

SAS: Segment Any 3D Scene with Integrated 2D Priors

Zhuoyuan Li^{1*}, Jiahao Lu^{1*}, Jiacheng Deng¹, Hanzhi Chang¹,
Lifan Wu¹, Yanzhe Liang¹, Tianzhu Zhang^{1,2†}

¹University of Science and Technology of China

²Deep Space Exploration Laboratory

Project Page: <https://peoplelu.github.io/SAS.github.io>

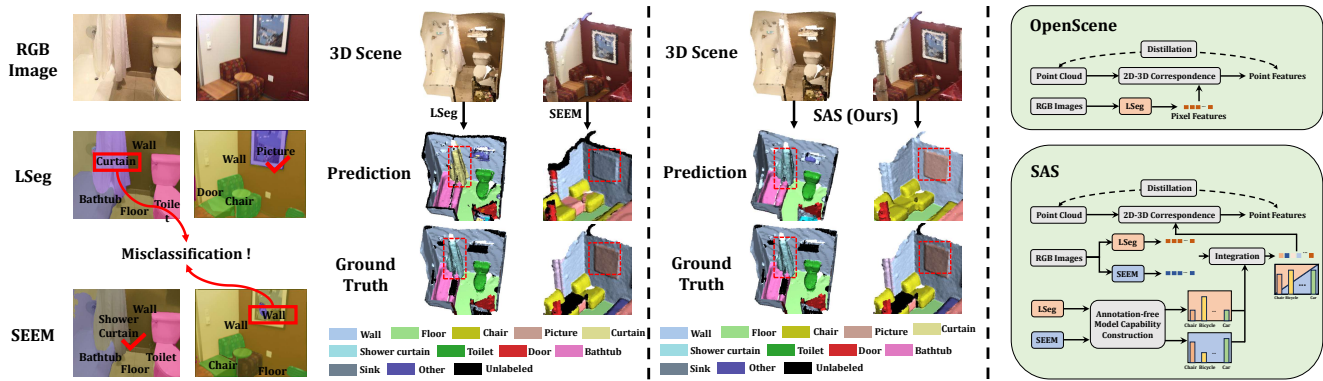


Figure 1. **Left:** The leading 2D open vocabulary models like LSeg [40] and SEEM [100] often misidentify objects, which makes the distilled 3D model perform the same misidentification. **Middle:** Our proposed SAS successfully correct the misidentified object. **Right:** SAS distills open vocabulary knowledge from multiple 2D models with novel designs, e.g., Annotation-free Model Capability Construction.

Abstract

The open vocabulary capability of 3D models is increasingly valued, as traditional methods with models trained with fixed categories fail to recognize unseen objects in complex dynamic 3D scenes. In this paper, we propose a simple yet effective approach, **SAS**, to integrate the open vocabulary capability of multiple 2D models and migrate it to 3D domain. Specifically, we first propose **Model Alignment via Text** to map different 2D models into the same embedding space using text as a bridge. Then, we propose **Annotation-Free Model Capability Construction** to explicitly quantify the 2D model’s capability of recognizing different categories using diffusion models. Following this, point cloud features from different 2D models are fused with the guide of constructed model capabilities. Finally, the integrated 2D open vocabulary capability is transferred to 3D domain through feature distillation. SAS outperforms previous methods by a large margin across multiple datasets, including ScanNet v2, Matterport3D, and nuScenes, while its generalizability is further validated on downstream tasks,

e.g., gaussian segmentation and instance segmentation.

1. Introduction

3D scene understanding is a fundamental task that aims to predict the semantics for every 3D point. Numerous real-world applications, such as autonomous driving [1, 20, 21], virtual reality [56], and robot manipulation [69], depend heavily on 3D scene understanding. Traditional methods in this area perform supervised training on labeled 3D datasets with a limited number of categories, which hinders the model’s ability to identify unseen objects. This prompts researchers to turn their attention to open vocabulary capabilities of 3D scene understanding models.

In 2D open vocabulary understanding, a typical approach is to align image features to language features by contrast learning on a large number of image-text pairs, e.g., CLIP [64]. By analogy, contrast learning on a large number of point-text pairs is supposed to enable 3D open vocabulary understanding. However, point-text pairs are much harder to acquire than image-text pairs due to the high cost and time-consuming nature of point cloud annotation [18, 57]. A compromise approach is to distill the open vocabulary ca-

*Equal Contribution

†Corresponding Author

pability of 2D open vocabulary models onto the 3D models [5, 18, 29, 57, 80, 88, 92, 98]. For example, OpenScene [57] adopts 2D open vocabulary segmenters to extract pixel features for point clouds to enable point-language alignment. These pixel features are distilled to transfer the 2D open vocabulary capability to 3D models.

Although distillation-based methods have become dominant in 3D scene understanding, the distilled model inherited its 3D open vocabulary capability from the adopted 2D models, making the performance of the distilled model vary with the performance of the adopted 2D model. As shown in Fig. 1, the 2D open vocabulary model may incorrectly recognize objects in the scene, which generates ambiguous supervision signals for the 3D model so that the 3D model will make the same mistake. For example, SEEM [100] misidentifies a picture as a wall (left), making the distilled model perform the same misidentification (left). Therefore, an intuitive idea is to integrate the capabilities of different 2D models to correct misidentified objects. However, there are two inherent difficulties in integrating the capability of different 2D models. First, different 2D models have unaligned image-text feature spaces. The unaligned features cannot be directly integrated. Second, testing the model on a batch of images reflects the capability of this model. However, since 3D open vocabulary tasks are inherently designed to recognize unseen categories in a zero-shot fashion, obtaining test images and their annotations is challenging. As a result, it is impractical to directly evaluate the model’s performance.

To address the above problems, we propose **SAS**, to learn better 3D representations from the integration of multiple 2D models. Our key idea is to explicitly model 2D models’ capability of identifying different categories in the scene, and use this to guide the integration of different 2D models, which produces better 3D representations than a single 2D model can do. First, we propose **Model Alignment via Text** to align features of different 2D models using text as a bridge. Specifically, a 2D open vocabulary model, e.g., SEEM model [100], predicts a label for each pixel (or mask), based on which we further generate a caption containing more complex information such as color, shape, etc., for each pixel (or mask), which enriches the semantic information and intra-class diversity. Now, different 2D open vocabulary models are aligned at the text-level. The caption is then inputted into a shared text encoder to get the aligned features. Second, we propose **Annotation-free Model Capability Construction** to quantitatively measure the capability of different 2D open vocabulary models. Specifically, it leverages the text-to-image diffusion model [66] to synthesize images of common types. We utilize the 2D open vocabulary model’s performance on these synthesized images to quantify the model’s capabilities, which later can guide the integration of different 2D models. Finally, SAS

transfers the integrated open vocabulary capability of 2D models to 3D domain via feature distillation.

We evaluate SAS on multiple datasets, including ScanNet v2 [10], Matterport3D [3] and nuScenes [1]. Overall, we make the following contributions:

- We propose SAS, the first approach to learn better 3D representations from multiple 2D models.
- We propose Model Alignment via Text and Annotation-free Model-capability Construction to address the aforementioned difficulties in integrating different 2D models. The former aligns the features of models with different embedding spaces, while the latter explicitly models the model capability via synthetic images, which is then used to guide the integration of different model features.
- We validate SAS on different 3D scene understanding tasks, including semantic segmentation, instance segmentation, and gaussian segmentation. The results show that SAS significantly outperforms previous methods while preserving strong generalizability.

2. Related Work

Closed-set 3D Scene Understanding. Recent work on closed-set 3D scene understanding has seen significant advancement, spanning tasks such as 3D semantic segmentation [36, 43, 44, 59, 61, 85, 94], object detection [13, 34, 37, 55, 70, 97], instance segmentation [27, 34, 35, 48–51, 72, 77], and shape correspondence [11, 12, 14, 15, 17, 24, 38, 91]. Some studies [16, 41, 42, 74, 79] are exploring the use of multimodal information, particularly 2D images, to enhance closed-set 3D scene understanding. However, these methods primarily focus on supervised 3D tasks rather than open-vocabulary problems. Our approach aims to better leverage and integrate existing 2D open-vocabulary foundation models to achieve more accurate 3D open-vocabulary scene understanding tasks.

Open-Vocabulary 2D Scene Understanding. The rapid development of vision-language models, such as CLIP [65], has driven zero-shot 2D scene understanding tasks. However, image-level recognition falls short for real-world applications. As a result, many works [6, 22, 39, 45, 87, 101] aim to correlate pixel embeddings with text embeddings to enable dense prediction tasks like segmentation. LSeg [39] and OpenSeg [22] are pioneering works in 2D open-vocabulary semantic segmentation via pixel-level text alignment. Following OpenScene [58], we use 2D open-vocabulary segmentation models as teacher models for distilling 3D models. However, unlike previous works [28, 58, 81, 99], we find that different 2D open-vocabulary segmentation models exhibit varying recognition capabilities across categories. To better utilize these models, our method quantitatively evaluates their performance on annotation-free categories and integrates their strengths for improved scene understanding.

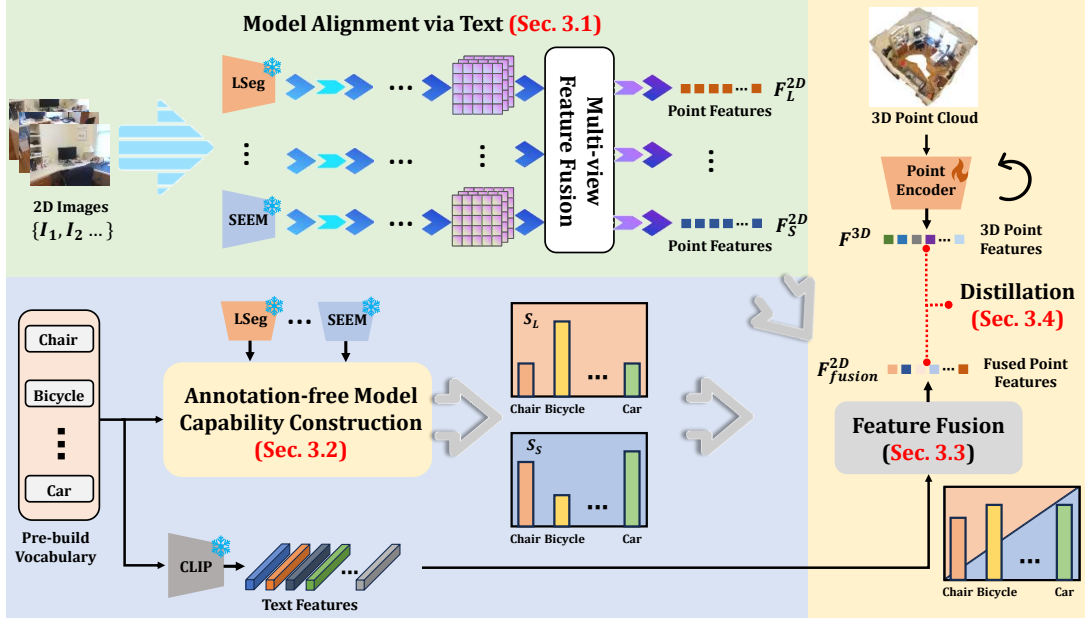


Figure 2. **Overview of our proposed SAS.** SAS first align features of different models in a unified embedding space (Sec. 3.1). Then SAS constructs models’ capability to recognize various objects (Sec. 3.2). With the constructed capability as guide, features from different 2D models are integrated (Sec. 3.3). Finally, a 3D network is distilled to enable 3D open vocabulary understanding (Sec. 3.4).

Open-Vocabulary 3D Scene Understanding. Recent works on open-vocabulary 3D scene understanding can be classified into two categories based on representation patterns. The first category [2, 31, 63, 90], represented by works like LERF [31] and LangSplat [63], distills 2D features into NeRF [52] or 3DGS [30]. However, generating NeRF or 3DGS in a feed-forward manner remains challenging. The second category [28, 58, 81, 99], including methods like OpenScene [58], bridges the gap between point clouds and images using camera parameters and depth. OV3D [28] enriches textual descriptions with contextual information from foundation models, enabling richer feature distillation during alignment. Diff2Scene [99] utilizes text-image generative models [67] for open-vocabulary 3D scene understanding, while GGSD [81] leverages the Mean-teacher paradigm [75] for improved distillation. According to these methods, we note that the current capability of open-vocabulary 3D scene understanding heavily depends on the scene understanding ability of 2D foundation models [22, 39, 45, 87, 101]. Therefore, exploring the potential of these 2D models is crucial for advancing 3D scene understanding. Based on this, our method tries to acquire a quantitative evaluation of their performance across different categories and integrate their strengths for better scene understanding.

3. Method

An overview of SAS is illustrated in Fig. 2. Our initial step involves aligning features of different 2D models to

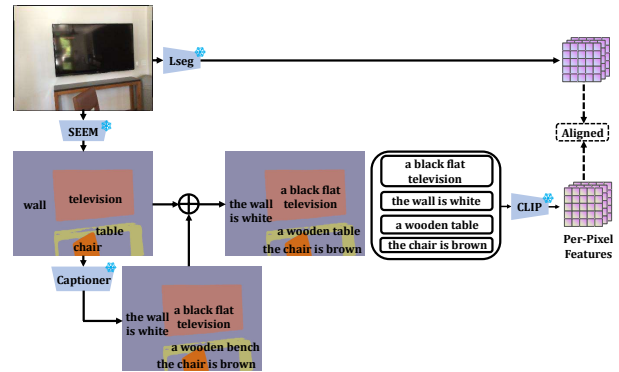


Figure 3. **Overview of Model Alignment via Text.** Features from different models are first aligned on text level, which are then encoded by a shared text encoder to produce aligned features.

the same feature space. The 2D features are then aggregated onto 3D points through pixel-point correspondence (Sec. 3.1). We then utilize diffusion models to quantify 2D models’ capabilities of identifying different categories in the scene (Sec. 3.2). Subsequently, we leverage the constructed model capability to guide the fusion of 2D point features to obtain fused 2D point features (Sec. 3.3). Finally, we distill a 3D MinkowskiNet from the fused 2D point features (Sec. 3.4). Details are in Sec. 3.5.

3.1. Model Alignment via Text

In this section, our objective is to align two 2D models into the same feature space. Here, we take LSeg [40] and SEEM

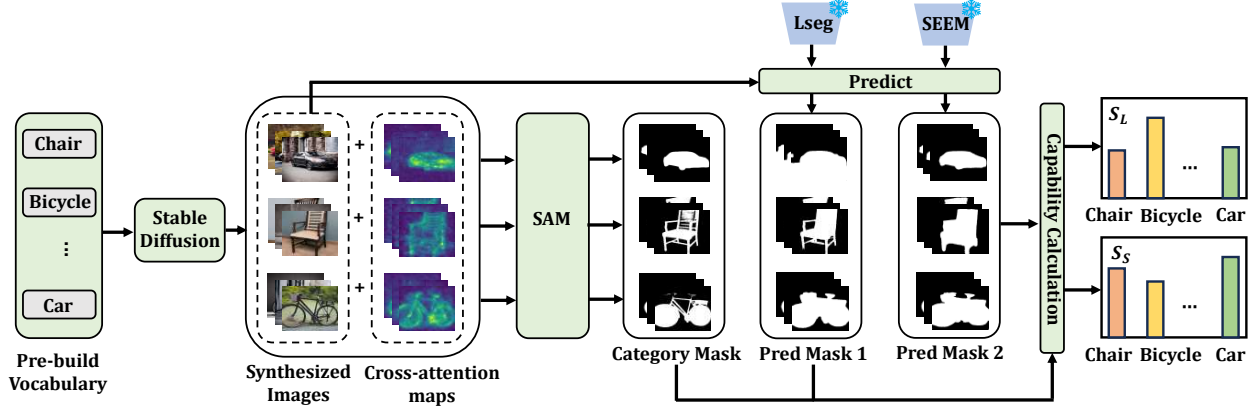


Figure 4. **Overview of Annotation-free Model Capability Construction.** Stable Diffusion model [66] is utilized to generate synthesized images with masks computed by SAM [33]. By assessing model’s performance on synthesized images, we construct model capabilities.

[100] as an example without loss of generalizability. 2D open vocabulary models usually consist of a text encoder and an image encoder that project text and images into a unified embedding space. LSeg [40] outputs dense per-pixel feature that is aligned with CLIP [64]. SEEM [100] is a maskformer [7] style model that predicts masks for an image and computes labels for each mask. However, SEEM [100]’s embedding space differs from LSeg’s due to variations in their text encoders. To address this, we align their embedding spaces.

Fig. 3 shows an example of the aligning process. Given an image $I_i \in \mathbb{R}^{H \times W}$, we first input it into LSeg to get per-pixel features $f_L^{2D} \in \mathbb{R}^{H \times W \times C}$. Similarly, we input I_i into SEEM to obtain the mask-label pairs. Specifically, we denote the j -th mask and its corresponding label as $m_{i,j}$ and $l_{i,j}$. Therefore, the output of SEEM can be formulated as $S_i = \{m_{i,j}, l_{i,j}\}_{j=1,2,\dots,N_i}$, where N_i is the total number of masks predicted in I_i . However, $l_{i,j}$ are simple nouns that lack intra-class diversity, e.g., color and shape. Therefore, we adopt a pre-trained captioner, TAP [54], to generate captions $\{c_{i,j}\}_{j=1,2,\dots,N_i}$ for each mask $m_{i,j}$, providing additional semantic information. We then replace the noun labels in $\{c_{i,j}\}$ with the none labels predicted by SEEM $l_{i,j}$ to obtain $\{\hat{c}_{i,j}\}_{j=1,2,\dots,N_i}$, e.g., $c_{i,j}$ ="a woodeow bench", $l_{i,j}$ ="table" and $\hat{c}_{i,j}$ ="a wooden table". Now we have obtained semantically-rich captions $\hat{c}_{i,j}$ for each mask, which can then be encoded by CLIP and mapped back to the image to form per-pixel features $f_S^{2D} \in \mathbb{R}^{H \times W \times C}$, which are aligned with f_L^{2D} in CLIP’s feature space.

We then follow OpenScene [57] to calculate the pixel-point correspondence, which is then utilized to map the pixel features onto point features with a multi-view fusion strategy [57]. The point features from LSeg and SEEM are denoted as $F_L^{2D} \in \mathbb{R}^{N \times C}$ and $F_S^{2D} \in \mathbb{R}^{N \times C}$, where N is the number of points in the point cloud.

3.2. Annotation-free Model Capability Construction

In this section, we aim to quantify model capabilities. However, obtaining test images and their annotations is difficult, making it impractical to directly assess the model’s performance. To overcome this, we utilize the Stable Diffusion (SD) [66] model and SAM [33] to establish a reference for evaluating model capabilities.

As shown in Fig. 4, given a pre-built vocabulary consisting of common classes in the scene, denoted as $C = \{C_1, C_2, \dots, C_K\}$ where K is the number of classes, we leverage the SD [66] model to generate m images $\hat{I}_q = \{\hat{I}_1^q, \hat{I}_2^q, \dots, \hat{I}_m^q\}$ for each class, where q indicates the q -th class and \hat{I}_j^q is the j -th generated image of the q -th class. In addition to synthesizing a picture of a certain class, we need to localize the corresponding object in it with the help of cross-attention maps inside SD [66] models. Considering a conditional diffusion model (e.g., SD [66] model), the input to a UNet layer are noisy image features $F_v \in \mathbb{R}^{H \times W \times C}$ and text prompt features $F_p \in \mathbb{R}^{N \times D}$. The noisy image features are squeezed and then projected into the query Q as $Q = F_v W^Q$, while text prompt features are projected into the key K and the value V similarly as $K = F_p W^K$ and $V = F_p W^V$, where W^Q, W^K, W^V are the matrices of corresponding linear layers. Following this, the cross-attention map M can be calculated as:

$$M = \text{Softmax}\left(\frac{QK^T}{\sqrt{D}}\right). \quad (1)$$

We denote $M_x^{y,z}$ as the cross-attention map of the x -th token in the y -th UNet layer at diffusion step z . To obtain more robust cross-attention maps, we follow [83] to aggregate all attention maps at all Z time steps in all Y UNet layers as:

$$\bar{M}_x = \frac{1}{Y \cdot Z} \sum_{y \in Y, z \in Z} \frac{M_x^{y,z}}{\max(M_x^{y,z})}. \quad (2)$$

We can now obtain a coarse mask for the target object in the image by binarizing the cross-attention map \bar{M}_x via thresholding. Further, we leverage SAM [33] to generate precise mask predictions. Specifically, we sample several points within the coarse mask as the point prompt for the SAM model [33] to generate accurate masks $\mathbf{M}_{i,j}^{Pseudo}$, where i and j denote the i -th image of the j -th class. Similarly, LSeg and SEEM can be applied to the synthesized images to obtain masks $\mathbf{M}_{i,j}^{LSeg}$ and $\mathbf{M}_{i,j}^{SEEM}$ respectively.

Following this, the metric mean intersection over union (mIoU) is adopted to measure model’s capability. Specifically, the capability of LSeg [40] model for a certain class C_j can be constructed as :

$$S_j^{LSeg} = \frac{1}{m} \sum_{i=1, \dots, m} \mathbf{mIoU}(\mathbf{M}_{i,j}^{Pseudo}, \mathbf{M}_{i,j}^{LSeg}), \quad (3)$$

where S_j^{LSeg} is the model capability of LSeg for the j -th class. $\mathbf{mIoU}(a, b)$ computes the mIoU between two binary masks. Similarly, the capability of SEEM model for a certain class C_j can be constructed as:

$$S_j^{SEEM} = \frac{1}{m} \sum_{i=1, \dots, m} \mathbf{mIoU}(\mathbf{M}_{i,j}^{Pseudo}, \mathbf{M}_{i,j}^{SEEM}). \quad (4)$$

Following Eq. 3 and Eq. 4, model capabilities for LSeg [40] and SEEM [100] can be established as S_L and S_S :

$$\begin{aligned} S_L &= [S_1^{LSeg}, \dots, S_K^{LSeg}], \\ S_S &= [S_1^{SEEM}, \dots, S_K^{SEEM}]. \end{aligned} \quad (5)$$

3.3. Feature Fusion

Now we have obtained aligned point features F_L^{2D} and F_S^{2D} and the corresponding constructed capability S_L and S_S . In this section, our objective is to fuse the point features with the guide of model capabilities. We first encode the pre-built vocabulary $C = \{C_1, \dots, C_K\}$ using CLIP [64] to obtain text features $F_{text} = \{f_1, \dots, f_K\}$. Next, we calculate the predicted category of each point for F_L^{2D} and F_S^{2D} :

$$\begin{aligned} \mathcal{P}_{LSeg} &= \underset{K}{\operatorname{argmax}}(F_L^{2D} \cdot F_{text}^T), \\ \mathcal{P}_{SEEM} &= \underset{K}{\operatorname{argmax}}(F_S^{2D} \cdot F_{text}^T). \end{aligned} \quad (6)$$

We assume the right prediction comes from either \mathcal{P}_{LSeg} or \mathcal{P}_{SEEM} . Hence, we adopt the sum of the model capabilities of identifying \mathcal{P}_{LSeg} and \mathcal{P}_{SEEM} to measure the probability of the model making correct predictions. Specifically, the probability of LSeg [40] and SEEM [100] making the right prediction can be formulated as:

$$\begin{aligned} \mathcal{P}_{LSeg} &= \frac{S_L[\mathcal{P}_{LSeg}] + S_L[\mathcal{P}_{SEEM}]}{2}, \\ \mathcal{P}_{SEEM} &= \frac{S_S[\mathcal{P}_{LSeg}] + S_S[\mathcal{P}_{SEEM}]}{2}. \end{aligned} \quad (7)$$

Following this, the probability \mathcal{P}_{LSeg} and \mathcal{P}_{SEEM} can be the weights guiding the fusion of F_L^{2D} and F_S^{2D} :

$$F_{fusion}^{2D} = \frac{\exp(\mathcal{P}_{LSeg}/\tau)}{\exp(\mathcal{P}_{LSeg}/\tau) + \exp(\mathcal{P}_{SEEM}/\tau)} F_L^{2D} + \frac{\exp(\mathcal{P}_{SEEM}/\tau)}{\exp(\mathcal{P}_{LSeg}/\tau) + \exp(\mathcal{P}_{SEEM}/\tau)} F_S^{2D}. \quad (8)$$

where τ is the temperature coefficient.

3.4. Distillation

Superpoint Distillation As 2D models could make potentially inconsistent predictions as shown in Fig. 1, we further introduce superpoint distillation to alleviate this problem, similar to GGSD [80].

Specifically, for a given point cloud $\mathbf{P} \in \mathbb{R}^{N \times 3}$ and a point cloud encoder ϕ , the encoded point features F^{3D} is:

$$F^{3D} = \phi(\mathbf{P}), \mathbb{R}^{N \times 3} \mapsto \mathbb{R}^{N \times C}, \quad (9)$$

where N is the number of points and C is the feature dimension. We extract L non-overlapping superpoints from \mathbf{P} as $\{p_1, p_2, \dots, p_L\}$. We assume superpoints are semantically coherent that each superpoint should have the same category. The mean feature for each superpoint for F_{fusion}^{2D} and F^{3D} can be computed as: $f_{SP}^{2D} = \{\hat{f}_1^{2D}, \hat{f}_2^{2D}, \dots, \hat{f}_L^{2D}\}$ and $f_{SP}^{3D} = \{\hat{f}_1^{3D}, \hat{f}_2^{3D}, \dots, \hat{f}_L^{3D}\}$ respectively. Then we map the superpoint features f_{SP}^{2D} onto all points to obtain per-point features F_{SP}^{2D} . The distillation is performed on both point level and superpoint level as:

$$\begin{aligned} \mathcal{L}_p &= 1 - \cos(F_{SP}^{2D}, F^{3D}), \\ \mathcal{L}_{sp} &= 1 - \cos(f_{SP}^{2D}, f_{SP}^{3D}), \\ \mathcal{L} &= \mathcal{L}_p + \mathcal{L}_{sp}, \end{aligned} \quad (10)$$

where \cos computes the cosine similarity, \mathcal{L}_p and \mathcal{L}_{sp} indicate the distillation on point and superpoint level, \mathcal{L} is the total loss.

Temporal Ensembling Self-Distillation The improvement in the distilled model inspires us to further exploit the model’s potential via self-distillation. GGSD [80] proposes to use the student model’s output to supervise the teacher model while updating the student model through EMA of the teacher model. However, we find training the teacher model with a trainable and variable student model leads to an unstable training process, even a model collapse. Therefore, we propose temporal ensembling self-distillation.

Specifically, for a 3D network ϕ with its output F^{3D} , we construct \hat{F}^{3D} to store ϕ ’s output in previous epochs to establish a more smooth optimization process:

$$\hat{F}^{3D} = \alpha \hat{F}^{3D} + (1 - \alpha) F^{3D}, \quad (11)$$

where α is a constant. By applying average pooling to superpoints of \hat{F}^{3D} , the pooled features \hat{f}_{SP}^{3D} is then obtained. The pseudo label for \hat{F}^{3D} and \hat{f}_{SP}^{3D} can then be computed:

$$\begin{aligned}\hat{\mathcal{P}}^{3D} &= \mathbf{argmax}(\hat{F}^{3D} \cdot \hat{F}_{text}^T), \\ \hat{\mathcal{P}}_{SP}^{3D} &= \mathbf{argmax}(\hat{f}_{SP}^{3D} \cdot \hat{F}_{text}^T),\end{aligned}\quad (12)$$

where \hat{F}_{text} is the class embedding of the dataset. $\hat{\mathcal{P}}^{3D}$ and $\hat{\mathcal{P}}_{SP}^{3D}$ then supervise ϕ by optimizing the losses:

$$\begin{aligned}\mathcal{L}_p^{ST} &= \mathbf{CE}(\mathbf{argmax}(F^{3D} \cdot \hat{F}_{text}^T), \hat{\mathcal{P}}^{3D}), \\ \mathcal{L}_{sp}^{ST} &= \mathbf{CE}(\mathbf{argmax}(f_{SP}^{3D} \cdot \hat{F}_{text}^T), \hat{\mathcal{P}}_{SP}^{3D}), \\ \mathcal{L} &= \mathcal{L}_p^{ST} + \mathcal{L}_{sp}^{ST},\end{aligned}\quad (13)$$

where \mathbf{CE} denotes the CrossEntropy loss, \mathcal{L}_p^{ST} and \mathcal{L}_{sp}^{ST} indicate the self distillation on point level and superpoint level respectively and \mathcal{L} is the total loss.

3.5. Training and Inference

Training The training process contains 100 epochs. In the first 70 epochs, we employ superpoint distillation. In the last 30 epochs, we employ both superpoint distillation and temporal ensembling self-distillation. Additionally, we designed two pre-build vocabularies for indoor scenes and outdoor scenes respectively. The detail is included in supplementary materials. While outdoor point clouds (e.g., nuScenes [1]) are typically dominated by “roads”, we do not compute superpoints for nuScenes. Instead, we treat every single point in nuScenes as a superpoint for simplicity.

Inference We directly use the output features from the 3D model to calculate the similarity with CLIP [64] features of different categories without any other post-processing, e.g., superpoint or 2D-3D ensemble [57].

4. Experiments

We conduct extensive experiments to demonstrate the effectiveness of SAS on 3D scene understanding tasks in a **zero-shot** fashion. We first introduce the experiment setup in Sec. 4.1. In Sec. 4.2, We evaluate SAS on zero-shot open vocabulary semantic segmentation tasks. We then perform comprehensive ablation studies to validate our design choices in Sec. 4.3. Further, we extend SAS to instance segmentation and gaussian segmentation in Sec. 4.4.

4.1. Experiment Setup

Dataset. We evaluate SAS on three commonly used datasets: ScanNet v2 [10], Matterport3D [3] and nuScenes [1]. ScanNet v2 [10] is an indoor dataset, containing 1513 room scans from 2.5 million RGB-D images, the average point number of which is 148k Matterport3D [3] is another complex indoor dataset consisting of 10800 scenes of the

Table 1. **Evaluations on zero-shot 3D semantic segmentation.** We compare SAS with both zero-shot and fully-supervised approaches on nuScenes, ScanNet and MatterPort3D using mIOU as metrics. Best results under each setting are shown bold.

| | nuScenes | ScanNet v2 | Matterport3D |
|------------------------------------|-------------|-------------|--------------|
| <i>Fully-supervised approaches</i> | | | |
| PointNet++ [60] | - | 53.5 | - |
| PointConv [82] | - | 61.0 | - |
| KPConv [76] | - | 69.2 | - |
| Mix3D [53] | - | 73.6 | - |
| LidarMultiNet [89] | 82.0 | - | - |
| PTv3 [84] | 80.4 | 77.5 | - |
| MinkowskiNet [9] | 78.0 | 69.0 | 54.2 |
| <i>Zero-shot approaches</i> | | | |
| OpenScene [57] | 42.1 | 54.2 | 43.4 |
| GGSD [80] | 46.1 | 56.5 | 40.1 |
| Diff2Scene [98] | - | 48.6 | 45.5 |
| Seal [46] | 45.0 | - | - |
| OV3D [29] | 44.6 | 57.3 | 45.8 |
| Ours | 47.5 | 61.9 | 48.6 |

building environment, equipped with 194K RGB-D images. nuScenes [1] is an outdoor dataset containing 1,000 driving sequences in total with 34k LiDAR point clouds. ScanNet v2 [10], Matterport3D [3], and nuScenes [1] follow the official data splits to generate corresponding training set, validation set and test set, with evaluations on 20, 21, and 23 categories respectively. ScanNet v2 [10] validation set, Matterport3D [3] test set, and, nuScenes [1] validation set are adopted for evaluation.

Implementation Details. We follow the settings specified in OpenScene [57], utilizing MinkowskiNet18A [9] as the 3D backbone network, and adopting a voxel size of 2cm and 5cm for indoor datasets and outdoor datasets respectively. Similarly with OpenScene [57], we employ a prompt engineering that modifies each class name “XX” to “a XX in a scene”. AdamW [47] is adopted for optimization. We use a batch size of 12 for ScanNet v2 [10] and Matterport3D [3] with 4 RTX 3090 GPUs. For nuScenes [1], we use a batch size of 8 with 4 A6000 GPUs. More details can be found in the supplementary file. Besides, based on the ablation table of OpenScene [57], we employ LSeg + SEEM for indoor datasets and OpenSeg + SEEM for outdoor dataset. Notably, OpenScene [57] adopts a 2D-3D ensemble strategy to further enhance the performance, while we directly report the performance of the pure 3D model.

4.2. Main Results

Evaluation on zero-shot 3D semantic segmentation.

Our proposed SAS is compared with both fully-supervised and zero-shot methods on ScanNet v2 [10], Matterport3D [3] and nuScenes [1]. SAS exhibits superior performance in zero-shot 3D semantic segmentation, as shown in Tab. 1. Specifically, SAS outperforms all previous methods, including the recent state-of-the-art OV3D [29]. Compared with

Table 2. **Evaluation in long-tail scenarios.** We evaluate SAO in MatterPort40, MatterPort80 and MatterPort160 using mIOU as the metric. Best results are shown bold.

| | $K = 21$ | $K = 40$ | $K = 80$ | $K = 160$ |
|----------------|-------------|-------------|-------------|------------|
| OpenScene [57] | 43.4 | 22.9 | 11.3 | 5.8 |
| Ours | 48.6 | 28.5 | 16.0 | 8.1 |

previous SOTAs, SAS shows a +1.4%, +4.6% and +2.8% improvement in mIOU on nuScenes [1], ScanNet v2 [10] and Matterport3D [3] respectively.

Evaluation in long-tail scenarios. Different from the standard benchmarks with dozens of categories, we further evaluate SAS under more complex long-tail scenarios with more categories (e.g., 160) to validate its open vocabulary capability. The Matterport3D [3] benchmark additionally provides K most common categories from the NYU label set for $K = 40, 80, 160$, which we evaluate on as shown in Tab. 2. Complying the zero-shot setting, we adopt the same distilled model to evaluate for all $K = 40, 80, 160$. The result suggests that our proposed SAS consistently outperforms OpenScene [57], demonstrating its strong open vocabulary capability in long-tail scenarios.

Visual Comparison Visual comparisons with OpenScene [57] on semantic segmentation in ScanNet v2 [10] are shown in Fig. 5, which shows that our approach effectively corrects some wrong predictions made by OpenScene [57]. For example, OpenScene [57] misidentifies a chair as a sofa, while our approach can easily tell the difference between a chair and a sofa. Additional visualization about querying different objects in a scene is provided in Fig. A in supplementary material. Visualization about performance on other datasets is provided in Fig. B, Fig. C and Fig. D.

4.3. Ablation Studies and Analysis

Different 2D Features. Top three lines of Tab. 3 show the performance obtained by directly projecting different 2D features onto 3D points. Based on the findings from OpenScene [57], we employ Lseg + SEEM for indoor scenes and OpenSeg + SEEM for outdoor scenes. Tab. 3 suggests that the performance of these models on different datasets varies randomly, e.g., SEEM [100] performs better than LSeg [40] on Matterport3D [3] (+1.6%) while SEEM [100] performs worse than LSeg [40] (-3.9%) on ScanNet v2 [10]. However, with our proposed fusion strategy, the performance after fusion is always better than that of two separate models.

Fusion strategy. Tab. 3 also performs ablations of different fusion strategies including adding, linear fusion and our proposed fusion strategy in Sec. 3.3. Adding indicates adding two features directly. Linear fusion computes the similarity between a point feature and text features and selects the largest similarity as the point feature’s weight,

Table 3. **Ablation study of SAS.** Zero-shot semantic segmentation performance of certain component of SAO being ablated. SP, TESD and Dis. indicates superpoint, temporal ensembling self distillation and distillation respectively.

| | ScanNet v2 | Matterport | nuScenes |
|--------------------------|-------------|-------------|-------------|
| 2D Fusion LSeg [40] | 51.2 | 38.6 | - |
| 2D Fusion OpenSeg [23] | - | - | 32.9 |
| 2D Fusion SEEM [100] | 47.3 | 40.2 | 37.8 |
| 2D Fusion Add | 48.6 | 39.1 | 34.1 |
| 2D Fusion Linear | 49.9 | 39.4 | 34.8 |
| 2D Fusion (Ours) | 55.5 | 43.6 | 40.0 |
| 3D Pixel-Point Dis. | 56.7 | 45.1 | 45.4 |
| 3D SP Dis. | 59.2 | 46.3 | 45.4 |
| 3D SP Dis. + TESD (Ours) | 61.9 | 48.6 | 47.5 |

based on which two point features are fused. As can be seen, among all datasets, our proposed strategy performs the best. This indicates simple summation or linear fusion fails to inject model capabilities as guides into the feature fusion process, while our approach solves this.

Distillation Pattern. Bottom half of Tab. 3 shows the performance of the model distilled from 2D fused features in different ways, including pixel-point distillation, superpoint distillation, and temporal ensembling self distillation. OpenScene [57] adopts pixel-point distillation to directly align point features and text features. Note that the superpoint distillation performs the same with pixel-point distillation on nuScenes, because we treat every point in nuScenes as a superpoint in Sec. 3.5. The pixel-point distillation has already exhibited better performance than 2D features without additional designs, resulting in a +1.2%, +1.5%, and +5.4% increase in mIOU on ScanNet v2 [10], Matterport3D [3] and nuScenes [1] respectively. Superpoint distillation brings extra inherent structural information of point clouds into distillation and achieves more accurate predictions. Temporal ensembling self distillation further exploits the potential of distilled models via self distillation, boosting the performance gap with 2D features to +6.4%, +5.0%, and +7.5% in mIOU.

Table 4. **Fusion of LSeg, SEEM and ODISE.** Dis. indicates distillation.

| | ScanNet v2 | Matterport3D |
|-----------------------------|-------------|--------------|
| 2D Fusion LSeg [40] | 51.2 | 43.5 |
| 2D Fusion SEEM [100] | 47.3 | 46.1 |
| 2D Fusion ODISE [86] | 46.9 | 46.7 |
| 3D SEEM + LSeg Dis. | 61.9 | 48.6 |
| 3D SEEM + LSeg + ODISE Dis. | 62.5 | 49.8 |

More 2D Models. We further evaluate the effectiveness of SAS by incorporating additional 2D open-vocabulary models. As detailed in Table 4, we integrate another 2D model, ODISE [86], alongside LSeg [40] and SEEM

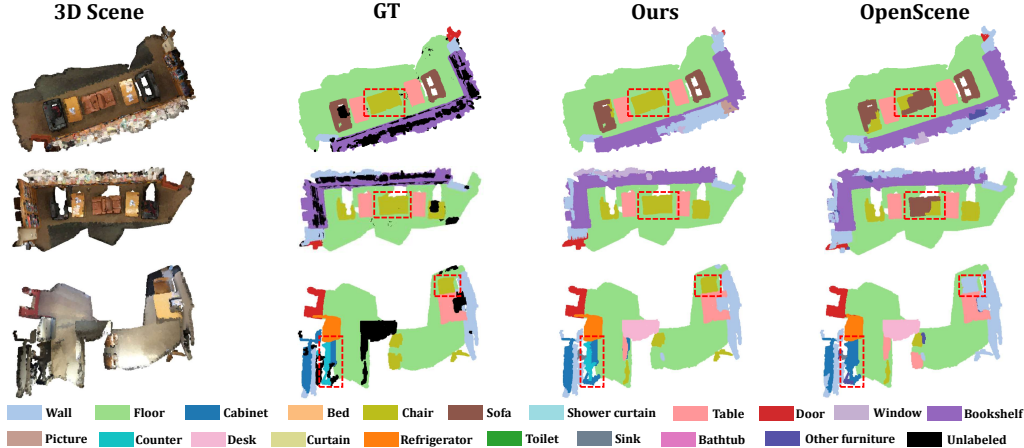


Figure 5. **Visualization results.** Semantic segmentation results of SAS on ScanNet v2.

Table 5. **Gaussian segmentation.** 2D semantic segmentation results of 3D gaussian splatting on 12 scenes from the ScanNet v2 validation set. Comparison focuses on NeRF/3DGS-based methods using mIoU and mAcc.

| Method | Backbone | mIoU | mAcc |
|------------------------------------|----------------------|-------------|-------------|
| <i>Fully-supervised approaches</i> | | | |
| Mask2Former [7] | ViT [19] | 46.7 | - |
| DM-NeRF [78] | NeRF [52] | 49.5 | - |
| SemanticNeRF [95] | NeRF | 59.2 | - |
| Panoptic Lifting [71] | NeRF | 65.2 | - |
| <i>Zero-shot approaches</i> | | | |
| OpenSeg [23] | EfficientNet [73] | 53.4 | 75.1 |
| LSeg [40] | ViT | 56.1 | 74.5 |
| LERF [32] | NeRF+CLIP [64] | 31.2 | 61.7 |
| PVLFF [4] | NeRF+LSeg | 52.9 | 67.0 |
| LangSplat [62] | 3DGS [30]+CLIP | 24.7 | 42.0 |
| Feature3DGS [96] | 3DGS+LSeg | 59.2 | 75.1 |
| Semantic Gaussians [25] | 3DGS+LSeg | 60.7 | 76.3 |
| Ours | 3DGS+LSeg+SEEM [101] | 63.9 | 79.9 |

[100], conducting experiments on the ScanNet v2 and Matterport3D. The results demonstrate that the inclusion of ODISE leads to a further performance enhancement, while also highlighting the promising scalability of our approach.

4.4. Further Exploration

We further apply SAS to other tasks, including 3D Gaussian segmentation and 3D instance segmentation.

3D Gaussian Segmentation 3D Gaussian segmentation aims to assign a semantic label to each Gaussian point, facilitating the rendering of arbitrary 2D views for a thorough understanding of 3D scenes. In this task, we use Semantic Gaussians as our baseline, harnessing the 2D foundation model LSeg to achieve zero-shot 3D Gaussian segmentation via knowledge distillation. As shown in Tab. 5, rather than relying exclusively on LSeg, we integrate the strengths of both LSeg and SEEM to improve scene understanding, achieving a performance increase of 3.2 in mIoU and 3.6 in mAcc. Notably, due to the distinct 3D representations

Table 6. **Instance segmentation.** 3D open-vocabulary instance segmentation results on ScanNet v2 validation set.

| Method | Semantic | mAP | AP@50 | AP@25 |
|-----------------------------|-----------------|-------------|-------------|-------------|
| PointClip [93] | Clip [65] | - | 4.5 | 14.4 |
| OpenIns3D [26] | Yoloworld [8] | 19.9 | 28.9 | 38.9 |
| OpenScene [57] (2D/3D Ens.) | LSeg [40] | 19.7 | 25.9 | 30.4 |
| OpenScene (3D Distill) | LSeg | 19.8 | 25.7 | 30.4 |
| Ours | LSeg+SEEM [101] | 25.1 | 31.7 | 36.7 |

(Gaussian vs point cloud), we distill a new 3D encoder tailored specifically for 3D Gaussian. During this process, we do not utilize the distillation strategy outlined in Sec. 3.4. This task further highlights the effectiveness and generalizability of our approach.

3D Instance Segmentation 3D instance segmentation seeks to detect and delineate multiple instances of specific object categories within a 3D space. In this task, building on previous approaches, we employ Mask3D [68] to generate mask proposals and then leverage the distilled point encoder, as presented in Tab. 1, to derive the corresponding semantics. As demonstrated in Tab. 6, our method achieves significant improvements over the baseline OpenScene, with an mAP of 5.3, AP@50 of 6.0, and AP@25 of 6.3. Compared to the current state-of-the-art OpenIns3D, our approach boosts mAP by 5.2 and AP@50 by 2.8.

5. Conclusion

In this paper, we present SAS, a simple yet effective approach to transfer the open vocabulary capabilities of multiple pre-trained 2D models to 3D domain. By aligning the embedding space of different 2D models and utilizing diffusion models to construct the model capability, which then guides the fusion of different features, SAS achieves superior performance on both indoor and outdoor datasets over previous methods. Additionally, SAS exhibits strong generalization by extending to other tasks, including gaussian segmentation and instance segmentation.

Acknowledgments

This work was supported by Open Fund of National Key Laboratory of Deep Space Exploration (Grant NKLDSE2023A009).

References

- [1] Holger Caesar, Varun Bankiti, Alex H Lang, Sourabh Vora, Venice Erin Liong, Qiang Xu, Anush Krishnan, Yu Pan, Giancarlo Baldan, and Oscar Beijbom. nuscenes: A multimodal dataset for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11621–11631, 2020. 1, 2, 6, 7
- [2] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Wei Shen, Lingxi Xie, Dongsheng Jiang, Xiaopeng Zhang, Qi Tian, et al. Segment anything in 3d with nerfs. *Advances in Neural Information Processing Systems*, 36:25971–25990, 2023. 3
- [3] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. *arXiv preprint arXiv:1709.06158*, 2017. 2, 6, 7
- [4] Haoran Chen, Kenneth Blomqvist, Francesco Milano, and Roland Siegwart. Panoptic vision-language feature fields. *IEEE Robotics and Automation Letters*, 2024. 8
- [5] Runnan Chen, Youquan Liu, Lingdong Kong, Xinge Zhu, Yuexin Ma, Yikang Li, Yuenan Hou, Yu Qiao, and Wenping Wang. Clip2scene: Towards label-efficient 3d scene understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7020–7030, 2023. 2
- [6] Yujia Chen, Wangkai Li, Zhaoyang Li, Rui Sun, Tianzhu Zhang, Zhiwei Xiong, and Feng Wu. Sam-glomeruli: Enhanced segment anything model for precise glomeruli segmentation. In *International Workshop on Medical Optical Imaging and Virtual Microscopy Image Analysis*, pages 182–191. Springer, 2024. 2
- [7] Bowen Cheng, Ishan Misra, Alexander G Schwing, Alexander Kirillov, and Rohit Girdhar. Masked-attention mask transformer for universal image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 1290–1299, 2022. 4, 8
- [8] Tianheng Cheng, Lin Song, Yixiao Ge, Wenyu Liu, Xinggang Wang, and Ying Shan. Yolo-world: Real-time open-vocabulary object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16901–16911, 2024. 8
- [9] Christopher Choy, JunYoung Gwak, and Silvio Savarese. 4d spatio-temporal convnets: Minkowski convolutional neural networks. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 3075–3084, 2019. 6
- [10] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 5828–5839, 2017. 2, 6, 7
- [11] Jiacheng Deng and Jiahao Lu. Adaptive siamese masked autoencoder with global optimization for unsupervised point cloud shape correspondence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2699–2707, 2025. 2
- [12] Jiacheng Deng, Chuxin Wang, Jiahao Lu, Jianfeng He, Tianzhu Zhang, Jiyang Yu, and Zhe Zhang. Se-or-net: Self-ensembling orientation-aware network for unsupervised point cloud shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5364–5373, 2023. 2
- [13] Jiacheng Deng, Jiahao Lu, and Tianzhu Zhang. Diff3detr: Agent-based diffusion model for semi-supervised 3d object detection. In *European Conference on Computer Vision*, pages 57–73. Springer, 2024. 2
- [14] Jiacheng Deng, Jiahao Lu, and Tianzhu Zhang. Unsupervised template-assisted point cloud shape correspondence network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5250–5259, 2024. 2
- [15] Jiacheng Deng, Jiahao Lu, Zhixin Cheng, and Wenfei Yang. Diffcorr: Conditional diffusion model with reliable pseudo-label guidance for unsupervised point cloud shape correspondence. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2690–2698, 2025. 2
- [16] Jiacheng Deng, Jiahao Lu, and Tianzhu Zhang. Quantity-quality enhanced self-training network for weakly supervised point cloud semantic segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025. 2
- [17] Theo Deprelle, Thibault Groueix, Matthew Fisher, Vladimir Kim, Bryan Russell, and Mathieu Aubry. Learning elementary structures for 3d shape generation and matching. *Advances in Neural Information Processing Systems*, 32, 2019. 2
- [18] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. Pla: Language-driven open-vocabulary 3d scene understanding. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 7010–7019, 2023. 1, 2
- [19] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 8
- [20] Scott Ettinger, Shuyang Cheng, Charles R Qi, Yin Zhou, et al. Large scale interactive motion forecasting for autonomous driving: The waymo open motion dataset. In *CVPR*, pages 9710–9719, 2021. 1
- [21] Andreas Geiger, Philip Lenz, Christoph Stiller, and Raquel Urtasun. Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research*, 32(11):1231–1237, 2013. 1
- [22] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *European Conference on Computer Vision*, pages 540–557. Springer, 2022. 2, 3

- [23] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *Eur. Conf. Comput. Vis.*, pages 540–557. Springer, 2022. 7, 8
- [24] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 3d-coded: 3d correspondences by deep deformation. In *Proceedings of the european conference on computer vision (ECCV)*, pages 230–246, 2018. 2
- [25] Jun Guo, Xiaojian Ma, Yue Fan, Huaping Liu, and Qing Li. Semantic gaussians: Open-vocabulary scene understanding with 3d gaussian splatting. *arXiv preprint arXiv:2403.15624*, 2024. 8
- [26] Zhening Huang, Xiaoyang Wu, Xi Chen, Hengshuang Zhao, Lei Zhu, and Joan Lasenby. Openins3d: Snap and lookup for 3d open-vocabulary instance segmentation. In *European Conference on Computer Vision*, pages 169–185. Springer, 2024. 8
- [27] Li Jiang, Hengshuang Zhao, Shaoshuai Shi, Shu Liu, Chi-Wing Fu, and Jiaya Jia. Pointgroup: Dual-set point grouping for 3d instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and Pattern recognition*, pages 4867–4876, 2020. 2
- [28] Li Jiang, Shaoshuai Shi, and Bernt Schiele. Open-vocabulary 3d semantic segmentation with foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21284–21294, 2024. 2, 3
- [29] Li Jiang, Shaoshuai Shi, and Bernt Schiele. Open-vocabulary 3d semantic segmentation with foundation models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21284–21294, 2024. 2, 6
- [30] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 3, 8
- [31] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lurf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 3
- [32] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lurf: Language embedded radiance fields. In *Int. Conf. Comput. Vis.*, pages 19729–19739, 2023. 8
- [33] Alexander Kirillov, Eric Mintun, Nikhila Ravi, Hanzi Mao, Chloe Rolland, Laura Gustafson, Tete Xiao, Spencer Whitehead, Alexander C Berg, Wan-Yen Lo, et al. Segment anything. *arXiv preprint arXiv:2304.02643*, 2023. 4, 5
- [34] Maxim Kolodiaznyy, Anna Vorontsova, Anton Konushin, and Danila Rukhovich. Oneformer3d: One transformer for unified point cloud segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20943–20953, 2024. 2
- [35] Xin Lai, Yuhui Yuan, Ruihang Chu, Yukang Chen, Han Hu, and Jiaya Jia. Mask-attention-free transformer for 3d instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3693–3703, 2023. 2
- [36] Loic Landrieu and Martin Simonovsky. Large-scale point cloud semantic segmentation with superpoint graphs. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4558–4567, 2018. 2
- [37] Alex H Lang, Sourabh Vora, Holger Caesar, Lubing Zhou, Jiong Yang, and Oscar Beijbom. Pointpillars: Fast encoders for object detection from point clouds. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12697–12705, 2019. 2
- [38] Itai Lang, Dvir Ginzburg, Shai Avidan, and Dan Raviv. Dpc: Unsupervised deep point correspondence via cross and self construction. In *2021 International Conference on 3D Vision (3DV)*, pages 1442–1451. IEEE, 2021. 2
- [39] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022. 2, 3
- [40] Boyi Li, Kilian Q. Weinberger, Serge J. Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. In *Int. Conf. Learn. Represent.*, 2022. 1, 3, 4, 5, 7, 8
- [41] Jiale Li, Hang Dai, Hao Han, and Yong Ding. Mseg3d: Multi-modal 3d semantic segmentation for autonomous driving. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21694–21704, 2023. 2
- [42] Yingwei Li, Adams Wei Yu, Tianjian Meng, Ben Caine, Jiquan Ngiam, Daiyi Peng, Junyang Shen, Yifeng Lu, Denny Zhou, Quoc V Le, et al. Deepfusion: Lidar-camera deep fusion for multi-modal 3d object detection. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17182–17191, 2022. 2
- [43] Zhuoyuan Li, Yubo Ai, Jiahao Lu, ChuXin Wang, Jiacheng Deng, Hanzhi Chang, Yanzhe Liang, Wenfei Yang, Shifeng Zhang, and Tianzhu Zhang. Mamba24/8d: Enhancing global interaction in point clouds via state space model. *arXiv preprint arXiv:2406.17442*, 2024. 2
- [44] Zhuoyuan Li, Yubo Ai, Jiahao Lu, ChuXin Wang, Jiacheng Deng, Hanzhi Chang, Yanzhe Liang, Wenfei Yang, Shifeng Zhang, and Tianzhu Zhang. Pamba: Enhancing global interaction in point clouds via state space model. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 5092–5100, 2025. 2
- [45] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yinan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7061–7070, 2023. 2, 3
- [46] Youquan Liu, Lingdong Kong, Jun Cen, Runnan Chen, Wenwei Zhang, Liang Pan, Kai Chen, and Ziwei Liu. Segment any point cloud sequences by distilling vision foundation models. *Advances in Neural Information Processing Systems*, 36:37193–37229, 2023. 6
- [47] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 6

- [48] Jiahao Lu and Jiacheng Deng. Relation3d: Enhancing relation modeling for point cloud instance segmentation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 8889–8899, 2025. 2
- [49] Jiahao Lu, Jiacheng Deng, Chuxin Wang, Jianfeng He, and Tianzhu Zhang. Query refinement transformer for 3d instance segmentation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18516–18526, 2023.
- [50] Jiahao Lu, Jiacheng Deng, and Tianzhu Zhang. Bsnet: Box-supervised simulation-assisted mean teacher for 3d instance segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20374–20384, 2024.
- [51] Jiahao Lu, Jiacheng Deng, and Tianzhu Zhang. Beyond the final layer: Hierarchical query fusion transformer with agent-interpolation initialization for 3d instance segmentation. *arXiv preprint arXiv:2502.04139*, 2025. 2
- [52] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 3, 8
- [53] Alexey Nekrasov, Jonas Schult, Or Litany, Bastian Leibe, and Francis Engelmann. Mix3d: Out-of-context data augmentation for 3d scenes. In *2021 international conference on 3d vision (3dv)*, pages 116–125. IEEE, 2021. 6
- [54] Ting Pan, Lulu Tang, Xinlong Wang, and Shiguang Shan. Tokenize anything via prompting. *arXiv preprint arXiv:2312.09128*, 2023. 4
- [55] Xuran Pan, Zhuofan Xia, Shiji Song, Li Erran Li, and Gao Huang. 3d object detection with pointformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7463–7472, 2021. 2
- [56] Kyeong-Beom Park, Minseok Kim, and Jae Yeol Lee. Deep learning-based smart task assistance in wearable augmented reality. *Robotics and Computer-Integrated Manufacturing*, page 101887, 2020. 1
- [57] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 815–824, 2023. 1, 2, 4, 6, 7, 8
- [58] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 815–824, 2023. 2, 3
- [59] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017. 2
- [60] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NeurIPS*, 2017. 6
- [61] 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. 2
- [62] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20051–20060, 2024. 8
- [63] Minghan Qin, Wanhua Li, Jiawei Zhou, Haoqian Wang, and Hanspeter Pfister. Langsplat: 3d language gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20051–20060, 2024. 3
- [64] 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. pages 8748–8763. PMLR, 2021. 1, 4, 5, 6, 8
- [65] 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. 2, 8
- [66] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models, 2021. 2, 4
- [67] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 3
- [68] Jonas Schult, Francis Engelmann, Alexander Hermans, Or Litany, Siyu Tang, and Bastian Leibe. Mask3d for 3d semantic instance segmentation. *arXiv preprint arXiv:2210.03105*, 2022. 8
- [69] Daniel Seita, Yufei Wang, Sarthak J Shetty, Zackory Erickson, and David Held. Toolflownet: Robotic manipulation with tools via predicting tool flow from point clouds. In *Conference on Robot Learning*, pages 1038–1049. PMLR, 2023. 1
- [70] Yichao Shen, Zigang Geng, Yuhui Yuan, Yutong Lin, Ze Liu, Chunyu Wang, Han Hu, Nanning Zheng, and Baining Guo. V-detr: Detr with vertex relative position encoding for 3d object detection. *arXiv preprint arXiv:2308.04409*, 2023. 2
- [71] Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Norman Müller, Matthias Nießner, Angela Dai, and Peter Kotschieder. Panoptic lifting for 3d scene understanding with neural fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 9043–9052, 2023. 8
- [72] Jiahao Sun, Chunmei Qing, Junpeng Tan, and Xiangmin Xu. Superpoint transformer for 3d scene instance segmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2393–2401, 2023. 2
- [73] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *In-*

- ternational conference on machine learning, pages 6105–6114. PMLR, 2019. 8
- [74] Zaizuo Tang, Guangzhu Chen, Yinhe Han, Xiaojuan Liao, Qingjun Ru, and Yuanyuan Wu. Bi-stage multi-modal 3d instance segmentation method for production workshop scene. *Engineering Applications of Artificial Intelligence*, 112:104858, 2022. 2
- [75] Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. *Advances in neural information processing systems*, 30, 2017. 3
- [76] Hugues Thomas, Charles R Qi, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *ICCV*, pages 6411–6420, 2019. 6
- [77] Thang Vu, Kookhoi Kim, Tung M Luu, Thanh Nguyen, and Chang D Yoo. Softgroup for 3d instance segmentation on point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2708–2717, 2022. 2
- [78] Bing Wang, Lu Chen, and Bo Yang. Dm-nerf: 3d scene geometry decomposition and manipulation from 2d images. *arXiv preprint arXiv:2208.07227*, 2022. 8
- [79] Li Wang, Xinyu Zhang, Ziyang Song, Jiangfeng Bi, Guoxin Zhang, Haiyue Wei, Liyao Tang, Lei Yang, Jun Li, Caiyan Jia, et al. Multi-modal 3d object detection in autonomous driving: A survey and taxonomy. *IEEE Transactions on Intelligent Vehicles*, 8(7):3781–3798, 2023. 2
- [80] Pengfei Wang, Yuxi Wang, Shuai Li, Zhaoxiang Zhang, Zhen Lei, and Lei Zhang. Open vocabulary 3d scene understanding via geometry guided self-distillation. In *European Conference on Computer Vision*, pages 442–460. Springer, 2024. 2, 5, 6
- [81] Pengfei Wang, Yuxi Wang, Shuai Li, Zhaoxiang Zhang, Zhen Lei, and Lei Zhang. Open vocabulary 3d scene understanding via geometry guided self-distillation. In *European Conference on Computer Vision*, pages 442–460. Springer, 2024. 2, 3
- [82] Wenxuan Wu, Zhongang Qi, and Li Fuxin. Pointconv: Deep convolutional networks on 3d point clouds. In *CVPR*, pages 9621–9630, 2019. 6
- [83] Weijia Wu, Yuzhong Zhao, Mike Zheng Shou, Hong Zhou, and Chunhua Shen. Diffumask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1206–1217, 2023. 4
- [84] Xiaoyang Wu, Li Jiang, Wanli Ouyang, Tong He, and Hengshuang Zhao. Point transformer v3: Simpler faster stronger. In *CVPR*, 2024. 6
- [85] Xiaoyang Wu, Li Jiang, Peng-Shuai Wang, Zhijian Liu, Xihui Liu, Yu Qiao, Wanli Ouyang, Tong He, and Hengshuang Zhao. Point transformer v3: Simpler faster stronger. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4840–4851, 2024. 2
- [86] Jiarui Xu, Sifei Liu, Arash Vahdat, Wonmin Byeon, Xiaolong Wang, and Shalini De Mello. Open-Vocabulary Panoptic Segmentation with Text-to-Image Diffusion Models. *arXiv preprint arXiv:2303.04803*, 2023. 7
- [87] Mengde Xu, Zheng Zhang, Fangyun Wei, Han Hu, and Xiang Bai. Side adapter network for open-vocabulary semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2945–2954, 2023. 2, 3
- [88] Jihan Yang, Runyu Ding, Weipeng Deng, Zhe Wang, and Xiaojuan Qi. Regionplc: Regional point-language contrastive learning for open-world 3d scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024. 2
- [89] Dongqiangzi Ye, Zixiang Zhou, Weijia Chen, Yufei Xie, Yu Wang, Panqu Wang, and Hassan Foroosh. Lidarmultinet: Towards a unified multi-task network for lidar perception. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3231–3240, 2023. 6
- [90] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. In *European Conference on Computer Vision*, pages 162–179. Springer, 2024. 3
- [91] Yiming Zeng, Yue Qian, Zhiyu Zhu, Junhui Hou, Hui Yuan, and Ying He. Corrnnet3d: Unsupervised end-to-end learning of dense correspondence for 3d point clouds. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6052–6061, 2021. 2
- [92] Junbo Zhang, Runpei Dong, and Kaisheng Ma. Clip-fo3d: Learning free open-world 3d scene representations from 2d dense clip. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2048–2059, 2023. 2
- [93] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8552–8562, 2022. 8
- [94] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021. 2
- [95] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *IEEE Conf. Comput. Vis. Pattern Recog.*, pages 15838–15847, 2021. 8
- [96] Shijie Zhou, Haoran Chang, Sicheng Jiang, Zhiwen Fan, Zehao Zhu, Dejia Xu, Pradyumna Chari, Suyu You, Zhangyang Wang, and Achuta Kadambi. Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21676–21685, 2024. 8
- [97] Yin Zhou and Oncel Tuzel. Voxelnet: End-to-end learning for point cloud based 3d object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4490–4499, 2018. 2
- [98] Xiaoyu Zhu, Hao Zhou, Pengfei Xing, Long Zhao, Hao Xu, Junwei Liang, Alexander Hauptmann, Ting Liu, and

Andrew Gallagher. Open-vocabulary 3d semantic segmentation with text-to-image diffusion models. In *European Conference on Computer Vision*, pages 357–375. Springer, 2024. [2](#), [6](#)

- [99] Xiaoyu Zhu, Hao Zhou, Pengfei Xing, Long Zhao, Hao Xu, Junwei Liang, Alexander Hauptmann, Ting Liu, and Andrew Gallagher. Open-vocabulary 3d semantic segmentation with text-to-image diffusion models. In *European Conference on Computer Vision*, pages 357–375. Springer, 2024. [2](#), [3](#)
- [100] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *Advances in neural information processing systems*, 36: 19769–19782, 2023. [1](#), [2](#), [4](#), [5](#), [7](#), [8](#)
- [101] Xueyan Zou, Jianwei Yang, Hao Zhang, Feng Li, Linjie Li, Jianfeng Wang, Lijuan Wang, Jianfeng Gao, and Yong Jae Lee. Segment everything everywhere all at once. *Advances in Neural Information Processing Systems*, 36, 2024. [2](#), [3](#), [8](#)