

ASSIST-3D: Adapted Scene Synthesis for Class-Agnostic 3D Instance Segmentation

Shengchao Zhou¹, Jiehong Lin^{1†}, Jiahui Liu¹, Shizhen Zhao¹, Chirui Chang¹ and Xiaojuan Qi^{1†}

¹The University of Hong Kong
{sczhou, liujh, zhaosz, xjqj}@eee.hku.hk, {mortimer.jh.lin, justinsheunghk}@gmail.com

Abstract

Class-agnostic 3D instance segmentation tackles the challenging task of segmenting all object instances, including previously unseen ones, without semantic class reliance. Current methods struggle with generalization due to the scarce annotated 3D scene data or noisy 2D segmentations. While synthetic data generation offers a promising solution, existing 3D scene synthesis methods fail to simultaneously satisfy geometry diversity, context complexity, and layout reasonability, each essential for this task. To address these needs, we propose an Adapted 3D Scene Synthesis pipeline for class-agnostic 3D Instance Segmentation, termed as **ASSIST-3D**, to synthesize proper data for model generalization enhancement. Specifically, ASSIST-3D features three key innovations, including 1) **Heterogeneous Object Selection** from extensive 3D CAD asset collections, incorporating randomness in object sampling to maximize geometric and contextual diversity; 2) **Scene Layout Generation** through LLM-guided spatial reasoning combined with depth-first search for reasonable object placements; and 3) **Realistic Point Cloud Construction** via multi-view RGB-D image rendering and fusion from the synthetic scenes, closely mimicking real-world sensor data acquisition. Experiments on ScanNetV2, ScanNet++, and S3DIS benchmarks demonstrate that models trained with ASSIST-3D-generated data significantly outperform existing methods. Further comparisons underscore the superiority of our purpose-built pipeline over existing 3D scene synthesis approaches.

1 Introduction

3D instance segmentation tackles the fundamental problem of identifying and delineating individual object instances in cluttered 3D scenes, with wide applications spanning autonomous driving (Guo et al. 2024; Jiang et al. 2024), robotic navigation (Xie et al. 2021; Yilmaz et al. 2024) and virtual reality (Liu et al. 2023). This field has been traditionally dominated by class-aware approaches (Jiang et al. 2020; Chen et al. 2021; Vu et al. 2022; Lu et al. 2023; Kolodiaznyi et al. 2024), which simultaneously segment and classify objects belonging to predefined categories. However, the heavy reliance on annotation restricts these methods to a limited set of object classes, often just dozens, while thousands of diverse, unseen categories exist in real-world scenarios.

The class-aware bottleneck has driven recent interest in class-agnostic 3D instance segmentation (Xu et al. 2023; Huang et al. 2024; Rozenberszki et al. 2024; Yang et al. 2024b), where the goal shifts to segment all object instances without categorical constraints, including those never encountered during training. Existing solutions primarily explore two strategies, as illustrated in Fig. 1 (a) and (b). One kind of approaches (Takmaz et al. 2023; Nguyen et al. 2024) modify traditional class-aware 3D architectures by replacing their multi-class classifiers with binary objectness predictors, though their generalization remains hampered by the limited diversity of training objects. Alternatively, methods like those based on 2D foundation models (e.g., SAM (Kirillov et al. 2023)) first perform class-agnostic segmentation on multi-view RGB images, then project and fuse the results into 3D according to depth information. While the merged 3D masks can be refined to be final outputs (Xu et al. 2023) or used as pseudo labels for training (Huang et al. 2024), they often suffer from inherent limitations including 2D segmentation inaccuracies and multi-view fusion inconsistencies, ultimately restricting performance and generalization potential.

This work establishes data diversity as the key driver for advancing class-agnostic 3D instance segmentation, where generalization capability directly correlates with the variety and volume of training instances. To address the core challenge of data scarcity, we present **3D scene synthesis** as a solution. Unlike real-world data collection, which faces costly acquisition, incomplete annotations, and limited scene complexity, synthetic data generation offers a scalable and cost-effective alternative with full annotation control. This paradigm shift enables the creation of large-scale, task-specific 3D datasets with rich supervision, paving the way for generalization.

To improve data effectiveness, we further identify three fundamental principles for synthetic data, including 1) **geometry diversity** to ensure broad shape variations, 2) **context complexity** to capture the variety of objects coexisting in scenes, and 3) **layout reasonability** to maintain physically plausible global arrangements. Existing synthesis methods (Bautista et al. 2022; Li et al. 2024; Liu et al. 2024; Zhang et al. 2024), however, fall short of these criteria, as they either violate the first two principles by biasing towards common real-world objects and their typical neighbors (Yang et al. 2024d), or break the third principle through physically implausible random arrangements (Rao et al. 2021).

[†]Corresponding authors.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

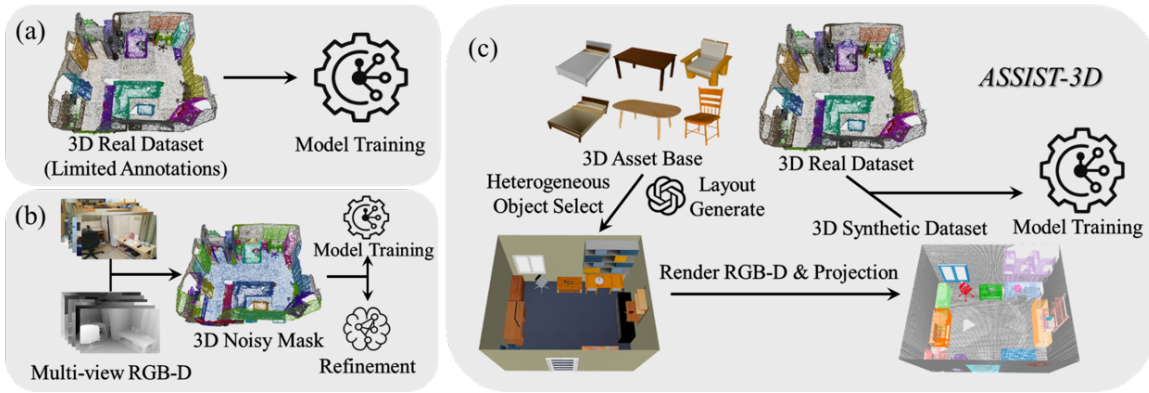


Figure 1: (a) Methods adapted from conventional class-aware approaches rely heavily on limited real-world datasets and face severe data scarcity. (b) Methods that use 2D foundation models for multi-view image segmentation (representing 3D scenes) struggle with 2D segmentation errors or cross-view inconsistency issues. (c) Our proposed **ASSIST-3D** addresses these limitations by generating high-quality, fully-annotated synthetic 3D scenes to improve model generalization.

To this end, we propose an **Adapted Scene Synthesis pipeline** for class-agnostic 3D Instance Segmentation (**ASSIST-3D**) that satisfies all outlined principles. Specifically, ASSIST-3D comprises three stages, including **Heterogeneous Object Selection, Scene Layout Generation, and Realistic Point Cloud Construction**. First, ASSIST-3D utilizes a large-scale collection of 3D CAD models, e.g. those from Objaverse (Deitke et al. 2023a,b), as its asset base, and maximizes selection randomness of objects to be placed in one scene to ensure diverse object compositions and contextual arrangements, satisfying the first two principles. Second, ASSIST-3D employs large language models for their human-aligned preference to infer plausible spatial relationships and applies a depth-first search strategy for sequential object placement, ensuring layout reasonability. Finally, ASSIST-3D mimics real-world point cloud acquisition by rendering multi-view RGB-D images of scenes and merging their partial point clouds to produce realistic point clouds along with instance mask annotations.

We evaluate ASSIST-3D by combining its generated synthetic data with existing real-world datasets to train Mask3D (Schult et al. 2023) adapted for class-agnostic segmentation. Experiments on ScanNetV2 (Dai et al. 2017), ScanNet++ (Yeshwanth et al. 2023), and S3DIS (Armeni et al. 2016) demonstrate consistent performance improvements and enhanced generalization capability with the integration of our synthetic data. The resulting model also significantly outperforms existing state-of-the-art methods. Furthermore, we conduct comprehensive ablation studies to validate the effectiveness of each individual component, as well as to assess the contribution of the three guiding principles. ASSIST-3D also outperforms existing 3D scene synthesis methods on the target task, including those that overly rely on real-world object selection and neighborhood configurations (e.g., Holodeck (Yang et al. 2024d)), and those lacking plausible scene layouts (e.g., RandomRooms (Rao et al. 2021)), both of which fall short of fully adhering to the proposed principles. In summary, our technical contributions are three-fold:

- We repurpose 3D scene synthesis as a solution to the data

scarcity challenge in class-agnostic 3D instance segmentation and introduce ASSIST-3D, a novel, purpose-built pipeline for generating fully annotated synthetic 3D scene data to enhance model generalization.

- ASSIST-3D consists of three stages, including Heterogeneous Object Selection, Scene Layout Generation, and Realistic Point Cloud Construction, to ensure geometry diversity, context complexity, and layout reasonability.
- Experiments on the ScanNetV2, ScanNet++, and S3DIS benchmarks demonstrate the effectiveness of the synthetic data generated by ASSIST-3D, with the resulting model significantly outperforming existing methods. Additional comparisons highlight the advantages of our purpose-built pipeline over existing 3D scene synthesis approaches.

2 Related Work

Class-agnostic 3D Instance Segmentation Class-agnostic 3D instance segmentation advances beyond the traditional class-aware setting by removing category dependencies, enabling segmentation of arbitrary objects in 3D scenes, including objects from categories not seen during training.

Current approaches address the generalization challenge through two main strategies. The first strategy (Takmaz et al. 2023; Yan et al. 2024; Nguyen et al. 2024; Boudjoghra et al. 2025) modifies conventional class-aware 3D architectures by eliminating category dependencies and instead learning the concept of "objectness" to differentiate object segments from non-object regions. For instance, OpenMask3D (Takmaz et al. 2023) adapts the framework of Mask3D by replacing its multi-class classifier with a binary objectness predictor. While these methods leverage existing annotated 3D scene datasets for training, their generalization capability remains constrained by the limited quantity of annotated objects available in current datasets. The second strategy (Xu et al. 2023; Huang et al. 2024; Rozenberszki et al. 2024; Yin et al. 2024; Yang et al. 2024b) capitalizes on the strong generalization performance of 2D foundation models, which benefit from large-scale training data and have demonstrated robust per-

formance across diverse downstream tasks. These methods typically employ 2D foundation models to segment objects in multi-view rendered images of a 3D scene and then backproject and fuse the results into 3D space. For example, SAI3D (Yin et al. 2024) introduces a hierarchical region-growing algorithm to refine 3D masks from SAM-generated (Kirillov et al. 2023) 2D segmentations, while Segment3D (Huang et al. 2024) uses the 2D results as pseudo ground-truth masks for 3D model pre-training. However, these approaches inherit challenges from 2D segmentation errors and multi-view inconsistencies, ultimately limiting their effectiveness.

In this work, we propose a novel alternative that enhances generalization capabilities through cost-effective synthesis of 3D scenes with diverse object assets, simultaneously addressing the data scarcity problem while circumventing the inherent limitations of 2D-based approaches.

3D Scene Synthesis Recently, 3D scene synthesis methods have made significant progress (Hu et al. 2024; Li et al. 2024; Öcal et al. 2024; Zhou et al. 2024). For example, Holodeck (Yang et al. 2024d) will prompt large language models (LLM) to select objects to be placed and design their positions, resulting in high visual quality and better alignment with user instructions. In contrast, RandomRooms (Rao et al. 2021) adopts a random object selection and placement strategy, leading to diverse visual appearances. However, these methods are not suitable for our targeted task, as LLM will bias towards and repeat selecting most common objects so that violating geometry diversity and context complexity while random selection and placement will break layout reasonability. By contrast, we propose a new synthesis pipeline that enhances geometry diversity and context complexity while ensuring physically reasonable layouts, making the synthetic dataset better aligned with our task and leading to better model performance as shown in experiments.

3 Methodology

3.1 Problem Formulation and Learning Paradigm

Class-agnostic 3D instance segmentation addresses the task of segmenting all object instances in a 3D scene (typically represented as a point cloud) without identifying its semantic category. Formally, given a training dataset $\mathcal{D} = \{(\mathcal{P}_i, \hat{\mathcal{M}}_i)\}_{i=1}^N$ consisting of N scenes, where \mathcal{P} represents a 3D point cloud and $\hat{\mathcal{M}}$ denotes the set of ground-truth instance masks for all objects, we train a model Φ to predict a set of instance masks \mathcal{M} for an input point cloud, supervised by the objective \mathcal{L} :

$$\min_{\Phi} \mathcal{L}(\mathcal{M}, \hat{\mathcal{M}}) = \mathcal{L}(\Phi(\mathcal{P}), \hat{\mathcal{M}}). \quad (1)$$

At inference time, Φ is expected to segment all object instances, including those unseen in \mathcal{D} , in arbitrary 3D scenes, posing significant generalization challenges.

Data diversity is a key factor to achieve generalization. However, real-world 3D datasets \mathcal{D}^r often lack sufficient diversity due to the difficulties in 3D data acquisition and labeling. This limitation leads to suboptimal performance in methods trained exclusively on such constrained datasets (Takmaz et al. 2023; Nguyen et al. 2024). Alternative approaches (Huang et al. 2024; Yin et al. 2024), which leverage

2D foundation models by segmenting objects in multi-view RGB-D images and lifting the results to 3D, face challenges like 2D segmentation errors and cross-view inconsistency. To address the data scarcity problem, we propose **ASSIST-3D**, an innovative 3D scene synthesis pipeline that procedurally synthesizes diverse, automatically annotated data \mathcal{D}^s . By combining \mathcal{D}^r and \mathcal{D}^s , we formulate the following objective to enhance data diversity and improve generalization:

$$\min_{\Phi} \mathcal{L}(\Phi(\mathcal{P}^r), \hat{\mathcal{M}}^r) + \alpha \mathcal{L}(\Phi(\mathcal{P}^s), \hat{\mathcal{M}}^s), \quad (2)$$

where superscripts ‘r’ and ‘s’ denote real and synthetic sources, and α is a balancing parameter.

3.2 ASSIST-3D: Adapted Scene Synthesis for Instance Segmentation in 3D

To generate the target synthetic 3D scene dataset \mathcal{D}^s , we propose **ASSIST-3D**, a purpose-built 3D scene generation pipeline for class-agnostic 3D instance segmentation.

ASSIST-3D comprises three main stages, including **Heterogeneous Object Selection**, **Scene Layout Generation**, and **Realistic Point Cloud Construction**. Fig. 2 gives a total illustration of the process. First, ASSIST-3D utilizes a large collection of 3D CAD models as its asset base and selects heterogeneous objects by introducing randomness in selection of objects to be placed, ensuring rich variation in placed objects and contextual arrangements. Second, to generate physically plausible scenes, ASSIST-3D leverages LLM-based guidance, prompting GPT-4 to infer spatial relationships among objects and employing a depth-first search strategy for sequential placement. Finally, ASSIST-3D simulates real-world point cloud acquisition process by rendering multi-view RGB-D images of the synthetic scenes and merging the resulting partial point clouds. This process produces realistic point clouds along with instance mask annotations, together constituting the synthetic dataset.

Heterogeneous Object Selection To ensure geometry diversity, we begin with a large collection of 3D models as our digital asset base, forming the foundation for data variety in our task. We select a subset of Objaverse (Deitke et al. 2023b), which contains 50,000 3D objects across 800 classes. Each object has been pre-processed to align its front view with the canonical coordinate axes. The asset is categorized into three groups, including 1) objects typically placed on the floor (e.g., furniture), 2) objects commonly mounted on or leaning against walls (e.g., photos, bookshelves), and 3) objects that can be placed on either of the first two groups. These groups are denoted as $\mathcal{O}_{\text{floor}}$, $\mathcal{O}_{\text{wall}}$, and \mathcal{O}_{obj} , respectively.

Leveraging the asset base, we can construct synthetic 3D scenes with high object diversity through simply randomized selection. Specifically, for each scene, we uniformly sample M_1 objects from $\mathcal{O}_{\text{floor}}$ and M_2 objects from $\mathcal{O}_{\text{wall}}$, denoted as $\mathbf{o}_{\text{floor}}$ and \mathbf{o}_{wall} , respectively. For each object in $\mathbf{o}_{\text{floor}}$ and \mathbf{o}_{wall} , we then sample 5 additional objects from \mathcal{O}_{obj} to place on top of it, yielding a total of $5(M_1 + M_2)$ objects, denoted as \mathbf{o}_{obj} . This sampling strategy effectively enhances the context complexity by disrupting conventional class co-occurrence patterns and increasing object diversity in one




Figure 2: An overview of our proposed ASSIST-3D. First, ASSIST-3D selects heterogeneous objects from a 3D asset base like Objaverse, satisfying geometry diversity and context complexity. Next, it leverages GPT-4 to design the arrangement of these objects within the scene, ensuring layout reasonability. Finally, it constructs realistic point clouds by mimicking the construction procedure of real dataset to reduce the domain gap and enhance performance.

scene. To compensate for annotation incompleteness in real-world dataset \mathcal{D}^r , we implement a complementary sampling strategy that preferentially selects object categories present in \mathcal{D}^r by adjusting their sampling probabilities. During scene synthesis, we alternate between these two strategies.

Scene Layout Generation With the selected objects, the next step is to decide their arrangements while ensuring reasonable scene layouts, bringing the synthetic scene closer to real-world environments. To achieve this target, we leverage a large language model (LLM) to infer physically plausible spatial relationships between objects and adopt a depth-first search (DFS) strategy for sequential object placement.

More specifically, we first prompt GPT-4 to derive reasonable spatial relationships among objects in $\mathcal{o}_{\text{floor}}$, $\mathcal{o}_{\text{wall}}$, and each group within \mathcal{o}_{obj} . For structured placement, we prompt GPT-4 to sequentially determine object positions within each group, with the location of each object conditioned only on those placed earlier. Due to the limited spatial reasoning capabilities of GPT-4, its role is restricted to provide only object orientations and positions relative to previously objects, rather than absolute coordinates of objects in the scene. Our GPT-4 prompt template is provided in the appendix.

The sequential prompting enables our subsequent DFS-based placement strategy. Taking $\mathcal{o}_{\text{floor}}$ as an example, we begin by discretizing the floor space into uniform grids and initializing placement with the first object. At each iteration, we identify feasible grids for the current object, discarding those that violate spatial relationships provided by GPT-4. The object is then placed in one of the valid grids, and the process recurses with the next object. If no valid grid exists for the current object, the configuration of placement so far

will be saved as a potential solution and backtrack to reposition the previous object. Among all solutions, we select the one that accommodates the most objects.

The same procedure is applied to objects in $\mathcal{o}_{\text{wall}}$ and to each group within \mathcal{o}_{obj} , with placement occurring either on the wall or on supporting objects from $\mathcal{o}_{\text{floor}} \cup \mathcal{o}_{\text{wall}}$, as appropriate. Once the scene layout is generated, we place the objects accordingly to construct the final scene mesh.

Realistic Point Cloud Construction After generating the mesh of a synthetic 3D scene, a straightforward way to obtain its corresponding point cloud data is to sample points directly from the mesh surface. However, such point clouds tend to be overly uniform and lack the noise, density variations and occlusion commonly existing in real datasets. In practice, real-world point clouds are typically reconstructed via Simultaneous Localization and Mapping (SLAM) systems, which fuse multi-view RGB-D or LiDAR measurements from different vantage points in the scene, resulting in non-uniform sampling and occlusion. To bridge this domain gap, we propose a realistic point cloud construction method that closely mimics real-world acquisition processes.

We first search for optimal scanning positions by uniformly partitioning the mid-height plane of the scene into $0.1 \times 0.1 \text{ m}^2$ regions while excluding areas occupied by objects. Using Farthest Point Sampling (FPS), we select five optimal vantage points for a comprehensive scene coverage. At each point, we render 12 RGB-D images by rotating camera in 30° increments (360° total), yielding 60 images in total.

The subsequent process involves projecting depth maps into partial point clouds using camera intrinsics, then transforming and aggregating them into a unified global coordi-

Method	ScanNet++			S3DIS			ScanNetV2		
	AP	AP ₅₀	AP ₂₅	AP	AP ₅₀	AP ₂₅	AP	AP ₅₀	AP ₂₅
Baseline	12.0	21.7	32.7	13.6	23.2	41.9	46.6	69.0	81.4
Unscene3D (Rozenberszki et al. 2024)	9.8	18.6	29.3	21.4	40.3	52.6	15.9	32.2	58.5
SAI3D (Yin et al. 2024)	17.1	31.1	49.5	24.8	42.4	57.9	30.8	50.5	70.6
Segment3D (Huang et al. 2024)	12.0	22.7	37.8	21.7	40.6	53.4	28.3	43.7	53.6
Open3DIS (Nguyen et al. 2024)	17.9	30.4	39.7	23.7	41.7	56.3	31.5	45.2	55.0
SA3DIP (Yang et al. 2024b)	19.6	32.4	42.5	25.7	42.4	58.6	41.6	64.6	81.3
ASSIST-3D (Ours)	22.2	35.5	45.2	29.0	43.9	61.6	48.1	70.7	83.4

Table 1: Performance comparison among different class-agnostic instance segmentation methods on ScanNet++ (Yeshwanth et al. 2023), S3DIS (Armeni et al. 2016), and ScanNetV2 (Dai et al. 2017).

nate system according to camera poses. To further enhance realism, we apply voxel-based downsampling to discretize point cloud space and retains one random point per voxel.

Finally, we generate semantic labels for the aggregated point cloud by assigning each point to its nearest object in the scene, producing a complete set of segmentation masks. Each point cloud paired with its corresponding mask set constitutes a training sample in the synthetic dataset.

4 Experiments

4.1 Experimental Settings

Real-world Datasets We conduct experiments on three widely used benchmarks, including ScanNetV2 (Dai et al. 2017), ScanNet++ (Yeshwanth et al. 2023), and S3DIS (Armeni et al. 2016). ScanNetV2 provides dense instance masks for 18 categories, with official splits of 1,201 training and 312 validation scans. ScanNet++ enhances ScanNetV2 with increased point cloud number and labeled categories. S3DIS covers six indoor areas, with 13 semantic categories. We incorporate the training set of ScanNetV2 with ASSIST-3D-generated dataset for training and evaluate the model’s in-domain performance on the validation set of ScanNetV2. To assess generalization, we test our method on ScanNet++ and S3DIS, which contain unseen object categories and follow different data distributions with ScanNetV2.

Evaluation Metrics We follow (Yin et al. 2024; Huang et al. 2024) to evaluate our method using the standard Average Precision (AP) metric. Specifically, we compute the mean AP score across Intersection-over-Union (IoU) thresholds ranging from 50% to 95% with 5% increments (denoted as AP). Additionally, we report performance at specific IoU thresholds of 50% and 25% (AP₅₀ and AP₂₅ respectively). Consistent with the class-agnostic nature of our task, all reported metrics assess only instance mask accuracy, independent of semantic category recognition.

We further introduce metrics to quantify the synthetic data in accordance with the three proposed principles. To assess geometry diversity, we use 3DShape2VecSet (Zhang et al. 2023) to extract the features of objects in the synthetic scenes. Since 3DShape2VecSet is a 3D VAE model trained for object reconstruction, its latent features effectively capture object

geometry information. We compute the entropy of the resulting feature distribution, where higher entropy indicates greater geometry diversity. To measure context complexity, given one class, we compute for each left class its occurrence probability in a scene containing the given class, and then average the maximal ones across all classes. Lower value suggests more complex and diverse contextual relationships among objects. For assessing layout reasonability, we follow prior works (Yang et al. 2024a,c) by rendering multi-view images of each scene and prompting GPT-4o to score the layout on a scale from 0 to 100. A higher score indicates a more reasonable and realistic spatial arrangement. The exact prompt used for GPT-4o is provided in the appendix.

Implementation Details For each scene, the selected number of objects on floor and wall are $M_1 = 100$ and $M_2 = 50$. Using our proposed pipeline, we synthesize a dataset of 2,000 scenes with approximately 134,000 object instances in total, averaging 67 objects per scene. Half scenes are generated according to complementary strategy with 0.7 probability to select object categories in \mathcal{D}^r . We combine this synthetic dataset with ScanNetV2 to train Mask3D (Schult et al. 2023) by replacing its multi-class classifier with a binary objectness classifier (following (Huang et al. 2024)) and using binary cross-entropy loss, dice loss and mask loss as optimization objectives. Mask3D trained purely on ScanNetV2 will serve as the baseline. Training runs for 600 epochs with a batch size of 36 (distributed across 6 A100 GPUs). To mitigate domain gaps, we set the scaling factor α to 0.5 for optimization.

4.2 Comparison with State-of-The-Arts

We first compare the performance of our ASSIST-3D method with other state-of-the-art methods, including Open3DIS (Nguyen et al. 2024), Unscene3D (Rozenberszki et al. 2024), Segment3D (Huang et al. 2024), SAI3D (Yin et al. 2024), and SA3DIP (Yang et al. 2024b). Among these methods, SAI3D and SA3D leverage SAM (Kirillov et al. 2023) to generate 2D masks from multi-view RGB-D images, then refine 3D masks through projection and aggregation of these 2D predictions. Other methods build upon Mask3D, trained on ScanNetV2 using either limited manual annotations (e.g., Open3DIS) or pseudo-3D masks derived from SAM-generated 2D masks (e.g., Unscene3D and Segment3D). For a fair comparison, we

Method	Geometry Diversity (\uparrow)	Context Complexity (\downarrow)	Layout Reasonability (\uparrow)	ScanNet++			S3DIS		
				AP	AP ₅₀	AP ₂₅	AP	AP ₅₀	AP ₂₅
Holodeck	× (0.85)	× (0.38)	✓ (72)	14.2	24.6	36.3	18.2	28.4	48.3
RandomRooms	✓ (4.37)	✓ (0.04)	× (23)	16.6	26.8	38.4	23.5	41.5	56.4
ASSIST-3D w/o Content	✓ (4.03)	× (0.54)	✓ (65)	19.4	31.9	41.8	26.4	42.8	58.6
ASSIST-3D (Ours)	✓ (4.15)	✓ (0.08)	✓ (62)	22.2	35.5	45.2	29.0	43.9	61.6

Table 2: Performance comparison among different 3D scene synthesis methods on ScanNet++ (Yeshwanth et al. 2023) and S3DIS (Armeni et al. 2016). Evaluation metrics of geometry diversity, context complexity, and layout reasonability are feature entropy, object co-occurrence probability, GPT-4o prompted layout score, respectively, as detailed in Sec. 4.1.

also train a version of Mask3D on a combination of annotated ScanNetV2 data and our synthetic dataset, and establish a baseline using only the original ScanNetV2 annotations.

As shown in Table 1, our method achieves the best performance across all the benchmarks. Notably, it enables Mask3D to retain strong accuracy on the in-domain ScanNetV2, where other methods struggle due to noisy 2D masks and limited 3D annotations. While existing methods exhibit decent generalization and outperform the baseline on ScanNet++ and S3DIS, our approach further enhances Mask3D with synthetic data generated by our proposed ASSIST-3D pipeline.

We also compare ASSIST-3D with other 3D scene synthesis methods, including Holodeck (Yang et al. 2024d), which differs in selecting objects for placement by leveraging GPT-4 and thus favors common objects and typical co-occurrence patterns, and RandomRooms (Rao et al. 2021), which relies on completely random object selection and placement. As shown in Table 2, synthetic data generated by Holodeck lacks geometry diversity and context complexity, while that generated by RandomRooms suffers from poor layout reasonability. To isolate the effects of geometry diversity and context complexity, we further introduce a variant of ASSIST-3D that disables context complexity. Specifically, each object class in the 3D asset base is paired with a co-occurring class. The paired class will also be selected with 66% probability when selecting objects from one class for placement. The performance gains achieved by the full ASSIST-3D pipeline on ScanNet++ and S3DIS underscore the importance of incorporating all three principles in generating high-quality synthetic data for class-agnostic 3D instance segmentation.

4.3 Ablation Studies

Effect of Geometry Diversity In Section 4.2, we have demonstrated the importance of geometry diversity in synthetic data generation. Here, we further validate its impact on model performance through controlled experiments. We begin by extracting features for all 3D assets using 3DShape2VecSet (Zhang et al. 2023), and apply the k-means clustering algorithm to group them into five distinct clusters based on geometric similarity. To vary geometry diversity, we synthesize training datasets using an increasing number of clusters to select objects from. As in Table 3, model performance consistently improves with greater geometry diversity.

To further assess the effectiveness of geometry diversity, we evaluate performance on out-of-domain objects, specifi-

Cluster Number	ScanNet++			S3DIS		
	AP	AP ₅₀	AP ₂₅	AP	AP ₅₀	AP ₂₅
1	14.6	24.7	36.1	16.7	28.2	47.2
2	16.8	27.3	38.9	21.1	34.7	52.3
3	17.8	29.6	40.3	24.1	39.7	56.4
4	20.1	32.4	43.2	27.3	41.6	58.7
5 (Ours)	22.2	35.5	45.2	29.0	43.9	61.6

Table 3: Performance comparison of ASSIST-3D with different numbers of object clusters for data synthesis.

Cluster Number	ScanNet++			S3DIS		
	AR	AR ₅₀	AR ₂₅	AR	AR ₅₀	AR ₂₅
1	3.2	6.7	8.4	13.3	23.8	35.2
2	6.7	11.3	15.2	16.1	28.1	39.2
3	8.6	15.2	20.9	18.6	32.4	44.6
4	10.6	18.4	25.3	21.2	36.4	48.0
5 (Ours)	12.4	21.6	29.8	23.6	40.2	51.8

Table 4: Performance comparison of ASSIST-3D using different numbers of object clusters for data synthesis on test objects with geometries distinct from those in ScanNetV2 (Dai et al. 2017).

cally, those in ScanNet++ and S3DIS that are geometrically distinct from those in ScanNetV2. We extract features for all objects in ScanNetV2 using 3DShape2VecSet and compute the mean feature vector \mathcal{F}_{avg} . We then calculate the maximum L2 distance d_{max} from \mathcal{F}_{avg} to any object feature in ScanNetV2. For each object in ScanNet++ or S3DIS, we extract its feature \mathcal{F} and compute its L2 distance d to \mathcal{F}_{avg} . An object is considered out-of-domain and selected for evaluation if $d > d_{max}$. We report the mean Average Recall (AR) over IoU thresholds from 50% to 95% (in 5% increments), as well as AR₅₀ and AR₂₅, which represent recall at IoU thresholds of 50% and 25%, respectively. The results, presented in Table 4, are consistent with those in Table 3, further confirming the effectiveness of geometry diversity in enhancing generalization to novel object geometries.




Figure 3: Performance curves of average precision on ScanNet++ (Yeshwanth et al. 2023) and S3DIS (Armeni et al. 2016) with varying numbers of object classes used for synthesis and different amounts of synthetic training data.

Level	Prob	ScanNet++			S3DIS		
		AP	AP ₅₀	AP ₂₅	AP	AP ₅₀	AP ₂₅
1	100%	17.2	29.4	36.3	24.3	42.0	57.6
2	66%	19.4	31.9	41.8	26.4	42.8	58.6
3	33%	21.6	34.7	44.6	28.2	43.4	60.7
4 (Ours)	0%	22.2	35.5	45.2	29.0	43.9	61.6

Table 5: Performance comparison of ASSIST-3D with different levels of context complexity. ‘Prob’ denotes the probability of simultaneously sampling objects from paired class of one class for placement to control context complexity.

Point Cloud Construction	ScanNet++			S3DIS		
	AP	AP ₅₀	AP ₂₅	AP	AP ₅₀	AP ₂₅
-	12.0	21.7	32.7	13.6	23.2	41.9
Mesh-sampled	14.2	24.6	36.3	18.8	30.5	51.5
Rendering-based	22.2	35.5	45.2	29.0	43.9	61.6

Table 6: Performance comparison of various strategies of point cloud construction in ASSIST-3D.

Effect of Context Complexity In this part, we design experiments to quantify the effect of context complexity. We begin by pairing each class in our 3D asset base with a co-occurring class. For classes present in ScanNetV2, we pair them with their most frequently co-occurring counterparts based on the ScanNetV2 training data, while the remaining classes are paired randomly. To control context complexity and co-occurrence probability, we define four levels by varying the probability of simultaneously selecting from paired class when selecting objects from one class for placement: 100%, 66%, 33%, and 0% (i.e., no co-occurrence enforcement). The corresponding object co-occurrence probabilities and performance results are reported in Table 5, which show that increasing context complexity consistently leads to improved performance on both ScanNet++ and S3DIS.

Effect of Realistic Point Cloud Construction As described in Section 3.2, we replicate the point cloud construc-

tion process commonly used in real-world datasets. This involves rendering RGB-D images from multiple viewpoints within each synthetic scene, projecting the depth pixels into 3D space, and applying voxel-based downsampling. This pipeline helps bridge the domain gap between synthetic and real data. To evaluate its effectiveness, we compare it against a baseline where point clouds are directly sampled from the synthetic meshes. As shown in Table 6, Mask3D trained on mesh-sampled point clouds offers limited improvement over the baseline without synthetic data, whereas our rendering-based approach yields significantly better performance.

Scalability of Synthetic Data In this part, we analyze the scalability of the synthetic data in terms of both the number of classes and the number of scenes. Specifically, we consider three levels for the number of classes used in scene synthesis: 100, 200, and 400; and four levels for the number of synthesized scenes: 500, 1000, 1500, and 2000. For each combination of class levels and scene levels, we generate a synthetic dataset and combine it with ScanNetV2 to train the adapted Mask3D. As illustrated in Fig. 3, the AP performance on both ScanNet++ and S3DIS consistently improves with an increasing number of classes. Furthermore, the benefit of synthesizing more scenes becomes more pronounced when the number of object classes is higher. These results demonstrate that scaling both the class diversity and dataset size enhances the model’s generalization ability.

5 Conclusion

In this paper, we propose ASSIST-3D for class-agnostic 3D instance segmentation. To address the challenges of limited annotations in real dataset and noisy 2D segmentations, we propose to synthesize 3D scene-level dataset for training. Rather than replicating previous 3D scene synthesis methods, we design a new process that simultaneously satisfies the principle of geometry diversity, context complexity and layout reasonability. Additionally, we mimic the point cloud construction of real datasets to reduce the domain gap between real and synthetic dataset. By training Mask3D on ScanNetV2 augmented with our synthetic dataset, it achieves the best performance on ScanNet++, S3DIS, and ScanNetV2.

Acknowledgments

The work has been supported by Hong Kong Research Grant Council - Early Career Scheme (Grant No. 27209621), General Research Fund Scheme (Grant No. 17202422, 17212923, 17215025), Theme-based Research (Grant No. T45-701/22-R) and Shenzhen Science and Technology Innovation Commission (SGDX20220530111405040). Part of the described research work is conducted in the JC STEM Lab of Robotics for Soft Materials funded by The Hong Kong Jockey Club Charities Trust.

References

- Armeni, I.; Sener, O.; Zamir, A. R.; Jiang, H.; Brilakis, I.; Fischer, M.; and Savarese, S. 2016. 3d semantic parsing of large-scale indoor spaces. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1534–1543.
- Bautista, M. A.; Guo, P.; Abnar, S.; Talbott, W.; Toshev, A.; Chen, Z.; Dinh, L.; Zhai, S.; Goh, H.; Ulbricht, D.; et al. 2022. Gaudi: A neural architect for immersive 3d scene generation. In *Advances in Neural Information Processing Systems*, volume 35, 25102–25116.
- Boudjoghra, M. E. A.; Dai, A.; Lahoud, J.; Cholakkal, H.; Anwer, R. M.; Khan, S.; and Khan, F. S. 2025. Open-YOLO 3D: Towards Fast and Accurate Open-Vocabulary 3D Instance Segmentation. In *The Thirteenth International Conference on Learning Representations*.
- Chen, S.; Fang, J.; Zhang, Q.; Liu, W.; and Wang, X. 2021. Hierarchical aggregation for 3d instance segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 15467–15476.
- Dai, A.; Chang, A. X.; Savva, M.; Halber, M.; Funkhouser, T.; and Nießner, M. 2017. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5828–5839.
- Deitke, M.; Liu, R.; Wallingford, M.; Ngo, H.; Michel, O.; Kusupati, A.; Fan, A.; Laforte, C.; Voleti, V.; Gadre, S. Y.; et al. 2023a. Objaverse-xl: A universe of 10m+ 3d objects. In *Advances in Neural Information Processing Systems*, volume 36, 35799–35813.
- Deitke, M.; Schwenk, D.; Salvador, J.; Weihs, L.; Michel, O.; VanderBilt, E.; Schmidt, L.; Ehsani, K.; Kembhavi, A.; and Farhadi, A. 2023b. Objaverse: A universe of annotated 3d objects. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 13142–13153.
- Guo, H.; Zhu, H.; Peng, S.; Wang, Y.; Shen, Y.; Hu, R.; and Zhou, X. 2024. Sam-guided graph cut for 3d instance segmentation. In *European Conference on Computer Vision*, 234–251.
- Hu, Z.; Iscen, A.; Jain, A.; Kipf, T.; Yue, Y.; Ross, D. A.; Schmid, C.; and Fathi, A. 2024. Scenecraft: An llm agent for synthesizing 3d scenes as blender code. In *Forty-first International Conference on Machine Learning*.
- Huang, R.; Peng, S.; Takmaz, A.; Tombari, F.; Pollefeys, M.; Song, S.; Huang, G.; and Engelmann, F. 2024. Segment3d: Learning fine-grained class-agnostic 3d segmentation without manual labels. In *European Conference on Computer Vision*, 278–295.
- Jiang, G.; Liu, J.; Wu, Y.; Liao, W.; He, T.; and Peng, P. 2024. Mwsis: Multimodal weakly supervised instance segmentation with 2d box annotations for autonomous driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 2507–2515.
- Jiang, L.; Zhao, H.; Shi, S.; Liu, S.; Fu, C.-W.; and Jia, J. 2020. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4867–4876.
- Kirillov, A.; Mintun, E.; Ravi, N.; Mao, H.; Rolland, C.; Gustafson, L.; Xiao, T.; Whitehead, S.; Berg, A. C.; Lo, W.-Y.; et al. 2023. Segment anything. In *Proceedings of the IEEE International Conference on Computer Vision*, 4015–4026.
- Kolodiazhnyi, M.; Vorontsova, A.; Konushin, A.; and Rukhovich, D. 2024. Oneformer3d: One transformer for unified point cloud segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 20943–20953.
- Li, H.; Shi, H.; Zhang, W.; Wu, W.; Liao, Y.; Wang, L.; Lee, L.-h.; and Zhou, P. Y. 2024. Dreamscene: 3d gaussian-based text-to-3d scene generation via formation pattern sampling. In *European Conference on Computer Vision*, 214–230.
- Liu, Y.; Hu, B.; Huang, J.; Tai, Y.-W.; and Tang, C.-K. 2023. Instance neural radiance field. In *Proceedings of the IEEE International Conference on Computer Vision*, 787–796.
- Liu, Y.; Li, X.; Li, X.; Qi, L.; Li, C.; and Yang, M.-H. 2024. Pyramid diffusion for fine 3d large scene generation. In *European Conference on Computer Vision*, 71–87.
- Lu, J.; Deng, J.; Wang, C.; He, J.; and Zhang, T. 2023. Query refinement transformer for 3d instance segmentation. In *Proceedings of the IEEE International Conference on Computer Vision*, 18516–18526.
- Nguyen, P.; Ngo, T. D.; Kalogerakis, E.; Gan, C.; Tran, A.; Pham, C.; and Nguyen, K. 2024. Open3dis: Open-vocabulary 3d instance segmentation with 2d mask guidance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4018–4028.
- Öcal, B. M.; Tatarchenko, M.; Karaoğlu, S.; and Gevers, T. 2024. SceneTeller: Language-to-3D Scene Generation. In *European Conference on Computer Vision*, 362–378.
- Rao, Y.; Liu, B.; Wei, Y.; Lu, J.; Hsieh, C.-J.; and Zhou, J. 2021. Randomrooms: Unsupervised pre-training from synthetic shapes and randomized layouts for 3d object detection. In *Proceedings of the IEEE International Conference on Computer Vision*, 3283–3292.
- Rozenberszki, D.; Litany, O.; Dai, A.; and Dai, A. 2024. Unscene3d: Unsupervised 3d instance segmentation for indoor scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 19957–19967.
- Schult, J.; Engelmann, F.; Hermans, A.; Litany, O.; Tang, S.; and Leibe, B. 2023. Mask3d: Mask transformer for 3d semantic instance segmentation. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 8216–8223. IEEE.

- Takmaz, A.; Fedele, E.; Sumner, R.; Pollefeys, M.; Tombari, F.; and Engelmann, F. 2023. OpenMask3D: Open-Vocabulary 3D Instance Segmentation. In *Advances in Neural Information Processing Systems*, volume 36, 68367–68390.
- Vu, T.; Kim, K.; Luu, T. M.; Nguyen, T.; and Yoo, C. D. 2022. Softgroup for 3d instance segmentation on point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2708–2717.
- Xie, C.; Xiang, Y.; Mousavian, A.; and Fox, D. 2021. Unseen object instance segmentation for robotic environments. *IEEE Transactions on Robotics*, 37(5): 1343–1359.
- Xu, M.; Yin, X.; Qiu, L.; Liu, Y.; Tong, X.; and Han, X. 2023. SAMPro3D: Locating SAM Prompts in 3D for Zero-Shot Scene Segmentation. *CoRR*.
- Yan, M.; Zhang, J.; Zhu, Y.; and Wang, H. 2024. Maskclustering: View consensus based mask graph clustering for open-vocabulary 3d instance segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 28274–28284.
- Yang, L.; Zhang, Z.; Han, J.; Zeng, B.; Li, R.; Torr, P.; and Zhang, W. 2024a. Semantic score distillation sampling for compositional text-to-3d generation. *arXiv preprint arXiv:2410.09009*.
- Yang, X.; Gu, X.; Yin, X.; and Gao, X. 2024b. SA3DIP: Segment Any 3D Instance with Potential 3D Priors. In *Advances in Neural Information Processing Systems*.
- Yang, Y.; Lu, J.; Zhao, Z.; Luo, Z.; Yu, J. J.; Sanchez, V.; and Zheng, F. 2024c. Llplace: The 3d indoor scene layout generation and editing via large language model. *arXiv preprint arXiv:2406.03866*.
- Yang, Y.; Sun, F.-Y.; Weihs, L.; VanderBilt, E.; Herrasti, A.; Han, W.; Wu, J.; Haber, N.; Krishna, R.; Liu, L.; et al. 2024d. Holodeck: Language guided generation of 3d embodied ai environments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 16227–16237.
- Yeshwanth, C.; Liu, Y.-C.; Nießner, M.; and Dai, A. 2023. Scannet++: A high-fidelity dataset of 3d indoor scenes. In *Proceedings of the IEEE International Conference on Computer Vision*, 12–22.
- Yilmaz, K.; Schult, J.; Nekrasov, A.; and Leibe, B. 2024. Mask4former: Mask transformer for 4d panoptic segmentation. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 9418–9425. IEEE.
- Yin, Y.; Liu, Y.; Xiao, Y.; Cohen-Or, D.; Huang, J.; and Chen, B. 2024. Sai3d: Segment any instance in 3d scenes. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3292–3302.
- Zhang, B.; Tang, J.; Niessner, M.; and Wonka, P. 2023. 3dshape2vecset: A 3d shape representation for neural fields and generative diffusion models. *ACM Transactions On Graphics (TOG)*, 42(4): 1–16.
- Zhang, Q.; Wang, C.; Siarohin, A.; Zhuang, P.; Xu, Y.; Yang, C.; Lin, D.; Zhou, B.; Tulyakov, S.; and Lee, H.-Y. 2024. Towards text-guided 3d scene composition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6829–6838.
- Zhou, S.; Fan, Z.; Xu, D.; Chang, H.; Chari, P.; Bharadwaj, T.; You, S.; Wang, Z.; and Kadambi, A. 2024. Dreamscene360: Unconstrained text-to-3d scene generation with panoramic gaussian splatting. In *European Conference on Computer Vision*, 324–342.