

Multi-modal Relation Distillation for Unified 3D Representation Learning

Huiqun Wang^{1,2*}, Yiping Bao³, Panwang Pan³, Zeming Li³, Xiao Liu³
Ruijie Yang^{1,2} and Di Huang^{1,2†}

¹SKLSDE, Beihang University, Beijing, China

²IRIP Lab, SCSE, Beihang University, Beijing, China

³PICO

Abstract. Recent advancements in multi-modal pre-training for 3D point clouds have demonstrated promising results by aligning heterogeneous features across 3D shapes and their corresponding 2D images and language descriptions. However, current straightforward solutions often overlook intricate structural relations among samples, potentially limiting the full capabilities of multi-modal learning. To address this issue, we introduce **M**ulti-modal **R**elation **D**istillation (**MRD**), a tri-modal pre-training framework, which is designed to effectively distill reputable large Vision-Language Models (VLM) into 3D backbones. MRD aims to capture both intra-relations within each modality as well as cross-relations between different modalities and produce more discriminative 3D shape representations. Notably, MRD achieves significant improvements in downstream zero-shot classification tasks and cross-modality retrieval tasks, delivering new state-of-the-art performance.

1 Introduction

Recently, 3D shape understanding has garnered increasing attention due to its wide range of applications, such as space calculation [2, 27, 30, 42], autonomous driving [23, 53] and robotic perception [3, 46]. Despite notable advances in 3D visual analysis, the limited availability of 3D data, characterized by constraints in scale and scarcity of annotations, remains a significant barrier.

To tackle this challenge, many researchers have delved into integrating auxiliary modalities within the self-supervised learning framework. Some attempts utilize priors from the image modality to craft more instructive pretext tasks [9, 58], thereby enhancing the discriminative power of the learned 3D representations. Meanwhile, others focus on distilling knowledge from pre-trained models in either image or text modalities to facilitate conceptual understanding within the 3D modality [57, 60].

Among these endeavors, tri-modal-based methodologies [16, 26, 49, 50, 59] have shown exceptional prowess by aligning point cloud representations with the pre-aligned image-text feature space. They employ a synthesized set of triplet data

* For this joint project, he is also an intern with PICO ARCH; † Corresponding author (dhuang@buaa.edu.cn)

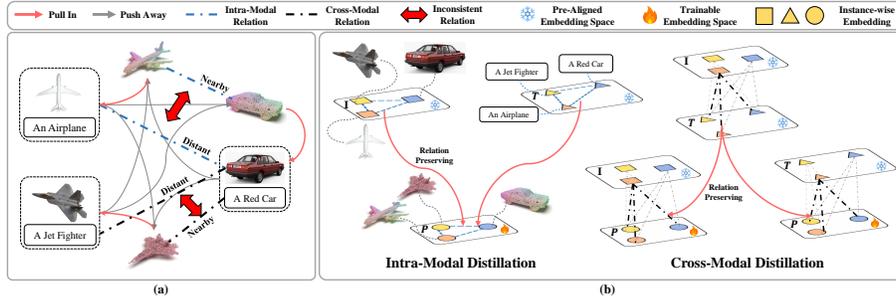


Fig. 1: Illustration of Multi-modal Relation Distillation (MRD). (a) Conventional contrastive learning focuses on instance-level alignment but disrupts the intra-modality and cross-modality relations established in previous image-text alignment. For example, the nearby and distant relations between the three samples are disturbed in the 3D modality due to naive alignment. (b) MRD distills structural knowledge from both intra- and cross-modality mutual relations, aiming to preserve the semantic relations in the pre-aligned embedding spaces, thereby delivering more discriminative and coherent distributions. Zoom in for better view.

to draw point cloud representations closer to their respective image-text pairs, simultaneously distancing them from non-associated examples. The strategic use of the impressive discriminative capabilities of CLIP [38] markedly boosts the zero-shot performance and effectively creates a unified representation spanning different modalities, which highlights the potential of harnessing complex, multi-modal interactions to improve understanding of 3D shapes.

As illustrated in Fig. 1, these methodologies concentrate exclusively on aligning point cloud representations with corresponding features in image and text modalities at the individual instance level, thereby neglecting the intricate structural relations within the pre-aligned image-text feature space. Unfortunately, these mutual relations between data samples are crucial for developing a discriminative representation [32]. For instance, intra-modal relations emphasize similarities such as shapes or textures in images, or connections among scene compositions in texts. Meanwhile, cross-modal relations reveal semantic relevance across modalities, underscoring the connections between diverse entities. The failure to consider these multifaceted relations results in the underutilization of the extensive priors learned by CLIP, culminating in a fragmented comprehension of multi-modal representations.

In this study, we aim to transfer mutual relational knowledge from the pre-aligned image-text domain into the 3D modality, thus improving the effectiveness of current multi-modal contrastive 3D pre-training approaches. Achieving this involves addressing two pivotal challenges. On the one hand, the method, which accurately models mutual relations across various modalities, has not yet been sufficiently explored. Unlike the dual-distribution scenarios of prior studies [12, 13, 32], adding the 3D modality to the pre-aligned image-text framework

incurs a complex array of mutual relations that demands substantial analysis and elaborate modeling. On the other hand, inherent semantic variances within each modality create a pronounced gap between image and text modalities [24, 45], potentially leading to conflicts when aligning 3D representations to pre-trained image or text features. Identifying an effective strategy to mitigate such conflicts for improved convergence is an area that warrants further efforts.

To address the issues aforementioned, we commence with an in-depth analysis of existing representations and associated constraints of mutual relations, examining both within individual modalities (intra-modal) and across different modalities (cross-modal). Following this, we develop a data-driven mechanism to dynamically reconcile these conflicting intra-modal and cross-modal discrepancies. Integrating both parts, we propose a Multi-modal Relation Distillation pre-training framework, namely **MRD**, for 3D representation learning. Since MRD adeptly transfers structural relational knowledge from the image-text modality to the 3D one, it successfully enhances 3D representations by effectively merging information from diverse modalities, consequently leading to improved performance in downstream tasks.

To sum up, our contributions can be summarized as follows:

- We introduce MRD, an innovative self-supervised learning framework, which effectively distills both intra-modal and cross-modal relations, significantly improving the discriminative capability of 3D point cloud representations.
- We perform comprehensive analysis and comparison of mutual relation representations within multi-modal 3D representation learning, and propose a data-driven strategy to resolve the discrepancies in the pre-aligned image-text space of CLIP.
- We present state-of-the-art pre-trained models across diverse applications and deliver superior performance on ModelNet40 [47], ScanObjectNN [41] and Objaverse [8].

2 Related Work

2.1 3D Self-supervised Learning

Self-Supervised Learning (SSL) for point cloud understanding initiates with the 3D modality. The researchers diligently craft a range of pretext tasks such as geometric reconstruction [22, 39, 40, 43], mask auto-modeling [6, 31, 54, 56], normal estimation [33], and contrastive learning [18, 48, 51] to enhance shape representations of deep learning backbones of 3D point clouds [34, 36, 44, 54]. Contemporary endeavors seek to augment 3D SSL by importing insights from other modalities. Attempts leveraging image-guided strategies [58], multi-modal reconstruction [9, 14, 35], and weight-sharing [5, 37] are integrated into conventional frameworks such as Masked Autoencoder [15] or contrastive learning, yielding significant improvements over single-modality methods. Despite the success in enhancing the performance of downstream applications, these methods are limited by their inability to forge more comprehensive cross-modal correlations.

Leveraging the groundbreaking achievements in learning visual concepts directly from textual descriptions through contrastive learning methods, some efforts are made to capitalize on the impressive zero-shot classification capability of CLIP [38] to facilitate the comprehension of 3D shapes. PointCLIP [57] converts 3D point clouds into multi-view images for zero-shot classification by using the pre-trained visual-text encoders of CLIP and PointCLIP-2 [60] advances the methodology by refining the image projection strategies and employing Large Language Models (LLMs) to optimize prompt design, reporting improved performance. However, these techniques suffer the challenges related to information loss during 3D to 2D projection and increased computational demands due to the use of images as an intermediary, which restrict their pervasion.

Simultaneously, some alternatives directly integrate multi-modal representation learning into the 3D domain. CG3D [16] generates ternary pairs of image-text-points from ShapeNet and aligns point cloud features with the text and image features of CLIP, aiming to distill image and text modalities into 3D representations. ULIP [49] enhances this approach by incorporating more intricate multi-view renderings and varied text descriptions. Subsequent developments like ULIP-2 [50] and OpenShape [26] expand the ternary dataset for training, leading to superior outcomes. Uni3D [59] scales up the approach by adapting a Vision Transformer (ViT) pre-trained on image datasets to 3D models, reaching a milestone of 1 billion parameters. While these advances markedly improve the zero-shot classification performance, they primarily align features at the instance level across modalities and neglect the potential of exploring inter-sample relations.

2.2 Relation Modeling in Contrastive Learning

Previous research demonstrates that a thorough understanding of the complex relations among samples significantly improves the acquisition of structural intricacies within representations [32]. Recently, this detailed relation modeling has been integrated into the realm of contrastive learning. RINCE [17] enhances contrastive learning by establishing a partial order among samples, thus contributing to more precise representations of 2D images. Similarly, Soft-InfoNCE [21] refines the margins of the sample distribution based on sample similarities, leading to a notably more distinct feature space.

In the context of bi-modal contrastive learning, CyCLIP [13] introduces the constraints based on pairwise Euclidean distances, applicable both within single modality and across different modalities. Meanwhile, CLIP-PSD [1] adopts similarity-based connections to mitigate the requirement of rigid one-to-one correspondence during training. Some alternative methodologies prioritize modeling complex relations through low-level similarities [11, 12], coherence of local patches [52], and diversity augmentation [55]. Similar enhancements also appear in the scenarios like downstream distillation [20] and CAD representation learning [29]. Both the strategies outperform the standard baseline, indicating their effectiveness.

Nevertheless, relational distillation has not yet been well-explored in the field of multi-modal unified representation learning for 3D point clouds. The integration of the 3D modality with the pre-aligned image-text modality incurs two significant challenges. Firstly, the expansion in the number of modalities results in an increase of inter-sample relations that need to be articulated. The methodology for modeling these intra-modality and cross-modality sample relations as well as establishing the corresponding constraints remains unclear. Secondly, the inconsistency in feature distribution of the image-text modality [24, 45] leads to conflicting relations among these distributions, making it a dilemma to leverage them for proficient representation learning in the 3D modality.

3 Method

3.1 Preliminaries

Previous studies align 3D shape representations with the pre-trained CLIP embedding space of images and texts by applying the contrastive learning loss. Given a set of synthesized image-text-3D triplet $X = \{(x_i^I, x_i^T, x_i^P)\}_{i=1}^N$ as training samples, dedicated encoders E^I, E^T, E^P for each modality process inputs in their respective domains to produce the corresponding feature embeddings f_i^I, f_i^T, f_i^P . Notably, E^I and E^T , derived from the extensive pre-training of CLIP on a large-scale image-text dataset, cover a data volume substantially greater than the number of available synthesized triplets. To circumvent potential model collapse, the weights for the image branch E^I and the text branch E^T are kept frozen, with only the point cloud branch E^P being actively trained to align 3D representations with corresponding image-text pairs. Upon acquiring the feature embeddings for all samples in the batch, alignment between 3D-to-Image and 3D-to-Text is executed following the contrastive learning procedure in CLIP:

$$\begin{cases} \mathcal{L}_{P2T} = \frac{1}{2} \left(\sum_i \frac{\exp(f_i^P \cdot f_i^T / \tau)}{\sum_j \exp(f_i^P \cdot f_j^T / \tau)} + \sum_i \frac{\exp(f_i^P \cdot f_i^I / \tau)}{\sum_j \exp(f_i^P \cdot f_j^I / \tau)} \right) \\ \mathcal{L}_{P2I} = \frac{1}{2} \left(\sum_i \frac{\exp(f_i^P \cdot f_i^I / \tau)}{\sum_j \exp(f_i^P \cdot f_j^I / \tau)} + \sum_i \frac{\exp(f_i^P \cdot f_i^T / \tau)}{\sum_j \exp(f_i^P \cdot f_j^T / \tau)} \right) \end{cases} \quad (1)$$

where i, j represent the indices of the samples, and τ denotes a learnable temperature parameter. The training minimizes the triplet contrastive loss for the weights θ^P in the 3D branch:

$$\mathcal{L}_{Align} = \min_{\theta^P} \frac{1}{2} (\mathcal{L}_{P2T} + \mathcal{L}_{P2I}) \quad (2)$$

3.2 Multi-modal Relation Representation (MRD)

To distill structural knowledge from the multi-modal unified representation space, it is essential to first investigate the form of mutual relation representations between samples within this multi-modal context.

With the image encoder E^I and the text encoder E^T fixed during the alignment phase, the learning of the 3D representation $\{f_i^P\}_{i=1}^N$ targets identifying

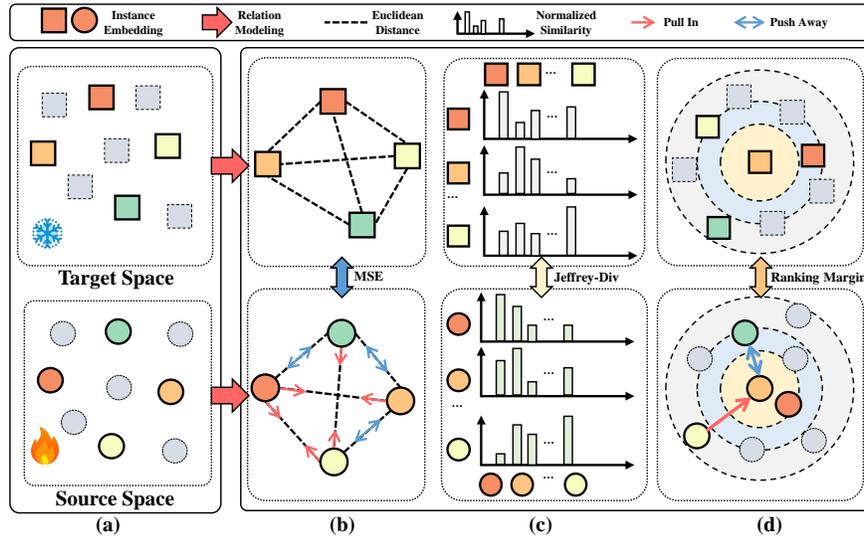


Fig. 2: Comparison of various relation representation forms as well as corresponding distillation strategies. (a) Different embedding spaces. (b) Euclidean distance-based; (c) normalized similarity-based; and (d) partial order-based.

the optimal distribution, based on the image features $\{f_i^I\}_{i=1}^N$ and text features $\{f_i^T\}_{i=1}^N$ extracted from the training set X . This process strives to ensure the maximal preservation of the inherent priori knowledge within the pre-aligned image-text embedding spaces.

Therefore, delineating two primary types of relations becomes crucial. The first type is the intra-modality mutual relation, symbolized as $\psi_{\mathcal{M}}(\cdot)$, concentrating on the sample distribution within a specific modality \mathcal{M} . This singular relation underscores the representational significance of various samples within the same modality. The second type, *i.e.* the cross-modal mutual relation, is denoted as $\phi_{(\mathcal{M}_1, \mathcal{M}_2)}(\cdot)$, delving into the binary relation to reveal the semantic links between samples across modalities \mathcal{M}_1 and \mathcal{M}_2 .

In this study, we explore three widely acknowledged methods for representing mutual relations: Euclidean Distance, Normalized Similarity, and Partial Order.

Euclidean Distance serves to quantify the relative difference between two samples. For intra-modal mutual relations, this difference is captured by the pairwise distance among all samples within the same modality, denoted as:

$$\psi_{\mathcal{M}}^D(X) = \left\{ \frac{1}{\mu} \|f_i^{\mathcal{M}} - f_j^{\mathcal{M}}\|_2^2 \right\}_{(f_i^{\mathcal{M}}, f_j^{\mathcal{M}}) \in \mathcal{M}^2} \quad (3)$$

Similarly, cross-modal relations between samples from different modalities can be expressed as follows, where μ denotes the normalization constant.

$$\phi_{\mathcal{M}_1, \mathcal{M}_2}^D(X) = \left\{ \frac{1}{\mu} \|f_i^{\mathcal{M}_1} - f_j^{\mathcal{M}_2}\|_2^2 \right\}_{f_i^{\mathcal{M}_1}, f_j^{\mathcal{M}_2} \in \mathcal{M}_1 \times \mathcal{M}_2} \quad (4)$$

Normalized Similarity assesses correlation of features of samples within a single batch, emphasizing the overall distribution of correlation between samples rather than the absolute distance between individual sample pairs, as seen with Euclidean distance modeling. This approach offers a perspective that prioritizes relational dynamics within and across modalities. For intra-modal relations, the normalized similarity can be articulated as follows:

$$\psi_{\mathcal{M}}^S(X) = \left\{ \frac{\exp(f_i^{\mathcal{M}} \cdot f_j^{\mathcal{M}} / \tau)}{\sum_{k=1}^N \exp(f_i^{\mathcal{M}} \cdot f_k^{\mathcal{M}} / \tau)} \right\}_{(f_i^{\mathcal{M}}, f_j^{\mathcal{M}}) \in \mathcal{M}^2} \quad (5)$$

Similarly, the normalized similarity of inter-modal relations can be described as:

$$\phi_{\mathcal{M}_1, \mathcal{M}_2}^S(X) = \left\{ \frac{\exp(f_i^{\mathcal{M}_1} \cdot f_j^{\mathcal{M}_2} / \tau)}{\sum_{k=1}^N \exp(f_i^{\mathcal{M}_1} \cdot f_k^{\mathcal{M}_2} / \tau)} \right\}_{f_i^{\mathcal{M}_1}, f_j^{\mathcal{M}_2} \in \mathcal{M}_1 \times \mathcal{M}_2} \quad (6)$$

Partial Order captures the relative ordering between samples, differentiating itself from Euclidean distance and normalized similarity by not imposing strict metric distance constraints. Instead, it is defined through binary relations among samples. For intra-modal relations, partial order can be represented as:

$$\psi_{\mathcal{M}}^P(X) = \{r(f_i^{\mathcal{M}}, f_j^{\mathcal{M}})\}_{(f_i^{\mathcal{M}}, f_j^{\mathcal{M}}) \in \mathcal{M}^2} \quad (7)$$

Likewise, cross-modal relations can be depicted as:

$$\phi_{\mathcal{M}_1, \mathcal{M}_2}^P(X) = \{r(f_i^{\mathcal{M}_1}, f_j^{\mathcal{M}_2})\}_{f_i^{\mathcal{M}_1}, f_j^{\mathcal{M}_2} \in \mathcal{M}_1 \times \mathcal{M}_2} \quad (8)$$

where $r(f_i^{\mathcal{M}_1}, f_j^{\mathcal{M}_2})$ indicates the rank of $f_j^{\mathcal{M}_2}$ after sorting all samples in $\{f_k^{\mathcal{M}_2}\}_{k=1}^N$ by the Euclidean distance.

3.3 Dynamic Relation Distillation

After establishing the structural relational representations, we proceed to distill intra-modality and cross-modality mutual relations. Specifically, in the context of intra-modality relations, our objective is to align the mutual relations within the 3D modality as closely as possible with those observed in the image and text modalities. Conversely, concerning cross-modal mutual relations, we aim for the 3D-to-Image and 3D-to-Text relations to emulate the correlations observed between Image-to-Text or Text-to-Image.

To tackle the issue of inconsistent mutual relations across different modalities, we introduce learnable weights for both the processes of intra-modal and cross-modal mutual relation distillation. These parameters are designed to dynamically

adjust the balance between various distillation objectives in the learning phase, thus aiding in more effective network convergence. The losses associated with intra- and cross-modality distillation are articulated as follows:

$$\begin{cases} \mathcal{L}_{Intra}^{\mathcal{P}} = \alpha L(\psi_{\mathcal{P}}(X), \psi_{\mathcal{I}}(X)) + (1 - \alpha)L(\psi_{\mathcal{P}}(X), \psi_{\mathcal{T}}(X)) \\ \mathcal{L}_{Cross}^{\mathcal{P}2\mathcal{T}} = \beta L(\phi_{\mathcal{P},\mathcal{T}}(X), \phi_{\mathcal{I},\mathcal{T}}(X)) + (1 - \beta)L(\phi_{\mathcal{P},\mathcal{T}}(X), \phi_{\mathcal{T},\mathcal{I}}(X)) \\ \mathcal{L}_{Cross}^{\mathcal{P}2\mathcal{I}} = \gamma L(\phi_{\mathcal{P},\mathcal{I}}(X), \phi_{\mathcal{I},\mathcal{T}}(X)) + (1 - \gamma)L(\phi_{\mathcal{P},\mathcal{I}}(X), \phi_{\mathcal{T},\mathcal{I}}(X)) \end{cases} \quad (9)$$

Since directly incorporating parameters such as α and $1 - \alpha$ can result in unstable optimization due to their non-smoothness, we draw inspiration from the architecture parameter settings in neural architecture search algorithms [25]. For each weight, we introduce three pairs of learnable parameters: $[w_{r_1}^\alpha, w_{r_2}^\alpha; w_{r_1}^\beta, w_{r_2}^\beta; w_{r_1}^\gamma, w_{r_2}^\gamma]$. During training, the values of α and $1 - \alpha$ are derived from softmax of their corresponding learnable weights $[w_{r_1}^\alpha, w_{r_2}^\alpha]$. Similarly, this is applied to β and γ , with these parameters being optimized iteratively. $L(\cdot, \cdot)$ denotes a loss function that penalizes the discrepancies between two mutual relations.

Given the distinct nature of the approaches used to characterize structural relations in Sec. 3.2, the required form of the loss function varies accordingly. Specifically, to mod by calculating Euclidean distances, we employ the Mean Squared Error (MSE) form for distillation. The intra-modal loss can be denoted as

$$L(\Gamma_1, \Gamma_2) = \frac{1}{N^2} \sum_{i,j \in (1,N)} \|\gamma_1^{i,j} - \gamma_2^{i,j}\|_2^2 \quad (10)$$

where Γ_1, Γ_2 represent two distinct mutual relations, and $\gamma_k^{i,j}$ denotes the Euclidean distance between the i -th and j -th samples in the relation modeling of Γ_k .

In the case of the normalized similarity form, we apply the Jeffrey divergence for imposing constraints.

$$L(\Gamma_1, \Gamma_2) = \frac{1}{N} \sum_{i \in (1,N)} (KL(\gamma_1^i | \gamma_2^i) + KL(\gamma_2^i | \gamma_1^i)) \quad (11)$$

where $KL(\cdot | \cdot)$ represents the Kullback-Leibler Divergence, and γ_k^i denotes the normalized similarity of the i -th sample with others in the relation modeling of Γ_j .

Lastly, for the partial order relation form, we utilize a margin-based ranking loss:

$$L(\Gamma_1, \Gamma_2) = \frac{1}{N^2} \sum_{i,j \in (1,N)} \max(0, -r_{i,j} * (f_i^{\mathcal{M}_1} - f_j^{\mathcal{M}_2}) + \eta) \quad (12)$$

where $r_{i,j}$ is the sign function indicating whether the rank order between i and j is consistent in Γ_1 and Γ_2 , and η represents the margin constraint.

We summarize the three relational representations and their respective constraints in Fig. 2, and assess their impacts in the subsequent experimental section. Finally, we select the normalized similarity-based relational representation and Jeffrey divergence to distill both intra- and inter-modal relations.

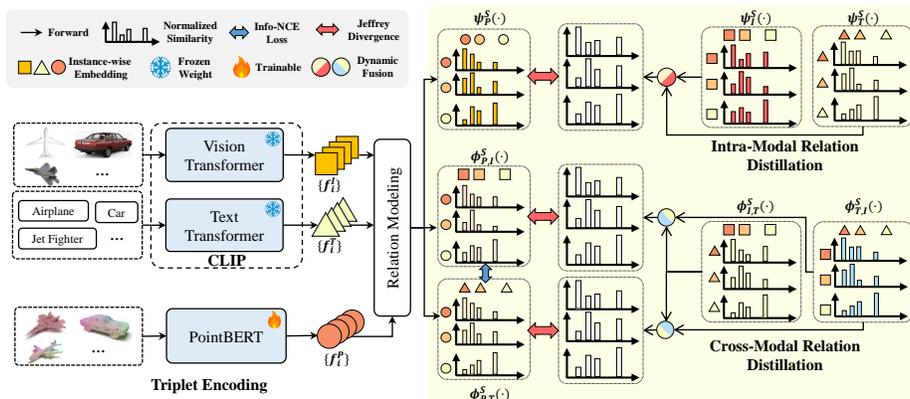


Fig. 3: Overall framework of MRD. With the triplet input, image-text pairs are processed by the pre-trained CLIP model, while the accompanying point clouds are encoded by the 3D encoder. MRD captures the intra-modal mutual relations $\psi(\cdot)$ within each modality and the cross-modal mutual relations $\phi(\cdot)$ across each modality pair. It then dynamically distills and transfers structural information from the pre-aligned image-text space of CLIP into the 3D representations.

3.4 Framework

The entire framework of MRD is depicted in Fig. 3. Its pre-training process involves two main components. Initially, it encodes triplet Image-Text-3D data using the pre-trained CLIP model in conjunction with a trainable 3D foundation model. Following this, both intra- and cross-modal relations between samples within each batch across the three modalities are calculated as described in Sec. 3.2. These structural mutual relations from the image and text modalities are dynamically distilled into the 3D modality as explained in Sec. 3.3, thereby facilitating representation learning.

During training, our objective is to minimize both the instance-level contrastive loss and the relational distillation loss across modalities. Thus, the overall loss function can be articulated as:

$$\mathcal{L} = \min_{\theta^P} \mathcal{L}_{Align} + \lambda(\mathcal{L}_{Intra}^P + \mathcal{L}_{Cross}^{P2T} + \mathcal{L}_{Cross}^{P2I}) \quad (13)$$

where λ is a tunable hyper-parameter used to balance the loss terms.

4 Experiments

4.1 Setup

Training Settings. In alignment with prior research [26, 50, 59], we compile triplets comprising 3D point clouds, 2D images, and textual descriptions from

two extensive datasets of 3D objects. The first integrates the comprehensive 3D collection featured in both OpenShape and Uni3D, encompassing the datasets from Objaverse [8], ShapeNet [4], 3D-FUTURE [10], and ABO [7], culminating in a total of 876K training samples. The second, ShapeNet, contains 52.5K 3D objects across 55 annotated categories. Image data for these datasets are generated through 12 preset camera angles, meticulously arranged to uniformly cover the entire spatial domain, while textual data originate from diverse sources, including curated descriptions, captions crafted by LLMs, and retrieval data. Consistent with preceding methodologies [26], we employ OpenCLIP ViT-BigG-14 as the foundational pre-trained model in this work.

Implementation Details. PointBERT is taken as our fundamental model to pre-train MRD, given its status as a transformer-based backbone that has demonstrated robust performance in previous studies. To assess the impact of scaling up, we incrementally increase its size from 5M (T) to 22M (S), 32M(M), 88M (B) and 307M (L) parameters and evaluate its performance under the same settings, following the way adopted in the counterparts.

We utilize a learning rate of 0.001 for all versions of PointBERT. The batch size is configured to 192 when the model is trained on ShapeNet and increased to 512 on other datasets. Specifically, for ShapeNet, the training duration is set to 70 epochs, while for Objaverse, we extend training to 300 epochs to ensure adequate convergence. We use a cosine annealing schedule with a 15-epoch warm-up period. The weight decay is set at 0.05, λ is set to 3, and the random seed is fixed at 0. The experiments are conducted on eight A800 GPUs, and the whole training process takes about 16 GPU-days.

4.2 Zero-Shot Shape Classification

For the zero-shot classification task, our evaluation considers three standard datasets: ModelNet40 [47], ScanObjectNN [41], and Objaverse [8]. ModelNet40 and ScanObjectNN feature 2,468 and 2,890 test samples across 40 and 15 target classes, respectively, while Objaverse includes 46,205 test samples spanning 1,156 target classes. We calculate and compare the top-1, top-3, and top-5 accuracies with those achieved by other 3D zero-shot classification approaches. For evaluation on ModelNet and Objaverse, the input consists of sampled coordinates and color information from 10,000 points on the mesh surface of each sample. For the ScanObjectNN dataset, the input is composed of 2,048 point-cloud data points derived from the OBJ_ONLY version.

We compare the performance of the proposed MRD framework with the representative state-of-the-art methods such as ULIP-2, OpenShape, and Uni3D. The results are presented in Tab. 1. From Tab. 1 we observe that MRD achieves highly competitive results in all the settings. Notably, despite that OpenShape, Uni3D, and MRD utilize exactly the same training data, MRD significantly surpasses both OpenShape and Uni3D when operating with backbones of comparable parameter scales. It is worth noting when scaled to 88M, MRD delivers the top-1 accuracies of 53.2% on Objaverse and 88.8% on ModelNet40, outperforming Uni3D-L, which employs a considerably larger backbone of 307M,

Table 1: Performance (%) of zero-shot classification on Objaverse-LVIS [8], ModelNet40 [47], and ScanObjectNN [41]. The best and second-best performing methods are highlighted in **bold** and underlined, respectively.

Method	Params (M)	Training Shape Source	Objaverse-LVIS [8]			ModelNet40 [47]			ScanObjectNN [41]			
			Top1	Top3	Top5	Top1	Top3	Top5	Top1	Top3	Top5	
PointCLIP [57]	–	2D Images	1.9	4.1	5.8	19.3	28.6	34.8	10.5	20.8	30.6	
PointCLIP v2 [60]	–	at Inference	4.7	9.5	12.9	63.6	77.9	85.0	42.2	63.3	74.5	
ReCon [35]	43.6	ShapeNet	1.1	2.7	3.7	61.2	73.9	78.1	42.3	62.5	75.6	
CG3D [16]	22.5		5.0	9.5	11.6	48.7	60.7	66.5	42.5	57.3	60.8	
CLIP2Point [19]	–		2.7	5.8	7.9	49.5	71.3	81.2	25.5	44.6	59.4	
ULIP+PointBERT [49]	22.6		6.2	13.6	17.9	60.4	79.0	84.4	51.5	71.1	80.2	
OpenShape+SparseConv [26]	33.7		11.6	<u>21.8</u>	<u>27.1</u>	72.9	87.2	93.0	<u>52.7</u>	72.7	83.6	
OpenShape+PointBERT	32.3		10.8	20.2	25.0	70.3	86.9	91.3	51.3	69.4	78.4	
MRD-T	5.1		<u>11.8</u>	21.2	25.8	<u>74.2</u>	88.3	90.9	55.7	75.5	83.0	
MRD-S	22.6		13.3	23.8	29.2	74.2	<u>88.2</u>	<u>92.5</u>	52.1	<u>73.8</u>	<u>81.2</u>	
ULIP+PointBERT	22.6		Ensembled (no LVIS)	21.4	38.1	46.0	71.4	84.4	89.2	46.0	66.1	76.4
OpenShape+SparseConv	33.7			37.0	58.4	66.9	82.6	95.0	97.5	54.9	76.8	87.0
OpenShape+PointBERT	32.3	39.1		60.8	68.9	85.3	96.2	97.4	47.2	72.4	84.7	
Uni3D-B [59]	88.4	45.9		67.4	<u>74.8</u>	86.1	<u>97.4</u>	<u>98.7</u>	61.7	82.0	89.5	
Uni3D-L	306.7	46.2		<u>67.6</u>	74.7	86.6	96.3	97.8	58.4	81.4	<u>90.1</u>	
MRD-S	22.6	45.2		67.2	74.7	88.5	98.2	99.1	63.0	84.2	91.9	
MRD-M	32.3	<u>45.9</u>		67.8	75.1	<u>87.8</u>	96.5	98.0	<u>62.3</u>	<u>82.4</u>	89.9	
ULIP+PointBERT	22.6	Ensembled		26.8	44.8	52.6	75.1	88.1	93.2	51.6	72.5	82.3
OpenShape+SparseConv	33.7		43.4	64.8	72.4	83.4	95.6	97.8	56.7	78.9	88.6	
OpenShape+PointBERT	32.3		46.8	69.1	77.0	84.4	96.5	98.0	52.2	79.7	88.7	
Uni3D-B [59]	88.4		51.7	74.0	80.8	86.3	96.5	97.9	63.8	82.7	90.2	
Uni3D-L	306.7		53.1	75.0	81.5	86.3	96.8	98.3	58.2	81.8	89.4	
MRD-B	88.4		<u>53.2</u>	<u>75.4</u>	<u>82.1</u>	88.8	97.6	<u>98.5</u>	63.7	<u>84.0</u>	<u>91.5</u>	
MRD-L	306.7		53.6	75.7	82.4	<u>87.3</u>	<u>97.4</u>	98.8	64.1	85.7	91.9	

by 0.1% and **2.5%** respectively. Upon further scaling MRD-B to MRD-L, its performance improves to 53.6%, surpassing all the competitors. Moreover, even when trained solely on ShapeNet, a much smaller dataset, MRD still consistently delivers superior results, outperforming its counterparts by a substantial margin of **1.7%** on the most challenging Objaverse dataset.

4.3 Cross-Modality Retrieval

Considering that previous research assess cross-modal retrieval capabilities primarily relying on qualitative sample visualizations, lacking quantitative comparisons, we propose a new quantitative evaluation protocol to fill this gap for validating the effectiveness in text-to-3D retrieval. Utilizing the recently proposed Cap3D [28] method, which is capable of generating detailed textual descriptions for 3D data, we produce precise descriptions for the samples in the Objaverse test set. These descriptions then serve as the foundation for retrieving the most relevant samples within the Objaverse test set. Subsequently, we compute both the retrieval accuracy for individual instances and that for corresponding categories, providing a robust metric to assess the external-text retrieval capabilities of multi-modal 3D representation learning frameworks. We compare our MRD

Table 2: Accuracy (%) of external text retrieval from Cap3D to Objaverse.

Method	Instance-wise Accuracy			Category-wise Accuracy		
	Top1	Top3	Top5	Top1	Top3	Top5
OpenShape	15.4	28.9	36.8	49.6	71.1	78.8
Uni3D	18.9	33.8	42.2	55.9	76.4	83.4
MRD	20.5	36.0	44.5	57.9	78.2	84.8

Table 3: Ablation results (%) on various candidate relation representation forms.

Method	Relation	Objaverse-LVIS [8]			ModelNet40 [47]			ScanObjectNN [41]		
	Source	Top1	Top3	Top5	Top1	Top3	Top5	Top1	Top3	Top5
Base	None	11.0	20.0	24.6	72.2	85.5	88.8	53.4	73.5	81.9
Euclidean Distance	Intra-Modal	11.3	20.3	24.9	73.0	87.2	90.8	52.1	72.8	84.3
Partial Order		10.0	17.7	21.7	70.6	86.8	90.9	53.1	70.8	79.0
Normalized Similarity		11.3	20.0	24.9	73.1	87.9	91.0	54.6	76.3	83.8
Euclidean Distance	Cross-Modal	11.0	20.2	25.1	72.9	86.8	90.1	53.8	73.1	81.6
Partial Order		10.6	18.6	23.0	71.6	87.0	90.5	55.2	74.0	84.4
Normalized Similarity		11.4	20.3	24.9	72.2	86.7	91.7	52.4	75.9	86.2

with recent state-of-the-art methods, including OpenShape and Uni3D, with the results presented in Tab. 2.

From Tab. 2, we can find that MRD consistently outperforms its counterparts across both metrics, delivering the best performance. Benefiting from its structural distillation capability, MRD effectively learns the relations within the textual feature space. This, in turn, enhances its ability to accurately retrieve corresponding instances when processing external text descriptions, thereby demonstrating its effectiveness.

4.4 Ablation Study

Candidate Relation Representation. We evaluate the effects of different candidate relation representations and their associated distillation losses, as outlined in Sec. 3.2, on various test sets, with the outcomes presented in Tab. 3. From Tab. 3, we can find that incorporating the relational distillation mechanism leads to improvements in accuracy for most settings. Notably, the most significant enhancement is achieved when modeling intra- and inter-modal relations using normalized similarities. This superiority may stem from the ability of normalized similarity-based relation representations, which strikes a more effective balance between the flexibility of the sample distribution and the consistency of relational constraints, in contrast to the stricter Euclidean distance and the more lenient partial order relation, thereby enhancing the performance.

Relation Distillation Strategies. Leveraging the normalized similarity and its corresponding distillation loss, we conduct ablation experiments on dy-

Table 4: Ablation results (%) on various relation distillation strategies within MRD (IR: Intra-modal Relation, CR: Cross-modal Relation, DD: Dynamic Distillation).

Setting			Objaverse-LVIS [8]			ModelNet40 [47]			ScanObjectNN [41]		
IR	CR	DD	Top1	Top3	Top5	Top1	Top3	Top5	Top1	Top3	Top5
			11.0	20.0	24.6	72.2	85.5	88.8	53.4	73.5	81.9
✓			11.3	20.0	24.9	73.1	87.9	91.0	54.6	76.3	83.8
	✓		11.4	20.3	24.9	72.2	86.7	91.7	52.4	75.9	86.2
✓	✓		11.4	20.7	25.4	73.0	87.1	90.3	53.5	73.3	81.2
✓	✓	✓	11.8	21.2	25.8	74.2	88.3	90.9	55.7	75.5	83.0

Table 5: MAE between the similarity matrices on ShapeNet.

IR	CR	DD	P2P/I2I	P2P/T2T	P2T/I2T	P2I/T2I
			0.36	0.20	0.21	0.17
✓			0.08	0.24	0.21	0.18
	✓		0.37	0.20	0.19	0.16
✓	✓		0.20	0.13	0.20	0.16
✓	✓	✓	0.15	0.14	0.20	0.16

Table 6: Performance (%) when scaling up the model size in MRD.

Model	Depth	Width	Heads	Params	MNet40	O-LVIS
MRD-T	6	256	4	5.1M	86.7	47.6
MRD-S	12	394	6	22.6M	88.6	51.4
MRD-M	12	512	8	32.3M	88.0	52.5
MRD-B	12	768	12	88.4M	88.8	53.2
MRD-L	24	1024	16	306.7M	87.3	53.6

dynamic distillation, with the results presented in Tab. 4. As observed from Tab. 4, simultaneously introducing both intra-modal and cross-modal relation distillation leads to the performance that is less favorable compared to independently implementing either of the intra-modal or cross-modal constraints, due to the inconsistency arising from the complex spectrum of mutual relations. Notably, upon introducing the Dynamic Distillation (DD) mechanism, there is a significant gain in performance, demonstrating the efficacy of this strategy.

To quantify the improvements provided by relational distillation more precisely, we compute the Mean Absolute Error (MAE) between the similarity matrices across both intra- and cross-modal samples on ShapeNet to examine the effects of incorporating various relations. In **Tab. 5**, it is evident that IR leads to a decrease in MAE between intra-modal similarities, while CR slightly reduces the divergence of cross-modal similarities. When DD is introduced, MRD achieves a better balance between multiple intra- and cross-modality relations.

Visualization of Dynamic Weights. We visualize the value changes of the dynamic weights of MRD-M when trained on ShapeNet and Objaverse in Fig. 4. As shown in Fig. 4, the intra-modal relations tend to mimic the image-image relations, as they share common information such as textures and geometries between the 3D and image modalities. In contrast, the cross-modal relations lean towards synthesizing the text-image relations, likely due to the inherent differences between textual descriptions and visual appearances. Meanwhile, because of the difference of data distributions of ShapeNet and Objaverse, the dynamic weights converge to distinct values.

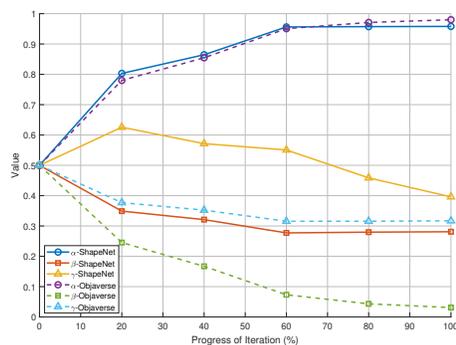


Fig. 4: Visualization of the value changes of dynamic weights along with the progress of iterations on ShapeNet and Objaverse.

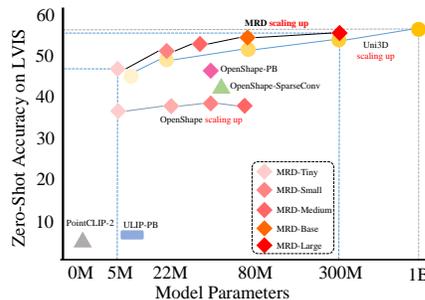


Fig. 5: Comparison of model parameters and zero-shot accuracy (%) on Objaverse, where MRD achieves the highest parameter efficiency.

Scaling up Model Size. We investigate the impact of scaling up the model size of PointBERT, similar to the way taken in Uni3D and OpenShape. The hyper-parameters of the model architecture and their corresponding accuracies are outlined in Tab. 6. Unlike the overfitting issue observed in OpenShape when scaling PointBERT parameters up to 72M, PointBERT trained with MRD successfully scales from 5M to 307M parameters. Additionally, as shown in Fig. 5, despite that Uni3D utilizes a large-scale image dataset for pre-training, MRD still outperforms Uni3D with comparable model sizes, demonstrating its superiority.

5 Conclusion and Discussion

We introduce MRD, an innovative approach to acquiring a unified multi-modal representation of 3D shapes. By distilling both intra-modal and cross-modal relations from the pre-aligned image and text modalities into the 3D modality, MRD facilitates the achievement of more discriminative representations. Our findings demonstrate notable improvements in the zero-shot shape classification and cross-modal retrieval tasks over previous research, underscoring the superior capability of MRD in 3D shape understanding.

Looking ahead, several avenues merit further investigation: a) developing advanced methods for relationship representation and distillation for improved 3D and multi-modal features; b) enhancing the granularity of relationship characterization to extract richer semantic insights for more robust 3D representations; and c) improving conflict removal mechanisms to integrate diverse relational representations for deeper insight into representational relations across modalities.

Acknowledgements

This work is partly supported by the National Natural Science Foundation of China (No. 62176012 and 62022011), the Research Program of State Key Laboratory of Software Development Environment, and the Fundamental Research Funds for the Central Universities, with additional in-kind contributions from PICO ARCH.

References

1. Andonian, A., Chen, S., Hamid, R.: Robust cross-modal representation learning with progressive self-distillation. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 16430–16441 (2022)
2. Armeni, I., Sener, O., Zamir, A.R., Jiang, H., Brilakis, I., Fischer, M., Savarese, S.: 3d semantic parsing of large-scale indoor spaces. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 1534–1543 (2016)
3. Cadena, C., Dick, A.R., Reid, I.D.: Multi-modal auto-encoders as joint estimators for robotics scene understanding. In: *Robotics: Science and systems*. vol. 5 (2016)
4. Chang, A.X., Funkhouser, T.A., Guibas, L.J., Hanrahan, P., Huang, Q., Li, Z., Savarese, S., Savva, M., Song, S., Su, H., Xiao, J., Yi, L., Yu, F.: Shapenet: An information-rich 3d model repository. *CoRR* **abs/1512.03012** (2015)
5. Chen, A., Zhang, K., Zhang, R., Wang, Z., Lu, Y., Guo, Y., Zhang, S.: Pimae: Point cloud and image interactive masked autoencoders for 3d object detection. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 5291–5301 (2023)
6. Chen, G., Wang, M., Yang, Y., Yu, K., Yuan, L., Yue, Y.: Pointgpt: Auto-regressively generative pre-training from point clouds. *Adv. Neural Inform. Process. Syst.* (2023)
7. Collins, J., Goel, S., Deng, K., Luthra, A., Xu, L., Gundogdu, E., Zhang, X., Vicente, T.F.Y., Dideriksen, T., Arora, H., Guillaumin, M., Malik, J.: ABO: dataset and benchmarks for real-world 3d object understanding. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 21094–21104 (2022)
8. Deitke, M., Schwenk, D., Salvador, J., Weihs, L., Michel, O., VanderBilt, E., Schmidt, L., Ehsani, K., Kembhavi, A., Farhadi, A.: Objaverse: A universe of annotated 3d objects. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 13142–13153 (2023)
9. Dong, R., Qi, Z., Zhang, L., Zhang, J., Sun, J., Ge, Z., Yi, L., Ma, K.: Autoencoders as cross-modal teachers: Can pretrained 2d image transformers help 3d representation learning? In: *Int. Conf. Learn. Represent.* (2023)
10. Fu, H., Jia, R., Gao, L., Gong, M., Zhao, B., Maybank, S.J., Tao, D.: 3d-future: 3d furniture shape with texture. *Int. J. Comput. Vis.* pp. 3313–3337 (2021)
11. Gao, Y., Liu, J., Xu, Z., Wu, T., Liu, W., Yang, J., Li, K., Sun, X.: Softclip: Softer cross-modal alignment makes clip stronger. In: *AAAI*. pp. 1860–1868 (2024)
12. Gao, Y., Liu, J., Xu, Z., Zhang, J., Li, K., Ji, R., Shen, C.: Pyramidclip: Hierarchical feature alignment for vision-language model pretraining. *Adv. Neural Inform. Process. Syst.* **35**, 35959–35970 (2022)
13. Goel, S., Bansal, H., Bhatia, S., Rossi, R., Vinay, V., Grover, A.: Cyclip: Cyclic contrastive language-image pretraining. *Adv. Neural Inform. Process. Syst.* **35**, 6704–6719 (2022)

14. Guo, Z., Li, X., Heng, P.A.: Joint-mae: 2d-3d joint masked autoencoders for 3d point cloud pre-training. In: IJCAI. pp. 791–799 (2023)
15. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 16000–16009 (2022)
16. Hegde, D., Valanarasu, J.M.J., Patel, V.: Clip goes 3d: Leveraging prompt tuning for language grounded 3d recognition. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 2028–2038 (2023)
17. Hoffmann, D.T., Behrmann, N., Gall, J., Brox, T., Noroozi, M.: Ranking info noise contrastive estimation: Boosting contrastive learning via ranked positives. In: AAAI. vol. 36, pp. 897–905 (2022)
18. Huang, S., Xie, Y., Zhu, S.C., Zhu, Y.: Spatio-temporal self-supervised representation learning for 3d point clouds. In: Int. Conf. Comput. Vis. pp. 6535–6545 (2021)
19. Huang, T., Dong, B., Yang, Y., Huang, X., Lau, R.W., Ouyang, W., Zuo, W.: Clip2point: Transfer clip to point cloud classification with image-depth pre-training. In: Int. Conf. Comput. Vis. pp. 22157–22167 (2023)
20. Kim, B., Choi, S., Hwang, D., Lee, M., Lee, H.: Transferring pre-trained multimodal representations with cross-modal similarity matching. Adv. Neural Inform. Process. Syst. (2022)
21. Li, H., Zhou, X., Tuan, L.A., Miao, C.: Rethinking negative pairs in code search. In: EMNLP. pp. 12760–12774 (2023)
22. Li, R., Li, X., Fu, C.W., Cohen-Or, D., Heng, P.A.: Pu-gan: a point cloud upsampling adversarial network. In: Int. Conf. Comput. Vis. pp. 7203–7212 (2019)
23. Li, Y., Yu, A.W., Meng, T., Caine, B., Ngiam, J., Peng, D., Shen, J., Lu, Y., Zhou, D., Le, Q.V., et al.: Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 17182–17191 (2022)
24. Liang, V.W., Zhang, Y., Kwon, Y., Yeung, S., Zou, J.Y.: Mind the gap: Understanding the modality gap in multi-modal contrastive representation learning. Adv. Neural Inform. Process. Syst. **35**, 17612–17625 (2022)
25. Liu, H., Simonyan, K., Yang, Y.: DARTS: differentiable architecture search. In: Int. Conf. Learn. Represent. (2019)
26. Liu, M., Shi, R., Kuang, K., Zhu, Y., Li, X., Han, S., Cai, H., Porikli, F., Su, H.: Openshape: Scaling up 3d shape representation towards open-world understanding. Adv. Neural Inform. Process. Syst. (2023)
27. Liu, Z., Zhang, Z., Cao, Y., Hu, H., Tong, X.: Group-free 3d object detection via transformers. In: Int. Conf. Comput. Vis. pp. 2949–2958 (2021)
28. Luo, T., Rockwell, C., Lee, H., Johnson, J.: Scalable 3d captioning with pretrained models. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) Adv. Neural Inform. Process. Syst. (2023)
29. Ma, W., Xu, M., Li, X., Zhou, X.: Multicad: Contrastive representation learning for multi-modal 3d computer-aided design models. In: ACM Int. Conf. Info. Know. Manage. (2023)
30. Misra, I., Girdhar, R., Joulin, A.: An end-to-end transformer model for 3d object detection. In: Int. Conf. Comput. Vis. pp. 2906–2917 (2021)
31. Pang, Y., Wang, W., Tay, F.E., Liu, W., Tian, Y., Yuan, L.: Masked autoencoders for point cloud self-supervised learning. In: Eur. Conf. Comput. Vis. pp. 604–621 (2022)
32. Park, W., Kim, D., Lu, Y., Cho, M.: Relational knowledge distillation. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 3967–3976 (2019)

33. Poursaeed, O., Jiang, T., Qiao, H., Xu, N., Kim, V.G.: Self-supervised learning of point clouds via orientation estimation. In: 3DV. pp. 1018–1028 (2020)
34. Qi, C.R., Yi, L., Su, H., Guibas, L.J.: Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In: Adv. Neural Inform. Process. Syst. (2017)
35. Qi, Z., Dong, R., Fan, G., Ge, Z., Zhang, X., Ma, K., Yi, L.: Contrast with reconstruct: Contrastive 3d representation learning guided by generative pretraining. In: Int. Conf. Mach. Learn. (2023)
36. Qian, G., Li, Y., Peng, H., Mai, J., Hammoud, H., Elhoseiny, M., Ghanem, B.: Pointnext: Revisiting pointnet++ with improved training and scaling strategies. In: Adv. Neural Inform. Process. Syst. (2022)
37. Qian, G., Zhang, X., Hamdi, A., Ghanem, B.: Pix4point: Image pretrained transformers for 3d point cloud understanding. In: 3DV. pp. 1280–1290 (2024)
38. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: Int. Conf. Mach. Learn. pp. 8748–8763 (2021)
39. Rao, Y., Lu, J., Zhou, J.: Global-local bidirectional reasoning for unsupervised representation learning of 3d point clouds. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 5376–5385 (2020)
40. Sauder, J., Sievers, B.: Self-supervised deep learning on point clouds by reconstructing space. Adv. Neural Inform. Process. Syst. **32**, 12942–12952 (2019)
41. Uy, M.A., Pham, Q.H., Hua, B.S., Nguyen, T., Yeung, S.K.: Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In: Int. Conf. Comput. Vis. pp. 1588–1597 (2019)
42. Vu, T., Kim, K., Luu, T.M., Nguyen, T., Yoo, C.D.: Softgroup for 3d instance segmentation on point clouds. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 2708–2717 (2022)
43. Wang, H., Liu, Q., Yue, X., Lasenby, J., Kusner, M.J.: Unsupervised point cloud pre-training via occlusion completion. In: Int. Conf. Comput. Vis. pp. 9782–9792 (2021)
44. Wang, H., Huang, D., Wang, Y.: Gridnet: efficiently learning deep hierarchical representation for 3d point cloud understanding. *Frontiers Comput. Sci.* **16**, 161301 (2022)
45. Wang, Z., Zhao, Y., Huang, H., Liu, J., Yin, A., Tang, L., Li, L., Wang, Y., Zhang, Z., Zhao, Z.: Connecting multi-modal contrastive representations. Adv. Neural Inform. Process. Syst. **36** (2024)
46. Wojek, C., Walk, S., Roth, S., Schiele, B.: Monocular 3d scene understanding with explicit occlusion reasoning. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1993–2000. IEEE (2011)
47. Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J.: 3d shapenets: A deep representation for volumetric shapes. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1912–1920 (2015)
48. Xie, S., Gu, J., Guo, D., Qi, C.R., Guibas, L., Litany, O.: Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In: Eur. Conf. Comput. Vis. pp. 574–591 (2020)
49. Xue, L., Gao, M., Xing, C., Martín-Martín, R., Wu, J., Xiong, C., Xu, R., Niebles, J.C., Savarese, S.: Ulip: Learning a unified representation of language, images, and point clouds for 3d understanding. In: IEEE Conf. Comput. Vis. Pattern Recog. pp. 1179–1189 (2023)
50. Xue, L., Yu, N., Zhang, S., Li, J., Martín-Martín, R., Wu, J., Xiong, C., Xu, R., Niebles, J.C., Savarese, S.: Ulip-2: Towards scalable multimodal pre-training for 3d understanding. In: IEEE Conf. Comput. Vis. Pattern Recog. (2024)

51. Yan, S., Yang, Z., Li, H., Song, C., Guan, L., Kang, H., Hua, G., Huang, Q.: Implicit autoencoder for point-cloud self-supervised representation learning. In: *Int. Conf. Comput. Vis.* pp. 14530–14542 (2023)
52. Yao, L., Huang, R., Hou, L., Lu, G., Niu, M., Xu, H., Liang, X., Li, Z., Jiang, X., Xu, C.: Filip: Fine-grained interactive language-image pre-training. In: *Int. Conf. Learn. Represent.* (2022)
53. Yin, T., Zhou, X., Krahenbuhl, P.: Center-based 3d object detection and tracking. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 11784–11793 (2021)
54. Yu, X., Tang, L., Rao, Y., Huang, T., Zhou, J., Lu, J.: Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 19313–19322 (2022)
55. Yuan, X., Lin, Z., Kuen, J., Zhang, J., Wang, Y., Maire, M., Kale, A., Faieta, B.: Multimodal contrastive training for visual representation learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.* (2021)
56. Zhang, R., Guo, Z., Gao, P., Fang, R., Zhao, B., Wang, D., Qiao, Y., Li, H.: Pointm2ae: multi-scale masked autoencoders for hierarchical point cloud pre-training. *Adv. Neural Inform. Process. Syst.* **35**, 27061–27074 (2022)
57. Zhang, R., Guo, Z., Zhang, W., Li, K., Miao, X., Cui, B., Qiao, Y., Gao, P., Li, H.: Pointclip: Point cloud understanding by clip. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 8552–8562 (2022)
58. Zhang, R., Wang, L., Qiao, Y., Gao, P., Li, H.: Learning 3d representations from 2d pre-trained models via image-to-point masked autoencoders. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 21769–21780 (2023)
59. Zhou, J., Wang, J., Ma, B., Liu, Y.S., Huang, T., Wang, X.: Uni3d: Exploring unified 3d representation at scale. In: *Int. Conf. Learn. Represent.* (2023)
60. Zhu, X., Zhang, R., He, B., Guo, Z., Zeng, Z., Qin, Z., Zhang, S., Gao, P.: Pointclip v2: Prompting clip and gpt for powerful 3d open-world learning. In: *Int. Conf. Comput. Vis.* pp. 2639–2650 (2023)