

Spherical Mask: Coarse-to-Fine 3D Point Cloud Instance Segmentation with Spherical Representation

Sangyun Shin Kaichen Zhou Madhu Vankadari Andrew Markham Niki Trigoni

Department of Computer Science
 University of Oxford

firstname.lastname@cs.ox.ac.uk

Abstract

Coarse-to-fine 3D instance segmentation methods show weak performances compared to recent Grouping-based, Kernel-based and Transformer-based methods. We argue that this is due to two limitations: 1) Instance size over-estimation by axis-aligned bounding box(AABB) 2) False negative error accumulation from inaccurate box to the refinement phase. In this work, we introduce **Spherical Mask**, a novel coarse-to-fine approach based on spherical representation, overcoming those two limitations with several benefits. Specifically, our coarse detection estimates each instance with a 3D polygon using a center and radial distance predictions, which avoids excessive size estimation of AABB. To cut the error propagation in the existing coarse-to-fine approaches, we virtually migrate points based on the polygon, allowing all foreground points, including false negatives, to be refined. During inference, the proposal and point migration modules run in parallel and are assembled to form binary masks of instances. We also introduce two margin-based losses for the point migration to enforce corrections for the false positives/negatives and cohesion of foreground points, significantly improving the performance. Experimental results from three datasets, such as ScanNetV2, S3DIS, and STPLS3D, show that our proposed method outperforms existing works, demonstrating the effectiveness of the new instance representation with spherical coordinates. The code is available at: <https://github.com/yunshin/SphericalMask>

1. Introduction

3D instance segmentation has gained immense attention with its wide range of applications for Indoor Scanning[38], Augmented Reality(AR)[18], and Autonomous Driving[22]. Similar to 2D instance segmentation, the goal of the task is to identify each object along with its class label. Nevertheless, the sparse and unordered nature of point clouds has led to the development of methods different from 2D image segmentation.

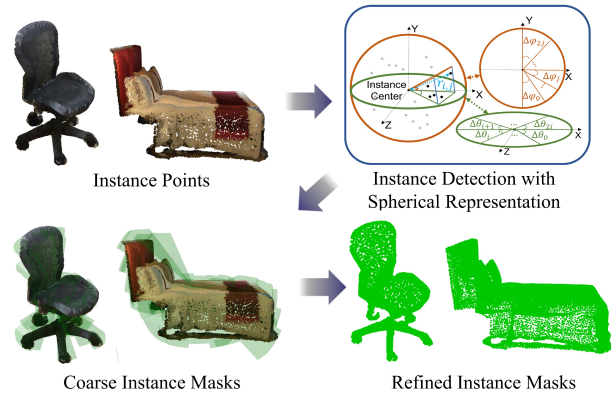


Figure 1. Pipeline of Spherical Mask with coarse-to-fine framework. Given point cloud, instances are detected with 3D polygons defined in spherical coordinates. In the refinement phase, the points virtually migrate based on the polygon to estimate fine instance masks.

Existing approaches for 3D instance segmentation are broadly categorized into coarse-to-fine based[12, 14, 20, 33, 34], grouping-based[4, 13, 19, 30, 36], kernel-based[10, 11, 13, 29, 31], and Transformer-based[1, 15, 21, 27, 28] approaches. Recent progress in clustering techniques and attention mechanisms have driven the performances of the last three approaches to state-of-the-art. Compared to these three approaches, the coarse-to-fine approach has received relatively less attention due to a low accuracy, caused by two limitations: 1) False negative error propagation from the coarse detection to the refinement stage. 2) Overestimation of instance size.

Coarse-to-fine instance segmentation first performs coarse detection, followed by refinement using the coarse detection as hard reference. The basic assumption of the approach is that the coarse detection stage always provides a neat detection for refinement. However, this assumption is often violated as the coarse detection stage cannot always produce neat outputs, which causes an issue of upper bound accuracy. For example, if the first coarse detection does not include all foreground points, the following

refinement(binary classification) step has no means to include them, only possibly accumulating the error from false negative prediction. The other limitation is that the axis-aligned-bounding-box(AABB) estimation, which is typically used for the coarse detection, has been claimed to be an ill-defined problem[30] because commonly used box regression losses(L1, L2) result in overestimation of object sizes. For instance, the target values of AABB are minimum and maximum values in x,y,z cartesian coordinates, making the box include redundant empty space, as points only lie on the surface of the object.

In this work, we address the aforementioned two limitations of coarse-to-fine instance segmentation. Our core intuition comes from the fact that the weakness of the coarse-to-fine approach is based on the structural disentanglement of coarse detection and fine refinement phase. Instead of assuming that the coarse part should be perfect, we take a relaxation approach, which regards the coarse detection as a soft reference during the refinement phase, providing more access for refinement yet restricting access to unnecessary background points. Specifically, to improve the coarse detection part, we estimate a 3D polygon in spherical coordinates instead of AABB, alleviating the issue of excessive object size estimation. To remove the error propagation from inaccurate coarse detection to the refinement stage, we virtually migrate points based on the 3D polygon with reduced complexity in spherical coordinates.

In summary, our method **Spherical Mask** finds each instance by estimating a 3D polygon fitting to an instance in spherical coordinates and migrating points inside or outside the polygon to produce the fine instance mask. Our contributions are:

- We introduce a new alternative instance representation based on spherical coordinates, which overcomes the limitations in existing coarse-to-fine approaches.
- To circumvent the issue of excessive estimation of instance size, we propose **Radial Instance Detection(RID)** that formulates an instance into a 3D polygon as a coarse detection.
- To cut the error propagation from coarse detection to the refinement phase, we introduce **Radial Point Migration(RPM)**, capable of refining both false positive and false negative points from RID.
- Extensive experiments on ScanNetV2 [7], S3DIS [2], and STPLS3D [3] show the effectiveness of our approach, pushing the boundary of current SOTA.

2. Related Work

Existing works on point cloud instance segmentation can be categorized into proposal-based, clustering-based, kernel-based and transformer-based methods.

2.1. Coarse to Fine(Proposal-based) Approach

Coarse-to-fine-based methods are built on a conceptually simple design, where they first perform coarse detection, followed by a refinement stage to acquire fine segmentation. Typically, 3D bounding box is employed for the coarse detection. 3D-SIS[12] performs instance segmentation by first detecting the 3D boxes and refining points inside. 3D-BoNet[33] matches query AABBs and ground-truth instance using Hungarian algorithm for the supervision. The predicted AABBs are then concatenated with point features to produce binary masks for each instance. GSPN[34] adopts set-abstraction[24] to get query points and infer AABBs. The features inside the AABBs are extracted and used for per-point mask segmentation. More recently, TD3D [14] proposes a fully sparse-convolutional approach for point instance segmentation. It first detects AABBs and extracts perpoint features inside the boxes to perform binary classification. Most of the works use coarse detection(bounding box) directly as geometric features for predicting per-point binary instance masks. Thus, their accuracies greatly depend on the precision of the coarse detection, as inaccurate detection could easily lead to a large number of false negative points.

2.2. Grouping-based Approach

Grouping-based methods learn latent embeddings to perform per-point predictions, such as semantic categories and clustering to acquire instances. PointGroup[13] predicts centroid offsets of each point and utilizes this shifted point cloud and original point cloud to obtain the clusters. Based on this concept, many studies improve the clustering technique with hierarchical intra-instance predictions[4], superpoint-based divisive grouping[19], soft grouping[30], and binary clustering[36]. The clustering-based methods have high expectations of the quality of per-point center prediction in 3D, which is challenging to generalize with various spatial extents of objects.

2.3. Kernel-based Approach

Kernel-based methods learn convolution kernels that aggregate point features to estimate instance masks. DyCo3D[10] proposes discriminative kernels by applying the clustering method from PointGroup[13]. Built on this, more recent works improved the performance by replacing the clustering part with farthest-point sampling[11], candidate localization[31], and instance-aware point sampling for the high-recall[29].

2.4. Transformer-based Approach

Recently, transformer-based methods have set the new SOTA. Based on mask-attention[5, 6], Mask3D[27] and SPFormer[28] present the pipeline that learns to output

instance prediction directly from a fixed number of object queries from voxel and superpoint features, respectively. Based on these works, recent studies improve the performance with auxiliary center regression[15], query initialization and set grouping[21], spatial and semantic supervision[1]. Although the powerful architectural advantage has driven the performance of Transformer-based approaches, low-recall and how to distribute initial queries remain challenges.

3. Method

3.1. Overview

Given an input point cloud $p_1 \in \mathbb{R}^{N_p \times 3}$ in 3-dimensional cartesian coordinates and the corresponding color information $p_{\text{rgb}} \in \mathbb{R}^{N_p \times 3}$, we aim to design a system that segments the point cloud into local binary masks of instances $\{o^{(i)} \in \mathbb{R}^{N_p \times 1}\}_{i=1}^{N_o}$ using a coarse to fine approach. Here, N_p, N_o are the total number of points and the number of instances, respectively. This is achieved using the proposed method depicted in Figure 2. Our system consists of mainly two modules: 3D backbone and proposed instance mask estimation. The details of these modules are in Section 3.2 and Section 3.3, respectively.

3.2. 3D backbone

Our 3D backbone is similar to [29] and is composed of two modules: a 3D encoder and a voting module. The 3D encoder uses U-Net [26] with sparse convolutions [9] to encode the given point cloud into deep features $F_1 \in \mathbb{R}^{N_p \times D}$. Then, F_1 and the respective input points p_1 are fed into the voting module. The voting module performs set abstraction [24], producing K votes with query points $p_2 \in \mathbb{R}^{K \times 3}$ and features $F_2 \in \mathbb{R}^{K \times D}$. Please find the original paper [25] for more details about this procedure. These votes are spread in the scene, providing features to be further processed through the proposed instance mask estimation module. The details are explained in the following sections.

3.3. Instance Mask Estimation

In this section, we estimate the instance masks from the votes predicted in the 3D backbone section using three modules: Radial Instance Detection, Radial Point Migration, and Mask Assembly. All of the modules are explained in the following sections. For the notation simplicity, we will explain how each vote feature is processed. Therefore, we write f_2 to refer to a single vote feature of F_2 from here onwards.

3.3.1 Radial Instance Detection(Coarse Mask)

Radial Instance Detection(RID) aims to detect instances for further refinement. Similar to PolarMask[32] for 2D segmentation, we define an instance as a 3D polygon with a

center f_{center} and multiple rays f_{ray} emitting from the center forming each spherical sectors. Here, the sectors are defined by preset angles. Each ray then determines the distance to be considered for their corresponding sectors, as shown in Figure 3.

We estimate the closest instances' center f_{center} using offsets predicted from an MLP network, *CenterHead*, which takes the respective f_2 as input and outputs offsets. The offsets are added to p_2 to infer f_{center} . After this, the input point cloud is converted into spherical coordinates using a transformation as $p_s = S(p_1)$ centered around f_{center} , where $S : (x, y, z) \rightarrow (r, \theta, \varphi)$ as follows:

$$\begin{aligned} r &= \sqrt{x^2 + y^2 + z^2}, \\ \theta &= \arctan \frac{y}{x}, \\ \varphi &= \arctan \frac{z}{\sqrt{x^2 + y^2 + z^2}}. \end{aligned} \quad (1)$$

Here, r, θ, φ refer to radius, horizontal, and vertical angles in spherical coordinates, respectively. The p_s is then divided uniformly with N_θ and N_φ separations for horizontal and vertical direction respectively, resulting in $N_\theta \cdot N_\varphi$ sectors, as shown in Figure 3 (b) and (c).

We consider points inside the same sector to have identical (θ, φ) , which enables us to close the sector using a boundary estimated by f_{ray} . To estimate $\{f_{\text{ray}}^{(i)}\}_{i=1}^{N_\theta N_\varphi}$, another MLP named *Ray Head* is employed, which takes f_2 as input.

At this point, every point inside the corresponding sector's boundary from f_{ray} is considered foreground. Using the boundary, RID offers tighter boundaries of instances than AABB in point cloud, as each sector is closed at the distance of the farthest foreground point in the sector, alleviating the problem of redundant space in AABB. Please refer to our supplementary material for additional visualizations.

Coarse Instance Loss: During training, the estimated f_{center} and f_{ray} are compared against their respective ground truth $\mathbf{g}_{\text{center}}$ and \mathbf{g}_{ray} to calculate L_{coarse} :

$$L_{\text{coarse}} = L_{\text{ray}} + L_{\text{center}}, \quad (2)$$

where L_{ray} and L_{center} are defined with L1 loss:

$$L_{\text{ray}} = \frac{1}{N_\theta N_\varphi} \sum_{i=1}^{N_\theta N_\varphi} \left| f_{\text{ray}}^{(i)} - \mathbf{g}_{\text{ray}}^{(i)} \right|_1 \quad (3)$$

$$L_{\text{center}} = |f_{\text{center}} - \mathbf{g}_{\text{center}}|_1, \quad (4)$$

Here, the ground-truth instance \mathbf{g} is matched with f_{center} by an injective mapping obtained using Hungarian algorithm as [29, 33]. For the details of the matching, please refer to Sec 3.4. $\mathbf{g}_{\text{ray}}^{(i)}$ is set to the distance between $\mathbf{g}_{\text{center}}$ and the furthest foreground point in the i_{th} sector. If there are no foreground points in the sector, $\mathbf{g}_{\text{ray}}^{(i)}$ is set to minimum as $1e-5$. The target center $\mathbf{g}_{\text{center}}$ is calculated as mean values

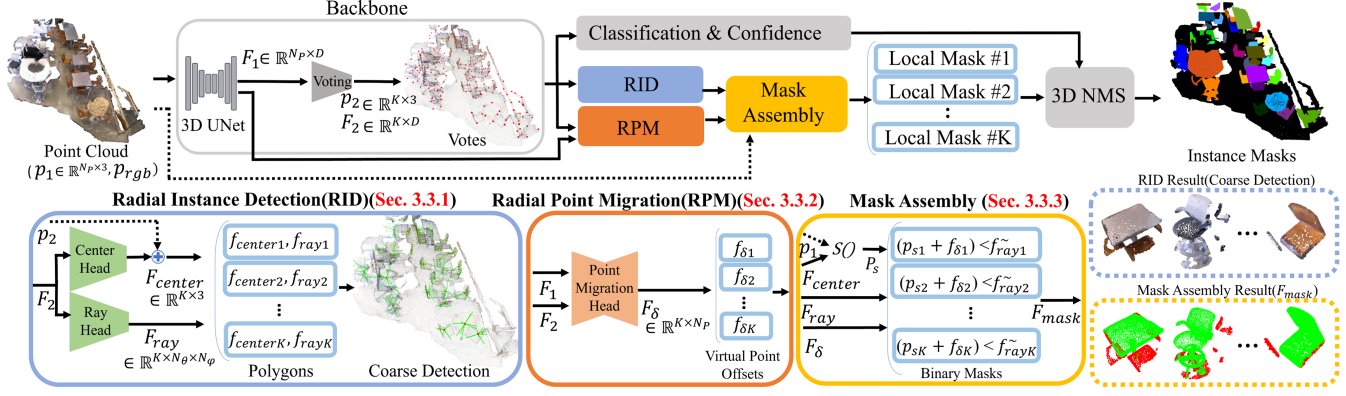


Figure 2. Overall pipeline of our proposed method based on coarse to fine approach. Given the point cloud, the backbone produces base features with 3D UNet and Voting module. Based on this, RID performs coarse detection while RPM produces the virtual point offsets to refine the coarse detection. In Mask Assembly, K local binary masks are generated, where each mask is a proposal for a single instance. 3D NMS is applied to acquire the final instance masks using local binary masks, classifications, and confidence scores.

of foreground points of the matched groundtruth instance in cartesian coordinates.

3.3.2 Radial Point Migration(Mask Refinement)

In this section, we introduce a refinement process to perform per-point fine-tuning. This is because the coarse detection will invariably include points belonging to the background or other instances (i.e. false positives) inside its boundary and neglect some instance points that fall outside (i.e. false negatives). We propose a conceptually simple yet effective dual that jitters individual points to belong to the correct instance. In particular, this is enabled by our innovative use of spherical coordinates - we only need to learn a single radial delta for each point to move it along the ray to the instance centroid while keeping angular quantities ϕ and θ constant. Note that this is a *virtual* point motion - we do not alter the final point cloud. We use this as a virtual offset to obtain clean instance labels without modifying the coarse sector.

By estimating an offset value for each point p_1 , these misclassified points can be virtually migrated to being in the correct region. Based on its good performance on per-point prediction, we adopt Dynamic Convolution in a similar manner to [10, 29] as our *Point Migration Head*, for predicting per-point offsets $F_\delta \in \mathbb{R}^{K \times N_p}$ using the vote features $F_2 \in \mathbb{R}^{K \times D}$ as queries against the point features $F_1 \in \mathbb{R}^{N_p \times D}$. For the coherence with the notations, we write $f_\delta \in \mathbb{R}^{N_p}$ to refer to an output of F_δ , corresponding to one vote. For the learning of f_δ , we divide points into two groups. The first is to learn the radial delta for the case of misclassification, and the second is to make instances more compact and cohesive by migrating points to the centroid of the sector.

Misclassification Correction Loss: This process aims to estimate a radial delta to move the misclassified points either

inside or outside the estimated coarse sector. There are two possible cases where the misclassification could occur, as shown in Figure 4 (a): Instance points could lie outside the sector boundary, acting as false negatives, or background points could incorrectly lie within the sector boundary, acting as false positives. The goal is to move these points to the correct region.

Formally, given the point indices of foreground points $\{j^{+(i)}\}_{i=1}^{N+}$ and background points $\{j^{-(i)}\}_{i=1}^{N-}$ from the groundtruth, the false negative points p_{fn} and the false positive points p_{fp} are defined as:

$$p_{fn} = \{p_s^{(j^+)} : (p_s^{(j^+)} + f_\delta^{(j^+)}) > f_{ray}^{(j^+)}\}, \quad (5)$$

$$p_{fp} = \{p_s^{(j^-)} : (p_s^{(j^-)} + f_\delta^{(j^-)}) < f_{ray}^{(j^-)}\}, \quad (6)$$

where $N+$ and $N-$ stand for the number of foreground and background points, respectively. Here, $\tilde{j} = findSector(p_s^{(j)})$ where $findSector(\cdot)$ is a function that takes $p_s^{(j)}$ as input and returns the index of the sector that $p_s^{(j)}$ belongs to. Please refer to the supplementary material for our implementation of $findSector(\cdot)$ function.

The union of p_{fp} and p_{fn} forms the misclassified points p_{miss} that we are interested: $p_{miss} = p_{fp} \cup p_{fn}$. Our aim is to push or pull them inside/outside of the ray with margins. Thus, the loss function L_{mc} is formulated with soft margin loss as:

$$L_{mc} = \frac{1}{N_{miss}} \sum_{i=1}^{N_{miss}} \log(1 + \exp(y * \tanh(p_{miss}^{(i)} + f_\delta^{(i)} - f_{ray}^{(i)}))) \quad (7)$$

where

$$y = \begin{cases} 1 & \text{if } p_{miss}^{(i)} \in p_{fp}, \\ -1 & \text{if } p_{miss}^{(i)} \in p_{fn}, \end{cases} \quad (8)$$

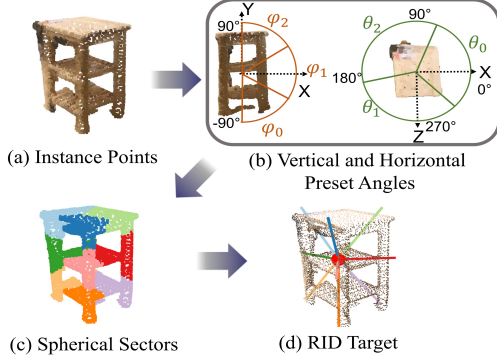


Figure 3. Process of RID. (a) Object points in cartesian coordinates (b) Converting points into a spherical coordinate system, using f_{center} , and preset angles θ and φ . (c) Assigning points to each sector defined by θ and φ . The example shows $3/3$ for θ/φ . (d) For each sector, the distance between the farthest point and the center becomes the target of f_{ray} . During inference, points with smaller distance than f_{ray} are considered foreground.

Here, \tanh is hyperbolic-tangent function and L_{mc} is calculated for all the votes that are assigned to the ground truth instance. N_{miss} stands for the number of element in p_{miss} and \hat{i} is index of f_{δ} corresponding to $p_{\text{miss}}^{(i)}$. f_{ray} is only used for reference and the gradient for learning f_{ray} is not calculated.

The misclassified points around the edge, g_{ray} , of an instance are provided with comparably easy learning targets. In contrast, misclassified points far from the predicted rays are assigned targets with large discrepancies, encouraging larger gradients during the training.

Sector Cohesion Loss The goal of this block is to move true-positive points to the centroid of the sector. By doing so, the sector becomes more cohesive, encouraging the learning of common and shared features of an instance as the foreground features are getting close to each other. In addition, this helps to provide a learning signal for true-positive points, as L_{mc} only considers false negatives/positives. This is shown more clearly in Figure 4 (b).

Similar to L_{mc} , using point indice of foreground points $j+$, we extract true positives p_{tp} as :

$$p_{\text{tp}} = \{p_s^{(j+)} : (p_s^{(j+)} + f_{\delta}^{(j+)}) < f_{\text{ray}}^{(j+)}\} \quad (9)$$

and formulate the loss with the soft margin calculation:

$$L_{sc} = \frac{1}{N_{\text{tp}}} \sum_{i=1}^{N_{\text{tp}}} \log(1 + \exp(\tanh(f_{\delta}^{(i)} + p_{\text{tp}}^{(i)} - f_{\text{center}}))), \quad (10)$$

where N_{tp} stands for the number of element in p_{tp} and \hat{i} is index of f_{δ} corresponding to $p_{\text{tp}}^{(i)}$. Since f_{center} is always 0 in centered spherical coordinate, the loss can be simplified as:

$$L_{sc} = \frac{1}{N_{\text{tp}}} \sum_{i=1}^{N_{\text{tp}}} \log(1 + \exp(\tanh(f_{\delta}^{(i)} + p_{\text{tp}}^{(i)}))), \quad (11)$$

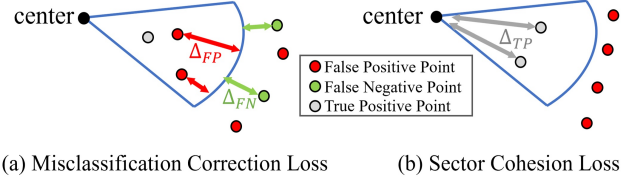


Figure 4. Conceptual diagram showing per-point migration following both (a) L_{mc} and (b) L_{sc} . Δ_{FP} and Δ_{FN} are distances penalized by L_{mc} with margin for misclassified points. Δ_{TP} is the distance that L_{sc} penalizes to enforce the learning of general features of an instance by making each sample close to the other around the center.

L_{mc} and L_{sc} together form the refinement loss L_{fine} as:

$$L_{\text{fine}} = L_{mc} + L_{sc}. \quad (12)$$

Our proposed virtual point migration brings three advantages for refinement over existing approaches that strictly disentangle coarse detection and refinement: 1) Instead of only focusing on points inside the coarse detection, predicting the offsets of all points allows the refinement of even false negative points outside of f_{ray} , sidestepping the error accumulation from the coarse detection. 2) By considering the sector radius, g_{ray} , it is possible to have a soft target for each point rather than a hard target which is the center of the sector/instance. For example, false negative points outside of the sector boundary only need to be migrated a small distance to being with the sector (soft), rather than being driven towards the center (hard). This makes it easier to learn how to perform the point migration. 3) We only need to learn a one-dimensional number to migrate the point virtually along the radial line. This is far simpler than having to learn a three-dimensional offset in cartesian coordinates.

3.3.3 Mask Assembly

Our final mask is assembled by comparing virtually migrated points and f_{ray} . Specifically, the local binary mask $\{f_{\text{mask}}^{(i)}\}_{i=1}^{N_p}$ is formed as:

$$f_{\text{mask}}^{(i)} = \begin{cases} 1 & \text{if } (p_s^{(i)} + f_{\delta}^{(i)}) < f_{\text{ray}}^{(i)} \\ 0 & \text{otherwise,} \end{cases} \quad (13)$$

3.4. Training

For the training, we also learn classification and confidence with respective MLPs. For classification, we apply cross-entropy loss, L_{cls} , for learning the classes of matched ground-truth instances. For the confidence scores of the proposals, we apply L2 loss to learn IoUs between the proposals and the groundtruth instances. Similar to [29], we also duplicate the number of groundtruth for 4 times and create a cost matrix C :

$$C(k, i) = L_{\text{coarse}}(k, i) + L_{\text{fine}}(k, i) + L_{\text{cls}}(k, i), \quad (14)$$

where $L(\cdot)$ refers to the loss value calculated using k_{th} vote and i_{th} ground-truth instance. Referring C , we apply Hungarian algorithm to find the least-cost injective mapping from each ground-truth instance to the votes. The final loss using the acquired ground-truths is:

$$L = \lambda_1 L_{cls} + \lambda_2 L_{conf} + \lambda_3 L_{coarse} + \lambda_4 L_{fine} \quad (15)$$

4. Experiment

4.1. Dataset

We evaluate our method on three datasets: ScanNetV2 [7], S3DIS [2], STPLS3D [3]. Following are descriptions of each dataset.

ScanNetV2 ScanNetV2 dataset consists of 1201, 312, and 100 scans with 18 object classes for training, validation, and testing, respectively. We report the evaluation results on the validation and hidden test sets as in the existing works.

S3DIS S3DIS dataset contains 271 scenes from 6 areas with 13 categories. We report evaluations for both Area 5. Additional evaluation results with 6-fold cross-validation can be found in the supplementary material.

STPLS3D The STPLS3D dataset is an aerial photogrammetry point cloud dataset from real-world and synthetic environments. It includes 25 urban scenes of $6km^2$ and 14 instance categories. Following [4, 29, 30], we use scenes 5, 10, 15, 20, and 25 for validation and the rest for training.

4.2. Evaluation Metric

We adopt average precision as our primary evaluation metric. Average precision is extensively used in vision tasks such as object detection and instance segmentation tasks. The metric calculates precisions by varying the IoU threshold. Following the existing works, we evaluate our model with three IoU thresholds: AP , AP_{50} , AP_{25} . AP_{50} and AP_{25} stand for average precisions with IoU threshold as 25% and 50%, respectively. AP is an averaged score by varying IoU thresholds from 50% to 95% by increasing the threshold with step size 5%. For S3DIS, we also evaluate our model with mean precision (mPrec50), and mean recall with IoU threshold as 50%(mRec50).

4.3. Implementation Detail

We build our model on PyTorch framework [23] and train it for 300 epochs with AdamW optimizer with a single NVIDIA A10 GPU. The batch size is set to 10. The learning rate and weight decay are initialized to 0.001 and 0.0001. Cosine annealing [35] is used for scheduling the learning rate. Following [29, 30], the voxel size is set to 0.02m for ScanNetV2 and S3DIS, and to 0.3m for STPLS3D. For the augmentation during training, we use random cropping for each scene with a maximum number of 250,000 points. During testing, a whole scene is used as an input to the network.

Our backbone is similar to [29, 30], which outputs features F_1 with hidden dimension D as 32 channels. For the voting, we use two set-abstraction [24, 25] layers with the ball query radius 0.2 and 0.4, respectively. The number seeds and votes are set to 1024 and 256, respectively. The number of neighbors is set to 32 for both layers, similar to [29]. For *Point Migration Head*, we use two layers of dynamic convolution [10], and their hidden dimensions are set to 32. λ_1 , λ_2 , λ_3 , and λ_4 are set to 0.5, 0.5, 1, and 1, respectively. For training and inference, we set N_θ and N_φ to 5 and 5, respectively. During inference, Non-Maximum-Suppression is applied to the K binary masks to delete redundant masks using a confidence score 0.2 as a threshold. Following [19, 21, 29, 31], we aggregate superpoints [16, 17] to align the final prediction masks on the ScanNetV2 dataset. Other details, such as the architecture of MLPs and runtime analysis, are included in the supplementary material.

4.4. Main Results

ScanNetV2 Table 1 and Table 2 show the quantitative result of instance segmentation on the test and validation sets. Our proposed method achieves the highest AP and AP_{50} surpassing the previous strongest method by the margin of 4.1% and 2.4% for the test set, and 6.7% and 5.7% for the validation set, respectively. Compared to the previous methods based on the coarse-to-fine approach, our method achieves 24.1% of the improvement in AP on the test set. In particular, for the test set, our method outperforms existing methods on instances that are typically located close to each other, such as pictures, desks, and bookshelves. This suggests that explicitly penalizing misclassified points around the edges of instances is helpful in RPM. Please refer to the supplementary material for more results about this.

S3DIS Table 3 illustrates the quantitative result on Area 5. Our proposed method outperforms the second best performing method with margins of 3.3 and 2.9 in mAP, AP_{50} , and mRec50, improving the performance of SOTA 5.7%, 4.2%, and 5.6% respectively.

STPLS3D Table 4 shows the quantitative comparison on the validation set of STPLS3D dataset. Our method outperforms all of the existing methods, improving SOTA performance in mAP and AP_{50} for 3.0 and 4.3, respectively.

4.5. Qualitative Results

Fig. 5 shows visual comparisons of ISBNet [29], MAFT [15], and our proposed Spherical Mask for challenging instances on ScanNetV2 validation set. Spherical Mask accurately segments large instances such as wall and book shelves (row 1), curtains and a window between them (row 2), a large sofa (row 3) and a circular shape sofa (row 4).

ISBNet [29] struggles to segment large instances (wall and bookshelves in row 1) and a disconnected instance (curtains and a window between them in row 2). On the other hand,

Method	mAP	mAP ₅₀	bath	bed	bk.shf	cabinet	chair	counter	curtain	desk	door	other	picture	fridge	s. cur.	sink	sofa	table	toilet	wind.
3D-BoNet(C)[33]	25.3	48.8	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
TD3D(C)[14]	48.9	75.1	85.2	51.1	43.4	32.2	73.5	10.1	51.2	35.5	34.9	46.8	28.3	51.4	67.6	26.8	67.1	51.0	90.8	32.9
SoftGroup(G)[30]	50.4	76.1	66.7	57.9	37.2	38.1	69.4	7.2	67.7	30.3	38.7	53.1	31.9	58.2	75.4	31.8	64.3	49.2	90.7	38.8
PBNet(G)[36]	57.3	74.7	92.6	57.5	61.9	47.2	73.6	23.9	48.7	38.3	45.9	50.6	53.3	58.5	76.7	40.4	71.7	55.9	96.9	38.1
DKNet(K)[31]	53.2	71.8	81.5	62.4	51.7	37.7	74.9	10.7	50.9	30.4	43.7	47.5	58.1	53.9	77.5	33.9	64.0	50.6	90.1	38.5
ISBNet(K)[29]	55.9	75.7	93.9	65.5	38.3	42.6	76.3	18.0	53.4	38.6	49.9	50.9	62.1	42.7	70.4	46.7	64.9	57.1	94.8	40.1
MAFT(T)[15]	57.8	78.6	88.9	72.1	44.8	46.0	76.8	25.1	55.8	40.8	50.4	53.9	61.6	61.8	85.8	48.2	68.4	55.1	93.1	45.0
QueryFormer(T)[21]	58.3	78.7	92.6	70.2	39.3	50.4	73.3	27.6	52.7	37.3	47.9	53.4	53.3	69.7	72.0	43.6	74.5	59.2	95.8	36.3
Ours(C)	61.6	81.2	94.6	65.4	55.5	43.4	76.9	27.1	60.4	44.7	50.5	54.9	69.8	71.6	77.5	48.0	74.7	57.5	92.5	43.6

Table 1. Quantitative comparison of top-performing methods for each approach on ScanNetV2 **hidden** test set. (C),(G),(K), and (T) next to the names of the methods refer to coarse-to-fine, grouping, kernel, and Transformer based methods, respectively. All the methods take the same input, such as point cloud and corresponding color information. The best results are in bold, and the second best ones are in underlined.

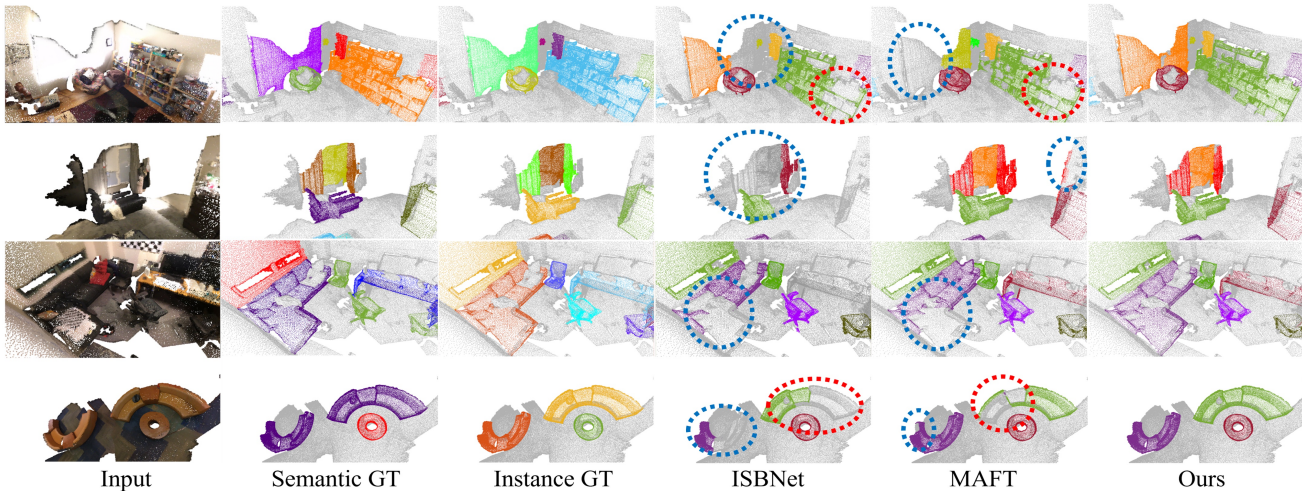


Figure 5. Qualitative comparison of ISBNet[29], MAFT[15], and ours on ScanNetV2 validation set.

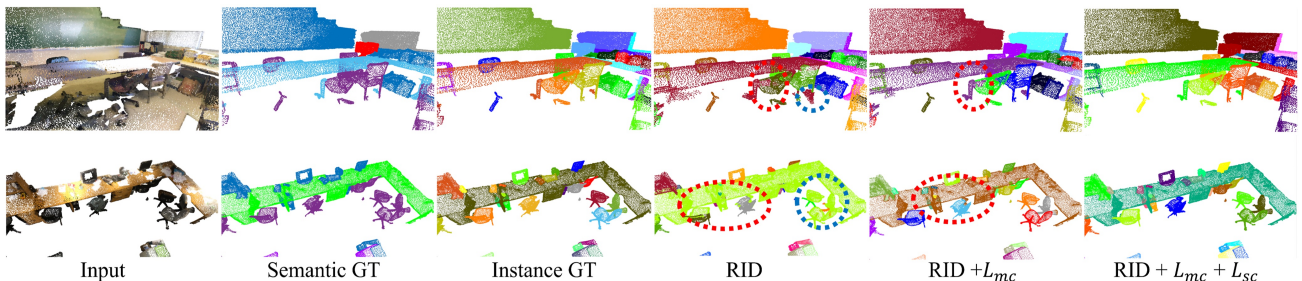


Figure 6. Visually comparing impacts of RID, L_{mc} , and L_{sh} .

MAFT[15] shows better generalization capability for learning semantics (curtains and a window between them (row 2)). However, it struggles to segment some parts of an instance that look different, as shown in sofas (row 2,3) and over-segments the instance by considering physically further away points as the same instance (wall in row 2 and sofa in row 4), probably due to the local queries that overfit to certain semantics.

4.6. Ablation Study

In this section, we investigate Spherical Mask with ablation studies designed for its core components.

Impact of θ and φ is shown at Table 6. For this experiment, we fixed the backbone and the RPM with L_{mc} and L_{sc} , and change N_θ and N_φ . In theory, increasing N_θ and N_φ should always produce better results. However, in practice, increasing them faces an issue of high complexity, as can be seen. Using 4/4 and 5/5 of N_θ and N_φ improves the mAP for 0.9% and 1.7%, respectively. However, increasing them from 5/5 to 6/6 results in 2.8% of the performance drop, suggesting that too large N_θ and N_φ make a negative impact on the performance due to increased complexity. For example, too large N_θ and N_φ would result in many sectors without any foreground points, leading to biased target values for learn-

Method	Venue	mAP	AP50	AP25
3D-SIS(C)[12]	CVPR19	-	18.7	35.7
GSPN(C)[34]	CVPR19	-	37.8	53.4
TD3D(C)[14]	WACV23	47.3	71.2	81.9
PointGroup(G)[13]	CVPR20	34.8	51.7	71.3
SSTNet(G)[19]	ICCV21	49.4	64.3	74.0
MaskGroup(G)[37]	ICME22	27.4	42.0	63.3
SoftGroup(G)[30]	CVPR22	46.0	67.6	78.9
RPGN(G)[8]	ECCV22	-	64.2	-
PBNet(G)[36]	ICCV23	54.3	70.5	78.9
PointInst3D(K)[11]	ECCV22	45.6	63.7	-
DKNet(K)[31]	ECCV22	50.8	66.9	76.9
ISBNet(K)[29]	CVPR23	54.5	73.1	82.5
Mask3D(T)[27]	ICRA23	55.2	73.7	82.9
3IS-ESSS(T)[1]	ICCV23	56.1	75.0	83.7
QueryFormer(T)[21]	ICCV23	56.5	74.2	83.3
MAFT(T)[15]	ICCV23	<u>58.4</u>	<u>75.6</u>	<u>84.5</u>
Ours(C)	-	62.3	79.9	88.2

Table 2. Quantitative 3D instance segmentation results on ScanNetV2 validation set. (C),(G),(K), and (T) next to the names of the methods refer to coarse-to-fine, grouping, kernel, and Transformer based methods, respectively. The best results are in bold, and the second best ones are in underlined.

Method	mAP	AP50	mPrec ₅₀	mRec ₅₀
GSPN [34]	-	-	36.0	28.7
PointGroup [13]	-	5.78	61.9	62.1
HAIS [4]	-	-	71.1	65.0
SSTNet [19]	42.7	59.3	65.6	64.2
SoftGroup [30]	51.6	66.1	73.6	66.6
Mask3D [27]	56.5	69.3	68.7	70.7
RPGN [8]	-	-	64.0	63.0
PointInst3D [11]	-	-	<u>73.1</u>	65.2
DKNet [31]	-	-	70.8	65.3
ISBNet [29]	56.3	67.5	70.5	72.0
PBNet [36]	53.5	66.4	74.9	65.4
QueryFormer [21]	<u>57.7</u>	<u>69.9</u>	70.5	<u>72.2</u>
MAFT [15]	-	69.1	-	-
Ours	60.5	72.3	71.3	76.3

Table 3. Quantitative 3D instance segmentation results on S3DIS Area 5. The best results are in bold, and the second best ones are in underlined.

Method	mAP	AP ₅₀	RID	L_{mc}	L_{sc}	mAP	AP ₅₀	AP ₂₅
HAIS [4]	35.0	46.7	✓			51.6	69.4	86.8
PointGroup [13]	23.3	38.5	✓	✓		58.5	77.6	87.5
ISBNet [29]	49.2	64.0	✓		✓	56.5	77.1	87.1
Ours	52.2	68.3	✓	✓	✓	62.3	79.9	88.2

Table 4. Quantitative instance segmentation results on STPLS3D

Table 5. Impact of each component on ScanNetV2 validation set

ing. Despite the variation, all configurations of N_θ and N_φ still outperform existing methods in Table 2, implying that RID with RPM always improves the performance.

Impact of Radial Point Migration is illustrated in Table 5. Here, we investigate the impact of the RPM and each loss for it. The baseline is the model trained with only RID, exclud-

N_θ/N_φ	mAP	AP ₅₀	AP ₂₅	Seed/Vote	mAP	AP ₅₀	AP ₂₅
1/1	59.6	77.1	85.3	1024/128	62.3	79.2	87.8
2/2	60.0	77.7	86.1	1024/256	62.3	79.9	88.2
3/3	60.6	78.4	87.1	1024/512	62.3	79.4	87.6
4/4	61.2	80.0	89.3	2048/128	62.0	79.7	88.0
5/5	62.3	79.9	88.2	2048/256	61.8	79.6	87.9
6/6	60.6	78.5	88.2	2048/512	61.7	79.1	87.5

Table 6. Impact of N_θ and N_φ on ScanNetV2 validation set

Table 7. Impact of seed and vote numbers on ScanNetV2 validation set

ing RPM. As the baseline model cannot refine the false positive points inside radial predictions or false negative points outside, as shown in Figure 6 (parts of chairs included as tables), its performance with high iou thresholds(AP, AP₅₀) are 17.1% and 13.1% worse than the full model. Including RPM with L_{mc} leads 13.3% of improvement in AP, suggesting refinement to push and pull misclassified points is crucial for the performance. L_{sc} shows 9.4% improvement in AP from the baseline, suggesting that focusing on true positive samples also contributes significantly to learning fine granularity and common features shared within an instance. As shown in Figure 6(column 5,6), adding L_{sc} improves the segmentation around the center of the objects, which was expected as the true positive samples near the instance centers could be neglected from L_{mc} as they are usually inside rays. As L_{mc} and L_{sc} target different samples, the full model combining both L_{mc} and L_{sc} shows 20.7% and 15.1% improvements for AP and AP₅₀, demonstrating the synergy of the two losses.

Impact of Voting Parameters is shown in Table 7. In this experiment, we change the seed and vote numbers inside the backbone while fixing RID and RPM. As can be seen, a seed number of 1024 produces more reliable results without variation than 2048, suggesting too many seed points negatively impact the system. Despite the slight difference, we observe that changing the seed and vote produces comparably small variations of ± 0.6 in mAP for all settings.

5. Conclusion

We present Spherical Mask, a novel coarse-to-fine approach for 3D instance segmentation in point cloud. As a coarse detection, the RID module finds instances as 3D polygons defined with center and rays. In contrast to existing coarse-to-fine approaches, the RPM module uses the polygons as soft references and migrates points efficiently in spherical coordinates to acquire final masks. We demonstrate how each module contributes to the instance segmentation and achieves state-of-the-art performances on public benchmarks ScanNet-V2, S3DIS, and STPLS3D.

Acknowledgement

This research was supported by ACE-OPS project (EP/S030832/1).

References

- [1] Salwa Al Khatib, Mohamed El Amine Boudjoghra, Jean Lahoud, and Fahad Shahbaz Khan. 3d instance segmentation via enhanced spatial and semantic supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 541–550, 2023. 1, 3, 8
- [2] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. 2, 6
- [3] Meida Chen, Qingyong Hu, Zifan Yu, Hugues Thomas, Andrew Feng, Yu Hou, Kyle McCullough, Fengbo Ren, and Lucio Soibelman. Stpls3d: A large-scale synthetic and real aerial photogrammetry 3d point cloud dataset. *arXiv preprint arXiv:2203.09065*, 2022. 2, 6
- [4] Shaoyu Chen, Jiemin Fang, Qian Zhang, Wenyu Liu, and Xinggang Wang. Hierarchical aggregation for 3d instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15467–15476, 2021. 1, 2, 6, 8
- [5] Bowen Cheng, Alex Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. *Advances in Neural Information Processing Systems*, 34:17864–17875, 2021. 2
- [6] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 2
- [7] 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. 2, 6
- [8] Shichao Dong, Guosheng Lin, and Tzu-Yi Hung. Learning regional purity for instance segmentation on 3d point clouds. In *European Conference on Computer Vision*, pages 56–72. Springer, 2022. 8
- [9] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9224–9232, 2018. 3
- [10] Tong He, Chunhua Shen, and Anton Van Den Hengel. Dyco3d: Robust instance segmentation of 3d point clouds through dynamic convolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 354–363, 2021. 1, 2, 4, 6
- [11] Tong He, Wei Yin, Chunhua Shen, and Anton van den Hengel. Pointinst3d: Segmenting 3d instances by points. In *European Conference on Computer Vision*, pages 286–302. Springer, 2022. 1, 2, 8
- [12] Ji Hou, Angela Dai, and Matthias Nießner. 3d-sis: 3d semantic instance segmentation of rgb-d scans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4421–4430, 2019. 1, 2, 8
- [13] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and Pattern recognition*, pages 4867–4876, 2020. 1, 2, 8
- [14] Maksim Kolodiaznyi, Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Top-down beats bottom-up in 3d instance segmentation. *arXiv preprint arXiv:2302.02871*, 2023. 1, 2, 7, 8
- [15] Xin Lai, Yuhui Yuan, Ruihang Chu, Yukang Chen, Han Hu, and Jiaya Jia. Mask-attention-free transformer for 3d instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3693–3703, 2023. 1, 3, 6, 7, 8
- [16] Loic Landrieu and Mohamed Boussaha. Point cloud oversegmentation with graph-structured deep metric learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7440–7449, 2019. 6
- [17] Loic 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*, pages 4558–4567, 2018. 6
- [18] Ville V Lehtola, Harri Kaartinen, Andreas Nüchter, Risto Kaijaluoto, Antero Kukko, Paula Litkey, Eija Honkavaara, Tomi Rosnell, Matti T Vaaja, Juho-Pekka Virtanen, et al. Comparison of the selected state-of-the-art 3d indoor scanning and point cloud generation methods. *Remote sensing*, 9(8):796, 2017. 1
- [19] Zhihao Liang, Zhihao Li, Songcen Xu, Minghui Tan, and Kui Jia. Instance segmentation in 3d scenes using semantic superpoint tree networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2783–2792, 2021. 1, 2, 6, 8
- [20] Hong Liu, Mingsheng Long, Jianmin Wang, and Yu Wang. Learning to adapt to evolving domains. *Advances in Neural Information Processing Systems*, 33:22338–22348, 2020. 1
- [21] Jiahao Lu, Jiacheng Deng, Chuxin Wang, Jianfeng He, and Tianzhu Zhang. Query refinement transformer for 3d instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 18516–18526, 2023. 1, 3, 6, 7, 8
- [22] Kyeong-Beom Park, Minseok Kim, Sung Ho Choi, and Jae Yeol Lee. Deep learning-based smart task assistance in wearable augmented reality. *Robotics and Computer-Integrated Manufacturing*, 63:101887, 2020. 1
- [23] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. 6
- [24] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017. 2, 3, 6
- [25] Charles R Qi, Or Litany, Kaiming He, and Leonidas J Guibas. Deep hough voting for 3d object detection in point clouds. In *proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9277–9286, 2019. 3, 6

- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015. [3](#)
- [27] Jonas Schult, Francis Engelmann, Alexander Hermans, Or Litany, Siyu Tang, and Bastian Leibe. Mask3d for 3d semantic instance segmentation. *arXiv preprint arXiv:2210.03105*, 2022. [1](#), [2](#), [8](#)
- [28] Jiahao Sun, Chunmei Qing, Junpeng Tan, and Xiangmin Xu. Superpoint transformer for 3d scene instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2393–2401, 2023. [1](#), [2](#)
- [29] Khoi Nguyen Tuan Duc Ngo, Binh-Son Hua. Isbnet: a 3d point cloud instance segmentation network with instance-aware sampling and box-aware dynamic convolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [30] Thang Vu, Kookhoi Kim, Tung M Luu, Thanh Nguyen, and Chang D Yoo. Softgroup for 3d instance segmentation on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2708–2717, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [31] Yizheng Wu, Min Shi, Shuaiyuan Du, Hao Lu, Zhiguo Cao, and Weicai Zhong. 3d instances as 1d kernels. In *European Conference on Computer Vision*, pages 235–252. Springer, 2022. [1](#), [2](#), [6](#), [7](#), [8](#)
- [32] Enze Xie, Peize Sun, Xiaoge Song, Wenhai Wang, Xuebo Liu, Ding Liang, Chunhua Shen, and Ping Luo. Polarmask: Single shot instance segmentation with polar representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12193–12202, 2020. [3](#)
- [33] Bo Yang, Jianan Wang, Ronald Clark, Qingyong Hu, Sen Wang, Andrew Markham, and Niki Trigoni. Learning object bounding boxes for 3d instance segmentation on point clouds. *Advances in neural information processing systems*, 32, 2019. [1](#), [2](#), [3](#), [7](#)
- [34] Li Yi, Wang Zhao, He Wang, Minhyuk Sung, and Leonidas J Guibas. Gspn: Generative shape proposal network for 3d instance segmentation in point cloud. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3947–3956, 2019. [1](#), [2](#), [8](#)
- [35] Biao Zhang and Peter Wonka. Point cloud instance segmentation using probabilistic embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8883–8892, 2021. [6](#)
- [36] Weiguang Zhao, Yuyao Yan, Chaolong Yang, Jianan Ye, Xi Yang, and Kaizhu Huang. Divide and conquer: 3d point cloud instance segmentation with point-wise binarization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 562–571, 2023. [1](#), [2](#), [7](#), [8](#)
- [37] Min Zhong, Xinghao Chen, Xiaokang Chen, Gang Zeng, and Yunhe Wang. Maskgroup: Hierarchical point grouping and masking for 3d instance segmentation. In *2022 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2022. [8](#)
- [38] Dingfu Zhou, Jin Fang, Xibin Song, Liu Liu, Junbo Yin, Yuchao Dai, Hongdong Li, and Ruigang Yang. Joint 3d instance segmentation and object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1839–1849, 2020. [1](#)