

2D-3D Interlaced Transformer for Point Cloud Segmentation with Scene-Level Supervision

Cheng-Kun Yang^{1,2} Min-Hung Chen³ Yung-Yu Chuang¹ Yen-Yu Lin^{4,5}

¹National Taiwan University ²MediaTek ³NVIDIA

⁴National Yang Ming Chiao Tung University ⁵Academia Sinica

Abstract

We present a *Multimodal Interlaced Transformer (MIT)* that jointly considers 2D and 3D data for weakly supervised point cloud segmentation. Research studies have shown that 2D and 3D features are complementary for point cloud segmentation. However, existing methods require extra 2D annotations to achieve 2D-3D information fusion. Considering the high annotation cost of point clouds, effective 2D and 3D feature fusion based on weakly supervised learning is in great demand. To this end, we propose a transformer model with two encoders and one decoder for weakly supervised point cloud segmentation using only scene-level class tags. Specifically, the two encoders compute the self-attended features for 3D point clouds and 2D multi-view images, respectively. The decoder implements interlaced 2D-3D cross-attention and carries out implicit 2D and 3D feature fusion. We alternately switch the roles of queries and key-value pairs in the decoder layers. It turns out that the 2D and 3D features are iteratively enriched by each other. Experiments show that it performs favorably against existing weakly supervised point cloud segmentation methods by a large margin on the S3DIS and ScanNet benchmarks. The project page will be available at https://jimmy15923.github.io/mit_web/.

1. Introduction

Point cloud segmentation offers rich geometric and semantic information of a 3D scene, thereby being essential to many 3D applications, such as scene understanding [5, 10, 16, 26, 36], augmented reality [2, 35], and autonomous driving [7, 8, 13]. However, developing reliable models is time-consuming and challenging due to the need for vast per-point annotations and the difficulty in capturing detailed semantic clues from textureless point clouds.

Research efforts have been made to address the aforementioned issues. Several methods are proposed to derive point cloud segmentation models using various weak supervisions, such as sparsely labeled points [24, 33, 58, 65],

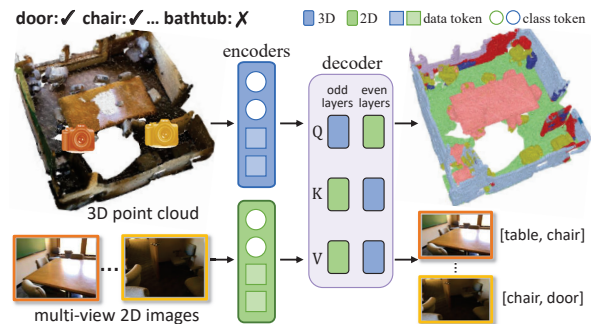


Figure 1: Overview of the Multimodal Interlaced Transformer (MIT). The input includes a 3D point cloud, multi-view 2D images, and class-level tags of a scene. Our method is a transformer model with two encoders and one decoder. The two encoders compute features for 3D voxel tokens and 2D view tokens, respectively. The decoder conducts interlaced 2D-3D attention and carries out 2D and 3D feature fusion. In its odd layers, 3D voxels serve as queries and are enriched by the semantic features of 2D views, acting as key-value pairs. In the even layers, the roles of 3D voxels and 2D views switch: 2D views are described by additional 3D geometric features.

bounding box labels [9], subcloud-level annotations [54], and scene-level tags [40, 61]. These weak annotations are cost-efficient and can significantly reduce the annotation burden. On the other hand, recent studies [19, 20, 23, 32, 41, 52, 53, 62] witness the remarkable success of 2D vision, and utilize 2D image features to enhance the 3D recognition task. They show promising results because 2D detailed texture clues are well complementary to 3D geometry features.

Although 2D-3D fusion is effective, current methods require extra annotation costs for 2D images. To the best of our knowledge, no prior work has explored fusing 2D-3D features under *extremely weak supervision*, where only scene-level class tags of the 3D scene are given. It is challenging to derive a segmentation model that leverages both 2D and 3D data under scene-level supervision, as no per-point/pixel annotations or per-image class tags are avail-

able to guide the learning process. Furthermore, existing 2D-3D fusion methods require camera poses or depth maps to establish pixel-to-point correspondences, adding extra burdens on data collection and processing. In this work, we address these difficulties by proposing a **Multimodal Interlaced Transformer (MIT)** that works with scene-level supervision and can implicitly fuse 2D, and 3D features without camera poses and depth maps.

Our MIT is a transformer model with two encoders and one decoder, and can carry out weakly supervised point cloud segmentation. As shown in Figure 1, the input to our method includes the 3D point cloud, multi-view images, and scene-level tags of a scene. The two encoders utilize the self-attention mechanism to extract the features of the 3D point cloud and the 2D multi-view images, respectively. The decoder computes the proposed interlaced 2D-3D attention and can implicitly fuse the 2D and 3D data.

Specifically, one encoder is derived for 3D feature extraction, where the voxels of the input point cloud yield the data tokens. The other encoder is for 2D multi-view images, where images serve as data tokens. Also, the multi-class tokens [57] are included to match the class-level annotations. The encoders capture long-range dependencies and aggregate class-specific features for their respective modalities.

The decoder comprises 2D-3D interlaced layers, and is developed to fuse 2D and 3D features, where the correspondences between 3D voxels and 2D views are implicitly computed via cross-attention. In odd layers of the decoder, 3D voxels are enriched by 2D image features, while in even layers, 2D views are augmented by 3D geometric features. Specifically, in each odd layer, each 3D voxel serves as a *query*, while 2D views act as *key-value pairs*. Through cross-attention, a query is a weighted combination of the values. Together with residual learning, this query (3D voxel) is characterized by the fused 3D and 2D features. In each even layer, the roles of 3D voxels and 2D views switch: 3D voxels and 2D views become key-value pairs and queries, respectively. This way, 2D views are described by the augmented 2D and 3D features.

By leveraging multi-view information without extra annotation effort, our proposed MIT effectively fuses the 2D and 3D features and significantly improves 3D point cloud segmentation. The main contribution of this work is three-fold. First, to the best of our knowledge, we make the first attempt to fuse 2D-3D information for point cloud segmentation under scene-level supervision. Second, we enable this new task by presenting a new model named Multimodal Interlaced Transformer (MIT) that implicitly fuses 2D-3D information via interlaced attention, which does not rely on camera pose information. Besides, a contrastive loss is developed to align the class tokens across modalities. Third, our method performs favorably against existing methods on the large-scale ScanNet [11] and S3DIS [3] benchmarks.

2. Related Work

Weakly supervised point cloud segmentation. This task aims at learning a point cloud segmentation model using weakly annotated data, such as sparsely labeled points [16, 18, 27, 28, 30, 33, 43, 45, 47, 49, 56, 58, 65, 66], box-level labels [9], subcloud-level labels [54, 61] and scene-level labels [24, 40]. Significant progress has been made in the setting of using sparsely labeled points: The state-of-the-art methods [18, 33, 63] show comparable performances with supervised ones. These methods usually utilize self-supervised pre-training [16, 33], graph propagation [33, 47], and contrastive learning [28, 66] to derive the models. Despite the effectiveness, they require at least one annotated point for each category in a scene. Hence, it is not straightforward to extend these methods to work with scene-level or subcloud-level annotations.

In this work, we aim to develop a segmentation method based on a more challenging setting of using scene-level annotations. The literature about point cloud segmentation with scene-level annotations is relatively rare. Yang *et al.* [61] derive a transformer by applying multiple instance learning to paired point clouds. However, their performance is much inferior to fully supervised methods. Kweon and Yoon [24] leverage 2D and 3D data for point cloud segmentation by introducing additional image-level class tags, which require extra annotation efforts. Our method compensates for the lack of point-level or pixel-level annotations by integrating additional 2D features while using scene-level annotation only.

2D and 3D fusion for point cloud applications. Given the accessible or syntheticable [23] 2D images in most 3D dataset, research studies [14, 17, 19, 20, 23, 24, 29, 32, 41, 46, 52, 53, 55, 59, 60, 62, 64] explore 2D data to enhance 3D applications. Hu *et al.* [19] and Robert *et al.* [41] construct a pixel-point mapping matrix to fuse 2D and 3D features for point cloud segmentation. Despite the effectiveness, existing methods rely on camera poses and/or depth maps to build the correspondences between the 2D and 3D domains. In contrast, our method learns a transformer with interlaced 2D-3D attention, enabling the implicit integration of 2D and 3D features without the need for camera poses or depth maps.

Query and key-value pair swapping. Cross-attention is widely used in the transformer decoder. It captures the dependency between queries and key-value pairs. Umam *et al.* [48] and Lim *et al.* [21] swap queries and key-value pairs for point cloud decomposition and generation, respectively. Different from their methods working with data in a domain, our method generalizes query and key-value pair swapping to cross-domain feature fusion. In addition, we develop a contrastive loss for 2D and 3D feature alignment.

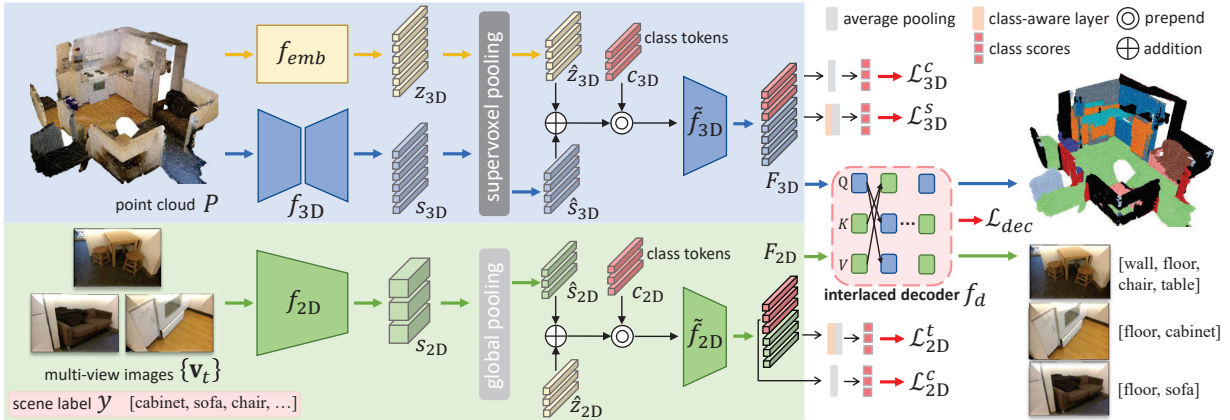


Figure 2: An overview of our **Multimodal Interlaced Transformer (MIT)** for weakly supervised point cloud segmentation. It is a transformer-based model with two encoders, \tilde{f}_{3D} and \tilde{f}_{2D} , for modality-specific feature extraction and one decoder, f_d , for feature fusion. The 2D and 3D pooled features, \hat{s}_{2D} and \hat{s}_{3D} , are added to each learnable position embedding (\hat{z}_{2D} and \hat{z}_{3D}), and further prepended with the class tokens and passed through the encoders to obtain self-attended features, F_{2D} and F_{3D} . The predicted class scores for each modality are obtained through average pooling and class-aware layers.

3. Proposed Method

We present the proposed method in this section. We first give the problem statement in Section 3.1. Then, we specify the developed MIT with an encoder-decoder architecture in Section 3.2 and Section 3.3. Finally, the implementation details are provided in Section 3.4.

3.1. Problem Statement

We are given a set of N point clouds as well as their corresponding RGB multi-view images and the class tag annotations, *i.e.*, $\{P_n, V_n, \mathbf{y}_n\}_{n=1}^N$, where P_n denotes the n th point cloud, V_n represents the multi-view images, and \mathbf{y}_n is the class-level labels. Note that P_n , V_n , and \mathbf{y}_n are acquired from the same scene. Without loss of generality, we assume that each point cloud consists of M points, *i.e.*, $P_n = \{\mathbf{p}_{nm}\}_{m=1}^M$, where each point $\mathbf{p}_{nm} \in \mathbb{R}^6$ is represented by its 3D coordinate and RGB color. The RGB multi-view images are grabbed in the same scene as P_n , and consist of a set of T images, *i.e.*, $V_n = \{\mathbf{v}_{nt}\}_{t=1}^T$. Each image $\mathbf{v}_{nt} \in \mathbb{R}^{H \times W \times 3}$ is of resolution $H \times W$ with RGB channels. The class tags of P_n , *i.e.*, $\mathbf{y}_n \in \{0, 1\}^C$, are a C -dimensional binary vector storing which categories are present, where C is the number of categories of interest.

With the weakly annotated dataset $\{P_n, V_n, \mathbf{y}_n\}_{n=1}^N$, we aim to derive a model for point cloud segmentation that classifies each point of a testing cloud into one of the C categories. Note that in this weakly supervised setting, neither points nor pixels are labeled, and camera poses are unavailable, making it challenging to enhance 3D point cloud segmentation by incorporating additional 2D features due to the absence of point/pixel supervision and explicit correspondences between 2D pixels and 3D points. Furthermore,

as multi-view images share the same scene-level class tag, the lack of individual class tag annotation for each view image may lead to an inaccurate semantic understanding of each image.

Method overview. Figure 2 illustrates the network architecture of MIT, which comprises two transformer encoders, \tilde{f}_{3D} and \tilde{f}_{2D} , and one decoder f_d . The two encoders extract features for 3D point clouds and 2D multi-view images, respectively. The decoder is developed for 2D-3D feature fusion, which utilizes cross-attention to connect 2D and 3D data implicitly. They are elaborated in the following.

3.2. Transformer Encoders

3D point cloud feature extraction. A 3D backbone f_{3D} , *e.g.*, MinkowskiNet [10] or PointNet++ [38], is applied to extract the point embedding $s_{3D} \in \mathbb{R}^{M \times D}$ for all M points of a point cloud P . Like WYPR [31], we perform supervoxel partition using an unsupervised off-the-shelf algorithm [25]. The 3D coordinates of P are fed into a coordinate embedding module f_{emb} , which is composed of two 1×1 convolution layers with ReLU activation, to get the positional embedding $z_{3D} \in \mathbb{R}^{M \times D}$, where D is the embedding dimension. We aggregate both the point features and point positional embedding through supervoxel average pooling [33], producing the supervoxel features $\hat{s}_{3D} \in \mathbb{R}^{S \times D}$ and pooled positional embedding $\hat{z}_{3D} \in \mathbb{R}^{S \times D}$, where S is the number of the supervoxels in P . The supervoxel features are added to the positional embedding.

To learn the class-specific representation for fitting the scene-level supervision, we prepend C learnable class tokens [57] $c_{3D} \in \mathbb{R}^{C \times D}$ with S supervoxel tokens. Total $(C + S)$ tokens are fed into the transformer encoder \tilde{f}_{3D} .

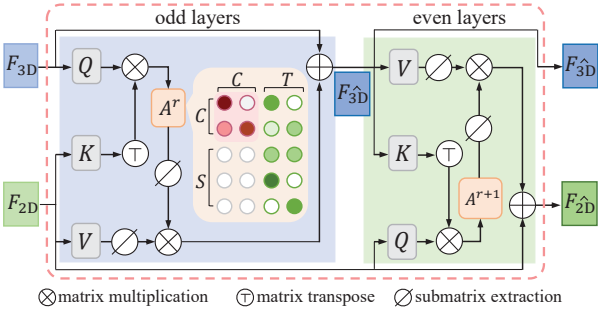


Figure 3: The architecture of an interlaced block. The multilayer perceptron with residual learning is not present for simplicity but is used in the block.

Through the self-attention mechanism, the dependencies of the class and supervoxel tokens are captured, producing the self-attended 3D features $F_{3D} \in \mathbb{R}^{(C+S) \times D}$.

2D multi-view images feature extraction. A 2D backbone network f_{2D} , e.g. ResNet [15], is employed to extract image features $s_{2D} \in \mathbb{R}^{T \times H' \times W' \times D}$, where $H' = H/32$ and $W' = W/32$. We apply global average pooling to image features s_{2D} along the spatial dimensions. The pooled image features $\hat{s}_{2D} \in \mathbb{R}^{T \times D}$ are added to the learnable positional embedding $\hat{z}_{2D} \in \mathbb{R}^{T \times D}$, producing T view tokens. Analogous to 3D feature extraction, another transformer encoder \tilde{f}_{2D} is applied to C class tokens $c_{2D} \in \mathbb{R}^{C \times D}$ and T view tokens, obtaining the self-attended 2D features $F_{2D} \in \mathbb{R}^{(C+T) \times D}$.

Encoder optimization. During training, we consider a point cloud P and its associated T multi-view images $\{\mathbf{v}_t\}$ and scene-level label \mathbf{y} . The 2D and 3D self-attended features F_{2D} and F_{3D} are compiled as specified above. We conduct the multi-label classification losses [40, 57] for optimization.

For 3D attended features $F_{3D} \in \mathbb{R}^{(C+S) \times D}$, we divide it into C class tokens $F_{3D}^c \in \mathbb{R}^{C \times D}$ and S supervoxel tokens $F_{3D}^s \in \mathbb{R}^{S \times D}$. For the class tokens F_{3D}^c , the C class scores are estimated by applying average pooling along the feature dimension. The multi-label classification loss \mathcal{L}_{3D}^c is computed based on the estimated class scores and the scene-level ground-truth labels \mathbf{y} . For the supervoxel tokens F_{3D}^s , we introduce a class-aware layer [44], i.e., a 1×1 convolution layer with C filters, which maps the supervoxel tokens F_{3D}^s into the class activation maps (CAM) $\tilde{F}_{3D}^s \in \mathbb{R}^{S \times C}$. The estimated class scores are obtained by applying global average pooling to \tilde{F}_{3D}^s along the dimension of supervoxels. The multi-label classification loss \mathcal{L}_{3D}^s is computed based on the class scores and label \mathbf{y} . The loss for the 3D modality is defined by $\mathcal{L}_{3D} = \mathcal{L}_{3D}^c + \mathcal{L}_{3D}^s$. For the self-attended 2D features $F_{2D} \in \mathbb{R}^{(C+T) \times D}$ of the C class tokens and T view tokens, the 2D loss is similarly defined by $\mathcal{L}_{2D} = \mathcal{L}_{2D}^c + \mathcal{L}_{2D}^t$.

In sum, both encoders are derived in a weakly-supervised manner using the objective function

$$\mathcal{L}_{enc} = \mathcal{L}_{2D} + \mathcal{L}_{3D}. \quad (1)$$

3.3. Transformer Decoder

The two encoders produce self-attended 3D features F_{3D} of $C + S$ tokens and 2D features F_{2D} of $C + T$ tokens, respectively. We propose a decoder that performs interlaced 2D-3D cross-attention for feature fusion. The decoder f_d in Figure 2 is a stack of R interlaced blocks. Each interlaced block is composed of two successive decoder layers, as shown in Figure 3. In the first layer of this block, 3D tokens are enriched by 2D features, while in the second layer, 2D tokens are enriched by 3D features.

In the odd/first layer (the blue-shaded region in Figure 3), the $C + S$ tokens in F_{3D} serve as the queries, while the $C + T$ tokens in F_{2D} act as the key-value pairs. Through scaled dot-product attention [50], the cross-modal attention matrix $A \in \mathbb{R}^{(C+S) \times (C+T)}$ (the yellow-shaded region) is computed to store the consensus between the 3D tokens and 2D tokens. As we focus on exploring the relationships between 3D tokens and merely 2D view tokens in this layer, we ignore the attention values related to the 2D class tokens in A . Specifically, only the query-to-view attention values $A_{q2v} \in \mathbb{R}^{(C+S) \times T}$ (green dots in Figure 3) are considered. This is implemented by applying submatrix extraction to the attention matrix A and the value matrix $V \in \mathbb{R}^{(C+T) \times D}$, i.e., $A_{q2v} = A[1: C + S, C + 1: C + T]$ and $V_d = V[C + 1: C + T, :]$.

After applying the softmax operation to A_{q2v} , we perform matrix multiplication between the query-to-view attention matrix A_{q2v} and the masked value matrix V_d . This way, each query (3D token) is a weighted sum of the values (2D view tokens). Together with a residual connection, the resultant 3D tokens $F_{3D}^{\hat{}}$ are enriched by 2D features. It turns out that implicit feature fusion from 3D features to 2D features is carried out without using annotated data.

In the even/second layer (the green-shaded region in Figure 3), the roles of $F_{3D}^{\hat{}}$ and F_{2D} switch: The former serves as the key-value pairs while the latter yields the queries. After a similar procedure, the resultant 2D tokens $F_{2D}^{\hat{}}$ $\in \mathbb{R}^{(C+T) \times D}$ are augmented with 3D information. $F_{2D}^{\hat{}}$ and $F_{3D}^{\hat{}}$ are the output of the interlaced block. By stacking R interlaced blocks, the proposed decoder is built to fuse 2D and 3D features iteratively.

Decoder optimization. In the last interlaced block, the 2D class scores and 3D class scores can be estimated by applying average pooling to the corresponding class tokens. The multi-label classification losses for 2D $\mathcal{L}_{2D}^{\hat{}}$ and 3D $\mathcal{L}_{3D}^{\hat{}}$ data can be computed between the ground truth and the estimated class scores.

To mine additional supervisory signals, we employ contrastive learning on the class-to-class attention matrix $A_{c2c} = A[1:C, 1:C] \in \mathbb{R}^{C \times C}$. Though the 2D class tokens and 3D class tokens attend to respective modalities, they share the same class tags. Hence, the attention value between a pair of class tokens belonging to the same class should be larger than those between tokens of different classes, which can be enforced by the N-pair loss [39]. We employ this regularization in all attention matrices in the decoder layers

$$\begin{aligned} \mathcal{L}_{con} = & \frac{1}{2R} \sum_{r=1}^{2R} \sum_{i=1}^C -\log \frac{A_{ii}^r}{\sum_{j=1}^C A_{ij}^r} \\ & + \frac{1}{2R} \sum_{r=1}^{2R} \sum_{j=1}^C -\log \frac{A_{jj}^r}{\sum_{i=1}^C A_{ij}^r}, \end{aligned} \quad (2)$$

where A^r is the attention matrix in the r th decoder layer.

The objective function of learning the decoder is

$$\mathcal{L}_{dec} = \mathcal{L}_{2D} + \mathcal{L}_{3D} + \alpha \mathcal{L}_{con}, \quad (3)$$

where α is a positive constant.

3.4. Implementation Details

The proposed method is implemented in PyTorch. ResNet-50 [15] pre-trained on ImageNet [42] serves as the 2D feature extractor while MinkowskiNet [10] works as the 3D feature extractor in the experiments. The numbers of heads, encoder layers, interlaced blocks, embedding dimension, and the width of FFN in the transformer are set to 4, 3, 2, 96, and 96, respectively. 16 multi-view images are randomly sampled for each scene. The model is trained on eight NVIDIA 3090 GPUs with 500 epochs. The batch size, learning rate, and weight decay are set to 32, 10^{-2} , and 10^{-4} , respectively. We use AdamW [22] as the optimizer. The weight α for \mathcal{L}_{con} is set to 0.5.

Inference. Given a point cloud P for inference, we feed it into our 3D encoder for feature extraction. The 3D CAM $\tilde{F}_{3D}^s \in \mathbb{R}^{S \times C}$, *i.e.*, the segmentation result, is then obtained by passing the extracted features into the class-aware layer, as specified in Section 3.2. In MCTformer [57], 3D CAM can be further refined by the class-to-voxel attention maps $A_{c2s} \in \mathbb{R}^{C \times S}$ from the last K transformer encoder layers, where $K = 3$. The refined 3D CAM is obtained through element-wise multiplication between CAM and the attention maps: $F = \tilde{F}_{3D}^s \odot A_{c2s}$, where \odot denotes Hadamard product. In addition, we consider the class-to-voxel attention maps in the interlaced decoder, if multi-view images are provided, which can be extracted from all the even layers, producing another refined 3D CAM \hat{F} . Finally, the segmentation results can be obtained by applying the element-wise max operation to F and \hat{F} .

We followed the common practice in [6,33,40,54,58,61] to generate pseudo-segmentation labels by running inference on the training set. Then use a segmentation model, *e.g.* Res U-Net [16], to train on the pseudo labels with high confidence, *i.e.* over 0.5, and derive the segmentation model with 150 epochs. No further post-processing is applied.

4. Experimental Results

This section evaluates the proposed method. We begin by introducing the datasets and evaluation metrics. The competing methods are then presented and compared. Finally, we present ablation studies for each proposed component and analysis of our method.

4.1. Datasets and Evaluation Metrics

The experiments were conducted using two large-scale point cloud datasets with multi-view images, S3DIS [3] and ScanNet [11]. S3DIS [3] contains 272 scenes from six indoor areas. A total of 70,496 RGB images are collected. Each scene is represented by a point cloud with 3D coordinates and RGB values. Each point and pixel is labeled with one of 13 categories. Following previous works [37, 38, 40, 51, 58], area 5 is taken as the test scene. ScanNet [11] includes 1,201 training scenes, 312 validation scenes, and 100 test scenes with 20 classes. Over 2.5 million RGB images are collected. Following [19], we sample one image out of every twenty to avoid redundancy in image selection. The mean intersection over Union (mIoU) is employed as the evaluation metric for both datasets.

4.2. Competing Methods and Comparisons

We compare our MIT with competing weakly supervised point cloud segmentation and 2D-3D feature fusion methods.

4.2.1 Point Cloud Segmentation Method Comparison

We compare our proposed method to state-of-the-art methods for segmenting point clouds with scene-level supervision. We also consider methods utilizing different supervision signals and extra data as input. To begin with, we present the fully supervised methods [10, 38, 41, 53] for point cloud segmentation as they suggest potential performance upper bounds. Next, we show the methods [24, 33, 63] that employ various types of weak labels. Finally, we compare the segmentation methods [40, 54, 61] utilizing scene-level labels that indicate whether each class appears in the scene.

Table 1 reports the mIoU results of the competing methods using different types of supervision or extra input data, such as RGB images, camera poses, or depth maps. Existing methods that fuse 2D images with 3D data have demonstrated superior performance compared to 3D-only meth-

Method	Sup.	Extra inputs			ScanNet		S3DIS
		RGB	Pose	Depth	Val	Test	Test
MinkUNet [10]	\mathcal{F} .	-	-	-	72.2	73.6	65.8
DeepViewAgg [41]	\mathcal{F} .	✓	✓	-	71.0	-	67.2
SemAffiNet [53]	\mathcal{F} .	✓	✓	✓	-	74.9	71.6
OTOC [33]	\mathcal{P} .	-	-	-	-	59.4	50.1
Yu <i>et al.</i> [63]	\mathcal{P} .	✓	✓	✓	-	63.9	-
Kweon <i>et al.</i> [24]	$\mathcal{S} + \mathcal{I}$.	✓	✓	✓	49.6	47.4	-
MPRM [54]	\mathcal{S} .	-	-	-	24.4	-	10.3
MIL-Trans [61]	\mathcal{S} .	-	-	-	26.2	-	12.9
WYPR [40]	\mathcal{S} .	-	-	-	29.6	24.0	22.3
MIT (3D-only)	\mathcal{S} .	-	-	-	31.6	26.4	23.1
MIT (Ours)	\mathcal{S} .	✓	-	-	35.8	31.7	27.7

Table 1: Quantitative results (mIoU) of several point-cloud segmentation methods with diverse supervisions and input data settings on the ScanNet and S3DIS datasets. “Sup.” indicates the type of supervision. “ \mathcal{F} .” represents full annotation. “ \mathcal{P} .” gives sparsely labeled points. “ \mathcal{S} .” denotes scene-level annotation. “ \mathcal{I} .” implies image-level annotation.

ods. However, the reliance on camera poses or depth maps limits their applicability. In contrast, our MIT can benefit from 2D images without such requirements, enhancing its generalizability.

By using efficient scene-level annotation, our MIT with 3D data only (the blue-shaded region in Figure 2) shows comparable results to the state-of-the-art weakly supervised method [40], demonstrating the effectiveness of transformer encoder with the multi-class token [57]. The proposed interlace decoder further enhances the performance of the MIT with 3D-only data by incorporating the 2D image information. Without introducing extra annotation costs, our method with 2D-3D fusion outperforms the existing methods by a large margin on both the ScanNet and S3DIS datasets. This result demonstrates once again that 2D and 3D data are complementary. More importantly, the proposed method is capable of utilizing their complementarity in a weakly supervised manner.

Kweon *et al.* [24] also confirms the effectiveness of combining 2D-3D data. However, their method requires non-negligible extra annotation costs on the images. According to [4, 61], their method incurs more than five times the annotation cost required for scene-level annotation and even more than the sparsely labeled points setting.

We summarize the advantages of the scene-level setting in three aspects: **1) Efficiency:** Scene-level supervision is much more efficient to collect than other weak supervision types. According to [40, 61], the labeling cost of sparse points (1% of points in ScanNet) is more than ten times higher than our scene-level setting. **2) Generalization:** Our method based on scene-level supervision can be extended to other forms of weak supervision. Section 4.3.2 evaluates our method trained with diverse weak supervision types. **3)**

Method	Fusion	Pose	Depth	mIoU
MIL-Trans [61]	MLP	✓	-	25.6
MIL-Trans [61]	BPM [19]	✓	✓	25.9
MIT	MLP	✓	-	32.6
MIT	BPM [19]	✓	✓	32.4
MIT	Interlaced	-	-	35.8
MIT	Interlaced	✓	✓	37.1

Table 2: Quantitative results (mIoU) of our method (interlaced decoder) and competing methods with different 2D-3D fusion strategies on the ScanNet validation set using scene-level annotations.

Potential: Existing weakly supervised point cloud segmentation methods focus on the sparse-point supervision setting and achieve performances almost as good as fully supervised ones. Therefore, working with lower annotation costs, such as scene-level tags, shows potential and is worth exploring. Our method effectively carries it out by utilizing information from unlabeled images.

4.2.2 2D-3D Fusion Method Comparison

As far as we are aware, our MIT is the first attempt at exploring 2D-3D fusion without *poses*, and the model is derived through *scene-level supervision*. Hence, there is no existing fusion method for performance comparison. To evaluate our method, we explore two approaches for 2D-3D feature fusion. First, we design a baseline method using a simple multi-layer perceptron for 2D-3D fusion. For each 3D voxel, we locate the nearest 2D pixel and concatenate the 3D feature with the 2D feature, followed by a 1×1 convolution to perform 2D-3D feature fusion. Second, we employ the bidirectional projection module [19] for 2D-3D fusion, which utilizes the pixel-to-point link matrix to fuse the 2D-3D features.

We apply the 2D-3D fusion methods on a weakly supervised point cloud segmentation method, MIL-transformer [61], as well as our proposed method. Table 2 provides the mIoU results of the competing 2D-3D fusion methods. Our proposed interlaced decoder achieves superior results compared to the two competing 2D-3D fusion methods. More importantly, the interlaced decoder implements 2D-3D fusion without using poses or depths and performs even better when camera information is available (more details in Section 4.3.2 and supplementary materials).

Our interlaced decoder offers the following advantages. *Multi-view aggregation:* The view quality differs in different views of the same 3D point, such as occlusion, or no 2D-3D correspondence. Through the attention mechanism, the decoder learns how to effectively aggregates the multi-view information based on the semantic information. *Global attention:* The decoder can capture long-range dependencies,

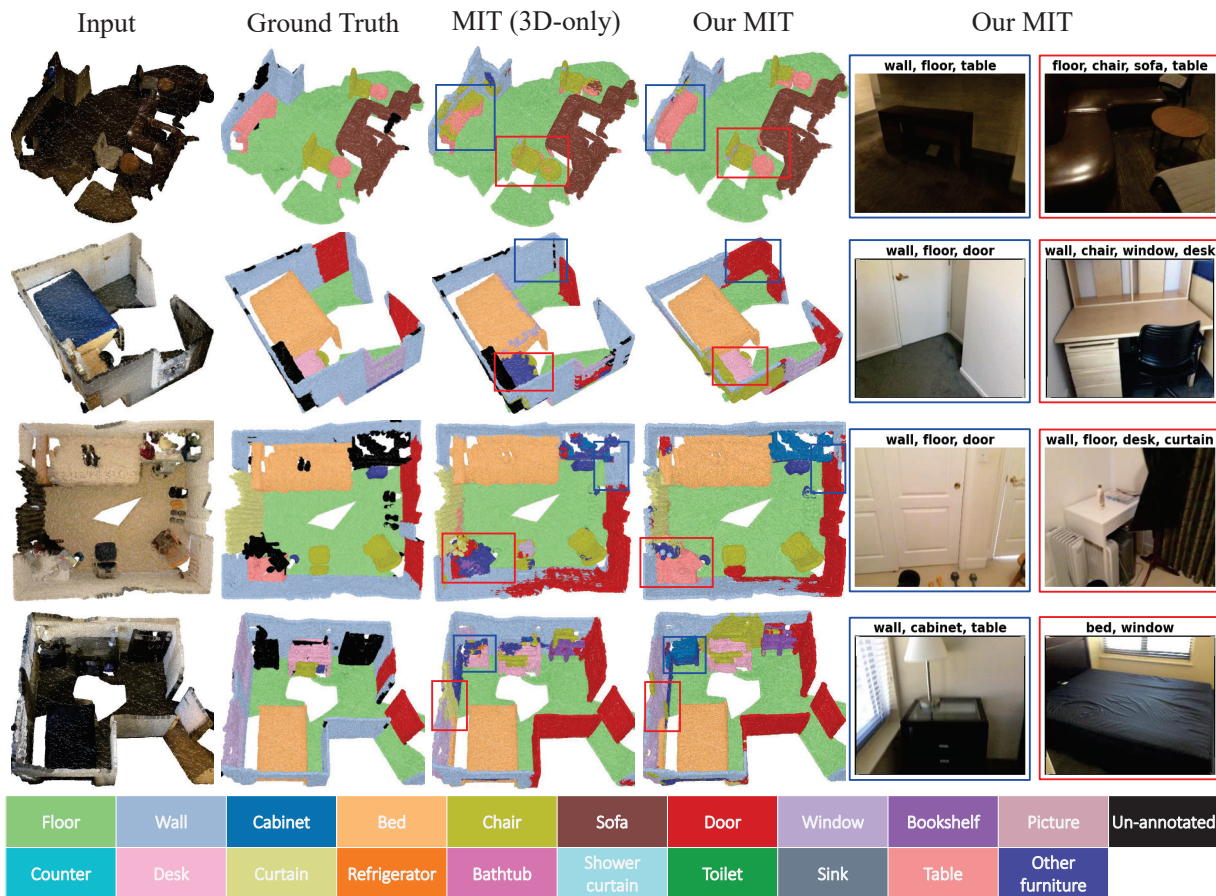


Figure 4: Qualitative results on the ScanNet dataset with scene-level supervision. The colored boxes highlight the differences between our MIT and MIT with 3D data only, and their corresponding views are shown on the right with outlines of the same color. For each view, the tags at the top indicate the results of the multi-label classification.

i.e., the receptive field is the whole scene. *Low overhead:* The computational bottleneck of the decoder lies in cross-attention, whose complexity is linear to $N_{2D} \times N_{3D}$, where N_{2D} and N_{3D} are the numbers of 2D and 3D tokens, respectively. Since we cast each 2D view into a token via global average pooling, $N_{2D} = C + T$, where C is the number of classes and T is the number of 2D views. As shown in Table 4, we can achieve good results by giving $T = 16$ views; hence N_{2D} can be small. To summarize, the proposed interlaced decoder introduces an acceptable cost but provides multi-view aggregation with global attention. Moreover, our interlaced decoder can further enrich features in 2D and 3D domains for better 3D segmentation.

4.2.3 Qualitative Results

Figure 4 shows the qualitative results of our MIT with and without using the complementary 2D data. By utilizing both 3D and 2D data, our method achieves promising segmenta-

tion results without using any point-level supervision. With the help of detailed texture features in the 2D image, our method is able to classify objects with very similar geometric shapes, for example, door and wall. Take the second row of Figure 4 as an example; MIT successfully segments the points belonging to the door by cooperating with the correct prediction from our 2D view (marked in blue), while the 3D only MIT fails to locate the points of the door, by only considering geometric and color features.

In addition, the category co-occurrence issue could hinder the optimization of the model with 3D data only. Since optimization is based on scene-level labels, it is difficult to learn discriminative features for those co-occurring categories. As demonstrated in the second and third rows of Figure 4, MIT (3D-only) often fails to classify the chairs and tables since these categories often co-occur in a scene. In contrast, our method leverages multi-view information during training. As each view only captures a small part of the scene, the issue of category co-occurrence could be alle-

Query	Key-Value	MCT	Interlaced	\mathcal{L}_{con}	mIoU
-	-				26.1
-	-	✓			31.6
3D	2D	✓			33.7
3D	2D	✓	✓		35.4
2D	3D	✓	✓		35.2
3D	2D	✓	✓	✓	35.8

Table 3: The mIoU performance of different combinations of proposed components on the validation set of the ScanNet dataset. “Query” and “Key-Value” denote the input to the decoder. “MCT” and “Interlaced” are the multi-class tokens encoder and decoder architectures respectively. \mathcal{L}_{con} denotes the contrastive loss on the class tokens.

viated, resulting in better segmentation performance. With the proposed interlaced decoder, the network can learn more corresponding features between view and voxel under weak supervision. Additionally, the data tokens with position embedding and class tokens with contrastive loss facilitate the linking of views and voxels.

4.3. Ablation Studies and Performance Analysis

To evaluate the effectiveness of the proposed components, we perform ablation studies and analyze their performance. We present ablation studies to evaluate the impacts of the proposed components and provide performance analysis.

4.3.1 Contributions of Components

To evaluate the effectiveness of each proposed component, we first construct the baseline by considering only 3D data and utilizing class activation maps [54,58]. Then, we assess the contributions of each component, including the multi-class token transformer encoder (MCT), the interlaced decoder (Interlaced), and the N-pair loss (\mathcal{L}_{con}), by successively adding each one to the baseline. In addition, we evaluate the roles for 2D and 3D, as query and key-value pairs, by switching them. The result of the standard transformer decoder is also reported (the third row of Table 3) by taking 3D as query and 2D as key-value. Table 3 illustrates the performance when using different combinations of the proposed modules and loss. The results validate that each component contributes to the performance of our method.

4.3.2 Performance Analysis

We discuss the extension of the proposed method and evaluate our method with different parameters and synthesized images in the following.

Number of views	4	16	32	64
mIoU	29.7	32.7	30.9	31.2

Table 4: Performance with different numbers of views on the mIoU of pseudo labels on the ScanNet.

R interlaced blocks	1	2	3	4
mIoU	31.4	32.7	32.1	32.4

Table 5: Performance with different numbers of interlaced blocks on the mIoU of pseudo labels on the ScanNet.

Extension with known poses and depths. When camera poses, and depth maps are available, the correspondence between 3D world coordinates and 2D pixels can be established. Therefore, we can explicitly construct the position correlation between 2D views and 3D voxels. To this end, we first generate the 3D world coordinate maps for each view by following Yu *et al.* [63]. All the 3D coordinate maps are fed into the coordinate embedding module f_{emb} to obtain positional embedding, which is then added to the 2D image features. Through explicit positional information between the 2D view and 3D voxel, we can further boost performance, as shown in the last row of Table 2.

Analysis of parameters. We explore the influence of the number of 2D views and interlaced blocks by evaluating the quality of pseudo labels on the training set. Table 4 shows the performance of our MIT with different numbers of views used in the transformer. We found that performance is stable when given a sufficient number of views, as also reported in [19]. Table 5 presents the performance by altering the number of the proposed interlaced blocks. The results indicate that stacking two interlaced blocks performs the best while being saturated by adding more blocks.

Experiments with virtual image rendering. Among the limitations of our method is the need for multi-view 2D images within the 3D dataset. A potential solution would be the virtual view rendering of the 3D data. Several studies [23,34] suggest that synthesized images can further improve 3D segmentation performance. With the help of virtual view rendering [23], our method still achieves competitive results (34.3% mIoU on ScanNet validation set) using the synthesized RGB images.

Extensions to other weak supervisions. Thanks to the flexibility of the transformer, our method can be easily adapted to other weakly supervised settings, such as additional image-level labels, subcloud-level annotation, or sparsely labeled points.

	Scene	Scene+Image	Subcloud	20pts
mIoU	35.8	45.4	46.8	61.9
label effort	< 1 min	5 min	3 min	2 min

Table 6: The mIoU performance of our MIT and its average annotation time per scene of different weak supervisions on ScanNet.

For the extra image-label annotation, it provides the class tag indicating the existing object category within each view image. Several methods [1, 44] are proposed to derive a 2D segmentation model based on this supervision and achieve promising results. Our method can easily train on image-level supervision by computing the multi-label classification loss on each image token.

Regarding the subcloud-level annotation, it sequentially crops a sphere point cloud from the scene and labels the existing objects within the sphere. This type of supervision alleviates the severe class imbalance issue in scene-level supervision. Our approach can be directly trained on subcloud-level supervision by considering the corresponding multi-view images in the subcloud.

For the setting with sparsely labeled points [33, 61], we can calculate the cross entropy loss on the self-attended voxel features \hat{F}_{3D} and the labeled points. Furthermore, we note that the sparsely labeled 3D points can be projected onto the 2D image pixels, generating 2D pixel annotation. In spite of this, we do not explore this operation in our experiments and leave it for future research.

Table 6 shows the performance of our method under different types of weak supervision and the corresponding annotation cost. While scene-level annotation is the most efficient [61], its performance has room for improvement. The extra image-level labels can improve the performance of scene-level supervision but introduce additional burdens due to the large number of view images in each scene. According to [4], the image-level labels cost 20 seconds per image. In line with [24], which utilized 17 multi-view images in their setting, we used 16 images per scene, resulting in an additional five-minute annotation time. Even though both image-level and subcloud-level supervision types do not require point-level annotation, they could require more annotation efforts due to the large number of views and subclouds that need to be annotated. Sparsely labeled points, on the other hand, may perform better with less annotation effort.

Our approach can work effectively with diverse weak supervision, allowing for flexible savings in annotation costs. More importantly, our MIT shows promising results by using the most efficient scene-level supervision, while other weakly supervised methods cannot be straightforwardly applied in this scenario.

3D Backbone	2D Backbone	mIoU
ResUNet-18	ResNet-50	32.7
ResUNet-18	ResNet-101	33.1
ResUNet-34	ResNet-101	32.9

Table 7: Performance with different backbones on the mIoU of pseudo labels on the ScanNet.

Experiments with different backbones. Table 7 provides the performance (pseudo-label quality in mIoU) of our method on ScanNet with different 2D and 3D backbones, including different versions of 2D ResNet and 3D ResUNet. Our method’s performance is consistent across different backbones.

5. Conclusion

This paper presents a multimodal interlaced transformer, MIT, for weakly supervised point cloud segmentation. Our method represents the first attempt at 2D and 3D information fusion with scene-level annotation. Through the use of the proposed interlaced decoder, which performs implicit 2D-3D feature fusion via cross-attention, we are able to effectively fuse 2D-3D features without using camera poses or depth maps. Our MIT achieves promising performance without using any point-level or pixel-level annotations. Furthermore, we develop class token consistency to align the multimodal features. MIT is end-to-end trainable. It has been extensively evaluated on two challenging real-world large-scale datasets. Experiments show that our method performs favorably against existing weakly supervised methods. We believe MIT has the potential to enhance other recognition tasks that involve both 2D and 3D observations, in an efficient manner.

Discussion and future work. Our current method has not utilized the spatial information conveyed in the images since global average pooling is applied to the image features. We attempted to flatten image features instead of using global average pooling to obtain patch tokens, similar to [12, 32], but achieved inferior results. One possible reason is that a large number of patch tokens introduces noise under scene-level supervision. A solution to this issue can achieve joint 2D-3D segmentation with weak supervision, which could be an interesting area for future research.

Acknowledgements. This work was supported in part by the National Science and Technology Council (NSTC) under grants 111-2628-E-A49-025-MY3, 112-2221-E-A49-090-MY3, 111-2634-F-002-023, 111-2634-F-006-012, 110-2221-E-002-124-MY3, and 111-2634-F-002-022. This work was funded in part by MediaTek, Qualcomm, NVIDIA, and NTU-112L900902.

References

- [1] Jiwoon Ahn and Suha Kwak. Learning pixel-level semantic affinity with image-level supervision for weakly supervised semantic segmentation. In *CVPR*, 2018. 9
- [2] Evangelos Alexiou, Evgeniy Upenik, and Touradj Ebrahimi. Towards subjective quality assessment of point cloud imaging in augmented reality. In *MMSP*, 2017. 1
- [3] Iro Armeni, Ozan Sener, Amir R Zamir, Helen Jiang, Ioannis Brilakis, Martin Fischer, and Silvio Savarese. 3D semantic parsing of large-scale indoor spaces. In *CVPR*, 2016. 2, 5
- [4] Amy Bearman, Olga Russakovsky, Vittorio Ferrari, and Li Fei-Fei. What’s the point: Semantic segmentation with point supervision. In *ECCV*, 2016. 6, 9
- [5] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *ICCV*, 2019. 1
- [6] Yu-Ting Chang, Qiaosong Wang, Wei-Chih Hung, Robinson Piramuthu, Yi-Hsuan Tsai, and Ming-Hsuan Yang. Weakly-supervised semantic segmentation via sub-category exploration. In *CVPR*, 2020. 5
- [7] Siheng Chen, Baoan Liu, Chen Feng, Carlos Vallespi-Gonzalez, and Carl Wellington. 3D point cloud processing and learning for autonomous driving: Impacting map creation, localization, and perception. *IEEE Signal Processing Magazine*, 2020. 1
- [8] Xiaozhi Chen, Huimin Ma, Ji Wan, Bo Li, and Tian Xia. Multi-view 3D object detection network for autonomous driving. In *CVPR*, 2017. 1
- [9] Julian Chibane, Francis Engelmann, Tuan Anh Tran, and Gerard Pons-Moll. Box2mask: Weakly supervised 3D semantic instance segmentation using bounding boxes. In *ECCV*, 2022. 1, 2
- [10] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4D spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, 2019. 1, 3, 5, 6
- [11] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. ScanNet: Richly-annotated 3D reconstructions of indoor scenes. In *ICCV*, 2017. 2, 5
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv:2010.11929*, 2020. 9
- [13] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The KITTI dataset. *IJRR*, 2013. 1
- [14] Kyle Genova, Xiaoqi Yin, Abhijit Kundu, Caroline Pantofaru, Forrester Cole, Avneesh Sud, Brian Brewington, Brian Shucker, and Thomas Funkhouser. Learning 3D semantic segmentation with only 2D image supervision. In *International Conference on 3D Vision (3DV)*, 2021. 2
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4, 5
- [16] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3D scene understanding with contrastive scene contexts. In *CVPR*, 2021. 1, 2, 5
- [17] Ji Hou, Saining Xie, Benjamin Graham, Angela Dai, and Matthias Nießner. Pri3D: Can 3D priors help 2D representation learning? In *ICCV*, 2021. 2
- [18] Qingyong Hu, Bo Yang, Guangchi Fang, Yulan Guo, Aleš Leonardis, Niki Trigoni, and Andrew Markham. SQN: Weakly-supervised semantic segmentation of large-scale 3D point clouds. In *ECCV*, 2022. 2
- [19] Wenbo Hu, Hengshuang Zhao, Li Jiang, Jiaya Jia, and Tien-Tsin Wong. Bidirectional projection network for cross dimension scene understanding. In *CVPR*, 2021. 1, 2, 5, 6, 8
- [20] Maximilian Jaritz, Jiayuan Gu, and Hao Su. Multi-view pointnet for 3D scene understanding. In *CVPR Workshop*, 2019. 1, 2
- [21] Jinwoo Kim, Jaehoon Yoo, Juho Lee, and Seunghoon Hong. Setvae: Learning hierarchical composition for generative modeling of set-structured data. In *CVPR*, 2021. 2
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2014. 5
- [23] Abhijit Kundu, Xiaoqi Yin, Alireza Fathi, David Ross, Brian Brewington, Thomas Funkhouser, and Caroline Pantofaru. Virtual multi-view fusion for 3D semantic segmentation. In *ECCV*, 2020. 1, 2, 8
- [24] Hyeokjun Kweon and Kuk-Jin Yoon. Joint learning of 2D-3D weakly supervised semantic segmentation. In *NeurIPS*, 2022. 1, 2, 5, 6, 9
- [25] Florent Lafarge and Clément Mallet. Creating large-scale city models from 3D-point clouds: A robust approach with hybrid representation. *IJCV*, 2012. 3
- [26] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *CVPR*, 2018. 1
- [27] Min Seok Lee, Seok Woo Yang, and Sung Won Han. Gaia: Graphical information gain based attention network for weakly supervised point cloud semantic segmentation. In *WACV*, 2023. 2
- [28] Mengtian Li, Yuan Xie, Yunhang Shen, Bo Ke, Ruizhi Qiao, Bo Ren, Shaohui Lin, and Lizhuang Ma. Hybridcr: Weakly-supervised 3D point cloud semantic segmentation via hybrid contrastive regularization. In *CVPR*, 2022. 2
- [29] Siqi Li, Changqing Zou, Yipeng Li, Xibin Zhao, and Yue Gao. Attention-based multi-modal fusion network for semantic scene completion. In *AAAI*, 2020. 2
- [30] Kangcheng Liu, Yuzhi Zhao, Zhi Gao, and Ben M Chen. Weaklabel3d-net: A complete framework for real-scene lidar point clouds weakly supervised multi-tasks understanding. In *ICRA*, 2022. 2
- [31] Qing Liu, Vignesh Ramanathan, Dhruv Mahajan, Alan Yuille, and Zhenheng Yang. Weakly supervised instance segmentation for videos with temporal mask consistency. In *CVPR*, 2021. 3
- [32] Yingfei Liu, Tiancai Wang, Xiangyu Zhang, and Jian Sun. PETR: Position embedding transformation for multi-view 3D object detection. In *ECCV*, 2022. 1, 2, 9
- [33] Zhengzhe Liu, Xiaojuan Qi, and Chi-Wing Fu. One thing one click: A self-training approach for weakly supervised 3D semantic segmentation. In *CVPR*, 2021. 1, 2, 3, 5, 6

- [34] John McCormac, Ankur Handa, Stefan Leutenegger, and Andrew J Davison. SceneNet RGB-D: Can 5M synthetic images beat generic ImageNet pre-training on indoor segmentation? In *ICCV*, 2017.
- [35] Youngmin Park, Vincent Lepetit, and Woontack Woo. Multiple 3D object tracking for augmented reality. In *ISAMR*, 2008. 1
- [36] Charles R Qi, Xinlei Chen, Or Litany, and Leonidas J Guibas. ImVoteNet: Boosting 3D object detection in point clouds with image votes. In *CVPR*, 2020. 1
- [37] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. PointNet: Deep learning on point sets for 3D classification and segmentation. In *CVPR*, 2017. 5
- [38] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. PointNet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017. 3, 5
- [39] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICLR*, 2021. 5
- [40] Zhongzheng Ren, Ishan Misra, Alexander G Schwing, and Rohit Girdhar. 3D spatial recognition without spatially labeled 3D. In *CVPR*, 2021. 1, 2, 4, 5, 6
- [41] Damien Robert, Bruno Vallet, and Loic Landrieu. Learning multi-view aggregation in the wild for large-scale 3D semantic segmentation. In *CVPR*, 2022. 1, 2, 5, 6
- [42] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. ImageNet large scale visual recognition challenge. *IJCV*, 2015. 5
- [43] Hanyu Shi, Jiacheng Wei, Ruibo Li, Fayao Liu, and Guosheng Lin. Weakly supervised segmentation on outdoor 4D point clouds with temporal matching and spatial graph propagation. In *CVPR*, 2022. 2
- [44] Guolei Sun, Wenguan Wang, Jifeng Dai, and Luc Van Gool. Mining cross-image semantics for weakly supervised semantic segmentation. In *ECCV*, 2020. 4, 9
- [45] Tianfang Sun, Zhizhong Zhang, Xin Tan, Yanyun Qu, Yuan Xie, and Lizhuang Ma. Image understands point cloud: Weakly supervised 3D semantic segmentation via association learning. *arXiv:2209.07774*, 2022. 2
- [46] Weixuan Sun, Jing Zhang, and Nick Barnes. 3D guided weakly supervised semantic segmentation. In *ACCV*, 2020. 2
- [47] An Tao, Yueqi Duan, Yi Wei, Jiwen Lu, and Jie Zhou. Seg-Group: Seg-level supervision for 3D instance and semantic segmentation. *TIP*, 2022. 2
- [48] Ardian Umam, Cheng-Kun Yang, Yung-Yu Chuang, Jen-Hui Chuang, and Yen-Yu Lin. Point mixSwap: Attentional point cloud mixing via swapping matched structural divisions. In *ECCV*, 2022. 2
- [49] Ozan Unal, Dengxin Dai, and Luc Van Gool. Scribble-supervised lidar semantic segmentation. In *CVPR*, 2022. 2
- [50] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 4
- [51] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E Sarma, Michael M Bronstein, and Justin M Solomon. Dynamic graph cnn for learning on point clouds. *TOG*, 2019. 5
- [52] Yikai Wang, TengQi Ye, Lele Cao, Wenbing Huang, Fuchun Sun, Fengxiang He, and Dacheng Tao. Bridged transformer for vision and point cloud 3D object detection. In *CVPR*, 2022. 1, 2
- [53] Ziyi Wang, Yongming Rao, Xumin Yu, Jie Zhou, and Jiwen Lu. SemAffiNet: Semantic-affine transformation for point cloud segmentation. In *CVPR*, 2022. 1, 2, 5, 6
- [54] Jiacheng Wei, Guosheng Lin, Kim-Hui Yap, Tzu-Yi Hung, and Lihua Xie. Multi-path region mining for weakly supervised 3D semantic segmentation on point clouds. In *CVPR*, 2020. 1, 2, 5, 6, 8
- [55] ZHENNAN WU, YANG LI, Yifei Huang, Lin Gu, Tatsuya Harada, and Hiroyuki Sato. 3D segmenter: 3D transformer based semantic segmentation via 2D panoramic distillation. In *ICLR*, 2023. 2
- [56] Zhonghua Wu, Yicheng Wu, Guosheng Lin, Jianfei Cai, and Chen Qian. Dual adaptive transformations for weakly supervised point cloud segmentation. In *ECCV*, 2022. 2
- [57] Lian Xu, Wanli Ouyang, Mohammed Bennamoun, Farid Boussaid, and Dan Xu. Multi-class token transformer for weakly supervised semantic segmentation. In *CVPR*, 2022. 2, 3, 4, 5, 6
- [58] Xun Xu and Gim Hee Lee. Weakly supervised semantic point cloud segmentation: Towards 10x fewer labels. In *CVPR*, 2020. 1, 2, 5, 8
- [59] Xu Yan, Jiantao Gao, Chaoda Zheng, Chao Zheng, Ruimao Zhang, Shuguang Cui, and Zhen Li. 2Dpass: 2D priors assisted semantic segmentation on lidar point clouds. In *ECCV*, 2022. 2
- [60] Chaolong Yang, Yuyao Yan, Weiguang Zhao, Jianan Ye, Xi Yang, Amir Hussain, and Kaizhu Huang. Towards deeper and better multi-view feature fusion for 3D semantic segmentation. *arXiv:2212.06682*, 2022. 2
- [61] Cheng-Kun Yang, Ji-Jia Wu, Kai-Syun Chen, Yung-Yu Chuang, and Yen-Yu Lin. An mil-derived transformer for weakly supervised point cloud segmentation. In *CVPR*, 2022. 1, 2, 5, 6, 9
- [62] Ze Yang and Liwei Wang. Learning relationships for multi-view 3D object recognition. In *ICCV*, 2019. 1, 2
- [63] Ping-Chung Yu, Cheng Sun, and Min Sun. Data efficient 3D learner via knowledge transferred from 2D model. In *ECCV*, 2022. 2, 5, 6, 8
- [64] Zhihao Yuan, Xu Yan, Yinghong Liao, Yao Guo, Guanbin Li, Shuguang Cui, and Zhen Li. X-trans2cap: Cross-modal knowledge transfer using transformer for 3D dense captioning. In *CVPR*, 2022. 2
- [65] Yachao Zhang, Zonghao Li, Yuan Xie, Yanyun Qu, Cuihua Li, and Tao Mei. Weakly supervised semantic segmentation for large-scale point cloud. In *AAAI*, 2021. 1, 2
- [66] Yachao Zhang, Yanyun Qu, Yuan Xie, Zonghao Li, Shanshan Zheng, and Cuihua Li. Perturbed self-distillation: Weakly supervised large-scale point cloud semantic segmentation. In *ICCV*, 2021. 2