

Rethinking Task and Metrics of Instance Segmentation on 3D Point Clouds

Kosuke Arase¹, Yusuke Mukuta^{1,2}, Tatsuya Harada^{1,2}

¹The University of Tokyo ²RIKEN AIP

{arase,mukuta,harada}@mi.t.u-tokyo.ac.jp

Abstract

Instance segmentation on 3D point clouds is one of the most extensively researched areas toward the realization of autonomous cars and robots. Certain existing studies have split input point clouds into small regions such as $1\text{ m} \times 1\text{ m}$; one reason for this is that models in the studies cannot consume a large number of points because of the large space complexity. However, because such small regions occasionally include a very small number of instances belonging to the same class, an evaluation using existing metrics such as mAP is largely affected by the category recognition performance. To address these problems, we propose a new method with space complexity $\mathcal{O}(N_p)$ such that large regions can be consumed, as well as novel metrics for tasks that are independent of the categories or size of the inputs. Our method learns a mapping from input point clouds to an embedding space, where the embeddings form clusters for each instance and distinguish instances using these clusters during testing. Our method achieves state-of-the-art performance using both existing and the proposed metrics. Moreover, we show that our new metric can evaluate the performance of a task without being affected by any other condition.

1. Introduction

3D environment recognition has been extensively researched toward the realization of autonomous cars and robots. In particular, instance segmentation, the task of not only labeling each point but also distinguishing each instance belonging to the same class, is one of the key tasks to such realization. Instance segmentation is challenging because the number of instances is not fixed, and thus, methods for categorical classification cannot be directly applied. Although there are several typical 3D data representations such as voxels, meshes, and point clouds, in this study, we focus on point clouds, which can be obtained directly from depth sensors such as Light Detection and Ranging (LiDAR).

The instance segmentation model learns the mapping

from each input point to the semantics of the corresponding point. When evaluating the instance segmentation model, the pairing of a prediction and a ground truth is considered true positive when the intersection over union (IoU) between them is higher than the threshold. In many cases, semantic segmentation can be solved simultaneously, and thus, the important issue is distinguishing objects in the same category.

There have been many studies on instance segmentation, where the input point clouds have been split into small regions such as 1 m square [21,22]; however, conducting evaluations on such small regions is somewhat complicated.

One solution is first merging small regions into one entire scene prediction and then evaluating the entire scene [22]. However, the final result is largely affected by the merging algorithm, and it is difficult to evaluate the pure instance segmentation performance.

Another way is evaluating the instance segmentation in small regions [21]; however, this is not desirable owing to the following reason. As shown in Figure 1, small regions often contain only one instance for a certain category, and in such cases, the resulting semantic segmentation is sufficient for instance segmentation because it is not necessary to distinguish objects belonging to the same class. When the input regions are too small and there is only one object in each region, it is unnecessary to distinguish the object, and thus, instance segmentation does not have to be conducted. Conversely, when there are many objects belonging to the same class, it is necessary to consume larger regions in order to evaluate an instance segmentation. Consuming large regions is also challenging because it is necessary to consume a large number of points to avoid a sparse input, which decreases the performance of certain models including PointNet [15, 16]. Handling dense point clouds is also helpful in the application of instance segmentation. However, as an example, the Similarity Group Proposal Network (SGPN) [22] calculates the similarities for each pair of points and its space complexity is $\mathcal{O}(N_p^2)$ for N_p number of points, which makes it difficult to consume large point clouds. Thus, a memory efficient method is required.

Moreover, there are certain problems in existing metrics.

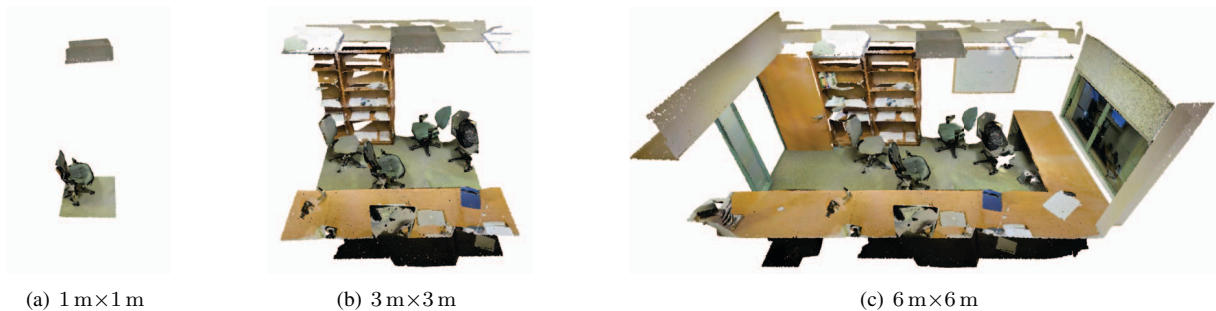


Figure 1. Objects within grids for various grid sizes (S3DIS [1], Area 6, Office 29)

The method in [22] was evaluated using the mean average precision (mAP), which has a characteristic in that the effect of false positives with low confidence scores is small. Although this property is appropriate for tasks such as object detection where multiple candidates with overlaps are allowed, or retrieval where the rank of the output is important, this property is not suitable for instance segmentation. Because outputs are objects without an overlap for each point, we need to equally evaluate the outputs of each point regardless of the confidence score. In addition, when evaluating instance segmentation, we focus on whether two objects are properly distinguished, and whether one object is incorrectly split. However, these failures cannot be distinguished from a misclassification when performing evaluations using existing metrics, and a misclassification is often the main factor of decreasing mAP.

In this study, we first experimentally show our claim that evaluating instance segmentation in small regions with existing metrics is inappropriate and reveal the problem using metrics that has not been investigated in previous studies. Then, we propose a novel instance segmentation method with small space complexity that enables the consumption of large regions. Our loss function learns a one-to-one mapping from an input feature space to an embedding space, where embeddings from the same instance form a cluster, and we can distinguish instances by clustering at the test time. Because our method does not have to handle point pairs, the space complexity is $\mathcal{O}(N)$ and is scalable to the number of points. We show that the proposed memory efficient method outperforms other state-of-the-art methods.

In addition, we demonstrate that consuming small regions and evaluating them by using existing metrics is not appropriate. This fact has been overlooked by previous research, so we propose a novel metric that can evaluate it correctly for the first time. Our metric is based on inclusion, which is the relationship of one set being a subset of another. Using the proposed metric, we can evaluate the pure performance regardless of the size of the regions, categories, or confidence scores. We can also analyze the types of errors quantitatively.

We conducted extensive experiments to reveal the effect of the size of the regions and the density of the points on the instance segmentation performance and showed that consuming a large number of points increases the performance for large regions.

The key contributions of this study are as follows:

- We propose a new loss function that learns to push embeddings for each instance to be clustered and is scalable to the number of points; we also experimentally demonstrated that the proposed method outperforms existing methods.
- We reveal the problems associated with existing metrics, including the fact that they are affected by the size of inputs or categories, which have been overlooked in previous research.
- We propose a novel metric that is not affected by these factors and can evaluate instance segmentation performance correctly.

2. Related Work

2.1. Feature Extraction on 3D Point Clouds

The effective handling of point clouds is challenging because they are unordered, non-uniformly distributed data.

Methods to extract features from point clouds can be roughly classified into two approaches, namely, describing local features [18, 19] and describing relationships among multiple points [2, 6].

PointNet [15], which addresses the problem of unordered data by using symmetric functions, and PointNet++ [16], which stacks PointNets and is able to handle local features, have made recent breakthroughs in deep learning on point clouds. We use PointNet and PointNet++ as feature extractors in this study.

2.2. Instance Segmentation

Segmentation is a task of labeling each minimum element in the data such as a pixel or point. In particular, labeling the category of each element and distinguishing objects

belonging to the same category are called instance segmentation against semantic segmentation.

2D Images Many studies on instance segmentation on images have been recently published [4, 5, 8, 12–14], and Novotny *et al.* [14] classified instance segmentation into two approaches, *propose & verify* (P&V) and *instance coloring* (IC). P&V is an approach that first proposes candidates of objects based on their objectness and then verifies whether an object is a candidate. This is currently a popular approach in the field of object detection [17] and instance segmentation [8] on images. Although P&V approaches have achieved significant success in image segmentation, they have weaknesses in that object candidates are approximations of the object shapes, and a second-stage to refine the candidates is necessary for segmentation, such as Mask R-CNN [8], and thus, the network architecture tends to be complex.

Approaches labeling an object identifier directly to each pixel are called IC, and some studies have been conducted in this area for image segmentation [5, 12, 14]. Brabandere *et al.* [5] proposed a discriminative loss function that learns a mapping to an embedding space where the embeddings form clusters for each object. The loss function is simple and efficient but has some shortcomings, as described in Section 2.3.

We choose an IC-based approach because the architecture tends to be simpler, and it is thus expected to be computationally efficient.

3D Point Clouds SGPN [22] and deep functional dictionaries (DFD) [21] have tackled instance segmentation on 3D point clouds. SGPN first predicts similarities for every pair of points that describes whether two points belong to the same object and then merges points to instance proposals by considering a pair of points with a similarity higher than a certain threshold as being contained in the same object. Although it is a pioneering work of instance segmentation on points clouds, the space complexity of the similarity matrix is proportional to the square of the number of points and cannot handle too many points. We discuss this problem in Section 3.2. Thus, input scenes are split into 1 m square regions, and the results are then aggregated for each region using a heuristic algorithm. However, the final performance depends on the merging algorithm, as described in Section 1, and applying the method for every small region is computationally inefficient.

Recently, Sung *et al.* [21] proposed a general method called DFD that produces a dictionary of the probe functions. The authors proposed a general framework that learns a mapping from the shape to the dictionary. Each atom of the dictionary can be associated with semantics, instances, or something else based on the task and constraint. A per-

formance comparable to that of state-of-the-art techniques was achieved on S3DIS, but the authors evaluated its performance for each small region. Thus, this evaluation has certain problems, as discussed in Section 1.

2.3. Embedding Learning

Our method performs instance segmentation by first learning the feature embeddings for each point such that the diameter of the embedding cluster corresponding to the same object is small compared to the distance among clusters from different objects; then, clustering is conducted in the embedding space. Such a feature learning method that trains the embedding to minimize the distance between embeddings with the same semantics while maximizing the distance between embeddings with different semantics is widely used in category classification [3, 23] and similarity learning [11, 20]. This concept has been used for recent instance segmentation studies on images such as those on discriminative loss [5], [12]. Inspired by this, we propose a novel instance segmentation method that overcomes the discriminative loss problem.

Discriminative loss L consists of L_{var} , which makes the distance between points and centroids of the corresponding cluster smaller than δ_v ; L_{dist} , which makes the distance between cluster centroids larger than δ_d ; and a regularizer L_{reg} , which prevents the feature norms from diverging. Here, L is written as follows:

$$L_{\text{var}} = \frac{1}{C} \sum_{c=1}^C \frac{1}{N_c} \sum_{i=1}^{N_c} [\|\mu_c - \mathbf{x}_i\| - \delta_v]_+^2 \quad (1)$$

$$L_{\text{dist}} = \frac{1}{C(C-1)} \sum_{c_A \neq c_B} [2\delta_d - \|\mu_{c_A} - \mu_{c_B}\|]_+^2 \quad (2)$$

$$L_{\text{reg}} = \frac{1}{C} \sum_{c=1}^C \|\mu_c\| \quad (3)$$

$$L = L_{\text{sem}} + \alpha L_{\text{var}} + \beta L_{\text{dist}} + \gamma L_{\text{reg}}, \quad (4)$$

where C denotes the number of clusters, and μ_c and N_c are the centroid and number of points of cluster c , respectively, \mathbf{x}_i is the embedding, L_{sem} is the softmax cross entropy loss of the category classification, $\|\cdot\|$ is the Euclidean norm in the feature space, and $[x]_+ = \max(0, x)$. When we conduct instance segmentation, we apply clustering on the learned embedding space. When we set $\delta_d \geq \delta_v$ and the learned embedding space satisfies $L_{\text{var}} = L_{\text{dist}} = 0$, we can guarantee that all points whose distances from a point are smaller than δ_v belong to the same object.

However, there are some drawbacks in this original formulation of discriminative loss. First, it is difficult to select the hyperparameter β, γ that balances the weights of L_{reg} and L_{dist} . The optimization is hyperparameter-sensitive because L_{reg} attempts to reduce the distances between points

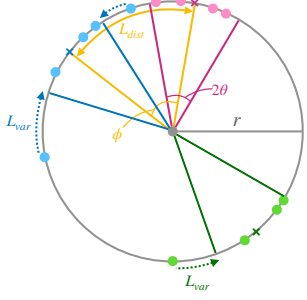


Figure 2. Overview of the proposed feature learning method. Each point (●) represents one feature embedding, and the points with the same color belong to the same object. The crosses (×) are cluster centroids. As the training progresses, points with the same colors move to a nearby spot and clusters move away from each other.

(i.e., make them closer), whereas L_{dist} attempts to increase the distance between points (i.e., make them more distant). Empirically, it turns out that γ should be about 100-times smaller than α, β , and seeking such balance is an cumbersome task. Moreover, when we concatenate the learned feature to other features such as the raw coordinates of a point, we need to arrange the scale of the features such that both features are effective for clustering. However, it is difficult to arrange the scale because the norms of the feature are different among feature spaces. In contrast, when we normalize each feature after we learn the feature space, we cannot distinguish between points with the same unit vector and a different norm.

In the following sections, we propose a novel embedding method that solves these problems.

3. Method

3.1. Proposed Feature Embedding

In this section, we describe the proposed feature learning method. As described in a previous section, Equation (2) in L attempts to increase the distances between different clusters while minimizing the norms of the feature using Equation (3). Thus, L is sensitive to the hyperparameters β and γ that balance these conflicting losses. Moreover, it is difficult to combine a learned feature with other features for clustering.

In this study, we overcome these difficulties by restricting the features to a unit hypersphere and the learning of the feature space based on a cosine similarity instead of the Euclidean loss. We present an overview of our method in Figure 2, where each point (●) represents one feature embedding, points with the same color belong to the same object, and the cross (×) indicates the cluster centroid. Moreover, θ and ϕ satisfy $\delta_v = \cos(\theta)$ and $\delta_d < \cos(\phi)$, respectively.

Using the cosine similarity between two embeddings

$\mathbf{x}_i, \mathbf{x}_j$, which is calculated as $s(\mathbf{x}_i, \mathbf{x}_j) = \frac{\mathbf{x}_i^T \mathbf{x}_j}{\|\mathbf{x}_i\| \|\mathbf{x}_j\|}$, the proposed loss function is written as follows:

$$L_{\text{var}} = \frac{1}{C} \sum_{c=1}^C \frac{1}{N_c} \sum_{i=1}^{N_c} [\delta_v - s(\boldsymbol{\mu}_c, \mathbf{x}_i)]_+ \quad (5)$$

$$L_{\text{dist}} = \frac{1}{C(C-1)} \sum_{c_A \neq c_B} [s(\boldsymbol{\mu}_{c_A}, \boldsymbol{\mu}_{c_B}) - \delta_d]_+ \quad (6)$$

$$L = L_{\text{sem}} + \alpha L_{\text{var}} + \beta L_{\text{dist}}, \quad (7)$$

where δ_d and δ_v satisfy $\delta_d \ll \delta_v \approx 1$ so that $s(\boldsymbol{\mu}_c, \mathbf{x}_i)$ becomes larger than $s(\boldsymbol{\mu}_{c_A}, \boldsymbol{\mu}_{c_B})$. In addition, we use the absolute error of the $[\cdot]_+$ terms instead of the squared error adopted in [5] because the norm of the $[\cdot]_+$ terms is smaller than 1 and the squared errors become considerably smaller when these terms are near zero.

When the angle between an embedding and its cluster centroid is larger than θ , L_{var} attempts to reduce the distance between the embedding and the centroid. In addition, L_{dist} attempts to increase distance between cluster centroids when the cosine similarity is larger than δ_d . In an image recognition study ([12]), the feature was also learned using the cosine similarity using a unit hypersphere. However, in that study, similarities between all pairs of points were calculated, whereas our method only considers the similarities between points and the corresponding cluster centroids. Thus, our method is considerably more computationally effective.

Compared to [5], the advantages of our method are as follows:

- We do not need to consider the scale of the feature space and thus, we can omit L_{reg} and do not need to consider the balance between β and γ .
- Because the embeddings are guaranteed to have a unit norm, it is easy to combine the learned embeddings to other features.

We learn the mapping from the feature space to the embedding space by adding one fully connected layer.

3.2. Computational Complexity

In Section 2.2, we discussed the fact that one of the problems of the existing IC-based instance segmentation method SGPN [22] is that it requires a large space complexity. Because our method and SGPN require only a few extra layers in the feature extractor, and thus, the number of iterations for training is nearly the same, we focus on analyzing the detailed computation complexity of the loss functions of SGPN and our method. In the following section, we denote the batch size as B , the number of points as N_p , the number of points in a cluster c as N_c , and the dimensions of the input feature space and embedding space as d_f and d_e ,

respectively. We also write the input and embedded features of the i -th point as $\mathbf{f}_i \in \mathbb{R}^{d_f}$ and $\mathbf{h}_i \in \mathbb{R}^{d_e}$, respectively.

In SGPN, the similarity S_{ij} between the i -th and j -th points is calculated as

$$S_{ij} = \|\mathbf{f}_i - \mathbf{f}_j\| = \sqrt{\|\mathbf{f}_i\|^2 - 2\mathbf{f}_i\mathbf{f}_j + \|\mathbf{f}_j\|^2}. \quad (8)$$

Because the method calculates S_{ij} for all pairs of points, the space complexity of the similarity matrix is $\mathcal{O}(BN_p^2d_f)$. As for the time complexity, because we need to evaluate $\|\mathbf{f}_i\|^2$ for each i and $\mathbf{f}_i\mathbf{f}_j$ for each pair (i, j) to calculate S_{ij} , the time complexity is $\mathcal{O}(BN_p^2d_f)$.

In contrast, the proposed loss function obtains d_e -dimensional embedded features and calculates the cosine similarity between each point and its cluster centroid, and between each pair of cluster centroids. Therefore, the space complexity for the embeddings of each point is $\mathcal{O}(BN_p d_e)$, and the computation complexity is $\mathcal{O}(B(N_p + C^2)d_e)$; however, this order is equivalent to $\mathcal{O}(BN_p d_e)$ because $C \ll N_p$ in most cases. Both complexities are linear in N_p . Because we use $N_p = 2^n$ ($n = 12, 13, 14$), $d_e = 2^5$ in the experiment, the proposed method can calculate the loss function with a smaller space/time complexity than SGPN.

3.3. Clustering

We describe our feature learning method in Section 3.1. In this section, we explain the clustering method applied to the learned feature space to conduct instance segmentation.

The requirements for the clustering method are as follows:

- The number of clusters is variable.
- The clustering result is robust to outliers.
- The clustering does not fail even when the number of points in each cluster has a large variety.

In this study, we adopt the density-based spatial clustering of applications with noise (DBSCAN) [7], which satisfies these requirements. DBSCAN is a density-based clustering method that first calculates the densities of points based on the number of neighboring points and then constructs clusters by considering a continuous region with a density of above a certain threshold as a single cluster. The number of clusters of the output of DBSCAN can vary, and DBSCAN is robust to outliers because it accepts the noise points that do not belong to any clusters.

We apply this clustering to the embeddings, which are predicted as the same category. We concatenate the learned embeddings using the normalized coordinates of the point as the input for the clustering method. In addition, some clusters consist of a very small number of points. Because such clusters are false positive in most cases, we handle such clusters as points in that they do not belong to any cluster in the evaluation.

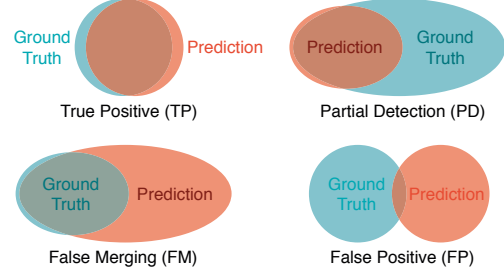


Figure 3. Error patterns for instance segmentation

4. Evaluation Metrics

As described in Section 1, existing metrics of instance segmentation are affected by the misclassification, the confidence of prediction, and the size of the regions. Therefore, we propose a novel evaluation metric that focuses on the distinct ability of the objects regardless of the confidence or semantics, and which can be used for any sized input region.

When we neglect semantic errors, we can observe four patterns for each prediction output:

- There is a corresponding ground truth (GT) for the prediction output (true positive (TP)).
- The prediction output covers some part of a GT (partial detection (PD)).
- The prediction output contains more than one GT (false merging (FM)).
- There is no corresponding GT (false positive (FP)).

Figure 3 shows a diagram of these four error patterns.

Note that one prediction output can fulfill more than one of the patterns even though each point corresponds to exactly one GT and one prediction. For example, one prediction output, 90% of which is contained in a GT, can cover other small GTs with the remaining 10%. In particular, PD and FM are characteristics of instance segmentation.

To formulate these patterns, we define "intersection over a set (IoS)", which describes the part of an object A that is contained in an object B as follows:

$$IoS(A, B) = \frac{N(A \cap B)}{N(A)} \quad (9)$$

where $N(X)$ denotes the number of points in X , and object A is considered to be contained in object B when $IoS(A, B)$ exceeds a certain threshold t . Note that $IoS()$ is an asymmetric function and A cannot be contained in more than one object when we set $t > 0.5$.

Using this IoS, we can establish the proposed metrics as follows. We first calculate a map from the GTs to the prediction outputs ($gt2pred$) that describes which prediction

outputs are contained in each GT, and conversely calculate a map from the prediction outputs to the GTs ($pred2gt$) that describes which GTs are contained in each prediction. Note that $gt2pred$ and $pred2gt$ are not exclusive. Then, we label each prediction for at least one of the patterns. As for $gt2pred$, for each GT g , we can obtain a list of prediction outputs corresponding to g ($g2p$). If a prediction p is on the list that also contains g itself, p is considered as TP; otherwise, p is considered as PD because p is a subset of g .

In contrast, for $pred2gt$, with each prediction output p , we can obtain a list of GTs corresponding to p ($p2g$). If a GT g on the list does not contain any data and p is not labeled as PD in the last process, p is considered as FP. Otherwise, if g also contains p itself, it must be labeled as TP in the last process owing to its symmetry. Here, g , which does not contain p , is considered as FM because g is a subset of g in this case. We define the ratio of TP to the number of predictions as *precision* and the ratio of TP to the number of GTs as *recall*; we define the F-score as their harmonic mean. We also evaluate the error patterns based on the ratio of PD, FM, and FP to the number of predictions. Because we ignore the semantic segmentation in the calculation, one prediction output can be TP even if its predicted semantics are incorrect. This proposed metric does not depend on the semantics, confidence, or size of the input regions. Therefore, we can evaluate the pure performance of the instance segmentation. The procedure explained above can be written as Algorithm 1.

Algorithm 1 Criteria for instance segmentation

```

1: procedure AggregateResults( $GT$ , prediction  $P$ )
2:   arr  $gt2pred[\text{len}(G)][]$   $\triangleright$  map from  $GT$  to  $preds$ 
3:   arr  $pred2gt[\text{len}(P)][]$   $\triangleright$  map from  $pred$  to  $GTs$ 
4:   for each  $g$  in  $G$  do
5:     for each  $p$  in  $P$  do
6:       if  $s(g \cap p)/s(g) > t$  then  $\triangleright g$  is included in  $p$ 
7:          $pred2gt.append(g)$ 
8:       if  $s(g \cap p)/s(p) > t$  then  $\triangleright p$  is included in  $g$ 
9:          $gt2pred.append(p)$ 
10:  Summarize( $gt2pred$ ,  $pred2gt$ )
11: procedure Summarize( $gt2pred$ ,  $pred2gt$ )
12:   $results[\text{len}(P)][]$   $\triangleright$  2D array to store predictions
13:  for each  $g2p$  in  $gt2pred$  do
14:    for each  $p$  in  $g2p$  do
15:      if  $g$  in  $pred2gt[p]$  then
16:         $results[p].append("TP")$   $\triangleright$  true positive
17:      else
18:         $results[p].append("PD")$   $\triangleright$  partial detection
19:  for each  $p2g$  in  $pred2gt$  do
20:    if  $\text{len}(p2g) == 0$  and  $results[p] == []$  then
21:       $results[p].append("FP")$   $\triangleright$  false positive
22:    for each  $g$  in  $p2g$  do
23:      if  $p$  not in  $gt2pred[g]$  then
24:         $results[p].append("FM")$   $\triangleright$  false merging

```

5. Experiments

In this section, we conduct experiments to compare our method with existing methods in order to demonstrate its effectiveness and to show that existing evaluation metrics of instance segmentation in small regions are inappropriate. We then clearly distinguish errors of misclassification and splitting instances by using our proposed evaluation metric, which cannot be achieved using existing evaluation metrics such as mAP. Moreover, we evaluate the relationships between the size of the split regions and the instance segmentation performance to validate our assumption that evaluating instance segmentation methods using existing evaluation metrics for small regions is inappropriate.

5.1. Datasets and Setup

We use the Stanford large-scale 3D Indoor Spaces Dataset (S3DIS) [1]. S3DIS consists of 270 indoor scenes scanned from six areas and 13 objects. We use 203 scenes for training and the remaining 67 scenes for evaluation.

Although PointNet++ [16] and SGPN [22] have been used to evaluate methods by splitting the input scene horizontally into small regions, such as 1 m square regions, we conducted additional experiments using larger regions as input. This is because one of our aims is to construct a method that can be applied to wide regions with a greater number of points. During each training iteration, we randomly sample subregions with a fixed size from each scene, and then randomly sample a fixed number of points from the sampled subregion as the input. In the following experiments, the region size is 1 m square, and the number of points is 4,096 unless otherwise noted. Each point has a nine-dimensional normalized feature consisting of RGB values, relative coordinates in the subregion, and absolute coordinates in the room. For data augmentation, we apply random noise to some of the input features.

The number of objects in the dataset differs significantly among categories. For example, the number of objects of the category with the largest number of objects is 55 times as large as the number for the category with the smallest number of objects. To eliminate the effect of this imbalance, we weight the miscategorization cross-entropy loss as the weight corresponding to the category with a small number of points weighted as a large value.

Although we can apply our embedding learning method to any feature extractors, we use PointNet (PN) [15] and PointNet++ (PN++) [16] as feature extractors for our experiments. We use 131-dimensional features consisting of 128-dimensional features extracted using the feature extractor and three-dimensional RGB features as the input for feature embedding. The 131-dimensional features are then passed through a fully connected layer, which produces 32-dimensional embeddings.

We use the Adam [10] optimizer with an initial learning

rate of 0.001 and a batch size of 32. We train our network for 6,000 steps, and the learning rate is divided by 10 at the 4,500th step. We set the hyperparameters for our method as $\delta_v = 0.9, \delta_d = 0.4, \alpha = \beta = 0.5$.

5.2. Evaluation on Existing Evaluation Metrics

In this section, we evaluate the proposed embedding learning method using existing instance segmentation metrics. We compare our method with SGPN [22] and DFD [21], which have been found to exhibit the highest accuracy for this task. The scores for these two methods are reported in [21].

Following [21], we chose the proposal recall [9] as the evaluation metric and used PointNet as a feature extractor for a fair comparison. The proposal recall is calculated as follows: first, for each GT object, we select the predicted object with the highest intersection over union (IoU) regardless of the category of the object and consider the output as a true positive when the IoU is higher than a certain threshold (we chose a value of 0.5). The ratio of the number of true positives is then calculated with respect to the number of GTs. Because the number of objects for each category is unbalanced, we evaluated both the mean of the proposal recall of 12 categories, except the 'clutter' class (mean), and the overall proposal recall regardless of the categories (total). Note that the overall proposal recall (total) can be high even if the model overfits some of the categories with many instances and ignores the categories with fewer instances, and thus, it may not be reliable. However, DFD, which does not use category information for training, cannot solve the imbalance problem between categories, and thus, it was necessary to add the total proposal recall.

Moreover, to validate our argument that instance segmentation on small regions is a substantially semantic segmentation because there is often only one instance in the region, we also evaluated the result obtained using semantic segmentation model (SemSeg), which never splits objects belonging to the same category. We also report the score obtained when using PointNet++ (PN++) instead of PN as the feature extractor; however, we do not compare PN++ with PN as it would not make for a fair comparison.

Table 1 shows a comparison of our methods with existing methods as well as the obtained semantic segmentation results. We can see that our method with PointNet (PN) outperforms existing methods in terms of the mean proposal recall by a large margin, and the use of PointNet++ leads to a considerably better score. As described earlier, DFD achieves a high total score; however, its mean score is low, which means that the model ignores categories with fewer instances. In addition, for some categories such as a ceiling, floor, and beam, a mere semantic segmentation result achieves a very high score because there is essentially only one instance of such categories. This result supports

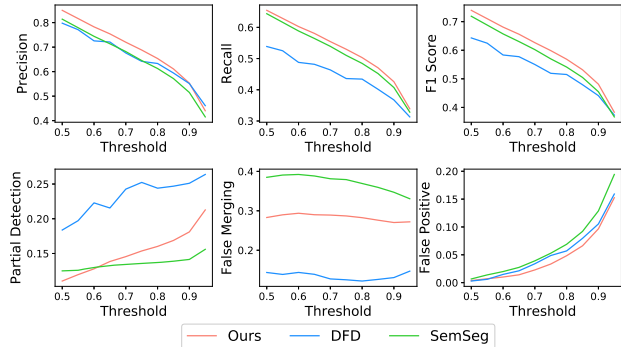


Figure 4. Comparison of our method and [21] using the proposed evaluation metrics

our argument that semantic segmentation results affect the instance segmentation performance, and an evaluation in small regions using existing metrics is inappropriate.

5.3. Evaluation of the Proposed Evaluation Metrics

We then evaluated our method with DFD, which outperforms SGPN, using the proposed evaluation metrics. Some predicted objects consist of a very small number of points. Because such predicted objects are often false positives, we set a threshold and use the predicted objects with a number of points larger than the threshold as the targets for evaluation. There is a trade-off between *precision* and *recall*, which we introduced in Section 4. As the threshold decreases, *precision* decreases while *recall* increases. We fix $t = 0.75$ for the IoS defined by Equation 9 and search for the threshold that can obtain the highest F-score. As a result, we use a threshold of 150 for the DFD and 35 for our method. The DFD shows a larger threshold, which implies that it outputs noisy small predicted objects that are false positives.

Furthermore, as discussed in Section 5.2, the output of the semantic segmentation model (SemSeg), which never splits objects belonging to the same category, achieves high scores using the existing metrics for instance segmentation. This occurs in some categories in which multiple objects seldom exist in a single subregion. We evaluated the results of semantic segmentation using our evaluation metrics to demonstrate whether this problem was solved.

We plot the results when varying t of IoS from 0.5 to 0.95 in Figure 4. Note that precision, recall, and f1-score evaluate performance, whereas partial detection, false merging, and false positive represent types of mistakes. We can see that, although the semantic segmentation model shows a high score for the existing metrics, specifically for *recall*, the false merging score is quite high. This is because semantic segmentation outputs one prediction per category at most; this is why the partial detection of the semantic segmentation model is low. To the best of our knowledge,

Table 1. Comparison with existing methods ([21, 22]) by proposal recall [%]

method	ceiling	floor	wall	beam	column	window	door	table	chair	sofa	book-case	board	mean	total
SGPN [22]	67.0	71.4	66.8	54.5	45.4	51.2	69.9	63.1	67.6	64.0	54.4	60.5	61.3	64.7
DFD [21]	95.4	99.2	77.3	48.0	39.2	68.2	49.2	56.0	53.2	35.3	31.6	42.2	57.9	69.1
SemSeg	95.8	95.2	61.7	89.3	50.0	76.6	65.7	60.2	44.1	16.6	40.6	45.6	61.8	59.4
Ours (PN)	95.9	94.6	64.5	89.3	61.3	83.3	75.0	64.0	55.0	70.8	55.6	50.8	71.7	68.8
Ours (PN++)	96.2	94.1	65.6	87.8	71.4	81.0	82.6	68.8	60.9	68.4	63.2	67.2	75.6	72.7

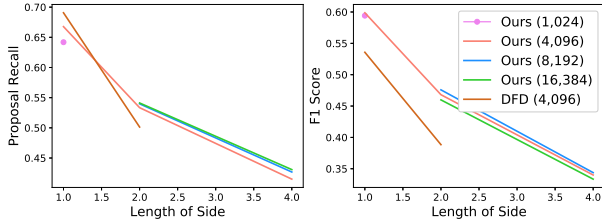


Figure 5. Effect of region size and number of points

this fact is revealed by the proposed metrics for the first time. We can also observe fine patterns and the property of instance segmentation errors when eliminating semantic errors, which cannot be obtained using existing evaluation metrics. For example, for DFD, most errors arise from partial detection whereas false merging is the dominant cause in our method. Such information is useful not only for analyzing and improving the model but also for applying an ensemble of models when considering the characteristics of each model.

5.4. Effect of Region Size and Number of Points

In this section, we analyze the effects of the input region size and the number of points. We varied the number of points from 4,096 to 16,384 and the size of the regions from 1 m square to 4 m square. Note that the instance segmentation results can be also affected by the density of the points. The settings with 1,024 points and 1 m square, 4,096 points and 2 m square, and 16,384 points with 4 m square have the same density.

Figure 5 shows the proposal recall and the F-score values for each setting. Even if the density is the same, the score of the instance segmentation decreases with the size of the region. As discussed in Section 1, when the input region is small, we do not need to distinguish the objects because the number of different objects of the same category is small. Therefore, the task becomes difficult as the input region increases and the apparent score decreases. In particular, the F1 score of DFD is significantly decreased compared with the proposal recall; one reason for this is that the proposal recall does not penalize false positives and cannot reveal the weakness of the DFD model that the instance segmentation results are quite noisy. Moreover, the figure shows that the

density of points does not considerably affect the instance segmentation performance.

Although we can apply instance segmentation on entire scenes by first applying instance segmentation on each subregion and then integrating the results through a post-processing technique, this approach has a significantly high computational complexity because we need to repeat instance segmentation on each subregion, making it unsuitable for practical use. Moreover, choosing an appropriate subregion size is difficult, and an integration procedure can add noise to the final result. Therefore, it is desirable to use as large a region as possible for the input. However, this figure implies that the instance segmentation task becomes significantly difficult when the input size is large, this difficulty has not been adequately investigated in existing works. Handling large regions, such as an entire scene, is a challenging task; however, it is of great importance in the application of instance segmentation.

6. Conclusion

We proposed a new method for instance segmentation on 3D point clouds. Our memory efficient loss function learns mapping to the embedding space, where the embeddings form clusters for each object. We experimentally showed that our method outperforms existing methods. Our method can handle a large number of points and performs well even when consuming large regions. Moreover, we claimed and experimentally demonstrated that existing metrics are not suitable for evaluating instance segmentation because they are considerably affected by the input size of the misclassification. We proposed novel metrics that are unaffected by such external conditions and can aid in evaluating instance segmentation performances correctly. Using the proposed metrics, we not only evaluated the instance segmentation task without being affected by external conditions but also analyzed the types of errors in an instance segmentation task for each method.

7. Acknowledgement

This work was partially supported by JST CREST Grant Number JPMJCR1403, and partially supported by JSPS KAKENHI Grant Number JP19H01115.

References

- [1] I. Armeni, O. Sener, A. R. Zamir, H. Jiang, I. Brilakis, M. Fischer, and S. 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](#), [6](#)
- [2] C. Choi, Y. Taguchi, O. Tuzel, M.-Y. Liu, and S. Ramalingam. Voting-based pose estimation for robotic assembly using a 3d sensor. In *2012 IEEE International Conference on Robotics and Automation*, pages 1724–1731. IEEE, 2012. [2](#)
- [3] S. Chopra, R. Hadsell, and Y. LeCun. Learning a similarity metric discriminatively, with application to face verification. In *null*, pages 539–546. IEEE, 2005. [3](#)
- [4] J. Dai, K. He, and J. Sun. Instance-aware semantic segmentation via multi-task network cascades. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3150–3158, 2016. [3](#)
- [5] B. De Brabandere, D. Neven, and L. Van Gool. Semantic instance segmentation for autonomous driving. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 7–9, 2017. [3](#), [4](#)
- [6] B. Drost, M. Ulrich, N. Navab, and S. Ilic. Model globally, match locally: Efficient and robust 3d object recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 998–1005. Ieee, 2010. [2](#)
- [7] M. Ester, H.-P. Kriegel, J. Sander, X. Xu, et al. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Kdd*, volume 96, pages 226–231, 1996. [5](#)
- [8] K. He, G. Gkioxari, P. Dollár, and R. Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. [3](#)
- [9] J. Hosang, R. Benenson, P. Dollár, and B. Schiele. What makes for effective detection proposals? *IEEE transactions on pattern analysis and machine intelligence*, 38(4):814–830, 2016. [7](#)
- [10] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. [6](#)
- [11] M. Koestinger, M. Hirzer, P. Wohlhart, P. M. Roth, and H. Bischof. Large scale metric learning from equivalence constraints. In *2012 IEEE conference on computer vision and pattern recognition*, pages 2288–2295. IEEE, 2012. [3](#)
- [12] S. Kong and C. C. Fowlkes. Recurrent pixel embedding for instance grouping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9018–9028, 2018. [3](#), [4](#)
- [13] Y. Li, H. Qi, J. Dai, X. Ji, and Y. Wei. Fully convolutional instance-aware semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2359–2367, 2017. [3](#)
- [14] D. Novotny, S. Albanie, D. Larlus, and A. Vedaldi. Semi-convolutional operators for instance segmentation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 86–102, 2018. [3](#)
- [15] C. R. Qi, H. Su, K. Mo, and L. 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. [1](#), [2](#), [6](#)
- [16] C. R. Qi, L. Yi, H. Su, and L. J. Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *Advances in Neural Information Processing Systems*, pages 5099–5108, 2017. [1](#), [2](#), [6](#)
- [17] S. Ren, K. He, R. Girshick, and J. Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pages 91–99, 2015. [3](#)
- [18] R. B. Rusu, N. Blodow, and M. Beetz. Fast point feature histograms (fpfh) for 3d registration. In *2009 IEEE International Conference on Robotics and Automation*, pages 3212–3217. IEEE, 2009. [2](#)
- [19] R. B. Rusu, N. Blodow, Z. C. Marton, and M. Beetz. Aligning point cloud views using persistent feature histograms. In *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*, pages 3384–3391. IEEE, 2008. [2](#)
- [20] F. Schroff, D. Kalenichenko, and J. Philbin. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 815–823, 2015. [3](#)
- [21] M. Sung, H. Su, R. Yu, and L. J. Guibas. Deep functional dictionaries: Learning consistent semantic structures on 3d models from functions. In *Advances in Neural Information Processing Systems*, pages 483–493, 2018. [1](#), [3](#), [7](#), [8](#)
- [22] W. Wang, R. Yu, Q. Huang, and U. Neumann. Sgpn: Similarity group proposal network for 3d point cloud instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2569–2578, 2018. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [23] K. Q. Weinberger and L. K. Saul. Distance metric learning for large margin nearest neighbor classification. *Journal of Machine Learning Research*, 10(Feb):207–244, 2009. [3](#)