SCENEFORGE: Enhancing 3D-text alignment with Structured Scene Compositions

Cristian Sbrolli

Department of
Electronics, Information and Bioengineering
Politecnico di Milano
Via Ponzio 34/5, 20133 Milan, Italy
cristian.sbrolli@polimi.it

Matteo Matteucci

Department of
Electronics, Information and Bioengineering
Politecnico di Milano
Via Ponzio 34/5, 20133 Milan, Italy
matteo.matteucci@polimi.it

Abstract

The whole is greater than the sum of its parts, even in 3D-text contrastive learning. We introduce SCENEFORGE, a novel framework that enhances contrastive alignment between 3D point clouds and text through structured multi-object scene compositions. SceneForge leverages individual 3D shapes to construct multiobject scenes with explicit spatial relations, pairing them with coherent multi-object descriptions refined by a large language model. By augmenting contrastive training with these structured, compositional samples, SCENEFORGE effectively addresses the scarcity of large-scale 3D-text datasets, significantly enriching data complexity and diversity. We systematically investigate critical design elements, such as the optimal number of objects per scene, the proportion of compositional samples in training batches, and scene construction strategies. Extensive experiments demonstrate that SCENEFORGE delivers substantial performance gains across multiple tasks, including zero-shot classification on ModelNet, ScanObjNN, Objaverse-LVIS, and ScanNet, as well as few-shot part segmentation on ShapeNetPart. SCENEFORGE's compositional augmentations are model-agnostic, consistently improving performance across multiple encoder architectures. Moreover, SCENEFORGE improves 3D visual question answering on ScanQA, generalizes robustly to retrieval scenarios with increasing scene complexity, and showcases spatial reasoning capabilities by adapting spatial configurations to align precisely with textual instructions.

1 Introduction

Large-scale contrastive learning has transformed vision-language modeling, with early breakthroughs like CLIP [19] and ALIGN [11] demonstrating the power of aligning visual and textual representations at scale. By leveraging vast image-text datasets, these models have achieved remarkable success in zero-shot recognition, retrieval, segmentation, and transfer learning. Following these advancements in 2D, researchers have increasingly turned to 3D, where richer geometric and spatial information is critical for robotics, virtual environments, and augmented reality. However, scaling contrastive learning to 3D remains challenging due to the limited availability of large-scale datasets. Recent works such as Uni3D [31] and OmniBind [21] have made significant strides by leveraging OpenShape [17] dataset, a large-scale ensemble of 3D-text data. These methods align 3D point clouds with 2D-text representation spaces using pretrained CLIP models, achieving strong zero-shot performance on single-object classification benchmarks like ModelNet and ScanObjNN. However, despite these advances, the available 3D text data remain limited, specially compared to image-text datasets, necessitating new strategies to enhance learning.

In this work, we propose a novel approach which allows us to both virtually increase the amount of 3D-text training data and introduce harder samples, improving the contrastive representation. Inspired by image compositions typically used for augmenting classification datasets and methods in previous works, we leverage compositional learning in 3D by constructing multi-object multimodal training samples for contrastive learning. Our method is motivated by two key insights. First, unlike 2D images, where objects are inherently tied to backgrounds, lighting conditions, and perspectives, individual 3D point clouds can be freely combined into structured scenes without visual artifacts. Second, the spatial flexibility of 3D data allows explicit control over object positioning, an ability difficult to achieve in 2D. In contrast to images, 3D objects exist independently of any scene context, enabling meaningful spatial configurations and natural textual descriptions with relational cues (e.g., "A on top of B"). We further refine these descriptions using a large language model, generating diverse and nuanced combinations. Using structured synthesis, we virtually leverage a large-scale, synthetic multi-object 3D-text dataset grounded in real-world captions, significantly expanding diversity and complexity. Training any 3D encoder on these compositional scenes to align with CLIP's representation space achieves consistent improvements across tasks including zero-shot classification, retrieval, segmentation, and VQA. We empirically demonstrate the model-agnostic nature of our compositional augmentations, confirming their effectiveness across multiple encoder architectures. Our contributions are threefold: (1) proposing a novel compositional data pipeline for 3D-text contrastive learning approach synthesizing multi-object 3D scenes; (2) demonstrating consistent performance gains across tasks and backbones; and (3) analyzing the impact of key design choices, including composited object counts, object ratios, and 3D composition strategies.

2 Related Work

Among the first successful works on contrastive alignment of 3D data, ULIP [23] aligned 3D features with CLIP and scaled to the Objaverse dataset in ULIP-2 [23], while Uni3D [31] further advanced this direction by scaling to billion-parameter models, leveraging 2D pretraining and using OpenShape [17] ensembled data. Other works proposed methods for improving the alignment: TAMM [30] mitigates the domain gap between rendered and natural images via adapter modules, while MixCon3D [8] sculpts holistic 3D representations by integrating multi-view rendered images and point clouds. OmniBind [21] adpots instead a differente approach, proposing to ensemble multiple pretrained models via a learnable routing mechanism, achieving state-of-the-art multimodal performance. Our work instead explores an orthogonal approach, scaling multimodal learning by virtually increasing dataset diversity through composition-based augmentation. This strategy aims to improve generalization by exposing the model to a richer distribution of multimodal data.

Compositionality in Multimodal Learning. Compositionality refers to the idea of constructing new concepts by combining simpler ones, a principle that is fundamental human-like concept generalization [12]. Multi-sample composition augmentations such as CutMix [26] and MixUp [28] have been demonstrated to be highly effective for robustness and generalization in 2D tasks such as classification, detection, and segmentation. In the 3D domain, adaptations of these augmentation strategies have been proposed and shown to be effective for unimodal 3D tasks, such as point cloud classification and segmentation [3, 29]. More recently, a stronger form of composition has been employed in image-text contrastive learning [1], demonstrating improved multimodal alignment. Unlike CutMix and MixUp, which blend two images by either mixing pixel values or pasting cut regions, thus being only a weak form of compositional learning, this approach vertically stacks centered crops of two images and combines their captions using the conjunction "and", creating stronger semantic compositions. Building on these ideas, we extend compositional augmentation to 3D data, where the structural properties of point clouds make them particularly well-suited for such techniques. Unlike images, which often contain noisy elements such as background or secondary objects, point clouds representing individual objects can be seamlessly merged into a unified scene without introducing visual artifacts. Moreover, 3D compositional augmentations enable explicit modeling of spatial relationships, which can be reflected in textual descriptions, facilitating relational reasoning alongside object-level recognition.

3 Method

We introduce SCENEFORGE (SF), a framework for composing multi-object 3D scenes by combining individual point clouds with their text descriptions according to explicit spatial relations. SCENE-

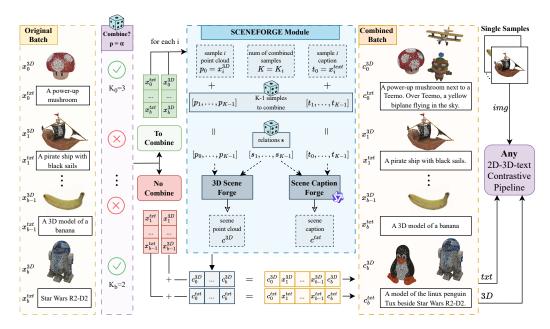


Figure 1: **Our multimodal scene composition framework for contrastive 3D-text learning.** Given a batch of point clouds and their captions, each sample is randomly kept as single object or combined into a synthetic 3D scene with a random number of objects. Objects designated for combination are passed to the SCENEFORGE module, which samples the additional objects and their captions. Spatial relationships among the selected objects are randomly assigned, and the *3D Scene Forge* and *Scene Caption Forge* generate the corresponding combined 3D scene and its composite caption. The newly formed scenes are merged with the unmodified objects to construct a final batch, which is then used for contrastive 3D-text alignment using a pretrained CLIP text encoder.

FORGE can be incorporated into any contrastive multimodal 3D–text learning pipeline, and we refer to the resulting approaches as SF-variants. An overview of our pipeline is shown in Figure 1.

3.1 SCENEFORGE module

SF takes as input a (point cloud, caption) sample (p_0,t_0) and the number K of objects to be combined. It randomly samples the required K-1 point clouds $\{p_1,p_2,\ldots,p_{K-1}\}$ and their corresponding textual captions $\{t_1,t_2,\ldots,t_{K-1}\}$ to generate a composite 3D scene c^{3D} and an associated scene caption c^{txt} . The framework consists of two core modules: the 3D Scene Forge, which spatially arranges the point clouds, and the Scene Caption Forge, which constructs the textual description of the generated scene. Both modules rely on a set of randomly sampled spatial relations $\{s_1,s_2,\ldots,s_{K-1}\}$, which define the relative placement of each object with respect to the one previously placed. These relations dictate how each object is positioned within the combined 3D scene and are used to generate a caption that accurately describes the spatial arrangement. The spatial relations are sampled once per scene and remain consistent across both the 3D and captioning processes.

Spatial Relations. We define three simple spatial relations: "over," "under," and "next to". Since objects in the OpenShape dataset are not consistently oriented along the horizontal axes, using directional terms such as "left" or "right" would be ambiguous unless interpreted as absolute displacements along the X or Z axis. However, this would limit compositions to strictly axis-aligned translations. Instead, we employ the "next to" relation, which allows flexible horizontal placement while preserving semantic coherence. In contrast, the vertical orientation of objects is consistent, making "over" and "under" well-defined, provided that appropriate constraints are imposed on rotation augmentations, as discussed later (Section 3.2). These relations serve as constraints for both the spatial placement of objects in the 3D scene and the construction of descriptive captions that reflect the compositions.

3D Scene Forge. The 3D Scene Forge arranges objects according to their assigned spatial relations, ensuring a semantically coherent composition. The high-level procedure is outlined in Algorithm 1. The module takes as input K point clouds, a set of K-1 pairwise spatial relations, and a target size

```
\begin{array}{l} \textbf{Input: Samples } p, \text{ Relations } s, \text{ Target count } P \\ \textbf{Output: } \text{ Composed 3D sample } c^{3D} \\ c^{3D}, p_{prev} \leftarrow \mathcal{A}^{3D}(p_0) \\ \textbf{for } i = 1 \text{ to } n \text{ do} \\ & \begin{vmatrix} p_i \leftarrow \mathcal{A}^{3D}(p_i) \\ \Delta_{pos} \leftarrow \mathcal{P}(p_i, p_{prev}, s_i) \\ p_i \leftarrow p_i + \Delta_{pos} + \delta + \epsilon \\ c^{3D} \leftarrow \text{cat}(c^{3D}, p_i) \\ p_{prev} \leftarrow p_i \\ \textbf{end} \\ c^{3D} \leftarrow \mathcal{A}^{3D}(\text{subsample}(c^{3D}, P)) \\ \textbf{return } c^{3D} \end{array}
```

Algorithm 1: 3D Scene Forge algorithm.

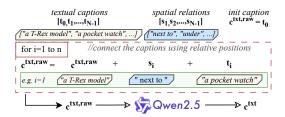


Figure 2: **Scene Caption Forge.** Starting from the initial caption (t_0) , each caption (t_i) is connected using its relative position (s_i) , creating a raw combined caption $c^{\text{txt,raw}}$. The raw caption $c^{\text{txt,raw}}$ is then refined to the final c^{txt} using Owen2.5.

for the final merged cloud. The first object p_0 is used to initialize the scene c^{3D} . Each subsequent object p_i is then placed relative to the previous object p_{i-1} according to its assigned spatial relation s_i , using a function $\mathcal{P}(\cdot)$ that computes the appropriate displacement. Specifically, we define $\mathcal{P}(\cdot)$ based on the spatial relation between objects. For "over", we align the minimum z-coordinate of p_i above the maximum z-coordinate of p_{i-1} :

$$\mathcal{P}(p_i, p_{i-1}, \text{``over''}) = \max_{\mathbf{z}}(p_{i-1}) - \min_{\mathbf{z}}(p_i),$$

and reverse the roles for "under". For "next to", we sample a horizontal unit vector \mathbf{d} in the xy-plane and compute:

$$\mathcal{P}(p_i, p_{i-1}, \text{``next to''}) = \left(\max_{\mathbf{x} \in p_{i-1}} \langle \mathbf{x}, \mathbf{d} \rangle - \min_{\mathbf{y} \in p_i} \langle \mathbf{y}, \mathbf{d} \rangle \right) \mathbf{d},$$

where $\langle \cdot, \cdot \rangle$ is the inner product.

In all relations, to prevent perfect alignment and introduce slight randomness, we add a fixed offset δ along the shift direction $(+\mathbf{z}, -\mathbf{z} \text{ or } \mathbf{d})$, along with a small Gaussian noise term $\epsilon \in \mathcal{R}^3$. After placing all K objects, we downsample the composite point cloud c^{3D} to the target number of points P, ensuring diversity while maintaining spatial consistency. Notice that normalization and augmentation \mathcal{A}^{3D} is performed on each sample before adding it to the scene, as well as on the final scene.

Scene Caption Forge This module constructs a textual description of the combined scene by sequentially incorporating individual object captions $\{t_0, t_1, \dots, t_{K-1}\}$ and their corresponding spatial relations $\{s_1, s_2, \dots, s_{K-1}\}$. The captioning process mirrors the spatial composition of the 3D scene, starting with the first object's caption and appending each subsequent caption preceded by its respective spatial relation. However, due to the method's simplicity, the generated caption may exhibit artifacts such as misplaced punctuation before spatial relations, incorrect capitalization following conjunctions, and disfluent sentence structures. Consequently, we refer to this as a raw caption, $c^{txt,raw}$. To enhance readability and coherence, we refine the raw caption using a large language model (Qwen2.5 [24]), obtaining the refined caption e^{txt} . The model corrects grammar, punctuation, and structure while preserving the original meaning and spatial relationships. It restructures the text into a fluent, human-like description, splitting overly long sentences when necessary. Beyond improving fluency, this rewriting process also enhances caption diversity and refines OpenShape [17] captions, originally generated by BLIP [14] and Microsoft Azure Cognitive Services (2023), by leveraging the more advanced linguistic capabilities of a recent language model. The full prompt used for refinement is provided in the supplementary, together with an ablaton on the LLM. An overview of the proposed method is provided in Figure 2.

3.2 Training Scheme

We mix single and multi-object samples in the same training batch with a predefined ratio α (ablated in Section 5). This allows the model to retain strong performance on single-object tasks while benefiting from the additional compositional training signals. For each combined sample, the number of objects to combine is randomized between 2 and the maximum number of combinable objects N, which we investigate in Section 4. This allows for $\sum_{k=1}^{N} \frac{D!}{(D-k)!} 3^{k-1}$ possible configurations with

up to k samples and 3 possible relations, with D being the dataset cardinality (extended derivation in the supplementary). With OpenShape data, this allows for 6E scenes.

Our framework operates entirely at the batch–generation level and can be plugged into *any* contrastive pipeline that aligns text, image, and 3D point-cloud embeddings. Because all previous methods include images, we likewise keep the 2D modality when deploying SCENEFORGE. In principle, we could rasterise every composed scene and align its image views, but real-time rendering is computationally prohibitive given our budget. Instead, we mask composed samples for the 2D–3D loss terms, evaluating those terms only on pre-rendered single-object views. This preserves baseline performance on image–3D tasks (see supplementary) while allowing our compositions to focus on strengthening text–3D alignment. Following established practice, the CLIP image and text encoders stay frozen; gradients flow only through the 3D encoder.

Loss Partitioning. We consider contrastive models employing the InfoNCE loss proposed in CLIP [19]. For modalities $m, n \in \{txt, 2D, 3D\}$ and a sample subset S, we define

$$\mathcal{L}_{m \to n}(\mathcal{S}) = -\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} \log \frac{\exp(\langle e_i^m, e_i^n \rangle / \tau)}{\sum_{j \in \mathcal{S}} \exp(\langle e_i^m, e_j^n \rangle / \tau)},\tag{1}$$

where e_i^m, e_i^n are ℓ_2 -normalised embeddings and τ is a learnable temperature.

Let S_c and S_s denote the composed and single-object samples in a batch, with $N = |S_c| + |S_s|$. Because each sample is composed with probability α , $\mathbb{E}[|S_s|] = (1 - \alpha)N$. We scale the image–3D block so that, *per batch*, it contributes the same total gradient budget as the text–3D block:

$$\mathcal{L} = \underbrace{\frac{1}{2} \left[\mathcal{L}_{3D \to txt} (\mathcal{S}_c \cup \mathcal{S}_s) + \mathcal{L}_{txt \to 3D} (\mathcal{S}_c \cup \mathcal{S}_s) \right]}_{\text{text-3D (all } N \text{ samples)}} + \frac{N}{|\mathcal{S}_s|} \underbrace{\frac{1}{2} \left[\mathcal{L}_{3D \to 2D} (\mathcal{S}_s) + \mathcal{L}_{2D \to 3D} (\mathcal{S}_s) \right]}_{\text{2D-3D (singles only)}}. \tag{2}$$

Considering the large batch sizes used in contrastive models, $|S_s|$ is tightly concentrated around its mean, so we simply replace the dynamic factor $N/|S_s|$ by its expectation $1/(1-\alpha)$.

Implementation Details. We instantiate our scene-composition pipeline on three approaches using different point-cloud encoders: OpenShape-PointBERT [17], Uni3D-G [31], and ViT-Lens-G (a frozen ViT-bigG/14 with trainable adapter lenses) [13]. This variety allows us to verify the encoderagnostic nature of our method and to select the approach which benefit the most from our pipeline. All variants are trained with their public available code and our modified loss for 200 epochs with a global batch size of 1152, $\alpha = 0.5$, and a target point-cloud resolution of P = 10 k points. This point budget was chosen as a trade-off between detail and efficiency, as detailed in the supplementary. For caption generation, we use the Qwen 2.5 7B-instruct [24] large language model. SceneForge only requires an additional GPU hosting the lightweight LLM for composition. We pre-generate the first M batches and then switch to a producer-consumer setup: while batch t trains, batch t+M is assembled in parallel. Though composition latency is not always fully hidden, parallelization significantly reduce it. For faster models, multiple SCENEFORGE instances and quantization can be used to further amortize overhead. Measured slowdown ranges from 0% to 50% depending on backbone, LLM quantization and replication (see supplementary). For the function \mathcal{A}^{3D} , we normalize point clouds to the unit sphere and apply random point dropout and scaling. Rotation and translation strategies differ based on whether we process single objects, objects intended for combination in a scene, or the final composed scene. For single objects, we allow both shifts and larger rotations. However, for objects that will be combined, shifts are disabled to avoid inconsistencies in composition, while full rotations around the vertical axis and slight rotations along other axes are permitted. The latter ensures that concepts such as "over" and "under" remain semantically meaningful. For the final combined point cloud, we adopt the same rotation constraints as in the previous case to preserve spatial semantics but additionally allow translations.

4 Experiments

4.1 Zero-Shot Classification

Datasets. We follow the standard zero-shot evaluation protocol on Objaverse LVIS [6], Model-Net40 [22] and ScanObjNN [20], where categories are mapped to text prompts by formatting a set of

Model	LV	/IS	Mod	elNet	Scan	ObjNN	Scannet	Avg	Model	LV	/IS	Mod	elNet	Scan	ObjNN	Scannet	
	T1	T5	T1	T5	T1	T5	T1	Δ		T1	T5	T1	T5	T1	T5	T1	Δ
ULIP 2	46.3	75.0	84.0	97.2	45.6	82.9	38.1	_	ULIP 2	50.6	79.1	84.7	97.1	51.5	89.3	38.9	_
TAMM	42.0	71.7	86.3	98.1	56.7	86.1	42.4	_	TAMM	50.7	80.6	85.0	98.1	55.7	88.9	41.8	_
MixCon3D	47.5	76.2	87.3	98.1	57.7	89.8	43.0	_	MixCon3D	52.5	81.2	86.8	98.3	58.6	89.2	44.1	_
OmniBind-L	_	_	-	-	-	_	_	_	OmniBind-L	54.0	82.9	86.6	99.0	64.7	94.2	46.3	-
OmniBind-F	-	-	-	-	-	-	-	-	OmniBind-F	53.6	81.8	87.1	99.0	64.7	94.4	46.1	_
OpenShape SF-OpenShape				97.4 98.1		84.7 85.9	40.3 41.5	+1.50	OpenShape SF-OpenShape			84.4 85.2			88.7 89.5	39.4 41.8	+1.43
ViT-Lens SF-ViT-Lens	50.1 50.9			97.8 98.0		87.7 89.1	43.8 44.5	+0.78	ViT-Lens SF-ViT-Lens	52.0 52.8		87.6 88.0			90.3 91.2	43.7 45.1	+0.85
Uni3D SF-Uni3D				98.4 99.0	66.5 67.3	90.1 91.5	43.9 47.6	+1.73	Uni3D SF-Uni3D			87.3 88.2		63.9 65.2	91.7 93.4	45.8 49.4	+1.75

⁽a) Trained on ensemble (no LVIS).

Table 1: **Zero-shot classification accuracy** (%). "SF-" denotes models trained with SCENEFORGE. Green cells (\blacksquare) are the best results, yellow (\blacksquare) the second best. The rightmost column reports the average Top-1 improvement (Δ) of the augmented model over its baseline.

templates (e.g., "a point cloud model of a) and the model is evaluated on the classification accuracy. Additionally, adopting the pipeline from CLIP² [27], we test our models on the Scannet [5] dataset to evaluate their zero-shot performance on object instances from real-world scenarios.

What is the optimal value of N? Figure 3 tracks zero-shot accuracy for three contrastive learners: OpenShape-PointBERT [17], Uni3D (EVA-Giant) [7, 31], and ViT-Lens-G [13], as we vary the maximum number of shapes that SCENEFORGE can merge during training. A consistent trend emerges. From N=1 to N=3, accuracy rises monotonically: LVIS gains +0.8-1.3 pp, ScanObjNN +0.8-1.3 pp, Scan-Net +1.4-3.6 pp, and ModelNet40 about +0.8pp. Improvements peak at N=3; Uni3D benefits most, while the lighter OpenShape also advances. ViT-Lens shows a more modest gain, plausibly because its frozen CLIP backbone trains only lightweight adapters, offering less plasticity when confronted with composed shapes. Increasing to N=4 plateaus on canonical datasets and already trims accuracy on LVIS and ScanNet, whose higher intra-class variability makes them clutter-sensitive. At N=5 the drop is universal, indicating that squeezing five

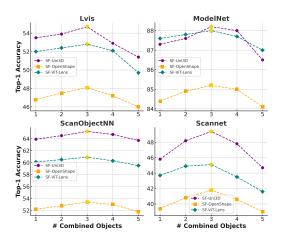


Figure 3: Top-1 accuracy across different datasets (Lvis, ModelNet, ScanObjectNN, and Scannet) as a function of the number of combined objects.

objects into a fixed 10 k-point budget fragments salient geometry and introduces caption noise, hampering alignment. Interestingly, performance variations are relative minor on ModelNet40 and ScanObjNN, but the effect is far more pronounced on LVIS and ScanNet. We attribute this to the latter's more complex object distributions: as additional shape combinations are introduced, the model must balance greater intra-class variability with the need for discriminative features, a trade-off that becomes harder to resolve in these richer, noisier datasets. Considering these results, we adopt $N{=}3$ as default SF-variants and report the full sweeps $(N{=}1{-}5)$ in the supplementary.

Detailed Quantitative Comparison. Table 1 benchmarks SCENEFORGE on the considered zeroshot 3D-text classification suites. For each backbone we report its *best* composition size (N=3) together with its single-shape baseline (N=1), and list recent and sota models for reference.

Effect of SceneForge. Across all three backbones SceneForge delivers consistent top-1 gains. With the "no-LVIS" training split the average improvement is +1.50 pp for OpenShape, +0.78 pp

⁽b) Trained on ensemble (with LVIS).

Method	1-sł	not	2-shot			
	mIoU	Δ	mIoU	Δ		
OmniBind-L OmniBind-F	77.2 77.8	-	79.9 80.3	-		
OpenShape SF-OpenShape	74.0 76.2	+2.2	76.5 79.1	+2.6		
ViT-Lens SF-ViT-Lens	75.5 77.0	+1.5	77.9 80.1	+2.2		
Uni3D SF-Uni3D	75.9 78.5	+2.6	78.2 81.2	+3.0		

Model	B-4	Δ B-4	CIDEr	Δ CIDEr	EM	$\Delta \mathrm{EM}$
OmniBind-L + BLIP2-FlanT5 OmniBind-F + BLIP2-FlanT5	8.5 8.3	_	62.9 62.1	-	17.1 17.6	-
OpenShape + BLIP2-FlanT5 SF-OpenShape + BLIP2-FlanT5	6.3	+1.8	54.8 61.5	+6.7	14.1 16.9	+2.8
ViT-Lens + BLIP2-FlanT5 SF-ViT-Lens + BLIP2-FlanT5	7.2 8.5	+1.3	57.5 63.4	+5.9	15.7 17.8	+2.1
Uni3D + BLIP2-FlanT5 SF-Uni3D + BLIP2-FlanT5	7.5 10.4	+2.9	58.3 66.7	+8.4	16.4 20.5	+4.1

Table 2: One-shot and two-shot part segmentation on ShapeNetPart.

Table 3: Performance on the ScanQA dataset using BLEU-4, CIDEr, and Exact Match.

for ViT-Lens and +1.73 pp for Uni3D; with the "+LVIS" split the corresponding gains are +1.43, +0.85 and +1.75 pp. Uni3D benefits most, likely a consequence of its larger capacity, which can better exploit the richer intra-sample diversity injected by multi-shape compositions, yet even the smaller OpenShape and the adapter-based ViT-Lens improve.

Comparison with prior works. All three SCENEFORGE variants surpass previous non-ensemble methods (ULIP-2 [23], TAMM [30], MixCon3D [8]) on every dataset. Moreover, SF-Uni3D outperforms OmniBind [21], the strongest published ensemble, despite using a single model: on the LVIS-ModelNet-ScanObjNN-ScanNet quartet it achieves absolute top-1 margins of +0.7, +1.6, +0.5 and +3.1 pp, respectively. These results underscore that structured multi-object augmentation offers a more inference-efficient strategy for enhancing representations than costly ensemble methods.

Why do multi-object compositions enhance single-object classification? While it might initially appear counterintuitive, training with multi-object compositions fosters improved representations that benefit single-object recognition tasks. This phenomenon aligns with established findings in representation learning literature, particularly in image classification, where multi-sample augmentations (e.g., CutMix, MixUp) are known to induce smoother decision boundaries and promote robust generalization by exposing models to more diverse feature combinations. Analogously, in our structured 3D scene compositions, the increased complexity and relational context implicitly regularize the learned representations, facilitating the emergence of discriminative features resilient to variations in single-object scenarios encountered at inference. By analyzing the positive–negative similarity margins in zero-shot classification, we find that SF-variants yield larger margins over baselines, indicating stronger inter-class separation and more robust decision boundaries. We report the quantitative analysis in the supplementary material.

4.2 Few-Shot Part Segmentation on PartNet

We follow the protocol of Uni3D [31] on the ShapeNetPart dataset [25], evaluating each backbone in one-shot and two-shot regimes. As in PointNet++ [18], we freeze the pretrained transformer encoder and attach lightweight feature-propagation heads that upsample intermediate representations to dense part predictions. Only these heads are fine-tuned on the few labeled part annotations.

Table 2 reports mean IoU (mIoU) and the improvement Δ over the corresponding single-shape baseline for all three backbones (*OpenShape*, *ViT-Lens*, *Uni3D*) with and without SCENEFORGE. SCENEFORGE yields consistent gains: SF-OpenShape improves by +2.2/ +2.6 pp in one-/two-shot, SF-ViT-Lens by +1.5/ +2.2 pp, and SF-Uni3D by +2.6/ +3.0 pp, with SF-Uni3D achieving the highest absolute mIoU (78.5 / 81.2). These improvements suggest that multi-shape pretraining encourages more fine-grained, group-based features that transfer effectively to part segmentation, structuring the feature space to better capture part-level relationships even under extreme label scarcity.

4.3 3D Question Answering on ScanQA

We further assess 3D-text alignment on the ScanQA benchmark [2, 9], which requires answering natural-language questions about ScanNet scenes. Following prior works, we freeze each 3D encoder and attach it to BLIP2-FlanT5 [4, 15], then fine-tune on ScanQA's question-answer pairs.

In addition to the sota contrastive baselines (OmniBind-Large, OmniBind-Full), we include SF-OpenShape, SF-ViT-Lens and SF-Uni3D as our multi-shape variants, alongside their single-shape counterparts. We report here BLEU-4 (B-4), CIDEr and Exact Match (EM); full metrics appear in the supplementary. Table 3 gives the results. All three backbones see substantial gains with SCENEFORGE: SF-OpenShape improves B-4 by +1.8 pp, CIDEr by +6.7 pp and EM by +2.8 pp; SF-ViT-Lens adds +1.3 pp, +5.9 pp and +2.1 pp, respectively; and SF-Uni3D leads with +2.9 pp, +8.4 pp and +4.1 pp. While all backbones gain, Uni3D yields the largest relative and absolute boosts, consistent with its greater capacity, whereas ViT-Lens, despite smaller gains, still surpasses OmniBind. A qualitative review shows that, while baseline encoders match our variants on attribute- or color-based questions, SCENEFORGE variants significantly outperform them on spatial reasoning queries, e.g. "What is over the brown chair?", where modeling inter-object relationships is essential. This suggests that our structured multi-shape augmentation not only sharpens local feature representations but also boosts the encoder's ability to infer complex spatial configurations, a critical capability for 3D scene understanding.

4.4 Supervised Fine-Tuning

To test if the benefits of SCENEFORGE pre-training extend to supervised settings, we evaluate full fine-tuning and Parameter-Efficient Fine-Tuning (PEFT) methods [10, 16, 32]. We perform experiments on three of the considered classification benchmarks: the synthetic ModelNet40, the challenging real-world ScanObjectNN, and ScanNet Instances extracted from complex indoor scenes.

Table 4: Supervised fine-tuning accuracy (%).

	14610 11	. Super vise a mine	tuning acc	#### (/e) .		
Model	Method	Trainable Params	ModelNet40	ScanObjectNN	ScanNet Inst.	
	Full Fine-Tuning	1016.5M (100%)	94.28	97.12	82.72	
Uni3D	Adapter	7.6M (0.74%)	94.35	96.80	81.42	
	DAPT	7.3M (0.72%)	94.33	96.78	82.65	
	PointGST	4.1M (0.40%)	94.83	97.68	83.04	
	Full Fine-Tuning	1016.5M (100%)	94.42	97.58	83.58	
SF-Uni3D	Adapter	7.6M (0.74%)	94.46	97.09	82.56	
	DAPT	7.3M (0.72%)	94.49	97.15	83.46	
	PointGST	4.1M (0.40%)	94.95	98.09	84.29	

As shown in Table 4, our SF-Uni3D backbone provides a superior initialization, consistently outperforming the baseline, especially on the complex ScanNet Instances. A strong synergy with PEFT is also evident: the PointGST method surpasses even full fine-tuning while using just 0.4% of the trainable parameters.

4.5 N-Objects cross-modal Retrieval

Following prior work [23, 31], we evaluate cross-modal retrieval on the unseen Objaverse-LVIS dataset, measuring top-k accuracy for both 3D-to-text and text-to-3D retrieval. Models are trained on the full ensemble excluding this set, and performance is assessed via cosine similarity retrieval across embedded samples. Beyond standard single-object retrieval, we aim to analyze how well our models understand increasingly multi-object scenes. To this end, we introduce the N-LVIS benchmark, where each shape is composed with N-1 additional objects. For clarity, we report results only for the best SF backbone (Uni3D with N = 1..5) and compare against OmniBind-F [21]; analogous results for other backbones are provided in the supplementary.

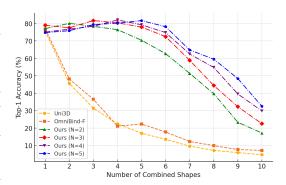


Figure 4: Top-1 averaged retrieval accuracy on the N-LVIS datasets as N increases.

We evaluate retrieval on N-LVIS from N=1 (standard retrieval) to N=10, reporting averaged top-1 accuracy (text to 3D and viceversa) in Figure 4. Prior models degrade sharply when faced with multi-object compositions, dropping below 50% at N=2. In contrast, our models, trained with

varying numbers of composed objects, exhibit strong generalization, each peaking near its training composition size. Notably, the $N{=}3$ model sustains over 70% accuracy at $N{=}6$ and around 60% at $N{=}7$, highlighting its robustness in complex scenes. This comes at a slight cost: models trained with higher N perform marginally worse at lower N, reinforcing our findings in Section 4.1 that simpler compositions better preserve single-object understanding. Full top-1 and top-5 results are provided in the supplementary.

4.6 Object Repositioning

To qualitatively assess our approach ability to improve reasoning about spatial relationships, we introduce a simple object repositioning task. Given two objects, we combine them in an initial configuration using our MSF. The combined caption is then modified to describe a new spatial relationship, and we optimize a three-parameter offset vector to reposition the second object, maximizing alignment with the updated description while keeping the 3D-text encoder frozen. Full optimisation details appear in the supplementary. As illustrated in Figure 5 encoders

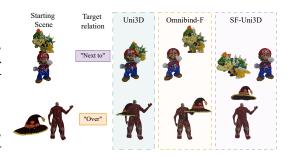


Figure 5: Object repositioning example.

trained with SCENEFORGE reliably relocate the object to the target position, while baseline encoders stall in semantically plausible yet misaligned configurations. We observe the same qualitative behaviour with all other variants, confirming that SCENEFORGE strengthens spatial and compositional reasoning.

5 Ablation Studies

5.1 Proportion of Multi-Object Samples

The hyperparameter α determines the fraction of multiobject samples in each training batch. We analyze its impact by varying α from 0 to 1 and evaluating the average zero-shot top-1 accuracy on the zero-shot datasets (Figure 6). Increasing α initially improves performance with respect to the baselines, with accuracy peaking around $\alpha=0.5$. However, for larger α , performance declines, likely due to an excessive focus on compositional relationships at the expense of single-object understanding.

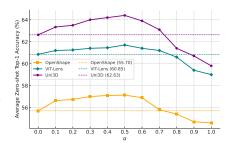


Figure 6: Effect of varying α on average zero-shot top-1 accuracy.

5.2 Do other composition functions work?

To assess the impact of our 3D scene composition method, we conduct experiments by fixing the training method to ours, the backbone to the strongest we found (Uni3D) and just ablate the composition method against object composition approaches: Point-CutMix [29] and PointMixUp [3]. PointCutMix replaces regions of one

Composition Method	Lvis Top-1	ModelNet Top-1	ScanObjNN Top-1	Scannet Top-1	
None (Uni3D)	53.5	87.3	63.9	45.8	
PointCutMix-R	53.5	87.1	64.1	47.5	
PointCutMix-K	44.7	83.0	45.1	34.8	
PointMixup	39.2	78.7	41.4	30.2	
SF-Uni3D (N=2)	53.9	87.6	64.5	48.2	

Table 5: Different 3D composition methods on zero-shot cls.

point cloud with those from another and has two variants: PointCutMix-R, which randomly replaces points, and PointCutMix-K, which preserves key structures. PointMixUp interpolates point coordinates and features between objects. For a fair comparison, since these methods combine two objects, we compare their results specifically to our N=2 configuration. Textual descriptions for PointCutMix and PointMixUp compositions are generated by concatenating object captions with "and" and subsequently refined using Qwen [24], following our pipeline.

Results in Table 5 show our method consistently outperforming all baselines. Although PointCut-Mix-R exceeds Uni3D on a most benchmarks, it still lags behind our method, even more considering the N=3 variant; PointCutMix-K and PointMixUp perform worse than Uni3D. This gap stems from how each approach handles object semantics and spatial coherence: our method builds structured scenes that align naturally with captions, whereas PointCutMix-R randomly mixes whole objects, creating overlaps, and PointCutMix-K and PointMixUp fragment or interpolate shapes, producing unrealistic, poorly described scenes. Overall, maintaining clear object semantics and spatial relationships through structured composition yields superior generalization.

5.3 Are simple relations enough?

To test if the learned spatial understanding generalizes beyond the simple pre-training relations, we isolate performance on ScanQA questions involving more complex, unseen spatial queries. Table 6 shows that our compositionally-trained models consistently outperform their baselines across all backbones and a variety of complex relations.

Relation Type	Metric	OpenShape			ViT-Lens			Uni3D		
110111011011 15 pc	1120210	Baseline	SF	Δ	Baseline	SF	Δ	Baseline	SF	Δ
Attached To (21)	CIDEr	54.5	61.5	+7.0	57.1	63.3	+6.2	57.9	66.6	+8.7
	EM	14.1	17.1	+3.0	15.6	17.9	+2.3	16.5	20.8	+4.3
Sitting On (59)	CIDEr	56.8	63.4	+6.6	59.0	65.1	+6.1	61.0	70.1	+9.1
	EM	15.2	17.7	+2.5	16.6	18.4	+1.8	17.5	22.6	+5.1
Between (112)	CIDEr	54.1	61.2	+7.1	56.8	62.9	+6.1	57.2	66.5	+9.3
	EM	14.0	17.0	+3.0	15.5	17.9	+2.4	15.8	20.5	+4.7
Closest To (112)	CIDEr EM	55.0 14.3	61.8 17.2	+6.8	57.5 15.8	63.5 18.0	+6.0	58.5 16.2	67.0 20.4	+8.5

Table 6: Generalization to unseen spatial relations on ScanQA for all backbones.

The comprehensive results in Table 6 show that SCENEFORGE provides consistent benefits across all three backbones. Each SF-variant significantly outperforms its respective baseline on all complex relation types. This indicates that our pre-training builds a robust spatial foundation that generalizes effectively to more nuanced relational queries, regardless of the underlying encoder architecture.

+6.4

+2.7

58.2

16.3

64.0

18.2

+5.8

+1.9

60.3

17.1

68.3

+8.0

+4.7

62.5

17.6

14.9

6 Limitations and Future Directions

CIDEr

In Front Of (246)

While SCENEFORGE consistently enhances 3D-text alignment across multiple backbones and tasks, we are aware of its limitations. First, our synthetic scene generation employs only three basic spatial relations, which, although diversified through LLM-based refinement, do not fully capture the complexity of natural environments. Future research could focus on more realistic and varied364compositions guided by learned object co-occurrence patterns and spatial priors. Second, due to computational constraints, we maintain a fixed 10 k-point budget for multi-object compositions, resulting in accuracy degradation for densely populated scenes (see Figure 3). Addressing this will require exploring larger point budgets or employing more sophisticated sampling techniques to preserve salient geometric features in complex scenarios. Third, although we leverage the lightweight Qwen2.5 model for refinement, the overhead is not always negligible and reducing it introduces a memory-time tradeoff. Finally, due to rendering costs, we could not pair each composition with a synthetic image, though extending the pipeline to incorporate rendered views for joint 2D–3D learning or studying alternative approaches, such as aligning 3D compositions to aggregations of the single image embeddings, remains a promising direction. Overall, addressing these limitations will significantly broaden the practical impact and robustness of compositional 3D-text learning methods.

Acknowledgments. This paper is supported by the PNRR-PE-AI FAIR project funded by the NextGeneration EU program. We acknowledge ISCRA for awarding this project access to the LEONARDO supercomputer, owned by the EuroHPC Joint Undertaking, hosted by CINECA (Italy).

References

- [1] Maxwell Aladago, Lorenzo Torresani, and Soroush Vosoughi. Semantic compositions enhance vision-language contrastive learning. *arXiv* preprint arXiv:2407.01408, 2024.
- [2] Daichi Azuma, Taiki Miyanishi, Shuhei Kurita, and Motoaki Kawanabe. Scanqa: 3d question answering for spatial scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022. URL https://github.com/ATR-DBI/ScanQA. This dataset is released under the Creative Commons Attribution-NonCommercial-ShareAlike 3.0 Unported License.
- [3] Yunlu Chen, Vincent Tao Hu, Efstratios Gavves, Thomas Mensink, Pascal Mettes, Pengwan Yang, and Cees GM Snoek. Pointmixup: Augmentation for point clouds. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III 16*, pages 330–345. Springer, 2020.
- [4] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Yunxuan Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *Journal of Machine Learning Research*, 25(70):1–53, 2024.
- [5] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. URL http://www.scan-net.org/. The ScanNet data is released under the ScanNet Terms of Use, and the code is released under the MIT license.
- [6] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13142–13153, 2023. URL https://huggingface.co/datasets/allenai/objaverse. The use of the dataset as a whole is licensed under the ODC-By v1.0 license. Individual objects in Objaverse are all licensed as creative commons distributable objects, and may be under the following licenses: CC-BY 4.0 721K objects CC-BY-NC 4.0 25K objects CC-BY-NC-SA 4.0 52K objects CC-BY-SA 4.0 16K objects CC0 1.0 3.5K objects.
- [7] Yuxin Fang, Wen Wang, Binhui Xie, Quan Sun, Ledell Wu, Xinggang Wang, Tiejun Huang, Xinlong Wang, and Yue Cao. Eva: Exploring the limits of masked visual representation learning at scale. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19358–19369, 2023.
- [8] Yipeng Gao, Zeyu Wang, Wei-Shi Zheng, Cihang Xie, and Yuyin Zhou. Sculpting holistic 3d representation in contrastive language-image-3d pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22998–23008, June 2024.
- [9] Yining Hong, Haoyu Zhen, Peihao Chen, Shuhong Zheng, Yilun Du, Zhenfang Chen, and Chuang Gan. 3d-llm: Injecting the 3d world into large language models. *Advances in Neural Information Processing Systems*, 36:20482–20494, 2023.
- [10] Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pages 2790–2799. PMLR, 2019.
- [11] 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 *International conference on machine learning*, pages 4904–4916. PMLR, 2021.
- [12] Brenden M. Lake, Ruslan Salakhutdinov, and Joshua B. Tenenbaum. Human-level concept learning through probabilistic program induction. *Science*, 350(6266):1332-1338, 2015. doi: 10.1126/science.aab3050. URL https://www.science.org/doi/abs/10.1126/science.aab3050.

- [13] Weixian Lei, Yixiao Ge, Kun Yi, Jianfeng Zhang, Difei Gao, Dylan Sun, Yuying Ge, Ying Shan, and Mike Zheng Shou. Vit-lens: Towards omni-modal representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 26647–26657, 2024. URL https://github.com/TencentARC/ViT-Len. This work is licensed under the Apache 2.0 License.
- [14] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. Blip: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *International conference on machine learning*, pages 12888–12900. PMLR, 2022.
- [15] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023.
- [16] Dingkang Liang, Tianrui Feng, Xin Zhou, Yumeng Zhang, Zhikang Zou, and Xiang Bai. Parameter-efficient fine-tuning in spectral domain for point cloud learning. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [17] Minghua Liu, Ruoxi Shi, Kaiming Kuang, Yinhao Zhu, Xuanlin Li, Shizhong Han, Hong Cai, Fatih Porikli, and Hao Su. Openshape: Scaling up 3d shape representation towards openworld understanding. *Advances in neural information processing systems*, 36, 2024. URL https://github.com/Colin97/OpenShape_code. This work is licensed under the Apache 2.0 License.
- [18] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Advances in neural information processing systems*, 30, 2017.
- [19] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [20] Mikaela Angelina Uy, Quang-Hieu Pham, Binh-Son Hua, Duc Thanh Nguyen, and Sai-Kit Yeung. Revisiting point cloud classification: A new benchmark dataset and classification model on real-world data. In *International Conference on Computer Vision (ICCV)*, 2019. URL https://hkust-vgd.github.io/scanobjectnn/. This dataset is released under the MIT License.
- [21] Zehan Wang, Ziang Zhang, Hang Zhang, Luping Liu, Rongjie Huang, Xize Cheng, Hengshuang Zhao, and Zhou Zhao. Omnibind: Large-scale omni multimodal representation via binding spaces, 2024. URL https://arxiv.org/abs/2407.11895.
- [22] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaoou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015. URL https://modelnet.cs.princeton.edu/. Academic research, see website for more details.
- [23] Le Xue, Mingfei Gao, Chen Xing, Roberto Martín-Martín, Jiajun Wu, Caiming Xiong, Ran Xu, Juan Carlos Niebles, and Silvio Savarese. Ulip: Learning a unified representation of language, images, and point clouds for 3d understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1179–1189, 2023.
- [24] An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, et al. Qwen2. 5 technical report. arXiv preprint arXiv:2412.15115, 2024. URL https://huggingface.co/Qwen/Qwen2.5-VL-32B-Instruct-AWQ. Qwen2.5 is licensed under the Apache 2.0 license.
- [25] Li Yi, Vladimir G. Kim, Duygu Ceylan, I-Chao Shen, Mengyan Yan, Hao Su, Cewu Lu, Qixing Huang, Alla Sheffer, and Leonidas Guibas. A scalable active framework for region annotation in 3d shape collections. *SIGGRAPH Asia*, 2016.

- [26] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In Proceedings of the IEEE/CVF international conference on computer vision, pages 6023–6032, 2019.
- [27] Yihan Zeng, Chenhan Jiang, Jiageng Mao, Jianhua Han, Chaoqiang Ye, Qingqiu Huang, Dit-Yan Yeung, Zhen Yang, Xiaodan Liang, and Hang Xu. Clip2: Contrastive language-image-point pretraining from real-world point cloud data. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15244–15253, 2023.
- [28] Hongyi Zhang. mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412, 2017.
- [29] Jinlai Zhang, Lyujie Chen, Bo Ouyang, Binbin Liu, Jihong Zhu, Yujin Chen, Yanmei Meng, and Danfeng Wu. Pointcutmix: Regularization strategy for point cloud classification. *Neurocomputing*, 505:58–67, 2022. ISSN 0925-2312. doi: https://doi.org/10.1016/j.neucom.2022.07.049. URL https://www.sciencedirect.com/science/article/pii/S0925231222009092.
- [30] Zhihao Zhang, Shengcao Cao, and Yu-Xiong Wang. Tamm: Triadapter multi-modal learning for 3d shape understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21413–21423, June 2024.
- [31] Junsheng Zhou, Jinsheng Wang, Baorui Ma, Yu-Shen Liu, Tiejun Huang, and Xinlong Wang. Uni3d: Exploring unified 3d representation at scale. In *International Conference on Learning Representations (ICLR)*, 2024. URL https://github.com/baaivision/Uni3D. This work is licensed under the MIT License.
- [32] Xin Zhou, Dingkang Liang, Wei Xu, Xingkui Zhu, Yihan Xu, Zhikang Zou, and Xiang Bai. Dynamic adapter meets prompt tuning: Parameter-efficient transfer learning for point cloud analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14707–14717, 2024.