

CLIP²: Contrastive Language-Image-Point Pretraining from Real-World Point Cloud Data

Yihan Zeng^{1*}, Chenhan Jiang^{2*}, Jiageng Mao³, Jianhua Han¹, Chaoqiang Ye¹,
Qingqiu Huang¹, Dit-Yan Yeung², Zhen Yang¹, Xiaodan Liang⁴, Hang Xu^{1†}

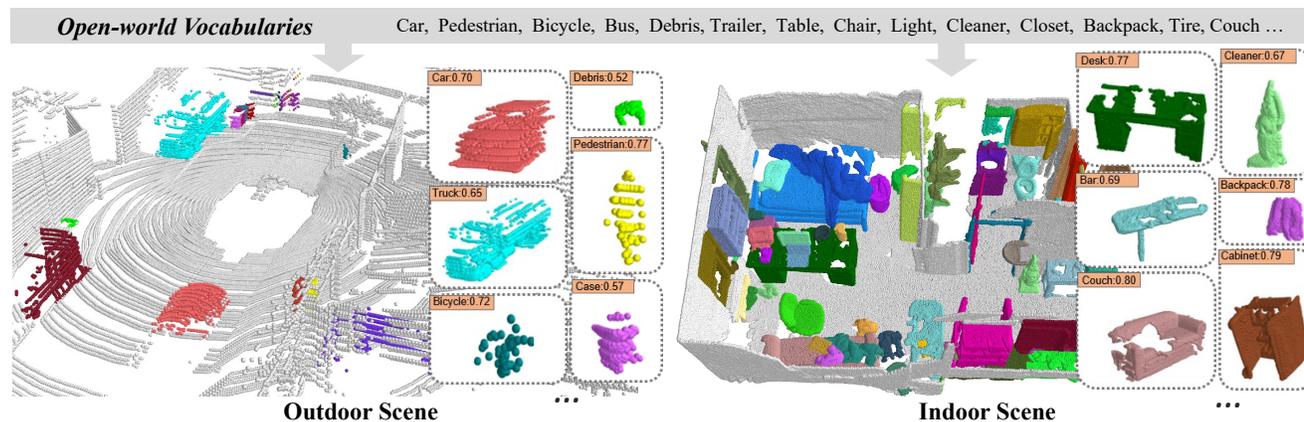


Figure 1. **Illustration of our open-world recognition results.** Benefiting from our CLIP², the 3D representation is aligned to the open-world language representation, which enables flexible zero-shot transfer. Best viewed in colors.

Abstract

Contrastive Language-Image Pre-training, benefiting from large-scale unlabeled text-image pairs, has demonstrated great performance in open-world vision understanding tasks. However, due to the limited Text-3D data pairs, adapting the success of 2D Vision-Language Models (VLM) to the 3D space remains an open problem. Existing works that leverage VLM for 3D understanding generally resort to constructing intermediate 2D representations for the 3D data, but at the cost of losing 3D geometry information. To take a step toward open-world 3D vision understanding, we propose *Contrastive Language-Image-Point Cloud Pretraining (CLIP²)* to directly learn the transferable 3D point cloud representation in realistic scenarios with a novel proxy alignment mechanism. Specifically, we exploit naturally-existed correspondences in 2D and 3D scenarios, and build well-aligned and instance-based text-image-point proxies from those complex scenarios. On top of that, we propose a cross-modal contrastive objective to learn semantic and instance-level aligned point cloud representa-

tion. Experimental results on both indoor and outdoor scenarios show that our learned 3D representation has great transfer ability in downstream tasks, including zero-shot and few-shot 3D recognition, which boosts the state-of-the-art methods by large margins. Furthermore, we provide analyses of the capability of different representations in real scenarios and present the optional ensemble scheme.

1. Introduction

Powerful 3D point cloud representation plays a crucial role in various real-world applications, e.g., 3D object recognition and detection [10, 20, 31, 40, 44]. Compared to 2D images, 3D point cloud provides specific information like accurate geometry that is robust to illumination changes. However, current methods [25, 40] that learn 3D representations generally rely on the predefined number of object categories and require plenty of labor-intensive annotations. Those learned 3D representations are insufficient for safety-critical scenarios like self-driving which includes a long-tail class distribution far beyond the predefined taxonomy. Therefore, it is highly demanded to learn a transferable 3D representation equipped with zero-shot recognition ability in vocabulary scalable real-world scenes. Figure 1 shows an open-world recognition example by our CLIP² in

*Equal contribution. ¹Huawei Noah's Ark Lab ²Hong Kong University of Science and Technology ³The Chinese University of Hong Kong ⁴Sun Yat-san University [†]Corresponding Author: xu.hang@huawei.com

outdoor and indoor scenes, where the 3D objects can be classified with the correlation alignment between 3D representations and open-world vocabularies.

The critical ingredient of open-world understanding is that the models learn sufficient knowledge to obtain general representations. To achieve this, recent Vision-Language Models (VLM) [14, 27, 38] leverage Internet-scale text-image pairs to conduct vision-language pretraining, which facilitates transferable 2D representation and demonstrates promising performance in 2D open-vocabulary tasks. However, 3D vision-language pretraining remains unexplored due to the limitation of existing 3D datasets in diversity and scale compared to the massive data sources in 2D counterparts [14, 15, 27, 38]. Though some recent works [12, 13, 43] try to avoid this problem by transferring the pretrained 2D VLM into the intermediate representation including projected image patches [12, 18] or depth maps [1, 43], those representations suffer from the loss of 3D geometric information and limited viewpoints under realistic scenarios. Especially the camera images are only sometimes available due to the sensor failure in 3D scenes. We believe the 3D representation based on original point cloud data retains most information and is the optimal solution for 3D real world understanding, which requires a rethink of learning the transferable 3D representation under realistic scenarios.

To this end, we propose a **Contrastive Language-Image-Point cloud Pretraining** framework, short for CLIP², which directly aligns 3D space with broader raw text and advances the 3D representation learning into an open-world era. Our learning process can be decomposed into two stages: **Firstly**, we introduce a *Triplet Proxy Collection* to alleviate the limitation of accessible pretraining data by constructing language-image-point triplets from real-world scenes. Since the large-scale realistic 3D datasets for outdoor driving [2, 19] and indoor scenarios [9, 32] are collected in open-world, it contains huge amounts of realistic objects that vary in semantics and diversity. Thus we consider them as potential pretraining data sources without extra human supervision. Specifically, we propose “Proxy” instances as the bridges between language descriptions, 2D images and 3D point clouds. Enabled by a well-aligned VLM, a scalable caption list and the geometry transformation between 2D and 3D, we automatically create more than 1 million triplets to facilitate pretraining. **Secondly**, we further propose a *Cross-Modal Pretraining* scheme to jointly optimize the feature space alignments of three modalities, *i.e.* point cloud, language and image. It contains both the contrastive learning objective of semantic-level text-3D correlation and instance-level image-3D correlation, which contributes to better transferability of learned 3D representation.

We study the transferable capability of CLIP² by benchmarking the zero-shot recognition performance on four popular indoor and outdoor real-world datasets, and find a significant improvement over current methods, achieving Top1 accuracy 61.3% on SunRGBD [32], 43.8% on ScanNet [9]),

28.8% on nuScenes [2] and 56.0% on ONCE [19]. For a fair comparison with existing methods [1, 13, 36, 43], we conduct zero-shot and few-shot classification on single object dataset ScanObjectNN [34] and find consistent dominance, 16.1% relative improvement on zero-shot classification over previous state-of-the-art method [13]. To validate the vocabulary-increasing ability of CLIP², we report the quantity results and visualizations to show the improved discovery of the long-tail categories. Moreover, we make ablations and analysis on different representations, and investigate ensembling alternatives to merge complementary knowledge of all available representations in realistic applications. Our contributions can be summarized as follows:

- We propose a novel CLIP² framework that aligns 3D space with open-world language representation, facilitating zero-shot transfer in realistic scenarios.
- We present a Triplet Proxies Collection scheme in real-world scenes, which alleviates the shortage of text-3D data sources and facilitates the pretraining methods.
- CLIP² jointly optimizes the correlation alignment between point cloud, language and image by proposed cross-modal pretraining mechanism, which enhances the transferability of learned 3D representation.
- Our CLIP² achieves the state-of-the-art zero-shot transfer performance on 5 datasets (indoor/outdoor scenes and single-object) and shows quality results on vocabulary-increasing discovery in real world.

2. Related Work

Vision-Language Model. Large vision language models (VLM) [14, 15, 27, 38] have demonstrated successful performance in downstream zero-shot tasks with the learned transferable 2D representations. CLIP [27] and ALIGN [14] push the limit by collecting Internet-scale image-text pairs and then learning the correlation alignment between image and language feature space with contrastive pretraining objectives. Those models can be directly transferred to zero-shot 2D recognition and achieve impressive results. Recent DetClip [38] learns to align image patches to test phrases after pretraining under hybrid supervision from detection, grounding and image-text pair data, which extends the ability to localize open-vocabulary 2D proposals in images. In this paper, we attempt to transfer the open-vocabulary ability of pre-trained VLM to the 3D domain, making language applicable to zero-shot point cloud recognition.

Zero-shot/Open-world Learning in 3D. Recognizing 3D objects with a large vocabulary is necessary for safety-critical autonomous driving and robotic tasks, yet remains under-explored. Cheraghian et al. [5–8] first attempt to associate PointNet [25] feature with category semantic information via a projection function, and separately proposed an unsupervised skewness loss [5] to mitigate the hubness

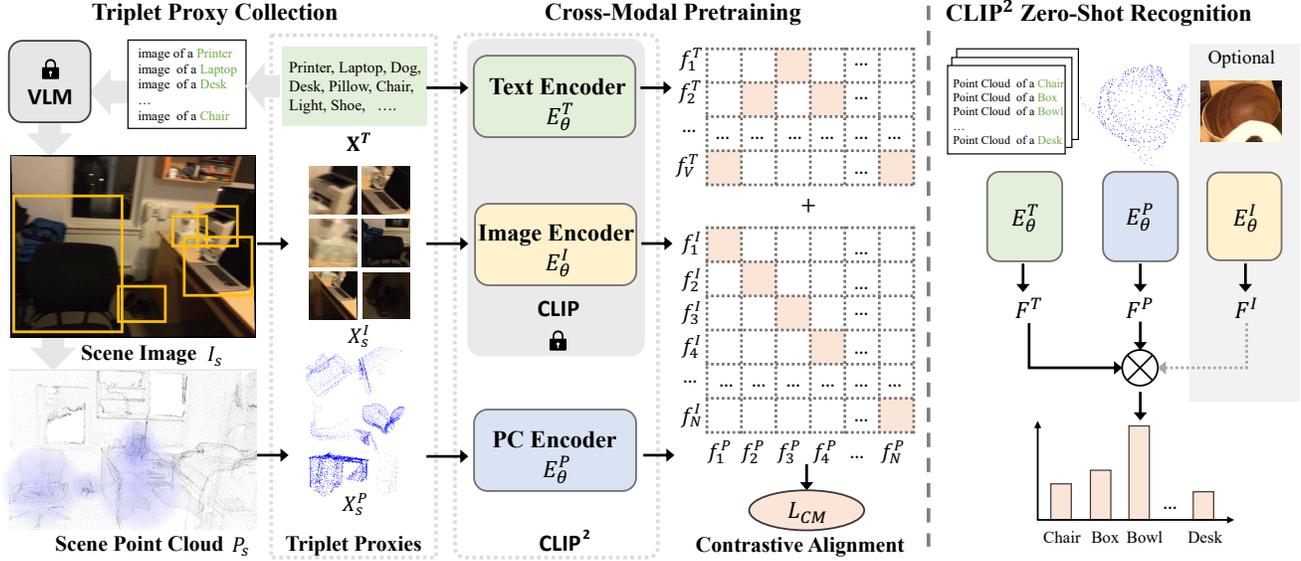


Figure 2. **Overview of CLIP² framework.** The main components contain two parts, the *Triplet Proxy Collection* and the *Cross-Modal Pretraining*. The defined Triplet Proxy set $\mathcal{D}_{\text{proxy}}$ consists of language captions \mathbf{X}^T , corresponding image instances \mathbf{X}^I and raw 3D point cloud instances \mathbf{X}^P , which come from the free data source under realistic scenarios without any labeling labor. On top of that, we pretrain a point cloud encoder E^P with the cross-modal contrastive learning objective. Equipped with CLIP², the learned 3D point cloud representation F^P is well aligned to the language representation, which facilitates downstream zero-shot 3D transfer tasks in the real world.

problem. The transductive case [6] is discussed in which extends [5] using a triplet loss. Notably, the above works conduct experiments on synthetic datasets and need to divide datasets into “seen” categories as training data and “unseen” categories as testing data. Thus they are not suitable for realistic scenarios due to the domain gap between synthetic and real-world data, as well as the limited vocabulary-increasing ability. Recently, inspired by the success of VLMs [14, 27] in 2D tasks, some works [13, 43] propose to transfer the zero-shot recognition ability of pretrained CLIP [27] into 3D area. PointCLIP [43] directly projects point cloud into multi-view depth maps as image-like data input for pretrained CLIP to make classification predictions. While CLIP2Point [13] trains an image-depth embedding on ShapeNet [39] to better align the depth representation to the pretrained image space of CLIP. However, depth maps lost plenty of geometry information of the original point cloud data structure, resulting in poor performance especially in realistic scenarios. By contrast, we aim to learn transferable 3D representation based on the original point cloud data structure in realistic scenarios.

3D Representation Learning. Much progress has been made in learning a comprehensive 3D representation in an unsupervised manner. Most works [1, 16, 17, 23, 30, 37, 41, 42] follow the paradigm that conducts pretraining on unlabeled datasets and then finetunes on the limited downstream annotations. Though the improved transferability of 3D representation, they can not be directly transferred to zero-shot tasks with open-world vocabularies. In this work, we con-

duct language-image-point cloud pretraining, which learns transferable 3D representation aligned to open-vocabulary language space to facilitate the zero-shot transfer.

3. Method

In this section, we introduce CLIP² to learn a transferable 3D point cloud representation with arbitrary category recognition ability under realistic scenarios, illustrated in Figure 2. We will first present the *Triplet Proxy Collection* in Section 3.1, which utilizes a pretrained VLM and geometric transformation to obtain language-image-point triplets from real-world scenes. Then we will elaborate *Cross-Modal Contrastive Pretraining* mechanism in Section 3.2, which jointly optimizes the alignment correlations between language, image and point cloud feature space.

3.1. Triplet Proxy Collection

Inspired by the significant performance of 2D VLMs on open-vocabulary tasks, we aim to develop 3D vision-language pretraining to facilitate category-increasing capacity for real-world scenarios. However, the core challenge is the shortage of pretraining data. Compared to the 2D vision-language pretraining framework CLIP [27], which takes more than 400M image-language pairs from the Internet, the largest 3D single-object dataset ShapeNet [39] only contains 50K CAD models with 55 categories. In addition to the insufficiency of data scale, pretraining on such synthetic data fails to transfer well in the real world due to the huge domain gap. Enlightened by the recent emer-

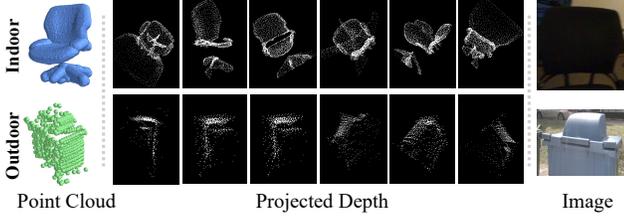


Figure 3. **Illustration of three representation modals** of two 3D objects examples under indoor and outdoor scenarios.

gence of large-scale point cloud datasets collected in indoor [9, 32] and outdoor scenarios [2, 19], we observe that those naturally-collected datasets potentially contain vast amounts of open-world objects that vary in semantics and diversity. Considering the data collection itself is cheap except for laborious annotation, we novelly take leverage of those available datasets without human annotations as a practical yet effective pretraining data source.

Specifically, given the realistic scene data $\mathcal{S} = \{(P_s, I_s)_{s=1}^{|\mathcal{S}|}\}$, where $P_s \in \mathbb{R}^{N_P \times 3}$ and $I_s \in \mathbb{R}^{N_I \times H \times W \times 3}$ are corresponding 3D point clouds and images of scene s , we propose a novel concept, *Proxy*, as the bridge between language, image and 3D point cloud. As illustrated in Figure 2, equipped by those proxy instances, we can automatically collect a massive number of language-image-point cloud pairs $\mathcal{D}_{\text{proxy}}$ in the format of proxies under open-world scenes. We detail the process as follows.

Language Proxy. We set the language proxies $\mathbf{X}^T \in \mathbb{R}^V$ as a raw text list from the 2D open-world dataset [11], where $V = 1206$ denotes the vocabulary size of language proxies.

Image Proxy. Next, we obtain the image proxies \mathbf{X}^I by an open vocabulary detector DetCLIP [38], denoted as M , which is trained with open-world data and performs open-set detection. Concretely, given language proxies \mathbf{X}^T and input scene image I_s , we extract corresponding image proposals as image proxies X_s^I with M by the similarity between input language embeddings and proposal features as

$$\{X_s^I\}_{s \in |\mathcal{S}|} = M(\{I_s\}_{s \in |\mathcal{S}|}, \mathbf{X}^T). \quad (1)$$

3D Proxy. We exploit the naturally-existed geometry relations between 2D and 3D scenes to obtain 3D proxies \mathbf{X}^P , which consists of point cloud instances corresponding to image proposals in \mathbf{X}^I . We simplify the geometry transformation as $G(\cdot)$ and formulate the relations as:

$$X_i^P = G(X_i^I). \quad (2)$$

Detailedly, for *indoor scenes* equipped with RGB-D sensors, we first remove the background pixels by unsupervised segmentation algorithm [28] for each image proxy $X_{s,i}^I$, $i \in |X_s^I|$. Since depth information is known, we then transform the segmented pixels from uvd coordinate $X_{s,i}^{I,uvd} \in \mathbb{R}^{n,3}$ to xyz coordinate $X_{s,i}^{P,xyz} \in \mathbb{R}^{n,3}$ as a 3D

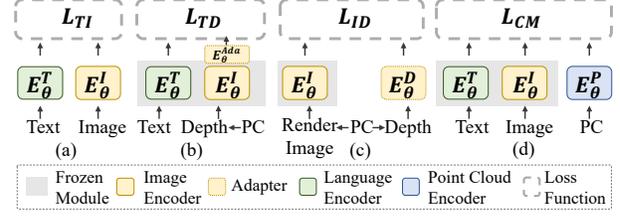


Figure 4. **Comparison of different pretraining strategies.** (a) CLIP aligns image and language embedding space [27] as L_{TI} based on large-scale text-image pairs. (b) PointClip [27] aligns projected depth map to CLIP language space as L_{TD} . (c) Clip2Point aligns depth map to CLIP image space as L_{ID} . (d) our CLIP² aligns original 3D point cloud to both CLIP language space and image space via cross-modal objective L_{CM} .

point cloud proxy with the given camera parameters. For *outdoor scenes* captured by LiDAR sensors, we first create a 3D frustum for each image proxy by extruding the 2D image proposal into 3D space following [22, 24]. Then we conduct DBSCAN algorithm [29] within frustum and select the point cloud cluster as the point proxy $X_{s,i}^{P,xyz}$.

Eventually, we construct Triplet Proxy $\mathcal{D}_{\text{proxy}} = \{\mathbf{X}^T, \mathbf{X}^I, \mathbf{X}^P\}_{s=1}^{|\mathcal{S}|}$ by combining corresponding language proxies \mathbf{X}^T , image proxies \mathbf{X}^I and 3D proxies \mathbf{X}^P , where $\mathbf{X}^I = \{X_s^I\}_{s=1}^{|\mathcal{S}|}$ and $\mathbf{X}^P = \{X_s^P\}_{s=1}^{|\mathcal{S}|}$. 220K and 1.4M proxy triplets are formed for indoor and outdoor scenes, respectively. More details can be found in the appendix.

3.2. Cross-Modal Contrastive Pretraining

With the triplet proxies $\mathcal{D}_{\text{proxy}}$, a straightforward pretraining objective is forcing the alignment between the embedding spaces of point cloud X_i^P and language X_i^T from scratch. However, it might not promise good transferability of learned representation, since the number of language-image-point pretraining data triplets remains two orders of magnitude smaller than the language-image pairs adopted by CLIP [27] and the vocabulary size is much more limited. Therefore, we design to learn the correlation alignment based on the pretrained embedding space of CLIP.

The comparison of current pretraining strategies [13, 43] is illustrated in Figure 4, which is a series of 3D variants of CLIP. Notably, both existing methods exploit projected depth map as the intermediate representation of point cloud, which are respectively learned to align to language space [43] and image space [13]. Intuitively, as illustrated in Figure 3, depth representation lost plenty of geometry information compared to the original point cloud, especially in outdoor scenarios. Moreover, images are sometimes unavailable for 3D objects. Thus we conduct pretraining on original 3D point cloud data as an optimal representation.

Toward learning more transferable representation, we introduce a cross-modal contrastive learning objective to jointly optimize the correlation alignment across lan-

guage, image and point cloud, including *Semantic-Level Language-3D Alignment* and *Instance-Level Image-3D Alignment*. Specifically, the overall architecture of CLIP², shown in Figure 2, contains language encoder E_θ^T , point cloud encoder E_θ^P and visual encoder E_θ^I , which respectively embed the triplet proxies into text feature $f^T \in \mathcal{R}^{1 \times C_T}$, point cloud feature $f^P \in \mathcal{R}^{1 \times C_P}$ and image feature $f^I \in \mathcal{R}^{1 \times C_I}$, where C is the embedding dimension.

Semantic-Level Language-3D Alignment. In order to inherit the open-world recognition ability from pretrained CLIP [27], we align the point cloud feature f^P with text embedding f^T from well-trained CLIP with Language-Point Proxy $\{X_i^T, X_i^P\}$ input. We replace *classname* in the prompts, like “point cloud of a $\{classname\}$.” with raw text in proxy X_i^T as language sentences. The core idea is to drive the feature centroids of 3D instances and the corresponding text prompt closer. We compute the loss function of between of language proxy and point cloud proxy as:

$$l(i, T, P) = -\log \frac{\exp(f_i^T \cdot f_i^P / \tau)}{\exp(f_i^T \cdot f_i^P / \tau) + \sum_{j \in N, X_j^T \neq X_i^T} \exp(f_i^T \cdot f_j^P / \tau)}, \quad (3)$$

where N is the mini-batch size, τ is the temperature coefficient. Within a training mini-batch, the language-3D alignment objective $L(T, P)$ can be described as:

$$L(T, P) = \frac{1}{N} \sum_{i \in N} l(i, T, P). \quad (4)$$

Instance-Level Image-3D Alignment. In addition to the alignment between semantic language and 3D proxy instances, we further introduce the contrastive alignment between instance-wise image proxy and 3D proxy instances. Note that the instance-aware visual concept has been well-studied in the embedding space of pretrained CLIP. We believe instance-sensitive learning can contribute to further correlation learning and benefits to the transferability of learned 3D representation. The contrastive aligned objective $L(I, P)$ across point cloud and image is formulated as:

$$l(i, I, P) = -\log \frac{\exp(f_i^I \cdot f_i^P / \tau)}{\exp(f_i^I \cdot f_i^P / \tau) + \sum_{j \in N, j \neq i} \exp(f_i^I \cdot f_j^P / \tau)}, \quad (5)$$

$$L(I, P) = \frac{1}{N} \sum_{i \in N} l(i, I, P). \quad (6)$$

Finally, we obtain the resultant cross-modal contrastive learning objective $L_{CM}(T, I, P)$ as the combination of $L(T, P)$ and $L(I, P)$, where both alignments of semantic-level text-3D correlation and instance-level image-3D correlation are injected:

$$L_{CM}(T, I, P) = \lambda_1 L(T, P) + \lambda_2 L(I, P), \quad (7)$$

where the hyper-parameters λ_1 and λ_2 are both set to 0.5.

4. Experiment

In this section, we evaluate CLIP² on realistic indoor and outdoor scenarios. We report the zero-shot transfer results on various datasets [2, 9, 32–34] and make further analysis on the designs of pretraining strategy.

4.1. Zero-shot Transfer

Setting. After pretraining, natural language is applied to reference learned 3D representation to enable following zero-shot transfer tasks. **(i) Zero-Shot Recognition:** we evaluate zero-shot recognition performance for realistic objects, where K category names are transferred to text prompt “point cloud of $\{CLASS\}$ ” to encode the text features $F_K \in \mathbb{R}^{K \times C}$. Then the classification logits are calculated with the 3D feature f^P and text features as:

$$\text{logits}_i = \text{softmax}(f_i^P (F_K)^T). \quad (8)$$

We present the results under both indoor and outdoor scenarios in Table 1, Table 2 and Table 5, as well as the object-level benchmark in Table 6. **(ii) Open-vocabulary recognition:** we enlarge the category vocabularies of ScanNet to 249 and 384 to study the open-vocabulary recognition ability in Table 3. **(iii) Open-vocabulary localization:** we study the open-vocabulary localization ability by localizing open-world 3D objects with our proxy generation process and then classifying them with our learned 3D representation, of which the visualization is illustrated in Figure 5 and evaluation results are reported in Table 4. Notably, we investigate representation ensembling alternatives to enable knowledge merging of all available representations for realistic applications, illustrated in Table 8.

4.1.1 Indoor Scenarios

Datasets and details. We adopt the widely used indoor 3D dataset SUN RGB-D [32] as the realistic indoor scenario that provides pretraining data source, a single-view RGB-D dataset consisting of $\sim 10K$ scenes. To validate the transferability of learned 3D representation, we also evaluate another popular indoor 3D dataset ScanNet [9], which contains $\sim 1.5K$ scenes of 3D reconstructed meshes. We remove objects in ScanNet with less than 5 points, leaving 384 noisy categories. For open-vocabulary recognition, we evaluate performance on the ScanNet 384-class set and a 249-class merged set. In addition to the scene-wise indoor dataset, we conduct evaluations on ScanObjectNN [34], which collects $\sim 3K$ individual realistic objects with 15 categories and is applied in the previous zero-shot evaluation [13, 43]. During the proxy collection process, we empirically set $\epsilon = 0.3$ in [38] as a tradeoff between filtering FPs and preserving TPs to generate image proxies. Considering the occurrence frequencies of different indoor categories vary a lot, we adopt the class balance strategy [4] to mitigate the class imbalance. During pretraining process, we adopt [26] as point cloud encoder and set the overall training epoch number to 100.

Method	Avg.	Bed	Bookshelf	Chair	Desk	Sofa	Table	toilet	Bathhub	Dresser	Night Stand
PointClip [43]	11.5	0.0	94.0	0.0	0.0	0.0	14.7	0.0	0.0	6.1	0.0
Clip2Point [13]	18.6	10.9	20.6	64.3	34.4	13.8	14.1	26.2	0.0	1.4	0.0
PointClip [43] w/ TP.	38.0	45.3	100.0	62.5	48.5	44.4	4.8	55.2	16.3	3.3	0
Clip2Point [13] w/ TP.	56.9	78.0	87.6	36.2	36.6	64.7	37.4	82.1	77.5	67.6	1.2
CLIP ²	61.3	84.0	75.5	70.7	47.3	75.5	33.8	86.2	65.3	71.8	2.4
CLIP ² w/ En.	69.6	87.3	94.3	70.7	54.1	79.3	47.0	91.7	85.7	82.2	3.6

Table 1. **Zero-shot recognition results in SUN RGB-D.** **Avg.:** the mean average Top1 accuracy across all categories. **TP.:** our Triplet Proxy set. **En.:** our optional ensembling scheme for inference.

Method	Avg.	Cab	Bed	Chair	Sofa	Tabl	Door	Wind	Bksf	Pic	Cntr	Desk	Curt	Fridg	Bath	Showr	Toil	Sink
PointClip [43]	6.3	0.0	0.0	0.0	0.0	0.7	0.0	0.0	91.8	0.0	0.0	0.0	15.0	0.0	0.0	0.0	0.0	0.0
Clip2Point [13]	24.9	0.0	20.8	85.1	43.3	26.5	69.9	0.0	20.9	1.7	31.7	27.0	0.0	1.6	46.5	0.0	22.4	25.6
PointClip [43] w/ TP.	26.1	55.7	0.0	72.8	5.0	5.1	1.7	0.0	77.2	0.0	0.0	51.7	0.3	0.0	0.0	40.3	85.3	49.2
Clip2Point [13] w/ TP.	35.2	3.0	11.8	45.1	27.6	10.5	61.5	2.6	71.9	0.3	33.6	29.9	4.7	11.5	72.2	92.4	86.1	34.0
CLIP ²	38.5	67.2	32.6	69.3	42.3	18.3	19.1	4.0	62.6	1.4	12.7	52.8	40.1	9.1	59.7	41.0	71.0	45.5

Table 2. **Zero-shot recognition in ScanNet.** **Avg.:** the mean average Top1 accuracy across all categories. **TP.:** our Triplet Proxy set.

Method	Top5 Acc.		Method	IN			OUT	
	384 cls.	249 cls.		mAP ₂₅	AR ₂₅	mIoU	P.	R.
[43]	0.3	0.4	[21]	1.3	-	-	-	-
[13]	6.4	7.0	[18]	13.0	37.7	-	-	-
CLIP ²	22.0	31.7	CLIP ²	12.7	43.0	52.8	6.7	87.4

Table 3. Recognition results of vocabulary expansion in ScanNet. Table 4. **Zero-shot localization results.** **IN:** indoor scenario SUN RGB-D. **OUT:** outdoor scenario nuScenes.

Quantity results. For zero-shot recognition task, we take two recent works as our baselines, *i.e.* PointClip [43] and Clip2Point [13], which study the zero-shot classification task on 3D object-level benchmarks [34, 35] by leveraging pretrained CLIP with projected depth maps. Focusing on the real-world scenarios, we conduct comparison not only on the realistic object-level [34] as illustrated in Table 6 but also on the scene-level datasets shown in Table 1 and Table 2, where the evaluation follows the common classes split in [10, 44] and reports the instance Top1 accuracy of each class. As shown in tables, our CLIP² can outperform baselines on all benchmarks by large margins. Besides, we apply our triplet proxy generation mechanism (TP.) to baseline methods, and achieve considerable improvements on SUN RGB-D and ScanNet by 26.5% and 19.8% for PointClip, 38.3% and 10.3% for Clip2Point. On the one hand, the contrasts demonstrate the effectiveness of our triplet proxies for open-world understanding. On the other hand, our learned 3D representation is superior in 3D object recognition by retaining more 3D-specific information than depth representation. Besides, we present the optional ensembling scheme (En.) when camera images are available, which can take advantage of multi-modal knowledge and further boost the performance by 8.3%. To further validate the open-vocabulary recognition ability, we conduct evaluation on a larger category set of ScanNet in Table 3 and report the instance Top5 accuracy, which illustrates the superior-

ity of our CLIP² when vocabulary increases. Beyond that, CLIP² is also equipped with zero-shot 3D localization ability by proxy generation. On indoor scenario SUN RGB-D, we compare with a SOTA indoor 3D detector 3DETR [21] and a recent work OV3D [18] that studies open-vocabulary detection, where evaluation is conducted on the same “unseen” split in [18]. Since CLIP² do not fit the tight bounding boxes of point cloud instances, we estimate the maximum bounding box of proxies and GT instances to conduct evaluation following the same metrics mAP₂₅ and AR₂₅ in [18], as shown in Table 4. Notably, compared to baseline works that train on “seen” 3D annotations and test on “unseen” categories, we have no access to any 3D annotations yet achieve comparable localization ability, which yields 5.3% AR₂₅ improvement over OV3D [18]. We further evaluate segmentation results in Table 4.

Quality results. The visualization results of CLIP² under a indoor scene of SUN RGB-D [32] is shown in Figure 5(a). Our triplet proxy generation process can localize open-world 3D objects in a point cloud scene. Moreover, the 3D representation learned from our cross-modal pre-training can provide more accurate classification results for 3D instances by exploiting original point cloud, which corrects the mistaken “People” prediction in image to “Picture” by considering the geometry information.

4.1.2 Outdoor Scenarios

Datasets and details. We exploit a prevalent large-scale 3D dataset nuScenes [2] as the outdoor data source and extra validate the performance on the ONCE dataset [19]. The nuScenes dataset consists of \sim 28K frames with 10 categories, while ONCE contains 6 annotated sequences with 5 categories. Similarly, we set the $\epsilon = 0.3$ for image proxies collection and adopt the class balance strategy [4].

Quantity results. Since the outdoor point cloud is collected by LiDAR sensors, it has a wider perception range

Method	Avg.	nuScenes										ONCE				
		Car	Truck	Bus	Ped.	Bicycle	Trailer	C.V.	Motor.	Barrier	T.C.	Car	Cyc.	Ped.	Truck	Bus
PointClip [43]	11.7	0.0	0.0	0.0	29.1	41.8	3.4	0.0	0.1	42.5	0.0	0.0	13.8	79.2	0.2	7.61
Clip2Point [13]	12.4	0.4	0.3	0.1	13.5	31.0	1.8	9.3	1.5	66.2	0.3	17.3	11.8	95.4	35.7	4.0
PointClip [43] w/ TP.	28.3	18.8	0	5.5	74.0	17.9	57.0	1.9	4.5	2.1	29.7	51.8	9.2	99.8	5.0	46.8
Clip2Point [13] w/ TP.	33.0	26.7	16.8	51.2	45.2	15.8	13.9	20.0	5.7	10.5	34.2	39.4	27.8	95.5	40.6	51.7
CLIP ²	37.8	41.9	41.3	22.5	40.3	21.1	20.6	24.8	22.4	17.3	35.3	52.7	27.3	77.7	78.5	44.0

Table 5. **Zero-shot recognition results in outdoor scenario:** nuScenes (Left) and ONCE (Right). **TP.:** our Triplet Proxy set. **Avg.:** the mean average Top1 accuracy across all categories of two benchmarks.

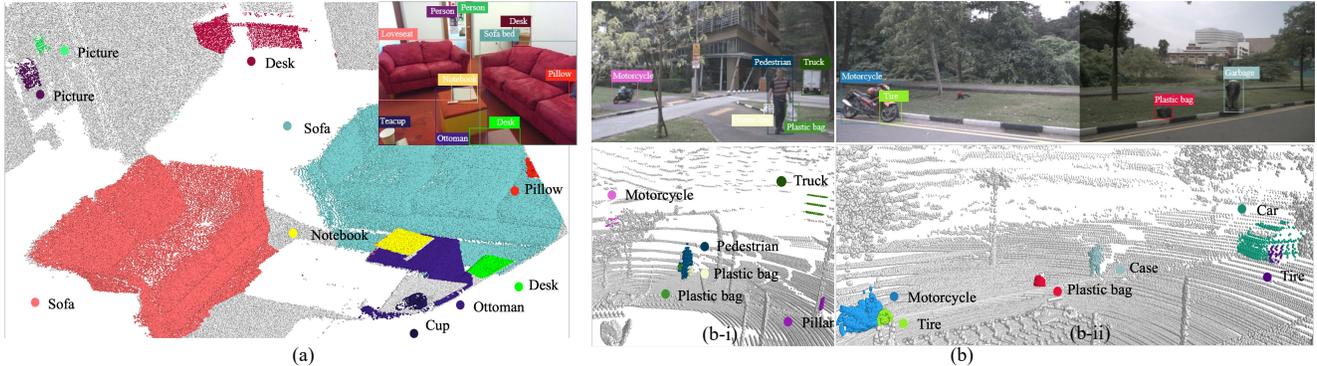


Figure 5. **Visualizations of the zero-shot localization and recognition results** by CLIP² under open-world (a) indoor realistic scene [32] and (b) outdoor scenes [2]. Notably, the whole pipeline of CLIP² not only has no access to human annotations, but also enables the open-world vocabularies beyond groundtruth annotations, such as 'Picture' in (a) and 'Plastic bag', 'Tire' in (b). Best viewed in colors.

Method	ZS	15-way			10-way	
		4-S	8-S	16-S	10-S	20-S
PointNet++ [25]	-	41.0	47.6	55.0	-	-
PointCLIP [43]	15.4	46.0	50.0	55.6	-	-
Clip2Point [13]	23.3	-	-	-	-	-
CrossPoint [1]	-	-	-	-	58.7±1.8	64.6±1.2
CLIP ²	39.1	51.3	59.6	62.5	60.6±2.5	66.3±3.2

Table 6. **Zero-shot and Few-shot classification results on ScanObjectNN.** ZS: zero-shot. K-way N-shot: few-shot settings.

than RGB-D but leads to sparse distribution. Thus the projected depth representation of baselines results in severer information lost, as illustrated in the second row in Figure 3. As shown in Table 5, our CLIP² considerably outperforms the baseline recognition results by more than 20%, and our triplet proxies respectively boost two baselines by 9.5% and 4.8%. Additionally, we evaluate the localization ability on the outdoor scenario nuScenes in Table 4. Due to the lack of works that tackle outdoor open-vocabulary localization problems, we follow classic detection accuracy metrics Precision(P.) and Recall(R.) as evaluation metrics. Specifically, we calculate the center distance between groundtruth bounding boxes and our 3D proxies that are predicted to belong to the same category as the groundtruth, and set the distance threshold as $\lambda = 2m$. For those matched pairs that are closer than λ , we count the proxies as TPs. Otherwise, for those unmatched proxies and groundtruth, we count them as FPs and FNs respectively, thus $P.=\frac{TPs}{TPs+FPs}$, $R.=\frac{TPs}{TPs+FNs}$. As shown in Table 4, our CLIP² pipeline can provide high

recall for outdoor objects. Since CLIP² is highly sensitive to open-world objects and can perceive categories beyond groundtruth list, it tends to create overmuch predictions thus the precision is comparably low. The perception ability of open-world objects can be viewed in Figure 5(b).

Quality results. We show off two outdoor scenes of nuScenes [2] in Figure 5(b-i) and Figure 5(b-ii). In addition to perceiving those common categories, our CLIP² surprisingly localizes and recognizes those uncommon 3D objects in 3D scenes such as the tires of vehicles, the plastic bag in the hand of pedestrian as well as the plastic bag on the road. We believe it contributes to auto-driving safety by providing the localization and recognition of universal obstacles to facilitate follow-up driving decisions.

4.2. Few-shot Classification

Setting. Lightweight few-shot learning is practical for application by finetuning the pretraining model with given limited data annotations, which can also validate the generalization capability of our learned representation. To make a fair comparison, we follow the existing methods [1, 43] to conduct experiments under “K-way N-shot” setting on the challenging realistic object-level dataset ScanObjectNN [34], where we randomly sample N point cloud objects from each of the randomly selected K classes.

Quantity results. As illustrated in Table 6, we compare with representative 3D networks PointNet++ [26], the recent zero-shot approach PointClip [43] as well as a state-

	Rep.	Obj.	Avg._IN	Avg._OUT	Avg._OBJ
(a)	Depth	Lang.-Depth	38.0	9.3	31.8
(b)	Depth	Image-Depth	56.4	21.7	39.0
(c)	PC.	Lang.-Point	52.6	26.8	37.7
(d)	PC.	Image-Point	56.5	24.4	37.8
(e)	PC.	Lang.-Image-Point	61.3	28.8	39.4

Table 7. Ablations on representation learning. **Rep.:** the recognition representations. **Obj.:** the learning objectives. **PC.:** point cloud. **Lang.:** language space.

of-the-art representation learning method CrossPoint [1], which conducts contrastive pretraining between point cloud and rendered images on CAD dataset ShapeNet [3]. As we can see, with a slight number of samples, our CLIP² can boost the classification results by a large margin, exceeding PointClip by 5.3%, 9.6% and 6.9% with 4, 8 and 16 shots. Besides, we outperform CrossPoint with considerable gain, illustrating our pretraining strategy on collected proxies can learn sufficient knowledge from realistic open world to generate transferable 3D representation, which is superior to the pretraining on a small-scale synthetic dataset.

4.3. Ablations and Analysis

Ablations on representation learning. To observe the transferability of different representations and the effect of different learning objectives, we conduct ablations and report the mean average Top1 accuracy across all classes of zero-shot recognition in indoor scenarios [32] (Avg._IN), outdoor scenarios [2] (Avg._OUT) and object-level benchmark [34] (Avg._OBJ), which is shown in Table 7. Firstly, for fair comparisons, we follow [13, 43] to project input point cloud into depth maps in N_V different views as alternative representation. Secondly, we adopt various objectives to learn different correlation alignments across language, image and point cloud feature space or depth space. Specifically, comparing (a) and (b), aligning depth space to image space yields better transfer performance due to the similar data structure of image and depth map. In (c) and (d), point cloud representation is better when aligning to image space in indoor scenes, while better to align with language space in outdoor scenes due to the data discrepancy between image-like RGB-D points and sparse LiDAR points. Generally, 3D point cloud representation outperforms depth representation in all benchmarks due to preserving the complete 3D structure and sufficient 3D-specific information. Comparing (e) and (d), the joint alignment between three feature spaces contributes to the best 3D point cloud representation transferability on all benchmarks.

Analysis of representation ensembling. Intuitively, different representations contain different perspectives of knowledge, which can be potentially merged to achieve the optimum results during inference. To validate the ensembling application, we adopt three optional representation modals, *i.e.* point clouds, projected depth maps and cor-

	PC.	Image	Depth	Avg._IN	Avg._OUT	Avg._OBJ
(i)	✓			61.3	28.8	39.4
(ii)		✓		64.2	41.1	-
(iii)			✓	56.9	23.9	39.0
(f)	✓	✓		68.7	43.9	-
(g)	✓		✓	64.8	30.4	43.2
(h)	✓	✓	✓	69.6	42.3	-

Table 8. Analysis on the representation ensembling schemes.

responding image patches, where depth representation is trained on our proxies and image representation is generated from pretrained image branch of CLIP. We ensemble their predicted logits by simple summation as the final output, and illustrate the separate recognition results and ensembling performance in Table 8. Benefiting from the sufficient knowledge learned from massive CLIP training data, image representation presents best performance in separate applications. By merging the complementary knowledge, our 3D representation leads to gains of 4.6% indoors [32] and 2.8% outdoors. Though further improving the indoor recognition performance with 0.9% when merged, depth representation yields 1.6% drop for outdoor objects, illustrating the information lost especially for outdoor scenarios. Since image representation is sometimes missing, such as in [34], our 3D representation is more robust for 3D applications.

5. Limitation

As a pilot work for the language-3D pretraining problem, though CLIP² enables zero-shot localization and recognition with proposed triplet proxy generation and learned transferable 3D representation, it can not provide the accurate tight bounding box for open-world 3D objects as a common detector does. We believe CLIP² can facilitate the development of open-world 3D detectors by introducing the recognition ability to general 3D detectors or providing presented 3D proxies to enable further training of 3D detectors.

6. Conclusion

In this paper, we present a novel contrastive language-image-point cloud pretraining framework, CLIP², which consists of a triplet proxy collection scheme and a cross-modal contrastive learning mechanism. Based on the observation that realistic scenarios contain a massive amount of open-world objects, we innovatively propose to collect triplet proxies from realistic scenes as pretraining data. We then conduct cross-modal contrastive alignment across language, image and point cloud feature space to learn transferable 3D representation. The zero-shot transfer results on various indoor and outdoor benchmarks validate the ability of CLIP² for 3D open-world understanding.

Acknowledgements We acknowledge the support of MindSpore¹, CANN, and Ascend AI Processor in this work.

¹<https://www.mindspore.cn/>

References

- [1] Mohamed Afham, Isuru Dissanayake, Dinithi Dissanayake, Amaya Dharmasiri, Kanchana Thilakarathna, and Ranga Rodrigo. Crosspoint: Self-supervised cross-modal contrastive learning for 3d point cloud understanding. In *CVPR*, 2022. [2](#), [3](#), [7](#), [8](#)
- [2] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nusscenes: A multi-modal dataset for autonomous driving. In *CVPR*, 2020. [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [3] Angel X Chang, Thomas Funkhouser, Leonidas Guibas, Pat Hanrahan, Qixing Huang, Zimo Li, Silvio Savarese, Manolis Savva, Shuran Song, Hao Su, et al. Shapenet: An information-rich 3d model repository. *arXiv preprint arXiv:1512.03012*, 2015. [8](#)
- [4] Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019. [5](#), [6](#)
- [5] Ali Cheraghian, Shafin Rahman, Dylan Campbell, and Lars Petersson. Mitigating the hubness problem for zero-shot learning of 3d objects. In *BMVC*, 2019. [2](#), [3](#)
- [6] Ali Cheraghian, Shafin Rahman, Dylan Campbell, and Lars Petersson. Transductive zero-shot learning for 3d point cloud classification. In *WACV*, 2020. [2](#), [3](#)
- [7] Ali Cheraghian, Shafin Rahman, Townim F Chowdhury, Dylan Campbell, and Lars Petersson. Zero-shot learning on 3d point cloud objects and beyond. *IJCV*, 2022. [2](#)
- [8] Ali Cheraghian, Shafin Rahman, and Lars Petersson. Zero-shot learning of 3d point cloud objects. In *MVA*. IEEE, 2019. [2](#)
- [9] 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. [2](#), [4](#), [5](#)
- [10] Zhipeng Ding, Xu Han, and Marc Niethammer. Votenet: A deep learning label fusion method for multi-atlas segmentation. In *MICCAI*, 2019. [1](#), [6](#)
- [11] Agrim Gupta, Piotr Dollar, and Ross Girshick. Lvis: A dataset for large vocabulary instance segmentation. In *CVPR*, pages 5356–5364, 2019. [4](#)
- [12] Huy Ha and Shuran Song. Semantic abstraction: Open-world 3d scene understanding from 2d vision-language models. In *CoRL*, 2022. [2](#)
- [13] 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. *arXiv preprint arXiv:2210.01055*, 2022. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [14] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*. PMLR, 2021. [2](#), [3](#)
- [15] Junnan Li, Ramprasaath Selvaraju, Akhilesh Gotmare, Shafiq Joty, Caiming Xiong, and Steven Chu Hong Hoi. Align before fuse: Vision and language representation learning with momentum distillation. *NeurIPS*, 2021. [2](#)
- [16] Hanxue Liang, Chenhan Jiang, Dapeng Feng, Xin Chen, Hang Xu, Xiaodan Liang, Wei Zhang, Zhenguo Li, and Luc Van Gool. Exploring geometry-aware contrast and clustering harmonization for self-supervised 3d object detection. In *ICCV*, 2021. [3](#)
- [17] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. In *ECCV*, 2022. [3](#)
- [18] Yuheng Lu, Chenfeng Xu, Xiaobao Wei, Xiaodong Xie, Masayoshi Tomizuka, Kurt Keutzer, and Shanghang Zhang. Open-vocabulary 3d detection via image-level class and de-biased cross-modal contrastive learning. *arXiv preprint arXiv:2207.01987*, 2022. [2](#), [6](#)
- [19] Jiageng Mao, Minzhe Niu, Chenhan Jiang, Hanxue Liang, Jingheng Chen, Xiaodan Liang, Yamin Li, Chaoqiang Ye, Wei Zhang, Zhenguo Li, et al. One million scenes for autonomous driving: Once dataset. *arXiv preprint arXiv:2106.11037*, 2021. [2](#), [4](#), [6](#)
- [20] 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. [1](#)
- [21] Ishan Misra, Rohit Girdhar, and Armand Joulin. An end-to-end transformer model for 3d object detection. In *ICCV*, 2021. [6](#)
- [22] Anshul Paigwar, David Sierra-Gonzalez, Özgür Ercent, and Christian Laugier. Frustum-pointpillars: A multi-stage approach for 3d object detection using rgb camera and lidar. In *ICCV*, 2021. [4](#)
- [23] 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. [3](#)
- [24] Charles R Qi, Wei Liu, Chenxia Wu, Hao Su, and Leonidas J Guibas. Frustum pointnets for 3d object detection from rgb-d data. In *CVPR*, 2018. [4](#)
- [25] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *CVPR*, 2017. [1](#), [2](#), [7](#)
- [26] 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. [5](#), [7](#)
- [27] 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. [2](#), [3](#), [4](#), [5](#)
- [28] Carsten Rother, Vladimir Kolmogorov, and Andrew Blake. ” grabcut” interactive foreground extraction using iterated graph cuts. *TOG*, 2004. [4](#)
- [29] Erich Schubert, Jörg Sander, Martin Ester, Hans Peter Kriegel, and Xiaowei Xu. Dbscan revisited, revisited: why and how you should (still) use dbscan. *TODS*, 42(3):1–21, 2017. [4](#)
- [30] Charu Sharma and Manohar Kaul. Self-supervised few-shot learning on point clouds. *NeurIPS*, 2020. [3](#)

- [31] Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. Pointcnn: 3D object proposal generation and detection from point cloud. In *CVPR*, 2019. 1
- [32] Shuran Song, Samuel P Lichtenberg, and Jianxiong Xiao. Sun rgb-d: A rgb-d scene understanding benchmark suite. In *CVPR*, 2015. 2, 4, 5, 6, 7, 8
- [33] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *CVPR*, 2020. 5
- [34] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *ICCV*, 2019. 2, 5, 6, 7, 8
- [35] Kashi Venkatesh Vishwanath, Diwaker Gupta, Amin Vahdat, and Ken Yocum. Modelnet: Towards a datacenter emulation environment. In *2009 IEEE Ninth International Conference on Peer-to-Peer Computing*. IEEE, 2009. 6
- [36] Hanchen Wang, Qi Liu, Xiangyu Yue, Joan Lasenby, and Matt J Kusner. Unsupervised point cloud pre-training via occlusion completion. In *ICCV*, 2021. 2
- [37] Saining Xie, Jiatao Gu, Demi Guo, Charles R Qi, Leonidas Guibas, and Or Litany. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *ECCV*, 2020. 3
- [38] Lewei Yao, Jianhua Han, Youpeng Wen, Xiaodan Liang, Dan Xu, Wei Zhang, Zhenguo Li, Chunjing Xu, and Hang Xu. Detclip: Dictionary-enriched visual-concept paralleled pre-training for open-world detection. In *NeurIPS*, 2022. 2, 4, 5
- [39] Yangyang Ye, Houjin Chen, Chi Zhang, Xiaoli Hao, and Zhaoxiang Zhang. Sarpnet: Shape attention regional proposal network for lidar-based 3d object detection. *Neurocomputing*, 2020. 3
- [40] Tianwei Yin, Xingyi Zhou, and Philipp Krähenbühl. Center-based 3d object detection and tracking. *arXiv preprint arXiv:2006.11275*, 2020. 1
- [41] 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. 3
- [42] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Pointm2ae: Multi-scale masked autoencoders for hierarchical point cloud pre-training. *arXiv preprint arXiv:2205.14401*, 2022. 3
- [43] 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. 2, 3, 4, 5, 6, 7, 8
- [44] Zaiwei Zhang, Bo Sun, Haitao Yang, and Qixing Huang. H3dnet: 3d object detection using hybrid geometric primitives. In *ECCV*, 2020. 1, 6