

SegDINO3D: 3D Instance Segmentation Empowered by Both Image-Level and Object-Level 2D Features

Jinyuan Qu^{1,5*}, Hongyang Li^{2,5*}, Xingyu Chen^{3,5}, Shilong Liu⁴,
Yukai Shi^{1,5}, Tianhe Ren⁵, Ruitao Jing^{1,5}, Lei Zhang^{5†}

¹Tsinghua University

²South China University of Technology

³Zhongguancun Academy

⁴Princeton University

⁵International Digital Economy Academy (IDEA)

{qjy23, jingrt24, shiyk22}@mails.tsinghua.edu.cn, ftwangyeunglei@mail.scut.edu.cn,
chenxingyu@bjzgc.a.edu.cn, sl8264@princeton.edu {rentianhe, leizhang}@idea.edu.cn,

Abstract

In this paper, we present SegDINO3D, a novel Transformer encoder-decoder framework for 3D instance segmentation. As 3D training data is generally not as sufficient as 2D training images, SegDINO3D is designed to fully leverage 2D representation from a pre-trained 2D detection model, including both image-level and object-level features, for improving 3D representation. SegDINO3D takes both a point cloud and its associated 2D images as input. In the encoder stage, it first enriches each 3D point by retrieving 2D image features from its corresponding image views and then leverages a 3D encoder for 3D context fusion. In the decoder stage, it formulates 3D object queries as 3D anchor boxes and performs cross-attention from 3D queries to 2D object queries obtained from 2D images using the 2D detection model. These 2D object queries serve as a compact object-level representation of 2D images, effectively avoiding the challenge of keeping thousands of image feature maps in the memory while faithfully preserving the knowledge of the pre-trained 2D model. The introducing of 3D box queries also enables the model to modulate cross-attention using the predicted boxes for more precise querying. SegDINO3D achieves the state-of-the-art performance on the ScanNetV2 and ScanNet200 3D instance segmentation benchmarks. Notably, on the challenging ScanNet200 dataset, SegDINO3D significantly outperforms prior methods by **+8.7** and **+6.8** mAP on the validation and hidden test sets, respectively, demonstrating its superiority.

Code — <https://github.com/IDEA-Research/SegDINO3D>

1 Introduction

3D environment understanding is a crucial capability for AI systems to interact with the physical world. While 3D reconstruction has seen remarkable progress in the past few decades from conventional multi-stage geometry-based

*These authors contributed equally. This work was done when Jinyuan Qu and Hongyang Li were interns at IDEA.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

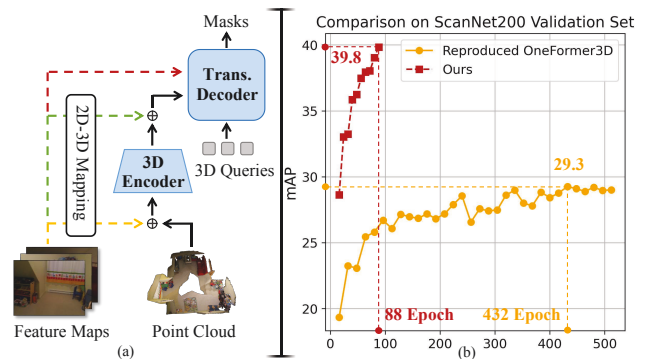


Figure 1: (a) The dashed lines indicate different ways of utilizing 2D image features. “2D-3D Mapping” denotes the association between 3D points and 2D features, and “⊕” indicates feature fusion. (b) SegDINO3D obtains significantly better performance and faster convergence speed.

pipelines (Mur-Artal and Tardós 2017; Schönberger et al. 2016) to more recent end-to-end dense reconstruction approaches (Wang et al. 2024b, 2025), 3D semantic understanding, such as 3D object instance detection and segmentation, still remains a challenge, which greatly hinders the development for downstream tasks like manipulation (Gu et al. 2024) and autonomous navigation (Xie et al. 2021).

3D instance segmentation typically takes a 3D point cloud as input and predicts a set of 3D instance masks. Recently, motivated by the success of DETECTION TRANSFORMERS (DETR) (Carion et al. 2020) in 2D detection tasks, Transformer-based frameworks (Schult et al. 2023; Sun et al. 2023) have increasingly dominated the field of 3D instance segmentation. They typically follow an encoder-decoder architecture, where a 3D encoder processes the input point cloud and a decoder predicts the object instances.

Despite having achieved great progress, most previous methods (black lines in Fig.1(a)) only consider point clouds as input, overlooking the accompanying 2D images used for reconstructing the point clouds. However, benefiting from

the abundant 2D data, 2D features trained on the massive 2D images can provide much richer and more reliable semantic information compared to 3D features trained on 3D point clouds. Therefore, recent works in open-vocabulary scene understanding explore how to leverage 2D foundation models for 3D perception tasks. For example, SAM3D (Yang et al. 2023) and SAI3D (Yin et al. 2024) directly use instance proposals generated by 2D perception models but overlook the semantic-rich 2D features. OpenScene (Peng et al. 2023) introduces 2D foundation models to extract independent 2D image features (green lines in Fig. 1(a)), yet lacks global 3D context fusion. Locate3D (Arnaud et al. 2025) thereby proposes to further train a 3D encoder through self-supervised learning, enhancing the point cloud with both 2D features and global context (yellow lines in Fig.1(a)).

However, training such a 3D encoder still requires a sufficiently large-scale 3D data set, which is far from comparable with the scale of 2D data available. As a result, the 2D features lifted to the 3D space are inadvertently degraded after passing through the 3D encoder and lead to suboptimal performance in 3D instance segmentation. Motivated by this, we hypothesize that integrating globally-enhanced 2D features in the encoder stage, while simultaneously preserving their semantic discrimination ability in the decoder stage, could yield significant improvements. Compared to the approach that maps raw 2D image features to discrete 3D points (green lines in Fig.1(a)), directly attending to 2D image feature maps to update 3D queries (red lines in Fig.1(a)), as done in BEVFormer (Li et al. 2022), is generally more effective. However, for generic 3D instance segmentation, a single scene often contains thousands of images, making it computationally expensive to cross-attend over all of them.

To more effectively leverage 2D image features, we develop SegDINO3D, a novel DETR-like 3D instance segmentation framework. In the encoder stage, to enhance 3D point features with semantically discriminative 2D information, we retrieve and aggregate 2D image-level features for each 3D point from suitable viewpoints through our Nearest View Sampling strategy. The decorated point clouds are then processed by a 3D encoder for global context fusion, effectively fusing 3D geometric structures with 2D semantic cues. As the 2D semantic features might be deteriorated after the 3D encoder due to limited 3D training data, we propose to reuse 2D features again in the decoder. More specifically, in the decoder stage, we propose a Distance-Aware Cross-Attention (DACA-2D) module. This module utilizes 2D object queries, instead of 2D image feature maps, from a 2D DETR-based detection model as a compact representation. This effectively addresses the challenge of keeping thousands of image feature maps in memory and attending to them. As these 2D object queries are directly provided by the well-trained 2D model, they retain strong semantic discrimination ability learned from large-scale 2D data. To make DACA-2D more effective, DACA-2D employs a carefully designed distance-aware attention mask to impose spatial constraints, allowing each 3D object query to cross-attend to the most relevant 2D object queries. Thanks to the symmetric design in dual-level 2D feature enhancement on both 3D points (encoder stage) and object queries (decoder

stage), their features are not only enriched but also aligned in a shared space, facilitating more accurate mask decoding via point-query similarity.

In addition, relying solely on the similarity of semantic features, as in previous methods (Schult et al. 2023; Sun et al. 2023), to query features for 3D object queries is inherently limited. Although some of the previous methods try to introduce positional embedding to provide spatial constraints (Lai et al. 2023), it is still isotropic and insufficient for capturing object scale variations, which are prevalent in 3D environments. Inspired by 2D DETR-like models (Liu et al. 2021; Zhang et al. 2022), we formulate 3D object queries as axis-aligned 3D boxes and introduce an auxiliary regression task. The predicted box sizes can modulate the positional attention map, enabling the 3D object queries to adjust their attention according to object size dynamically.

As shown in Fig.1(b), SegDINO3D not only outperforms previous methods by a large margin on the challenging ScanNet200 (Rozenberszki, Litany, and Dai 2022) dataset, but also achieves significantly faster training convergence.

In summary, our contributions are threefold:

- We propose SegDINO3D, a novel DETR-like encoder-decoder framework for 3D instance segmentation, which effectively utilizes both image-level features and object-level queries from a 2D detection model to improve 3D representations, leading to superior performance.
- We introduce 3D box queries to 3D segmentation models and use them to modulate the positional attention map, further enhancing the precision of perception ability.
- SegDINO3D achieves the state-of-the-art performance on the 3D instance segmentation benchmarks ScanNet200 and ScanNetV2. Notably, SegDINO3D surpasses prior methods by **+8.7** and **+6.8** mAP on the validation and hidden test sets of ScanNet200, respectively.

2 Related Works

3D Instance Segmentation. 3D instance segmentation aims to identify, segment, and semantically categorize each object instance within a 3D scene. Existing approaches generally fall into three categories: proposal-based, grouping-based, and Transformer-based methods. Proposal-based methods (Yang et al. 2019; Hou, Dai, and Nießner 2019; Yi et al. 2019; Engelmann et al. 2020; Kolodiazny et al. 2024b) are similar to their successful counterparts in 2D object detection, where the core idea is first to detect 3D bounding boxes and then refine them. In contrast, grouping-based methods (Liang et al. 2021; Chen et al. 2021; Vu et al. 2022; Jiang et al. 2020b; Wang et al. 2019; Jiang et al. 2020a; Zhang and Wonka 2021) do not explicitly generate object proposals. Instead, they directly cluster 3D points into distinct instances based on certain similarity criteria.

Recently, Transformer-based methods (Sun et al. 2023; Al Khatib et al. 2023; Wang et al. 2024a) have emerged and started to dominate this field. Similar to DETR-like architectures in 2D perception, they directly predict segmentation results and offer advantages in both accuracy and inference speed. Mask3D (Schult et al. 2023) first proposes a Transformer architecture for 3D instance segmentation.

QueryFormer (Lu et al. 2023) and MAFT (Lai et al. 2023) further improved query initialization and refinement. OneFormer3D (Kolodiazhnyi et al. 2024a) unifies semantic, instance, and panoptic segmentation within a single framework. Relation3D (Lu and Deng 2025) explicitly models relationships through a geometric-biased self-attention for queries and a contrastive learning scheme for scene features. These methods operate directly on 3D point clouds, while ignoring the use of 2D images, whereas ODIN (Jain et al. 2024) leverages RGB-D images and employs a unified model to perform both 2D and 3D segmentation.

Leverage 2D Models for 3D Segmentation. Recent open-vocabulary approaches attempt to incorporate 2D models into 3D segmentation. However, some methods (Yang et al. 2023; Yin et al. 2024) simply merge the mask predictions from 2D models such as SAM (Kirillov et al. 2023), without leveraging intermediate 2D features. Others (Zeid et al. 2025; Takmaz et al. 2023; Peng et al. 2023) utilize image features provided by CLIP (Radford et al. 2021) or DINOv2 (Oquab et al. 2023) from independent views, but lack global 3D context fusion across these features.

3 Method

3.1 Overview

As shown in Fig. 2 (a), in addition to taking the point cloud $\mathbf{P} \in \mathbb{R}^{N \times 3}$ with N points as input, SegDINO3D also considers the point cloud’s corresponding V posed RGB images $\mathbf{I} \in \mathbb{R}^{V \times H \times W \times 3}$ and depth maps $\mathbf{D} \in \mathbb{R}^{V \times H \times W}$, where $H \times W$ denotes the image resolution.

In the encoder stage, we use a frozen 2D DETR-based model to process the images \mathbf{I} and obtain the corresponding 2D image feature maps $\mathbf{X} \in \mathbb{R}^{V \times h \times w \times C}$, where $h \times w$ is the resolution of the feature maps¹ and C is the number of channels. Then, as shown in Fig. 2 (b), we employ Nearest view Sampling to extract 2D features $\mathbf{F}^{2D} \in \mathbb{R}^{N \times C}$ from \mathbf{X} for each 3D point. After fusing \mathbf{F}^{2D} with the point cloud \mathbf{P} , we adopt a 3D encoder to perform global 3D context fusion, resulting in 3D point features $\mathbf{F}^{3D} \in \mathbb{R}^{N \times C}$ which encapsulate both rich 2D semantic information and detailed 3D geometric structure. Finally, following the previous work (Sun et al. 2023), we aggregate the 3D point features \mathbf{F}^{3D} into superpoint features $\mathbf{S}^{3D} \in \mathbb{R}^{S \times C}$ using average pooling based on the pre-computed superpoints (Landrieu and Simonovsky 2018). Considering that the insufficiently trained 3D encoder may contaminate the high-quality 2D features, we further leverage the 2D object queries output by the 2D model. We project the 2D object queries into the 3D space and perform downsampling to provide O 2D object-level features, denoted as $\mathbf{R} \in \mathbb{R}^{O \times C}$, as illustrated in Fig. 2 (c).

In the decoder stage, we send the 3D superpoint features and the 2D object-level features to a Transformer decoder with L layers for 3D instance segmentation. Similar to the previous work (Kolodiazhnyi et al. 2024a), we randomly select M superpoints as the initial 3D object queries during training, and initialize their content part with the corresponding superpoint features \mathbf{S}^{3D} , denoted as $\mathbf{Q}^C \in \mathbb{R}^{M \times C}$.

¹For clarity, we assume the feature maps are of single scale.

While during evaluation, we regard all superpoints as initial 3D object queries, resulting in $M \equiv S$. Meanwhile, different from prior works, we additionally utilize the positions of these M superpoints to initialize their positional part $\mathbf{Q}^P \in \mathbb{R}^{M \times 6}$, where each represents an axis-aligned 3D bounding box in the scene. The content and positional components together constitute each 3D object query.

In each decoder layer, the 3D object queries first perform Box-Modulated Cross-Attention (BMCA-3D) with the 3D superpoint features \mathbf{S}^{3D} . Then, based on the Distance-Aware attention mask, they cross-attend to the 2D object queries \mathbf{R} , which preserve the original 2D information, enabling direct injection of object-level 2D cues (DACA-2D). Finally, the 3D object queries are refined by a self-attention layer followed by a feed-forward network to further enhance their feature representations. After each decoder layer, the positional part of the 3D object queries is updated by linear regression. After each decoder layer, the positional part of the 3D object queries is updated by linear regression. And the 3D segmentation results $\mathbf{M} \in \{0, 1\}^{M \times S}$ are updated by computing the similarity between the updated content feature and the 3D superpoint features \mathbf{S}^{3D} .

3.2 Feature Preparation in the Encoder Stage

Before introducing the novel modules in the decoder, we first describe how to construct the required 3D superpoint features \mathbf{S}^{3D} and 2D object queries \mathbf{R} in the encoder stage.

3D Superpoint Feature Construction. We first inject semantic information from 2D images into the 3D point cloud through Nearest View Sampling. Specifically, given known camera parameters, we can project each 3D point onto multiple views and sample features at the projected location to provide a 2D feature for each 3D point. However, a 3D point may be visible in hundreds of views, selecting a subset is crucial for computational efficiency. Considering that the closer a 3D point is to the camera’s optical center, the more details will be captured, we propose to find the top- k (in our implementation, $k = 3$) nearest views for each 3D point. Without loss of generality, given the i -th 3D point’s world coordinate as $[X_i^w, Y_i^w, Z_i^w]^\top$, its distance to V views’ camera centers $\boldsymbol{\rho}_i = [\rho_i^0, \rho_i^1, \dots, \rho_i^{V-1}]$ can be computed as

$$\rho_i^v = \|\mathbf{T}^v [X_i^w, Y_i^w, Z_i^w, 1]^\top\|_2, \quad (1)$$

where $\mathbf{T}^v \in \mathbb{R}^{4 \times 4}$ is the extrinsic parameter of the v -th camera view, $\rho_i^v \in \mathbb{R}$ is the distance between the i -th 3D point and the v -th camera. The i -th 3D point’s visibility $\boldsymbol{\delta}_i = [\delta_i^0, \delta_i^1, \dots, \delta_i^{V-1}]$ in V views can be obtained by

$$[u_i^v, v_i^v, d_i^v] = \mathbf{K}^v \mathbf{T}^v [X_i^w, Y_i^w, Z_i^w, 1]^\top, \quad (2)$$

$$\tilde{d}_i^v = \text{Bili}(\mathbf{D}_v, [u_i^v, v_i^v]),$$

$$\delta_i^v = (u_i^v \in (0, W)) \wedge (v_i^v \in (0, H)) \wedge (d_i^v \leq \tilde{d}_i^v),$$

where $\mathbf{K}^v \in \mathbb{R}^{3 \times 4}$ is the intrinsic parameter of the v -th camera view, augmented with a zero column for compatibility with homogeneous coordinates, $[u_i^v, v_i^v, d_i^v]$ is the coordinate of the i -th point in the v -th view’s image coordinate system, $\tilde{d}_i^v \in \mathbb{R}$ is the depth value sampled on the v -th view’s depth

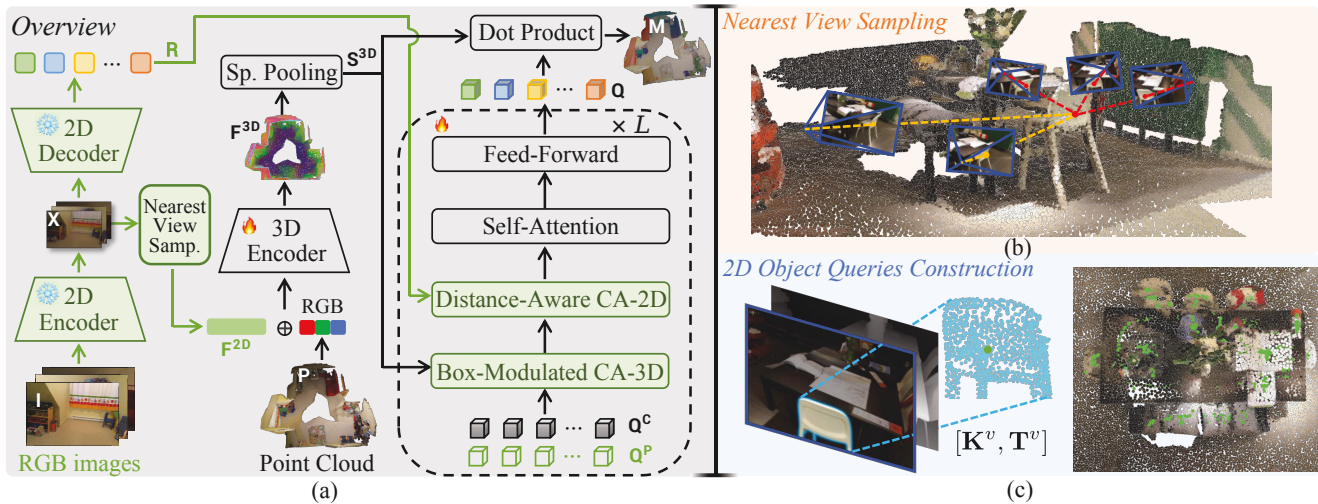


Figure 2: (a) Overview of SegDINO3D, our contributions are highlighted in green. (b) Visual illustration of the Nearest View Sampling operation. Each 3D point calculates its distance to all the views that it is visible to finds the top- k nearest views (red dash lines). (c) Visual illustration of 2D object queries construction. Each 2D object query is assigned a 3D center computed as the medoid of its corresponding 3D points, which are derived by projecting the depth values within its 2D mask using the camera parameters. The green points on the right side show the distribution of the 2D object queries’ 3D centers in the scene.

map, $\delta_i^v \in \{0, 1\}$ is the visibility of the i -th 3D point in the v -th view, and $\text{Bili}()$ is the bilinear interpolation. Based on the distances and the visibilities, we obtain the indices of the top- k nearest views $\text{Nst}(\rho_i, \delta_i, k)$ for the i -th 3D point, where $\text{Nst}()$ is the top- k nearest selection function.

After identifying the k nearest views for each 3D point, we sample and aggregate the corresponding 2D features by

$$\mathbf{F}^{2D} = [\mathbf{f}_0^{2D}, \mathbf{f}_1^{2D}, \dots, \mathbf{f}_{N-1}^{2D}],$$

$$\mathbf{f}_i^{2D} = \frac{1}{k} \sum_{v \in \text{Nst}(\rho_i, \delta_i, k)} \text{Bili}(\mathbf{X}_v, [u_i^v, v_i^v]), \quad (3)$$

where $\mathbf{f}_i^{2D} \in \mathbb{R}^C$ is the i -th 3D point’s 2D feature, and $\mathbf{X}_v \in \mathbb{R}^{h \times w \times C}$ is the v -th view’s 2D feature map.

Subsequently, we enhance the 3D point cloud \mathbf{P} with rich semantic cues from 2D by concatenating \mathbf{F}^{2D} along the feature dimension. However, since the 2D features \mathbf{F}^{2D} are extracted from different and independent 2D images, they only encode information visible from specific viewpoints, lacking a global understanding of the scene and the spatial relationships among its elements. To address this, we employ a 3D encoder to establish global contextual information for \mathbf{F}^{2D} , resulting in enhanced 3D features \mathbf{F}^{3D} . Finally, we group \mathbf{F}^{3D} into superpoint features \mathbf{S}^{3D} according to pre-computed superpoints by average pooling.

2D Object Query Construction. Since each scene corresponds to a large number of 2D images, it is impractical to take all 2D objects from every 2D image into consideration. To address this, we select only a subset. We first only retain high-confidence 2D objects, and then project them into 3D space based on the depth values within their 2D masks and the corresponding camera parameters, resulting in a 3D point cloud for each 2D object. To alleviate the effect of the depth map’s noise, we adopt the Partitioning Around

Medoids (PAM) algorithm (Kaufman and Rousseeuw 2009) to determine the 3D center of each 2D object. Then, we apply the Farthest Point Sampling algorithm to downsample the 2D objects, ensuring that they are evenly distributed in 3D space and cover as many distinct instances as possible. This results in O objects’ content features $\mathbf{R} \in \mathbb{R}^{O \times C}$ and their corresponding 3D centers $\mathbf{L}^o \in \mathbb{R}^{O \times 3}$.

It is worth highlighting that, instead of directly retrieving information from the original image feature maps via cross-attention, our constructed 2D object queries $\mathbf{R} \in \mathbb{R}^{O \times C}$ provide a significantly more memory-efficient implementation. In our implementation, for each 3D scene, we retain only $O = 2048$ queries, while the scene typically contains thousands of images. Although using the coarsest feature maps may reduce memory at the cost of information loss, the computational overhead remains high. By contrast, the queries provided by 2D models serve as both compact and semantic-rich object-level representations.

3.3 Box-Modulated Cross-Attention with 3D Superpoint

We propose box-modulated cross-attention, which modulates the attention map based on estimated object size. Specifically, we introduce a positional part $\mathbf{Q}^P \in \mathbb{R}^{M \times 6}$ for 3D object queries, each of them explicitly encodes the center coordinate and dimension of an axis-aligned 3D bounding box for a 3D instance, denoted as (x, y, z, l, w, h) . We first introduce how conventional mask cross-attention with 3D superpoints is performed. Then, we describe the implementation of our proposed box-modulated cross-attention.

Mask Cross-Attention with 3D Superpoint. Following the previous method (Kolodiazny et al. 2024a), we employ mask attention here. We first calculate the segmentation re-

sults $\mathbf{M} \in \{0, 1\}^{M \times S}$ for all 3D object queries, and then transform the segmentation results to the attention masks $\mathbf{M}_\infty \in \mathbb{R}^{M \times S}$ to constrain the attention allocation. The process of obtaining the attention mask can be formulated as

$$\mathbf{M} = \left[\sigma(\mathbf{Q}^C \mathbf{S}^{3D \top}) > \tau_{sim} \right], \mathbf{M}_\infty = (\mathbf{M} - 1) \cdot (\infty), \quad (4)$$

where σ is the sigmoid function and τ_{sim} is a similarity threshold. If the similarity between the i -th content query and the j -th superpoint is larger than τ_{sim} , the superpoint will be assigned to the i -th object. After that, we share the attention mask when querying features from the superpoint features \mathbf{S}^{3D} in cross attention, which can be formulated as

$$\begin{aligned} SimCont &= \mathbf{Q}^C \mathbf{S}^{3D \top} / \sqrt{C}, \\ \mathbf{A} &= \text{SoftMax}(SimCont + \mathbf{M}_\infty), \\ \mathbf{Q}^C &\Leftarrow \mathbf{Q}^C + \mathbf{A} \mathbf{S}^{3D}, \end{aligned} \quad (5)$$

where $SimCont$ indicates the content similarity between query and key, $\mathbf{A} \in \mathbb{R}^{M \times S}$ is the attention weights. However, the attention mask \mathbf{M}_∞ solely relies on the similarity between the content embeddings of the 3D object queries and the superpoint features, without considering any spatial constraints. This may result in high similarity scores between different instances of the same category that are spatially far apart, potentially leading to confusion. Some methods (Lai et al. 2023; Lu and Deng 2025) attempt to incorporate positional embeddings to provide spatial constraints. However, it is still isotropic and insufficient for capturing object scale variations, which are prevalent in 3D environments. Therefore, we further propose a box-modulated cross-attention mechanism that breaks the isotropy of positional embeddings, enabling more precise attention control.

Box-Modulated Cross-Attention. Since each 3D object query contains a positional component that explicitly represents the center coordinate and size of the corresponding instance, we leverage the center coordinate to add positional embedding and use the size information to modulate the positional attention map. First, by computing the dot product between the positional embeddings of queries and keys, we introduce an additional positional part to the query-to-key similarity, resulting in a change of attention weights:

$$\begin{aligned} SimPos &= \text{PE}(\mathbf{Q}_{[xyz]}^P) \text{PE}(\mathbf{L}^{sp})^\top / \sqrt{C}, \\ \mathbf{A} &= \text{SoftMax}(SimCont + SimPos + \mathbf{M}_\infty), \end{aligned} \quad (6)$$

where $\mathbf{Q}_{[xyz]}^P \in \mathbb{R}^{M \times 3}$ are the center coordinates (x, y, z) of \mathbf{Q}^P , $\mathbf{L}^{sp} \in \mathbb{R}^{S \times 3}$ are the center locations of superpoints, $SimPos$ is the positional similarity, and $\text{PE}()$ is the positional encoding function. In implementation, we use sinusoidal positional encoding with a temperature of $T = 20$.

Furthermore, we employ the dimensions (l, w, h) of \mathbf{Q}^P to modulate the positional attention maps. We define the modulation function $\text{Mod}(a, b)$ as follows

$$\begin{aligned} \mathbf{Q}^{ref} &= \sigma(\text{MLP}(\mathbf{Q}^C)), \\ \text{Mod}(x, l) &= (\text{PE}(\mathbf{Q}_{[x]}^P) \text{PE}(\mathbf{L}_{[x]}^{sp})^\top) \odot (\mathbf{Q}_{[l]}^{ref} / \mathbf{Q}_{[l]}^P), \end{aligned} \quad (7)$$

where $\mathbf{Q}^{ref} \in \mathbb{R}^{M \times 3}$ is the reference value for the dimension of the bounding box, which is predicted by an MLP from \mathbf{Q}^C . $\mathbf{Q}_{[x]}^P \in \mathbb{R}^{M \times 1}$ and $\mathbf{L}_{[x]}^{sp} \in \mathbb{R}^{S \times 1}$ are the x -coordinate of \mathbf{Q}^P and \mathbf{L}^{sp} , respectively. \odot denotes element-wise multiplication with broadcasting. $\mathbf{Q}_{[l]}^{ref} \in \mathbb{R}^{M \times 1}$ and $\mathbf{Q}_{[l]}^P \in \mathbb{R}^{M \times 1}$ are the l -dimension of \mathbf{Q}^{ref} and $\mathbf{Q}_{[l]}^P$. The modulated positional similarity is computed as

$$SimPos = \frac{\text{Mod}(x, l) + \text{Mod}(y, w) + \text{Mod}(z, h)}{\sqrt{C}}. \quad (8)$$

This modulation allows the attention map to adapt based on the size of the objects, enhancing the model’s ability to focus on relevant spatial regions.

3.4 Distance-Aware Cross-Attention with 2D Object Queries

Given that a well-trained 2D DETR-like model is capable of providing semantic-rich 2D object queries, we propose to conduct cross-attention between 2D and 3D object queries.

Distance-Aware Attention Mask Construction. Cross-attending to all 2D object queries from every 3D object query incurs extra computational cost and may cause confusion. Therefore, we introduce a distance-aware attention mask based on spatial proximity, allowing each 3D object query to only attend to nearby 2D object queries, which are more likely to belong to the same 3D instance. However, since objects captured in 2D images are often partial observations, the 3D centers of the 2D objects may deviate significantly from their corresponding 3D objects’ centers. As a result, directly computing the attention mask based on the 3D spatial distance between 3D object queries and 2D objects may bring noise. We propose to first find the spatial relationship between 2D objects and superpoints, then, according to the assignment relationship between superpoints and 3D object queries (from segmentation results \mathbf{M}), we obtain the attention mask between 3D object queries and 2D objects $\mathbf{M}_\infty^o \in \mathbb{R}^{M \times O}$. The process can be formulated as

$$\begin{aligned} dist &= \text{Dist}(\mathbf{L}^o, \mathbf{L}^{sp}), \quad \mathbf{M}^d = [dist < \tau_{dist}], \\ \mathbf{M}^o &= \mathbf{M} \mathbf{M}^d, \quad \mathbf{M}_\infty^o = (\mathbf{M}^o - 1) \cdot (\infty), \end{aligned} \quad (9)$$

where $\text{Dist}()$ is an operator that calculates the $L2$ distance $dist \in \mathbb{R}^{S \times O}$ between every 3D superpoint and 2D object, τ_{dist} is the distance threshold to determine whether a 3D superpoint and a 2D object are related.

Cross-Attention with 2D Object Queries. Based on the attention mask, we further optimize the 3D content queries with the 2D object-level features through cross-attention

$$\begin{aligned} \mathbf{A}^o &= \text{SoftMax}(\mathbf{Q}^C \mathbf{R}^\top / \sqrt{C} + \mathbf{M}_\infty^o), \\ \mathbf{Q}^C &\Leftarrow \mathbf{Q}^C + \mathbf{A}^o \mathbf{R}, \end{aligned} \quad (10)$$

where $\mathbf{A}^o \in \mathbb{R}^{M \times O}$ is the attention weight.

Method	All			Head			Common			Tail		
	mAP	mAP ₅₀	mAP ₂₅	mAP	mAP ₅₀	mAP ₂₅	mAP	mAP ₅₀	mAP ₂₅	mAP	mAP ₅₀	mAP ₂₅
Validation Set												
OpenMask3D (2023)	15.4	19.9	23.1	-	-	-	-	-	-	-	-	-
Mask3D (2023)	27.4	37.0	42.3	40.3	55.0	62.2	22.4	30.6	35.4	18.2	23.2	27.0
MAFT (2023)	29.2	38.2	43.3	-	-	-	-	-	-	-	-	-
SAI3D (2024)	12.7	18.8	24.1	12.1	-	-	10.4	-	-	16.2	-	-
Open3DIS (2024)	23.7	29.4	32.8	27.8	-	-	21.3	-	-	21.8	-	-
Open-YOLO 3D (2024)	24.7	31.7	36.2	27.8	-	-	24.3	-	-	21.6	-	-
OneFormer3D (2024a)	30.2	40.9	44.6	42.0	57.7	63.9	27.0	36.3	39.8	20.1	26.6	27.7
ODIN (2024)	31.5	45.3	53.1	37.5	54.2	66.1	31.6	43.9	50.2	24.1	36.6	41.2
Relation3D (2025)	31.6	41.2	45.6	-	-	-	-	-	-	-	-	-
Ours-Grounding DINO	36.2	47.6	54.2	44.8	60.9	69.6	32.0	42.3	47.3	31.1	38.2	44.4
Ours-DINO-X	40.2	52.4	58.6	46.3	63.7	71.5	37.4	47.6	52.3	36.2	44.9	51.0
Hidden Test Set												
Mask3D (2023)	27.8	38.8	44.5	38.3	54.2	65.3	26.3	35.7	39.2	16.8	23.7	25.4
TD3D (2024b)	21.1	32.0	37.9	33.2	50.1	60.3	17.7	26.4	30.6	10.3	16.4	19.0
ODIN (2024)	26.5	38.1	45.1	34.9	50.7	63.7	26.8	37.5	40.7	16.3	23.7	27.7
Ours-DINO-X	34.6	45.4	51.1	43.7	58.7	68.5	35.3	45.3	48.4	22.9	29.6	33.1

Table 1: Comparison of SegDINO3D with prior methods on validation set and hidden test set of ScanNet200.

3.5 Updating of Mask and Box

The mask update follows Equation 4 in Sec. 3.3. For each box query, the center coordinate (x, y, z) is initialized using the coordinates of its corresponding superpoint, while the dimension (l, w, h) is initialized with fixed values. After each decoder layer, we apply an MLP to predict a residual offset, enabling layer-by-layer refinement of \mathbf{Q}^P . The update process at the l -th decoder layer can be formulated as

$$\mathbf{Q}_{l+1}^P = \mathbf{Q}_l^P + \text{MLP}(\mathbf{Q}_{l+1}^C). \quad (11)$$

To supervise the 3D positional queries \mathbf{Q}^P , we set up an additional regression task. We construct axis-aligned box labels for all the instances based on the ground truth masks, and the predicted boxes are supervised by an L1 loss. Moreover, we incorporate a bounding box regression cost into the matching strategy during training.

Finally, it is worth noting that with negligible computational overhead, we can use the predicted box of each 3D object query to filter out mask regions that are far from the box, thereby further improving segmentation accuracy.

4 Experiments

4.1 Experiment Setup

Dataset. We conduct experiments on the ScanNetv2 (Dai et al. 2017) and ScanNet200 (Rozenberszki, Litany, and Dai 2022) datasets. ScanNetv2 is a richly annotated 3D indoor scene reconstruction dataset, containing 20 semantic classes and annotations for 18 object instances. It provides point clouds, RGB-D images, and corresponding camera intrinsics and extrinsics. The dataset is divided into training, validation, and hidden test splits, consisting of 1201, 312, and 100 scenes, respectively. ScanNet200 extends the original ScanNetv2 by significantly increasing the number of instance categories, providing 198 instance-level classes and 2 semantic classes. According to the number of annotations, these

200 classes are further categorized into “head”, “common”, and “tail” splits, with 66, 68, and 66 categories respectively. This long-tailed distribution, with the head classes being more frequent and the tail classes more sparse, makes ScanNet200 a more challenging benchmark, especially for evaluating model performance on real-world long-tail distributions. To validate the superiority of SegDINO3D, we report the detailed performance on the 3 sub-splits of ScanNet200.

Evaluation Metrics. We adopt the standard evaluation metric for 3D instance segmentation, mean Average Precision (mAP), which measures the area under the precision-recall curve across different Intersection-over-Union (IoU) thresholds. Specifically, mAP is computed by averaging the precision values at IoU thresholds ranging from 50% to 95%, with an interval of 5%, providing a comprehensive assessment of segmentation quality. In addition to the overall mAP, we also report mAP₅₀ and mAP₂₅, which reflect model performance at fixed IoU thresholds of 50% and 25%, respectively. These metrics are particularly useful for analyzing performance under different levels of localization tolerance.

Implementation Details. For the hyperparameter settings, we primarily follow OneFormer3D (Kolodiaznyy et al. 2024a) and adopt it as our baseline. Specifically, we voxelize the 3D points with a voxel size of 2cm and employ a U-Net-like sparse convolutional network as our 3D encoder (Choy, Gwak, and Savarese 2019) to globally encode the point cloud enhanced by image-level features. During training, we use the AdamW (Loshchilov 2017) optimizer with an initial learning rate of 1×10^{-4} , a weight decay of 0.05, a batch size of 4, and a polynomial learning rate scheduler with a base of 0.9. We adopt the disentangled matching strategy proposed in OneFormer3D, and retain the cross-entropy loss \mathcal{L}_{cls} for classification, the binary cross-entropy loss \mathcal{L}_{bce} and Dice loss \mathcal{L}_{dice} for the superpoint mask, as well as the binary cross-entropy loss \mathcal{L}_{sem} for semantic. We

also employ the same data augmentation strategies, including horizontal flipping, random rotations around the z-axis, elastic distortion, and random scaling.

Distinct from previous methods, SegDINO3D achieves significantly faster convergence, reducing the number of training epochs from 432 to 88, which results in about a $5\times$ reduction in training time (see Fig. 1(b)). In addition, we develop two versions, incorporating either Grounding DINO (Liu et al. 2024) or DINO-X (Ren et al. 2024) as the 2D model to provide 2D features and queries. For the box regression task, we further introduce an additional L1 loss \mathcal{L}_{box} . Therefore, the overall loss function is:

$$\mathcal{L} = \lambda_1 \cdot \mathcal{L}_{cls} + \mathcal{L}_{bce} + \mathcal{L}_{dice} + \lambda_2 \cdot \mathcal{L}_{sem} + \lambda_3 \cdot \mathcal{L}_{box} \quad (12)$$

where $\lambda_1 = \lambda_2 = \lambda_3 = 0.5$ in our implementation. Both training and inference of SegDINO3D can be performed on a single NVIDIA RTX 3090 GPU.

4.2 Comparison with Prior Work

We compare our method with previous approaches on the ScanNet200 validation and hidden test set. Our method offers two versions that utilize either Grounding DINO or DINO-X as the 2D model to provide features. In the following discussions, we refer to the DINO-X-based version for comparison. As shown in Table 1, on the validation set, SegDINO3D outperforms all existing methods by a large margin in all sub-splits, achieving a +8.7 mAP higher mAP than the previous best method. Especially, SegDINO3D’s performance on the tail-split even surpasses previous methods on the common-split (36.2 mAP vs. 31.6 mAP). This indicates that the incorporation of 2D features significantly alleviates the issue of insufficient 3D training data for long-tailed categories. Meanwhile, on the hidden test set, SegDINO3D also shows significant improvement, establishing a new state of the art. These consistent and substantial improvements demonstrate the superiority of SegDINO3D.

Method	mAP	mAP ₅₀	mAP ₂₅
SSTNet (2021)	49.4	64.3	74.0
SoftGroup (2022)	46.0	67.6	78.9
Mask3D (2023)	55.2	73.7	83.5
SPFormer (2023)	56.3	73.9	82.9
QueryFormer (2023)	56.5	74.2	83.3
MAFT (2023)	59.9	76.5	-
OneFormer3D (2024a)	59.3	78.1	86.4
TD3D (2024b)	47.3	71.2	81.9
ODIN (2024)	50.0	71.0	83.6
Spherical Mask (2024)	62.3	79.9	88.2
Relation3D (2025)	62.5	80.2	87.0
Ours-DINO-X	64.0	81.5	88.9

Table 2: Comparison of SegDINO3D with prior methods on the validation set of ScanNetv2.

Moreover, we compare SegDINO3D with prior methods on the ScanNetv2 validation set. As shown in Table 2, our method outperforms all existing methods across all metrics. Specifically, it outperforms our baseline by +4.7 mAP and

Row	Img. F.	Obj. F.	Mod.	mAP	mAP ₅₀	mAP ₂₅
1	✗	✗	✗	26.8	36.9	42.1
2	✓	✗	✗	35.2	46.9	54.6
3	✗	✓	✗	33.0	46.6	54.9
4	✓	✓	✗	38.3	50.5	57.4
5	✓	✓	✓	39.8	52.1	58.6
6 [†]	✓	✓	✓	40.2	52.4	58.6

Table 3: Ablation on our key components. Img. F. denotes the use of 2D image-level features in the encoder, Obj. F. refers to the incorporation of 2D object-level features in the decoder, and Mod. indicates the use of box modulated cross-attention. [†] denotes using predicted boxes to filter masks.

Methods	mAP	mAP ₅₀	mAP ₂₅
w/o Global Context Fusion	35.9	48.7	55.4
w Global Context Fusion	39.8	52.1	58.6

Table 4: Ablation on Global Context Fusion.

surpasses the previous state-of-the-art method by +1.5 mAP. Compared with the improvement achieved on the more challenging ScanNet200, the improvement on ScanNetv2 is relatively smaller. We attribute this to the fact that 2D representations provide more discriminative semantic features, which are particularly advantageous under the fine-grained 200-class setting of ScanNet200.

Sampling Methods	mAP	mAP ₅₀	mAP ₂₅
Farthest Sampling	36.2	48.5	56.3
Nearest Sampling	39.8	52.1	58.6

Table 5: Ablation on Nearest View Sampling.

4.3 Ablation Studies and Analysis

We conduct ablations on the challenging ScanNet200 validation set to assess the effectiveness of our key components.

Key Components. We progressively add components to the baseline to validate their effectiveness. As shown in rows 1~3 of Table 3, incorporating globally enhanced 2D image-level features and leveraging information directly from raw 2D object-level features lead to improvements of +8.4 and +6.2 mAP, respectively. Moreover, thanks to their complementarity, where the image-level features enhance the 3D point cloud scene with 2D semantic information, while the object-level features directly inject the rich information from the original 2D object queries into the 3D object queries, combining both brings an additional gain of +3.1 mAP, as shown in rows 2 and 4. Subsequently, as illustrated in rows 4 and 5, introducing box-modulated cross-attention using the estimated boxes yields a further gain of +1.5 mAP. Finally, as shown in rows 5 and 6, using the estimated boxes to filter out distant masks also brings a slight improvement.

Global Context Fusion. After establishing the correspondence between each 3D point and 2D image-level features,

Methods	mAP	mAP ₅₀	mAP ₂₅
w/o Distance-Aware Mask	36.3	48.7	54.9
w Distance-Aware Mask	39.8	52.1	58.6

Table 6: Ablation on Distance-Aware Mask.

Row	Center PE.	Box Reg.	Box Mod.	mAP
1	✗	✗	✗	38.3
2	✓	✗	✗	38.1
3	✓	✓	✗	37.8
4	✓	✓	✓	39.8

Table 7: Ablation on box queries. Center PE. indicates using positional embedding with center coordinates, Box Reg. indicates adding box regression loss, and Box Mod. indicates using box to modulate positional attention map.

we apply a 3D encoder for global context fusion. Alternatively, if we simply fuse the corresponding 2D image-level features to the point cloud through an MLP, as reported in Table 4, there will be a -3.9 mAP performance drop, highlighting the effectiveness of global contextualization.

Nearest View Sampling. As discussed in Sec.3.2, we compare two sampling strategies: selecting the nearest k views and selecting the farthest k views. As shown in Table 5, the “nearest sampling” strategy yields superior performance, suggesting that 2D information from views closer to the camera’s optical center tends to be more informative.

Distance-Aware Attention Mask in DACA-2D. As shown in Table 6, when extracting information from 2D object queries, employing our proposed distance-aware mask significantly improves performance by selectively attending to the most relevant queries based on spatial relationships, leading to a +3.5 mAP improvement.

Box Modulation in BMCA-3D. To validate the specific contributions of our proposed box-modulated cross-attention, we progressively integrate components in the order of implementation. As shown in Row 2 of Table 7, we first introduce positional queries and add a center-point regression task, using the center coordinates to incorporate position embedding into the cross-attention. In row 3, we further add a box size regression task. It is observed that neither of these additions yields a significant performance improvement. Finally, we modulate the positional part of the cross-attention using the estimated box, resulting in a 1.5 mAP improvement, demonstrating its effectiveness.

5 Conclusion

In this paper, we have presented SegDINO3D, a novel method for 3D instance segmentation. SegDINO3D comprehensively leverages both 2D image-level and object-level features from pre-trained 2D DETR-based detection models, effectively transferring the strong capabilities of 2D models into the 3D domain. For the 2D image-level features, we select the most suitable views for each 3D point to extract features, aggregate them, and then perform global context

fusion. For the 2D object-level features, through a carefully designed distance-aware mask, we can effectively extract semantic-rich 2D information into 3D queries in a memory-efficient manner. Finally, we introduce 3D box queries and modulate the positional attention maps using the estimated boxes to further enhance the precision of perception ability. SegDINO3D achieves state-of-the-art results on the 3D instance segmentation benchmarks ScanNetv2 and ScanNet200, and converges much faster during training, demonstrating its superiority.

References

- Al Khatib, S.; El Amine Boudjoghra, M.; Lahoud, J.; and Khan, F. S. 2023. 3D Instance Segmentation via Enhanced Spatial and Semantic Supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 541–550.
- Arnaud, S.; McVay, P.; Martin, A.; Majumdar, A.; Jatavallabhula, K. M.; Thomas, P.; Partsey, R.; Dugas, D.; Gejji, A.; Sax, A.; et al. 2025. Locate 3D: Real-World Object Localization via Self-Supervised Learning in 3D. *arXiv preprint arXiv:2504.14151*.
- Boudjoghra, M. E. A.; Dai, A.; Lahoud, J.; Cholakkal, H.; Anwer, R. M.; Khan, S.; and Khan, F. S. 2024. Open-YOLO 3D: Towards Fast and Accurate Open-Vocabulary 3D Instance Segmentation. *arXiv preprint arXiv:2406.02548*.
- Carion, N.; Massa, F.; Synnaeve, G.; Usunier, N.; Kirillov, A.; and Zagoruyko, S. 2020. End-to-End Object Detection with Transformers. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I 16*, 213–229. Springer.
- Chen, S.; Fang, J.; Zhang, Q.; Liu, W.; and Wang, X. 2021. Hierarchical Aggregation for 3D Instance Segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15467–15476.
- Choy, C.; Gwak, J.; and Savarese, S. 2019. 4D Spatio-Temporal ConvNets: Minkowski Convolutional Neural Networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 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 *Proceedings of the IEEE conference on computer vision and pattern recognition*, 5828–5839.
- Engelmann, F.; Bokeloh, M.; Fathi, A.; Leibe, B.; and Nießner, M. 2020. 3D-MPA: Multi-Proposal Aggregation for 3D Semantic Instance Segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9031–9040.
- Gu, Q.; Kuwajerwala, A.; Morin, S.; Jatavallabhula, K. M.; Sen, B.; Agarwal, A.; Rivera, C.; Paul, W.; Ellis, K.; Chellappa, R.; et al. 2024. ConceptGraphs: Open-Vocabulary 3D Scene Graphs for Perception and Planning. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 5021–5028. IEEE.
- Hou, J.; Dai, A.; and Nießner, M. 2019. 3D-SIS: 3D Semantic Instance Segmentation of RGB-D Scans. In *Proceedings*

- of the *IEEE/CVF conference on computer vision and pattern recognition*, 4421–4430.
- Jain, A.; Katara, P.; Gkanatsios, N.; Harley, A. W.; Sarch, G.; Aggarwal, K.; Chaudhary, V.; and Fragkiadaki, K. 2024. ODIN: A Single Model for 2D and 3D Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3564–3574.
- Jiang, H.; Yan, F.; Cai, J.; Zheng, J.; and Xiao, J. 2020a. End-to-End 3D Point Cloud Instance Segmentation Without Detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12796–12805.
- Jiang, L.; Zhao, H.; Shi, S.; Liu, S.; Fu, C.-W.; and Jia, J. 2020b. PointGroup: Dual-Set Point Grouping for 3D Instance Segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and Pattern recognition*, 4867–4876.
- Kaufman, L.; and Rousseeuw, P. J. 2009. *Finding groups in data: an introduction to cluster analysis*. John Wiley & Sons.
- 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 *Proceedings of the IEEE/CVF international conference on computer vision*, 4015–4026.
- Kolodiazhnyi, M.; Vorontsova, A.; Konushin, A.; and Rukhovich, D. 2024a. OneFormer3D: One Transformer for Unified Point Cloud Segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20943–20953.
- Kolodiazhnyi, M.; Vorontsova, A.; Konushin, A.; and Rukhovich, D. 2024b. Top-Down Beats Bottom-Up in 3D Instance Segmentation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3566–3574.
- Lai, X.; Yuan, Y.; Chu, R.; Chen, Y.; Hu, H.; and Jia, J. 2023. Mask-Attention-Free Transformer for 3D Instance Segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3693–3703.
- Landrieu, L.; and Simonovsky, M. 2018. Large-Scale Point Cloud Semantic Segmentation With Superpoint Graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4558–4567.
- Li, Z.; Wang, W.; Li, H.; Xie, E.; Sima, C.; Lu, T.; Qiao, Y.; and Dai, J. 2022. BEVFormer: Learning Bird’s-Eye-View Representation from Multi-Camera Images via Spatiotemporal Transformers. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part IX*, 1–18. Springer.
- Liang, Z.; Li, Z.; Xu, S.; Tan, M.; and Jia, K. 2021. Instance Segmentation in 3D Scenes Using Semantic Superpoint Tree Networks. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2783–2792.
- Liu, S.; Li, F.; Zhang, H.; Yang, X.; Qi, X.; Su, H.; Zhu, J.; and Zhang, L. 2021. DAB-DETR: Dynamic Anchor Boxes are Better Queries for DETR. In *International Conference on Learning Representations*.
- Liu, S.; Zeng, Z.; Ren, T.; Li, F.; Zhang, H.; Yang, J.; Jiang, Q.; Li, C.; Yang, J.; Su, H.; et al. 2024. Grounding DINO: Marrying DINO with Grounded Pre-training for Open-Set Object Detection. In *European Conference on Computer Vision*, 38–55. Springer.
- Loshchilov, I. 2017. Decoupled Weight Decay Regularization. *arXiv preprint arXiv:1711.05101*.
- Lu, J.; and Deng, J. 2025. Relation3D: Enhancing Relation Modeling for Point Cloud Instance Segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 8889–8899.
- Lu, J.; Deng, J.; Wang, C.; He, J.; and Zhang, T. 2023. Query Refinement Transformer for 3D Instance Segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 18516–18526.
- Mur-Artal, R.; and Tardós, J. D. 2017. ORB-SLAM2: An Open-Source SLAM System for Monocular, Stereo, and RGB-D Cameras. *IEEE transactions on robotics*, 33(5): 1255–1262.
- Nguyen, P.; Ngo, T. D.; Kalogerakis, E.; Gan, C.; Tran, A.; Pham, C.; and Nguyen, K. 2024. Open3DIS: Open-Vocabulary 3D Instance Segmentation with 2D Mask Guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4018–4028.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. DINOv2: Learning Robust Visual Features without Supervision. *arXiv preprint arXiv:2304.07193*.
- Peng, S.; Genova, K.; Jiang, C.; Tagliasacchi, A.; Pollefeys, M.; Funkhouser, T.; et al. 2023. OpenScene: 3D Scene Understanding With Open Vocabularies. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 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 *International conference on machine learning*, 8748–8763. PmLR.
- Ren, T.; Chen, Y.; Jiang, Q.; Zeng, Z.; Xiong, Y.; Liu, W.; Ma, Z.; Shen, J.; Gao, Y.; Jiang, X.; et al. 2024. DINO-X: A Unified Vision Model for Open-World Object Detection and Understanding. *arXiv preprint arXiv:2411.14347*.
- Rozenberszki, D.; Litany, O.; and Dai, A. 2022. Language-Grounded Indoor 3D Semantic Segmentation in the Wild. In *European Conference on Computer Vision*, 125–141. Springer.
- Schönberger, J. L.; Zheng, E.; Pollefeys, M.; and Frahm, J.-M. 2016. Pixelwise View Selection for Unstructured Multi-View Stereo. In *European Conference on Computer Vision (ECCV)*.
- Schult, J.; Engelmann, F.; Hermans, A.; Litany, O.; Tang, S.; and Leibe, B. 2023. Mask3D: Mask Transformer for 3D Semantic Instance Segmentation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 8216–8223. IEEE.

- Shin, S.; Zhou, K.; Vankadari, M.; Markham, A.; and Trigoni, N. 2024. Spherical Mask: Coarse-to-Fine 3D Point Cloud Instance Segmentation with Spherical Representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4060–4069.
- Sun, J.; Qing, C.; Tan, J.; and Xu, X. 2023. Superpoint Transformer for 3D Scene Instance Segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2393–2401.
- Takmaz, A.; Fedele, E.; Sumner, R. W.; Pollefeys, M.; Tombari, F.; and Engelmann, F. 2023. OpenMask3D: Open-Vocabulary 3D Instance Segmentation. *arXiv preprint arXiv:2306.13631*.
- Vu, T.; Kim, K.; Luu, T. M.; Nguyen, T.; and Yoo, C. D. 2022. SoftGroup for 3D Instance Segmentation on Point Clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2708–2717.
- Wang, D.; Liu, J.; Gong, H.; Quan, Y.; and Wang, D. 2024a. CompetitorFormer: Competitor Transformer for 3D Instance Segmentation. *arXiv preprint arXiv:2411.14179*.
- Wang, J.; Chen, M.; Karaev, N.; Vedaldi, A.; Rupprecht, C.; and Novotny, D. 2025. VGGT: Visual Geometry Grounded Transformer. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, 5294–5306.
- Wang, S.; Leroy, V.; Cabon, Y.; Chidlovskii, B.; and Revaud, J. 2024b. DUS3R: Geometric 3D Vision Made Easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20697–20709.
- Wang, X.; Liu, S.; Shen, X.; Shen, C.; and Jia, J. 2019. Associatively Segmenting Instances and Semantics in Point Clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 4096–4105.
- Xie, C.; Xiang, Y.; Mousavian, A.; and Fox, D. 2021. Unseen Object Instance Segmentation for Robotic Environments. *IEEE Transactions on Robotics*, 37(5): 1343–1359.
- Yang, B.; Wang, J.; Clark, R.; Hu, Q.; Wang, S.; Markham, A.; and Trigoni, N. 2019. Learning Object Bounding Boxes for 3D Instance Segmentation on Point Clouds. *Advances in neural information processing systems*, 32.
- Yang, Y.; Wu, X.; He, T.; Zhao, H.; and Liu, X. 2023. SAM3D: Segment Anything in 3D Scenes. *arXiv preprint arXiv:2306.03908*.
- Yi, L.; Zhao, W.; Wang, H.; Sung, M.; and Guibas, L. J. 2019. 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*, 3947–3956.
- Yin, Y.; Liu, Y.; Xiao, Y.; Cohen-Or, D.; Huang, J.; and Chen, B. 2024. SAI3D: Segment Any Instance in 3D Scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3292–3302.
- Zeid, K. A.; Yilmaz, K.; de Geus, D.; Hermans, A.; Adrian, D.; Linder, T.; and Leibe, B. 2025. DINO in the Room: Leveraging 2D Foundation Models for 3D Segmentation. *arXiv preprint arXiv:2503.18944*.
- Zhang, B.; and Wonka, P. 2021. Point Cloud Instance Segmentation Using Probabilistic Embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8883–8892.
- Zhang, H.; Li, F.; Liu, S.; Zhang, L.; Su, H.; Zhu, J.; Ni, L.; and Shum, H.-Y. 2022. DINO: DETR with Improved De-Noising Anchor Boxes for End-to-End Object Detection. In *The Eleventh International Conference on Learning Representations*.