zhu., 2021), when training with single-FOV snapshots, the classification accuracy on the Oakland dataset using all training samples is 90.6% and it is 96.4% on Semantic3D. This confirms that the segmentation results are consistent with the classification accuracy, despite being benefited from the new multi-FOV training. So the performance difference between the two experiment setups reflects our model’s reaction towards these two datasets.

Here some interpretations are made to account for this performance discrepancy: Both experiments have shown the advantage of our model on feature learning with the use of snapshots. However, the way we generate these snapshots is fairly straightforward and fully statistical by randomly taking k nearest neighbors. This poses a problem when some classes are loosely distributed and outnumbered by other classes in the background, such as the default wire and utility pole in this scenario. These two classes are barely sampled as a snapshot by our definition of “snapshots”, and therefore exacerbates the inequality of the training samples than it appears in the entire scene.

Results on cluster-based pseudo labeling: To mitigate the impacts of the aforementioned two potential issues when using only a few labels, tests are conducted on the proposed cluster-based pseudo labeling for more training samples. Here on the Semantic3D, 120 clusters are randomly selected out of 300 and each cluster center is assigned with one label. In addition, a collection of 10 labeled sampled from each class is joined into the pseudo labeled data, consuming a total of 180 labels. Table 6 illustrates a trade-off effect between the number of pseudo labels and the labeling accuracy from the choice of threshold. A larger threshold causes a heavier constraint when selecting the samples near the cluster center, which leads to fewer samples to be pseudo-labeled but they are essentially much more likely to share the same semantic label with the cluster center. According to Table 7, comparing to our method using 180 labels without pseudo labeling, a boost of 7% on the overall accuracy and 11% on the F-score is seen, when thresholding at 0.8. Our cluster-based pseudo labeling method has outperformed the SOA deep learning method using pseudo-labeling on the F-score by 7.4%, particularly ours has an edge on segmenting the smaller objects such as artefacts and cars, where increases of 23.5% and 11.4% are gained respectively. When comparing to the seg-aided method, we also have an advantage on the overall accuracy by 8.9%.

It is worth explaining on the decision of the random selection when choosing 120 clusters for pseudo labeling. It was realized that for those larger clusters, despite being able to generate more pseudo labels, the included samples tend to be homogeneous. In other words, their features are less representative, so it is not always a good idea to go for larger clusters in the pursuit of more pseudo labels. For instance, two buildings looking completely different might have their samples far from each other in the feature space, and it would be more helpful to have each of their samples being chosen during the pseudo labeling. On the other hand, the very small clusters often contain few samples from the minority classes, and these clusters are an important source of acquiring distinct features from those minority classes. These factors pose a difficult decision on the cluster selection, to obtain as many pseudo labels as possible while maintain the diversity of the features. Thanks

Table 7

Parameter studies on the cluster-based pseudo labeling threshold selection on Semantic3D. Experiments are performed on our method with the multi-FOV contrasting to compare against the state of the art methods. Three threshold levels are tested on 120 random clusters for pseudo labeling samples. An test case using 30 labels/class without pseudo labeling is included for comparison.

| Method | Overall Accuracy (%) | Average F-score (%) | Per-class F-scores (%) | | | | | |
|--|----------------------|---------------------|------------------------|------------|----------|-----------|-----------|------|
| | | | Terrain | Vegetation | Building | Hardscape | Artefacts | Cars |
| Seg-aided (Guinard and Landrieu, 2017) | 83.3 | 82.3 | 98.1 | 67.0 | 98.8 | 91.5 | 51.3 | 82.3 |
| Pseudo-labeling (Xu et al., 2019) | 95.6 | 66.7 | 94.2 | 61.2 | 97.7 | 84.6 | 9.0 | 53.3 |
| 180 labels without pseudo labeling | 85.4 | 63.4 | 84.0 | 57.5 | 94.3 | 89.2 | 11.5 | 44.3 |
| 120 clusters t0.9 + 10 labels/class | 91.6 | 70.6 | 91.6 | 61.8 | 96.9 | 87.2 | 32.6 | 53.4 |
| 120 clusters t0.8 + 10 labels/class | 92.2 | 74.1 | 91.6 | 71.8 | 96.9 | 87.4 | 32.5 | 64.7 |
| 120 clusters t0.75 + 10 labels/class | 92.0 | 73.5 | 91.1 | 67.3 | 97.3 | 85.3 | 31.1 | 69.1 |

Table 8

Parameter studies on the cluster-based pseudo labeling threshold selection and segmentation results on the Oakland dataset. Four threshold levels are tested on 150 random clusters for pseudo labeling samples. An additional test case using 150 random labels without pseudo labeling is included for comparison.

| Method | Number of Pseudo Labels | Labeling Accuracy (%) | Overall Accuracy (%) | Per-class F-scores (%) | | | | | |
|--------------------|-------------------------|-----------------------|----------------------|------------------------|------|---------|------|-------------|-------------|
| | | | | Scatter | Misc | Default | Wire | Utility | |
| 150 random labels | 0 | 100 | 88.5 | 76.0 | – | – | – | 95.1 | 68.4 |
| 150 clusters t0.99 | 312 | 97.4 | 78.2 | 59.4 | – | – | – | 87.7 | 59.4 |
| 150 clusters t0.97 | 762 | 90.6 | 81.4 | 65.5 | – | – | – | 89.0 | 66.6 |
| 150 clusters t0.95 | 1479 | 89.3 | 78.7 | 58.5 | – | – | – | 88.4 | 64.4 |
| 150 clusters t0.90 | 3121 | 87.7 | 76.8 | 56.1 | – | – | – | 86.0 | 64.6 |

to the large number of clusters, we believe that random selection is a better way to evenly include both large and small clusters for pseudo labeling.

The results on the Oakland dataset provide a better viewpoint on the label number vs. labeling accuracy dilemma when the dataset is highly unbalanced. **Table 8** shows this trade-off and the leading segmentation performance using 150 labels. The baseline is when randomly choosing 150 labeled samples without pseudo labeling, and it is worth noting that since each sample has three FOVs, so there are a total of 450 training snapshots in this case. When loosing the threshold on the pseudo labeling, more pseudo labels are obtained at the cost of lower labeling accuracy. At 97%, the segmentation performance reaches the best among all four pseudo labeling tests, where we have 762 pseudo labels at the accuracy of 90.6%. By comparing this result with the baseline, it is seen that 1.7 times more training samples cannot makeup for the 9.4% loss on labeling accuracy, and it is not a problem on the Semantic3D where thousands of pseudo labels are obtained with high accuracy.

Segmentation across scenes: Having shown the advantages of proposed method on a single scene point cloud when the learning is governed by the multi-FOV contrasting. To verify the strength of our model on quickly adapting to other data, the following experiments are designed to test on cross-scene segmentation. The goal is to find out if this model is capable of producing decent segmentation performance by only fine-tuning the model on minimal amount of fine-tuning data. Taking the model obtained from previous experiments, which is trained on the scene ‘Untermaederbrunnen3’ from the Semantic3D. This model is fine-tuned with a series of number of samples from the scene ‘Bildstein3’, such as 1600, 800, 400, 0 samples, which take up to 20%, 10%, 5%, and 0% (no fine-tuning) of the total samples. Recall that the model has two networks working in sequence. The clusternet is trained with the pseudo-labels acquired from the features extracted using a well trained contrastnet. So to fine-tune our model with new data, the pseudo-labeling process is carried through first. This involves extracting features of the fine-tuning data using a pre-trained contrastnet and predicting new pseudo-labels for them with the converged KMeans from our previous experiments. Then the pre-trained clusternet is fine-tuned on the new data with their pseudo-labels before starting segmenting the new scene.

As demonstrated in **Table 9**, experiments are tested with different amount of fine-tune data from 1600 to 0 samples. It is seen that the OA and average F-scores are close to each other among the four fine-tuned

models. Compared to the other five classes, artefacts are prone to large fluctuations on F-score. To compare with the single-scene performance trained with 8000 labels (**Table 4**), a drop of 27% on the artefacts attracts most attention among all other results, which are considerably close to or even surpasses the single-scene results. Despite that the results have shown an edge of our method when adopting to new data even without any fine-tuning, there is no significant improvement by adding in more fine-tuning data. One interpretation of this particular behavior is again related to the statistical distribution of the classes. As discussed on the snapshot purity, the number of points in each class are greatly uneven, leading to a large disparity on the number of snapshots being captures in each class. This is particularly the case when it comes to the artefacts, where as shown in **Fig. 1** that no snapshots from this class are picked up during sampling, meaning that the fine-tuning was largely conducted on samples from bigger classes such as terrain or building. **Fig. 8** illustrates the visualization of the semantic labeling compared with the ground truth. It can be observed that the hardscapes away from the center are mislabeled as terrain, and this also happens on the lower part of the church’s tower. The cars are correctly located but the labeling precision is not as satisfactory because the surrounding terrains are mislabeled into cars. These observations again align with our hypothesis, that some classes are under-trained due to the lack of samples or the training is contaminated with low-quality samples.

5. Conclusion and discussion

In summary, we have proposed the SnapshotNet for self-supervised feature learning on the complex scene point cloud, including a new pre-text task that joins the part contrasting and the proposed scale contrasting for stronger features. We have also designed a weakly-supervised method for point cloud semantic segmentation by training with fewer coarse-grained labels. While reducing the labels involved in the downstream tasks, a cluster-based pseudo labeling technique is implemented to obtain more training data. The proposed methods are evaluated and verified on two real life complex scene datasets and the experimental results indicate that our method is capable of learning effective features from unlabeled scene point cloud data. Compared to the state of the art methods, our methods still show several advantages. This model is able to produce comparable results at a slightly higher cost on label collection. When the cluster-based pseudo labeling is enabled on representative snapshots, our model is capable of producing comparable results with the state-of-the-art methods using only 180

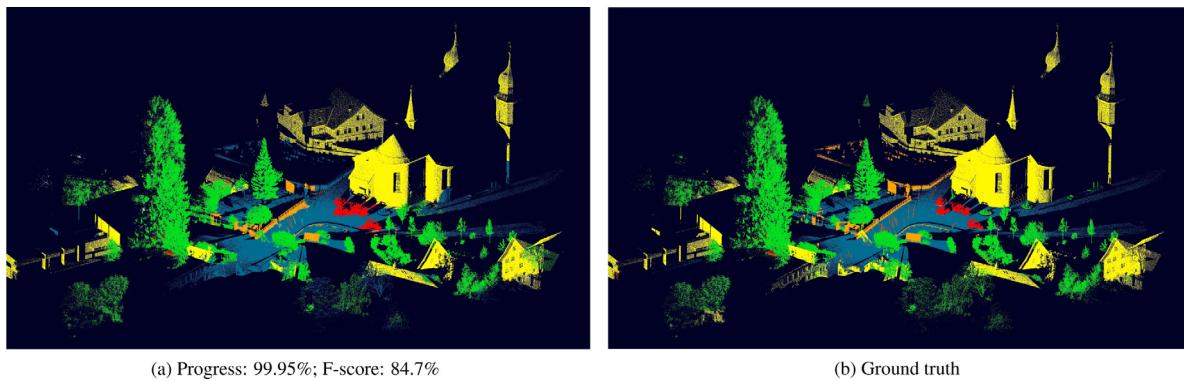


Fig. 8. Semantic labeling of the scene 'Bildstein3' using the model trained on 'Untermaederbrunnen3'. Colors correspond to semantic classes as the following: terrain is cyan, vegetation is green, buildings are yellow, hardscapes are orange, scanning artefacts are orange red, cars are red, and the unlabeled points are gray. It can be seen that our method is capable of recognizing the rough outline of the smaller items such as cars, but lacks precise semantic labeling. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 9

Cross scene segmentation performance of our method. The model is pre-trained on the scene 'Untermaederbrunnen3' and fine-tuned on different numbers of samples from 'Bildstein3'. The decreasing numbers of fine-tune data make up to 20%, 10%, 5%, and 0% of the total 8000 samples captured from 'Bildstein3'.

| Fine-tune data | Overall accuracy (%) | Average F-score (%) | Per-class F-scores (%) | | | | | |
|----------------|----------------------|---------------------|------------------------|------------|----------|-----------|-----------|------|
| | | | Terrain | Vegetation | Building | Hardscape | Artefacts | Cars |
| 1600 | 93.2 | 84.7 | 94.3 | 94.9 | 94.7 | 87.8 | 47.8 | 88.7 |
| 800 | 96.2 | 84.8 | 97.1 | 95.4 | 97.5 | 94.7 | 36.9 | 87.9 |
| 400 | 94.5 | 85.5 | 95.6 | 95.1 | 94.3 | 90.7 | 47.4 | 89.6 |
| 0 | 95.9 | 84.7 | 96.8 | 95.4 | 95.2 | 94.4 | 33.4 | 93.4 |

labels. As a deep learning model, our method does not rely on hand crafted features, and it has proven to be robust to be directly applied to cross-scene segmentation without or with a small dose of fine-tuning within the same dataset, saving the effort of training a model on every new scene.

There are also some weaknesses of the proposed method, particularly that the quality of snapshots capturing are greatly influenced by the statistical distribution of the semantic classes. We have tried to resolve this issue by designing the multi-FOV snapshots and have gained some significant improvement, but the performance on smaller items still needs further improvement. In the future, the snapshot capturing could be further investigated, such as utilizing the surface normal or other local geometrical information, to potentially improve the sampling quality and enhancing the local semantic labeling precision. The noises in the snapshot sampling could also be turned into certain advantages for a multi-level contrastive learning: noisy snapshots might contain parts from other objects such that not all points from both samples agree on each other when forming a positive pair. If the dissimilarity between two parts can be quantified and well measured, a contrastive pair can be then defined in one of the multiple levels, instead of choosing from only positive or negative.

CRediT authorship contribution statement

Xingye Li: Conceptualization, Methodology, Software, Data curation, Investigation, Writing – original draft. **Ling Zhang:** Software. **Zhigang Zhu:** Conceptualization, Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

The research is supported by the National Science Foundation, USA through Awards PFI #1827505 and SCC-Planning, USA #1737533, the Air Force Office of Scientific Research (AFOSR), USA via Award #FA9550-21-1-0082, and Bentley Systems, USA, Incorporated, through a CUNY-Bentley CRA 2017-2020. Additional support is provided by a CCNY CEN Course Innovation Grant from Moxie Foundation, USA, and the Intelligence Community Center of Academic Excellence (IC CAE), USA at Rutgers University (Awards #HHM402-19-1-0003 and #HHM402-18-1-0007).

Appendix A. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.cviu.2021.103339>.

References

- Achlioptas, P., Diamanti, O., Mitliagkas, I., Guibas, L., 2018. Learning representations and generative models for 3d point clouds. In: International Conference on Machine Learning. PMLR, pp. 40–49.
- Alonso, I., Riazuelo, L., Montesano, L., Murillo, A.C., 2020. 3D-MiniNet: Learning a 2D representation from point clouds for fast and efficient 3D LiDAR semantic segmentation. IEEE Robot. Autom. Lett. 5 (4), 5432–5439. <http://dx.doi.org/10.1109/lra.2020.3007440>.
- Apple Inc., 2020. Ipad pro. <https://www.apple.com/ipad-pro/specs/>.
- Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., Savarese, S., 2016. 3D semantic parsing of large-scale indoor spaces. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition. CVPR, pp. 1534–1543. <http://dx.doi.org/10.1109/CVPR.2016.170>.
- Bhargava, P., Kim, T., Poulton, C.V., Notaros, J., Yaacobi, A., Timurdogan, E., Baiocco, C., Fahrenkopf, N., Kruger, S., Ngai, T., Timalsina, Y., Watts, M.R., Stojanović, V., 2019. Fully integrated coherent LiDAR in 3D-integrated silicon photonics/65nm CMOS. In: 2019 Symposium on VLSI Circuits, pp. C262–C263.
- Boulch, A., Le Saux, B., Audebert, N., 2017. Unstructured point cloud semantic labeling using deep segmentation networks. In: 3DOR, Vol. 2, p. 7.

Caron, M., Bojanowski, P., Joulin, A., Douze, M., 2018. Deep clustering for unsupervised learning of visual features. In: Proceedings of the European Conference on Computer Vision, ECCV, pp. 132–149.

Chang, A.X., Funkhouser, T., Guibas, L., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F., 2015. ShapeNet: An information-Rich 3D model repository. Technical Report 1512.03012. arXiv preprint.

Chen, N., Liu, L., Cui, Z., Chen, R., Ceylan, D., Tu, C., Wang, W., 2020. Unsupervised learning of intrinsic structural representation points. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 9121–9130.

Doersch, C., Gupta, A., Efros, A.A., 2015. Unsupervised visual representation learning by context prediction. In: The IEEE International Conference on Computer Vision, ICCV.

Gong, J., 2013. Mobile lidar data collection and analysis for post-sandy disaster recovery. In: Computing in Civil Engineering (2013), pp. 677–684.

Gong, J., Zhou, H., Gordon, C., Jalayer, M., 2012. Mobile terrestrial laser scanning for highway inventory data collection. In: Computing in Civil Engineering (2012), pp. 545–552.

Guinard, S., Landrieu, L., 2017. Weakly supervised segmentation-aided classification of urban scenes from 3D lidar point clouds. *ISPRS - Int. Arch. Photogramm. Remote Sens. Spat. Inf. Sci.* 151–157.

Guo, Y., Wang, H., Hu, Q., Liu, H., Liu, L., Bennamoun, M., 2020. Deep learning for 3d point clouds: A survey. *IEEE Trans. Pattern Anal. Mach. Intell.*

Hackel, T., Savinov, N., Ladicky, L., Wegner, J.D., Schindler, K., Pollefeys, M., 2017. SEMANTIC3D.Net: A new large-scale point cloud classification benchmark. In: ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. IV-1-W1, pp. 91–98.

Hu, X., Gong, J., 2018. Advancing smart and resilient cities with big spatial disaster data: Challenges, progress, and opportunities. In: Data Analytics for Smart Cities. Auerbach Publications, pp. 53–90.

Huang, J., You, S., 2016. Point cloud labeling using 3d convolutional neural network. In: 2016 23rd International Conference on Pattern Recognition. ICPR, IEEE, pp. 2670–2675.

Jing, L., Yang, X., Liu, J., Tian, Y., 2018. Self-supervised spatiotemporal feature learning via video rotation prediction. arXiv preprint arXiv:1811.11387.

Lawin, F.J., Danielljan, M., Tostberg, P., Bhat, G., Khan, F.S., Felsberg, M., 2017. Deep projective 3D semantic segmentation. In: International Conference on Computer Analysis of Images and Patterns. Springer, pp. 95–107.

Li, X., Zhu, Z., 2021. A snapshot-based approach for self-supervised feature learning and weakly-supervised classification on point cloud data. In: Proceedings of the 16th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications Vol. 5. VISAPP: VISAPP, SciTePress, INSTICC, pp. 399–408. <http://dx.doi.org/10.5220/0010230103990408>.

Munoz, D., Bagnell, J.A.D., Vandapel, N., Hebert, M., 2009. Contextual classification with functional max-margin Markov networks. In: Proceedings of IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR.

Noroozi, M., Favaro, P., 2016. Unsupervised learning of visual representations by solving jigsaw puzzles. *Lecture Notes in Comput. Sci.* 69–84. http://dx.doi.org/10.1007/978-3-319-46466-4_5.

Qi, C.R., Su, H., Mo, K., Guibas, L.J., 2017a. Pointnet: Deep learning on point sets for 3d classification and segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 652–660.

Qi, C.R., Yi, L., Su, H., Guibas, L.J., 2017b. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Advances in Neural Information Processing Systems. pp. 5099–5108.

Rethage, D., Wald, J., Sturm, J., Navab, N., Tombari, F., 2018. Fully-convolutional point networks for large-scale point clouds. In: Proceedings of the European Conference on Computer Vision, ECCV, pp. 596–611.

Sauder, J., Sievers, B., 2019. Context prediction for unsupervised deep learning on point clouds. arXiv preprint arXiv:1901.08396.

Tchapmi, L., Choy, C., Armeni, I., Gwak, J., Savarese, S., 2017. Segcloud: Semantic segmentation of 3d point clouds. In: 2017 International Conference on 3D Vision, 3DV, IEEE, pp. 537–547.

Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M., 2019. Dynamic graph cnn for learning on point clouds. *ACM Trans. Graph.* 38 (5), 1–12.

Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J., 2015. 3d shapenets: A deep representation for volumetric shapes. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 1912–1920.

Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O., 2020a. Pointcontrast: Unsupervised pre-training for 3D point cloud understanding. In: European Conference on Computer Vision, Springer, pp. 574–591.

Xie, Y., Tian, J., Zhu, X.X., 2020b. Linking points with labels in 3D: A review of point cloud semantic segmentation. *IEEE Geosci. Remote Sens. Mag.* <http://dx.doi.org/10.1109/MGRS.2019.2937630>.

Xu, X., Lee, G.H., 2020. Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 13706–13715.

Xu, K., Yao, Y., Murasaki, K., Ando, S., Sagata, A., 2019. Semantic segmentation of sparsely annotated 3D point clouds by pseudo-labelling. In: 2019 International Conference on 3D Vision, 3DV, pp. 463–471.

Yang, Y., Feng, C., Shen, Y., Tian, D., 2018. Foldingnet: Point cloud auto-encoder via deep grid deformation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 206–215.

Zhang, L., Zhu, Z., 2019. Unsupervised feature learning for point cloud understanding by contrasting and clustering using graph convolutional neural networks. In: 2019 International Conference on 3D Vision, 3DV, IEEE, pp. 395–404.