

Open-Vocabulary Point-Cloud Object Detection without 3D Annotation

Yuheng Lu^{1*}, Chenfeng Xu^{2*}, Xiaobao Wei¹, Xiaodong Xie¹,
Masayoshi Tomizuka², Kurt Keutzer², Shanghang Zhang^{1†}

¹National Key Laboratory for Multimedia Information Processing,
School of Computer Science, Peking University ²University of California Berkeley
{yuhenglu, xiaobaowei, donxie, shanghang}@pku.edu.cn
{xuchenfeng, tomizuka, keutzer}@berkeley.edu,

Abstract

The goal of open-vocabulary detection is to identify novel objects based on arbitrary textual descriptions. In this paper, we address open-vocabulary 3D point-cloud detection by a dividing-and-conquering strategy, which involves: 1) developing a point-cloud detector that can learn a general representation for localizing various objects, and 2) connecting textual and point-cloud representations to enable the detector to classify novel object categories based on text prompting. Specifically, we resort to rich image pre-trained models, by which the point-cloud detector learns localizing objects under the supervision of predicted 2D bounding boxes from 2D pre-trained detectors. Moreover, we propose a novel de-biased triplet cross-modal contrastive learning to connect the modalities of image, point-cloud and text, thereby enabling the point-cloud detector to benefit from vision-language pre-trained models, i.e., CLIP. The novel use of image and vision-language pre-trained models for point-cloud detectors allows for open-vocabulary 3D object detection without the need for 3D annotations. Experiments demonstrate that the proposed method improves at least 3.03 points and 7.47 points over a wide range of baselines on the ScanNet and SUN RGB-D datasets, respectively. Furthermore, we provide a comprehensive analysis to explain why our approach works. Code is available at <https://github.com/lyhdet/OV-3DET>

1. Introduction

Current state-of-the-art point-cloud detectors are typically trained on a limited number of classes, which fails to account for the cornucopia of object classes that exist in the real world. As a result, these detectors often fail to generalize to unseen object classes that were not present

*Equal Contribution.

†Corresponding Author.

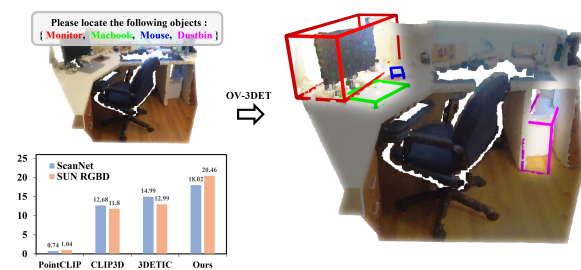


Figure 1. OV-3DET takes point-cloud and text as input, and detects objects according to the text description. Results on two point-cloud datasets demonstrate the effectiveness of OV-3DET. Note that RGB color is not used during inference. (we use it here only for better visualization.)

in the training data. To address this issue, we focus on the under-explored problem of open-vocabulary point-cloud detection, which involves detecting 3D objects that are outside the training vocabulary and accompanied by arbitrary text descriptions.

Achieving open-vocabulary detection requires the model to learn general representations and relate those representations to text cues. However, the typical approach of collecting data with a larger number of classes is challenging in the point-cloud field due to the difficulty of both data collection and annotation. Moreover, acquiring large-scale point-cloud-caption data is currently infeasible, which further hinders point-cloud detectors from learning to connect the representation with text prompts. Given these challenges, we work around the problem of open-vocabulary 3D object detection by leveraging existing well-established pre-trained models for images.

Indeed, the field of image/vision-language has made remarkable progress in learning general representations [8, 12, 19, 34], enabling various applications such as detection [34], vision-language question answering [26], and even image generation [22]. In particular, the CLIP pre-

trained model [19] connects visual and text representations, significantly advancing the development of 2D open-vocabulary detection/segmentation [10, 16, 31]. Rather than relying on large-scale point-cloud data with plentiful class labels and text pairs, we propose OV-3DET, which leverages advanced image/vision-language pre-trained models to achieve **Open-Vocabulary 3D point-cloud DETection**, as shown in Fig. 1. Notably, unlike previous open-vocabulary object detection/segmentation works [16, 31] that require supervisions of seen classes, inheriting general representations from image pre-trained models allows us to directly work around open-vocabulary point-cloud detection *without any 3D annotations of the bounding boxes and the corresponding class labels*.

Specifically, our proposed OV-3DET is a divide-and-conquer method that tackles two issues in two stages, respectively. In the first stage, the point-cloud detector learns to localize the unknown objects, and then in the second stage the point-cloud detector learns to name them according to the text prompts. Without cherry-picking seen categories and their ground truth, we directly take 2D pre-trained detectors, *e.g.*, Detic [34], to generate a series of 2D bounding boxes or 2D instance masks in the corresponding images. Note that we do not use any class labels predicted by 2D detectors, and we use the coarse 2D bounding boxes or 2D instance masks to supervise 3D point-cloud detectors to learn localizing 3D objects. Then in the second stage, we make novel use of CLIP by proposing a de-biased triplet cross-modal contrastive learning method to connect the modalities among point-cloud, image, and text, such that the point-cloud detector is able to relate the objects with corresponding text descriptions. During inference, only the point-cloud detector and the text prompts are used.

Experiments demonstrate the effectiveness of the proposed method, outperforming a wide range of open-vocabulary-related method [33] and generalizability-related methods [30, 34] over at least 7.47 mAP_{25} and 3.03 mAP_{25} on the SUN RGB-D dataset and ScanNet dataset, respectively. More importantly, we conduct sufficient ablation studies to explore how different components influence the performance of the method, and shed light on why the proposed method works.

2. Related Work

2.1. Open-Vocabulary 2D and 3D Detection.

Open-vocabulary (or known as zero-shot [31]) object detection targets to detect the novel classes that are not provided labels during training [10, 20, 21, 32], which usually accompanied by arbitrary text description [10]. Recent approaches [10, 19, 32] leverage image-text pairs to extract rich semantics from text, broadening the vocabulary size of the detector and advancing the embedding layer. Another

common solution is to replace the classifier with pre-trained vision-language embedding [20, 21, 31, 32] thus the detector can directly utilize the open-vocabulary classifier and perform open-vocabulary detection.

In the 3D point-cloud domain, it is not straightforward to obtain large-scale point-cloud text pairs, nor to directly exploit the vision-language embedding layer. Consequently, there has been limited exploration of open-vocabulary point-cloud detection in the point-cloud field. In particular, while images consist of dense RGB pixels, point-clouds are composed of sparse xyz points. This significant gap makes it challenging to transfer image or vision-language pre-trained models to point-clouds [30]. Point-CLIP [33] achieves open-vocabulary point-cloud recognition via projecting point-cloud into multi-view images and treating the images by CLIP. However, this approach is difficult to apply to point-cloud detection, as it does not address the issue of localizing unknown objects. Furthermore, we find that simply replacing the class embedding layer compromises performance in large-scale scenes. We propose to achieve open-vocabulary point-cloud detection by tackling the issues of localization and classification via image pre-trained models and de-biased triplet cross-modal contrastive learning.

2.2. Weakly-Supervised Detection.

Weakly-supervised object detection typically trains models using image-level class labels rather than full bounding box annotations [14, 15, 24, 25]. For instance, CASD [14] employs feature-level attention and self-distillation, while Li *et al.* [15] propose a generative adversarial segmentation module that collaborates with the conventional detection module. However, most weakly supervised methods rely on low-level proposal techniques, resulting in suboptimal localization quality.

In the first stage of our method, we utilize the 2D prediction results from the pre-trained 2D model as the pseudo label for the 3D point-cloud detector, which aims to enable the point-cloud detector to learn general representations that can localize diverse 3D bounding objects. This step represents a form of cross-domain, cross-modal weakly-supervised learning. The cross-domain aspect arises because the pre-trained 2D models have never been trained on images from the 3D dataset, resulting in domain gaps. For the cross-modal weakly-supervised learning, the 2D bounding boxes predicted by 2D pre-trained models serve as weak supervision for 3D point-cloud detectors. More importantly, in addition to teaching the point-cloud detector to localize objects, we further connect the point-cloud detector with a vision-language pre-trained model to enable open-vocabulary point-cloud detection.




Figure 2. OV-3DET learns to localize 3D objects from 2D pre-trained detectors. Specifically, for each predicted 2D box, we first extract the 3D points inside the frustum box, and then perform clustering on these points to remove background and outlier points. Finally, 3D pseudo box (center, size, and the orientation) is computed based on the remaining points.

2.3. Open-Set 3D Detection.

Open-set 3D detection [1, 3, 4, 28] is a similar but different setting, compared with our open-vocabulary 3D point-cloud detection. The former is to identify unknown objects from known ones. It does not classify each unknown object into specific categories. However, in our setting, we target at localizing and classifying each object with a specific bounding box by introducing text prompting that includes the unseen classes. In open-vocabulary 3D object detection, the model is not only required to be capable of learning general representation, but also needs to relate the point-cloud representation with text embedding.

3. Method

3.1. Overview

The proposed OV-3DET consists of two steps, where the first step deals with localization and the second step with classification. Specifically, the localization capability is learned from the 2D pre-trained models (e.g., Detic [34], Mask R-CNN [13], Fast R-CNN [9], etc.). We use the predicted 2D bounding boxes as the pseudo bounding box of the point-cloud detector after transforming the frustum into relatively tight bounding box according to the point-cloud geometry, as shown in Fig. 2. Regarding classification, we connect text, image, and point-cloud modalities via a de-biased triplet cross-modal contrastive learning, thus the model is able to recognize which localized objects belong to the description from the text. It is worth mentioning that throughout the pipeline, we only utilize existing well-established 2D pre-trained models and do not use any human annotations in the 3D point-cloud.

3.2. Notation and Preliminaries

We use \mathbf{T} , \mathbf{I} , \mathbf{P} to represent text, image and point-cloud, respectively, in which $\mathbf{I} \in \mathcal{R}^{3 \times H \times W}$, $\mathbf{P} = \{\mathbf{p}_i \in \mathcal{R}^3, i = 1, 2, 3, \dots, N\}$, where N is the point number in the point-cloud. During training, the unlabeled point-clouds dataset

with its paired image is used, denotes as $\mathcal{D}^{pc} = \{\mathbf{P}_j\}_{j=1}^{|\mathcal{D}^{pc}|}$, $\mathcal{D}^{img} = \{\mathbf{I}_j\}_{j=1}^{|\mathcal{D}^{img}|}$, respectively.

A point-cloud detector typically deals with both classification and localization, where the localization module predicts 3D bounding boxes $\hat{\mathbf{b}}_{3D} \in \mathcal{R}^7$ and we are able to obtain the corresponding point-cloud ROI features \mathbf{f}_{3D} . Similarly, the 3D bounding boxes can be projected into 2D, i.e., $\bar{\mathbf{b}}_{2D} \in \mathcal{R}^4$, via projection matrix K , and we can also index the corresponding patch features \mathbf{f}_{2D} . Then we can perform open-vocabulary classification by comparing between \mathbf{f}_{1D} and \mathbf{f}_{3D} , where \mathbf{f}_{1D} represents text feature [19].

3.3. Learn to Localize 3D Objects from 2D Pre-trained Detector

Localization is an important sub-task of detection. Yet it is not easy for point-cloud detector to learn localizing diverse objects, due to limited class annotations in the existing point-cloud dataset. In contrast, 2D detectors enjoy the advantages of rich data and annotations, already achieving detecting objects in a large variety [10, 16, 34]. To this end, we are motivated to make use of pre-trained 2D detectors to guide the 3D point-cloud detector to learn localization.

Specifically, for a pair of image and point-cloud from \mathcal{D}^{pc} and \mathcal{D}^{img} , 2D pre-trained detectors first predict a series of 2D bounding boxes or instance masks, if available. Afterwards, we back-project the 2D bounding box into 3D space, which results in the frustum 3D box that could not tightly enclose the 3D object, as shown in Fig. 2. To mitigate this issue, we shrink the 3D bounding box by leveraging the geometry of the point-cloud. Specifically, we first extract the 3D points inside the frustum box and then perform clustering on these points to remove background and outlier points. Finally, the center, size, and the orientation are computed based on the remaining points, which are further used to supervise the 3D detector, given by

$$L_{loc} = L_{box}^{3D}(\bar{\mathbf{b}}_{3D}, \hat{\mathbf{b}}_{3D}), \bar{\mathbf{b}}_{3D} = cluster(\bar{\mathbf{b}}_{2D} \circ K^{-1}), \quad (1)$$

where $\bar{\mathbf{b}}_{2D} \in \mathcal{R}^4$ represents the 2D bounding box predicted from the pre-trained 2D detector, and $\bar{\mathbf{b}}_{3D} \in \mathcal{R}^7$ is the 3D pseudo bounding box clustered from back-projected 2D bounding box, L_{box}^{3D} denotes the bounding box regression loss used in 3DETR [17]. Note that we do not use any classification prediction of the pre-trained 2D detectors by default as our main focus is enforcing the 3D detector to learn localizing from the pre-trained 2D detector at this stage.

3.4. Learn to Classify 3D Objects from 2D Pre-trained vision-language Model

Next, we aim at guiding the model to find the objects of interest from localized bounding boxes according to the text prompting. In recent 2D open-vocabulary detection, connecting the image modality and text modality

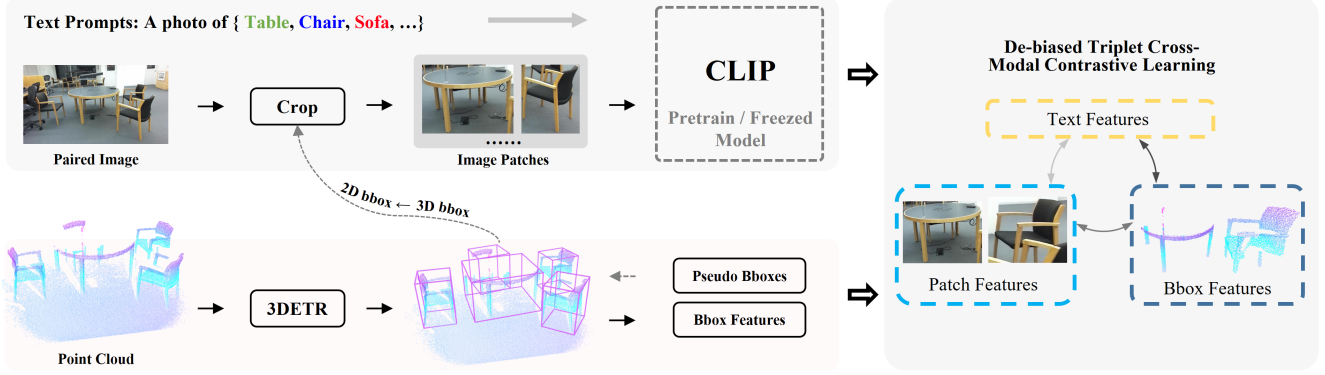


Figure 3. OV-3DET learns to classify 3D objects from CLIP pre-trained model. Specifically, for a paired image and point-cloud training input, we first localize objects and collect bounding box ROI features f_{3d} , and then crop the corresponding 2D patches from the image. After that, CLIP pre-trained model extracts image patch features f_{2d} and text features f_{1d} , and finally achieves open-vocabulary detection via connecting the embeddings of f_{1d} , f_{2d} and f_{3d} .




Figure 4. DTCC first classifies each image patch and 3D object via semantic propagation, and then samples positive and negative points in a de-biased manner.

is achieved by leveraging 2D pre-trained vision-language model [10, 16, 19]. Yet there is no such pre-trained model in point-cloud field. This further motivates us to leverage 2D pre-trained vision-language model [19] in open-vocabulary point-cloud detection. Inspired by recent works [30, 33] that connect image and point-cloud, we take image modality as the intermediary, and propose a **De-biased Triplet Cross-Modal Contrastive Learning (DTCC)** to connect text and point-cloud, which are intrinsically different modalities.

Specifically, during training, as shown in Fig. 3, we can get the point-cloud prediction as a series of 3D bounding boxes $\hat{\mathbf{b}}_{3D}$ with ROI features f_{3D} , then we crop the corresponding image patches by projecting the predicted 3D bounding box $\hat{\mathbf{b}}_{3D}$ based on the projection matrix K , which is further sent into a pre-trained vision-language model (*i.e.*, CLIP), with text prompts. Then we get the text features f_{1D} , and image patch features f_{2D} . Finally, we utilize **DTCC** to connects f_{1D} , f_{2D} and f_{3D} . It is noteworthy to mention that we directly reuse the hand-crafted prompts provided by CLIP which is originally designed for image classification. In practice, we just inject the category names into the blank position of the provided sentence.

De-biased Triplet Cross-Modal Contrastive Learning.

Before in-depth illustrating the proposed **DTCC**, we first revisit contrastive learning [5], which contrasts similar (positive) and dissimilar (negative) pairs of samples. In our cases, dissimilar (negative) samples are typically randomly sampled data points due to lack of annotation, and that may implicitly accept the false negative points, finally resulting in biased contrastive learning [6].

Based on this observation, **DTCC**, which belongs to De-biased Contrastive Learning [6], tries to correct the false-negative sampling. Specifically, as shown in Fig. 4a, **DTCC** first associates each 2D patch and 3D object with a distinct category (text prompts) via semantic propagation, and then perform negative sampling (as shown in Fig. 4b) across different categories. More specifically, semantic propagation utilizes vision-language pre-trained model to associate image patch with semantic label, and then propagate this semantic to 3D object according to the paired relationship between 2D patch and 3D box. The loss function of **DTCC** is given by

$$L_{DTCC} = L_{CL}^{P2T} + L_{CL}^{P2I} \quad (2)$$

where L_{DTCC} consists of two parts, L_{CL}^{P2T} , L_{CL}^{P2I} denote the contrastive loss between point-cloud and text / image, respectively. And the contrastive loss [18] is given by

$$L_{CL} = -\frac{1}{M} \sum_{i=1}^M \log \frac{\sum_{t=0}^m e^{h_i^\top h_t / \tau}}{\sum_{j=0}^M e^{h_i^\top h_j / \tau}}, \quad (3)$$

where h denotes the samples involved in contrastive learning, M is the number of samples, m is the number of positive samples corresponding to i^{th} sample h_i , and τ is a temperature hyper-parameter.

Note that in this stage, only the weights of 3D detector are trainable, while freezing CLIP. Therefore, there is no

Method		toilet	bed	chair	bathroom	sofa	dresser	scanner	fridge	lamp	desk	table	stand	cabinet	counter	bin	bookshelf	pillow	microwave	sink	stool	mean
Open-Vocabulary Baselines	OV-PointCLIP [33]	4.76	0.91	4.41	0.07	4.11	0.15	0.05	0.19	0.08	1.11	2.13	0.10	0.24	0.02	0.96	0.06	0.43	0.03	0.91	0.05	1.04
	3DETIC [34]	47.68	41.35	4.46	24.41	18.58	10.42	3.72	5.74	12.60	4.89	1.54	0.00	1.24	0.00	19.33	4.58	12.30	23.78	22.02	1.17	12.99
	CLIP-3D [19]	38.05	34.45	16.26	20.07	12.72	8.03	2.61	14.62	10.02	5.26	12.31	4.02	1.26	0.09	26.50	7.78	6.52	4.28	10.00	1.19	11.80
Fix-Vocabulary Methods	I2P-3D-DET [30]	47.98	50.93	6.40	24.58	17.63	3.75	0.08	2.36	1.75	2.32	5.32	0.13	0.45	0.06	1.90	0.11	4.95	0.65	3.87	0.04	8.80
	MN-3D-DET [29]	3.05	0.56	2.91	0.09	0.85	0.06	0.11	4.08	0.07	0.53	0.79	0.07	0.18	0.04	1.99	0.04	0.43	0.13	0.34	0.02	0.82
Upper Bound	3DETR [17]	89.62	82.08	65.91	74.20	57.06	24.49	12.49	24.43	24.94	28.17	49.74	59.71	18.18	28.86	43.76	31.58	19.45	10.07	31.61	13.55	39.50
Ours	OV-3DET	72.64	66.13	34.80	44.74	42.10	11.52	0.29	12.57	14.64	11.21	23.31	2.75	3.40	0.75	23.52	9.83	10.27	1.98	18.57	4.10	20.46
	Improvement	+24.66	+15.20	+18.54	+20.16	+23.52	+1.10	-3.43	-2.05	+2.04	+5.95	+11.00	-1.27	+2.14	+0.66	-2.98	+2.05	-2.03	-21.80	-3.45	+2.91	+7.47

Table 1. Detection results (AP_{25}) on SUN RGB-D.

Method		toilet	bed	chair	sofa	dresser	table	cabinet	bookshelf	pillow	sink	bathroom	refrigerator	desk	night stand	counter	door	curtain	box	lamp	bag	mean
Open-Vocabulary Baselines	OV-PointCLIP [33]	1.04	1.85	4.79	1.18	0.19	1.61	0.41	0.03	0.40	0.29	0.51	0.10	1.66	0.16	0.02	0.24	0.04	0.15	0.03	0.05	0.74
	3DETIC [34]	53.26	24.88	15.77	31.36	11.54	9.14	2.10	9.39	17.00	29.21	27.45	19.96	13.68	0.01	0.00	0.00	17.73	4.80	3.04	9.51	14.99
	CLIP-3D [19]	44.78	23.84	17.52	12.62	4.92	13.24	1.95	3.97	11.37	17.64	32.24	14.87	11.38	2.37	0.51	14.46	8.58	7.45	5.14	4.70	12.68
Fix-Vocabulary Methods	I2P-3D-DET [30]	22.97	16.42	2.77	9.97	0.53	1.34	0.32	0.03	3.84	0.24	7.20	0.15	5.91	0.11	0.00	0.08	0.10	0.32	0.08	0.29	3.63
	MN-3D-DET [29]	0.45	1.81	2.48	1.94	0.09	0.80	0.16	0.02	0.91	0.70	0.27	0.11	1.86	0.07	0.02	1.96	0.06	0.10	0.03	0.04	0.69
Upper Bound	3DETR [17]	90.75	62.97	67.56	68.42	34.29	53.48	38.90	44.22	43.66	62.56	77.27	43.67	52.54	50.61	28.15	44.63	36.53	11.96	17.66	9.13	46.95
Ours	OV-3DET	57.29	42.26	27.06	31.50	8.21	14.17	2.98	5.56	23.00	31.60	56.28	10.99	19.72	0.77	0.31	9.59	10.53	3.78	2.11	2.71	18.02
	Improvement	+4.03	+17.38	+9.52	+0.14	-3.33	+0.93	+0.88	-3.83	+6.00	+2.39	+24.04	-8.97	+6.04	-1.60	-0.20	-4.90	-7.20	-3.67	-3.03	-6.80	+3.03

Table 2. Detection results (AP_{25}) on ScanNet.

need to compute the contrastive loss between text and image. In other words, **DTCC** connects the embedding of text, image and point-cloud by pushing the point-cloud embedding towards that of the pre-trained vision-language model. The overall optimization objective is given by

$$L_{Connect} = L_{DTCC} + \lambda \cdot L_{loc}, \quad (4)$$

where L_{loc} ensures the localization does not degrade during cross-modal connection, and λ is a balance parameter.

4. Experiment

In this section, we compare the proposed OV-3DET with popular baselines on two widely used 3D detection datasets, SUN RGB-D [27] and ScanNetV2 [7]. Then we conduct sufficient analysis and ablation studies to explore why OV-3DET works.

4.1. Datasets and Evaluation Metric

3D Point-cloud Datasets. **SUN RGB-D** [27] and **ScanNetV2** [7] are two widely used 3D object detection datasets. As our method requires the paired point-cloud and image, we directly take the raw RGB-D frames from the ScanNet and SUN RGB-D dataset to conduct the experiments. Additionally, OV-3DET is evaluated on 20 ~ 200 common classes for both datasets.

The main metrics in the experiments are Average Precision (AP) and mean Average Precision (mAP) at IoU thresholds of 0.25, denoted as AP_{25} , mAP_{25} , respectively.

4.2. Main Results

Since there is no baseline directly working around open-vocabulary 3D point-cloud detection, we mainly compare

OV-3DET with well-known works [19, 30, 33, 34] that study either transfer-ability in 3D point-cloud or 2D open-vocabulary. Specifically, the **baselines** we used include:

1) *PointCLIP* [33] extends CLIP into 3D point-cloud and attempts to bridge the embedding space of text and point-cloud. However, it only deals with the point-cloud classification task. To address the localization problem, we use pseudo-box pre-trained 3DETR to locate objects, and then extract the points inside the predicted 3D bounding box, finally use the pre-trained PointCLIP to generate classification prediction for each predicted 3D bounding box. We term this baseline as *OV-PointCLIP*, which satisfies the definition of open-vocabulary 3D detection.

2) Detic [34] is a pioneer work that detects 21K classes in 2D image, which is used as the second baseline. We project the 2D prediction into 3D space with the same processing method to shrink the 3D bounding box frustum. This baseline is denoted as *3DETIC*. Note that *3DETIC* requires image and point-cloud pairs during inference, while the proposed *OV-3DET* only uses point-cloud during inference.

3) Similar to 3DETIC that locates and names 3D object directly using 2D pre-trained model, another straightforward baseline is using Detic [34] to locate 3D object, then using pre-trained CLIP [19] to perform open-vocabulary classification. This baseline is denoted as *CLIP-3D* and it satisfies the definition of open-vocabulary 3D detection setting. But *CLIP-3D* also requires image and point-cloud pairs during inference.

4) Image2Point [30] transfers the image pre-trained transformer to the point-cloud by copying or inflating the weights. Similarly, we copy the weights of the transformer from the pre-trained DETR [2] to 3DETR and fine-tune the

Setting	mAP_{25}	Transfer direction	mAP_{25}
OV-3DET	18.02	SUN RGB-D \rightarrow ScanNet	12.30
OV-3DET- <i>clean</i>	16.75	ScanNet \rightarrow SUN RGB-D	12.57

(a) Analysis of vocabulary overlap.

(b) Analysis of transferability.

Table 3. Analysis of vocabulary overlap and transferability.

set-aggregation module, 3D box and classification header. We refer to this baseline as *I2P-3D-DET*. This baseline can only classify predefined vocabulary sets and it is not an open-vocabulary 3D detector.

5) ModelNet [29] is a 3D classification dataset that contains 40 categories. Similar to OV-PointCLIP, we use pseudo-label pre-trained 3DETR to localize objects, and replace the classification head with a ModelNet pre-trained classifier. This baseline is denoted as *MN-3D-DET*. Due to the limited vocabulary of ModelNet, *MN-3D-DET* is not an open-vocabulary 3D detector.

Particularly, *OV-PointCLIP*, *3DETR* and *CLIP-3D* are open-vocabulary methods, *I2P-3D-DET* is a transferability-related method. Additionally, we provide the upper bound for OV-3DET, that is, *3DETR* trained with ground truth annotations in a fully supervised manner.

We evaluate all the aforementioned baselines and the proposed OV-3DET on both ScanNet and SUN RGB-D, AP_{25} and mAP_{25} on 20 common classes and report in Tabs. 1 and 2. The results reveal that our approach, which does not rely on paired images during inference, achieves at least 7.47% and 3.03% improvement in mAP_{25} on SUN RGB-D and ScanNet, respectively, compared to baselines that satisfy the open-vocabulary setting. Specifically, both *3DETR* and *CLIP-3D* surpass the other baselines and are comparable to each other. Moreover, the proposed OV-3DET not only transfers the open-vocabulary detection capability from images to point-clouds, but also outperforms these two baselines by a large margin. The reason is that the explicit connection between text and point-cloud brings the performance gains. Moreover, *OV-PointCLIP* and *MN-3D-DET* learn to classify point-cloud from ModelNet, but both approaches perform poorly due to the large gap between CAD handcrafted 3D models and point-clouds collected via RGB-D sensors. Furthermore, the results of *I2P-3D-DET* show that directly replicating learnable weights from pre-trained model on image to point-cloud detector does work, which further demonstrates the potentials of making use of 2D pre-trained models for the point-cloud field.

4.3. Analysis of Vocabulary Overlap

The connection between text and point-cloud relies on sampling text prompts as anchor points from the text embedding space, yet the sampling texts may lead to early leakage of vocabulary due to the unintentional overlap be-

Setting	mAP_{25}
OV-3DET- <i>rare</i>	9.90
OV-3DET- <i>common</i>	6.46
OV-3DET- <i>frequent</i>	16.05
OV-3DET- <i>all</i>	18.02

Table 4. Analysis of text prompts used in pseudo box generation.

tween training text prompts and test vocabulary. We therefore conduct the experiment to analyze the effect of vocabulary overlap.

Specifically, we use LVIS’s [11] categories as the text prompts during training. The relationship between LVIS and test category is shown in Fig. 6a. In this experiment, we remove the overlapping words from LVIS and dub the remaining words by the suffix *-clean*. Later, we re-train the OV-3DET-*clean* on ScanNet and present the results in Tab. 3a, which indicates that removing overlapping words results in a 1.27% drop in mAP_{25} , yet still outperforms the mAP_{25} of the baselines in Sec. 4.2, even though *3DETR* also suffers from the overlap issue.

4.4. Analysis of Transferability

The transferability of traditional 3D detectors is limited due to the difference between training and testing classes. Therefore, it is natural to ask whether the ScanNet pre-trained model can generalize to SUN RGB-D and vice versa. To answer this question, we evaluate OV-3DET with ScanNet and SUN RGB-D pre-trained model on SUN RGB-D and ScanNet, respectively. The results are reported in Tab. 3b, which demonstrates that cross dataset evaluation leads to 7.89% and 5.72% drop in mAP_{25} on SUN RGB-D and ScanNet, even though both SUN RGB-D and ScanNet are collected with RGB-D sensors, the resolution and noise of different sensors are considerably different, which results in the domain gap between these two datasets. In the presence of domain gap, the results further demonstrate the transferability of the proposed OV-3DET.

4.5. Analysis of the Influence of Text Prompts

Throughout OV-3DET, the usage of text prompts are threefold: 1) Pseudo 3D bounding box generation, 2) De-biased triplet cross-modal contrastive learning, 3) Evaluation. To further investigate the impact of text prompts, we conduct experiments for each usage of text prompts, and these analysis and ablation studies are performed on ScanNet by default.

Text Prompts Used in Pseudo Bounding Box Generation

In this experiment, we investigate the effect of the text prompts used in pseudo bounding box generation based on LVIS dataset. LVIS contains more than 1.2k categories




Figure 5. Results on the test category number vs. mAP.

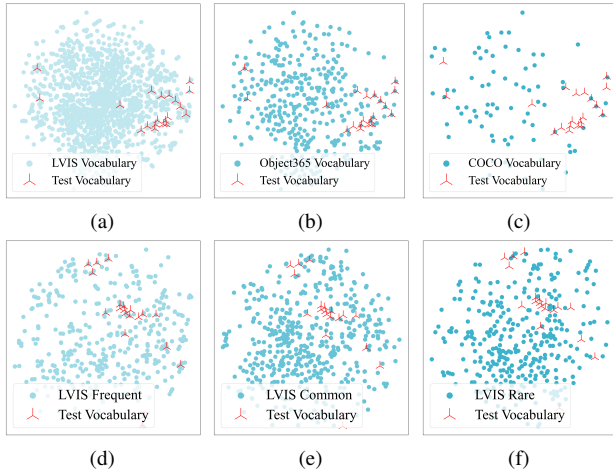


Figure 6. TSNE visualization to analyze the vocabulary overlap between test words and LVIS, Object365, COCO.

and is collected for long-tail object recognition, which can be divided into three groups: frequent, common, and rare based on the number of training images. The relationship between these three groups and test categories are illustrated in Figs. 6d to 6f. We re-train OV-3DET three times with pseudo-bounding boxes of frequent, common, and rare objects, respectively. We refer to these three re-trained models as OV-3DET-frequent, OV-3DET-common, and OV-3DET-rare. Results are presented in Tab. 4, compared with OV-3DET-all that uses all of the LVIS vocabulary to generate pseudo-bounding box, OV-3DET-frequent, OV-3DET-common, OV-3DET-rare drops 1.97%, 11.56%, 8.12% in terms of mAP_{25} , which demonstrate text prompts used in pseudo bounding box generation do make a difference. Moreover, OV-3DET-rare outperforms OV-3DET-common, as we can see in Figs. 6e and 6f, both LVIS-common and LVIS-rare are barely overlap with test classes, and we find that the boxes of LVIS-rare are more diverse regarding the objects in images in terms of scales, positions, and shapes, leading to more possibilities of localizing novel classes. Experimentally, we find the box recall of OV-3DET-common is 51.48%, and OV-3DET-rare is 58.72%.

Setting	mAP_{25}	Setting	mAP_{25}
OV-3DET-coco	11.75	OV-3DET-p2i	14.43
OV-3DET-lvis	18.02	OV-3DET-p2t	16.99
OV-3DET-365	19.01	OV-3DET-p2it	18.02

(a) Analysis of text prompts used in cross-modal connection.

(b) Ablation on de-biased contrastive learning.

Table 5. Analysis of prompts and de-biased contrastive learning.

Text Prompts Used in Cross-Modal Connection We have analysed in Sec. 4.3 the overlap between the text prompts used for cross-modal connection and test classes. Here, we perform a further ablation study on the sampling of text prompts and investigate the effect of the text prompts used in cross-modal connection. Specifically, the category of **Objects365v2** and **COCO** are used to replace **LVIS** and re-train OV-3DET. We refer to them as OV-3DET-lvis, OV-3DET-365, and OV-3DET-coco, and the relationships between these three sets of categories and test categories are illustrated in Figs. 6a to 6c. We report mAP_{25} in Tab. 5a, and the results show that OV-3DET-365 outperforms OV-3DET-lvis by 0.99%, indicating that the number of text prompts is not directly related to the final performance. In addition, the mAP_{25} of OV-3DET-coco drops severely, Fig. 6c provides an intuitive explanation that COCO under-samples the text embedding space and is insufficient for cross-modal connection.

Text Prompts Used in Evaluation In the above experiments, we evaluate OV-3DET on 20 common classes of indoor scenes. To further verify the effectiveness of OV-3DET, we resort to the vocabulary of ScanNet200 [23], and analyse the relationship between test vocabulary size and final performance. The results are illustrated in Fig. 5, as the test vocabulary size grows, mAP_{10} , mAP_{15} , and mAP_{25} gradually decrease and converge to 3.41%, 3.09% and 2.5%, respectively, when all of the classes of ScanNet200 are used in the evaluation. Due to the sparsity and incompleteness, 3D point-clouds are less discriminative than images, so as the test vocabulary size grows, it is inevitable that there will be many objects with similar shapes need to be distinguished, yet this is too hard even for ground truth supervised 3D detectors.

4.6. Ablation on De-biased Contrastive Learning

In this section, we focus on investigating the effect of the different strategy that used in the proposed de-biased cross-modal contrastive learning. Specifically, on the one hand, point-cloud can align with images using pair relationship, on the other hand, point-cloud can align with text via paired image and vision-language pre-trained model. A third option is to align the point-cloud with both the image and text




Figure 7. TSNE visualization to analyze the influence of the de-biased cross-modal contrastive learning.

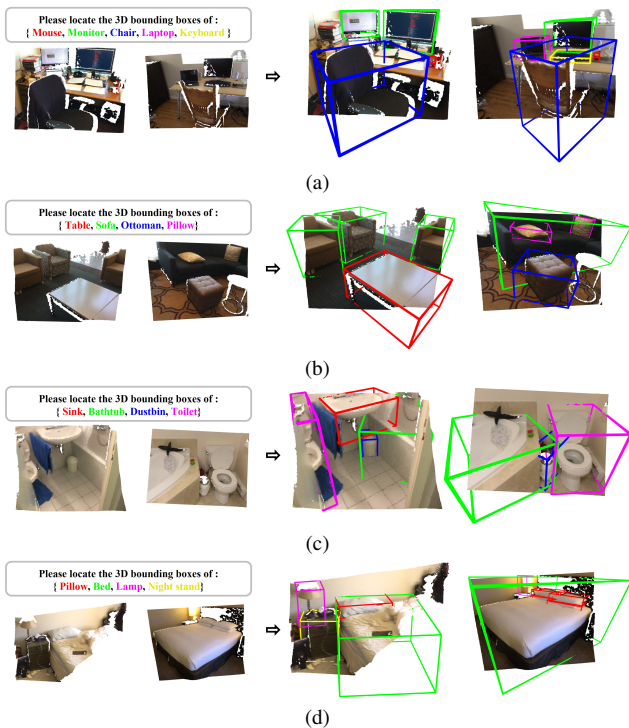


Figure 8. Qualitative results to visualize the 3D detection results. For each case, the text prompts and input point-clouds are shown on the left and the results are shown on the right. The colors of bounding boxes correspond to the classes in the text prompts.

simultaneously. We denote these three strategies (*point to image, point to text, and point to both image and text*) by *-p2i*, *-p2t*, and *-p2it*, respectively. Note that OV-3DET-*p2i* is biased contrastive learning, whereas OV-3DET-*p2t* and OV-3DET-*p2it* are de-biased contrastive learning. Results are shown in Tab. 5b, OV-3DET-*p2it* outperforms the other two strategies by at least 1.03%, which indicates that: 1) both OV-3DET-*p2i* and OV-3DET-*p2t* contribute to cross-modal connection. 2) OV-3DET-*p2t* (de-biased) is better than OV-3DET-*p2i* (biased).

To uncover why the de-biased contrastive learning works, we visualize the embedding of OV-3DET-*p2i*, OV-3DET-*p2t*, and OV-3DET-*p2it* via TSNE feature reduction. As shown in Fig. 7, our findings are: 1) CLIP pre-trained

model does not precisely connect image and text, as their features are clustered separately. 2) OV-3DET-*p2i* can connect image and point-cloud. 3) OV-3DET-*p2t* connects point-cloud with a local area of text embedding. 4) OV-3DET-*p2it* reaches a balance between text and image embedding. Regarding 3), SUN RGB-D and ScanNet only cover indoor scenes, while LVIS includes both indoor and outdoor objects, thus the connection between text and point-cloud should only occur in overlapping indoor objects. Regarding 4), since text and image are separated in the fixed CLIP embedding, and OV-3DET-*p2it* has to align point-cloud with both modalities, therefore, OV-3DET-*p2it* converges to a balance, which avoids biasing to either modality, thus leading to better results.

4.7. Qualitative Results

To further illustrate how OV-3DET works, we provide qualitative results in Fig. 8. For each case, we specify the detection targets by text prompts, and the figures under the text prompts are the input point-clouds, detection results are displayed to the right of the arrow, where the color of the predicted bounding box corresponds to the text prompt. The results demonstrate that OV-3DET is able to locate and name objects from small (mouse) to large (bed), and only the objects described in the text are returned. Further, if multiple instances of the same category appear in one point-cloud (pillows in Fig. 8d), all instances can be detected. Moreover, some objects are severely occluded and only partially visible, OV-3DET can detect them either (toilet in Fig. 8c). Nevertheless, we find that the orientation of predicted bounding box sometimes is not in line with human intuition (bathtub in Fig. 8c). Comparing the predicted bounding boxes, the incorrect orientations mostly appear on occluded or incomplete objects. We further examine the pseudo 3D bounding boxes and find the same problem. This phenomenon indicates that occlusions and incomplete points result in the failure of the orientation calculation.

5. Conclusion

In this paper, we propose an open-vocabulary point-cloud detector, dubbed OV-3DET, which is capable of localizing and naming 3D objects based on arbitrary text descriptions. More importantly, the training of OV-3DET does not require any 3D human annotations. Specifically, we achieve this by resorting to well-established 2D pre-trained detectors and vision-language models. More specifically, we first learn to localize 3D objects from 2D pre-trained detectors, and then classify the detected objects by connecting text and point-cloud embeddings. Experiments on ScanNet [7] and SUN RGB-D [27] demonstrate the effectiveness of the OV-3DET. We hope that this work will inspire more future research to further explore this area.

References

- [1] Antonio Alliegro, Francesco Cappio Borlino, and Tatiana Tommasi. 3dos: Towards 3d open set learning-benchmarking and understanding semantic novelty detection on point clouds. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. 3
- [2] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European conference on computer vision*, pages 213–229. Springer, 2020. 5
- [3] Jun Cen, Peng Yun, Junhao Cai, Michael Yu Wang, and Ming Liu. Open-set 3d object detection. In *2021 International Conference on 3D Vision (3DV)*, pages 869–878. IEEE, 2021. 3
- [4] Jun Cen, Peng Yun, Shiwei Zhang, Junhao Cai, Di Luan, Michael Yu Wang, Ming Liu, and Mingqian Tang. Open-world semantic segmentation for lidar point clouds. In *ECCV*. Springer, 2022. 3
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR, 2020. 4
- [6] Ching-Yao Chuang, Joshua Robinson, Yen-Chen Lin, Antonio Torralba, and Stefanie Jegelka. Debaised contrastive learning. *Advances in neural information processing systems*, 33:8765–8775, 2020. 4
- [7] Angela Dai, Angel X. Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proc. Computer Vision and Pattern Recognition (CVPR)*, IEEE, 2017. 5, 8
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 1
- [9] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 3
- [10] Xiuye Gu, Tsung-Yi Lin, Weicheng Kuo, and Yin Cui. Zero-shot detection via vision and language knowledge distillation. *arXiv e-prints*, pages arXiv–2104, 2021. 2, 3, 4
- [11] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5356–5364, 2019. 6
- [12] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 1
- [13] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 3
- [14] Zeyi Huang, Yang Zou, BVK Kumar, and Dong Huang. Comprehensive attention self-distillation for weakly-supervised object detection. *Advances in neural information processing systems*, 33:16797–16807, 2020. 2
- [15] Xiaoyan Li, Meina Kan, Shiguang Shan, and Xilin Chen. Weakly supervised object detection with segmentation collaboration. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9735–9744, 2019. 2
- [16] Matthias Minderer, Alexey Gritsenko, Austin Stone, Maxim Neumann, Dirk Weissenborn, Alexey Dosovitskiy, Aravindh Mahendran, Anurag Arnab, Mostafa Dehghani, Zhuoran Shen, et al. Simple open-vocabulary object detection with vision transformers. *arXiv preprint arXiv:2205.06230*, 2022. 2, 3, 4
- [17] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2906–2917, 2021. 3, 5
- [18] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 4
- [19] 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 *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 1, 2, 3, 4, 5
- [20] Shafin Rahman, Salman Khan, and Nick Barnes. Improved visual-semantic alignment for zero-shot object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11932–11939, 2020. 2
- [21] Shafin Rahman, Salman H Khan, and Fatih Porikli. Zero-shot object detection: joint recognition and localization of novel concepts. *International Journal of Computer Vision*, 128(12):2979–2999, 2020. 2
- [22] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10684–10695, 2022. 1
- [23] David Rozenberszki, Or Litany, and Angela Dai. Language-grounded indoor 3d semantic segmentation in the wild. *arXiv preprint arXiv:2204.07761*, 2022. 7
- [24] Yunhang Shen, Rongrong Ji, Yan Wang, Zhiwei Chen, Feng Zheng, Feiyue Huang, and Yunsheng Wu. Enabling deep residual networks for weakly supervised object detection. In *European Conference on Computer Vision*, pages 118–136. Springer, 2020. 2
- [25] Yunhang Shen, Rongrong Ji, Yan Wang, Yongjian Wu, and Liujuan Cao. Cyclic guidance for weakly supervised joint detection and segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 697–707, 2019. 2
- [26] Haoyu Song, Li Dong, Wei-Nan Zhang, Ting Liu, and Furu Wei. Clip models are few-shot learners: Empirical studies on vqa and visual entailment. *arXiv preprint arXiv:2203.07190*, 2022. 1

- [27] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 567–576, 2015. [5](#), [8](#)
- [28] Kelvin Wong, Shenlong Wang, Mengye Ren, Ming Liang, and Raquel Urtasun. Identifying unknown instances for autonomous driving. In *Conference on Robot Learning*, pages 384–393. PMLR, 2020. [3](#)
- [29] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. [5](#), [6](#)
- [30] Chenfeng Xu, Shijia Yang, Tomer Galanti, Bichen Wu, Xiangyu Yue, Bohan Zhai, Wei Zhan, Peter Vajda, Kurt Keutzer, and Masayoshi Tomizuka. Image2point: 3d point-cloud understanding with 2d image pretrained models. In *European Conference on Computer Vision*, pages 638–656. Springer, 2022. [2](#), [4](#), [5](#)
- [31] Mengde Xu, Zheng Zhang, Fangyun Wei, Yutong Lin, Yue Cao, Han Hu, and Xiang Bai. A simple baseline for zero-shot semantic segmentation with pre-trained vision-language model. *arXiv preprint arXiv:2112.14757*, 2021. [2](#)
- [32] Alireza Zareian, Kevin Dela Rosa, Derek Hao Hu, and Shih-Fu Chang. Open-vocabulary object detection using captions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14393–14402, 2021. [2](#)
- [33] 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. *arXiv preprint arXiv:2112.02413*, 2021. [2](#), [4](#), [5](#)
- [34] Xingyi Zhou, Rohit Girdhar, Armand Joulin, Phillip Krähenbühl, and Ishan Misra. Detecting twenty-thousand classes using image-level supervision. *arXiv preprint arXiv:2201.02605*, 2022. [1](#), [2](#), [3](#), [5](#)