

DBGroup: Dual-Branch Point Grouping for Weakly Supervised 3D Semantic Instance Segmentation

Xuexun Liu¹, Xiaoxu Xu¹, Qiudan Zhang^{1*}, Lin Ma², Xu Wang¹

¹College of Computer Science and Software Engineering, Shenzhen University, Shenzhen, 518060, China

²Meituan Inc., China

2019043026@email.szu.edu.cn

Abstract

Weakly supervised 3D instance segmentation is essential for 3D scene understanding, especially as the growing scale of data and high annotation costs associated with fully supervised approaches. Existing methods primarily rely on two forms of weak supervision: one-thing-one-click annotations and bounding box annotations, both of which aim to reduce labeling efforts. However, these approaches still encounter limitations, including labor-intensive annotation processes, high complexity, and reliance on expert annotators. To address these challenges, we propose **DBGroup**, a two-stage weakly supervised 3D instance segmentation framework that leverages scene-level annotations as a more efficient and scalable alternative. In the first stage, we introduce a Dual-Branch Point Grouping module to generate pseudo labels guided by semantic and mask cues extracted from multi-view images. To further improve label quality, we develop two refinement strategies: Granularity-Aware Instance Merging and Semantic Selection and Propagation. The second stage involves multi-round self-training on an end-to-end instance segmentation network using the refined pseudo-labels. Additionally, we introduce an Instance Mask Filter strategy to address inconsistencies within the pseudo labels. Extensive experiments demonstrate that DBGroup achieves competitive performance compared to sparse-point-level supervised 3D instance segmentation methods, while surpassing state-of-the-art scene-level supervised 3D semantic segmentation approaches.

Code — <https://github.com/liuxuexun/DBGroup>

Introduction

3D instance segmentation is a fundamental task in 3D scene understanding that aims to predict both masks and semantic categories for individual objects within point cloud scenes. This field has garnered significant attention in recent years due to its extensive applications in the emerging domains such as embodied AI, robotics, and AR/VR technologies (Liu et al. 2024; Xu et al. 2025c; Li et al. 2022a). Recently, learning-based 3D instance segmentation methods (Schult et al. 2023; Zhao et al. 2023) have achieved remarkable performance. However, these methods heavily rely on time-

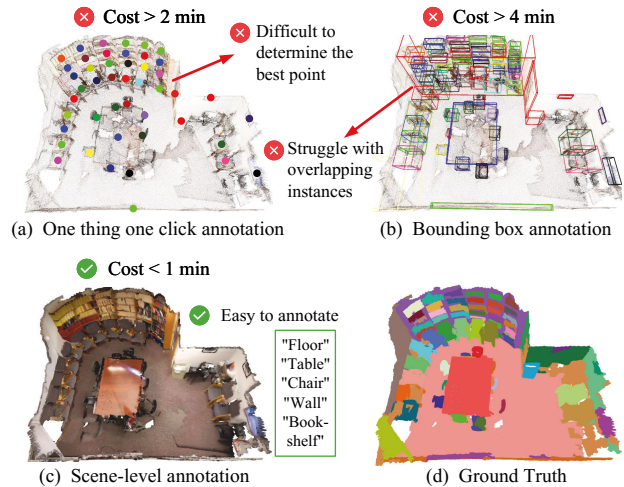


Figure 1: Comparison between conventional weak annotation formats and our proposed scene-level annotation in 3D instance segmentation task.

consuming and labor-intensive point-wise annotations for both semantic and instance labels.

To alleviate this limitation, numerous methods (Tao et al. 2022; Tang, Hui, and Xie 2022a; Chibane et al. 2022; Lu, Deng, and Zhang 2024) have explored weakly supervised 3D instance segmentation. As shown in Fig.1, these approaches primarily utilize two types of weak annotations: “one thing one click” (OTOC) annotation and bounding box (BBox) annotation. OTOC indicates each instance will be annotated at least one point while BBox need to annotate each instance with bounding box. While these annotation methods substantially reduce annotation costs, there still exist some issues. First, they still require distinguishing individual instances, which fails to significantly reduce the time-consuming annotation process. Second, they often demand specialized annotator training. For OTOC, annotators must precisely determine click point positions for complex structures; for BBox, untrained annotators struggle with overlapping instances with precise boundaries.

Therefore, inspired by recent weakly supervised 3D semantic segmentation approaches (Wei et al. 2020; Xu et al.

*Corresponding author.

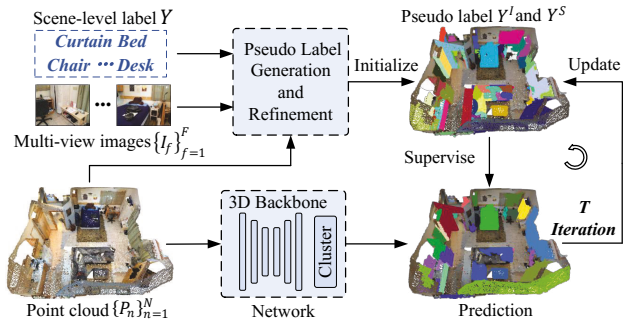


Figure 2: Overview of our weakly supervised 3D segmentation pipeline.

2024), we propose to employ scene-level annotation for weakly supervised 3D instance segmentation. As illustrated in Fig.1(c), scene-level annotation merely requires labeling the categories of objects present within a scene. Compared to OTOC and BBox annotation, scene-level annotation offers several significant advantages. On the one hand, it substantially reduces annotation time, requiring less than 1 minute per scene on average, whereas OTOC and BBox annotations typically require more than 2 minutes and 4 minutes (Yang et al. 2022). On the other hand, it eliminates the need for professional training of annotators, as they only need to distinguish the categories of objects in the scene.

However, compared with OTOC and BBox annotations, scene-level annotation lacks instance-level information, making it difficult to guide the model in segmenting individual instances. To address this issue, as illustrated in Fig.2, we propose a two-stage paradigm comprising Pseudo Label Generation and Refinement, followed by a multi-round self-training 3D instance segmentation network. In the first stage (see Fig.3), we introduce a Dual-Branch Point Grouping module consisting of a Semantic Guidance Branch (SGB) and Mask Guidance Branch (MGB). The SGB leverages a pre-trained vision-language model for pseudo semantic labels and applies radius-based BFS clustering to generate coarse-grained pseudo instances labels. Concurrently, the MGB employs superpoint prompts and SAM (Kirillov et al. 2023) to produce 2D masks that guide 3D point grouping into fine-grained clusters. We further enhance pseudo label quality through Granularity-Aware Instance Merging (GAIM) and Semantic Selection and Propagation (SSP): GAIM performs instance splitting or merging based on mask intersection ratios, while SSP filters low-confidence predictions and propagates high-confidence ones to broader regions. In the second stage (Fig.4), an Instance Mask Filter strategy mitigates inconsistencies between instance and semantic pseudo labels, enabling better focus on individual instances.

To summarize, our contributions are as follows:

- We propose a new weakly supervised 3D semantic instance segmentation method DBGroup, which use only scene-level labels. To the best of our knowledge, DBGroup is the first work investigating 3D instance segmentation with scene-level labels.

- Our DBGroup leverages a Dual-Branch Point Grouping module to produce both pseudo instance and semantic label and utilizes the Pseudo Label Refinement to generate high-quality pseudo labels. Moreover, an Instance Mask Filter strategy is introduced to improve the consistency of instance predictions.
- Extensive experiments demonstrate that our approach achieves competitive performance against sparse-point-level supervised methods while using more label-efficient scene-level annotations, and surpasses current state-of-the-art scene-level supervised approaches, validating our framework’s effectiveness and superiority.

Related Work

Weakly Supervised 3D Segmentation. In response to the challenge of manual annotation, weakly supervised 3D segmentation has garnered significant research interest in recent years. Researchers typically leverage multiple weakly supervised signals. Box-level annotation offers the most comprehensive representation of instance and geometric details, the primary challenge of these methods (Chibane et al. 2022; Du et al. 2023; Ngo, Hua, and Nguyen 2023; Lu, Deng, and Zhang 2024) is accurately assigning points within overlapping regions of 3D bounding boxes. In sparse-points-level annotation, some methods (Li et al. 2022b; Wu et al. 2023; Xu et al. 2023) adopt the mean teacher paradigm, while others (Tao et al. 2022; Tang, Hui, and Xie 2022b; Liu, Qi, and Fu 2021; Dong et al. 2023) employ graph-based approaches to propagate sparse supervised signals. Methods (Wei et al. 2020; Ren et al. 2021; Yang et al. 2022) based on scene-level annotation predominantly employ 3D Class Activation Mapping (CAM) (Zhou et al. 2016) for semantic object localization. Subsequent approaches (Kweon and Yoon 2022; Yang et al. 2023a) enhance this framework by incorporating 2D features for cross-modal alignment and fusion. Recent state-of-the-art advance (Xu et al. 2024, 2025b) leverage the robust generalization capabilities of 2D visual-language models to align textual information of scene label with 3D visual features. Although scene-level annotation is cost-effective and easily implementable, current instance segmentation approaches are unable to effectively utilize such coarse-grained labeling schemes.

2D Pre-trained Model in 3D Segmentation. With the rapid advancement of 2D pre-trained models, researchers have increasingly explored their application to 3D segmentation tasks. Some approaches (Yang et al. 2023b; Yin et al. 2024; Xu et al. 2025a) leverage SAM (Kirillov et al. 2023) for class-agnostic segmentation of 3D point clouds, while others (Peng et al. 2023; Li et al. 2025; Takmaz et al. 2023) utilize CLIP (Radford et al. 2021) to align point cloud features with text embeddings for open-vocabulary segmentation. Recent developments have extended these 2D pre-trained frameworks to weakly supervised 3D segmentation. 3DSSVLG (Xu et al. 2024) pioneered the use of CLIP models to guide weakly supervised semantic segmentation through textual semantics, while REAL (Kweon, Kim, and Yoon 2024) employs SAM as an active learning annotator to iteratively refine segmentation masks. However, most of

these methods require 2D images as input during training or inference, which inadvertently increases computational demands and limits potential applications. In contrast, our two-stage method only requires processing the dataset containing 2D images once.

Methodology

Problem Definition. In our weakly supervised setting, as shown in Fig.3, the inputs comprise three components: **1)** a point cloud scene $\{\mathbf{P}_n\}_{n=1}^N \in \mathbb{R}^{N \times 6}$ containing N points with its spatial coordinates and color values attributes (x, y, z, r, g, b) , **2)** its corresponding RGB multi-view images $\{\mathbf{I}_f\}_{f=1}^F$ containing F frames with a resolution of $H \times W$, **3)** scene-level label $Y \in \{Floor, Curtain, Bed, \dots, Desk\}^K$, K denotes the number of categories in the scene. Our objective is to produce high quality pseudo instance and semantic label $Y^I \in \mathbb{R}^N$ and $Y^S \in \mathbb{R}^N$ to facilitate the T -iteration self-training.

Semantic Guidance Branch

Inspired by previous works (Peng et al. 2023; Xu et al. 2024), we first leverage pre-trained vision-language models (Radford et al. 2021; Ghiasi et al. 2022) to obtain multi-modal embeddings. By performing matrix multiplication between these embeddings, we then derive pseudo semantic labels for the point cloud. Finally, we propose a BFS Grouping strategy to aggregate points that belong to the same instance and generate the predicted coarse-grained instance masks.

Feature Extraction. Following (Peng et al. 2023; Xu et al. 2024), we utilize the pre-trained image encoder and the text encoder from OpenSeg (Ghiasi et al. 2022) to obtain the embeddings $\mathbf{F}_{2D} \in \mathbb{R}^{F \times H \times W \times C}$ and $\mathbf{F}_{1D} \in \mathbb{R}^{K \times C}$, respectively, where C denotes the embedding dimension.

Projection and Accumulation. For each point \mathbf{P}_n in the point cloud $\{\mathbf{P}_n\}_{n=1}^N$, we first project it onto all multi-view images using geometric camera calibration matrices. After applying visibility filtering, we obtain the corresponding 2D positions of \mathbf{P}_n on the relevant images. These positions are then used to extract the corresponding embeddings from \mathbf{F}_{2D} . Since a single point \mathbf{P}_n may be visible in multiple images, we aggregate its multiple associated embeddings by computing their mean. This embedding extraction and aggregation process is repeated for all points in the point cloud, resulting in a aggregated 3D embeddings $\mathbf{F}_{3D} \in \mathbb{R}^{N \times C}$.

BFS Grouping. Given the scene-level label embeddings $\mathbf{F}_{1D} \in \mathbb{R}^{K \times C}$ and the aggregated 3D embeddings $\mathbf{F}_{3D} \in \mathbb{R}^{N \times C}$, we first compute classification scores $\mathbf{S} \in \mathbb{R}^{N \times K}$ through matrix multiplication. Based on \mathbf{S} , we adopt a widely-used Breadth-First Search (BFS) algorithm (Jiang et al. 2020) to group points belonging to foreground categories. Specifically, we select an initial foreground anchor point \mathbf{P}_n and search for neighboring points within a spherical region with radius r centered at \mathbf{P}_n . Neighboring points that sharing the same semantic prediction as the anchor point are grouped together into the same cluster. Subsequently, the newly grouped points serve as anchors to iteratively repeat the above procedure until no additional points satisfying the criteria can be included. Upon completion of this iterative

process, we obtain a complete instance cluster, which is then removed from the original point cloud. We subsequently select another anchor point to perform clustering for the next instance, repeating this procedure until all foreground points have been clustered into their respective instances.

After preceding steps, our SGB produces final classification scores \mathbf{S} and coarse-grained instance masks $\{\mathbf{M}_q\}_{q=1}^Q$, where each mask is a binary mask $\mathbf{M}_q \in \{0, 1\}^N$.

Mask Guidance Branch

While the coarse-grained instance masks $\{\mathbf{M}_q\}_{q=1}^Q$ from SGB capture complete object instances, they often lack sufficient precision. This is primarily due to the radius-based clustering performed by the BFS algorithm, which may mistakenly merge two distinct instances of the same category into a single cluster when their spatial distance is smaller than the predefined radius. To overcome this issue, we introduce a Mask Guidance Branch that leverages superpoint prompts and SAM (Kirillov et al. 2023) to produce fine-grained instance masks $\{\mathbf{O}_w\}_{w=1}^W$ for enhancing the coarse-grained masks.

Prompt Initialization. Given a point cloud scene $\{\mathbf{P}_n\}_{n=1}^N$, we first oversegment the points into superpoints $\{\mathbf{Q}_m\}_{m=1}^M$ using a normal-based graph cut algorithm (Felzenszwalb and Huttenlocher 2004), where M denotes the number of superpoints in $\{\mathbf{P}_n\}_{n=1}^N$. For each superpoint \mathbf{Q}_m , we compute its centroid $\mathbf{c}_m \in \mathbb{R}^3$ as the point closest to the average coordinates of all points within \mathbf{Q}_m . We project the resulting centroid set $\{\mathbf{c}_m\}_{m=1}^M$ onto the multi-view images, serving as prompts for SAM to perform preliminary mask segmentation. Compared to point-level prompt initialization (Xu et al. 2025a), which treats every point in the point cloud as an individual prompt, this superpoint-based strategy significantly reduces computational overhead. Moreover, using a small number of geometrically meaningful prompts enables SAM to generate more accurate and fine-grained instance masks. Finally, we can obtain 2D segmentation masks for all frames $\{\mathbf{MS}_f\}_{f=1}^F$.

Projection and Grouping. Given the superpoint centroid set $\{\mathbf{c}_m\}_{m=1}^M$ with previous $\{\mathbf{MS}_f\}_{f=1}^F$, we conduct a multi-view projection to achieve comprehensive 3D scene segmentation. Each 3D point is projected onto all masked frames $\{\mathbf{MS}_f\}_{f=1}^F$, with labels assigned through a robust voting mechanism. Specifically, if a point \mathbf{P}_n falls within the mask of prompt \mathbf{c}_m in a given frame \mathbf{MS}_f , \mathbf{P}_n is assigned the instance ID m at frame f . The final instance ID of \mathbf{P}_n is determined as the most frequently assigned prompt ID across all F frames. Through this process, we obtain fine-grained instance masks $\{\mathbf{O}_w\}_{w=1}^W$, where each mask is represented as a binary vector $\mathbf{O}_w \in \{0, 1\}^N$.

Granularity Aware Instance Merging

Granularity Aware Assignment. Through the SGB and MGB, we obtain the coarse-grained instance masks $\{\mathbf{M}_q\}_{q=1}^Q$ and fine-grained instance masks $\{\mathbf{O}_w\}_{w=1}^W$. The coarse-grained masks are better suited for capturing complete instance objects, whereas the fine-grained masks are more effective at accurately delineating object boundaries.

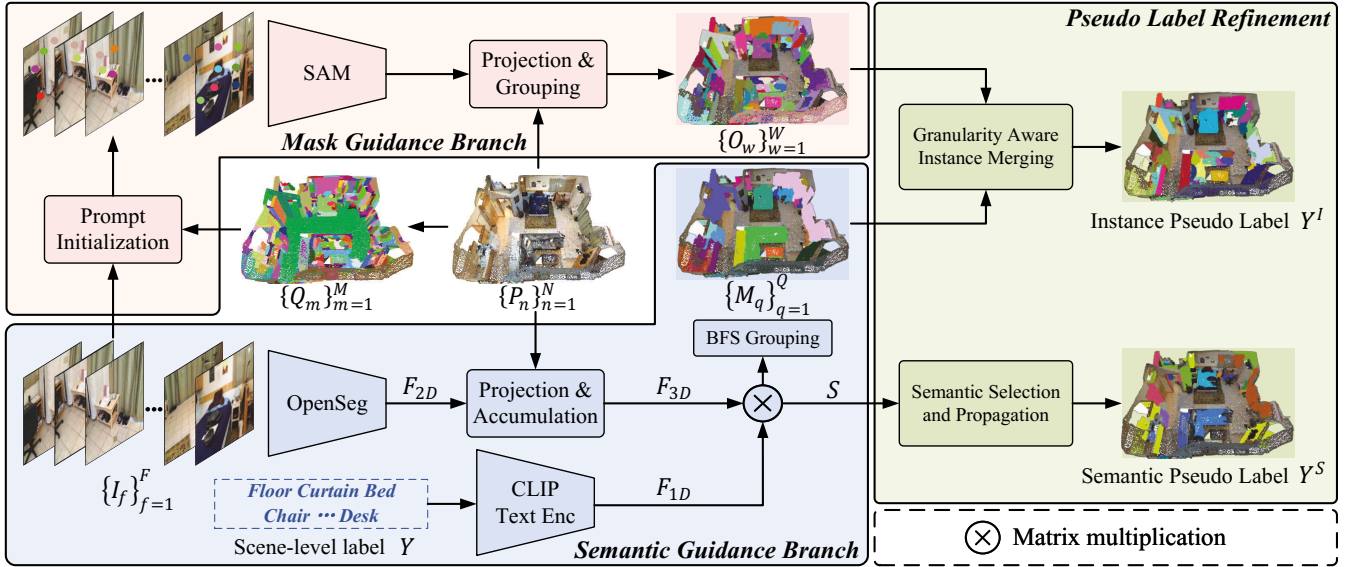


Figure 3: The workflow of our Pseudo Label Generation and Refinement comprises a dual-branch point grouping architecture and two refinement strategies. The SGB extracts multi-view image features via a pre-trained 2D model, projects them into 3D point clouds, computes text encoder similarities, and applies BFS clustering for coarse-grained masks. The MGB projects superpoints onto multi-view images as SAM prompts, creating masks that group 3D points into fine-grained masks. GAIM merges or splits masks for instance pseudo labels, while SSP filters semantic scores for semantic pseudo labels.

To integrate these complementary representations, we propose a Granularity Aware Assignment strategy. We first construct an overlap matrix $\mathbf{A} \in \mathbb{Z}^{Q \times W}$ to quantify the intersections between the two sets of masks. For each coarse-grained mask \mathbf{M}_q , we identify the fine-grained mask \mathbf{O}_w with the highest overlap ratio ρ . If ρ exceeds a predefined threshold θ , the coarse-grained mask is retained. Otherwise, we consider the coarse-grained mask \mathbf{M}_q to be under-segmented and likely to contain multiple distinct instances. To address this, we compute the intersection between the coarse-grained mask and all fine-grained masks \mathbf{O}_w , and retain only the overlapping regions with fine-grained masks to refine and decompose the under-segmented coarse-grained region. After that, we can obtain the ensemble instance masks Γ . The pseudo code is in our supplementary material.

Small Instances Merging. Although Granularity Aware Assignment can effectively decompose under segmented masks, fine-grained masks also have the problem of over segmentation. To resolve this challenge, we leverage spatial adjacency relationships between instances to facilitate appropriate merging. Our method first evaluates inter-instance adjacency relationship by employing the K-Nearest Neighbors (KNN) algorithm for all points in the scene. The adjacency relationship between instances A and B is established based on the proximity of their points. Specifically, the nearest neighboring instance of instance A is defined as the instance containing the greatest number of points that are nearest neighbors to points within instance A. Subsequently, we calculate the point number within each instance and merges instances containing fewer points than a predefined threshold γ with their nearest large neighboring instance. This pro-

cess ultimately yields the instance pseudo label $\mathbf{Y}^I \in \mathbb{R}^N$. Given that instance centroids are subject to spatial offset, with smaller instances particularly vulnerable to such displacement. We avoid using centroids for determining inter-instance adjacency relationship.

Semantic Selection and Propagation

Semantic Selection. As is well known, low-confidence pseudo labels are more likely to be inaccurate compared to high-confidence ones, which also holds true for the classification score matrix \mathbf{S} generated by the SGB. To address this issue, we propose a semantic selection algorithm. Specifically, for each category present in a scene, we retain only the top $\alpha\%$ of points with the highest classification confidence and discard predictions for the remaining points in that category. Unlike conventional global top- $\alpha\%$ filtering based solely on classification scores, our class-wise selection strategy maintains semantic balance between majority and minority classes, effectively alleviating performance degradation caused by the under-representation of minority classes.

Superpoint Propagation. After the selection process, the semantic pseudo labels become sparse with limited coverage rate and lack spatial continuity. Therefore, we propose a Superpoint Propagation algorithm that leverages geometric priors for effective label propagation. Specifically, we utilize superpoints $\{\mathbf{Q}_m\}_{m=1}^M$ to partition the point set $\{\mathbf{P}_n\}_{n=1}^N$ into M distinct groups. Within each group, we determine the most frequently predicted category and propagate this label to all points within the corresponding superpoint yielding the semantic pseudo label $\mathbf{Y}^S \in \mathbb{R}^N$.

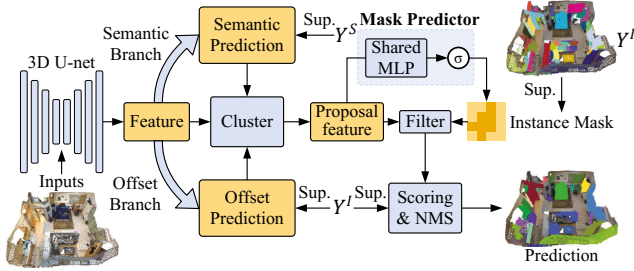


Figure 4: Framework of our 3D instance segmentation network. Following a grouping-based paradigm, we use a 3D U-Net for feature extraction with semantic and offset branches to predict semantic labels and offsets. After clustering, an Instance Mask Filter removes irrelevant points from proposals, with final predictions obtained via a scoring network and non-maximum suppression (NMS).

Self-Training

To gradually improve the quality of pseudo label and enhance the performance on the validation set, we adopt a T -iteration self-training strategy, as illustrated in Fig.2. This iterative approach allows the model to generate increasingly accurate and robust predictions by using its progressively refined outputs as supervisory signals.

The network architecture is shown in Fig.4. We adopt a grouping-based 3D instance segmentation framework following (Jiang et al. 2020; Vu et al. 2022; Zhao et al. 2023). More details are shown in our supplementary material.

However, due to the independence of the SGB and MGB, the generated semantic pseudo labels Y^S and instance pseudo labels Y^I may be inconsistent, introducing noise into the clustering process. To address this issue, we propose an *Instance Mask Filter* strategy. Specifically, for each proposal feature $PF_i \in \mathbb{R}^{N_i \times C}$, where N_i denotes the number of points in the i -th proposal, we apply a shared multi-layer perceptron (MLP) followed by binary thresholding at 0.5 to predict a proposal-specific mask $PM_i \in \{0, 1\}^{N_i}$. This mask is used to filter the proposal embeddings PF_i , refining the instance representation. Following (Zhao et al. 2023; Wu et al. 2022), we supervise PM_i using a combination of binary cross-entropy loss L_{bce} and Dice loss (Milletari, Navab, and Ahmadi 2016) L_{dice} .

Experiment

Datasets. We evaluate our method on two widely-used indoor point cloud datasets: ScanNetV2 (Dai et al. 2017) and S3DIS (Armeni et al. 2017). ScanNetV2 comprises 1,513 training samples across 20 semantic categories, with instance segmentation required for 18 categories. Following the default setting, we use 1,201 scenes for training and 312 scenes for validation. S3DIS consists of 271 scenes distributed across 6 areas, encompassing 13 semantic categories that all require instance segmentation. We follow the common evaluation setting, using Area 5 as the validation set and the remaining areas for training. Both datasets provide RGB-D sequences along with intrinsic and extrinsic

Method	Annotation	AP	AP ₅₀	AP ₂₅
GSPN	mask	19.3	37.8	53.4
3D-SIS	mask	-	18.7	35.7
MTML	mask	20.3	40.2	55.4
PointGroup	mask	34.8	56.9	71.3
SoftGroup	mask	46.0	67.6	78.9
Mask3D	mask	55.2	73.7	-
<hr/>				
SPiB	box	-	38.6	61.4
Box2Mask	box	39.1	59.7	71.8
BSNet	box	53.3	72.7	83.4
<hr/>				
CSC-50	0.034% point	22.9	41.4	62.0
TWIST	5% point	27.0	44.1	56.2
SegGroup	0.02% point	23.4	43.4	62.9
3D-WSIS	0.02% point	28.1	47.2	67.5
<hr/>				
DBGroup (Ours)	scene	28.6	46.2	59.6

Table 1: Comparison of different methods for 3D instance segmentation on ScanNetV2 dataset. “mask” represents fully annotation. “box” indicates bounding box annotation. “ $x\%$ point” denotes the percentage of annotated points relative to all points in the entire dataset. “scene” denotes scene-level annotation.

camera parameters, enabling 3D-to-2D projection.

Evaluation Metrics. We evaluate our approach on both semantic segmentation and instance segmentation tasks. For semantic segmentation, we employ mean Intersection over Union (mIoU) as the evaluation metric. For instance segmentation, we utilize mean Average Precision (mAP) at various Intersection over Union (IoU) thresholds: 0.25 (AP₂₅), 0.5 (AP₅₀), and the average over the range [0.5:0.95:0.05] (AP). Additionally, for the S3DIS dataset, we report mean Recall (mRec) and mean Precision (mPre) at an IoU threshold of 0.5.

Implementations Details. For the hyper-parameters of our pseudo label refinement module, we set the BFS Grouping radius r to 0.04, overlap threshold θ to 0.4, the point threshold γ to 200 and the top- $\alpha\%$ threshold to 30% respectively. In our training process, we set the number of self-training iterations T to 3, employed the Adam optimizer (Kingma 2014) with a batch size of 6 and initialized the learning rate at 0.001. We implement a cosine annealing schedule (Loshchilov and Hutter 2016), beginning with 50 epochs and applying per-step decay, for a total of 256 training epochs. Our architecture utilized MinkowskiNet34C (Choy, Gwak, and Savarese 2019) as the 3D network backbone, and all experimental evaluations were conducted using three NVIDIA RTX 4090 GPUs.

Comparison with SOTAs

Evaluation on 3D Instance Segmentation. As shown in Tab.1 and 2, we evaluate our method against a range of approaches utilizing different levels of supervision. The results demonstrate that our DBGroup achieves competitive performance on both benchmarks, outperforming the majority of point-annotation-based methods. Specifically, when

Method	Annotation	AP	AP ₅₀	mPrec	mRec
PointGroup	mask	-	57.8	61.9	62.1
Mask3D	mask	56.6	68.4	68.7	66.3
Box2Mask	box	43.6	54.6	66.7	65.5
BSNet	box	53.0	64.3	-	-
SegGroup	0.02% point	21.0	29.8	47.2	34.9
3D-WSIS	0.02% point	23.3	33.0	50.8	38.9
DBGGroup (Ours)	scene	27.1	40.5	38.5	44.8

Table 2: Comparison of different methods for 3D instance segmentation on S3DIS dataset.

Method	Annotation	ScanNetV2 mIoU	S3DIS mIoU
PointNet	mask	-	41.1
MinkowskiNet	mask	72.2	65.4
PTv3	mask	78.6	74.3
HybridCR	1% point	56.9	65.3
OTOC++	0.02% point	70.5	56.6
MPRM	scene	24.4	10.3
MIL-Trans	scene	26.2	12.9
WYPR	scene	29.6	22.3
MIT	scene	35.8	27.7
3DSS-VLG	scene	49.7	45.3
DBGGroup (Ours)	scene	56.9	48.2

Table 3: Comparison of different methods for 3D semantic segmentation on ScanNetV2 and S3DIS dataset.

compared to 3D-WSIS (Tang, Hui, and Xie 2022a), a state-of-the-art method based on point annotations, DBGGroup achieves comparable results on the ScanNetV2 dataset, while showing substantial gains on the S3DIS dataset surpassing 3D-WSIS by 3.8% and 7.5% in terms of *AP* and *AP*₅₀, respectively. It is worth noting that scene-level annotations, required by DBGGroup, are significantly more cost-efficient than point-level annotations, underscoring the effectiveness of our approach under simpler supervision.

Furthermore, DBGGroup demonstrates lower performance in the *mPrec* metric compared to point-annotation-based methods, while significantly outperforming them in *mRec*. This trade-off arises from the nature of scene-level supervision: lacking fine-grained spatial cues, pseudo labels generated from scene-level annotations tend to include irrelevant regions, thereby lowering precision. In contrast, point-level annotations offer precise localization but limited coverage, causing the model to overlook unlabeled object parts and thus suffer in recall. Our results highlight DBGGroup’s ability to strike a better balance between annotation cost and segmentation performance.

Evaluation on 3D Semantic Segmentation. We further evaluate DBGGroup on the 3D semantic segmentation task and compare it with existing methods under different annotation settings. As shown in Tab.3, when trained with

Setting	AP	AP ₅₀	AP ₂₅
(a) Only SGB	16.3	30.7	53.3
(b) Only MGB	15.5	30.2	56.2
(c) MGB+SGB (GAIM)	17.4	32.7	58.7
(d) MGB+SGB (BM)	16.3	30.8	54.0

Table 4: Quantitative ablation results of SGB, MGB and GAIM.

only scene-level annotations, DBGGroup surpasses the previous state-of-the-art by 7.2% on ScanNetV2 and 2.9% on the S3DIS dataset. These results clearly demonstrate the effectiveness of our method for 3D semantic segmentation.

Ablation Study

In this section, we conduct a series of ablation studies to validate the effectiveness of each component in our proposed framework. All experiments are performed on the ScanNetV2 dataset, using the same experimental settings as described in previous sections, except for the specific module under evaluation. To ensure training efficiency and facilitate fair comparisons, all models are trained with a single iteration ($T = 1$) unless otherwise stated.

Effect of SGB, MGB, and GAIM. In Tab.4, we present a comparative analysis of how different branches and merging strategies affect the quality of pseudo instance labels. Settings (a) and (b) independently employ the SGB and MGB to generate pseudo instance labels. Setting (c) utilizes both branches to produce instance masks of varying granularities, which are then integrated using our proposed GAIM strategy. Setting (d) implements the Bidirectional Merging (BM) approach from SAM3D (Yang et al. 2023b) for label integration, which is a widely adopted technique for combining diverse mask predictions into coherent results. Among all approaches examined, our GAIM strategy demonstrates superior performance, exceeding BM by 1.1% in AP metric.

Qualitative results corresponding to these settings are presented in Fig.5. As previously discussed, pseudo labels generated by SGB tend to suffer from under-segmentation (e.g., *multiple chairs in rows (II) and (IV) are merged into a single instance*), while those from MGB often lead to over-segmentation (e.g., *individual chairs in rows (I) and (III) are split into multiple instances*). Although the BM strategy attempts to reconcile these two extremes, it fails to effectively balance the trade-off. In contrast, our GAIM strategy produces results that are visually closer to the ground truth and exhibit fewer segmentation artifacts.

Both the quantitative and qualitative comparisons clearly demonstrate that GAIM successfully integrates complementary mask information from two granular levels, achieving a superior balance between under-segmentation and over-segmentation.

Effect of SSP. Tab.5 presents an ablation study evaluating the effectiveness of each component in the Semantic Selection and Propagation (SSP) strategy. In the baseline setting (a), we directly apply the *max* operation operation on the classification score matrix \mathbf{S} to select the top- $\alpha\%$ points




Figure 5: Qualitative ablation results of SGB, MGB and GAIM. From left to right: input point clouds, ground truth, setting (a) to (d) in Tab.4.

Setting	mIoU*	mIoU	AP	AP ₅₀	AP ₂₅
(a) Baseline	60.6	50.9	23.6	39.7	52.5
(b) Sel.	63.1	52.5	25.0	41.3	55.6
(c) Sel. + Prop.	67.2	54.2	26.5	43.3	55.8

Table 5: Ablation results of SSP. * indicates the results of the pseudo labels on the training set, while the remaining represent the prediction results on the validation set. "Sel." denotes "Selection" and "Prop." indicates "Propagation".

with the highest scores. Compared with (a), setting (b) enhances label quality by selecting the top- $\alpha\%$ points separately for each category. Experimental results demonstrate that setting (b) outperforms setting (a) by 1.6% and 1.4% in terms of mIoU and AP metrics on the validation set, respectively. This indicates that category-wise selection can effectively mitigate the performance degradation arising from the under-representation of minority classes. Building upon setting (b), setting (c) further improves label quality by propagating the selected semantic labels within each superpoint. Results show that setting (c) surpasses setting (b) by 1.7% and 1.5% in mIoU and AP, respectively. These gains suggest that leveraging geometric priors via superpoint-based propagation enhances spatial consistency and reduces label fragmentation, thereby promoting more effective model training.

Effect of Instance Mask Filter. To explore the impact of the Instance Mask Filter module, we conduct ablation studies as shown in Tab.6. In setting (a), the proposal features PF_i are directly fed into the scoring network without any filtering. In setting (b), we incorporate the Instance Mask Filter

Setting	AP	AP ₅₀	AP ₂₅
(a) w/o IMF	25.4	42.1	54.8
(b) w IMF	26.5	43.3	55.8

Table 6: Ablation results of Instance Mask Filter (IMF).

module into the pipeline. The results show consistent improvements across all metrics, indicating that Instance Mask Filter effectively suppresses the noise caused by inconsistencies in pseudo labels. This leads to more reliable instance representations and facilitates better model learning.

Conclusion

In this paper, we introduce DBGroup, a Dual-Branch Point Grouping framework designed to project both semantic and mask representations from multi-view images into 3D space. Our framework generates multi-granularity instance masks under scene-level label constraints, and we developed two specialized pseudo-label refinement strategies to effectively integrate instance masks and enhance semantic accuracy. Furthermore, we propose an Instance Mask Filtering mechanism to mitigate training noise resulting from label inconsistencies. Extensive experimental evaluations demonstrate that our approach outperforms existing point-level annotation methods in instance segmentation and surpasses scene-level annotation methods in semantic segmentation. We contend that leveraging scene-level labels for weakly supervised segmentation represents a new and promising research direction with significant potential for future development.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grants 62371310, 62501403 and W2412099, in part by the Guangdong Basic and Applied Basic Research Foundation under Grant 2023A1515011236, in part by the Shenzhen Science and Technology Program (JCYJ20241202124415021 and KJZD20230923114605011), in part by the Stable Support Project of Shenzhen (Project No.20231122122722001) and in part by the Scientific Foundation for Youth Scholars of Shenzhen University.

References

- Armeni, I.; Sax, S.; Zamir, A. R.; and Savarese, S. 2017. Joint 2d-3d-semantic data for indoor scene understanding. *Proc. ICCV*.
- Chibane, J.; Engelmann, F.; Anh Tran, T.; and Pons-Moll, G. 2022. Box2mask: Weakly supervised 3d semantic instance segmentation using bounding boxes. In *Proc. ECCV*, 681–699.
- Choy, C.; Gwak, J.; and Savarese, S. 2019. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *Proc. CVPR*, 3075–3084.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. CVPR*, 5828–5839.
- Dong, S.; Li, R.; Wei, J.; Liu, F.; and Lin, G. 2023. Collaborative Propagation on Multiple Instance Graphs for 3D Instance Segmentation with Single-point Supervision. In *Proc. ICCV*, 16665–16674.
- Du, H.; Yu, X.; Hussain, F.; Armin, M. A.; Petersson, L.; and Li, W. 2023. Weakly-supervised point cloud instance segmentation with geometric priors. In *Proc. WACV*, 4271–4280.
- Felzenszwalb, P. F.; and Huttenlocher, D. P. 2004. Efficient graph-based image segmentation. *Int. J. Comput. Vis.*, 59: 167–181.
- Ghiasi, G.; Gu, X.; Cui, Y.; and Lin, T.-Y. 2022. Scaling open-vocabulary image segmentation with image-level labels. In *Proc. ECCV*, 540–557.
- Jiang, L.; Zhao, H.; Shi, S.; Liu, S.; Fu, C.-W.; and Jia, J. 2020. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proc. CVPR*, 4867–4876.
- Kingma, D. P. 2014. Adam: A method for stochastic optimization.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proc. ICCV*, 4015–4026.
- Kweon, H.; Kim, J.; and Yoon, K.-J. 2024. Weakly supervised point cloud semantic segmentation via artificial oracle. In *Proc. CVPR*, 3721–3731.
- Kweon, H.; and Yoon, K.-J. 2022. Joint learning of 2D-3D weakly supervised semantic segmentation. *Advances in NeurIPS*, 35: 30499–30511.
- Li, J.; Jie, Z.; Wang, X.; Wei, X.; and Ma, L. 2022a. Expansion and shrinkage of localization for weakly-supervised semantic segmentation. *Advances in neural information processing systems*, 35: 16037–16051.
- Li, J.; Saltori, C.; Poiesi, F.; and Sebe, N. 2025. Cross-modal and uncertainty-aware agglomeration for open-vocabulary 3d scene understanding. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 19390–19400.
- Li, M.; Xie, Y.; Shen, Y.; Ke, B.; Qiao, R.; Ren, B.; Lin, S.; and Ma, L. 2022b. Hybridcr: Weakly-supervised 3d point cloud semantic segmentation via hybrid contrastive regularization. In *Proc. CVPR*, 14930–14939.
- Liu, X.; Xu, X.; Li, J.; Zhang, Q.; Wang, X.; Sebe, N.; and Ma, L. 2024. LESS: Label-Efficient and Single-Stage Referring 3D Segmentation. In Globerson, A.; Mackey, L.; Belgrave, D.; Fan, A.; Paquet, U.; Tomczak, J.; and Zhang, C., eds., *Advances in Neural Information Processing Systems*, volume 37, 11164–11185. Curran Associates, Inc.
- Liu, Z.; Qi, X.; and Fu, C.-W. 2021. One thing one click: A self-training approach for weakly supervised 3d semantic segmentation. In *Proc. CVPR*, 1726–1736.
- Loshchilov, I.; and Hutter, F. 2016. Sgdr: Stochastic gradient descent with warm restarts.
- Lu, J.; Deng, J.; and Zhang, T. 2024. BSNet: Box-Supervised Simulation-assisted Mean Teacher for 3D Instance Segmentation. In *Proc. CVPR*, 20374–20384.
- Milletari, F.; Navab, N.; and Ahmadi, S.-A. 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In *Proc. 3DV*, 565–571. Ieee.
- Ngo, T. D.; Hua, B.-S.; and Nguyen, K. 2023. Gapro: Box-supervised 3d point cloud instance segmentation using gaussian processes as pseudo labelers. In *Proc. ICCV*, 17794–17803.
- Peng, S.; Genova, K.; Jiang, C.; Tagliasacchi, A.; Pollefeys, M.; Funkhouser, T.; et al. 2023. Openscene: 3d scene understanding with open vocabularies. In *Proc. CVPR*, 815–824.
- Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; et al. 2021. Learning transferable visual models from natural language supervision. In *Proc. ICML*, 8748–8763. PMLR.
- Ren, Z.; Misra, I.; Schwing, A. G.; and Girdhar, R. 2021. 3d spatial recognition without spatially labeled 3d. In *Proc. CVPR*, 13204–13213.
- Schult, J.; Engelmann, F.; Hermans, A.; Litany, O.; Tang, S.; and Leibe, B. 2023. Mask3d: Mask transformer for 3d semantic instance segmentation. In *Proc. ICRA*, 8216–8223. IEEE.
- Takmaz, A.; Fedele, E.; Sumner, R. W.; Pollefeys, M.; Tombari, F.; and Engelmann, F. 2023. Openmask3d: Open-vocabulary 3d instance segmentation. *Advances in NeurIPS*.
- Tang, L.; Hui, L.; and Xie, J. 2022a. Learning inter-superpoint affinity for weakly supervised 3d instance segmentation. In *Proc. ACCV*, 1282–1297.
- Tang, L.; Hui, L.; and Xie, J. 2022b. Learning inter-superpoint affinity for weakly supervised 3d instance segmentation. In *Proc. ACCV*, 1282–1297.

Tao, A.; Duan, Y.; Wei, Y.; Lu, J.; and Zhou, J. 2022. Seg-group: Seg-level supervision for 3d instance and semantic segmentation. *IEEE Trans. Image Process.*, 31: 4952–4965.

Vu, T.; Kim, K.; Luu, T. M.; Nguyen, T.; and Yoo, C. D. 2022. Softgroup for 3d instance segmentation on point clouds. In *Proc. CVPR*, 2708–2717.

Wei, J.; Lin, G.; Yap, K.-H.; Hung, T.-Y.; and Xie, L. 2020. Multi-path region mining for weakly supervised 3d semantic segmentation on point clouds. In *Proc. CVPR*, 4384–4393.

Wu, Y.; Shi, M.; Du, S.; Lu, H.; Cao, Z.; and Zhong, W. 2022. 3d instances as 1d kernels. In *Proc. ECCV*, 235–252.

Wu, Y.; Yan, Z.; Cai, S.; Li, G.; Han, X.; and Cui, S. 2023. Pointmatch: A consistency training framework for weakly supervised semantic segmentation of 3d point clouds. *Computers & Graphics*, 116: 427–436.

Xu, M.; Yin, X.; Qiu, L.; Liu, Y.; Tong, X.; and Han, X. 2025a. SAMPro3D: Locating SAM Prompts in 3d for Zero-Shot Instance Segmentation. In *Proc. 3DV*.

Xu, X.; Liu, X.; Li, J.; Yuan, Y.; Zhang, Q.; Ma, L.; Sebe, N.; and Wang, X. 2025b. 3D Weakly Supervised Semantic Segmentation via Class-Aware and Geometry-Guided Pseudo-Label Refinement. *arXiv:2510.17875*.

Xu, X.; Yuan, Y.; Li, J.; Zhang, Q.; Jie, Z.; Ma, L.; Tang, H.; Sebe, N.; and Wang, X. 2024. 3d weakly supervised semantic segmentation with 2d vision-language guidance.

Xu, X.; Yuan, Y.; Zhang, Q.; Wu, W.; Jie, Z.; Ma, L.; and Wang, X. 2025c. Weakly-Supervised 3D Visual Grounding based on Visual Language Alignment. *IEEE Transactions on Multimedia*.

Xu, Z.; Yuan, B.; Zhao, S.; Zhang, Q.; and Gao, X. 2023. Hierarchical point-based active learning for semi-supervised point cloud semantic segmentation. In *Proc. ICCV*, 18098–18108.

Yang, C.-K.; Chen, M.-H.; Chuang, Y.-Y.; and Lin, Y.-Y. 2023a. 2D-3D interlaced transformer for point cloud segmentation with scene-level supervision. In *Proc. ICCV*, 977–987.

Yang, C.-K.; Wu, J.-J.; Chen, K.-S.; Chuang, Y.-Y.; and Lin, Y.-Y. 2022. An mil-derived transformer for weakly supervised point cloud segmentation. In *Proc. CVPR*, 11830–11839.

Yang, Y.; Wu, X.; He, T.; Zhao, H.; and Liu, X. 2023b. Sam3d: Segment anything in 3d scenes. *arXiv preprint arXiv:2306.03908*.

Yin, Y.; Liu, Y.; Xiao, Y.; Cohen-Or, D.; Huang, J.; and Chen, B. 2024. Sai3d: Segment any instance in 3d scenes. In *Proc. CVPR*, 3292–3302.

Zhao, W.; Yan, Y.; Yang, C.; Ye, J.; Yang, X.; and Huang, K. 2023. Divide and conquer: 3d point cloud instance segmentation with point-wise binarization. In *Proc. ICCV*, 562–571.

Zhou, B.; Khosla, A.; Lapedriza, A.; Oliva, A.; and Torralba, A. 2016. Learning Deep Features for Discriminative Localization. In *Proc. CVPR*.