

# Open Vocabulary 3D Scene Understanding via Geometry Guided Self-Distillation

Pengfei Wang<sup>1,2</sup>, Yuxi Wang<sup>2</sup>, Shuai Li<sup>1</sup>, Zhaoxiang Zhang<sup>1,2,3,4</sup>, Zhen Lei<sup>1,2,3,4</sup>, and Lei Zhang<sup>1</sup>

<sup>1</sup> The Hong Kong Polytechnic University

<sup>2</sup> Centre for Artificial Intelligence and Robotics, HKISI, CAS

<sup>3</sup> State Key Laboratory of Multimodal Artificial Intelligence Systems, CASIA

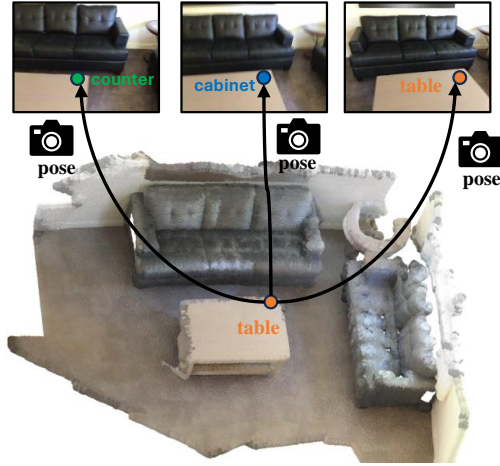
<sup>4</sup> School of Artificial Intelligence, University of Chinese Academy of Sciences (UCAS)  
pengfei.wang@connect.polyu.hk, zhaoxiang.zhang@ia.ac.cn,  
cslzhang@comp.polyu.edu.hk

**Abstract.** The scarcity of large-scale 3D-text paired data poses a great challenge on open vocabulary 3D scene understanding, and hence it is popular to leverage internet-scale 2D data and transfer their open vocabulary capabilities to 3D models through knowledge distillation. However, the existing distillation-based 3D scene understanding approaches rely on the representation capacity of 2D models, disregarding the exploration of geometric priors and inherent representational advantages offered by 3D data. In this paper, we propose an effective approach, namely Geometry Guided Self-Distillation (GGSD), to learn superior 3D representations from 2D pre-trained models. Specifically, we first design a geometry guided distillation module to distill knowledge from 2D models, and then leverage the 3D geometric priors to alleviate the inherent noise in 2D models and enhance the representation learning process. Due to the advantages of 3D representation, the performance of the distilled 3D student model can significantly surpass that of the 2D teacher model. This motivates us to further leverage the representation advantages of 3D data through self-distillation. As a result, our proposed GGSD approach outperforms the existing open vocabulary 3D scene understanding methods by a large margin, as demonstrated by our experiments on both indoor and outdoor benchmark datasets. Codes are available at <https://github.com/Wang-pengfei/GGSD>.

**Keywords:** 3D scene understanding · Open vocabulary · Self-distillation

## 1 Introduction

3D scene understanding plays a vital role in various real-world applications, including robot manipulation, virtual reality, and autonomous driving, *etc.* Traditional 3D scene understanding methods mostly employ supervised training on labeled 3D datasets with a limited number of classes [18, 42, 56]. These approaches, however, face difficulties when they encounter unknown categories in open environments. Such a limitation in practical applications motivates researchers to



**Fig. 1:** Example of the problems when utilizing 2D models, *e.g.*, LSeg [34], for pixel-to-point distillation. We can see the noise caused by the factors such as occlusion, lighting variations and view angles in 2D images. Therefore, only employing the 2D model in distillation brings errors to the 3D model, limiting its 3D representation capacity.

explore new 3D scene understanding models with open vocabulary capability. Inspired by the success of vision-language foundation models like CLIP [48], which aligns language features to image features, in various image understanding tasks [19, 28, 34, 50, 61, 67, 69], an intuitive strategy is to leverage a substantial amount of paired 3D point cloud and text data for training. Unfortunately, the high cost of annotating 3D datasets (*e.g.*, annotating a scene with 20 categories takes approximately 22.3 minutes [12]) makes it impractical to manually label enough data with numerous real-world categories.

An alternative yet more practical approach is to leverage internet-scale 2D data to acquire open vocabulary capability, and then transfer this capability to 3D scene understanding models via distillation [6, 12, 45, 62, 66]. For instance, PLA [12] utilizes a 2D caption model [58] to generate textual descriptions for images, which are then refined using GPT2 [1]. Subsequently, these descriptions are projected onto the 3D point cloud, and the vision-language knowledge can be passed to the 3D model via 3D-text contrastive learning. The recent state-of-the-art OpenScene [45] utilizes 2D open-vocabulary segmentors such as LSeg [34] and OpenSeg [17] to extract pixel-level 2D embeddings for the 3D point cloud. The ability of pixel-text alignment is then transferred to align 3D point to text via feature distillation, enabling 3D semantic-level scene understanding.

Though the above mentioned distillation-based methods have shown promising results in 3D scene understanding, they basically imitate the 2D models and inherit their limitations, restricting the upper bound of representation power. As shown in Fig. 1, the 2D model may suffer from errors due to issues such as occlusion, lighting variations, and viewing angles. The misclassification of ex-

tracted 2D features leads to erroneous supervisory signals during the distillation process of 3D point cloud learning. Instead of solely imitating 2D models, our objective is to harness both the geometric priors in 3D point clouds and the representational advantages of 3D data to guide the knowledge distillation process from pre-trained 2D models.

To achieve the above mentioned objectives, in this paper we propose an effective approach, namely Geometry Guided Self-Distillation (GGSD), to learn superior 3D representations from 2D pre-trained open vocabulary segmentors. Given a sparse point cloud composed of multiple semantic categories, geometrically homogeneous neighborhoods typically correspond to the same semantic information. Based on this observation, we first design a geometry guided distillation module to learn knowledge from 2D pre-trained models. Specifically, we generate geometrically homogeneous neighborhoods as superpoints and utilize their semantic consistency to guide the distillation process. The utilization of 3D geometric priors could mitigate the inherent noise of the 2D model and enhance the representation learning capability. Furthermore, due to the representation advantages of 3D data, which are less affected by factors such as lighting and viewing angles, the distilled 3D student model can exhibit superior performance compared to the 2D teacher model. This motivates us to further distill knowledge within the 3D network itself through self-distillation. We employ an exponential moving average (EMA) model to predict each 3D point cloud and perform voting within the superpoint to provide supervision signals for the online 3D model. This approach ensures that the online model could receive reliable and accurate guidance by leveraging the consensus among multiple-point predictions. Finally, our proposed GGSD method could learn robust 3D representations from the 3D distillation model itself, achieving significantly superior performance to the 2D teacher model.

The contributions of this work are summarized as follows. First, we analyze the limitations of pixel-point distillation methods and propose leveraging the advantages of 3D data representation through self-distillation, instead of relying solely on 2D representations. Secondly, we utilize 3D geometric priors to mitigate the noise in pixel-point distillation process, thereby enhancing the capability of representation learning. Third, we conduct experiments on both indoor and outdoor datasets, including ScanNet [10], Matterport [3], and nuScenes [2]. The results demonstrate that GGSD sets a new state-of-the-art for 3D open vocabulary scene understanding in both indoor and outdoor datasets.

## 2 Related Work

**Closed-set 3D Scene Understanding.** Previous studies in closed-set 3D scene understanding [9, 21, 24, 26, 27, 36, 43, 46, 52] have led to significant advancements across diverse 3D scene understanding benchmarks, encompassing tasks such as 3D object classification, 3D object detection, 3D semantic and instance segmentation, *etc.* Some researches [4] have explored the use of 2D pre-trained models, such as MAE [22], for pre-training point encoders. However, they mainly em-

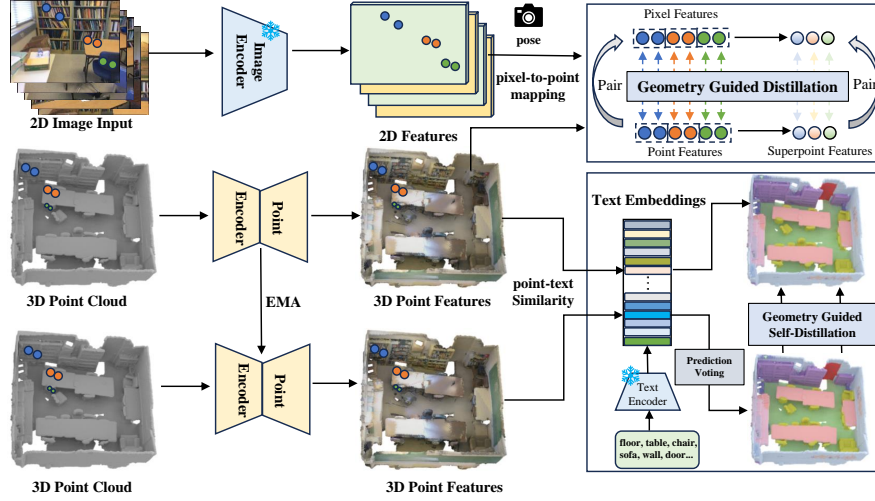
phasize on leveraging pre-training to enhance supervised 3D downstream tasks, rather than addressing open-vocabulary query scenarios. Existing closed-set 3D scene understanding methods excel in specific categories but often struggle to recognize new categories in open-set scenarios. Our objective is to develop a solution to effectively recognize and comprehend new categories in open-set environments, thus overcoming this limitation.

**Open-Vocabulary 2D Scene Understanding.** Traditional closed-set scene understanding has seen significant performance improvements [8, 14, 57, 59, 60, 65], while open-vocabulary scene understanding is also beginning to flourish. Prior studies have explored knowledge distillation from visual-language models, such as CLIP, to transfer their open vocabulary capabilities in 2D vision tasks [13, 15, 16, 20, 35, 37, 40, 51, 64, 68]. Recent research endeavors have aimed to establish correlations between dense image features and embeddings derived from language models [17, 19, 32, 34, 38, 49]. Due to their pixel-level text alignment capabilities, 2D open vocabulary scene understanding models, such as OpenSeg [17] and LSeg [34], can serve as teacher models for distillation of 3D models compared to CLIP. In this study, we rely on pre-trained 2D models with open vocabulary capabilities, specifically OpenSeg and LSeg, and propose a novel geometry-guided self-distillation strategy to conduct 3D scene-level understanding tasks without the need for any ground-truth 2D or 3D training data.

**Open-Vocabulary 3D Scene Understanding.** Recent works on open-vocabulary 3D scene understanding can be classified into two categories. The first category [25, 30, 39, 47], represented by methods like LERF [30], focuses on optimization-based approaches that utilize techniques such as NeRF [41] or 3DGS [29] to distill 2D features into 3D scenes. However, these methods lack generalization capability when confronted with new scenes. The second category [6, 12, 45, 54], represented by methods like OpenScene [45], is based on learning-based approaches. This paper falls in the second category. CLIP2Scene [6] leverages CLIP image model to extract 2D image features and utilizes textual category features to classify pixels, which are then projected onto 3D point clouds. OpenScene [45] aggregates pixel features from multiple view images on every 3D point, achieving state-of-the-art open vocabulary 3D scene understanding performance. Despite the promising results of these methods, their simple imitation of 2D models limits the representational capacity of 3D models.

### 3 Method

An overview of our approach is illustrated in Fig. 2. Firstly, we follow OpenScene [45] to create pixel-point feature pairs for the 3D point cloud (see Section 3.1). Subsequently, we analyze the limitations of pixel-point distillation methods (see Section 3.2). Then, we propose a geometry guided distillation module to enhance the learning of 3D representations from the 2D model (see Section 3.3). Finally, we introduce a geometry guided self-distillation module to further distill the knowledge within the 3D network itself (see Section 3.4).



**Fig. 2:** Overview of our Geometry Guided Self-Distillation (GGSD) framework, which consists of two main components: Geometry Guided Distillation and Self-Distillation. The first module leverages 3D geometric priors to mitigate the inherent noise in 2D models. In the second module, the 3D network learns from its own predictions. This self-distillation step allows the 3D network to continuously enhance its understanding and representation capabilities. Overall, the GGSD framework combines geometry guidance and self-distillation to fully exploit the valuable information hidden in 3D data, rather than solely relying on 2D pre-trained models.

### 3.1 Creating Pixel-Point Feature Pairs

Following OpenScene [45], our approach begins by extracting dense per-pixel embeddings from RGB images using a pre-trained 2D vision language segmentation model, such as LSeg [34] or OpenSeg [17]. After obtaining the pixel features, we can establish correspondences between each point in the point cloud and the pixels in different RGB images by utilizing the intrinsic matrix and the world-to-camera extrinsic matrix. Next, we employ a multi-view fusion strategy to combine the features from multiple views into a single feature vector, denoted as  $\mathbf{f}^{2D} = \phi(\mathbf{f}_1, \dots, \mathbf{f}_K)$ , for each point cloud  $\mathbf{p}$ . This fusion is achieved using an average pooling operator  $\phi$  over the  $K$ -view features. After obtaining the fused 2D features for each of the  $M$  points in the scene, we obtain a 2D pixel feature set  $\mathbf{F}^{2D} = \mathbf{f}_1^{2D}, \dots, \mathbf{f}_M^{2D} \in \mathbb{R}^{M \times C}$  corresponding to the point cloud scene. Finally, these embedded features are back-projected into the 3D point cloud scene. For more details, please refer to OpenScene [45].

### 3.2 Limitations of Pixel-Point 3D Distillation

Due to the text-image alignment capability of CLIP, the 2D features  $\mathbf{F}^{2D}$  can be directly utilized for language-driven open vocabulary scene understanding.

To transfer 2D vision-language knowledge to a 3D point cloud network, the distillation approach has been popularly used. One up-to-date pixel-point 3D distillation work is OpenScene [45]. Specifically, OpenScene aims to learn a 3D encoder that extracts 3D features from each point cloud, which are expected to be similar to their corresponding 2D pixel features. This allows the 3D encoder to distill the open vocabulary capability of the 2D model. To this end, OpenScene utilizes cosine distance to perform pixel-point distillation:

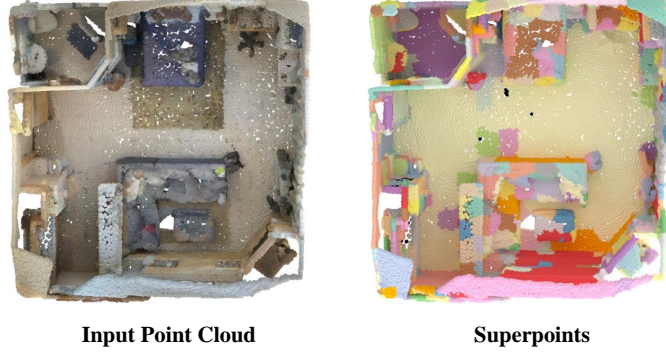
$$\mathcal{L}_p = 1 - \cos(\mathbf{F}^{2D}, \mathbf{F}^{3D}). \quad (1)$$

However, the current pixel-to-point distillation methods are mostly a simplistic imitation of the 2D models, which are limited in extracting useful information from 3D point data. Firstly, relying solely on the projection relationship between point clouds and pixels, these approaches fail to fully exploit the geometric priors inherent in 3D data, leading to suboptimal performance. In this paper, we propose to incorporate 3D geometric priors throughout the entire distillation process to address the issue of noise in 2D models by constraining semantic consistency. Secondly, the current pixel-point distillation methods fail to fully exploit the representational advantage offered by 3D point clouds. It has been observed that the distilled 3D model can outperform the directly projected 2D models. This can be attributed to the inherent capability of 3D data to handle challenges like lighting variations, occlusion, and viewpoint changes, which are prevalent in 2D representations. Therefore, the accuracy of 3D model segmentation results, particularly for foreground classes, can be significantly improved. This observation provides a strong motivation for us to further leverage the advantages of 3D representation within the 3D network itself via self-distillation, rather than relying solely on information from the 2D model.

### 3.3 Geometry Guided Distillation

We first design a geometry guided distillation method to learn knowledge from 2D models. Given a sparse point cloud composed of multiple semantic categories, we can decompose it into many simple parts, *i.e.*, superpoints, according to their geometric position in the space. As shown in Fig. 3, the superpoints for an indoor room are constructed by VCCS [44]. The details for creating superpoints can be found in the **supplementary file**. For a specific input point cloud  $\mathbf{P} \in \mathbb{R}^{M \times 3}$ , its superpoints are denoted as  $\{\tilde{\mathbf{p}}_1 \cdots \tilde{\mathbf{p}}_n \cdots \tilde{\mathbf{p}}_N\}$ , where each superpoint  $\tilde{\mathbf{p}}_n$  consists of a small subset of original point cloud  $\mathbf{p}$ . These superpoints are geometrically simple and usually cannot cover objects of different classes. Ideally, each superpoint as a whole should belong to the same semantic category. The simplicity and coherence of superpoints naturally facilitate the implementation of semantic consistency constraints.

To mitigate the noise issue in 2D models, we can leverage the semantic consistency provided by 3D geometric priors. In particular, for a specific input point cloud  $\mathbf{P}$ , we have its neural features  $\mathbf{F}^{3D} = \{\mathbf{f}_1^{3D}, \dots, \mathbf{f}_M^{3D}\}$ , the corresponding pixel fusion features  $\mathbf{F}^{2D} = \{\mathbf{f}_1^{2D}, \dots, \mathbf{f}_M^{2D}\} \in \mathbb{R}^{M \times C}$  and superpoints



**Fig. 3:** Example of the superpoints generated using VCCS [44] on ScanNet scene point clouds. Each colored patch corresponds to a distinct superpoint, showcasing a simple geometric structure. Additionally, the points within each superpoint consistently exhibit a shared semantic category.

$\{\tilde{\mathbf{p}}_1 \cdots \tilde{\mathbf{p}}_n \cdots \tilde{\mathbf{p}}_N\}$ . Firstly, we compute the mean 2D pixel features for superpoints, denoted as  $\mathbf{F}_{sp}^{2D} = \{\tilde{\mathbf{f}}_1^{2D}, \dots, \tilde{\mathbf{f}}_N^{2D}\}$ :

$$\tilde{\mathbf{f}}_n^{2D} = \frac{1}{Q} \sum_{q=1}^Q \mathbf{f}_q^{2D}, \quad \mathbf{f}_q^{2D} \in \mathbb{R}^{1 \times C}, \quad (2)$$

where  $Q$  is the total number of 3D points within a superpoint  $\tilde{\mathbf{p}}_n$ , and  $\mathbf{f}_q^{2D}$  is the feature vector for the  $q^{th}$  pixel that corresponds to the 3D point in the superpoint. Similarly, we can employ the same approach to calculate the mean 3D point features  $\mathbf{F}_{sp}^{3D} = \{\tilde{\mathbf{f}}_1^{3D}, \dots, \tilde{\mathbf{f}}_N^{3D}\}$ .

To enforce the superpoint feature  $\mathbf{F}_{sp}^{3D}$  to be consistent with the 2D fused features  $\mathbf{F}_{sp}^{2D}$ , we use a cosine similarity loss as that in pixel-point distillation:

$$\mathcal{L}_{sp} = 1 - \cos(\mathbf{F}_{sp}^{2D}, \mathbf{F}_{sp}^{3D}). \quad (3)$$

The final optimization objective for geometry guided distillation is as follows:

$$\mathcal{L}_d = \mathcal{L}_p + \mathcal{L}_{sp}. \quad (4)$$

### 3.4 Geometry Guided Self-Distillation

With geometry guided distillation, the performance of the 3D model can surpass that of the directly projected 2D model. This inspires us to further design a self-distillation module to fully leverage the inherent advantages of 3D data in representation. Recent methods such as CNS [5] and CLIP2Scene [6] typically transfer the pixel-level pseudo-labels from the 2D image to the 3D point cloud using a 2D-3D calibration matrix. In contrast to existing methods that rely on projecting from 2D, we introduce a geometry guided self-distillation approach to distill knowledge from the 3D model itself. Specifically, we generate predictions

directly on 3D point clouds instead of projecting 2D predictions back to 3D, as our 3D model possesses enhanced feature representation capabilities. Specifically, given the 3D point features  $\mathbf{F}^{3D} = \{\mathbf{f}_1^{3D}, \dots, \mathbf{f}_M^{3D}\} \in \mathbb{R}^{M \times C}$ , we assign pseudo-labels to each point  $\mathbf{f}_n^{3D}$  as follows:

$$\mathbf{f}^{\hat{i}} = \arg \max_l \psi(\mathbf{f}_n^{3D}, \mathbf{f}_l^t), \quad (5)$$

where  $\mathbf{f}_l^t$  is CLIP’s text embedding of the  $l$ -th class name with a pre-defined template.  $\psi$  is the dot product operation. Furthermore, in the self-distillation stage, we also incorporate 3D geometric priors to enhance the semantic consistency of the predictions.

However, the research findings from self-supervised methods [7, 23] have highlighted that directly using the network’s predictions as supervision signals can lead to model collapse due to overconfidence. To address this issue, we leverage the predictions from the exponential moving average (EMA) model as supervision signals to train the 3D point cloud network, aiming to enhance stability and mitigate the impact of overconfidence. Specifically, we use the features extracted by the EMA model to calculate the similarity with the text feature according to Equation 5, resulting in generating a point-text pair for each point cloud.

Although using EMA model can prevent mode collapse, it is still inevitable to have noise in the prediction results. To mitigate this issue, we incorporate the semantic consistency of 3D geometric priors during the self-distillation process. Specifically, we implement a voting mechanism within each superpoint to refine the prediction results. After assigning the most similar textual feature to each point cloud according to Equation 5, we calculate the proportion of textual features within each superpoint. We then modify the text corresponding to all point clouds within that superpoint to match the textual feature with the highest proportion. This correction process effectively mitigates the impact of noise and improves the accuracy of prediction. The voting outcome generates point-language pairs for each point cloud, then we can pull paired point and language features together while pushing unmatched features away through contrastive loss. The geometry guided self-distillation loss is as follows:

$$\mathcal{L}_{sd} = -\log \frac{\exp(\mathbf{f}^{3D} \cdot \mathbf{f}^{\hat{i}}/\tau)}{\sum_{i=1}^{n_t} \exp(\mathbf{f}^{3D} \cdot \mathbf{f}_i^t/\tau)}, \quad (6)$$

where  $\mathbf{f}^{\hat{i}}$  represents the textual features corresponding to the point cloud  $\mathbf{f}^{3D}$ ;  $n_t$  is the number of point-language pairs and  $\tau$  is a temperature factor.

### 3.5 Training and Inference

**Training.** The training process comprises two stages. In the initial stage, we employ geometry guided distillation to train the 3D network for 70 epochs, enabling it to acquire knowledge from the 2D model. Then, we incorporate the GGSD module into an additional 30 training epochs. The total number of training epochs is consistent with OpenScene.

**Inference.** Using any per-point feature  $\mathbf{f}^{3D}$  and CLIP text features obtained from a set of text prompts, we can estimate their similarities by calculating the cosine similarity score between them. This similarity score is utilized in all of our scene understanding tasks. In all experiments, unless otherwise stated, we exclusively use pure 3D point clouds for predictions during the inference stage, without employing the 2D-3D ensemble strategy.

## 4 Experiments

### 4.1 Experiment Setup

**Datasets.** To evaluate the effectiveness of our method, we conduct experiments on three widely used public benchmarks: ScanNet [10], Matterport [3], and nuScenes [2]. The first two datasets provide RGBD images and 3D meshes of indoor scenes, while the last one offers LIDAR scans of outdoor scenes. Additionally, Matterport3D is a challenging dataset with intricately detailed scenes, which allows us to thoroughly examine the performance of open-vocabulary queries across different datasets.

**Implementation Details.** We adopt the settings outlined in OpenScene [45], utilize MinkowskiNet18A [9] as the 3D backbone network, and employ a simple prompt engineering technique to extract CLIP text features. Specifically, we modify the text prompts for each object class “XX” to “a XX in a scene”, such as “a chair in a scene”. On the ScanNet and Matterport3D datasets, we employ a voxel size of 2cm with MinkowskiNet while utilizing a voxel size of 5cm on the nuScenes dataset. Based on the findings of OpenScene, we employ LSeg [34] for indoor datasets and OpenSeg [17] for outdoor dataset as the 2D teacher model. The temperature  $\tau$  in Eq. 6 is empirically set as 0.01. We utilize the Adam optimizer with an initial learning rate of  $1e-4$  for the entire training process. We use a batch size of 8 on ScanNet and Matterport3D with a single NVIDIA A100 GPU (80G). On nuScenes, we use a batch size of 16 with 4 A100 GPUs. It takes around 24 hours to train and 0.1 second for inference. More details can be found in the **supplementary file**.

### 4.2 Main Results

We conducted a series of experiments, including traditional closed-set scenarios, in-domain open vocabulary scenarios, and out-of-domain open vocabulary scenarios to test the effectiveness of the proposed methods in various 3D scene understanding tasks.

**Comparison on Standard Benchmarks.** We compare our approach with both fully-supervised and open vocabulary methods on all classes of the ScanNet [10] validation set and nuScenes [2] validation set. As shown in the bottom panel of Table 1, our approach, GGSD, consistently outperforms recent open vocabulary methods PLA [12], CLIP2Scene [6], CNS [5] and CLIP-FO3D [66] by a large margin on both datasets. Additionally, our method exhibits significant performance improvements compared to the baseline model OpenScene [45].

	ScanNet [10]		nuScenes [2]	
	mIoU	mAcc	mIoU	mAcc
<i>Fully-supervised methods</i>				
TangentConv [55]	40.9	-	-	-
TextureNet [27]	54.8	-	-	-
ScanComplete [11]	56.6	-	-	-
DCM-Net [53]	65.8	-	-	-
Mix3D [43]	73.6	-	-	-
VMNet [26]	73.2	-	-	-
LidarMultiNet [63]	-	-	82.0	-
MinkowskiNet [9]	69.0	77.5	78.0	83.7
<i>Open vocabulary methods</i>				
MSeg [33] Voting	45.6	54.4	31.0	36.9
PLA [12]	17.7	33.5	-	-
CLIP2Scene [6]	25.1	-	20.8	-
CNS [5]	26.8	-	33.5	-
CLIP-FO3D [66]	30.2	49.1	-	-
OpenScene <sup>†</sup> [45]	54.2	66.6	42.1	61.8
OpenScene <sup>‡</sup> [45]	52.9	63.2	42.9	57.1
<b>GGSD (Ours)</b>	<b>56.5</b>	<b>68.6</b>	<b>46.1</b>	<b>59.2</b>

**Table 1: Comparisons on Standard Benchmarks.** The numbers in gray are obtained from previous papers. The symbol <sup>†</sup> indicates the use of the 2D-3D ensemble strategy, while <sup>‡</sup> denotes the pure 3D results obtained from the official open-source model. Our method surpasses the current state-of-the-art by a large margin in both indoor and outdoor scenarios without using the 2D-3D ensemble strategy.

Specifically, on the ScanNet dataset, GGSD achieves a 3.6% increase in mIoU and a 5.4% increase in mACC. On the nuScenes dataset, GGSD demonstrates a 3.2% increase in mIoU and a 2.1% increase in mAcc. Except for OpenScene, the other methods in Table 1 do not employ the 2D-3D ensemble strategy due to the significant computational and storage costs associated with maintaining a 2D segmentation network and processing a large number of images per scene during testing. It is worth noting that our method using only point clouds still outperforms OpenScene with the ensemble strategy. Given that our method is evaluated on both indoor and outdoor scenes, and it even surpasses OpenScene with ensemble strategy, these results clearly demonstrate the effectiveness and superiority of GGSD.

**Comparison on Open-Vocabulary Scenarios.** The standard ScanNet benchmark only contains a small vocabulary of 20 classes. To evaluate the open vocabulary capability of GGSD, we expand the original vocabulary size by using the NYU-40 label set. We remove NYU-40 labels that lack specific semantics, such as “other structure”, “other furniture”, and “other prop”. The remaining categories are evenly divided into *Head*, *Common*, and *Tail* based on the number of samples in each category (see **supplementary file** for specific division). We present the quantitative results on these three categories in Table 2.

	Head		Common		Tail	
	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc
<i>Inference by feature projection</i>						
MaskCLIP-3D	19.8	33.1	13.3	26.8	7.5	15.3
OpenScene <sup>‡</sup> [45]	64.4	76.8	27.8	35.1	14.9	21.4
<b>GGSD (Ours)</b>	<b>66.8</b>	<b>78.8</b>	<b>30.1</b>	<b>40.9</b>	<b>16.0</b>	<b>24.2</b>

**Table 2: Comparison on open-vocabulary scenarios.** Open-vocabulary semantic segmentation on ScanNet with extended labels from the NYU label set. Categories are divided into *Head*, *Common* and *Tail* by point numbers. The numbers in gray are obtained from previous papers. The symbol <sup>‡</sup> denotes the pure 3D point results obtained from the official open-source model.

	Matterport21		Matterport40		Matterport80		Matterport160	
	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc	mIoU	mAcc
OpenScene <sup>‡</sup> [45]	36.0	48.0	21.1	27.5	10.8	13.9	6.0	8.1
<b>GGSD (Ours)</b>	<b>40.1</b>	<b>54.4</b>	<b>22.8</b>	<b>31.6</b>	<b>11.9</b>	<b>16.2</b>	<b>6.3</b>	<b>9.6</b>

**Table 3: Comparison on cross-domain scenarios.** Both GGSD and OpenScene are trained on ScanNet, and zero-shot tested on the Matterport3D dataset. The symbol <sup>‡</sup> denotes the pure 3D results obtained from the official open-source model.  $K = 21$  is derived from the original Matterport3D benchmark. For  $K = 40, 80, 160$ , we use the  $K$  most common categories from the NYU label set provided in the benchmark.

GGSD achieves superior results over OpenScene in all categories, with a specific improvement of 2.4% and 2.3% in mIOU for the *Head* and *Common* categories, respectively. Since the categories in the *Common* set typically consist of objects with smaller sizes (*e.g.*, night stand and floor mat), the 2.3% mIOU improvement demonstrates the stronger open vocabulary capability of our method. However, for the *Tail* category, although GGSD surpasses OpenScene by 2.8% in mACC, the performance of both OpenScene and GGSD is not satisfactory. This can be attributed to the small sizes and the limited number of samples in these categories (*e.g.*, lamp, box, bag, and book).

**Comparison on Cross-Domain Scenarios.** Our GGSD method has demonstrated excellent potential in solving open vocabulary scene understanding tasks within a specific domain. However, the exploration of open vocabulary learners that can transfer across datasets is also worth investigation, as they face the dual challenges of category and data distribution. We conduct zero-shot domain transfer experiments by training the model on ScanNet [10] and testing it on the Matterport3D [3] dataset without fine-tuning. It should be noted that as  $K$  increases, the test set will contain more small-sized objects. As shown in Table 3, our GGSD method consistently outperforms OpenScene in Matterport21, Matterport40, Matterport80, and Matterport160, demonstrating its strong open vocabulary capability and generalization ability.

	ScanNet [10]		Matterport3D [3]	
	mIoU	mAcc	mIoU	mAcc
2D Fusion Projection	50.0	62.7	32.3	40.0
Pixel-Point Distillation	52.9	63.2	36.1	48.0
Geometry Guided Distillation	53.5	65.0	36.7	49.3
Self-Distillation	56.1	68.2	39.0	53.3
Geometry Guided Self-distillation	56.5	68.6	40.1	54.4

**Table 4: Ablation study of GGSD.** Comparison of semantic segmentation performance of each component of our GGSD method.

	ScanNet [10]		Matterport3D [3]	
	mIoU	mAcc	mIoU	mAcc
Geometry Guided Distillation	54.0	65.0	36.7	49.3
Fixed 2D model	51.9	64.6	34.2	45.7
Fixed 3D model	55.0	65.8	37.3	49.8
EMA model	56.1	68.2	39.0	53.3

**Table 5:** Comparisons of the 3D supervision signals provided by different models.

### 4.3 Ablation Studies and Analysis

In this section, we investigate our key components through an in-depth ablation study. The experiments are conducted on the ScanNet dataset and subsequently tested on both ScanNet and Matterport3D datasets.

**Component Analysis.** As shown in Table 4, even without utilizing 3D geometric priors, pixel-point distillation can outperform 2D direct projection, demonstrating the advantage of 3D representation. By incorporating 3D geometric priors, the geometry guided distillation method further enhances performance, resulting in a 0.6% increase in mIoU and a 1.8% increase in mAcc on the ScanNet dataset. Through self-distillation, we achieve a notable enhancement in performance. Specifically, on the ScanNet dataset, there is a substantial increase of 2.6% in mIoU and 3.2% in mAcc. On the Matterport3D dataset, the mIoU is increased by 2.3%, and the mAcc is improved by 4.0%. Finally, the utilization of the voting mechanism within the superpoints leads to improved performance by mitigating the impact of noise through the refinement model’s predictions. It is noteworthy that our GGSD method demonstrates substantial improvements compared to the 2D teacher model, *i.e.*, 2D Fusion Projection, achieving a +6.5% mIoU and +5.9% mAcc on the ScanNet dataset, as well as a +7.8% mIoU and +14.4% mAcc on the Matterport3D dataset. These results provide strong evidence for the effectiveness of our method.

**The Effectiveness of the EMA Model.** In the second stage, we employ the EMA model to provide supervised signals for self-distillation. In Table 5, we investigate the performance of different models in providing supervised signals. When using the fixed 2D model for pixel prediction and projecting it onto

	ScanNet [10]		Matterport3D [3]	
	mIoU	mAcc	mIoU	mAcc
Pixel-Point Distillation	52.9	63.2	36.1	48.0
SAM [31] Refine	53.0	63.3	35.9	48.4
Geometry Guided Distillation	53.5	65.0	36.7	49.3

**Table 6:** Comparisons between the utilization of geometric priors and SAM [31].

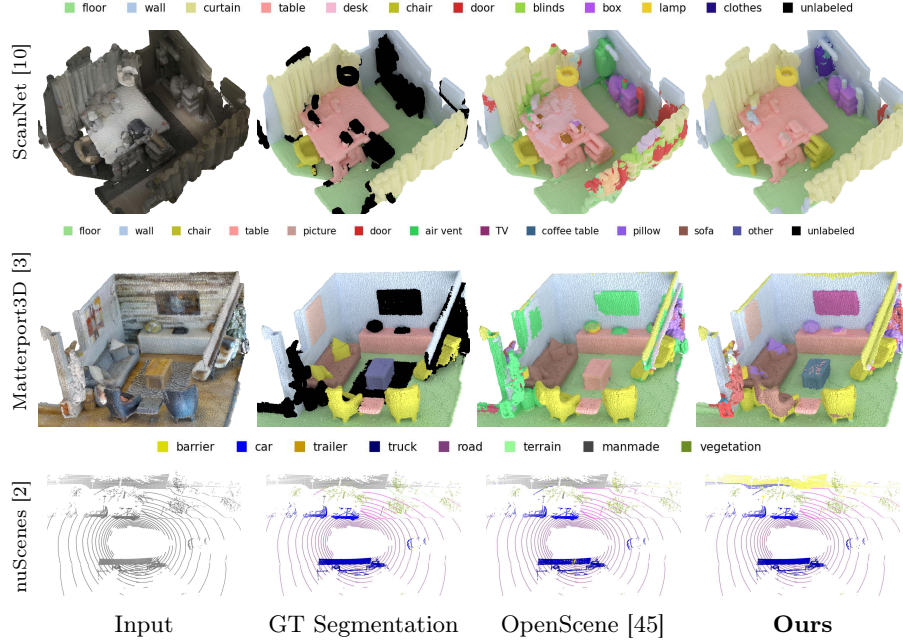
points, we observe lower performance compared to geometry-guided distillation. This can be attributed to the fact that the 2D model provides more erroneous supervised signals. Furthermore, using the fixed first-stage distilled 3D model only results in marginal performance improvement. However, when utilizing the EMA model, we achieve the best performance, primarily because the EMA model gradually obtains more accurate prediction results as training progresses.

**The Effectiveness of 3D Geometric Priors.** The recent work CNS [5] utilized the foundation model SAM [31] to refine pseudo-labels from the 2D model, achieving performance improvements on a lower baseline (27.3% mIoU in ScanNet). We also apply this approach to LSeg. Specifically, we follow the official implementation of SAM to generate masks for all images and perform average pooling within each mask to refine them. Finally, we utilize the extracted refined features for distillation. However, as shown in Table 6, this approach yields minimal improvements. Conversely, leveraging 3D geometric priors results in a 0.6% mIoU improvement in ScanNet.

**Qualitative Results.** As shown in Fig. 4, our method successfully segments all categories that exist in closed-set ScanNet benchmarks, such as wall and table. Additionally, our method successfully segments some categories that are not annotated in traditional benchmarks, such as box, lamp, and clothes, demonstrating our model’s excellent open vocabulary capability. On the Matterport3D, the GGSD method successfully segments unlabeled television and also identifies mislabeled pillow and coffee table in the ground truth. Furthermore, there have been instances of segmentation failures in Matterport3D. Due to linguistic ambiguity, sofa chairs are sometimes segmented as separate entities, such as sofas and chairs, while coffee tables can be segmented as coffee tables and tables. The misclassification of certain pillow points as sofas may be attributed to their strong color resemblance to sofas. In the case of nuScenes, GGSD also successfully segments unannotated barrier. These qualitative results obtained from GGSD demonstrate its excellent open vocabulary capability. More visualization results can be found in the **supplementary file**.

## 5 Conclusion

In this study, we proposed an effective approach called Geometry Guided Self-Distillation (GGSD) to learn superior 3D representations from 2D pre-trained



**Fig. 4: Qualitative results.** We present the qualitative results of 3D semantic segmentation on public indoor and outdoor benchmarks.

models. We first introduced a geometry guided distillation module that leveraged 3D geometric priors to alleviate noise in 2D models. Moreover, we harnessed the representation advantages of 3D data through geometry guided self-distillation, resulting in significant improvement in the model performance for 3D scene understanding. Our GGSD method significantly outperformed existing open-vocabulary 3D scene understanding methods on both indoor and outdoor datasets. This highlighted the crucial role of geometric priors and of 3D data representation in distillation-based approaches for 3D scene understanding.

**Limitations.** Though GGSD has achieved promising results in 3D open vocabulary scene understanding tasks, it has two main limitations. Firstly, while geometric priors can significantly improve the model performance on objects with pronounced geometric structures, the improvement can be limited on objects with fewer prominent geometric characteristics. Secondly, similar to other open vocabulary methods, our approach is susceptible to the influence of language ambiguity, as illustrated by the qualitative results in Fig. 4.

**Acknowledgement.** This work was supported by the InnoHK program.

## References

1. Vit-gpt2 image captioning. <https://huggingface.co/nlpconnect/vit-gpt2-image-captioning/discussions>
2. Caesar, H., Bankiti, V., Lang, A.H., Vora, S., Liong, V.E., Xu, Q., Krishnan, A., Pan, Y., Baldan, G., Beijbom, O.: nuscenes: A multimodal dataset for autonomous driving. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 11621–11631 (2020)
3. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3d: Learning from rgb-d data in indoor environments. arXiv preprint arXiv:1709.06158 (2017)
4. Chen, A., Zhang, K., Zhang, R., Wang, Z., Lu, Y., Guo, Y., Zhang, S.: Pimae: Point cloud and image interactive masked autoencoders for 3d object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5291–5301 (2023)
5. Chen, R., Liu, Y., Kong, L., Chen, N., Xinge, Z., Ma, Y., Liu, T., Wang, W.: Towards label-free scene understanding by vision foundation models. In: Thirty-seventh Conference on Neural Information Processing Systems (2023)
6. Chen, R., Liu, Y., Kong, L., Zhu, X., Ma, Y., Li, Y., Hou, Y., Qiao, Y., Wang, W.: Clip2scene: Towards label-efficient 3d scene understanding by clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7020–7030 (2023)
7. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
8. Cheng, B., Misra, I., Schwing, A.G., Kirillov, A., Girdhar, R.: Masked-attention mask transformer for universal image segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 1290–1299 (2022)
9. Choy, C., Gwak, J., Savarese, S.: 4d spatio-temporal convnets: Minkowski convolutional neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 3075–3084 (2019)
10. Dai, A., Chang, A.X., Savva, M., Halber, M., Funkhouser, T., Nießner, M.: Scannet: Richly-annotated 3d reconstructions of indoor scenes. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5828–5839 (2017)
11. Dai, A., Ritchie, D., Bokeloh, M., Reed, S., Sturm, J., Nießner, M.: Scancomplete: Large-scale scene completion and semantic segmentation for 3d scans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4578–4587 (2018)
12. Ding, R., Yang, J., Xue, C., Zhang, W., Bai, S., Qi, X.: Pla: Language-driven open-vocabulary 3d scene understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7010–7019 (2023)
13. Du, Y., Wei, F., Zhang, Z., Shi, M., Gao, Y., Li, G.C.: Learning to prompt for open-vocabulary object detection with vision-language model. In: CVPR (2022)
14. Fan, J., Wang, Y., Guan, H., Song, C., Zhang, Z.: Toward few-shot domain adaptation with perturbation-invariant representation and transferable prototypes. Frontiers of Computer Science **16**(3), 163347 (2022)
15. Feng, C., Zhong, Y., Jie, Z., Chu, X., Ren, H., Wei, X., Xie, W., Ma, L.: Promptdet: Towards open-vocabulary detection using uncurated images. In: ECCV (2022)

16. Ghiasi, G., Gu, X., Cui, Y., Lin, T.Y.: Open-vocabulary image segmentation. In: ECCV (2022)
17. Ghiasi, G., Gu, X., Cui, Y., Lin, T.Y.: Scaling open-vocabulary image segmentation with image-level labels. In: European Conference on Computer Vision. pp. 540–557. Springer (2022)
18. Graham, B., Engelcke, M., Van Der Maaten, L.: 3d semantic segmentation with submanifold sparse convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9224–9232 (2018)
19. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. arXiv preprint arXiv:2104.13921 (2021)
20. Gu, X., Lin, T.Y., Kuo, W., Cui, Y.: Open-vocabulary object detection via vision and language knowledge distillation. In: ICLR (2021)
21. Han, L., Zheng, T., Xu, L., Fang, L.: Occuseg: Occupancy-aware 3d instance segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2940–2949 (2020)
22. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)
23. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
24. Hu, W., Zhao, H., Jiang, L., Jia, J., Wong, T.T.: Bidirectional projection network for cross dimension scene understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14373–14382 (2021)
25. Hu, X., Wang, Y., Fan, L., Fan, J., Peng, J., Lei, Z., Li, Q., Zhang, Z.: Semantic anything in 3d gaussians. arXiv preprint arXiv:2401.17857 (2024)
26. Hu, Z., Bai, X., Shang, J., Zhang, R., Dong, J., Wang, X., Sun, G., Fu, H., Tai, C.L.: Vmnet: Voxel-mesh network for geodesic-aware 3d semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15488–15498 (2021)
27. Huang, J., Zhang, H., Yi, L., Funkhouser, T., Nießner, M., Guibas, L.J.: Texturenet: Consistent local parametrizations for learning from high-resolution signals on meshes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4440–4449 (2019)
28. Huang, T., Dong, B., Yang, Y., Huang, X., Lau, R.W., Ouyang, W., Zuo, W.: Clip2point: Transfer clip to point cloud classification with image-depth pre-training. arXiv preprint arXiv:2210.01055 (2022)
29. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics **42**(4) (2023)
30. Kerr, J., Kim, C.M., Goldberg, K., Kanazawa, A., Tancik, M.: Lerf: Language embedded radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19729–19739 (2023)
31. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. arXiv preprint arXiv:2304.02643 (2023)
32. Kuo, W., Cui, Y., Gu, X., Piergiovanni, A., Angelova, A.: F-vlm: Open-vocabulary object detection upon frozen vision and language models. arXiv preprint arXiv:2209.15639 (2022)
33. Lambert, J., Liu, Z., Sener, O., Hays, J., Koltun, V.: Mseg: A composite dataset for multi-domain semantic segmentation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2879–2888 (2020)

34. Li, B., Weinberger, K.Q., Belongie, S., Koltun, V., Ranftl, R.: Language-driven semantic segmentation. In: International Conference on Learning Representations (2022), <https://openreview.net/forum?id=RriDjddCLN>
35. Li, B., Weinberger, K.Q., Belongie, S.J., Koltun, V., Ranftl, R.: Language-driven semantic segmentation. In: ICLR (2022)
36. Li, J., He, X., Wen, Y., Gao, Y., Cheng, X., Zhang, D.: Panoptic-phnet: Towards real-time and high-precision lidar panoptic segmentation via clustering pseudo heatmap. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11809–11818 (2022)
37. Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted clip. arXiv **abs/2210.04150** (2022)
38. Liang, F., Wu, B., Dai, X., Li, K., Zhao, Y., Zhang, H., Zhang, P., Vajda, P., Marculescu, D.: Open-vocabulary semantic segmentation with mask-adapted clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7061–7070 (2023)
39. Liu, K., Zhan, F., Zhang, J., Xu, M., Yu, Y., El Saddik, A., Theobalt, C., Xing, E., Lu, S.: Weakly supervised 3d open-vocabulary segmentation. *Advances in Neural Information Processing Systems* **36** (2024)
40. Lüddecke, T., Ecker, A.S.: Image segmentation using text and image prompts. In: CVPR (2021)
41. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
42. Misra, I., Girdhar, R., Joulin, A.: An End-to-End Transformer Model for 3D Object Detection. In: ICCV (2021)
43. Nekrasov, A., Schult, J., Litany, O., Leibe, B., Engelmann, F.: Mix3d: Out-of-context data augmentation for 3d scenes. In: 2021 International Conference on 3D Vision (3DV). pp. 116–125. IEEE (2021)
44. Papon, J., Abramov, A., Schoeler, M., Worgotter, F.: Voxel cloud connectivity segmentation-supervoxels for point clouds. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2027–2034 (2013)
45. Peng, S., Genova, K., Jiang, C., Tagliasacchi, A., Pollefeys, M., Funkhouser, T., et al.: Openscene: 3d scene understanding with open vocabularies. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 815–824 (2023)
46. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems* **30** (2017)
47. Qin, M., Li, W., Zhou, J., Wang, H., Pfister, H.: Langsplat: 3d language gaussian splatting. arXiv preprint arXiv:2312.16084 (2023)
48. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
49. Rao, Y., Zhao, W., Chen, G., Tang, Y., Zhu, Z., Huang, G., Zhou, J., Lu, J.: Dense-clip: Language-guided dense prediction with context-aware prompting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18082–18091 (2022)

50. Rasheed, H., Maaz, M., Khattak, M.U., Khan, S., Khan, F.S.: Bridging the gap between object and image-level representations for open-vocabulary detection. In: 36th Conference on Neural Information Processing Systems (NIPS) (2022)
51. Rasheed, H., Maaz, M., Khattak, M.U., Khan, S., Khan, F.S.: Bridging the gap between object and image-level representations for open-vocabulary detection. In: NeurIPS (2022)
52. Robert, D., Vallet, B., Landrieu, L.: Learning multi-view aggregation in the wild for large-scale 3d semantic segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5575–5584 (2022)
53. Schult, J., Engelmann, F., Kontogianni, T., Leibe, B.: Dualconvmesh-net: Joint geodesic and euclidean convolutions on 3d meshes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8612–8622 (2020)
54. Takmaz, A., Fedele, E., Sumner, R.W., Pollefeys, M., Tombari, F., Engelmann, F.: Openmask3d: Open-vocabulary 3d instance segmentation. arXiv preprint arXiv:2306.13631 (2023)
55. Tatarchenko, M., Park, J., Koltun, V., Zhou, Q.Y.: Tangent convolutions for dense prediction in 3d. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3887–3896 (2018)
56. Vu, T., Kim, K., Luu, T.M., Nguyen, X.T., Yoo, C.D.: Softgroup for 3d instance segmentation on 3d point clouds. In: CVPR (2022)
57. Wang, H., Shen, Y., Fei, J., Li, W., Wu, L., Wang, Y., Zhang, Z.: Pulling target to source: A new perspective on domain adaptive semantic segmentation. arXiv preprint arXiv:2305.13752 (2023)
58. Wang, P., Yang, A., Men, R., Lin, J., Bai, S., Li, Z., Ma, J., Zhou, C., Zhou, J., Yang, H.: Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. CoRR **abs/2202.03052** (2022)
59. Wang, Y., Liang, J., Xiao, J., Mei, S., Yang, Y., Zhang, Z.: Informative data mining for one-shot cross-domain semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1064–1074 (2023)
60. Wang, Y., Peng, J., Zhang, Z.: Uncertainty-aware pseudo label refinery for domain adaptive semantic segmentation. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 9092–9101 (2021)
61. Xu, M., Zhang, Z., Wei, F., Lin, Y., Cao, Y., Hu, H., Bai, X.: A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. arXiv preprint arXiv:2112.14757 (2021)
62. Yang, J., Ding, R., Wang, Z., Qi, X.: Regionplc: Regional point-language contrastive learning for open-world 3d scene understanding. arXiv preprint arXiv:2304.00962 (2023)
63. Ye, D., Zhou, Z., Chen, W., Xie, Y., Wang, Y., Wang, P., Foroosh, H.: Lidarmulti-net: Towards a unified multi-task network for lidar perception. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 3231–3240 (2023)
64. Zang, Y., Li, W., Zhou, K., Huang, C., Loy, C.C.: Open-vocabulary detr with conditional matching. In: ECCV (2022)
65. Zhang, H., Li, F., Xu, H., Huang, S., Liu, S., Ni, L.M., Zhang, L.: Mp-former: Mask-piloted transformer for image segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18074–18083 (2023)
66. Zhang, J., Dong, R., Ma, K.: Clip-fo3d: Learning free open-world 3d scene representations from 2d dense clip. arXiv preprint arXiv:2303.04748 (2023)

- 67. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point cloud understanding by clip. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8552–8562 (2022)
- 68. Zhong, Y., Yang, J., Zhang, P., Li, C., Codella, N.C.F., Li, L.H., Zhou, L., Dai, X., Yuan, L., Li, Y., Gao, J.: Regionclip: Region-based language-image pretraining. In: CVPR (2021)
- 69. Zhou, C., Loy, C.C., Dai, B.: Extract free dense labels from clip. In: European Conference on Computer Vision (ECCV) (2022)