

UniM-OV3D: Uni-Modality Open-Vocabulary 3D Scene Understanding with Fine-Grained Feature Representation

Qingdong He¹, Jinlong Peng¹, Zhengkai Jiang¹, Kai Wu¹, Xiaozhong Ji¹,
Jiangning Zhang¹, Yabiao Wang¹, Chengjie Wang¹, Mingang Chen², Yunsheng Wu¹

¹YouTu Lab, Tencent

²Shanghai Development Center of Computer Software Technology

{yingcaihe, jeromepeng, zhengkaijiang, lloydwu, xiaozhongji, vtzhang, caseywang,
jasoncjwang, simonwu}@tencent.com, cmg@sscenter.sh.cn

Abstract

3D open-vocabulary scene understanding aims to recognize arbitrary novel categories beyond the base label space. However, existing works not only fail to fully utilize all the available modal information in the 3D domain but also lack sufficient granularity in representing the features of each modality. In this paper, we propose a unified multimodal 3D open-vocabulary scene understanding network, namely UniM-OV3D, which aligns point clouds with image, language and depth. To better integrate global and local features of the point clouds, we design a hierarchical point cloud feature extraction module that learns comprehensive fine-grained feature representations. Further, to facilitate the learning of coarse-to-fine point-semantic representations from captions, we propose the utilization of hierarchical 3D caption pairs, capitalizing on geometric constraints across various view-points of 3D scenes. Extensive experimental results demonstrate the effectiveness and superiority of our method in open-vocabulary semantic and instance segmentation, which achieves state-of-the-art performance on both indoor and outdoor benchmarks such as ScanNet, ScanNet200, S3IDS and nuScenes. Code is available at <https://github.com/hithqd/UniM-OV3D>.

1 Introduction

Accurate and resilient comprehension of 3D scenes plays a vital role in various practical applications, including autonomous driving [Mao *et al.*, 2022], virtual reality [Park *et al.*, 2008] and robot navigation [Shafiullah *et al.*, 2022]. Despite significant progress in recognizing closed-set categories on standard datasets [Graham *et al.*, 2018; Vu *et al.*, 2022; Misra *et al.*, 2021], existing models frequently struggle to recognize novel categories that are not present in the training data label space. The limited scalability of existing models in open-set scenarios hinders their practical applicability in the real world and motivates us to explore the open-vocabulary 3D scene understanding capability.




Figure 1: Open-vocabulary 3D scene understanding. Different colors represent the confidence of matching the user-specified query. (a) Comparison of different methods using the same query. (b) Results of our method in complex reasoning or content that requires extensive world knowledge.

Benefit from the 2D foundation model, such as CLIP [Radford *et al.*, 2021] and DINO [Caron *et al.*, 2021], previous methods [Peng *et al.*, 2023; Chen *et al.*, 2023; Liu *et al.*, 2023a; Liu *et al.*, 2023b] attempt to distill knowledge from 2D pixels to 3D features, aligning point clouds with 2D images. Another line of works [Takmaz *et al.*, 2023; Huang *et al.*, 2023c] focuses on points-only mechanism. The existing methods inadequately investigate data from multiple modalities, such as point clouds, images, linguistic input, and depth information, which are crucial for effective open-vocabulary 3D scene understanding. Despite their effectiveness, the absence of alignment and learning for other

modalities makes them miss a lot of semantic information and hinders their ability to effectively handle fine-grained point cloud object instances, as shown in Figure 1 (a).

Considering the properties of 3D scene, we should not overlook the inherent characteristics of 3D data itself. Firstly, depth information serves as a crucial modality for depth-invariant feature aggregation, which is often neglected. Early works [Zhang *et al.*, 2022; Zhu *et al.*, 2023; Huang *et al.*, 2023b] in 3D recognition have attempted to transfer depth modality to 3D understanding. However, mere projection or rendering, and merely aligning with a single modality such as image or text, cannot fully exploit all the characteristics of depth information, nor can it perform well in the task of open-vocabulary 3D scene understanding. The second aspect to consider is the exploration of point cloud information itself. Aiming at the point clouds caption learning, previous approaches either solely rely on simplistic templates derived from CLIP [Peng *et al.*, 2023; Liu *et al.*, 2023a; Zhang *et al.*, 2023] or utilize images as a bridge to generate corresponding textual descriptions from 2D images [Ding *et al.*, 2023b; Yang *et al.*, 2023]. However, point clouds captioning can make alignment easier and more comprehensive, particularly when multiple modalities are incorporated during the alignment training. And the point clouds encoder almost comes from mature 3D extractors [Choy *et al.*, 2019; Graham *et al.*, 2018], which often keep frozen during training. The frozen backbone and bridging alignment impede the generalization in complex 3D open vocabulary scenarios.

To this end, in order to fully leverage the synergistic advantages of various modalities, we propose a comprehensive multimodal alignment approach that co-embeds 3D points, image pixels, depth, and text strings into a unified latent space for open-vocabulary 3D scene understanding, namely UniM-OV3D. We thoroughly explore and utilize the characteristics of point clouds from two perspectives. First, to extract fine-grained geometric features of different levels from the original irregular 3D point clouds, we design a hierarchical point cloud feature extraction module. Second, in terms of generating point-caption pairs, we make the first attempt to generate corresponding text directly from point clouds instead of using images as a bridge. And we build hierarchical point-semantic caption pairs, including global-, eye- and sector-view captions, which can offer fine-grained language supervisions.

Furthermore, we introduce depth information into the entire network and the modified depth encoder inspired by CLIP2Point [Huang *et al.*, 2023b] keeps learning during training. As for image modality, we employ the pre-trained model PointBIND [Guo *et al.*, 2023] to extract image feature. Given the multi-modalities and their dense representations, we finally conduct multimodal contrastive learning for robust 3D understanding. In this way, we are capable of optimizing their reciprocal benefits by adeptly amalgamating multiple modalities.

Extensive experiments on ScanNet [Dai *et al.*, 2017], ScanNet200 [Rozenberszki *et al.*, 2022], S3DIS [Armeni *et al.*, 2016] and nuScenes [Caesar *et al.*, 2020] show the effectiveness of our method on 3D open-vocabulary tasks, surpassing previous state-of-the-art methods by 3.2%-7.8% hIoU on semantic segmentation and 3.8%-10.8% hAP₅₀ on instance seg-

mentation. More interestingly, UniM-OV3D demonstrates a robust ability to understand complex language queries, even those that involve intricate reasoning or necessitate extensive world knowledge, as illustrated in Figure 1 (b).

Our main contributions are summarized as follows:

- Within a joint embedding space, UniM-OV3D firstly aligns 3D point clouds with multi-modalities, including 2D images, language, depth, for robust 3D open-vocabulary scene understanding.
- In order to acquire more comprehensive and intricate fine-grained geometric features from point clouds, we propose a hierarchical point cloud extractor that effectively captures both local and global features.
- We innovatively build hierarchical point-semantic caption pairs that offer coarse-to-fine supervision signals, facilitating learning adequate point-caption representations from various 3D viewpoints directly.
- UniM-OV3D outperforms previous state-of-the-art methods on 3D open-vocabulary semantic and instance segmentation tasks by a large margin, covering both indoor and outdoor scenarios.

2 Related Work

2.1 3D Recognition

To process point clouds directly, PointNet [Qi *et al.*, 2017a] and PointNet++ [Qi *et al.*, 2017b] are the two pioneering studies in this field that focus on developing parallel MLPs for feature extraction from unstructured data which have significantly enhanced accuracy [Huang *et al.*, 2023a; Huang *et al.*, 2022]. Further, while the VLM [Radford *et al.*, 2021; Caron *et al.*, 2021] has achieved remarkable results in zero-shot or few-shot learning of 2D images by utilizing extensive image data, the availability of such internet-scale data for 3D point clouds is limited. Consequently, extending the open-vocabulary mechanism to 3D perception is a nontrivial task due to the difficulty in gathering sufficient data for point-language contrastive training, similar to the approach employed by CLIP [Radford *et al.*, 2021] in 2D perception. Therefore, PointCLIPs [Zhang *et al.*, 2022; Zhu *et al.*, 2023] take the first step to convert point clouds into CLIP-recognizable images and align the point cloud feature with the language feature by projecting 3D point cloud into 2D space. Similarly, CLIP2Point [Huang *et al.*, 2023b] attempts to align the projected depth map with the images by designing a trainable depth encoder and cross-modality learning losses. Different from point projecting, a series of works [Xue *et al.*, 2023a; Xue *et al.*, 2023b; Zeng *et al.*, 2023; Liu *et al.*, 2023b] try to collect multi-modalities triplets to train the 3D backbone. To enable zero-shot capability and improve standard 3D recognition, they align the 3D feature with the CLIP-aligned visual and textual features. This alignment facilitates the integration of 3D features with CLIP, enhancing both the zero-shot capability and the standard 3D recognition capability. However, existing methods cannot integrate the four modal information types, namely point cloud, image, depth, and text, which is the focus of this paper. We propose




Figure 2: **Architecture of our proposed UniM-OV3D.** The input point clouds are processed by a hierarchical point cloud extraction module to fuse the local and global features. To fulfill coarse-to-fine text supervision signal, the point-semantic caption learning is designed to acquire representations from various 3D viewpoints. The overall framework takes point clouds, 2D image, text and depth map as input to establish a unified multimodal contrastive learning for open-vocabulary 3D scene understanding.

a unified multimodal network that incorporates fine-grained feature representation from these four modalities, effectively leveraging their strengths.

2.2 Open-Vocabulary 3D Scene Understanding

Facing the challenge of the lack of diverse 3D open-vocabulary datasets for training generalizable models, one solution is to distill knowledge between pre-trained 2D open-vocabulary models and 3D models [Wu *et al.*, 2023a]. OpenScene [Peng *et al.*, 2023] and CLIP-FO3D [Zhang *et al.*, 2023] propose to distill the knowledge of 2D CLIP into 3D feature without any extra annotation. OpenMask3D [Takmaz *et al.*, 2023] first generates class-agnostic instance mask proposals with the 3D point cloud, which is used to choose the 2D views and masks. Meanwhile, 3D-OVS [Liu *et al.*, 2023a] distills the open-vocabulary multimodal knowledge and object reasoning capability of CLIP into a neural radiance field. OpenIns3D [Huang *et al.*, 2023c] employs a “Mask-Snap-Lookup” scheme to learn class-agnostic mask proposals in 3D point clouds. Another line focuses on point-language contrastive training. To obtain 3D backbone-language alignment, PLA-family [Ding *et al.*, 2023b; Yang *et al.*, 2023; Ding *et al.*, 2023a] uses 2D images as a bridge to generate descriptions of point cloud data. Although carefully designed, describing point clouds through images may not fully capture the details and point-semantic information due to differences between point clouds and images. Therefore, methods that describe point cloud data more directly and accurately need to be further explored. So we establish a point-semantic caption learning mechanism, in which captions can be obtained directly from point clouds with more accurate coarse-to-fine

languages supervisions.

3 Method

3.1 Preliminary

The objective of open-vocabulary 3D scene understanding is to localize and recognize unseen categories without the need for corresponding manual annotations. To put it formally, annotations on semantic and instance levels, denoted as $\mathcal{Y} = \{y^B, y^N\}$, are split into base categories $\mathcal{C}^B$ and novel categories $\mathcal{C}^N$. During the training phase, the 3D model has access to all point clouds $\mathcal{P} = \{\mathbf{p}\}$, but only to annotations for base classes $\mathcal{C}^B$. It remains unaware of both annotations y^N and category names related to novel classes $\mathcal{C}^N$. However, during the inference process, the 3D model is required to localize objects and classify points belonging to both basic and novel categories $\mathcal{C}^B \cup \mathcal{C}^N$. This inference step is crucial for achieving comprehensive scene understanding and enabling the model to generalize to unseen classes. During inference, the model can predict any desired category by computing the similarity between point-wise features and the embedding of queried categories, enabling open-world inference.

3.2 Hierarchical Feature Extractor with Local and Global Fusion

Taking the sparse point clouds as input, we propose a trainable hierarchical point cloud extractor to capture fine-grained local and global features, instead of just utilizing the frozen 3D extractors. As illustrated in Figure 2, the input is channeled into a transformer network, which employs attention-based layers to regress a 4×4 transformation matrix. This ma-

trix comprises the elements that represent the learned affine transformation values, which are used for aligning point clouds. Following alignment, the points are introduced into multiple stacked spatial-aware layers, which serve to produce a permutation-invariant embedding of these points. Within this structure, the residual attention module functions as a linking bridge between two adjacent layers, facilitating the transfer of information. After the information has been processed through the attention-based layers, the outputs from all these N -dimensional layers are concatenated. Finally, the segmentation head can be added to output the global information aggregation for the point clouds, providing a comprehensive summary of the data.

Spatial-aware Layers. The architecture of the spatial-aware layers is shown in Figure 2, where the current i layer acts as an intermediate layer and the features are not only transmitted by the mainstream information flow but the residual attention linking from $i - 1$ layer. Consider a R^D dimensional embedding point clouds set with n points in the i layer $\mathbf{p} = \{p_{i,1}, p_{i,2}, \dots, p_{i,n}\}$. In each attention-based flow, the local PointBERT [Yu *et al.*, 2022] layer and the global PointMAE [Pang *et al.*, 2022] layer process these points in parallel. In the branch of the local PointBERT layer, a transformation is employed on the points to get the feature H_L^i :

$$H_L^i = \{h_{i,1}^M, h_{i,2}^M, \dots, h_{i,g}^M\} = B(p_{i,1}, p_{i,2}, \dots, p_{i,n}) \quad (1)$$

where $B : R^D \rightarrow R^M$ and g is the local patch. After that, the attention map from $i - 1$ layer by the residual linking is added to the output H_L^i in a point-wise manner which can be written as:

$$f_L^i = H_L^i + \omega_1 \otimes H_L^{i-1} \quad (2)$$

where $\otimes$ denotes element-wise multiplication. ω_1 functions as a feature filter, dampening the impact of noisy features while amplifying the beneficial ones. Similarly, global PointMAE layer branch applies global transformation E on the points $\mathbf{p}$ and the output from decoder denotes H_G^i . And the final output after the linking of attention block is:

$$f_G^i = H_G^i + \omega_2 \otimes H_G^{i-1} \quad (3)$$

where ω_2 is similar to ω_1 . The outputs f_L^i and f_G^i from the two branches are of the same dimension and are transferred to the attention block of the next layer. Ultimately, they are concatenated together, and these combined embedding are used as the input for the subsequent layer.

The attention block consists of a series of residual units [He *et al.*, 2016], inserted with interpolation to form a symmetrical top-down architecture. Two consecutive 1×1 convolution layers are connected at the end to balance the dimensions. More details are presented in the supplementary.

3.3 Point-semantic Caption Learning

Recent success of open-vocabulary works in 2D [Bangalath *et al.*, 2022; Wu *et al.*, 2023b] and 3D [Ding *et al.*, 2023b; Yang *et al.*, 2023] vision has shown the effectiveness of introducing language supervision to guide the model to access abundant semantic concepts with a large vocabulary size. However, utilizing images as a bridge to associate the points




Figure 3: Point-caption pairs comparison. The middle column represents the captions generated by the point-captioning model, PointLLM [Xu *et al.*, 2023], and the third column represents the corresponding image captions.

with language is the core idea in these 3D works. And the final textual descriptions in the so-called point-caption pairs are all sourced from image descriptions and do not possess semantic information specific to the point cloud. To address this problem, we propose to provide direct language supervision and a coarse-to-fine point-semantic caption generation mechanism is designed.

Global-view Point-semantic Caption. Due to the similarity of point-captioning and image captioning [Mokady *et al.*, 2021; Wang *et al.*, 2022] as fundamental tasks, various base models [Xu *et al.*, 2023; Guo *et al.*, 2023] trained on a large number of 3D point samples are already available for perceiving 3D scenes. Given the i^{th} 3D point cloud scene $\mathbf{p}_i$, the simplest and most rudimentary manner is to link language supervision to all points. Specifically, the corresponding global language description $\mathbf{t}_i^{global}$ can be generated by the pre-trained point-captioning model $\mathcal{G}$ as:

$$\mathbf{t}_i^{global} = \mathcal{G}(\mathbf{p}_i). \quad (4)$$

Surprisingly, with the semantic understanding on the overall points, the entities covered in the generated captions have already encompassed the entire 3D semantic label space. As illustrated in Figure 3, the middle column represents the captions generated by the point-captioning model, and the third column represents the corresponding image captions. It is observable that the global point-captions not only offer a more precise and holistic depiction of the scene, but they also represent the orientation information of objects more accurately within the scene and the interrelationships among them. Despite the simplicity of global-view caption, we find that it can significantly improve the whole open vocabulary capabilities of understanding the scene and locating objects.

Hierarchical Point-semantic Caption. Diving into the exploration of more fine-grained point-semantic caption learning, global-view point cloud captions are still too coarse to capture fine-grained details between objects in point clouds, making them suboptimal for 3D scene understanding tasks. Hierarchical point captions can further refine the semantic information of point clouds without ignoring all object features in the scene. Therefore, we propose two other

association patterns on point sets at different 3D viewpoints.

For each 3D point cloud scene $\mathbf{p}_i$, we initiate from the x-axis and segment the entire point cloud space into sector-shaped regions with an angular magnitude of θ degree, where the intersecting angle between adjacent sectors is ϕ degree. The segmented 3D scene $\mathbf{p}_i^s$ can be denoted as:

$$\mathbf{p}_i^s = \{p_{ij}^{\theta, \phi} \mid 0 \leq \phi < \theta, \theta > 0, 0 < j < \lceil \frac{360}{\theta} \rceil\}. \quad (5)$$

For the partitioned sector-shaped regions $p_{ij}^{\theta, \phi}$, we categorize them into two distinct types of viewpoints based on the magnitude of the angle, named as *sector-view* and *eye-view*. The corresponding angle size of the two views is θ_1 and θ_2 , where $0 < \theta_1 \leq 90$ and $90 < \theta_2 \leq 180$. This implies that the *sector-view* focuses on the minutiae of local point cloud information, while the *eye-view* concentrates on a broader point cloud area, yet the extent of this area is within an obtuse angle range, akin to human eyes.

Contrastive Point-semantic Caption Training. With the obtained global-, eye- and sector-view point-semantic captions from different views, we can conduct point-language feature contrastive learning to guide the 3D network to learn from vocabulary-rich language supervisions.

Specifically, we extract the text embeddings $\mathbf{f}_i^t$ based on the pre-aligned and fixed text encoder in PointBIND [Guo *et al.*, 2023]. And the 3D feature $\mathbf{f}_i^p$ can be acquired by the trainable 3D point cloud encoder. We formulate the point-to-text alignment using the contrastive loss as:

$$\mathcal{L}_{capt}^m = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(\mathbf{f}_i^p \cdot \mathbf{f}_i^t / \tau)}{\sum_{j=0}^N \exp(\mathbf{f}_i^p \cdot \mathbf{f}_j^t / \tau)} \quad (6)$$

where $m \in \{global, eye, sector\}$, N is the total number of the point-semantic caption pairs and τ is a learnable temperature parameter. With Equation 6, we can easily compute different caption losses on global-view $\mathcal{L}_{capt}^{global}$, eye-view $\mathcal{L}_{capt}^{eye}$ and sector-view $\mathcal{L}_{capt}^{sector}$. The final caption loss is a weighted combination between different views as follows:

$$\mathcal{L}_{capt}^{total} = \alpha \mathcal{L}_{capt}^{global} + \beta \mathcal{L}_{capt}^{eye} + \gamma \mathcal{L}_{capt}^{sector} \quad (7)$$

where α , β and γ are used to balance the relative importance of different parts and are set to 1, 0.8, 0.8 by default. The effect of different views and the comparison of different combinations are shown in ablation studies.

3.4 Aligning Multimodal Representations

With the construction of the four modalities, the subsequent objective of UniM-OV3D is to conduct dense alignment between the 3D points with their corresponding image, depth and text.

For the paired image-text ($\mathbf{I}$, $\mathbf{t}$) data, we utilize their corresponding encoders from PointBIND [Guo *et al.*, 2023] for feature extraction, which are kept frozen during training. For each depth map $\mathcal{D}$, we set the gated aggregator weights to 0 in CLIP2Point [Huang *et al.*, 2023b] to extract the depth feature, formulated as

$$\begin{aligned} \mathbf{f}^t, \mathbf{f}^{img} &= PointBIND(\mathbf{t}, \mathbf{I}) \\ \mathbf{f}^d &= E_d(\mathcal{D}) \end{aligned} \quad (8)$$

where $\mathbf{f}^t$, $\mathbf{f}^{img}$, $\mathbf{f}^d$ denote the text, image and depth embeddings. And the depth encoder E_d keeps learnable during training. For the 3D point clouds, we adopt the designed hierarchical point cloud feature extractor module to extract 3D features for the point cloud. In order to transform the encoded 3D feature into multi-modal embedding space, we append a projection network after the 3D encoder E_p . So the final 3D features $\mathbf{f}^p$ can be formulate as:

$$\mathbf{f}^p = Proj(E_p(\mathbf{p})). \quad (9)$$

Dense Associations across Modalities. As shown in Figure 2, with a 3D scene and its corresponding features $\mathbf{f}^p$, $\mathbf{f}^t$, $\mathbf{f}^{img}$, $\mathbf{f}^d$, we employ a contrastive loss between point clouds and other modalities, which effectively compels the 3D embeddings to align with the joint representation space, formulated as:

$$\begin{aligned} \mathcal{L}_{(M_1, M_2)} &= \sum_{i,j} -\frac{1}{2} \log \frac{\exp(\mathbf{f}_i^{M_1} \cdot \mathbf{f}_j^{M_2} / \epsilon)}{\sum_k \exp(\mathbf{f}_i^{M_1} \cdot \mathbf{f}_k^{M_2} / \epsilon)} \\ &\quad -\frac{1}{2} \log \frac{\exp(\mathbf{f}_i^{M_1} \cdot \mathbf{f}_j^{M_2} / \epsilon)}{\sum_k \exp(\mathbf{f}_k^{M_1} \cdot \mathbf{f}_j^{M_2} / \epsilon)} \end{aligned} \quad (10)$$

where M_1 and M_2 represent two modalities in (*point, image, depth*) and (i, j) indicates a positive pair in each training batch. And ϵ is a learnable temperature parameter, similar to CLIP [Radford *et al.*, 2021].

Finally, the overall objective function to perform unified multimodal alignment is defined by:

$$\mathcal{L}_{overall} = \mathcal{L}_{(P, I)} + \mathcal{L}_{(P, D)} + \mathcal{L}_{(D, I)} + \mathcal{L}_{capt}^{total} \quad (11)$$

where the language modality delivers comprehensive and scalable textual descriptions, while the image modality provides accurate guidance on object edges and contextual data. Moreover, the depth and 3D modality exposes vital structural details of objects. By aligning these modalities jointly in a common space, our approach can maximize the synergistic advantages among them, resulting in superior open-vocabulary scene understanding performance.

4 Experiments

4.1 Setup

Datasets and Metrics. To validate the effectiveness of our proposed UniM-OV3D, we conduct extensive experiments on four popular public 3D benchmarks: ScanNet [Dai *et al.*, 2017], ScanNet200 [Rozenberszki *et al.*, 2022], S3DIS [Armeni *et al.*, 2016] and nuScenes [Caesar *et al.*, 2020]. The first three provide RGBD images and 3D meshes of indoor scenes, while the last one supplies Lidar scans of outdoor scenes. We use all four datasets to compare to alternative methods, covering both semantic and instance segmentation. For semantic segmentation, we employ the commonly used 3D segmentation metric mean intersection over union (mIoU^B, mIoU^N) and harmonic mean IoU (hIoU) for evaluating base, novel categories and their harmonic mean separately. For instance segmentation, We apply mean average precision under 50% IoU threshold (mAP₅₀^B, mAP₅₀^N, hAP₅₀) as major indicators.

Method	ScanNet									S3DIS					
	B15/N4			B12/N7			B10/N9			B8/N4			B6/N6		
	hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N
LSeg-3D [Li <i>et al.</i> , 2022]	0.0	64.4	0.0	0.9	55.7	0.1	1.8	68.4	0.9	0.1	49.0	0.1	0.0	30.1	0.0
3DGenZ [Michele <i>et al.</i> , 2021]	20.6	56.0	12.6	19.8	35.5	13.3	12.0	63.6	6.6	8.8	50.3	4.8	9.4	20.3	6.1
3DTZSL [Cheraghian <i>et al.</i> , 2020]	10.5	36.7	6.1	3.8	36.6	2.0	7.8	55.5	4.2	8.4	43.1	4.7	3.5	28.2	1.9
PLA [Ding <i>et al.</i> , 2023b]	65.3	68.3	62.4	55.3	69.5	45.9	53.1	76.2	40.8	34.6	59.0	24.5	38.5	55.5	29.4
OpenScene [Peng <i>et al.</i> , 2023] [†]	67.1	68.8	62.8	56.8	61.5	51.7	55.7	71.8	43.6	39.1	58.6	33.2	41.2	56.2	36.4
RegionPLC [Yang <i>et al.</i> , 2023]	69.4	68.2	70.7	68.2	69.9	66.6	64.3	76.3	55.6	-	-	-	-	-	-
OpenIns3D [Huang <i>et al.</i> , 2023c] [†]	72.6	71.8	73.4	69.3	70.7	68.8	64.5	80.2	49.8	46.8	63.2	38.6	53.6	60.8	47.5
UniM-OV3D (Ours)	75.8	75.3	76.6	74.7	75.2	74.1	69.9	83.5	57.3	54.6	68.3	45.2	59.1	66.7	54.2

Table 1: Results for open-vocabulary 3D semantic segmentation on ScanNet and S3DIS. The evaluation metrics are hIoU, mIoU^B and mIoU^N. † denotes results reproduced by us on its official implementation. Best open-vocabulary results are highlighted in **bold**.

Method	ScanNet200					
	B170/N30			B150/N50		
	hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N
3DGenZ [Michele <i>et al.</i> , 2021]	2.6	15.8	1.4	3.3	14.1	1.9
3DTZSL [Cheraghian <i>et al.</i> , 2020]	0.9	4.0	0.5	0.7	3.8	0.4
LSeg-3D [Li <i>et al.</i> , 2022]	1.5	21.1	0.8	3.0	20.6	1.6
PLA [Ding <i>et al.</i> , 2023b]	11.4	20.9	7.8	10.1	20.9	6.6
OpenScene [Peng <i>et al.</i> , 2023] [†]	15.3	22.5	10.4	16.8	23.5	11.2
RegionPLC [Yang <i>et al.</i> , 2023]	16.6	21.6	13.9	14.6	22.4	10.8
OpenIns3D [Huang <i>et al.</i> , 2023c] [†]	17.2	24.6	13.6	18.9	25.8	16.4
UniM-OV3D (Ours)	22.3	29.5	17.1	25.8	30.7	21.6

Table 2: Results for open-vocabulary 3D semantic segmentation on ScanNet200 on hIoU, mIoU^B and mIoU^N. † denotes results reproduced by us on its official implementation.

Category Partition. ScanNet densely annotates 20 classes, ScanNet200 owns 200 classes, S3DIS contains 13 classes and nuScenes involves 16 classes. Following [Ding *et al.*, 2023b; Yang *et al.*, 2023], we disregard the “otherfurniture” class in ScanNet and randomly partition the rest 19 classes into 3 base/novel partitions, i.e. B15/N4 (15 base and 4 novel categories), B12/N7 and B10/N9, for semantic segmentation. Following SoftGroup [Vu *et al.*, 2022] to exclude two background classes, we acquire B13/N4, B10/N7, and B8/N9 partitions for instance segmentation on ScanNet. Similarly, we ignore the “clutter” class in S3DIS and get B8/N4, B6/N6 for both semantic and instance segmentation. For ScanNet200, we split 200 classes to B170/N30 and B150/N50. For nuScenes, we drop the “otherflat” class and obtain B12/N3 and B10/N5.

4.2 Main Results

To validate the effectiveness of our proposed UniM-OV3D, we compare it with previous state-of-the-art 3D open-vocabulary methods on semantic and instance segmentation tasks.

3D Semantic Segmentation. As shown in Table 1, compared to previous captioning-learning based method RegionPLC [Yang *et al.*, 2023], we obtain 5.1%-13% mIoU performance gains among different partitions on ScanNet. Compared with previous zero-shot state-of-the-art method OpenIns3D [Huang *et al.*, 2023c], our method still obtains 3.2%-5.4% and 5.5%-7.8% improvements in terms of hIoU across various datasets and partitions. It is noteworthy that our approach consistently improves performance across different datasets and partitions. Moreover, for partitions with

Method	nuScenes					
	B12/N3			B10/N5		
	hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N
3DGenZ [Michele <i>et al.</i> , 2021]	1.6	53.3	0.8	1.9	44.6	1.0
3DTZSL [Cheraghian <i>et al.</i> , 2020]	1.2	21.0	0.6	6.4	17.1	3.9
LSeg-3D [Li <i>et al.</i> , 2022]	0.6	74.4	0.3	0.0	71.5	0.0
PLA [Ding <i>et al.</i> , 2023b]	47.7	73.4	35.4	24.3	73.1	14.5
OpenScene [Peng <i>et al.</i> , 2023] [†]	55.6	75.4	39.6	28.5	75.8	19.2
RegionPLC [Yang <i>et al.</i> , 2023]	64.4	75.8	56.0	49.0	75.8	36.3
OpenIns3D [Huang <i>et al.</i> , 2023c] [†]	65.4	77.3	59.6	49.7	78.7	39.2
UniM-OV3D (Ours)	70.2	80.5	64.6	55.1	81.7	44.8

Table 3: Results for open-vocabulary 3D semantic segmentation on nuScenes in terms of hIoU, mIoU^B and mIoU^N. † denotes results reproduced by us on its official implementation.




Figure 4: Comparisons on open-vocabulary 3D instance segmentation. (a) B13/N4 on ScanNet, (b) B8/N4 on S3DIS, (c) B12/N3 on nuScenes and (d) B170/N30 on ScanNet200.

more novel categories, our method shows even greater enhancement. This further demonstrates the robustness and effectiveness of our approach in open vocabulary scene understanding tasks.

When facing the long-tail problems in ScanNet200, as shown in Table 2, our method still surpasses corresponding zero-shot method by 5.1%-6.9% hIoU on both splits and 3.5%-5.2% mIoU on novel categories. Furthermore, when evaluating the performance of our UniM-OV3D on outdoor setting, our method lifts the hIoU and mIoU of novel categories by 4.8%-6.4% and 5%-5.6% as shown in Table 3. In this regard, UniM-OV3D demonstrates superior generalizability to complex 3D open-world scenarios.

Uni	HFE	PCL	B15/N4			B12/N7		
			hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N
-	-	-	46.2	57.1	42.3	43.5	58.6	41.0
-	-	✓	62.1	66.2	56.4	59.6	67.9	55.3
-	✓	-	60.0	65.8	56.2	59.2	66.5	55.1
✓	-	-	60.3	65.4	57.6	58.7	66.2	56.5
✓	✓	-	72.9	72.8	75.3	72.6	73.5	71.3
✓	-	✓	72.6	72.1	75.1	71.9	73.6	71.8
-	✓	✓	74.8	74.5	76.0	73.9	74.6	73.1
✓	✓	✓	75.8	75.3	76.6	74.7	75.2	74.1

Table 4: Ablations of different components of our network. Uni denotes the uni-modalities, HFE is the hierarchical point cloud feature extractor and PCL represents the point-semantic caption learning.

Fusion Method	B15/N4			B12/N7		
	hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N
Local PB	72.8	72.4	73.0	71.6	71.8	71.3
Global PM	73.1	72.6	73.2	71.7	72.0	71.4
Local PB + Global PM	75.0	74.8	75.3	73.7	73.8	73.2
Local PB + Attn	74.8	74.6	75.0	73.6	73.8	73.1
Global PM + Attn	74.9	74.8	75.2	73.8	74.3	73.5
Local PB + Global PM + Attn	75.8	75.3	76.6	74.7	75.2	74.1

Table 5: Ablations of different fusion methods. Local PB denotes local PointBERT, Global PM is the global PointMAE and Attn represents the attention block.

3D Instance Segmentation. Our pipeline not only offers point-semantic language descriptions for fine-grained point clouds, but also encourages points to learn discriminative features, which in turn benefits the instance-level localization task. As shown in Figure 4, by learning a uni-modalities feature representations, our method outperforms Region-PLC [Yang *et al.*, 2023] by 4.1% hAP₅₀, 5.3% mAP₅₀^B, 3% mAP₅₀^N on B13/N4 on the ScanNet. On the other datasets, our UniM-OV3D surpasses OpenIns3D [Huang *et al.*, 2023c] by 5%-10.6% hAP₅₀, 5%-5.3% mAP₅₀^B and 3.1%-13.2% mAP₅₀^N on their respective B8/N4, B12/N3 and B170/N30 partitions (RegionPLC only reports their instance segmentation results on ScanNet without open-source code). This illustrates the efficacy of our UniM-OV3D in empowering the model to identify unseen instances without human annotations.

4.3 Ablation Studies

In this section, we change components and variants of our proposed UniM-OV3D by conducting extensive ablation studies on ScanNet dataset for semantic segmentation by default. And the partitions are mainly B15/N4 and B12/N7.

Analysis of Different Components. We illustrate the importance of different components of our network by removing some parts and keeping all the others unchanged. The baseline setting is to only use three modalities of image, text and point cloud. The 3D point cloud encoder uses sparse convolutional UNet [Choy *et al.*, 2019]. The point captions come from corresponding image captions, and we use CLIP as the image and text encoder. As shown in Table 4, using the depth map to construct the uni-modality learning contributes much to the whole performance which proves the necessity of establishing a unified modal architecture. Without the point-semantic captions from different view, the performance drops dramatically which shows that coarse-to-fine supervi-

Angle		B15/N4			B12/N7		
		hIoU	mIoU ^B	mIoU ^N	hIoU	mIoU ^B	mIoU ^N
sector-view	30	66.5	68.4	62.1	67.7	69.3	61.3
	45	69.1	69.4	68.6	68.3	70.2	66.3
	60	72.3	72.3	74.2	70.2	73.3	69.2
	90	72.6	72.3	73.9	70.1	73.2	69.1
eye-view	120	73.1	72.4	74.3	71.3	73.4	70.5
	180	71.9	71.3	73.6	71.1	73.1	70.0
global-view	360	72.8	72.7	75.3	72.6	73.5	71.3
combined-view	360+60	73.3	73.2	75.6	73.5	74.2	72.6
	360+120	74.2	74.1	75.7	73.6	74.3	72.7
	360+60+120	75.8	75.3	76.6	74.7	75.2	74.1

Table 6: Performance of splitting point clouds into different angles and combining different views.

sion signal from fine-grained point-semantic caption learning is indeed a crucial factor. And the performance degradation caused by the absence of hierarchical point cloud extractor proves that only the combination of local and global features can produce the best results.

Comparison of Different Fusion Methods in Hierarchical Feature Extractor.

To validate the effect of different fusion methods in the hierarchical point cloud feature extractor, we design six fusion patterns as shown in Table 5. With the attention block to transfer features between connected layers, the fused models override the original ones by 0.8%-2.3% IoU. And our final fusion method of global and local with attention block outperforms the alternative by 0.5%-1.4% mIoU.

Effect of Different Point-semantic Caption Views. In order to explore the effect of dividing the point cloud into different angles in Section 3.3 to obtain the corresponding fine-grained point-semantic description, we split the point cloud into six angles for experiments. As shown in Table 6, the 120 degree in eye-view and 60 degree in sector-view perform relatively well, but they are both slightly inferior to global-view (360 degree). So, in order to further obtain the coarse-to-fine supervision, we combine 60 degree and 120 degree with 360 degree respectively, and finally adopt the method of combining the three in the last row.

5 Conclusion

In this paper, we propose UniM-OV3D, a unified multimodal network for open-vocabulary 3D scene understanding. We jointly establish dense alignment between four modalities, *i.e.*, point clouds, images, depth and text, to leverage their respective advantages. To fully capture local and global features of point clouds, we stack spatial-aware layers to construct a hierarchical point cloud extractor. Furthermore, we deeply explore the possibility of directly using point clouds to generate corresponding captions, and the established point-semantic learning mechanism can provide language supervision signals more effectively. Extensive experiments on both indoor and outdoor datasets demonstrate the superiority of our model in 3D open-vocabulary scene understanding task.

Contribution Statement

All authors conceive the paper’s methodology and experiments. Qingdong He and Jinlong Peng contribute equally to

this work and should be considered co-first authors. Yabiao Wang and Jiangning Zhang are the co-corresponding authors of this paper.

References

- [Armeni *et al.*, 2016] 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.
- [Bangalath *et al.*, 2022] Hanoona Bangalath, Muhammad Maaz, Muhammad Uzair Khattak, Salman H Khan, and Fahad Shahbaz Khan. Bridging the gap between object and image-level representations for open-vocabulary detection. *NeurIPS*, 2022.
- [Caesar *et al.*, 2020] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *CVPR*, 2020.
- [Caron *et al.*, 2021] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021.
- [Chen *et al.*, 2023] Runnan Chen, Youquan Liu, Lingdong Kong, Xinge Zhu, Yuexin Ma, Yikang Li, Yuenan Hou, Yu Qiao, and Wenping Wang. Clip2scene: Towards label-efficient 3d scene understanding by clip. In *CVPR*, 2023.
- [Cheraghian *et al.*, 2020] Ali Cheraghian, Shafin Rahman, Dylan Campbell, and Lars Petersson. Transductive zero-shot learning for 3d point cloud classification. In *CACV*, 2020.
- [Choy *et al.*, 2019] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *CVPR*, 2019.
- [Dai *et al.*, 2017] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, 2017.
- [Ding *et al.*, 2023a] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. Lowis3d: Language-driven open-world instance-level 3d scene understanding. *arXiv preprint arXiv:2308.00353*, 2023.
- [Ding *et al.*, 2023b] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. Pla: Language-driven open-vocabulary 3d scene understanding. In *CVPR*, 2023.
- [Graham *et al.*, 2018] Benjamin Graham, Martin Engelcke, and Laurens Van Der Maaten. 3d semantic segmentation with submanifold sparse convolutional networks. In *CVPR*, 2018.
- [Guo *et al.*, 2023] Ziyu Guo, Renrui Zhang, Xiangyang Zhu, Yiwen Tang, Xianzheng Ma, Jiaming Han, Kexin Chen, Peng Gao, Xianzhi Li, Hongsheng Li, et al. Point-bind & point-llm: Aligning point cloud with multi-modality for 3d understanding, generation, and instruction following. *arXiv preprint arXiv:2309.00615*, 2023.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *ECCV*, 2016.
- [Huang *et al.*, 2022] Tianxin Huang, Jiangning Zhang, Jun Chen, Zhonggan Ding, Ying Tai, Zhenyu Zhang, Chengjie Wang, and Yong Liu. 3qnet: 3d point cloud geometry quantization compression network. *ACM TOG*, 2022.
- [Huang *et al.*, 2023a] Tianxin Huang, Zhonggan Ding, Jiangning Zhang, Ying Tai, Zhenyu Zhang, Mingang Chen, Chengjie Wang, and Yong Liu. Learning to measure the point cloud reconstruction loss in a representation space. In *CVPR*, 2023.
- [Huang *et al.*, 2023b] Tianyu Huang, Bowen Dong, Yunhan Yang, Xiaoshui Huang, Rynson WH Lau, Wanli Ouyang, and Wangmeng Zuo. Clip2point: Transfer clip to point cloud classification with image-depth pre-training. In *ICCV*, 2023.
- [Huang *et al.*, 2023c] Zhening Huang, Xiaoyang Wu, Xi Chen, Hengshuang Zhao, Lei Zhu, and Joan Lasenby. Openins3d: Snap and lookup for 3d open-vocabulary instance segmentation. *arXiv preprint arXiv:2309.00616*, 2023.
- [Li *et al.*, 2022] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *ICLR*, 2022.
- [Liu *et al.*, 2023a] Kunhao Liu, Fangneng Zhan, Jiahui Zhang, Muyu Xu, Yingchen Yu, Abdulmotaleb El Saddik, Christian Theobalt, Eric Xing, and Shijian Lu. Weakly supervised 3d open-vocabulary segmentation. In *NeurIPS*, 2023.
- [Liu *et al.*, 2023b] Minghua Liu, Ruoxi Shi, Kaiming Kuang, Yinhao Zhu, Xuanlin Li, Shizhong Han, Hong Cai, Fatih Porikli, and Hao Su. Openshape: Scaling up 3d shape representation towards open-world understanding. *arXiv preprint arXiv:2305.10764*, 2023.
- [Mao *et al.*, 2022] Jiageng Mao, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. 3d object detection for autonomous driving: A review and new outlooks. *arXiv preprint arXiv:2206.09474*, 2022.
- [Michele *et al.*, 2021] Björn Michele, Alexandre Boulch, Gilles Puy, Maxime Bucher, and Renaud Marlet. Generative zero-shot learning for semantic segmentation of 3d point clouds. In *3DV*, 2021.
- [Misra *et al.*, 2021] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *ICCV*, 2021.
- [Mokady *et al.*, 2021] Ron Mokady, Amir Hertz, and Amit H Bermano. Clipcap: Clip prefix for image captioning. *arXiv preprint arXiv:2111.09734*, 2021.
- [Pang *et al.*, 2022] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *ECCV*, 2022.

- [Park *et al.*, 2008] Youngmin Park, Vincent Lepetit, and Woontack Woo. Multiple 3d object tracking for augmented reality. In *ISMAR*, 2008.
- [Peng *et al.*, 2023] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *CVPR*, 2023.
- [Qi *et al.*, 2017a] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, pages 652–660, 2017.
- [Qi *et al.*, 2017b] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 2017.
- [Radford *et al.*, 2021] 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 *ICML*, 2021.
- [Rozenberszki *et al.*, 2022] David Rozenberszki, Or Litany, and Angela Dai. Language-grounded indoor 3d semantic segmentation in the wild. In *ECCV*, 2022.
- [Shafiullah *et al.*, 2022] Nur Muhammad Mahi Shafiullah, Chris Paxton, Lerrel Pinto, Soumith Chintala, and Arthur Szlam. Clip-fields: Weakly supervised semantic fields for robotic memory. *arXiv preprint arXiv:2210.05663*, 2022.
- [Takmaz *et al.*, 2023] Ayça Takmaz, Elisabetta Fedele, Robert W Sumner, Marc Pollefeys, Federico Tombari, and Francis Engelmann. Openmask3d: Open-vocabulary 3d instance segmentation. *arXiv preprint arXiv:2306.13631*, 2023.
- [Vu *et al.*, 2022] Thang Vu, Kookhoi Kim, Tung M Luu, Thanh Nguyen, and Chang D Yoo. Softgroup for 3d instance segmentation on point clouds. In *CVPR*, 2022.
- [Wang *et al.*, 2022] Peng Wang, An Yang, Rui Men, Junyang Lin, Shuai Bai, Zhikang Li, Jianxin Ma, Chang Zhou, Jingren Zhou, and Hongxia Yang. Ofa: Unifying architectures, tasks, and modalities through a simple sequence-to-sequence learning framework. In *ICML*, 2022.
- [Wu *et al.*, 2023a] Jianzong Wu, Xiangtai Li, Haobo Yuan, Shilin Xu, Henghui Ding, Yibo Yang, Xia Li, Jiangning Zhang, Yunhai Tong, Xudong Jiang, Bernard Ghanem, et al. Towards open vocabulary learning: A survey. *arXiv preprint arXiv:2306.15880*, 2023.
- [Wu *et al.*, 2023b] Weijia Wu, Yuzhong Zhao, Mike Zheng Shou, Hong Zhou, and Chunhua Shen. Diffumask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models. *arXiv preprint arXiv:2303.11681*, 2023.
- [Xu *et al.*, 2023] Runsen Xu, Xiaolong Wang, Tai Wang, Yilun Chen, Jiangmiao Pang, and Dahua Lin. Pointllm: Empowering large language models to understand point clouds. *arXiv preprint arXiv:2308.16911*, 2023.
- [Xue *et al.*, 2023a] Le Xue, Mingfei Gao, Chen Xing, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, and Silvio Savarese. Ulip: Learning a unified representation of language, images, and point clouds for 3d understanding. In *CVPR*, 2023.
- [Xue *et al.*, 2023b] Le Xue, Ning Yu, Shu Zhang, Junnan Li, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, and Silvio Savarese. Ulip-2: Towards scalable multimodal pre-training for 3d understanding. *arXiv preprint arXiv:2305.08275*, 2023.
- [Yang *et al.*, 2023] Jihan Yang, Runyu Ding, Zhe Wang, and Xiaojuan Qi. Regionplc: Regional point-language contrastive learning for open-world 3d scene understanding. *arXiv preprint arXiv:2304.00962*, 2023.
- [Yu *et al.*, 2022] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *CVPR*, 2022.
- [Zeng *et al.*, 2023] Yihan Zeng, Chenhan Jiang, Jiageng Mao, Jianhua Han, Chaoqiang Ye, Qingqiu Huang, Dit-Yan Yeung, Zhen Yang, Xiaodan Liang, and Hang Xu. Clip2: Contrastive language-image-point pretraining from real-world point cloud data. In *CVPR*, 2023.
- [Zhang *et al.*, 2022] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *CVPR*, 2022.
- [Zhang *et al.*, 2023] Junbo Zhang, Runpei Dong, and Kaisheng Ma. Clip-fo3d: Learning free open-world 3d scene representations from 2d dense clip. *arXiv preprint arXiv:2303.04748*, 2023.
- [Zhu *et al.*, 2023] Xiangyang Zhu, Renrui Zhang, Bowei He, Ziyu Guo, Ziyao Zeng, Zipeng Qin, Shanghang Zhang, and Peng Gao. Pointclip v2: Prompting clip and gpt for powerful 3d open-world learning. In *ICCV*, 2023.