

# 3D Instances as 1D Kernels<sup>\*</sup>

Yizheng Wu<sup>1</sup>, Min Shi<sup>1</sup>, Shuaiyuan Du<sup>1</sup>, Hao Lu<sup>1</sup>, Zhiguo Cao<sup>1</sup>, and  
Weicai Zhong<sup>2</sup>

<sup>1</sup> Key Laboratory of Image Processing and Intelligent Control, Ministry of Education  
School of AIA, Huazhong University of Science and Technology, China

<sup>2</sup> Huawei CBG Consumer Cloud Service Search & Maps BU  
{yzwu21,min.shi,sydu,hlu,zgcao}@hust.edu.cn   zhongweicai@huawei.com

**Abstract.** We introduce a 3D instance representation, termed *instance kernels*, where instances are represented by one-dimensional vectors that encode the semantic, positional, and shape information of 3D instances. We show that instance kernels enable easy mask inference by simply scanning kernels over the entire scenes, avoiding the heavy reliance on proposals or heuristic clustering algorithms in standard 3D instance segmentation pipelines. The idea of instance kernel is inspired by recent success of dynamic convolutions in 2D/3D instance segmentation. However, we find it non-trivial to represent 3D instances due to the disordered and unstructured nature of point cloud data, *e.g.*, poor instance localization can significantly degrade instance representation. To remedy this, we construct a novel 3D instance encoding paradigm. First, potential instance centroids are localized as candidates. Then, a candidate merging scheme is devised to simultaneously aggregate duplicated candidates and collect context around the merged centroids to form the instance kernels. Once instance kernels are available, instance masks can be reconstructed via dynamic convolutions whose weights are conditioned on instance kernels. The whole pipeline is instantiated with a dynamic kernel network (DKNet). Results show that DKNet outperforms the state of the arts on both ScanNetV2 and S3DIS datasets with better instance localization. Code is available: <https://github.com/W1zheng/DKNet>.

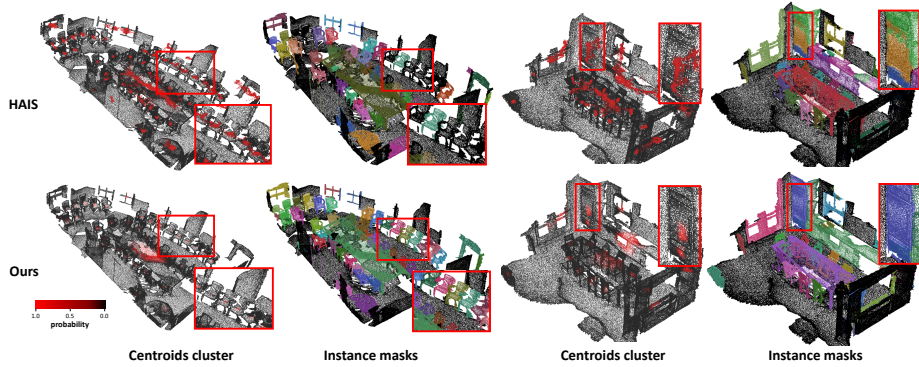
**Keywords:** Instance kernel, point cloud, instance segmentation

## 1 Introduction

3D Instance segmentation aims to predict point-level instance labels [8,12]. Standard approaches heavily rely on proposals [28,4,17] or heuristic clustering algorithms [12,2]. In this work, we show that instance masks can be reconstructed by scanning a scene with *instance kernels*, a representation for 3D instances, which simultaneously encodes the positional, semantic, and shape information of 3D instances.

---

<sup>\*</sup> Y. Wu and M. Shi contributed equally. Z. Cao is the corresponding author.



**Fig. 1. Comparison of inferred centroid clusters and instance masks.** Compared with HAIS [2], our DKNet generates more focused centroid clusters that can guide precise localization such that small and close instances can be discriminated and large instances have consistent predictions. Best viewed by zooming in and in color.

3D instance representation addresses two fundamental problems: i) how to localize an instance precisely and ii) how to aggregate features effectively to depict the instance. Unlike 2D instances that can be directly encoded via grid sampling [26] or dynamic kernel assigning [31], in the 3D domain, the disordered and unstructured nature of point cloud data renders difficulties for precise instance localization and reliable representation; and top-performing approaches [2,16,12] implicitly localize instances with centroid offsets [22], which only provides coarse information for instance representation, as shown in Fig. 1. Our instance kernel also draws inspiration from DyCo3D [9], which first applies dynamic convolution into 3D instance segmentation. However, DyCo3D is built upon the existing bottom-up segmentation pipeline [12], leaving the fundamental problems of instance encoding unsolved. To alleviate the difficulties above, we design a novel instance encoding paradigm that efficiently localizes different instances and encodes the semantic, positional, and shape information of instances into *instance kernels* for mask generation.

We further incorporate the kernel encoding paradigm into a dynamic kernel network (DKNet) for 3D instance segmentation. To localize instances, as shown in Fig. 1, DKNet predicts a centroid map for instances and extracts centroids via a customized Non-Maximum Suppression (NMS) operator with local normalization. Observing that duplicated candidates may be predicted for a single instance (especially for large ones), we design an iterative aggregation mechanism to merge duplicated candidates guided by a predicted merging score map. The score map indicates the probability whether each paired candidates should be merged. Afterwards, the merged instances are encoded into instance kernels by adaptively fusing the point features around the localized instance centroids. Finally, instance masks can be reconstructed with a few convolution layers, whose weights are conditioned on the generated instance kernels.

We evaluate DKNet on two popular 3D instance segmentation datasets, including ScanNetV2 [3] and S3DIS [1]. The results show that DKNet outperforms previous state-of-the-art approaches, ranking the first AP among published methods on the ScanNetV2 online leaderboard.<sup>1</sup> Thanks to the instance kernels and the specially designed instance localization pipeline, DKNet can better distinguish instances from dense areas than current top-performing approaches, as shown in Fig. 1. A series of ablation studies also demonstrate that the proposed instance localization and aggregation pipeline can greatly enhance the instance representation.

Our contributions are two-fold:

- We extend the idea of dynamic convolution into instance kernel, a comprehensive representation for 3D instances in point clouds;
- We propose a dynamic kernel network for 3D instance segmentation, with a novel instance kernel encoding paradigm;

## 2 Related Work

Here we briefly review the 3D instance segmentation approaches and kernel-based instance segmentation.

**Proposal-Based 3D Instance Segmentation.** Proposal-based approaches [29] assign instances with proposals and the instance masks are generated upon proposals. 3D-BoNet [28] directly predicts the bounding boxes, within instance masks are generated. 3D-MPA [4] samples proposals from predicted centroids; masks of proposals are then clustered to form the instance masks. GICN [17] simultaneously predicts the centroids and sizes of instances to obtain bounding box proposals. Predictions from proposal-based approaches show good objectness, while two major drawbacks exist: 1) the multi-stage training and the proposal generation process introduce large computational overhead; 2) the results highly rely on the proposals.

**Proposal-Free 3D Instance Segmentation.** Proposal-free approaches cluster points into instances in a bottom-up manner. SSTNet [16] models the entire scene by constructing a tree of superpoints and uses top-bottom traversal to aggregate nodes and form instance masks. PointGroup [12] clusters points using semantic labels and centroid offsets as clues. PE [30] encodes points into an embedding space where points from the same instances are close. Then clustering are performed in this embedding space. Considering that clustering points into instances with various sizes in one shot is difficult, HAIS [2] proposes a novel hierarchical clustering pipeline to gradually refine the aggregation results. However, even with the implicit guide of object signals, the objectness of predictions is still low, as shown in Fig. 1. Hence, directly adding kernel-based dynamic convolution modules upon existing proposal-free approaches cannot bring out the best of kernel-based instance segmentation paradigm.

<sup>1</sup> [http://kaldir.vc.in.tum.de/scannet\\_benchmark/semantic\\_instance\\_3d](http://kaldir.vc.in.tum.de/scannet_benchmark/semantic_instance_3d)

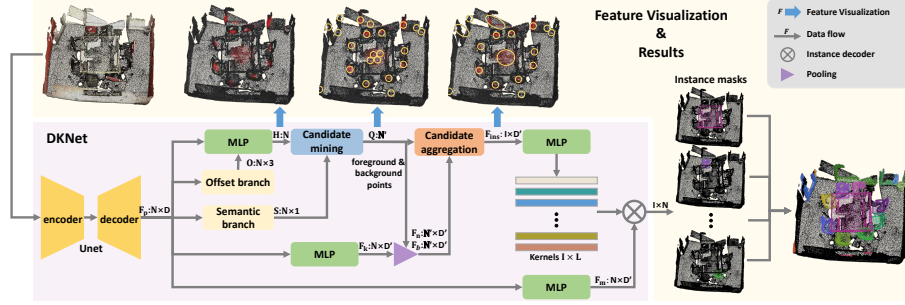


Fig. 2. Pipeline of the Dynamic Kernel Network.

**Kernel-Based Instance Segmentation.** Kernel-based instance segmentation uses instance-aware kernels to scan the whole scene to reconstruct instance masks, the pivot of which is to represent or associate instances with different kernels. After obtaining the kernels, the common solutions are scanning the scene via dot product or dynamic convolution [11,19]. CondInst [24] predicts instance proposals by object detection and encodes the proposal features into kernels. K-Net [31] associates a fixed number of kernels with instances in certain regions and dynamically updates them. SOLOv2 [26] partitions the feature maps into grids and generates a kernel for each grid. However, representing instances as kernels in the 3D domain is non-trivial due to the disordered and unstructured nature of point cloud data. DyCo3D [9] first introduce the kernel-based paradigm in 3D instance segmentation, which is built upon existing bottom-up approach [12]. However, they focus on the concrete implementation of dynamic convolutions, and bypass the core of the kernel-based paradigm: how to encode instances into kernels? In this work, we further explore the underlying relation between the discriminative representation of instances and effective kernel-based segmentation, resulting in a novel localize-then-aggregate instance kernel encoding paradigm.

### 3 Dynamic Kernel Network for 3D Instance Segmentation

#### 3.1 Overview

As illustrated in Fig. 2, at the core of our dynamic kernel network (DKNet) is to encode instances into discriminative instance kernels. The encoding process consists of three key stages: 1) processing raw point clouds with a UNet-like backbone and predicting point features, centroid offsets, and semantic masks; 2) localizing centroids for instances with a candidate mining branch; 3) merging duplicated candidates and collecting context around instance centroids to form instance kernels. Once the instance kernels are acquired, the instance masks can be obtained by processing the point cloud features with a few convolution layers, whose weights are conditioned on instance kernels.

### 3.2 Point-wise Feature Extraction

Following recent proposal-free approaches [12,2,16], we adopt the backbone from PointGroup [12] for feature extraction. Given the raw point cloud  $P \in \mathbb{R}^{N \times 6}$  with  $N$  points, a 3D UNet-like [23] backbone with sparse convolutions [5] outputs point features  $F_p \in \mathbb{R}^{N \times D}$ .  $F_p$  is then fed to a semantic branch which predicts semantic mask  $S \in \mathbb{R}^{N \times C}$ , and additionally, a centroid offset branch which infers the offset  $O \in \mathbb{R}^{N \times 3}$  of each point to the corresponding instance centroid. The semantic branch is a Multi-Layer Perceptron (MLP) with **softmax** activation at the output layer. Cross-entropy loss and multi-class dice loss [20] are used to supervise the training of this branch. Similar to the semantic branch, the centroid offset branch maps  $F_p$  into offsets  $O$ . For each point  $P_i$ ,  $O_i \in \mathbb{R}^3$  is a vector pointing to the centroid of instance that covers this point. Further details of the backbone can be referred to the supplementary.

### 3.3 Finding Instances

To generate the kernel for each instance, we should first find all the instances. However, the top-performing proposal-free approaches predict centroid offsets as implicit object signals, which is rather coarse to precisely localize instances, as shown in Fig. 1. Hence, learning from proposal-based approaches, we propose a candidate mining branch to generate centroid maps, followed by a searching algorithm to localize the instance candidates.


As shown in Fig. 3, point features  $F_p \in \mathbb{R}^{N \times D}$  and centroid offsets  $O \in \mathbb{R}^{N \times 3}$  are jointly concatenated to form the input  $F_c \in \mathbb{R}^{N \times (D+3)}$  for the centroid mining branch. Then  $F_c$  is fed into an MLP with **softmax** activation at the output layer to obtain sharp centroid heatmap  $H \in \mathbb{R}^N$ . Each element  $H_i$  indicates the probability of the  $i^{th}$  point being an instance centroid.

During training, we place a 3D Gaussian kernel on every instance centroid to form a pseudo ground truth heatmap as  $\hat{H}_i = \exp(-\alpha \cdot d_i^2 / r_i^2)$ , where  $d_i$  denotes the distance between point  $i$  to the centroid of the instance covering it.  $r_i$ , which controls the variance of the Gaussian kernel, equals the maximum side length of the axis-aligned bounding box of the corresponding instance. Hence, the Gaussian kernels are geometry adaptive w.r.t. the size of different instances.  $\alpha$  is set to 25 to keep the average of  $\hat{H}$  around 0.1. To supervise the training, the loss function  $\mathcal{L}_{center}$  for the candidate mining branch is defined as:

$$\mathcal{L}_{center} = \frac{1}{\sum_{i=1}^N \mathbb{I}(P_i)} \sum_{i=1}^N |H_i - \hat{H}_i| \cdot \mathbb{I}(P_i), \quad (1)$$

where  $\mathbb{I}(P_i)$  is an indicator function that outputs 1 when  $i^{th}$  point belongs to an instance, otherwise outputs 0.

With the predicted heatmap  $H$ , we iteratively search the local maximum as instance candidates with a customized local normalized NMS (LN-NMS) strategy. During each iteration, the algorithm localizes the point with the highest centroid score among the foreground points; the centroid scores of other points



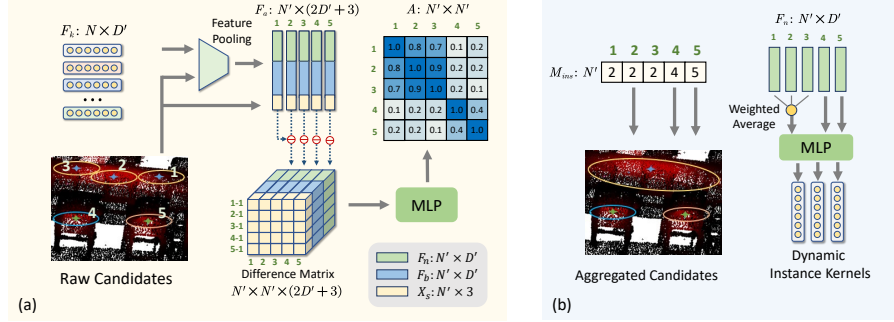
**Fig. 3. Centroid mining branch.** (a) The input of candidate mining branch; (b) The customized Non-Maximum Suppression with local normalization.

in its neighbor with radius  $R$  are then normalized via the division of the maximum value in this  $R$ -radius neighbor. If the normalized centroid score is larger than a threshold  $T_\theta$ , this point will be considered a candidate and all other points within its  $R$ -radius neighbor are suppressed and excluded in the next iteration, no matter whether the point is chosen as the candidate. We set  $T_\theta = 0.5$  and the radius  $R = 0.3m$  according to the average size of instances in ScanNet [3]. The iteration ends when no point remains or  $N_\theta$  candidates have been found.  $N_\theta$  is empirically set to 200. Finally, a candidate set  $Q \in \mathbb{R}^{N'}$  can be collected, where  $N'$  denotes the number of candidates. Refer to supplementary for more details.

### 3.4 Representing Instances as Kernels

After localizing the instance centroids, we represent these candidates as instance kernels. We expect that one kernel is extracted for one instance, and the kernel should be discriminative. Therefore, we design a duplicated candidate aggregation strategy that simultaneously eliminates extra candidates and adaptively fuses features around candidates for instance representation.

**Aggregating Duplicate Candidates.** We judge whether two candidates should be aggregated based on the context of each candidate. For each raw candidate, we use the features from its “foreground points”, “background points” to describe the context. The “foreground points” denote points with the same semantic label within a  $R$ -radius neighbor of each candidate, while the “background points” denote all the points with different semantic labels within a  $2R$ -radius neighbor of each candidate. To aggregate features from the foreground and background points, we first process point feature  $F_p$  with an MLP for dimensionality reduction. Then, the output features of the MLP w.r.t. “foreground points” and “background points” are averaged and respectively form the descriptive feature  $F_n \in \mathbb{R}^{N' \times D'}$  and the background feature  $F_b \in \mathbb{R}^{N' \times D'}$  for each candidate. As the above two features only encode semantic and shape information, we concatenate them with the shifted coordinate (add the raw coordinates with the centroid offset vectors) of each candidate as positional information to form the aggregation feature  $F_a \in \mathbb{R}^{N' \times (2D'+3)}$  for duplicated candidates aggregation.



**Fig. 4. Candidates aggregation module.** (a) The process of predicting the merging score map; (b) The process of generating instance kernels.

As shown in Fig. 4, for each candidate, its aggregation feature  $F_{a,i}$  is subtracted from the aggregation feature of all other candidates. By repeating this process for every instance, a candidate difference matrix reflecting how similar each pair of candidates will be generated. Taking the absolute values of the difference matrix as inputs, an MLP with **sigmoid** function outputs the merging score map  $A \in \mathbb{R}^{N' \times N'}$ , where  $A_{ij}$  indicates the probability that the  $i^{th}$  and  $j^{th}$  candidates shall be merged because they can belong to the same instance.

Once the merging score map  $A$  is obtained, a simple greedy algorithm will be used to iteratively merge candidates. We first initialize an instance centroid map  $M_{ins} \in \mathbb{R}^N$  where  $M_{ins,i} = i$ .  $M_{ins}$  records the indices of instance centroids that each candidate belongs to, and we define candidates with the same index as an instance *group*. Before aggregation, the instance centroids of candidates are themselves, and each candidate is an instance *group*. During each iteration, if  $A_{ij}$  is the maximum in  $A$  excluding diagonal elements, all candidates in the  $i^{th}$  and  $j^{th}$  instance *groups* will be merged. The instance centroid indices of these candidates will also be unified to the index of candidate with the highest centroid score among them. Since candidates within the same instance group can no longer be merged, we then update all merging scores between them to be 0. The iteration ends when all merging scores are below a predefined threshold, which is set to 0.5. The candidates with the same index are treated as a predicted instance. After aggregation, supposing  $I$  instance groups are generated, the centroid coordinates  $C_{ins} \in \mathbb{R}^{I \times 3}$  of new instances are set to the coordinates of the center candidates, while the features of new instances  $F_{ins} \in \mathbb{R}^{I \times D'}$  are obtained by a weighted average upon descriptive features of grouped candidates w.r.t. their neighbor sizes. The average can help dynamically aggregate the features of instances with different sizes. As shown in Fig. 4, if multiple candidates are predicted on one large instance, information from all the candidates will be propagated to the instance centroids to describe instances.

During training, it is easy to figure out which pairs of candidates are duplicated, and accordingly a ground truth merging map  $\hat{A}$  can be generated. The



standard binary cross entropy loss (BCELoss) is adopted as the loss function  $\mathcal{L}_{aggre}$  for candidate aggregation, which is defined as:

$$\mathcal{L}_{aggre} = BCELoss(A, \hat{A}), \quad (2)$$

where  $\hat{A}_{i,j} = 1$  if the candidate  $q_i$  and  $q_j$  belong to the same ground truth instance, and 0 otherwise.

Aside from aggregating duplicated candidates, by guiding the network to distinguish whether two candidates need to be merged, the representation ability of point feature  $F_p$  can also be enhanced, as the learning of aggregation encourages points from the same instance to be close in the feature space, and vice versa, similar to the idea of contrast learning [7].

**Encoding Instance Kernels.** After candidate aggregation, all instances in the scene are assigned with centroids, denoted by  $C_{ins}$ , and corresponding features, denoted by  $F_{ins}$ .  $F_{ins}$  is fed into an MLP to generate instance kernels  $\mathcal{W} \in \mathbb{R}^{I \times L}$ , where  $L$  is the length of instance kernels. In analogous to CondInst [24] and DyCo3D [9], the instance kernels are transformed into the weights for a few convolution layers in the instance decoding stage. Hence,  $L$  depends on the specific configurations of the convolution layers, which can be computed by Eq. 4.

### 3.5 Generating Masks with Instance Kernels

The instance kernels, denoted by  $\mathcal{W}$ , have encoded the positional, semantic, and shape characteristics of instances. To decode instances, the instance kernels are transformed into the weights with a few convolution layers, which are applied to augmented point cloud features to reconstruct instance masks.

To augment the point cloud features, an MLP further extracts the mask feature  $F_m \in \mathbb{R}^{N \times D'}$  from the point feature  $F_p$ . To inject instance-aware positional information into  $F_m$ , inspired by DyCo3d [9], the offset between each point to the instance centroids are added to  $F_m$  before convolution. *E.g.*, for each point  $P_i$ , we compute its offset to the centroid of the  $k^{th}$  instance as  $Z_{k,i} = C_{ins,k} - X_i$ . Then, the point decoding feature  $F_d \in \mathbb{R}^{N \times (D'+3)}$  for the  $k^{th}$  instance is generated by concatenating  $F_m$  and  $Z_k$  along the channel dimension. Although DyCo3D also generates instance masks via instance-specific kernels, the kernels are only applied to the points within the same semantic category. In contrast, the instance kernels in our approach scan the entire scene, which avoids the reliance of semantic prediction. Hence, the decoding process can correct some errors in semantic prediction. The bottom-up approaches [12,2] cannot correct such errors. The decoding process outputs the instance mask  $M \in \mathbb{R}^{I \times N}$  by

$$M_i = Conv(F_d, \mathcal{W}_i), \quad i \in [1, I], \quad (3)$$

where  $\mathcal{W}_i \in \mathbb{R}^L$  is transformed into the weights and biases via two  $1 \times 1$  convolution layers. The first layer has 16 output channels with ReLU activation function and the second one has 1 output channel with `sigmoid` for mask decoding. To



fit the number of parameters of convolution layers, the length of the instance kernel  $L$  can be computed by

$$L = (16 + 3) \times 16(\text{weight}) + 16(\text{bias}) + 16 \times 1(\text{weight}) + 1(\text{bias}) = 337. \quad (4)$$

The use of dynamic convolution is the same as [9,24]. Implementation details are depicted in the supplementary.

To supervise the generation of instance masks, we first match the predicted instances with actual instances with Hungarian algorithm [13] according to a cost matrix. Then, we apply the BCELoss and dice loss [20] for supervision. Supposing  $M \in \mathbb{R}^{I \times N}$  with  $I$  instance masks is generated, and  $\hat{M} \in \mathbb{R}^{G \times N}$  with  $G$  ground-truth instance masks is provided, the cost matrix  $\mathcal{C} \in \mathbb{R}^{I \times G}$  for the Hungarian algorithm is obtained by:

$$\mathcal{C}_{i,j} = \|C_{ins,i} - C_{gt,j}\|_2 + \mathbb{I}(S_{ins,i} == S_{gt,j}), \quad (5)$$

where  $C_{ins}$  and  $C_{gt}$  are centroid coordinates of predicted and ground truth instances, respectively, and  $S_{ins}$  and  $S_{gt}$  are corresponding semantic labels. We determine the instance semantic label  $S_{ins}$  by voting within the predicted instance. With the cost matrix, one predicted instance mask is expected to be matched to a instance with the closet centroid and identical semantic label.

After the matching process, the predicted instance mask  $M$  is assigned with the ground truth instance masks  $\hat{M} \in \mathbb{R}^{G \times N}$ . Then, BCELoss  $\mathcal{L}_{bce}$  and dice loss  $\mathcal{L}_{dice}$  are computed by:

$$\mathcal{L}_{mask} = \frac{1}{I'} \sum_{k=1}^I (BCE(M_k, \hat{M}_k) + (1 - 2 \frac{M_k \cdot \hat{M}_k}{|M_k| + |\hat{M}_k|})) \cdot \mathbb{I}(iou_k > 0.25), \quad (6)$$

where  $\hat{M}_k$  denotes the ground truth instance mask for the  $k^{th}$  instance, and  $M_k$  denotes the predicted one.  $iou_k$  denotes the Intersection-over-Union (IoU) between  $M_k$  and  $\hat{M}_k$ , and  $\mathbb{I}$  is an indicator function. We add the constraints so that the loss will only be computed when instances are correctly matched and  $I'$  is defined by  $I' = \sum_{k=1}^I \mathbb{I}(iou_k > 0.25)$ .

**Inference Post-processing.** During inference, to convert the soft instance masks  $M_k$ 's into hard instance labels and filter out potentially wrong predictions, a simple yet effective two-stage refinement pipeline is proposed. First, some small duplicated fragments or noise is removed. In the second stage, superpoints [15] are applied to refine the shapes of generated instance masks. Unlike proposal-free approaches [12,9,2], our approach do not need NMS or ScoreNet in post-processing, which is efficient.

In the first stage, given the predicted soft instance mask  $M \in \mathbb{R}^{I \times N}$ , we first generate the raw instance label by selecting the label of the instance with the highest score in  $M$ . Then, we define a coverage score  $S_c \in \mathbb{R}^I$  by:

$$S_{c,k} = N_{inter,k} / N_{intra,k}, \quad (7)$$

**Table 1. Quantitative results on the ScanNetV2 test set.** Refer to supplementary full results with all the 20 categories.

approaches	$mAP$	$AP@50$	bed	booksh.	cabinet	chair	curtain	desk	door	otherfu.	picture	refrige.	sofa	table	window
3D-BoNet[28]	25.3	48.8	67.2	59.0	30.1	48.4	62.0	30.6	34.1	25.9	12.5	43.4	49.9	51.3	43.9
3D-SIS[10]	16.1	38.2	43.2	24.5	19.0	57.7	26.3	3.3	32.0	24.0	7.5	42.2	69.9	27.1	23.5
MTML[14]	28.2	54.9	80.7	58.8	32.7	64.7	81.5	18.0	41.8	36.4	18.2	44.5	68.8	57.1	39.6
3D-MPA[4]	35.5	61.1	83.3	76.5	52.6	75.6	58.8	47.0	43.8	43.2	35.8	65.0	76.5	55.7	43.0
PointGroup[12]	40.7	63.6	76.5	62.4	50.5	79.7	69.6	38.4	44.1	55.9	47.6	59.6	75.6	55.6	51.3
GICN[17]	34.1	63.8	<b>89.5</b>	80.0	48.0	67.6	73.7	35.4	44.7	40.0	36.5	70.0	83.6	59.9	47.3
DyCo3D[9]	39.5	64.1	84.1	<b>89.3</b>	53.1	80.2	58.8	44.8	43.8	53.7	43.0	55.0	76.4	65.7	56.8
Occuseg[6]	48.6	67.2	75.8	68.2	57.6	84.2	50.4	52.4	56.7	58.5	45.1	55.7	79.7	56.3	46.7
PE[30]	39.6	64.5	77.3	79.8	53.8	78.6	79.9	35.0	43.5	54.7	54.5	64.6	76.1	55.6	50.1
SSTNet[16]	50.6	69.8	69.7	88.8	55.6	80.3	62.6	41.7	55.6	58.5	70.2	60.0	72.0	69.2	50.9
HAIS[2]	45.7	69.9	84.9	82.0	67.5	80.8	75.7	46.5	51.7	59.6	55.9	60.0	76.7	67.6	56.0
SoftGroup[25]	50.4	<b>76.1</b>	80.8	84.5	<b>71.6</b>	86.2	<b>82.4</b>	<b>65.5</b>	62.0	<b>73.4</b>	69.9	<b>79.1</b>	<b>84.4</b>	<b>76.9</b>	59.4
<b>Ours</b>	<b>53.2</b>	71.8	81.4	78.2	61.9	<b>87.2</b>	75.1	56.9	<b>67.7</b>	58.5	<b>72.4</b>	63.3	81.9	73.6	<b>61.7</b>

where  $N_{inter,k}$  and  $N_{intra,k}$  denote the number of “inter-point” and “intra-point” for the  $k^{th}$  instance, respectively. The “inter-point” denotes the number of points being assigned to the  $k^{th}$  instance in the raw instance label, while the “intra-point” denotes the number of points in the soft instance mask of the  $k^{th}$  instance that are above a threshold  $T_{m,k}$ .  $T_{m,k}$  is determined by the Otsu algorithm [21] that is adaptive to different instances. The cover score indicates the completeness and independence of each instance prediction. We then multiply the coverage score  $S_c$  with  $M$  to generate the refined soft instance mask. By taking the instance with the highest score in the refined mask, the final hard instance labels can be obtained. Refer to supplementary for more details on computing  $T_m$ .

In the second stage, we first aggregate the raw point clouds into superpoints [15]. Points within a certain superpoint should belong to the same instance. Hence, we unify the instance label in each superpoint to be the one that most points belong to. Aside from the instance label  $R$ , we need to assign a confidence score for each predicted instance that indicates the prediction quality for evaluation. This score is obtained by multiplying the average of instance score and semantic score of the “intra-point”.

## 4 Experiments

In this section, we first compare the proposed Dynamic Kernel Network (DKNet) with other state-of-the-art approaches on two 3D instance segmentation benchmarks: ScanNetV2[3] and S3DIS [1]. Then, we verify the effectiveness of different components in DKNet via a controlled ablation study.

### 4.1 Implementation Details

**Training Details.** For data preparation, the coordinates and colors are concatenated together to form 6D vectors for each point. The network is trained

**Table 2. Object detection results on ScanNetV2 validation set.** We report per-class mAP with an IoU of 25% and 50%. The IoU is computed on bounding boxes. We evaluate the performance of HAIS with the provided model. “Ours<sup>-</sup>” denotes the model without candidates aggregation part.

approach	$AP@25$	$AP@50$
VoteNet [22]	58.6	33.5
3DSIS [10]	40.2	22.5
3D-MPA [4]	64.2	49.2
DyCo3D [9]	58.9	45.3
HAIS [2]	66.0	54.2
Ours <sup>-</sup>	65.4	57.9
<b>Ours</b>	<b>67.4</b>	<b>59.0</b>

**Table 3. Quantitative results on S3DIS dataset.** We report mCov, mWCov, mPre, and mRec. approaches with ‡ marks are evaluated on scenes in Area-5. The others are evaluated via 6-fold cross validation.

approach	mCov	mWCov	mPre	mRec
PointGroup <sup>‡</sup> [12]	-	-	61.9	62.1
DyCo3D <sup>‡</sup> [9]	63.5	64.6	64.3	64.2
SSTNet <sup>‡</sup> [16]	-	-	65.5	64.2
HAIS <sup>‡</sup> [2]	64.3	<b>66.0</b>	<b>71.1</b>	65.0
<b>Ours<sup>‡</sup></b>	<b>64.7</b>	65.6	70.8	<b>65.3</b>
OccuSeg [6]	-	-	72.8	60.3
GICN [17]	-	-	68.5	50.8
PointGroup [12]	-	-	69.6	69.2
SSTNet [16]	-	-	73.5	<b>73.4</b>
HAIS [2]	67.0	70.4	73.2	69.4
SoftGroup [25]	69.3	71.7	<b>75.3</b>	69.8
<b>Ours</b>	<b>70.3</b>	<b>72.8</b>	<b>75.3</b>	71.1


on a single RTX 3090 GPU with a batch size of 4 for 400 epochs. We use the AdamW [18] optimizer with an initial learning rate of 0.001, which is adjusted by a cosine scheduler [27] during training. Weight decay is set to  $1e-5$ . Following previous methods [12], we voxelize the point clouds with the size of  $0.02m$  for ScanNetV2 and  $0.05m$  for S3DIS.

The overall training loss combines the loss from the semantic prediction, offset prediction, candidate mining, candidate aggregation branch, and the mask generation process, which can be defined as:

$$\mathcal{L} = \mathcal{L}_{sem} + \mathcal{L}_{off} + \mathcal{L}_{center} + \mathcal{L}_{aggre} + \mathcal{L}_{mask}, \quad (8)$$

where  $\mathcal{L}_{sem}$  and  $\mathcal{L}_{off}$  are the losses for semantic segmentation and centroid offsets prediction, respectively.

**Datasets.** We use ScanNetV2 and S3DIS for training and evaluation. ScanNetV2 includes 1,613 scenes with 20 different semantic categories. 1,201, 312 and 100 scenes are selected as the training, validation and test set, respectively. Note that the labels for test set are hidden for a fair comparison. Following official evaluation protocol, we use mean average precisions (mAPs) under different IoU thresholds as the evaluation metrics.  $AP@25$  and  $AP@50$  denote the average precision scores with IoU thresholds set to 25% and 50%.  $mAP$  denotes the average of all the  $AP$ s with IoU thresholds ranging from 50% to 95% with a step size of 5%. S3DIS dataset consists of 271 scenes collected from 6 different areas with 13 different object categories. Following previous approaches [12,9,2,16], we train and evaluate our approach in two ways: 1) scenes from area-5 are used for testing while scenes in other areas are used for training; 2) 6-fold cross validation where each area is used in turn for testing. On S3DIS, with the threshold IoU set



**Fig. 5. Qualitative results of our approaches on ScanNetV2 validation set.** We highlight the key details with red marks. Best viewed by zooming in and in color.

to 0.5, we report coverage (mCov), weighted coverage (mWCov), mean precision (mPrec), and mean recall (mRec) as evaluation metrics.

## 4.2 Comparison with the state of the arts.

**ScanNetV2.** Comparisons with the state of the arts on the ScanNetV2 test set are shown in Table 1. Our approach achieves an  $mAP$  of 53.2%, outperforming previous state-of-the-art approaches. The proposed DKNet obtains significant improvement on small instances like chairs or pictures, and competitive results on large instances like beds or tables. We also notice that, compared with a recent well-designed bottom-up method SoftGroup [25], DKNet shows inferior  $AP@50$ . The plausible reason are two fold. First, we find the DKNet retains relative high AP under strict IoU thresholds ( $>0.6$ ), indicating that the predicted masks can well preserve the instance shapes. However, under lower IoU thresholds ( $<0.6$ ), DKNet becomes less advantageous. Second, our pure data-driven candidate merging process shows mistakes on some difficult scenes, such as bookshelves with vague boundaries, which can be better tackled by well-designed bottom-up clustering. These results suggest that DKNet can be improved with careful design to further boost the potential of the proposed kernel-based paradigm.

To evaluate the instance localization performance of 3D instance segmentation approaches, we compare different 3D instance approaches under object detection metrics on ScanNetV2 validation set in Table 2. Predicted masks are converted into axis-aligned bounding boxes following DyCo3D [9]. Our approach achieves the best performance in  $AP@25$  of 67.4% and  $AP@50$  of 59.0%, which demonstrate that kernels are extracted from solid instance localization results.

**Table 4.** Comparison of different candidate mining algorithms on ScanNetV2 validation set.

Strategy	$AP$	$AP@50$	$AP@25$
Random	48.0	63.7	73.5
NMS	49.7	64.6	75.7
LN-NMS	50.8	66.7	76.9

**Table 5.** Comparison of different instance aggregation strategies on ScanNetV2 validation set.

Phase	$AP$	$AP@50$	$AP@25$
W/o	47.7	62.6	74.9
In training	48.4	64.5	76.1
All phases	50.8	66.7	76.9

**S3DIS.** As shown in Table 3, DKNet is comparable on scenes in area-5, while outperforms other approaches in 3 out of 4 metrics on 6-fold cross validation. Since most categories in S3DIS are large ones like ceiling, wall, and bookcase, the results can reflect the robustness of the proposed instance encoding paradigm on large instances.

### 4.3 Qualitative Evaluation

We visualize the predicted masks in Fig. 5. As is shown in column 2, excluding candidate aggregation (baseline) leads to severe over-segmentation as multiple kernels will be generated for one instance. The over-segmentation can be effectively alleviated with the candidate aggregation, which will simultaneously guides the generation of instance-aware features. Comparing with the common NMS in GICN [17], the proposed LN-NMS can better localize instances with different sizes, while common NMS omits some small instances (garbage bins in row 2 and chairs in row 3).

### 4.4 Ablation Study

Here, we first compare different candidate mining and aggregation strategies. Then, we analyze how to better represent instances as kernels.

**Candidate Mining.** Here we verify the effectiveness of the proposed LN-NMS algorithm for candidate mining. As in Table 4, three different ways are compared.

**Random** (row 1) denotes points above a threshold is randomly selected (at most 200) as instance candidates. **NMS** denotes the searching algorithm with common NMS used in GICN [17]. And the proposed **LN-NMS** algorithm shown in Sec. 3.3. Results in row 1 show that: 1) poor instance localization can significantly degrade instance segmentation performance; 2) even selecting points randomly in the heatmaps as candidates can yield competitive results, which demonstrates the robustness of the subsequent candidate aggregation module. Comparing results in row 2 with row 3, LN-NMS improves  $AP@50$  by 2.1%.

**Instance Candidates Aggregations.** Comparison of different candidate aggregation strategies are shown in Table 5. “W/o” means no candidate aggregation are preformed. “In training” means only optimizing aggregation loss during training, while the raw candidates will not be aggregated for inference. “All phases” denotes our full approach. When adding the aggregation loss only for

**Table 6.** Comparisons of different candidates aggregation approaches on ScanNetV2 validation set.

Strategy	$AP$	$AP@50$	$AP@25$
Candidates elimination	50.4	65.2	75.8
Candidates aggregation	50.8	66.7	76.9
Full instances	51.5	67.0	77.0

supervision (row 2),  $AP@50$  increases by 1.9%; Motivated by the aggregation loss, there is an instance-clustering trend for the identical instances in the feature space, which is similar to contrast learning [7]. By aggregating the duplicated candidates, the full approach promotes the  $AP$ ,  $AP@50$  and  $AP@25$  by 3.1%, 4.1% and 2.0% comparing with baseline, which demonstrates the effectiveness of candidate aggregation.

**Generating Instance Kernels.** With the raw instance candidates and merging map  $M_{ins}$  marking which candidates shall be aggregated, naturally, there are 2 different ways to represent each instance: 1) only using the feature from the candidate with the highest centroids scores; 2) aggregating features from all the merged candidates. As in Table 6, the latter way (default) obtains better results, which demonstrates that aggregating features from different candidates benefits the representation. In addition, we test an ideal representation in row 3 where features are collected from all the points within each instance. It can be observed that the performances of our instance representation approach are comparable with this ideal representation. This indicates that representative and discriminative instance contexts are obtained in our kernel generation process.

## 5 Conclusion

We introduce a 3D instance representation, termed *instance kernels*, which encodes the positional, semantic, and shape information of instances into a 1D vector. We find that the difficulty in representing 3D instances lies in precisely localizing instances and collecting discriminative features. Accordingly, we design a novel instance encoding paradigm that first mines centroids candidates for localization. Then, an aggregation process simultaneously eliminates duplicated candidates and gathers features around each instance for representation. We incorporate the instance kernel into a Dynamic Kernel Network (DKNet), which outperforms previous state-of-the-art approaches on public benchmarks.

## Acknowledgment

This work was supported in part by the National Key R&D Program of China (No.2018YFB1305504) and the DigiX Joint Innovation Center of Huawei-HUST.

## References

1. Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., Savarese, S.: 3d semantic parsing of large-scale indoor spaces. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 1534–1543 (2016). <https://doi.org/10.1109/CVPR.2016.170>
2. Chen, S., Fang, J., Zhang, Q., Liu, W., Wang, X.: Hierarchical aggregation for 3d instance segmentation. In: Proceedings of IEEE International Conference on Computer Vision (ICCV). pp. 15447–15456 (2021). <https://doi.org/10.1109/ICCV48922.2021.01518>
3. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: ScanNet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 2432–2443 (2017). <https://doi.org/10.1109/CVPR.2017.261>
4. Engelmann, F., Bokeloh, M., Fathi, A., Leibe, B., Nießner, M.: 3d-mpa: Multi-proposal aggregation for 3d semantic instance segmentation. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 9028–9037 (2020). <https://doi.org/10.1109/CVPR42600.2020.00905>
5. Graham, B., Engelcke, M., Maaten, L.v.d.: 3d semantic segmentation with sub-manifold sparse convolutional networks. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR) (2018)
6. Han, L., Zheng, T., Xu, L., Fang, L.: Occuseg: Occupancy-aware 3d instance segmentation. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 2937–2946 (2020). <https://doi.org/10.1109/CVPR42600.2020.00301>
7. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 9726–9735 (2020). <https://doi.org/10.1109/CVPR42600.2020.00975>
8. He, K., Gkioxari, G., Dollár, P., Girshick, R.: Mask r-cnn. In: Proceedings of IEEE International Conference on Computer Vision (ICCV). pp. 2980–2988 (2017). <https://doi.org/10.1109/ICCV.2017.322>
9. He, T., Shen, C., van den Hengel, A.: Dyc3d: Robust instance segmentation of 3d point clouds through dynamic convolution. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 354–363 (2021)
10. Hou, J., Dai, A., Nießner, M.: 3d-sis: 3d semantic instance segmentation of RGB-D scans. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 4421–4430 (2019). <https://doi.org/10.1109/CVPR.2019.00455>
11. Jia, X., De Brabandere, B., Tuytelaars, T., Gool, L.V.: Dynamic filter networks. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS). pp. 667–675 (2016)
12. Jiang, L., Zhao, H., Shi, S., Liu, S., Fu, C.W., Jia, J.: Pointgroup: Dual-set point grouping for 3d instance segmentation. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 4866–4875 (2020). <https://doi.org/10.1109/CVPR42600.2020.00492>
13. Kuhn, H.W.: The hungarian method for the assignment problem. Naval research logistics quarterly (1955)
14. Lahoud, J., Ghanem, B., Oswald, M.R., Pollefeys, M.: 3d instance segmentation via multi-task metric learning. In: Proceedings of IEEE International Conference on Computer Vision (ICCV). pp. 9255–9265 (2019). <https://doi.org/10.1109/ICCV.2019.00935>



15. Landrieu, L., Simonovsky, M.: Large-scale point cloud semantic segmentation with superpoint graphs. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 4558–4567 (2018). <https://doi.org/10.1109/CVPR.2018.00479>
16. Liang, Z., Li, Z., Xu, S., Tan, M., Jia, K.: Instance segmentation in 3d scenes using semantic superpoint tree networks. In: Proceedings of IEEE International Conference on Computer Vision (ICCV). pp. 2763–2772 (2021). <https://doi.org/10.1109/ICCV48922.2021.00278>
17. Liu, S., Yu, S., Wu, S., Chen, H., Liu, T.: Learning gaussian instance segmentation in point clouds. arXiv Computer Research Repository (2020)
18. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv Computer Research Repository (2017)
19. Lu, H., Dai, Y., Shen, C., Xu, S.: Index networks. IEEE Transactions on Pattern Analysis and Machine Intelligence **44**(1), 242–255 (2022)
20. Milletari, F., Navab, N., Ahmadi, S.A.: V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: International Conference on 3D Vision (3DV). pp. 565–571 (2016). <https://doi.org/10.1109/3DV.2016.79>
21. Otsu, N.: A threshold selection method from gray-level histograms. IEEE Transactions on Systems, Man, and Cybernetics (1979)
22. Qi, C.R., Litany, O., He, K., Guibas, L.: Deep hough voting for 3d object detection in point clouds. In: Proceedings of IEEE International Conference on Computer Vision (ICCV). pp. 9276–9285 (2019). <https://doi.org/10.1109/ICCV.2019.00937>
23. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Proceedings of International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI). pp. 234–241. Springer (2015). [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
24. Tian, Z., Shen, C., Chen, H.: Conditional convolutions for instance segmentation. In: Proceedings of European Conference on Computer Vision (ECCV). pp. 282–298 (2020). [https://doi.org/10.1007/978-3-030-58452-8\\_17](https://doi.org/10.1007/978-3-030-58452-8_17)
25. Vu, T., Kim, K., Luu, T.M., Nguyen, X.T., Yoo, C.D.: Softgroup for 3d instance segmentation on point clouds. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR) (2022)
26. Wang, X., Zhang, R., Kong, T., Li, L., Shen, C.: Solov2: Dynamic and fast instance segmentation. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS) (2020)
27. Xu, G., Cao, H., Dong, Y., Yue, C., Zou, Y.: Stochastic gradient descent with step cosine warm restarts for pathological lymph node image classification via pet/ct images. In: International Conference on Signal and Image Processing (ICSIP) (2020)
28. 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: Proceedings of Advances in Neural Information Processing Systems (NeurIPS). pp. 6737–6746 (2019)
29. Yi, L., Zhao, W., Wang, H., Sung, M., Guibas, L.J.: GSPN: generative shape proposal network for 3d instance segmentation in point cloud. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 3947–3956 (2019). <https://doi.org/10.1109/CVPR.2019.00407>
30. Zhang, B., Wonka, P.: Point cloud instance segmentation using probabilistic embeddings. In: Proceedings of IEEE Conference on Computer Vision Pattern Recognition (CVPR). pp. 8883–8892 (2021)

31. Zhang, W., Pang, J., Chen, K., Loy, C.C.: K-net: Towards unified image segmentation. In: Proceedings of Advances in Neural Information Processing Systems (NeurIPS). pp. 10326–10338 (2021)