

Few-shot 3D Point Cloud Semantic Segmentation

Na Zhao Tat-Seng Chua Gim Hee Lee

Department of Computer Science, National University of Singapore

{nazhao, chuats, gimhee.lee}@comp.nus.edu.sg

Abstract

Many existing approaches for 3D point cloud semantic segmentation are fully supervised. These fully supervised approaches heavily rely on large amounts of labeled training data that are difficult to obtain and cannot segment new classes after training. To mitigate these limitations, we propose a novel attention-aware multi-prototype transductive few-shot point cloud semantic segmentation method to segment new classes given a few labeled examples. Specifically, each class is represented by multiple prototypes to model the complex data distribution of labeled points. Subsequently, we employ a transductive label propagation method to exploit the affinities between labeled multi-prototypes and unlabeled points, and among the unlabeled points. Furthermore, we design an attention-aware multi-level feature learning network to learn the discriminative features that capture the geometric dependencies and semantic correlations between points. Our proposed method shows significant and consistent improvements compared to baselines in different few-shot point cloud semantic segmentation settings (i.e. 2/3-way 1/5-shot) on two benchmark datasets. Our code is available at <https://github.com/Na-Z/attMPTI>.

1. Introduction

Point cloud semantic segmentation is a fundamental computer vision problem, which aims to estimate the category of each point in the 3D point cloud representation of a scene. The outcome of 3D semantic segmentation can benefit various real-world applications, including autonomous driving, robotics, and augmented/virtual reality. However, point cloud semantic segmentation is a challenging task due to the unstructured and unordered characteristics of point clouds. Recently, a number of fully supervised 3D semantic segmentation approaches [7, 8, 10, 11, 17, 24, 28, 31] have been proposed and have achieved promising performance on several benchmark datasets [1, 2]. Nonetheless, their success relies heavily on the availability of large amounts of labeled training data that are time-consuming and expensive to collect. Moreover, these approaches follow the *closed set*

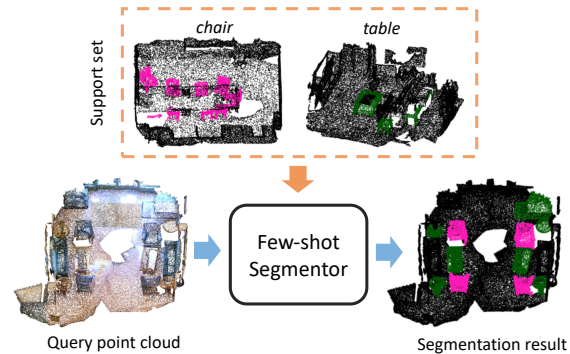


Figure 1. Few-shot point cloud semantic segmentation task is to learn a segmentor that segments the query point cloud in terms of new classes with learned knowledge from the support examples. This figure illustrates an example with 2-way 1-shot setting.

assumption which states that the training and testing data are drawn from the same label space. However, the closed set assumption is not strictly adhered to the dynamic real world, where new classes can easily occur after training. As a result, these fully supervised approaches suffer from poor generalization to new classes with only few examples.

Although several existing works used self-[26], weakly-[6, 27] and semi-supervised [14] learning to mitigate the data hungry bottleneck in fully supervised 3D semantic segmentation, these approaches are still under the closed set assumption, where the generalization ability to new classes is overlooked. The increasingly popular few-shot learning is a promising direction that allows the model to generalize to new classes with only a few examples. In few-shot point cloud segmentation, our goal is to train a model to segment new classes given a few labeled point clouds, as illustrated in Figure 1. We adopt the commonly used meta-learning strategy, *i.e.* episodic training [22], that learns over a distribution of similar few-shot tasks instead of only one target segmentation task. Each few-shot task consists of a few labeled samples (*support set*) and unlabeled samples (*query set*), and the model segments the query with learned knowledge from the support. Due to the consistency between the training few-shot task and the testing task, the model is endowed with better generalization ability that makes it less susceptible

to overfitting to rare support samples. Despite the benefit of episodic training, few-shot point cloud segmentation still faces two major challenges on how to: 1) distill discriminative knowledge from scarce support that can represent the distributions of novel classes; and 2) leverage this knowledge to effectively perform segmentation.

In this paper, we propose a novel attention-aware multi-prototype transductive inference method for few-shot point cloud semantic segmentation. Specifically, our approach is able to model the complex distributions of the points within the point clouds of the support set, and perform the segmentation via transductive inference with the discriminative features extracted under the few-shot constraint. We are motivated by the prototypical network [21], which represents each class with a single prototype obtained from averaging the embeddings of labeled samples in the support. We postulate that such uni-modal distribution assumption can be violated in point cloud segmentation due to the complex data distribution of points. In particular, the geometric structures of the points can vary largely within the same semantic class. Consequently, we propose to represent each class with multiple prototypes to better capture the complex distribution. Furthermore, it is important to learn discriminative features for the few-shot 3D point cloud semantic segmentation setting. To this end, we meticulously design an attention-aware multi-level feature learning network to learn the point-wise features by capturing the geometric dependencies and semantic correlations between the points. Subsequently, we perform the segmentation step in a transductive manner with the multiple prototypes in the learned feature space. In contrast to the conventional prototypical network [21] that matches unlabeled instances with the class prototypes by computing their Euclidean distances, our transductive inference not only considers the relationships between the unlabeled query points and the multi-prototypes, but also exploits the relationships among the unlabeled query points.

The main contributions of this work are: 1) We are the first to study the promising few-shot 3D point cloud semantic segmentation task, which allows a model to segment new classes given a few or even one example(s). 2) We propose a novel attention-aware multi-prototype transductive inference method. Our designs of the attention-aware multi-level feature learning, and the affinity exploitation between multi-prototypes and unlabeled query points enable our model to obtain highly discriminative features and accomplish more precise segmentation in the few-shot scenario. 3) We conduct comprehensive experiments on the S3DIS and ScanNet datasets to demonstrate the superior performance of the proposed approach over baselines in different (*i.e.* 2-/3-way 1-/5-shot) few-shot point cloud segmentation settings. Specifically, our method improves over the fine-tuning baseline in the challenging 3-way 1-shot setting by 52% and 53% on the S3DIS and ScanNet dataset, respectively.

2. Related Work

3D Semantic Segmentation. Many deep learning based approaches [7, 8, 10, 11, 17, 24, 28, 31] are proposed to tackle 3D semantic segmentation using full supervisions, *i.e.* point-wise ground truths. PointNet [17] is the first work that designs an end-to-end deep neural network to segment raw point clouds instead of their transformed representations, *e.g.* voxel grids and multi-view images. Despite its simplicity and efficiency, PointNet overlooks the important local information embedded in the neighboring points. DGCNN [24] addresses this issue by designing the EdgeConv module that can capture local structures. In our work, we make use of DGCNN as the backbone of our feature extractor to extract local geometric features and semantic features. Although these fully supervised approaches achieved promising segmentation performance, their requirement for large amounts of training data precludes their use in many real-world scenarios where training data is costly or hard to acquire. Moreover, these approaches can only segment a set of pre-defined classes that are seen during training. To alleviate these limitations, we explore the direction of few-shot learning for 3D semantic segmentation. This enables the model to segment new classes by seeing just a few labeled samples.

Few-shot Learning. The goal of few-shot learning is to develop a classifier that is able to generalize to new classes with very few examples (*e.g.* one example for the one-shot case). To address this challenging few-shot learning, several meta-learning approaches [4, 5, 15, 18, 19, 21, 22] have proposed to learn transferable knowledge from a collection of learning tasks and made significant progress. In particular, metric-based method [5, 21, 22] is notable because of its effectiveness in directly inferring labels for unseen classes during inference. The key idea in metric-based method is to learn a good metric function which is able to produce a similarity embedding space representing the relationship between labeled and unlabeled samples. *Matching network* [22] and *Prototypical Network* [21] are two representative metric-based methods. Both methods utilize deep neural network to map the support and query sets into an embedding space, and then apply a non-parametric method to predict classes for the query based on the support. Specifically, matching network leverages the weighted nearest neighbor method that represents a class by all its support samples, while prototypical network leverages the prototypical method that represents a class by the mean of its support samples. These two non-parametric methods become two extreme ends of the spectrum of complicated to simple data distribution modeling when applied to few-shot point cloud semantic segmentation. This is because the support samples for one class in the point cloud counterpart can contain a large number of points. In this paper, we represent each class in point clouds somewhere between the two extremes with multiple prototypes and perform segmentation in a transductive manner.

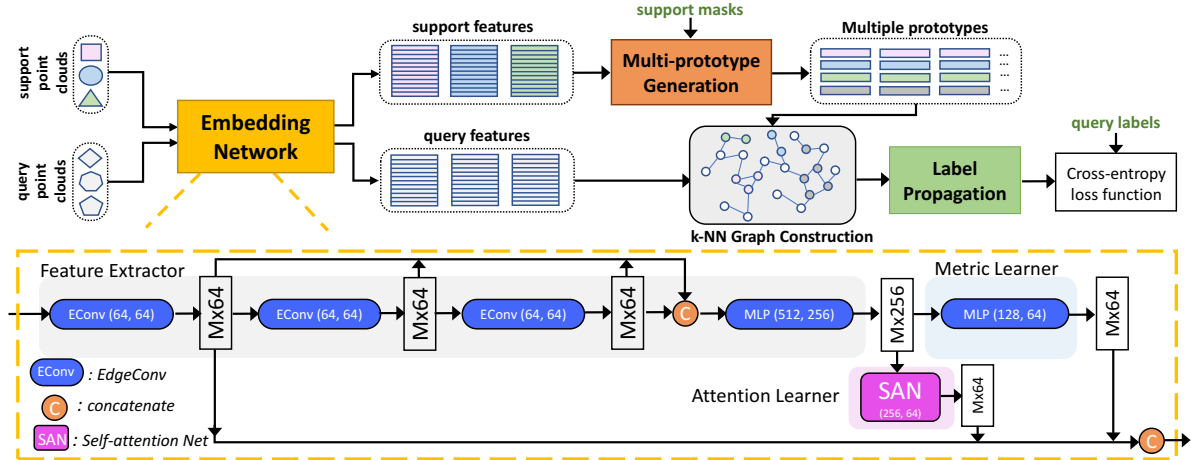


Figure 2. The architecture of our proposed method. This figure illustrates a 3-way 1-shot setting.

Few-shot Image Segmentation. All approaches mentioned previously focused on the few-shot image classification task. Only recently, several works [3, 12, 16, 20, 23, 29, 30] started to study few-shot learning on image segmentation by extending these meta-learning techniques to pixel levels. Most existing approaches [3, 16, 23, 30] leverage on metric-based techniques to solve a one-to-many matching problem between the support and query branch, where the support sample(s) of each class is represented as one global vector. On the contrary, Zhang *et al.* [29] considers the problem as many-to-many matching, where the support branch is represented as a graph with each element in the feature map of the support sample(s) as a node. However, these few-shot image segmentation approaches learn image features by using convolution neural network (CNN) based architectures, which are not applicable to point cloud data due to the irregular structures of point clouds. Moreover, the properties of a good embedding space are different for point clouds (*c.f.* Section 3.2.1) and images. In view of the differences, we design an attention-aware multi-level feature learning network and propose a novel attention-aware multi-prototype transductive inference method for the task of few-shot 3D point cloud semantic segmentation.

3. Our Methodology

3.1. Problem Definition

We align the training and testing of few-shot point cloud semantic segmentation with the episodic paradigm [22] that is commonly used in few-shot learning. Specifically, we train our model on a group of few-shot tasks sampled from a data set with respect to a training class set C_{train} , and then we test the trained model by evaluating it on another group of tasks sampled from a different data set with respect to new classes C_{test} , where $C_{test} \cap C_{train} = \emptyset$. Each few-shot task, a.k.a. an *episode*, is instantiated as an N -

way K -shot point cloud semantic segmentation task. In each N -way K -shot episode, we are given a *support* set, denoted as $S = \{(\mathbf{P}_s^{1,k}, \mathbf{M}^{1,k})_{k=1}^K, \dots, (\mathbf{P}_s^{N,k}, \mathbf{M}^{N,k})_{k=1}^K\}$, with K labeled pairs of support point cloud $\mathbf{P}_s^{n,k}$ and its corresponding binary mask $\mathbf{M}^{n,k}$ for each of the N unique classes. Each point cloud $\mathbf{P} \in \mathbb{R}^{M \times (3+f_0)}$ contains M points associated with the coordinate information $\in \mathbb{R}^3$ and an additional feature $\in \mathbb{R}^{f_0}$, *e.g.* color. We are also given a *query* set, denoted as $Q = \{(\mathbf{P}_q^i, \mathbf{L}^i)\}_{i=1}^T$, which contains T pairs of query point cloud \mathbf{P}_q^i and its corresponding label $\mathbf{L}^i \in \mathbb{R}^{M \times 1}$. Note that the ground-truth label \mathbf{L} is only available during training. The goal of N -way K -shot point cloud semantic segmentation is to learn a model $f_\Phi(\mathbf{P}_q, S)$ that predicts the label distribution $\mathbf{H} \in \mathbb{R}^{M \times (N+1)}$ for any *query* point cloud \mathbf{P}_q based on S . Formally, our training objective is to find the optimal parameters Φ^* of $f_\Phi(\mathbf{P}_q, S)$ by computing:

$$\Phi^* = \underset{\Phi}{\operatorname{argmin}} \mathbb{E}_{(S,Q) \sim \mathcal{T}_{train}} \left[\sum_{(\mathbf{P}_q^i, \mathbf{L}^i) \in Q} J(\mathbf{L}^i, f_\Phi(\mathbf{P}_q^i, S)) \right], \quad (1)$$

where \mathcal{T}_{train} denotes the training set containing all the episodes sampled from C_{train} , and $J(\cdot)$ is the loss function that will be defined in Section 3.2.4.

3.2. Attention-aware Multi-prototype Transductive Inference Method

Figure 2 illustrates our attention-aware multi-prototype transductive inference framework. It consists of five components: 1) the **embedding network** that learns the discriminative features for the support and query point clouds; 2) the **multi-prototype generation** that produces multiple prototypes for each of the $N + 1$ classes (N semantic classes and one background class); 3) the **k -NN graph construction** that encodes both the cross-set (support-query) and intra-set (support-support, query-query) relationships within the em-

bedding space; 4) the **label propagation** that diffuses labels through the whole graph along high density areas formed by the unlabeled query points; and 5) the **cross-entropy loss function** that computes the loss between the predicted labels and ground-truth labels of all the query points.

3.2.1 Embedding Network

The embedding network is the most important part of our model since both multi-prototype generation and k -NN graph construction are dependent on the learned embedding space. We expect this space to possess three properties: it can 1) encode the geometric structures of points based on local context; 2) encode the semantic information of points and their semantic correlation based on global context; and 3) quickly adapt to different few-shot tasks. To this end, we design an attention-aware multi-level feature learning network that incorporates three levels of features: local geometric features, global semantic features, and metric-adaptive features. Specifically, our embedding network is composed of three modules: *feature extractor*, *attention learner*, and *metric learner*. We adopt DGCNN [24], a dynamic graph CNN architecture, as the backbone of our feature extractor to respectively produce local geometric features (outputs of the first EdgeConv layer) and semantic features (outputs of the feature extractor). To further explore semantic correlation between points in the global context, we apply a self-attention network (SAN) on the generated semantic features. SAN allows the point-wise feature to aggregate the global contextual information of the corresponding point cloud in a flexible and adaptive manner. The architecture of SAN is illustrated in Figure 3. In addition, we introduce the metric learner, *i.e.* a stack of multi-layer perceptrons (MLP) layers to enable faster adaptability of the embedding space to different few-shot tasks since the feature extractor is updated with a slower learning rate (*c.f.* the training details in Section 4.2). The metric learner maps all point-wise features of support and query sets into a manifold space, where common distance functions (*e.g.* euclidean distance or cosine distance) can be directly used to measure proximity between points. Finally, we concatenate the three levels of learned features together as the output of the embedding network.

3.2.2 Multi-prototype Generation

For each of the $N + 1$ classes in the support set, we generate n^1 prototypes to model the complex data distribution according to the few labeled samples in the episode. We cast the generation procedure as a clustering problem. While there can be different ways to cluster support points into multiple prototypes, we employ a simple strategy: sampling seed

¹Although we can vary n for different classes, we keep it uniform for simplicity.

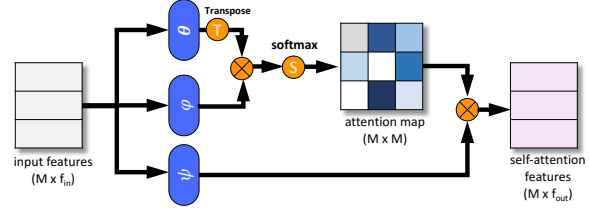


Figure 3. The architecture of Self Attention Network (SAN). θ , φ , and ψ are linear embedding functions with trainable parameters.

points and point-to-seed assignment based on the learned embedding space. Specifically, we sample a subset of n seed points from a set of support points in one class using the farthest point sampling based on the embedding space. Intuitively, the farthest points in this space can inherently represent different perspectives of one class if the embedding space is learned well. Let $\{\mathbf{s}_i^c\}_{i=1}^n$ and $\{\mathbf{f}_i^c\}_{i=1}^{m^c}$, where $\{\mathbf{s}_i^c\}_{i=1}^n \subset \{\mathbf{f}_i^c\}_{i=1}^{m^c}$, denote the sampled seeds and all the m^c support points belonging to the class c , respectively. We compute the point-to-seed distance and take the index of the closest seed as the assignment of a point. Formally, the multi-prototypes μ^c of class c is given by:

$$\mu^c = \left\{ \mu_1^c, \dots, \mu_n^c \mid \mu_i^c = \frac{1}{|\mathcal{I}_i^{c*}|} \sum_{\mathbf{f}_j^c \in \mathcal{I}_i^{c*}} \mathbf{f}_j^c \right\} \quad (2)$$

$$\text{s.t. } \mathcal{I}^{c*} = \underset{\mathcal{I}^c}{\operatorname{argmin}} \sum_{i=1}^n \sum_{\mathbf{f}_j^c \in \mathcal{I}_i^c} \|\mathbf{f}_j^c - \mathbf{s}_i^c\|_2,$$

where $\{\mathbf{f}_i^c\}_{i=1}^{m^c}$ is partition into n sets $\mathcal{I}^{c*} = \{\mathcal{I}_1^{c*}, \dots, \mathcal{I}_n^{c*}\}$ such that $\mathbf{f}_j^c \in \mathcal{I}_i^{c*}$ is assigned to \mathbf{s}_i^c .

3.2.3 Transductive Inference

In addition to the similarity relations between each unlabeled query point and the labeled multi-prototypes, we also consider the similarity relations between pairs of unlabeled query points to exploit the “smoothness” constraints² among neighboring query points in our few-shot point cloud semantic segmentation task. To this end, we leverage on transductive inference to reason cross-set and intra-set relationships based on the embedding space. Concretely, we propose the use of transductive label propagation to construct a graph on the labeled multi-prototypes and the unlabeled query points, and then propagate the labels in the graph with random walk.

k -NN graph construction. To mitigate the large number of query points, we construct a k Nearest Neighbor (NN) graph instead of a fully-connected graph for computational efficiency. Specifically, we take both the $n \times (N + 1)$ multi-prototypes and $T \times M$ query points as nodes of a graph with size $V = n \times (N + 1) + T \times M$. We construct a sparse affinity matrix, denoted as $\mathbf{A} \in \mathbb{R}^{V \times V}$, by computing

²Nearby points are most likely to have the same label.

the Gaussian similarity between each node and its k nearest neighbors in the embedding space:

$$\mathbf{A}_{ij} = \exp\left(-\frac{\|\mathbf{v}_i - \mathbf{v}_j\|_2^2}{2\sigma^2}\right), \text{ for } \mathbf{v}_j \in \mathcal{N}_k(\mathbf{v}_i), \quad (3)$$

where \mathbf{v}_i represents the node feature and σ^2 is the variance of the distance between two nodes. We follow [9] to let $\mathbf{W} = \mathbf{A} + \mathbf{A}^\top$, this assures the adjacency matrix is non-negative and symmetric. Subsequently, we symmetrically normalize \mathbf{W} to yield $\mathbf{S} = \mathbf{D}^{-1/2}\mathbf{W}\mathbf{D}^{-1/2}$, where \mathbf{D} is the diagonal degree matrix with its diagonal value to be the sum of the corresponding row of \mathbf{W} . Furthermore, we define a label matrix $\mathbf{Y} \in \mathbb{R}^{V \times (N+1)}$, where the rows corresponding to labeled prototypes are one-hot ground truth labels and the rest are zero.

Label propagation. Given \mathbf{S} and \mathbf{Y} , label propagation iteratively diffuses labels through the graph according to:

$$\mathbf{Z}_{t+1} = \alpha\mathbf{S}\mathbf{Z}_t + (1 - \alpha)\mathbf{Y}. \quad (4)$$

$\mathbf{Z}_t \in \mathbb{R}^{V \times (N+1)}$ represents the predicted label distributions at iteration t , and $\alpha \in (0, 1)$ is a parameter that controls the relative probability of the information from its adjacency nodes or its initial labels. In [32], Zhou *et al.* show that sequence $\{\mathbf{Z}_t\}$ converges to a closed-form solution:

$$\mathbf{Z}^* = (\mathbf{I} - \alpha\mathbf{S})^{-1}\mathbf{Y}. \quad (5)$$

We adopt the closed-form solution to directly compute the predictions \mathbf{Z}^* of label propagation.

3.2.4 Loss Function

Once \mathbf{Z}^* is obtained, we first take the predictions corresponding to the T query point clouds denoted as $\{\mathbf{z}^i\}_{i=1}^T$, where $\mathbf{z}^i \in \mathbb{R}^{M \times (N+1)}$ represents the predictions of the point cloud \mathbf{P}_q^i . The prediction of each point in \mathbf{z}^i is then normalized into a probability distribution using the softmax function:

$$\mathbf{H}_{m,n}^i = \frac{\exp(\mathbf{z}_{m,n}^i)}{\sum_{j=1}^{N+1} \exp(\mathbf{z}_{m,j}^i)}, \quad (6)$$

Finally, we compute the *cross-entropy loss* between $\{\mathbf{H}^i\}_{i=1}^T$ and the ground truth labels $\{\mathbf{L}^i\}_{i=1}^T$ as:

$$J_\Phi = -\frac{1}{T} \frac{1}{M} \sum_{i=1}^T \sum_{m=1}^M \sum_{n=1}^{N+1} \mathbb{1}[\mathbf{L}_m^i = n] \log(\mathbf{H}_{m,n}^i), \quad (7)$$

where Φ is the set of parameters of our model $f_\Phi(\mathbf{P}_q, S)$. More precisely, $f_\Phi(\mathbf{P}_q, S) = h(g_\Phi(\mathbf{P}_q, \mathbf{P}_s), \mathbf{M})$ is a composite function of the embedding network $g_\Phi(\cdot)$, and the multi-prototypes generation and transductive inference operations $h(\cdot)$. It becomes apparent that the minimization of J over the parameters Φ is governed by the affinity properties among the labeled multi-prototypes and unlabeled query points since the gradients have to flow through the parameter-less $h(\cdot)$ into the embedding network $g_\Phi(\cdot)$.

4. Experiments

We conduct experiments to evaluate the effectiveness of our method on two benchmark datasets. To the best of our knowledge, there is no prior study of few-shot point cloud semantic segmentation. Thus, we design the setup of the dataset, implementation details, and baselines for evaluation.

4.1. Datasets and Setup

Datasets. We evaluate on two datasets: 1) **S3DIS** [1] consists of 272 point clouds of rooms with various styles (*e.g.* lobby, hallway, office, pantry). The annotation of the point clouds corresponds to 12 semantic classes plus one for the clutter. 2) **ScanNet** [2] consists of 1,513 point clouds of scans from 707 unique indoor scenes. The annotation of the point clouds corresponds to 20 semantic classes plus one for the unannotated space.

Setup. To customize the dataset to the few-shot learning setting, we evenly split the semantic classes in each dataset into two non-overlapping subsets based on the alphabetical order of the class names. The splitting details are listed in Table 1 of the supplementary material. For each dataset, we perform cross-validation on the corresponding two subsets by selecting one split as the test class set C_{test} , while taking the remaining split as the training class set C_{train} .

Since the number of points in the original rooms is large, we follow the data pre-processing strategy used in [17, 24] to divide the rooms into blocks using a non-overlapping sliding window of $1\text{m} \times 1\text{m}$ on the xy plane. As a result, S3DIS and ScanNet give 7,547 and 36,350 blocks, respectively. From each block, $M = 2,048$ points are randomly sampled.

The training set \mathcal{T}_{train} is constructed by including all the blocks that contain at least 100 points for any target class in C_{train} . During training, we randomly sample one episode from \mathcal{T}_{train} using the following procedure: we first randomly choose N classes from C_{train} , where $N < |C_{train}|$; and then a support set S and a query set Q are randomly sampled based on the chosen N classes. The mask \mathbf{M} in the support set and the label \mathbf{L} in the query set are modified from the original point annotations accordingly to correspond to the chosen classes. The testing episodes are formed in a similar fashion, with the exception that we exhaustively iterate all the combinations of N classes out of C_{test} rather than randomly choosing N classes. Specifically, we sample 100 episodes for each combination and use them as the \mathcal{T}_{test} for evaluating each of the methods in our experiments. It is worth highlighting that the same point cloud can appear in both \mathcal{T}_{train} and \mathcal{T}_{test} , but the annotations of this point cloud are different due to the difference in the classes of interest.

Evaluation metric. For the evaluation metric, we adopt the widely used metric in point cloud semantic segmentation, *i.e.* mean Intersection over Union (mean-IoU). In our few-shot setting, the mean-IoU is obtained by averaging over the set of testing classes C_{test} .

Method	2-way						3-way					
	1-shot			5-shot			1-shot			5-shot		
	S ⁰	S ¹	mean	S ⁰	S ¹	mean	S ⁰	S ¹	mean	S ⁰	S ¹	mean
FT	36.34	38.79	37.57	56.49	56.99	56.74	30.05	32.19	31.12	46.88	47.57	47.23
ProtoNet	48.39	49.98	49.19	57.34	63.22	60.28	40.81	45.07	42.94	49.05	53.42	51.24
AttProtoNet	50.98	51.90	51.44	61.02	65.25	63.14	42.16	46.76	44.46	52.20	56.20	54.20
MPTI	52.27	51.48	51.88	58.93	60.56	59.75	44.27	46.92	45.60	51.74	48.57	50.16
Ours	53.77	55.94	54.86	61.67	67.02	64.35	45.18	49.27	47.23	54.92	56.79	55.86

Table 1. Results on **S3DIS** dataset using mean-IoU metric (%). Sⁱ denotes the split *i* is used for testing.

Method	2-way						3-way					
	1-shot			5-shot			1-shot			5-shot		
	S ⁰	S ¹	mean	S ⁰	S ¹	mean	S ⁰	S ¹	mean	S ⁰	S ¹	mean
FT	31.55	28.94	30.25	42.71	37.24	39.98	23.99	19.10	21.55	34.93	28.10	31.52
ProtoNet	33.92	30.95	32.44	45.34	42.01	43.68	28.47	26.13	27.30	37.36	34.98	36.17
AttProtoNet	37.99	34.67	36.33	52.18	46.89	49.54	32.08	28.96	30.52	44.49	39.45	41.97
MPTI	39.27	36.14	37.71	46.90	43.59	45.25	29.96	27.26	28.61	38.14	34.36	36.25
Ours	42.55	40.83	41.69	54.00	50.32	52.16	35.23	30.72	32.98	46.74	40.80	43.77

Table 2. Results on **ScanNet** dataset using mean-IoU metric (%). Sⁱ denotes the split *i* is used for testing.

4.2. Implementation Details

Framework details. We illustrate the architecture and configuration of the embedding network in Figure 2 (bottom). Following [32], the hyper-parameter α in label propagation is set to 0.99. The settings of the other three hyperparameters (*i.e.* n in multi-prototype generation, k and σ in the k -NN graph construction) are discussed in Section 4.4 and supplementary Section B.1.

Training. We pre-train the feature extractor module on training set \mathcal{T}_{train} by adding three MLP layers at the end of feature extractor as the segmentor over C_{train} . During pre-training, we set the batch size to 32 and train for 100 epochs. The pre-trained model is optimized by Adam with a learning rate of 0.001. After initializing the feature extractor with the pre-trained weights, we use the Adam optimizer with an initial learning rate of 0.0001 for the feature extractor module, and an initial learning rate of 0.001 for the metric learner and attention learner modules, respectively. Both learning rates are decayed by half after 5,000 iterations. In each iteration, one episode is randomly sampled, and all the point clouds in the support and query set are augmented by Gaussian jittering and random rotation around z-axis.

4.3. Baselines

We design four baselines for comparison with our method.

Fine-tuning (FT). We take the architecture of our pre-trained segmentation network as the backbone of this baseline. For fair comparison, we use the same pre-trained weights for this segmentation network and our method. Following the strategy in [20], we fine-tune the trained segmentation network on samples from the support set and test on the query samples for each N -way K -shot task. To avoid overfitting, we only fine-tune the last three MLP layers.

Prototypical Learning (ProtoNet). We adapt the prototypical network [5] utilized in the few-shot image segmentation [3, 23] task to few-shot point cloud segmentation. To extract the point-wise features for the support and query sets, we use similar architecture as our embedding network but replace SAN with a linear mapper that maps the features into the same dimension as SAN. Similarly, the feature extractor is initialized by the same pre-trained weights. We represent each class by one prototype given by the mean feature of its support points. The predictions of query points are from its squared Euclidean distance with respect to the prototypes.

Attention-aware Prototypical Learning (AttProtoNet). This baseline is an upgraded version of ProtoNet, where the self-attention mechanism is added into the embedding network. In other words, it uses the same architecture as our embedding network.

Multi-prototype Transductive Inference (MPTI). This can be considered as a degraded version of our method, where the attention learner module (SAN) in the embedding network is replaced by a linear mapper similar to ProtoNet.

4.4. Results and Analyses

Comparison with baselines. Table 1 and 2 summarize the results of comparing our method to the baselines on S3DIS and ScanNet, respectively. It is not surprising that using more labeled samples, *i.e.* larger K -shot leads to significant improvements for all the methods. We also observe that the performance of 3-way is generally lower than 2-way segmentation due to its higher difficulty. As can be seen from the two tables, our proposed method consistently and significantly outperforms the baselines in all four settings, *i.e.* 2/3-way 1/5-shot on both datasets. Particularly, our method improves upon FT under the challenging 3-way 1-

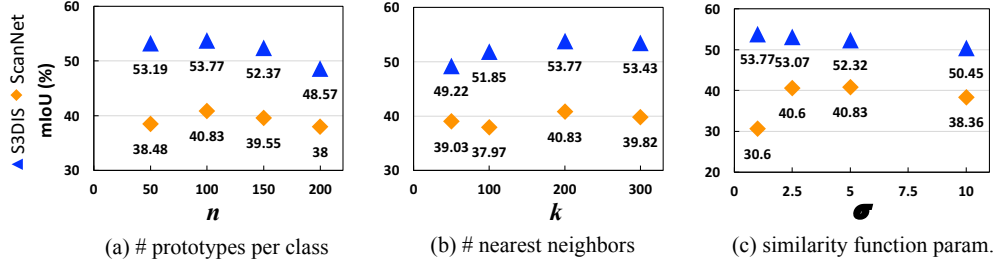


Figure 4. Effects of three hyper-parameters under 2-way 1-shot setting on S3DIS (S^0) and ScanNet (S^1) datasets.

$f_{geometric}$	$f_{semantic}$	f_{metric}	S3DIS	ScanNet
✓	✗	✗	40.31	26.91
✗	✓	✗	44.43	34.51
✗	✗	✓	48.24	35.07
✓	✓	✗	47.82	38.69
✓	✗	✓	52.21	36.12
✗	✓	✓	50.12	39.81
✓	✓	✓	53.77	40.83

Table 3. Effects of different levels of features under 2-way 1-shot setting on S3DIS (S^0) and ScanNet (S^1) datasets.

shot setting by around 52% and 53% on S3DIS and ScanNet dataset, respectively. Compared to ProtoNet, our method gains around 10% and over 20% improvements on S3DIS and ScanNet, respectively, when using only one sample, *i.e.* one-shot. These improvements shows that our proposed method can obtain more useful knowledge from very limited data during inference. The superiority of our method as compared to AttProtoNet shows the contribution of our proposed multi-prototype transductive inference mechanism. Additionally, both the improvements of AttProtoNet over ProtoNet and the improvements of our method over MPTI demonstrate the capacity of self-attention network in exploiting semantic correlations between the points, which is very important in inferring point-wise semantic labels.

An interesting observation is that the degraded version of our method, *i.e.* MPTI clearly outperforms ProtoNet under the one-shot settings, but loses the gain under five shots. This is probably due to the naive multi-prototype generation of MPTI made it difficult to extract accurate multi-prototypes for a large number of support points if the embedding space is not learned well. This phenomenon also indicates the importance of incorporating the self-attention network, which helps in learning a more representative embedding space.

Ablation study of multi-level features. We study the effects of various designs of the embedding network since it is one of the most important components of our method. We denote the levels of features, *i.e.* local geometric feature, global semantic feature, and metric-adaptive feature as $f_{geometric}$, $f_{semantic}$, and f_{metric} , respectively. We select one or two level(s) of feature(s) as our embedded feature³ for the estimation of its(their) contribution(s). The results

³Specifically, the output of the embedding network will be the selected feature or the concatenation of the two selected features.

of six variants are listed in Table 3. From the perspective of individual feature, $f_{semantic}$ and f_{metric} contribute more than $f_{geometric}$. This is reasonable since the embedding space are supposed to be semantic. By combining any two levels of features, we achieve varying improvements on the two datasets. Eventually, the integration of the three levels of features gives us the best performance on both datasets.

Effects of hyper-parameters. In Figure 4, we illustrate the effects of three hyper-parameters (*i.e.*, n , k , σ) for 2-way 1-shot point cloud semantic segmentation on one split of each dataset. As can be seen from Figure 4(a), increasing the number of prototypes per class n can achieve better results, but overly large n can lead to the over-fitting problem and cause adverse impact on the performance. As Figure 4(b) shows, there is a slight difference on performance between the two datasets when choosing a smaller k , *i.e.* $k = 50$. However, the overall trend is similar, and the selection of $k = 200$ gives the best result on both datasets. As reported in [13, 25], σ in the Gaussian similarity function used in the construction of the affinity matrix (see Eq. 3) plays a role in the performance of label propagation. We empirically find that σ in different datasets has different optimal values. Specifically, $\sigma = 1$ on S3DIS and $\sigma = 5$ on ScanNet enable us to achieve the best performance, respectively.

4.5. Qualitative Results

Figure 5 and 6 show the qualitative results of our proposed method for 2-way 1-shot point cloud semantic segmentation on the S3DIS and ScanNet dataset, respectively. We compare the predictions of one query point cloud from our method with the ground truths and predictions from ProtoNet. As we can see from Figure 5, the S3DIS dataset is very challenging in many scenarios, *e.g.*, “the white columns that are very similar to the white wall and the window frame” (first row of Figure 5), “the doors that only have visible door frames” (second row of Figure 5), “the table that has a lot of clutter on it” (last row of Figure 5). The accuracy of the predictions from our method drops for these challenging scenarios due to the limitation of having only one labeled sample as support. Nonetheless, our method still generally gives more accurate segmentation results than ProtoNet in all cases (*e.g.* our method nicely segments the ‘ceiling’, ‘floor’,

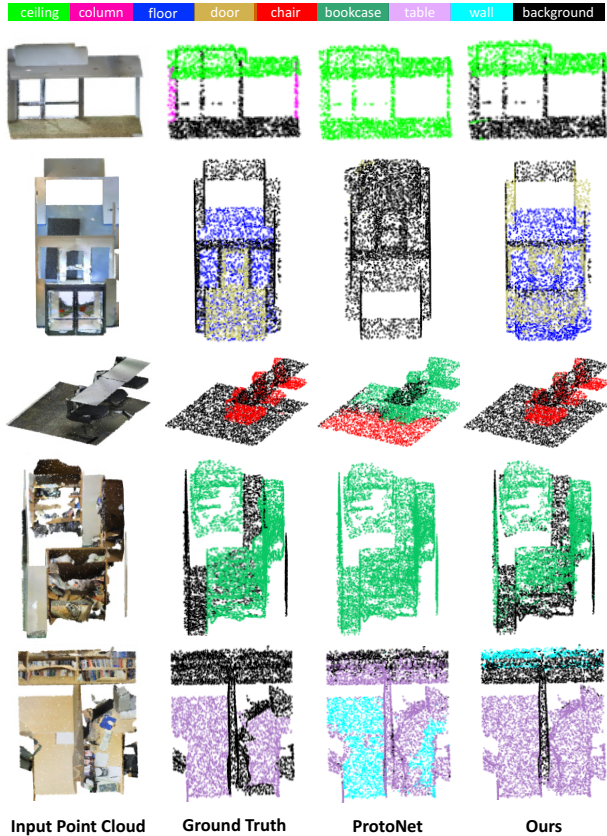


Figure 5. Qualitative results of our method in 2-way 1-shot point cloud semantic segmentation on the **S3DIS** dataset in comparison to the ground truth and ProtoNet. Four combinations of 2-way are illustrated from the top to bottom rows, i.e., “ceiling, column” (first row), “floor, door” (second row), “chair, bookcase” (third and fourth rows), “table, wall” (last row).

‘chair’, ‘bookcase’, ‘table’ in each scene from top to bottom in Figure 5).

In contrast to the S3DIS dataset, the ScanNet dataset contains more diverse room types, such as *bathroom* (see first and last rows of Figure 6), *bedroom/hotel* (see second row of Figure 6), *living room/lounge* (see third and fifth rows of Figure 6), etc. Our proposed method is able to correctly segment most of semantic classes within these new room types, while ProtoNet gives poor segmentation results that mix the background class with the semantic classes. We believe that our correct segmentations are consequences of integrating the attention-aware feature embedding and multi-prototype transductive inference, which facilitates the smoothness among neighboring points and the distinction between different semantic classes.

5. Conclusion

This paper investigates the unexplored yet important few-shot point cloud semantic segmentation problem. We propose a novel solution: the attention-aware multi-prototype

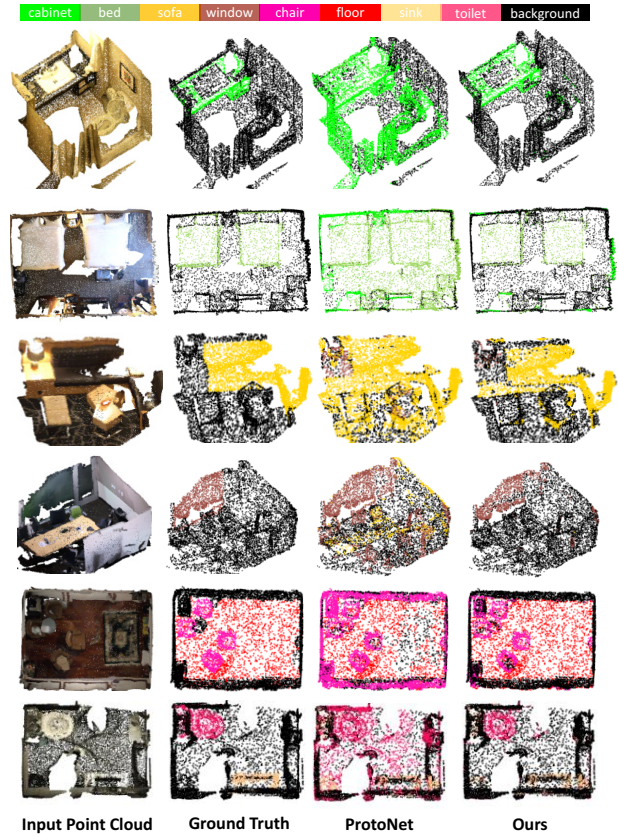


Figure 6. Qualitative results of our method in 2-way 1-shot point cloud semantic segmentation on the **ScanNet** dataset in comparison to the ground truth and ProtoNet. Four combinations of 2-way are illustrated from the top to bottom rows, i.e., “cabinet, bed” (first and second rows), “sofa, window” (third and fourth rows), “chair, floor” (fifth row), “sink, toilet” (last row).

transductive inference method, which achieves clear and consistent improvements over baselines on a variety of few-shot point cloud semantic segmentation tasks. Furthermore, this work offers several key insights on few-shot 3D point cloud semantic segmentation. Firstly, the learning of the discriminative features that encode both geometric and semantic context is the core of the metric-based few-shot point cloud semantic segmentation method. Secondly, the data distributions of 3D point clouds are complex and cannot be sufficiently modeled by a uni-modal distribution. Thirdly, the exploitation of intrinsic relationships in the embedding space is necessary for the point cloud segmentation task. Future work could investigate an adaptive generation of multi-prototypes to efficiently infer the number of prototypes based on data complexity.

Acknowledgements. This research is supported in part by the National Research Foundation, Singapore under its IRC@SG Funding Initiative and partially performed on resources of the National Supercomputing Centre, Singapore. It is also partially supported by the Tier 2 grant MOE-T2EP20120-0011 from the Singapore Ministry of Education.

References

- [1] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1534–1543, 2016.
- [2] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Niessner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5828–5839, 2017.
- [3] Nanqing Dong and Eric Xing. Few-shot semantic segmentation with prototype learning. In *Proceedings of the British Machine Vision Conference*, volume 3, 2018.
- [4] Chelsea Finn, Pieter Abbeel, and Sergey Levine. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 1126–1135. JMLR. org, 2017.
- [5] Victor Garcia and Joan Bruna Estrach. Few-shot learning with graph neural networks. In *6th International Conference on Learning Representations, ICLR 2018*, 2018.
- [6] S Guinard and L Landrieu. Weakly supervised segmentation-aided classification of urban scenes from 3d lidar point clouds. *ISPRS-International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pages 151–157, 2017.
- [7] Qingyong Hu, Bo Yang, Linhai Xie, Stefano Rosa, Yulan Guo, Zhihua Wang, Niki Trigoni, and Andrew Markham. Randla-net: Efficient semantic segmentation of large-scale point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11108–11117, 2020.
- [8] Qiangui Huang, Weiyue Wang, and Ulrich Neumann. Recurrent slice networks for 3d segmentation of point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2626–2635, 2018.
- [9] Ahmet Iscen, Giorgos Toliás, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5070–5079, 2019.
- [10] Loïc Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [11] Yangyan Li, Rui Bu, Mingchao Sun, Wei Wu, Xinhan Di, and Baoquan Chen. Pointcnn: Convolution on x-transformed points. In *Advances in neural information processing systems*, pages 820–830, 2018.
- [12] Weide Liu, Chi Zhang, Guosheng Lin, and Fayao Liu. Crnet: Cross-reference networks for few-shot segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4165–4173, 2020.
- [13] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. In *7th International Conference on Learning Representations, ICLR 2019*, 2019.
- [14] Jilin Mei, Biao Gao, Donghao Xu, Wen Yao, Xijun Zhao, and Huijing Zhao. Semantic segmentation of 3d lidar data in dynamic scene using semi-supervised learning. *IEEE Transactions on Intelligent Transportation Systems*, 2019.
- [15] Tsendsuren Munkhdalai and Hong Yu. Meta networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2554–2563. JMLR. org, 2017.
- [16] Khoi Nguyen and Sinisa Todorovic. Feature weighting and boosting for few-shot segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 622–631, 2019.
- [17] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 652–660, 2017.
- [18] Sachin Ravi and Hugo Larochelle. Optimization as a model for few-shot learning. In *5th International Conference on Learning Representations, ICLR 2017*, 2017.
- [19] Adam Santoro, Sergey Bartunov, Matthew Botvinick, Daan Wierstra, and Timothy Lillicrap. Meta-learning with memory-augmented neural networks. In *International conference on machine learning*, pages 1842–1850, 2016.
- [20] Amirreza Shaban, Shray Bansal, Zhen Liu, Irfan Essa, and Byron Boots. One-shot learning for semantic segmentation. In *Proceedings of the British Machine Vision Conference*, 2017.
- [21] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017.
- [22] Oriol Vinyals, Charles Blundell, Timothy Lillicrap, Daan Wierstra, et al. Matching networks for one shot learning. In *Advances in neural information processing systems*, pages 3630–3638, 2016.
- [23] Kaixin Wang, Jun Hao Liew, Yingtian Zou, Daquan Zhou, and Jiashi Feng. Panet: Few-shot image semantic segmentation with prototype alignment. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9197–9206, 2019.
- [24] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 38(5):1–12, 2019.
- [25] Zhu Xiaojin and Ghahramani Zoubin. Learning from labeled and unlabeled data with label propagation. *Tech. Rep., Technical Report CMU-CALD-02-107, Carnegie Mellon University*, 2002.
- [26] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European Conference on Computer Vision*, pages 574–591. Springer, 2020.
- [27] Xun Xu and Gim Hee Lee. Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2020.

- [28] Xiaoqing Ye, Jiamao Li, Hexiao Huang, Liang Du, and Xiaolin Zhang. 3d recurrent neural networks with context fusion for point cloud semantic segmentation. In *European Conference on Computer Vision*, pages 415–430. Springer, 2018.
- [29] Chi Zhang, Guosheng Lin, Fayao Liu, Jiushuang Guo, Qingyao Wu, and Rui Yao. Pyramid graph networks with connection attentions for region-based one-shot semantic segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 9587–9595, 2019.
- [30] Chi Zhang, Guosheng Lin, Fayao Liu, Rui Yao, and Chunhua Shen. Canet: Class-agnostic segmentation networks with iterative refinement and attentive few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5217–5226, 2019.
- [31] Na Zhao, Tat-Seng Chua, and Gim Hee Lee. Ps²-net: A locally and globally aware network for point-based semantic segmentation. In *Proceedings of 25th International Conference on Pattern Recognition (ICPR)*, pages 723–730, 2020.
- [32] Dengyong Zhou, Olivier Bousquet, Thomas N Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. In *Advances in neural information processing systems*, pages 321–328, 2004.