PointInst3D: Segmenting 3D Instances by Points

Tong He¹, Wei Yin¹, Chunhua Shen², Anton van den Hengel¹

¹ The University of Adelaide ² Zhejiang University

Abstract. The current state-of-the-art methods in 3D instance segmentation typically involve a clustering step, despite the tendency towards heuristics, greedy algorithms, and a lack of robustness to the changes in data statistics. In contrast, we propose a fully-convolutional 3D point cloud instance segmentation method that works in a per-point prediction fashion. In doing so it avoids the challenges that clustering-based methods face: introducing dependencies among different tasks of the model. We find the key to its success is assigning a suitable target to each sampled point. Instead of the commonly used static or distance-based assignment strategies, we propose to use an Optimal Transport approach to optimally assign target masks to the sampled points according to the dynamic matching costs. Our approach achieves promising results on both ScanNet and S3DIS benchmarks. The proposed approach removes intertask dependencies and thus represents a simpler and more flexible 3D instance segmentation framework than other competing methods, while achieving improved segmentation accuracy.

Keywords: Clustering-free, Dependency-free, 3D instance segmentation, Dynamic target assignment, Optimal Transport

1 Introduction

3D instance segmentation describes the problem of identifying a set of instances that explain the locations of a set of sampled 3D points. It is an important step in a host of 3D scene-understanding challenges, including autonomous driving, robotics, remote sensing, and augmented reality. Despite this fact, the performance of 3D instance segmentation lags that of 2D instance segmentation, not least due to the additional challenges of 3D representation, and variable density of points.

Most of the top-performing 3D instance segmentation approaches [17, 7, 14, 21, 4, 10] involve a clustering step. Despite their great success, clustering-based methods have their drawbacks: they are susceptible to the performance of the clustering approach itself, and its integration, due to either (1) error accumulation caused by the inter-task dependencies [17, 14, 4] or (2) non-differentiable processing steps [21, 10]. For example, in PointGroup [17], instance proposals are generated by searching homogenous clusters that have identical semantic predictions and close centroid predictions. However, the introduced dependencies on both tasks make the results sensitive to the heuristics values chosen.



Fig. 1. A comparison of the instance segmentation results achieved by DyCo3D [14] and our method. The subpar performance of instance segmentation for DyCo3D [14] is caused by the dependency on semantic segmentation. Our method addresses the task in a per-point prediction fashion and removes the dependencies between different tasks of the model. Thus, it is free from the error accumulation introduced by the intermediate tasks. Best viewed in colors.

DyCo3D [14] addressed the issue by encoding instances as continuous functions. But the accuracy is still constrained by the semantic-conditioned convolution. As a result, it can be impossible to recover from errors in intermediate stages, particularly given that many methods greedily associate points with objects (which leaves them particularly susceptible to early clustering errors). Even with careful design, because of the diversity in the scales of instances, and the unbalanced distribution of semantic categories, the performance of these intermediate tasks is often far from satisfactory. This typically leads to fragmentation and merging of instances, as shown in Fig. 1.

In this paper, we remove the clustering step and the dependencies within the model and propose a much simpler pipeline working in a per-point prediction fashion. Every sampled point will generate a set of instance-related convolutional parameters, which are further applied for decoding the binary masks of the corresponding instances. However, building such a clustering-free and dependency-free pipeline is non-trivial. For example, removing the clustering step and conditional convolution in DyCo3D causes mAP to drop by more than 8% and 6%, respectively. We conduct comprehensive experiments and find the reason for the huge drop in performance is the ambiguity of the targets for the sampled points. In 2D instance segmentation and object detection, the center prior, which assumes the predictions from the central areas of an instance are more likely to provide accurate results, offers a guideline to select well-behaved samples [32, 31, 8]. This distance-based prior is hard to apply in 3D, however, as the distribution of high-quality samples in 3D point clouds is irregular and unpredictable. The fact that objects can be arbitrarily close together in real 3D scenes adds additional complexity. Thus, the resulting ambiguity in point-instance associations can contaminate the training process and impact final performance. Instead of applying a static or widely used distance-based strategy, we propose to optimally assign instances to samples via an Optimal Transport (OT) solution. It is defined in terms of a set of suppliers and demanders, and the costs of transportation between them. We thus associated a demander with each instance prediction of

the sampled point, and a supplier with each potential instance ground truth. The cost of transport reflects the affinity between each pair thereof. The OT algorithm identifies the optimal strategy by which to supply the needs of each demander, given the cost of transport from each supplier. The points will then be associated with the target corresponding to the demander to which it has allocated the greatest proportion of its supply. The costs of transporting are determined by the Dice Coefficient, which is updated dynamically based on the per-point predictions. The OT solution not only minimizes the labor for heuristics tuning but allows it to make use of the sophisticated tools that have been developed for solving such problems. In particular, it can be efficiently solved by the off-the-shelf Sinkhorn-Knopp Iteration algorithm [5] with limited computation in training.

To summarise, our contributions are listed as follows.

- We propose a clustering-free framework for 3D instance segmentation, working in a per-point prediction fashion. In doing so it removes the dependencies among different tasks and thus avoids error accumulation from the intermediate tasks.
- For the first time, we address the target assignment problem for 3D instance segmentation, which has been overlooked in the 3D community. Our proposed Optimal Transport solution is free from heuristics with improved accuracy.
- We achieve promising results on both ScanNet and S3DIS, with a much simpler pipeline.

2 Related Work

Target Assignment in 2D Images. The problem of associating candidates to targets arises commonly in 2D object detection. Anchor-based detectors [30, 23, 22] apply a hard threshold to an intersection-over-union measure to divide positive and negative samples. This approach can also be found in many other methods [3, 11]. Anchor-free detectors [32, 42, 18] have drawn increasing attention due to their simplicity. These methods observe that samples around the center of objects are more likely to provide accurate predictions. Inspired by this center prior, some methods [31, 33, 18, 38] introduce a classifier by treating these central regions as positive samples. ATSS [39], in contrast, is adaptive in that it sets a dynamic threshold according to the statistics of the set of closest anchors. Free-Anchor [40] frames detector training as a maximum likelihood estimation (MLE) procedure and proposes a learning-based matching mechanism. Notably, OTA [8] formulates the task of label assigning as Optimal Transport problem.

Instance Segmentation on 3D Point Cloud. The task of instance segmentation in the 3D domain is complicated by the irregularity and sparsity of the point cloud. Unlike instance segmentation of images, in which top-down methods are the state-of-the-art, the leader board in instance segmentation of 3D point clouds has been dominated by bottom-up approaches due to unsatisfactory 3D



Fig. 2. The framework of our proposed method. The 'inst head' is designed to generate instance masks by applying dynamic convolution. K points are sampled via the farthest point sampling strategy. Each sampled point is responsible for one specific instance mask or background. The targets are calibrated dynamically via an Optimal Transport solution, which takes as input the mask prediction from the auxiliary head and outputs the calibrated ground truth for the main instance head. The targets for the auxiliary instance prediction 'pred_a' are consistent with the instance label of the sampled points.

detection results. SGPN [34], for instance, predicts an $N \times N$ matrix to measure the probability of each pair of points coming from the same instance, where Nis the number of total points. ASIS [35] applies a discriminative loss function from [2] to learn point-wise embeddings. The mean-shift algorithm is used to cluster points into instances. Many works (e.g. [41, 13, 12, 27]) follow this metricbased pipeline. However, these methods often suffer from low accuracy and poor generalization ability due to their reliance on pre-defined hyper-parameters and complex post-processing steps. Interestingly, PointGroup [17] exploits the voids between instances for segmentation. Both original and center-shifted coordinates are applied to search nearby points that have identical semantic categories. The authors of DvCo3D [14, 15] addressed the sensitivity of clustering methods to the grouping radius using dynamic convolution. Instead of treating clusters as individual instance proposals, DyCo3D utilized them to generate instance-related convolutional parameters for decoding masks of instances. Chen et al. proposed HAIS [4], which is also a clustering-based architecture. It addressed the problem of the over- and under-segmentation of PointGroup [17] by deploying an intrainstance filtering sub-network and adapting the grouping radius according to the size of clusters. SSTN [21] builds a semantic tree with superpoints [20] being the leaves of the tree. The instance proposals can be obtained when a non-splitting decision is made at the intermediate tree node. A scoring module is introduced to refine the instance masks.

3 Methods

4

T. He et al.

The pipeline of the proposed method is illustrated in Fig. 2, which is built upon a sparse convolution backbone [9]. It maintains a UNet-like structure and takes as input the coordinates and features, which have a shape of $N \times 3$ and $N \times I$, respectively. N is the total number of input points and I is the dimension of input features. There is one output branch of mask features, which is used to decode binary masks of instances. It is denoted as $F_m \in \mathbb{R}^{N \times d'}$, where d' is the dimension of the mask features. Inspired by DyCo3D [14], we propose to encode instance-related knowledge into a set of convolutional parameters and decode the corresponding masks with several 1×1 convolutions. Different from DyCo3D, which requires a greedy clustering algorithm and a conditioned decoding step, our proposed method, on the other hand, removes the clustering step and the dependencies among different tasks, simplifying the network in a point-wise prediction pipeline.

3.1 Preliminary on DyCo3D

DyCo3D [14] has three output branches: semantic segmentation, centroid offset prediction, and mask features. The breadth-first-searching algorithm is used to find out the homogenous points that have identical semantic labels and close centroid predictions. Each cluster is sent to the instance head and generates a set of convolution parameters for decoding the mask of the corresponding instance. Formally, the mask \hat{M}_k predicted by the k-th cluster can be formulated as:

$$\hat{M}_{k} = Conv_{1x1}(feature, weight) = Conv_{1x1}(F_{m} \oplus C_{rel}^{k}, mlp(G(P_{s}, P_{c})_{k})) \odot \mathbb{1}(P_{s} = s_{k})$$
(1)

The input features to convolution contains two parts: F_m and C_{rel}^k . F_m is the mask features shared by all instances. $C_{\text{rel}}^k \in \mathbb{R}^{N \times 3}$ is the instance-specific relative coordinates, which are obtained by computing the difference between the center of the k-th cluster and all input points. F_m and C_{rel}^k are concatenated $(`\oplus`)$ along the feature dimension. The convolutional weights are dynamically generated by an mlp layer, whose input is the feature of the k-th cluster. The clustering algorithm $G(\cdot)$ takes the semantic prediction $P_s \in \mathbb{R}^N$ and centroid prediction $P_c \in \mathbb{R}^N$ as input and finds out a set of homogenous clusters. The k-th cluster is denoted as $G(\cdot)_k$. Besides, the dynamic convolution in DyCo3D is conditioned on the results of semantic segmentation. For example, DyCo3D can only discriminate one specific 'Chair' instance from all points that are semantically categorized as 'Chair', instead of the whole point set. It is implemented by an element-wise production (` \odot ') with a binary mask (' $\mathbb{1}(\cdot)$ '). s_k is the semantic label of the k-th cluster. Finally, the target mask for \hat{M}_k is decided by the instance label of the k-th cluster. More details can be found in [14].

3.2 Proposed Method

Although promising, DyCo3D [14] involves a grouping step to get the instancerelated clusters, depending on the accuracy of semantic segmentation and offset prediction. Besides, the conditional convolution also forces the instance decoding to rely on the results of semantic segmentation. These inter-task dependencies cause error accumulation and lead to sub-par performance (See Fig. 1). In



Fig. 3. The left image is an indoor scene with three instances of 'Chair'. The right image is the quality of instance predictions by each point. The brighter the color, the more accurate the mask predicted by the point. Different from the 2D image, the distribution of the positive samples in 3D point cloud is irregular, making it hard to learn a criterion to select informative samples for each instance. In addition, the ambiguity of target assignment is widespread in the 3D scenes. Some samples in instance 'C' show high-quality predictions of the instance 'D'. Best viewed in color.

this paper, we propose a clustering-free and dependency-free framework in a per-point prediction fashion. Total K points are selected via the farthest point sampling strategy. The instance head takes as input both the mask feature F_m and point-wise feature f_b^k . The k-th mask \hat{M}_k predicted by the instance head can be formulated as:

$$M_{k} = Conv_{1x1}(feature, weight)$$

= $Conv_{1x1}(F_{m} \oplus C_{rel}^{k}, mlp(f_{b}^{k}))$ (2)

where f_b^k is the feature of the k-th sampled point from output of the backbone. $C_{\text{rel}}^k \in \mathbb{R}^{N \times 3}$ is the relative position embedding, obtained by computing the difference between the coordinate of the k-th point and all other points.

However, building such a simplified pipeline is non-trivial. Removing the clustering step and conditional convolution causes the mAP of DyCo3D to drop dramatically.

Observation To find out the reasons that cause the failure of this point-wise prediction pipeline, we visualize the quality of masks predicted by each point (according to Eq. 2). For training, the target mask for each point is consistent with its instance label. As shown in Fig. 3, the distribution of high-quality samples is irregular and can be influenced by many factors: (1) disconnection, (2) distance to the instance center, and (3) spatial relationships with other objects. Besides, the fact that objects can be arbitrarily close together in real 3D scenes adds additional complexity. As illustrated in Fig. 3(c,d), the poorly behaved samples in 'chair c' can accurately predict the mask of the 'desk'. Such ambiguity introduced by the static assigning strategy contaminates the training process, leading to inferior performance.

Target Assignment Although the task of target assignment has shown its significance in 2D object detection and instance segmentation [40, 39, 8], to the best of our knowledge, there is very little research in the 3D domain. One of the most straightforward ways is to define a criterion to select a set of informative samples for each instance. For example, thanks to the center prior [32], many approaches [31, 42, 18, 38] in the 2D domain treat the central areas of the instance as positive candidates. However, such regularity is hard to define for the 3D point cloud, as shown in Fig. 3. Quantitative results can be found in Tab. 1.

Instead of applying a static strategy or learning an indicative metric, we propose to assign a suitable target for each sample based on its prediction. A background mask (*i.e.* all zeros) is added to the target set to address the poorly-behaved points.

Optimal Transport Solution Given K sampled points (via farthest point sampling) and their corresponding mask predictions $\{\hat{M}_k\}^K$ (using Eq. 2), the goal of target assignment is to find a suitable target for each prediction in training. There are T+1 targets in total, including T instance masks $\{M_t\}^T$ and one background mask M_{T+1} (zero mask). Inspired by [8], we formulate the task as an Optimal Transport problem, which seeks a plan by transporting the 'goods' from suppliers (*i.e.* Ground Truth and Background Mask) to demanders (*i.e.* predictions of the sampled points) at a minimal transportation cost.

Supposing the t-th target has μ_t unit of goods and each prediction needs one unit of goods, we denote the cost for transporting one unit of goods from the t-th target to the k-th prediction as C_{tk} . By applying Optimal Transport, the task of the target assignment can be written as:

$$\boldsymbol{U}^{*} = \operatorname*{arg\,min}_{\boldsymbol{U} \in \mathbb{R}^{(T+1) \times K}_{+}} \sum_{t,k} C_{tk} U_{tk}$$

s.t. $\boldsymbol{U} \boldsymbol{1}_{K} = \boldsymbol{\mu}_{T+1}, \ \boldsymbol{U}^{\mathsf{T}} \boldsymbol{1}_{T+1} = \boldsymbol{1}_{K},$ (3)

where U^* is the optimal assignment plan, U_{tk} is the amount of labels transported from the *t*-th target to the *k*-th prediction. μ_{T+1} is the label vector for all T+1targets. The transportation cost C_{tk} is defined as:

$$C_{tk} = \begin{cases} \mathcal{L}_{\text{dice}}(M_t, \hat{M}_k) & t \leq T\\ \mathcal{L}_{\text{dice}}(1 - M_t, 1 - \hat{M}_k) & t = T + 1 \end{cases}$$
(4)

where $\mathcal{L}_{\text{dice}}$ denotes the dice loss. To calculate the cost between the background target and the prediction, we use $1 - M_t$ and $1 - \hat{M}_k$ for a numerically stable training. The restriction in Eq. 3 describes that (1) the total supply must be equal to the total demand and (2) the goods demand for each prediction is 1 (*i.e.* each prediction needs one target mask). Besides, the label vector $\boldsymbol{\mu}_{T+1}$, indicating the total amount of goods held by each target, is updated by:

$$\mu_t = \begin{cases} int(\sum_k IoU(\hat{M}_k, M_t)) & t \le T \\ K - \sum_{i=1}^T \mu_i & t = T+1 \end{cases}$$
(5)

8 T. He et al.

where μ_{T+1} refers to the target amount maintained in the background target and $int(\cdot)$ is the rounding operation. According to Eq. 5, the amount of supplied goods for each target is dynamically decided, depending on its IoU with each prediction. Due to the restriction in Eq. 3, we set μ_{T+1} equal to $K - \sum_{t=1}^{T}$. The efficient Sinkhorn-Knopp algorithm [5] allows it to obtain U^* with limited computation overhead. After getting the optimal assignment U^* , the calibrated targets for the K sampled points can be determined by assigning each point with the target that transports the largest amount of goods to it.

Compared with [8], the number of the demanders is much fewer. Thus, the minimum supply of each target can be zero in training. Doing so may make the model fall into a trivial solution when K is small: all predictions are zero masks and assigned to the background target due to the lowest transportation cost in Eq. 4. To this end, we propose a simple yet effective way by introducing an auxiliary instance head, whose targets are consistent with the instance labels of the sampled points. We use the predictions from this auxiliary head to calculate the cost matrix in Eq. 4. The dynamically calibrated targets are used for the main instance head. To alleviate the impact of the wrongly assigned samples in the auxiliary head, the loss weight for this auxiliary task is decreasing in training.

3.3 Training

To summarize, the loss function includes two terms for training, including the auxiliary loss term \mathcal{L}_{a} and the main task loss term \mathcal{L}_{m} :

$$\mathcal{L} = w_a \sum_{k=1}^{K} \mathcal{L}_{\mathbf{a}}(M_{\mathbf{k}}^a, \hat{M}_{\mathbf{k}}^a) + \sum_{k=1}^{K} \mathcal{L}_{\mathbf{m}}(M_{\mathbf{k}}^m, \hat{M}_{\mathbf{k}}^m)$$
(6)

where $\{M_k^a\}^K \in \{0,1\}^{K \times N}$ is the ground truth masks for the K predictions. These targets are static and decided by the instance labels of the K sampled points. $\{M_k^m\}^K \in \{0,1\}^{K \times N}$ is the set of the calibrated targets for the main instance head. $\{\hat{M}_k^a\}^K$ and $\{\hat{M}_k^m\}^K$ are the predictions from auxiliary and main instance heads, respectively. w_a is the loss weight for the auxiliary task. We set w_a to 1.0 with a decaying rate of 0.99. Early in the training phase, the static targets for the auxiliary task play a significant role in stabilizing the learning process. The loss of the main task is involved until the end of a warming-up period, which is set to 6k steps. So far, we have obtained a set of binary masks. There are many ways to obtain the corresponding categories, for example, adding a classification head for each mask proposal. In our paper, we implement it by simply introducing a semantic branch. The category c_k of the k-th instance is the majority of the semantic predictions within the foreground mask of \hat{M}_k^m . Instances with a number of points less than 50 are ignored.

Method	CP	DT	AUX	mAP	AP@50	AP@25
Baseline				33.7	52.4	65.0
	\checkmark			34.1	53.2	65.4
		\checkmark		36.8	54.8	65.9
			\checkmark	36.5	54.3	65.7
Ours		\checkmark	\checkmark	39.6	59.2	70.4

Table 1. Component-wise analysis on ScanNetV2 validation set. **CP**: the center prior tailored for 3D point cloud. **DT**: dynamic targets assignment using Optimal Transportation. **AUX**: the auxiliary loss used in Eq. 6.

4 Experiments

We conduct comprehensive experiments on two standard benchmarks to validate the effectiveness of our proposed method: ScanNet [6] and Stanford 3D Indoor Semantic Dataset (S3DIS) [1].

4.1 Datasets

ScanNet has 1613 scans in total, which are divided into training, validation, and testing with a size of 1201, 312, and 100, respectively. The task of instance segmentation is evaluated on 18 classes. Following [14], we report the results on the validation set for ablation study and submit the results on the testing set to the official evaluation server. The evaluation metrics are mAP (mean average precision) and AP@50.

S3DIS contains more than 270 scans, which are collected on 6 large indoor areas. It has 13 categories for instance segmentation. Following the previous method [35], the evaluation metrics include: mean coverage (mCov), mean weighed coverage (mWCov), mean precision (mPrec), and mean recall (mRec).

4.2 Implementation Details

The backbone model we use is from [9], which maintains a symmetrical UNet structure. It has 7 blocks in total and the scalability of the model is controlled by the channels of the block. To prove the generalization capability of our proposed method, we report the performance with both small and large backbones, denoted as **Ours-S** and **Ours-L**, respectively. The small model has a channel unit of 16, while the large model is 32. The default dimension of the mask features is 16 and 32, respectively.

For each input scan, we concatenate the coordinates and RGB values as the input features. All experiments are trained for 60K iteration with 4 GPUS. The batch size for each GPU is 3. The learning rate is set to 1e-3 and follows a polynomial decay policy. In testing, the computation related to the auxiliary

10 T. He et al.

head is ignored. Only Non-Maximum-Suppression (NMS) is required to remove the redundant mask predictions for inference, with a threshold of 0.3.

4.3 Ablation Studies

In this section, we verify the effectiveness of the key components in our proposed method. For a fair comparison, all experiments are conducted on the validation set of ScanNet [6] with the smaller model.

Baseline. We build a strong baseline by tailoring CondInst [31] for the 3D point cloud. It works in a per-point prediction fashion and each sampled point has a static target, which is consistent with the corresponding instance label. As shown in Tab. 1, our method achieves 33.7% 52.4%, and 65.0% in terms of mAP, AP@50, and AP@25, respectively. With a larger number of sampled points and longer iterations, our baseline model surpasses the implementation of DyCo3D [14] by a large margin.

Center Prior in 3D. To demonstrate the difficulty of selecting informative samples in 3D, we tailor the center prior [32] to 3D point cloud. As points are collected from the surface of the objects, centers of 3D instances are likely to be in empty space. To this end, we first predict the offset between each point and the center of the corresponding object. If the distance between the centershifted point and the ground truth is close (≤ 0.3 m), the point is regarded as positive and responsible for the instance. If the distance is larger than 0.6m, the point is defined as negative. Other points are ignored for training. As presented in Tab. 1, selecting positive samples based on the 3D center prior only boosts 0.4% and 0.8% in terms of mAP and mAP@50, respectively. The incremental improvement demonstrates the difficulty of selecting informative samples in 3D. In contrast, we propose to apply a dynamic strategy, by which the target for each candidate is determined based on its prediction.



Fig. 4. Ablation study on the number of the sampling point.

Dynamic Targets. To show the effectiveness of the dynamic strategy, we implement an experiment by removing the auxiliary head. As the predictions are

3D Object Det	ection
ScanNetV2	$\mathrm{AP}@50\%$
MRCNN 2D-3D [11]	10.5
F-PointNet [29]	10.8
GSPN [37]	17.7
3D-SIS [16]	22.5
VoteNet [28]	33.5
PointGroup [17]	42.3
DyCo3D [14]	45.3
3D-MPA [7]	49.2
Ours	51.0

Table 2. The performance of 3D object detection, tested on ScanNet validation set.AP@50 is reported.

basically random guesses in the early stage of the training, we first warm up the model for 12k iterations with a static assignment to avoid the trivial solution. In the remaining steps, targets are calibrated by the Optimal Solution. As shown in Tab. 1, our approach boosts the performance of the baseline model by 3.1%, 2.4%, and 0.9%, in terms of mAP, AP@50, and AP@25, respectively.

Auxiliary Supervision. As illustrated in Fig. 2, we propose to regularize the intermediate layers by introducing an auxiliary instance head for decoding the instance masks. The targets for this task are static and consistent with the instance labels. Besides, as the generated parameters are convolving with the whole point set, large context and instance-related knowledge are encoded in the point-wise features. To remove the influence of the dynamic assignment, both auxiliary and the main task are applying a static assignment strategy. As shown in Tab. 1, the auxiliary supervision brings 2.8%, 1.9%, and 0.7% improvement in terms of mAP, mAP@50, and mAP@25, respectively. In addition to the encoded large context, the predicted instance masks are also applied to the Optimal Solution to obtain calibrated targets. Combining with the proposed dynamic assignment strategy, it further boosts mAP, AP@50, and AP@25 for 3.1%, 4.4%, and 4.5%, respectively, achieving 39.6% in terms of mAP with a small backbone.

Analysis on Efficiency. Our method takes the whole scan as input, without complex pre-processing steps. Similar to DyCo3D [14], the instance head is implemented in parallel. To make a fair comparison, we set K equal to the average number of clusters in DyCo3D. Using the same GPU, the mAP of our proposed method is 1.8% higher than DyCo3D and the inference time is 26% faster than DyCo3D.

Number of Random Selected Samples. We randomly select K points, each of which is responsible for one specific instance or the background (all zeros). In this part, we study the influence of the value of K. The performance is shown in Fig. 4. We set K to 256 for its highest mAP.



Fig. 5. Comparison with the results of DyCo3D [14] and PointGroup [17]. The ellipses highlight specific over-segmentation/joint regions. Instances are presented with different colors. Best viewed in color.

The Dimension of the Mask Feature. The mask feature contains the knowledge of instances. We conduct experiments to show the influence of different dimensions of the mask feature. We find the fluctuation of the performance is relatively small when the dimension is greater than 8, showing the strong robustness of our method to the variation of d'. We set d' to 16 in our experiments.

4.4 Comparison with State-of-the-art Methods

We compare our method with other state-of-the-art methods on both S3DIS and ScanNet datasets.

3D Detection. Following [14, 7], we evaluate the performance of 3D detection on the ScanNet dataset. The results are obtained by fitting axis-aligned bounding boxes for predicted masks, as presented in Tab. 2. Our method surpasses DyCo3D [14] and 3D-MPA [7] by 4.8% and 1.8% in terms of mAP, respectively. The promising performance demonstrates the compactness of the segmentation results.

Instance Segmentation on S3DIS. Following the evaluation protocols that are widely applied in the previous approaches, experiments are carried out on both Area-5 and 6-Fold cross-validation. As shown in Tab. 3, our proposed method achieves the highest performance and surpasses previous methods with a

Method	mCov	mWCov	mPrec	mRec
Tes	st on A	rea 5		
SGPN'18 [34]	32.7	35.5	36.0	28.7
ASIS'19 [35]	44.6	47.8	55.3	42.4
3D-BoNet'19 [36]	-	-	57.5	40.2
3D-MPA'20 [7]	-	-	63.1	58.0
MPNet'20 [12]	50.1	53.2	62.5	49.0
InsEmb'20 [13]	49.9	53.2	61.3	48.5
PointGroup'20 [17]	-	-	61.9	62.1
DyCo3D'21 [14]	63.5	64.6	64.3	64.2
HAIS'21 [4]	64.3	66.0	71.1	65.0
SSTNet'21 [21]	-	-	65.5	64.2
Ours	64.3	65.3	73.1	65.2
Te	st on 6	6-fold		
SGPN'18 [34]	37.9	40.8	31.2	38.2
MT-PNet'19 [27]	_	-	24.9	-
MV-CRF'19 [27]	-	-	36.3	-
ASIS'19 [35]	51.2	55.1	63.6	47.5
3D-BoNet'19 [36]	-	-	65.6	47.6
PartNet'19 [25]	-	-	56.4	43.4
InsEmb'20[13]	54.5	58.0	67.2	51.8
MPNet'20 [12]	55.8	59.7	68.4	53.7
PointGroup'20 [17]	-	-	69.6	69.2
3D-MPA'20 [7]	-	-	66.7	64.1
HAIS'21 [4]	67.0	70.4	73.2	69.4
SSTNet'21 [21]	-	-	73.5	73.4
Ours	71 5	7/1	76 /	74 0

Table 3. Instance segmentation results on S3DIS. The performance on both Area-5and 6-fold cross-validation is reported.

much simpler pipeline. With 6-fold validation, our method improves HAIS [4] by 4.5%, 3.7%, 3.2%, and 4.6% in terms of mConv, mWConv, mPrec, and mRec, respectively. The proposed approach works in a fully end-to-end fashion, removing the error accumulation caused by the inter-task dependencies.

Instance Segmentation on ScanNet. The performance of instance segmentation on the validation and testing sets of ScanNet [6] is reported in Tab. 4 and Tab. 5, respectively. On the validation set, we report the performance with both small and large backbones, denoted as Ours-S and Ours-L, respectively. It surpasses previous top-performing methods on both architectures in terms of mAP, demonstrating strong generalization capability. Compared with DyCo3D [14], our approach exceeds it by 4.2% in terms of mAP. The qualitative

14 T. He et al.

Table 4. Quantitative comparison on the validation set of ScanNetV2. To make a fair comparison, we report the performance with different model scalability. The performance of HAIS-S is obtained by using the official training code.

	AP@50	mAP	cabinet	bed	chair	sofa	table	door	window	bookshe.	picture	counter	desk	curtain	fridge	s.curtain	toilet	sink	$_{\rm bath}$	otherfu.
SGPN [34]	11.3	-	10.1	16.4	20.2	20.7	14.7	11.1	11.1	0.0	0.0	10.0	10.3	12.8	0.0	0.0	48.7	16.5	0.0	0.0
3D-SIS [16]	18.7	-	19.7	37.7	40.5	31.9	15.9	18.1	0.0	11.0	0.0	0.0	10.5	11.1	18.5	24.0	45.8	15.8	23.5	12.9
3D-MPA [7]	59.1	35.3	51.9	72.2	83.8	66.8	63.0	43.0	44.5	58.4	38.8	31.1	43.2	47.7	61.4	80.6	99.2	50.6	87.1	40.3
PointGroup [17]	56.9	34.8	48.1	69.6	87.7	71.5	62.9	42.0	46.2	54.9	37.7	22.4	41.6	44.9	37.2	64.4	98.3	61.1	80.5	53.0
DyCo3D-S [14]	57.6	35.4	50.6	73.8	84.4	72.1	69.9	40.8	44.5	62.4	34.8	21.2	42.2	37.0	41.6	62.7	92.9	61.6	82.6	47.5
HAIS-S [4]	59.1	38.0	54.4	76.0	87.7	69.4	66.5	47.5	48.5	53.1	43.6	24.0	50.9	55.8	45.1	58.5	94.7	53.6	80.8	53.0
Ours-S	59.2	39.6	51.1	75.9	86.5	72.8	67.3	45.2	52.3	57.2	43.8	25.7	40.5	53.7	37.2	59.4	98.2	58.9	87.0	52.9
DyCo3D-L [14]	61.0	40.6	52.3	70.4	90.2	65.8	69.6	40.5	47.2	48.4	44.7	34.9	52.3	47.5	51.5	70.3	94.8	74.3	77.4	56.4
HAIS-L [4]	64.0	43.5	55.4	70.2	82.5	67.7	75.3	48.1	51.5	49.4	48.7	47.8	58.5	55.7	53.0	76.1	100.0	69.2	87.1	56.3
Ous-L	63.7	45.6	58.5	78.5	93.6	63.2	76.5	55.6	48.5	59.4	38.3	36.9	54.2	50.7	46.2	72.3	98.3	68.8	87.1	59.5

Table 5. Quantitative results on ScanNetV2 testing set.

	mAP	bathtub	bed	bookshe.	cabinet	chair	counter	curtain	desk	door	otherfu.	picture	refrige.	s.curtain	sink	sofa	table	toilet	window
R-PointNet [37]	15.8	35.6	17.3	11.3	14.0	35.9	1.2	2.3	3.9	13.4	12.3	0.8	8.9	14.9	11.7	22.1	12.8	56.3	9.4
3D-SIS [16]	16.1	40.7	15.5	6.8	4.3	34.6	0.1	13.4	0.5	8.8	10.6	3.7	13.5	32.1	2.8	33.9	11.6	46.6	9.3
MASC [24]	25.4	46.3	24.9	11.3	16.7	41.2	0.0	37.4	7.3	17.3	24.3	13.0	22.8	36.8	16.0	35.6	20.8	71.1	13.6
PanopticFusion [26]	21.4	25.0	33.0	27.5	10.3	22.8	0.0	34.5	2.4	8.8	20.3	18.6	16.7	36.7	12.5	22.1	11.2	66.6	16.2
3D-BoNet [36]	25.3	51.9	32.4	25.1	13.7	34.5	3.1	41.9	6.9	16.2	13.1	5.2	20.2	33.8	14.7	30.1	30.3	65.1	17.8
MTML [19]	28.2	57.7	38.0	18.2	10.7	43.0	0.1	42.2	5.7	17.9	16.2	7.0	22.9	51.1	16.1	49.1	31.3	65.0	16.2
3D-MPA [7]	35.5	45.7	48.4	29.9	27.7	59.1	4.7	33.2	21.2	21.7	27.8	19.3	41.3	41.0	19.5	57.4	35.2	84.9	21.3
DyCo3D [14]	39.5	64.2	51.8	44.7	25.9	66.6	5.0	25.1	16.6	23.1	36.2	323.2	33.1	53.5	22.9	58.7	43.8	85.0	31.7
PointGroup [17]	40.7	63.9	49.6	41.5	24.3	64.5	2.1	57.0	11.4	21.1	35.9	21.7	42.8	66.0	25.6	56.2	34.1	86.0	29.1
HAIS [4]	45.7	70.4	56.1	45.7	36.4	67.3	4.6	54.7	19.4	30.8	42.6	28.8	45.4	71.1	26.2	56.3	43.4	88.9	34.4
Ours	43.8	81.5	50.7	33.8	35.5	70.3	8.9	39.0	20.8	31.3	37.3	28.8	40.1	66.6	24.2	55.3	44.2	91.3	29.3
$OccuSeg^*$ [10]	44.3	85.2	56.0	38.0	24.9	67.9	9.7	34.5	18.6	29.8	33.9	23.1	41.3	80.7	34.5	50.6	42.4	97.2	29.1
SSTN [*] [21]	50.6	73.8	54.9	49.7	31.6	69.3	17.8	37.7	19.8	33.0	46.3	57.6	51.5	85.7	49.4	63.7	45.7	94.3	29.0

result is illustrated in Fig. 5. We also make a fair comparison with HAIS [4], the highest mAP is achieved on the validation set.

5 Conclusion and Future Works

In this paper, we propose a novel pipeline for 3D instance segmentation, which works in a per-point prediction fashion and thus removes the inter-task dependencies. We show that the key to its success is the target assignment, which is addressed by an Optimal Transport solution. Without bells and whistles, our method achieves promising results on two commonly used datasets.

The sampling strategy used in our method is fps, which is slightly better than random sampling. We believe there exist other informative strategies that can further improve the performance. In addition, due to the continuity representation capability, our method offers a simple solution to achieve instance-level reconstruction with the sparse point cloud. We leave these for future works.

References

- Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., Savarese, S.: 3d semantic parsing of large-scale indoor spaces. In: CVPR (2016)
- Brabandere, B.D., Neven, D., Gool, L.V.: Semantic instance segmentation with a discriminative loss function. arXiv preprint arXiv:1708.02551 (2017)
- Cai, Z., Vasconcelos, N.: Cascade R-CNN: Delving into high quality object detection. In: CVPR (2018)
- Chen, S., Fang, J., Zhang, Q., Liu, W., Wang, X.: Hierarchical aggregation for 3d instance segmentation. In: ICCV (2021)
- 5. Cuturi, M.: Sinkhorn distances: Lightspeed computation of optimal transport. In: NeurIPS (2013)
- Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nie
 ßner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: CVPR (2017)
- Engelmann, F., Bokeloh, M., Fathi, A., Leibe, B., Nießner, M.: 3D-MPA: Multi proposal aggregation for 3d semantic instance segmentation. In: CVPR (2020)
- Ge, Z., Liu, S., Li, Z., Yoshie, O., Sun, J.: Ota: Optimal transport assignment for object detection. In: CVPR (2021)
- 9. Graham, B., Engelcke, M., van der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: CVPR (2018)
- Han, L., Zheng, T., Xu, L., Fang, L.: Occuseg: Occupancy-aware 3d instance segmentation. In: CVPR (2020)
- 11. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask R-CNN. In: ICCV (2017)
- 12. He, T., Gong, D., Tian, Z., Shen, C.: Learning and memorizing representative prototypes for 3d point cloud semantic and instance segmentation. In: ECCV (2020)
- He, T., Liu, Y., Shen, C., Wang, X., Sun, C.: Instance-aware embedding for point cloud instance segmentation. In: ECCV (2020)
- 14. He, T., Shen, C., van den Hengel, A.: DyCo3d: Robust instance segmentation of 3d point clouds through dynamic convolution. In: CVPR (2021)
- He, T., Shen, C., van den Hengel, A.: Dynamic Convolution for 3D Point Cloud Instance Segmentation. arXiv preprint arXiv:2107.08392 (2021)
- Hou, J., Dai, A., Nießner, M.: 3D-SIS: 3d semantic instance segmentation of rgb-d scans. In: CVPR (2019)
- 17. Jiang, L., Zhao, H., Shi, S., Liu, S., Fu, C.W., Jia, J.: Pointgroup: Dual-set point grouping for 3d instance segmentation. In: CVPR (2020)
- Kong, T., Sun, F., Liu, H., Jiang, Y., Li, L., Shi, J.: Foveabox: Beyond anchor-based object detector. IEEE TIP (2020)
- Lahoud, J., Ghanem, B., Pollefeys, M., Oswald, M.R.: 3d instance segmentation via multi-task metric learning. In: ICCV (2019)
- 20. Landrieu, L., Simonovski, M.: Large-scale point cloud semantic segmentation with superpoint graphs. In: CVPR (2018)
- Liang, Z., Li, Z., Xu, S., Tan, M., Jia, K.: Instance segmentation in 3d scenes using semantic superpoint tree networks. In: ICCV (2021)
- 22. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: CVPR (2017)
- Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollár, P.: Focal loss for dense object detection. In: ICCV (2017)
- Liu, C., Furukawa, Y.: MASC: Multi-scale affinity with sparse convolution for 3d instance segmentation. arXiv preprint arXiv:1902.04478 (2019)

- 16 T. He et al.
- Mo, K., Zhu, S., Chang, A.X., Yi, L., Tripathi, S., Guibas, L.J., Su, H.: PartNet: A large-scale benchmark for fine-grained and hierarchical part-level 3D object understanding. In: CVPR (2019)
- 26. Narita, G., Seno, T., Ishikawa, T., Kaji, Y.: Panopticfusion: Online volumetric semantic mapping at the level of stuff and things. In: IROS (2019)
- Pham, Q.H., Nguyen, D.T., Hua, B.S., Roig, G., Yeung, S.K.: JSIS3D: Joint semantic-instance segmentation of 3d point clouds with multi-task pointwise networks and multi-value conditional random fields. In: CVPR (2019)
- Qi, C.R., Litany, O., He, K., Guibas, L.J.: Deep hough voting for 3d object detection in point clouds. In: ICCV (2019)
- Qi, C.R., Liu, W., Wu, C., Su, H., Guibas, L.J.: Frustum pointnets for 3d object detection from rgb-d data. In: CVPR (2018)
- Ren, S., He, K., Girshick, R., Sun, J.: Faster R-CNN: Towards real-time object detection with region proposal networks. In: NeurIPS (2015)
- Tian, Z., Shen, C., Chen, H.: Conditional convolutions for instance segmentation. In: ECCV (2020)
- 32. Tian, Z., Shen, C., Chen, H., He, T.: FCOS: Fully convolutional one-stage object detection. In: ICCV (2019)
- Tian, Z., Shen, C., Chen, H., He, T.: FCOS: A simple and strong anchor-free object detector. IEEE TPAMI (2021)
- Wang, W., Yu, R., Huang, Q., Neumann, U.: SGPN: Similarity group proposal network for 3d point cloud instance segmentation. In: CVPR (2018)
- Wang, X., Liu, S., Shen, X., Shen, C., Jia, J.: Associatively segmenting instances and semantics in point clouds. In: CVPR (2019)
- 36. Yang, B., Wang, J., Clark, R., Hu, Q., Wang, S., Markham, A., Trigoni, N.: Learning object bounding boxes for 3d instance segmentation on point clouds. In: NeurIPS (2019)
- 37. Yi, L., Zhao, W., Wang, H., Sung, M., Guibas, L.J.: GSPN: Generative shape proposal network for 3d instance segmentation in point cloud. In: CVPR (2018)
- Yu, J., Jiang, Y., Wang, Z., Cao, Z., Huang, T.: Unitbox: An advanced object detection network. In: ACM MM (2016)
- Zhang, S., Chi, C., Yao, Y., Lei, Z., Li, S.Z.: Bridging the gap between anchorbased and anchor-free detection via adaptive training sample selection. In: CVPR (2020)
- 40. Zhang, X., Wan, F., Liu, C., Ji, R., Ye, Q.: FreeAnchor: Learning to match anchors for visual object detection. In: NeurIPS (2019)
- 41. Zhao, L., Tao, W.: JSNet: Joint instance and semantic segmentation of 3d point clouds. In: AAAI (2020)
- Zhou, X., Wang, D., Krähenbühl, P.: Objects as points. In: arXiv preprint arXiv:1904.07850 (2019)