
Weakly Supervised 3D Open-vocabulary Segmentation

Kunhao Liu¹ Fangneng Zhan² Jiahui Zhang¹ Muyu Xu¹ Yingchen Yu¹
Abdulmotaleb El Saddik^{3,5} Christian Theobalt² Eric Xing^{4,5} Shijian Lu^{1*}

¹Nanyang Technological University ²Max Planck Institute for Informatics
³University of Ottawa ⁴Carnegie Mellon University ⁵MBZUAI

Abstract

Open-vocabulary segmentation of 3D scenes is a fundamental function of human perception and thus a crucial objective in computer vision research. However, this task is heavily impeded by the lack of large-scale and diverse 3D open-vocabulary segmentation datasets for training robust and generalizable models. Distilling knowledge from pre-trained 2D open-vocabulary segmentation models helps but it compromises the open-vocabulary feature as the 2D models are mostly finetuned with close-vocabulary datasets. We tackle the challenges in 3D open-vocabulary segmentation by exploiting pre-trained foundation models CLIP and DINO in a weakly supervised manner. Specifically, given only the open-vocabulary text descriptions of the objects in a scene, we distill the open-vocabulary multimodal knowledge and object reasoning capability of CLIP and DINO into a neural radiance field (NeRF), which effectively lifts 2D features into view-consistent 3D segmentation. A notable aspect of our approach is that it does not require any manual segmentation annotations for either the foundation models or the distillation process. Extensive experiments show that our method even outperforms fully supervised models trained with segmentation annotations in certain scenes, suggesting that 3D open-vocabulary segmentation can be effectively learned from 2D images and text-image pairs. Code is available at <https://github.com/Kunhao-Liu/3D-OVS>.

1 Introduction

Semantic segmentation of 3D scenes holds significant research value due to its broad range of applications such as robot navigation [1], object localization [2], autonomous driving [3], 3D scene editing [4], augmented/virtual reality, etc. Given the super-rich semantics in 3D scenes, a crucial aspect of this task is achieving open-vocabulary segmentation that can handle regions and objects of various semantics including those with long-tail distributions. This is a grand challenge as it necessitates a comprehensive understanding of natural language and the corresponding objects in the 3D world.

The main challenge in open-vocabulary 3D scene segmentation is the lack of large-scale and diverse 3D segmentation datasets. Existing 3D segmentation datasets like ScanNet [5] primarily focus on restricted scenes with limited object classes, making them unsuitable for training open-vocabulary models. An alternative is to distill knowledge from pre-trained 2D open-vocabulary segmentation models to 3D representations as learned with NeRF [6] or point clouds, by fitting the feature maps or segmentation probability outputs from the 2D models [4, 7]. Though this approach circumvents the need for the 3D datasets, it inherits the limitations of the 2D models which are usually finetuned with close-vocabulary datasets of limited text labels [8, 9], thereby compromising the open-vocabulary property, especially for text labels with long-tail distributions [2, 3].

*Corresponding author

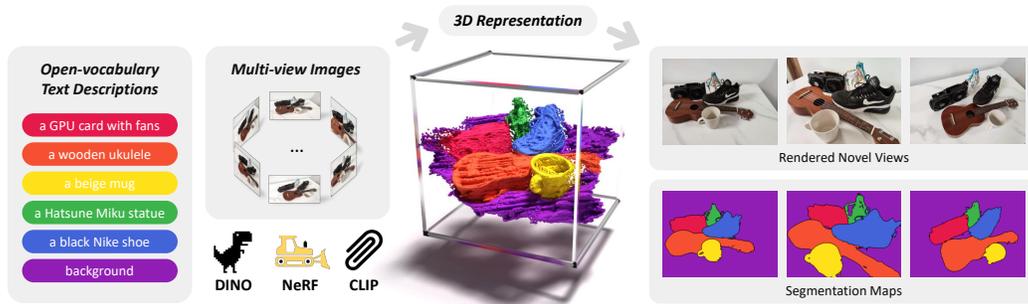


Figure 1: **Weakly Supervised 3D Open-vocabulary Segmentation.** Given the multi-view images of a 3D scene and the open-vocabulary text descriptions, our method distills open-vocabulary multimodal knowledge from CLIP and object reasoning ability from DINO into the reconstructed NeRF, producing accurate object boundaries for the 3D scene without requiring any segmentation annotations during training.

We achieve precise and annotation-free 3D open-vocabulary segmentation by distilling knowledge from two pre-trained foundation models into NeRF in a weakly supervised manner, supervised only by the open-vocabulary text descriptions of the objects in a scene, as illustrated in Fig. 1. One foundation model is CLIP [10] which is trained with Internet-scale text-image pairs [11] capturing extensive open-vocabulary multimodal knowledge. The other is DINO [12, 13] which is trained with large-scale unlabelled images capturing superb scene layout and object boundary information. However, CLIP yields image-level features which are not suitable for pixel-level semantic segmentation. Thus certain mechanisms should be designed to extract pixel-level CLIP features without fine-tuning. Additionally, the image patches' CLIP features may have ambiguities for segmentation, which need to be regularized for accurate open-vocabulary segmentation. At the other end, DINO produces feature maps instead of explicit segmentation maps. Certain distillation techniques should be designed to extract the necessary information from DINO features to facilitate precise segmentation.

We construct a hierarchical set of image patches to extract pixel-level features from image-level CLIP features and design a 3D *Selection Volume* to identify the appropriate hierarchical level for each 3D point, effectively aligning CLIP features with pixel-level features without fine-tuning. In addition, we introduce a *Relevancy-Distribution Alignment (RDA)* loss to address CLIP feature ambiguities, aligning segmentation probability distribution with class relevancies that capture similarities between class text features and corresponding CLIP features. Moreover, we propose a novel *Feature-Distribution Alignment (FDA)* loss to distill object boundary information from DINO features. The FDA loss encourages close segmentation probability distributions for points with similar DINO features and distant distributions for dissimilar features. To address the training instability due to diverse distribution shapes, we further re-balance weights associated with similar and dissimilar DINO features.

Our method enables weakly supervised open-vocabulary segmentation of 3D scenes with accurate object boundaries. By distilling knowledge from CLIP without fine-tuning, our approach preserves its open-vocabulary knowledge and effectively handles text labels with long-tail distributions. A notable aspect of our approach is that it does not require any manual segmentation annotations for either the foundation models or the distillation process. Remarkably, our experiments demonstrate that our method surpasses fully supervised models trained with segmentation annotations in certain scenes, highlighting the possibility that 3D open-vocabulary segmentation can be effectively learned from large amounts of 2D images and text-image pairs.

In summary, the contributions of this work are three-fold. *Firstly*, we propose an innovative pipeline for weakly supervised 3D open-vocabulary segmentation by distilling knowledge from pre-trained foundation models into NeRF without requiring any annotations in training. *Secondly*, we introduce a Selection Volume to align image-level CLIP features with pixel-level features, supplemented by novel Relevancy-Distribution Alignment and Feature-Distribution Alignment losses that respectively resolve CLIP features' ambiguities and effectively distill DINO features for 3D scene segmentation. *Lastly*, extensive experiments demonstrate that our method effectively recognizes long-tail classes and produces accurate segmentation maps, even with limited input data.

2 Related Work

Open-vocabulary Segmentation. In recent years, the field of 2D open-vocabulary segmentation has garnered significant attention, driven by the availability of extensive text-image datasets and vast computational resources. Predominant approaches [8, 14–19] typically distill knowledge from large-scale pre-trained models, such as image-text contrastive learning models [10, 20–22] and diffusion models [23]. However, the distillation process requires fine-tuning on close-vocabulary datasets, contrasting with massive datasets used for large-scale pre-trained models [11]. This leads to limited performance in recalling infrequent classes with long-tail distributions [2, 3], compromising the open-vocabulary property. OpenSeg [24] is not finetuned on a closed set of classes but is weakly supervised via image captions. However, OpenSeg has a smaller vocabulary and knowledge than CLIP as it is trained on a much smaller dataset. Our method, without fine-tuning CLIP, effectively handles such classes.

3D Scenes Segmentation. 3D scene segmentation has been a long-standing challenge in computer vision. Traditional approaches focus on point clouds or voxels with limited class variety in datasets, restricting generalizability to unseen classes [5, 25–36]. Recently, numerous point-cloud-based techniques have emerged to explore open-vocabulary 3D scene segmentation by encoding 2D open-vocabulary models’ features into 3D scene points [3, 7, 37–40]. However, these methods are also mostly evaluated on datasets with restricted scenes and limited class ranges [5, 27, 28, 41], not fully exhibiting the open-vocabulary property. Moreover, point clouds have compromised geometric details, making them less suitable for precise segmentation compared to NeRF representations [6, 42]. Consequently, there has been a surge in NeRF-based 3D segmentation techniques that mainly address interactive segmentation [4, 43, 44], panoptic segmentation [45, 46], moving part segmentation [47], object part segmentation [48], object co-segmentation [49], unsupervised object segmentation [50, 51], etc. FFD [4] attempts to segment unseen text labels during training by fitting LSeg’s [8] feature maps to a NeRF, but inherits LSeg’s limitations, hindering generalization to long-tail distribution classes. Our method overcomes these challenges by directly using CLIP image features and distilling them into a NeRF representation [42] without fine-tuning on close-vocabulary datasets.

Foundation Models. Pre-trained foundation models [52, 53] have become a powerful paradigm in computer science due to their ability to capture general knowledge and adapt to various downstream tasks [10, 12, 13, 20, 23, 54–57]. These models are trained using various paradigms in natural language processing, such as masked language modeling [57, 58], denoising autoencoder [59], replaced token detection [60], and sentence prediction tasks [61], as well as in computer vision, including data generation [23, 62, 63], data reconstruction [64], and data contrastive learning [10, 12, 13, 20–22]. Foundation models acquire emergent capabilities for exceptional performance on downstream tasks, either in a zero-shot manner or with fine-tuning. In this work, we harness the capabilities of two prominent foundation models, CLIP [10] and DINO [12, 13]. CLIP learns associations between images and texts by mapping them to a shared space, facilitating applications in tasks like image classification, object detection, visual question-answering, and image generation [9, 10, 23, 62, 65, 66]. DINO, trained in a self-supervised manner, extracts scene layout information, particularly object boundaries, and has been successfully employed in tasks such as classification, detection, segmentation, keypoint estimation, depth estimation, and image editing [12, 13, 49, 67–69].

3 Method

We propose a novel method for weakly supervised open-vocabulary segmentation of reconstructed NeRF. Given the multi-view images of a scene and the open-vocabulary text description for each class, we aim to segment the reconstructed NeRF such that every 3D point is assigned a corresponding class label.

To achieve this, we exploit the CLIP model’s multimodal knowledge by mapping each 3D point to a CLIP feature representing its semantic meaning. As CLIP only generates image-level features, we extract a hierarchy of CLIP features from image patches and learn a 3D *Selection Volume* for pixel-level feature extraction, as described in Sec. 3.1.

Algorithm 1: Extracting pixel-level features of an image from CLIP

Input: RGB image $I \in \mathbb{R}^{3 \times H \times W}$, number of scale N_s , CLIP image encoder E_i

Output: multi-scale pixel-level features $F_I \in \mathbb{R}^{N_s \times D \times H \times W}$

Initialize: the patch size of each scale $patch_sizes \in \mathbb{R}^{N_s}$, $F_I = \text{zeros}(N_s, D, H, W)$

```
/* Loop over all the scales */
for scale_idx, patch_size in enumerate(patch_sizes) do
    stride = patch_size/4
    count = zeros(1, 1, H, W) /* Record the patch count for each pixel */
    F_multi_spatial = zeros(1, D, H, W) /* Multi-spatial feature of current scale */
    /* Loop over all the patches */
    for x_idx in range((H - patch_size)/stride + 1) do
        start_x = x_idx * stride
        for y_idx in range((W - patch_size)/stride + 1) do
            start_y = y_idx * stride
            /* Get image patch's coordinates with randomness */
            (left, upper, right, lower) = (max(start_y - randint(0, stride), 0), max(start_x - randint(0, stride), 0),
            min(start_y + patch_size + randint(0, stride), W), min(start_x + patch_size + randint(0, stride), H))
            /* Get image patch's CLIP feature */
            F_patch = E_i(I.crop(left, upper, right, lower))
            F_multi_spatial[:, :, upper : lower, left : right] += F_patch
            count[:, :, upper : lower, left : right] += 1
        end
    end
    F_multi_spatial /= count
    F_I[scale_idx] = F_multi_spatial
end
```

However, the image patches' CLIP features may have ambiguities for segmentation, showing inaccurate absolute relevancy values which lead to misclassification. Take an image patch capturing an apple lying on a lawn as an example. The corresponding CLIP feature contains both apple and lawn information. If the apple is relatively small, the patch could be classified as lawn because lawn dominates the patch's CLIP features. Then the class apple would be ignored. To address this issue, we introduce a *Relevancy-Distribution Alignment (RDA)* loss, aligning the segmentation probability distribution with each class's normalized relevancy map, as described in Sec. 3.2. For precise object boundaries, we align the segmentation probability distribution with the images' DINO features. Previous segmentation approaches utilizing DINO features have focused on unsupervised segmentation [49, 68], lacking semantic meaning for segmented parts. In the open-vocabulary context, assigning accurate text labels to segmented regions requires aligning DINO feature clusters with correct text labels. To overcome this challenge, we propose a *Feature-Distribution Alignment (FDA)* loss, segmenting the scene based on DINO features' distribution and assigning appropriate text labels, as described in Sec. 3.3.

3.1 Distilling Pixel-level CLIP Features with a 3D Selection Volume

We propose a method based on multi-scale and multi-spatial strategies to adapt CLIP's image-level features for pixel-level segmentation, motivated by the observation that a pixel's semantic meaning should remain invariant to its surrounding pixels. The multi-scale component extracts features from patches of varying sizes around each pixel, and the multi-spatial component extracts from patches in which each pixel is at different positions. We average the multi-spatial features in each scale, attributing the primary direction of these features to the pixel's semantic meaning. We utilize a sliding-window algorithm for multi-spatial feature extraction. To prevent checkerboard patterns in the features and potential segmentation artifacts, we introduce randomness in the window size. The algorithm's pseudo-code is in Alg. 1.

After extracting pixel-level features from CLIP, we now have the multi-view RGB images and their corresponding multi-scale pixel-level features. Each ray \mathbf{r} is assigned its multi-scale feature $F(\mathbf{r}) \in \mathbb{R}^{N_s \times D}$ for supervision. Rather than simply averaging each ray's multi-scale features across scales, we introduce a 3D Selection Volume S to determine the most suitable scale indicative of the object size within a patch. Each 3D point $\mathbf{x} \in \mathbb{R}^3$ yields a selection vector $S_{\mathbf{x}} \in \mathbb{R}^{N_s}$ from S . Following [4, 45], we introduce an additional branch to render the CLIP feature. We can then render the RGB value, the CLIP feature, and the selection vector of each ray \mathbf{r} using volume rendering [6]:

$$\hat{C}(\mathbf{r}) = \sum_i T_i \alpha_i C_i \in \mathbb{R}^3, \quad \hat{F}(\mathbf{r}) = \sum_i T_i \alpha_i F_i \in \mathbb{R}^D, \quad (1)$$

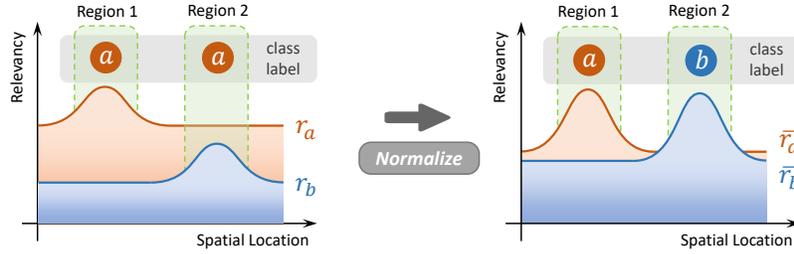


Figure 2: **Mitigating CLIP features' ambiguities with normalized relevancy maps.** For original relevancy maps r_a, r_b of classes a and b , we note a higher relevancy for class b in Region 2 than in other image regions. Despite this, the ambiguities of CLIP features lead to Region 2's classification as a due to the higher absolute relevancy of a in Region 2, even as a is located in Region 1. To rectify this, we normalize each class's relevancy maps to a fixed range. These normalized relevancy maps, \bar{r}_a and \bar{r}_b , reduce such ambiguities, facilitating accurate region-class assignments.

$$S(\mathbf{r}) = \text{Softmax} \left(\sum_i T_i \alpha_i S_i \right) \in [0, 1]^{N_s}, \quad (2)$$

where C_i, F_i, S_i are the color, feature, and selection vector of each sampled point along the ray, $T_i = \prod_{j=0}^{i-1} (1 - \alpha_j)$ is the accumulated transmittance and $\alpha_i = 1 - \exp(-\delta_i \sigma_i)$ is the opacity of the point. We apply a Softmax function to the selection vector of each ray such that the sum of the probability of each scale is equal to 1.

For a set of rays \mathcal{R} in each training batch, the supervision loss can then be formulated as the combination of the L2 distance between rendered and ground truth RGB values and the cosine similarities $\cos(\cdot, \cdot)$ between the rendered features and the selected multi-scale CLIP features:

$$\mathcal{L}_{supervision} = \sum_{\mathbf{r} \in \mathcal{R}} \left(\left\| \hat{C}(\mathbf{r}) - C(\mathbf{r}) \right\|_2^2 - \cos \langle \hat{F}(\mathbf{r}), S(\mathbf{r}) F(\mathbf{r}) \rangle \right). \quad (3)$$

Given a set of text descriptions $\{[\text{CLASS}]_i\}_{i=1}^C$ of C classes and the CLIP text encoder E_t , we can get the classes' text features $T = E_t([\text{CLASS}]) \in \mathbb{R}^{C \times D}$. Then we can get the segmentation logits $z(\mathbf{r})$ of the ray \mathbf{r} by computing the cosine similarities between the rendered CLIP feature and the classes' text features:

$$z(\mathbf{r}) = \cos \langle T, \hat{F}(\mathbf{r}) \rangle \in \mathbb{R}^C. \quad (4)$$

We can then get the class label of the ray $l(\mathbf{r}) = \text{argmax}(z(\mathbf{r}))$.

3.2 Relevancy-Distribution Alignment for Ambiguity Mitigation

To mitigate the ambiguities of the CLIP features, we propose to align the segmentation probability distribution with the spatially normalized relevancy maps of each class, enabling our method to identify specific image regions described by each class text, as illustrated in Fig. 2. The segmentation probability of each ray $P(\mathbf{r})$ can be derived from the segmentation logits with a Softmax function:

$$P(\mathbf{r}) = \text{Softmax}(z(\mathbf{r})) \in [0, 1]^C. \quad (5)$$

The relevancy of a given class is determined by the similarity between the class's text feature and the selected feature from the hierarchy of image patches' CLIP features. Given an image I , we can get its multi-scale pixel-level CLIP feature $F_I \in \mathbb{R}^{N_s \times D \times H \times W}$ using Alg. 1 and selection vector $S_I \in \mathbb{R}^{N_s \times H \times W}$ using Eq. (2). And then we can get the image's relevancy map $R_I \in \mathbb{R}^{C \times H \times W}$ as:

$$R_{I_{hw}} = S_{I_{hw}} \cos \langle T, F_{I_{hw}} \rangle, \quad (6)$$

where where h, w denotes the index in the H and W channel. We normalize each class's relevancy independently within an input view to $[0, 1]$ to mitigate the ambiguities of CLIP features, making our method discern image regions described by each class text:

$$\bar{R}_I = (R_I - \min(R_I)) / (\max(R_I) - \min(R_I)) \in [0, 1]^{C \times H \times W}, \quad (7)$$

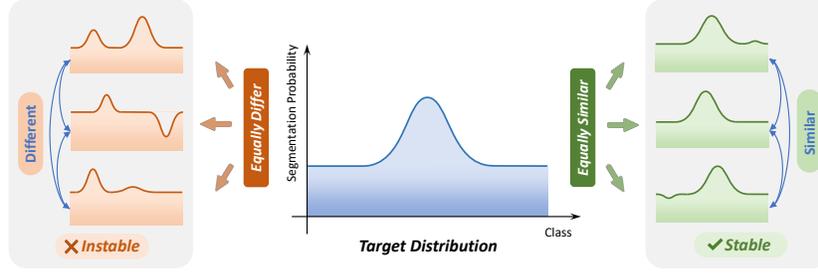


Figure 3: **Difference between similar and distant distributions.** Distributions having large divergence from the target distribution exhibit significantly diverse shapes, increasing the training instability (left). Conversely, distributions displaying low divergence with the target distribution consistently demonstrate a similar shape (right).

where $\min()$ and $\max()$ are the functions getting the lowest and highest values across the spatial dimensions (i.e. H and W). We apply a Softmax function to \bar{R}_I to make it a probability vector. Then we can assign each ray \mathbf{r} its normalized relevancy with all the classes $\bar{R}(\mathbf{r}) \in [0, 1]^C$. We employ the Jensen-Shannon (JS) divergence to measure the discrepancy between the normalized relevancy $\bar{R}(\mathbf{r})$ and the segmentation probability distribution $P(\mathbf{r})$ of each ray, formulating the Relevancy-Distribution Alignment (RDA) loss:

$$\mathcal{L}_{RDA} = \sum_{\mathbf{r} \in \mathcal{R}} \sum_{c \in C} \left(P(\mathbf{r})_c \log \left(\frac{P(\mathbf{r})_c}{M_{P\bar{R}}(\mathbf{r})_c} \right) + \bar{R}(\mathbf{r})_c \log \left(\frac{\bar{R}(\mathbf{r})_c}{M_{P\bar{R}}(\mathbf{r})_c} \right) \right) / 2, \quad (8)$$

where $M_{P\bar{R}}(\mathbf{r}) = (P(\mathbf{r}) + \bar{R}(\mathbf{r}))/2$ is the average of the two distributions, and the subscript c denotes the probability of the c th class. By aligning the normalized relevancies and the segmentation probability distributions, our method can effectively identify the specific region corresponding to the text description of each class.

3.3 Feature-Distribution Alignment for Precise Object Boundary Segmentation

To ensure the segmentation exhibits precise object boundaries, we align the segmentation probability distribution with the images' DINO features, which have been shown to capture superb scene layouts and object boundary information [12, 13]. Following [49, 68], we extract the scene layout information with a DINO feature correlation tensor. Given a patch of size $H_p \times W_p$, we can get the correlation tensor $Corr_F \in \mathbb{R}^{H_p W_p \times H_p W_p}$ as:

$$Corr_F_{hwi j} = \cos \langle f_{hw}, f_{ij} \rangle, \quad (9)$$

whose entries represent the cosine similarity between the DINO features f at spatial positions (h, w) and (i, j) of the patch. In order to construct the correlation tensor for the segmentation probability distribution, we propose utilizing the JS divergence to assess the similarity between segmentation probabilities at two distinct spatial positions. The choice of JS divergence offers several advantages, including its symmetric nature and a bounded range of $[0, 1]$, which contribute to improved numerical stability. However, since we only care about the class label of each point, i.e. the entry with the highest probability, we use a low temperature $\tau < 1$ to get a sharper version of the segmentation probability distribution \dot{P} to let the model focus on the entry with the largest probability:

$$\dot{P} = \text{Softmax}(z/\tau) \in [0, 1]^C. \quad (10)$$

The distribution correlation tensor $Corr_D \in \mathbb{R}^{H_p W_p \times H_p W_p}$ can thus be computed with:

$$Corr_D_{hwi j} = \sum_{c \in C} \left(\dot{P}_{hwc} \log \left(\frac{\dot{P}_{hwc}}{M_{\dot{P}\dot{P}}(\mathbf{r})_c} \right) + \dot{P}_{ijc} \log \left(\frac{\dot{P}_{ijc}}{M_{\dot{P}\dot{P}}(\mathbf{r})_c} \right) \right) / 2, \quad (11)$$

where $\dot{P}_{hwc}, \dot{P}_{ijc}$ are the segmentation probabilities of the c th class at spatial locations (h, w) and (i, j) of the patch, $M_{\dot{P}\dot{P}}(\mathbf{r}) = (\dot{P}_{hw} + \dot{P}_{ij})/2$ is the average of the two distributions. Thus the correlation loss [68] can be expressed as:

$$\mathcal{L}_{corr} = \sum_{hwi j} (Corr_F_{hwi j} - b) \times Corr_D_{hwi j}, \quad (12)$$

Table 1: **Quantitative comparisons.** We report the mIoU(\uparrow) scores and the Accuracy(\uparrow) scores of the following methods in 6 scenes and highlight the **best**, **second-best**, and **third-best** scores. Our method outperforms both 2D and 3D methods without any segmentation annotations in training.

	Methods	<i>bed</i>		<i>sofa</i>		<i>lawn</i>		<i>room</i>		<i>bench</i>		<i>table</i>	
		mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy
2D	LSeg [8]	56.0	87.6	04.5	16.5	17.5	77.5	19.2	46.1	06.0	42.7	07.6	29.9
	ODISE [16]	52.6	86.5	48.3	35.4	39.8	82.5	52.5	59.7	24.1	39.0	39.7	34.5
	OV-Seg [19]	79.8	40.4	66.1	69.6	81.2	92.1	71.4	49.1	88.9	89.2	80.6	65.3
3D	FFD [4]	56.6	86.9	03.7	09.5	42.9	82.6	25.1	51.4	06.1	42.8	07.9	30.1
	Sem(ODISE) [45]	50.3	86.5	27.7	22.2	24.2	80.5	29.5	61.5	25.6	56.4	18.4	30.8
	Sem(OV-Seg) [45]	89.3	96.7	66.3	89.0	87.6	95.4	53.8	81.9	94.2	98.5	83.8	94.6
	LERF [2]	73.5	86.9	27.0	43.8	73.7	93.5	46.6	79.8	53.2	79.7	33.4	41.0
	Ours	89.5	96.7	74.0	91.6	88.2	97.3	92.8	98.9	89.3	96.3	88.8	96.5

where b is a hyper-parameter denoting that we consider the segmentation probabilities of two spatial locations (h, w) and (i, j) to be similar if their DINO features’ similarity is larger than b and distant if less than b . Nonetheless, the correlation loss \mathcal{L}_{corr} introduces significant instability due to the diverse shapes of distributions with large divergence from a target distribution, making the loss assign wrong labels to the segmented parts. Conversely, when a distribution displays a low JS divergence with the target distribution, it consistently demonstrates a similar shape to the target distribution, as shown in Fig. 3. Based on this observation, we propose re-balancing the weights associated with similar and dissimilar DINO features. Specifically, we allocate a much greater weight to the correlation loss arising from similar DINO features and a smaller weight to that of dissimilar DINO features, thereby mitigating the instability caused by the correlation loss. Thus the Feature-Distribution Alignment (FDA) loss can be formulated with:

$$pos_F = \text{clamp}(Corr_F - b, \min = 0), \quad neg_F = \text{clamp}(Corr_F - b, \max = 0), \quad (13)$$

$$\mathcal{L}_{FDA} = \lambda_{pos} \sum_{hwi j} (pos_F_{hwi j} \times Corr_D_{hwi j}) / \text{count_nonzero}(pos_F) + \lambda_{neg} \sum_{hwi j} (neg_F_{hwi j} \times Corr_D_{hwi j}) / \text{count_nonzero}(neg_F), \quad (14)$$

where $\text{clamp}(\cdot, \min/\max = 0)$ is to clamp all the elements smaller/greater than 0 to 0, thus making $pos_F \geq 0$ and $neg_F \leq 0$, $\text{count_nonzero}(\cdot)$ is to count the number of non zero elements, and $\lambda_{pos}, \lambda_{neg}$ are the weights associated with similar and dissimilar DINO features, which are set as $\lambda_{pos} = 200 \gg \lambda_{neg} = 0.2$ by default.

The total training loss is: $\mathcal{L} = \mathcal{L}_{supervision} + \mathcal{L}_{RDA} + \mathcal{L}_{FDA}$, where $\mathcal{L}_{supervision}, \mathcal{L}_{RDA}$ are calculated with randomly sampled rays and \mathcal{L}_{FDA} is calculated with randomly sampled patches.

4 Experiments

We evaluate our method on 3D open-vocabulary segmentation, showing that our method can recognize long-tail classes and produce highly accurate object boundaries even with limited input data. We employ TensorRF [42] as the backbone and extract 3 scales of pixel-level CLIP features. More implementation details and experiments are in the appendix.

Dataset. Existing 3D segmentation datasets predominantly focus on either restricted scenes with a narrow range of object classes [5, 70], or individual objects [71–73], thereby limiting their capacity to fully assess the task of 3D open-vocabulary segmentation. Thus following [2], we create a dataset comprising 10 distinct scenes. Each scene features a set of long-tail objects situated in various poses and backgrounds. Ground truth masks for the test views are manually annotated, enabling both qualitative and quantitative evaluation of our segmentation methods. We also evaluate our method on more diverse datasets which include human body, human head, indoor scenes with low-quality images [70, 74], and a complex scene from LERF datasets [2], as shown in the appendix.

Baselines. We benchmark our method with three NeRF-based methods capable of 3D open-vocabulary segmentation: FFD [4], Semantic-NeRF (abbreviated as Sem) [45], and LERF [2].



Figure 4: **Qualitative comparisons.** Visualization of the segmentation results in 3 scenes. Our method successfully recognizes long-tail classes and produces the most accurate segmentation maps.

FFD, a pioneering effort in 3D open-vocabulary segmentation, applies the feature maps of the 2D open-vocabulary segmentation method, LSeg [8], to achieve its segmentation results. For Sem, we adapt it to the open-vocabulary segmentation context by distilling segmentation results from two state-of-the-art 2D open-vocabulary segmentation models: the diffusion-model-based ODISE [16], and the CLIP-based OV-Seg [19]. LERF closely aligns with our proposed approach due to its use of knowledge distillation from CLIP and DINO. However, its primary focus is on object localization rather than segmentation. We use the same scale level number and patch sizes in LERF for fair comparisons. We also include results obtained by independently segmenting each test view using the aforementioned 2D models. Note that under our settings, FFD, Sem are fully supervised methods using segmentation annotations.

4.1 Results

We present the qualitative results in Fig. 4 and quantitative results in Tab. 1. Our proposed method outperforms all other techniques, including those which heavily rely on extensive segmentation annotations, such as LSeg, ODISE, OV-Seg. In particular, ODISE and FFD underperform in our evaluation, as they are unable to identify many long-tail classes, suggesting that the fine-tuning process of LSeg and ODISE may have led to a significant loss of the open-vocabulary knowledge originally encapsulated by CLIP and Stable Diffusion [23]. OV-Seg attempts to retain CLIP’s knowledge by leveraging a mined dataset, however, it requires a mask proposal model which produces inconsistent segmentation across views, making Sem(OV-Seg) produce noisy and imprecise segmentation. LERF

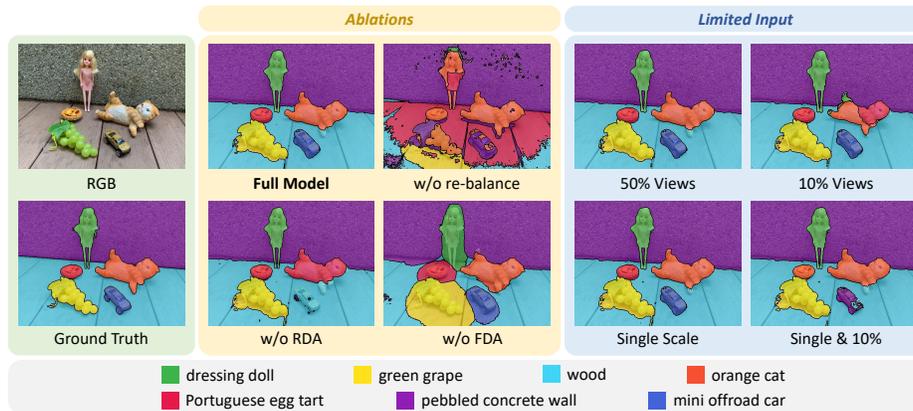


Figure 5: **Studies.** Visualization of the studies on ablations and limited input.

also fails to capture precise object boundaries due to its usage of a relatively naïve regularization loss, which fails to fully exploit the object boundary information within the DINO features. In contrast, our method exhibits robust performance, successfully recognizing long-tail classes and generating accurate and well-defined boundaries for each class. However, LERF allows querying any object without the need for running optimization again, which is an advantage over our method.

4.2 Studies

Ablations. We conduct ablation studies to evaluate the individual contributions of the RDA loss and the FDA loss to the overall performance of our proposed method. As shown in Tab. 2, both RDA loss and FDA loss are crucial to our method, without each of which can result in severe performance degradation. As illustrated in Fig. 5, without the RDA loss, the model does not resolve the ambiguities of the CLIP features, leading to misclassifications. For instance, it fails to distinguish between an *orange cat* and a *Portuguese egg tart*, and confuses a *mini offroad car* with *wood*. Without the FDA loss, although our method can correctly locate each class, it fails to segment precise object boundaries. When discarding the re-balancing in the FDA loss, i.e. using the correlation loss [68], the model produces accurate boundaries but assigns each cluster the wrong label due to the instability brought by diverse distribution shapes.

Table 2: **Studies.** We report the mean mIoU(\uparrow) scores and the Accuracy(\uparrow) scores in our studies.

	mIoU	Accuracy
w/o RDA loss	57.2	79.4
w/o FDA loss	58.2	82.7
w/o re-balance	44.9	74.3
50% views	85.7	95.7
10% views	79.1	94.6
single scale	85.2	95.5
single & 10%	77.1	94.6
full model	86.2	95.8

Limited input. Given the substantial computational and storage demands of extracting hierarchical CLIP features for each view (exceeding 1GB for 3 scales in the full model), we explore whether reducing input CLIP features would yield similar results, as shown in Tab. 2 and Fig. 5. We test two modifications: reducing views for feature extraction and using a single scale of CLIP features rather than three. Halving the input views for feature extraction leads to negligible performance degradation ($< 1\%$) and minor visual differences compared to the full model. When reducing to only 10% of input views, equivalent to 2-3 views in our dataset, we observe a modest 9% drop in the mIoU score and a 1% decrease in the Accuracy score, while retaining accurate segmentation across most classes. Using a single scale of CLIP features also only incurs minimal degradation ($< 1\%$). Even under extreme conditions, i.e., extracting a single scale of features from 10% of input views (total only $1/30$ input of the full model), performance degradation is limited to 10%. This efficient approach even outperforms LERF [2] which utilizes all input views and scales, highlighting our method’s robustness.

5 Limitations

The limitations of our method are twofold. First, unlike LERF [2], our method requires text labels before training. To perform segmentation with new text labels, our method needs to be retrained.

Inferring accurate boundaries with open vocabularies is challenging for implicit representations like NeRF, as NeRF learns a continuous representation rather than a discrete one. It is promising to learn object-level discrete representation using NeRF in future work.

Second, since our method has never seen any segmentation maps during training (only weakly supervised by the text labels), it fails to segment complex scenes like indoor datasets [70, 74] with high precision, as shown in the appendix. Our method distills pixel-level CLIP features in a patch-based fashion with a strong inductive bias for compact objects with an aspect ratio close to 1. For objects with large complex shapes and unobvious textures, our method would fail to recognize them.

6 Conclusion

In this study, we address the challenge of 3D open-vocabulary segmentation by distilling knowledge from the pre-trained foundation models CLIP and DINO into reconstructed NeRF in a weakly supervised manner. We distill the open-vocabulary multimodal knowledge from CLIP with a Selection Volume and a novel Relevancy-Distribution Alignment loss to mitigate the ambiguities of CLIP features. In addition, we introduce a novel Feature-Distribution Alignment loss to extract accurate object boundaries by leveraging the scene layout information within DINO features. Our method successfully recognizes long-tail classes and produces precise segmentation maps, even when supplied with limited input data, suggesting the possibility of learning 3D segmentation from 2D images and text-image pairs.

Acknowledgement

We sincerely thank Zuhao Yang, Zeyu Wang, Weijing Tao, and Kunyang Li for collecting the dataset. This project is funded by the Ministry of Education Singapore, under the Tier-1 project scheme with project number RT18/22.

References

- [1] Nur Muhammad Mahi Shafiullah, Chris Paxton, Lerrel Pinto, Soumith Chintala, and Arthur Szlam. Clip-fields: Weakly supervised semantic fields for robotic memory. *arXiv preprint arXiv:2210.05663*, 2022. 1
- [2] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lurf: Language embedded radiance fields. *arXiv preprint arXiv:2303.09553*, 2023. 1, 3, 7, 9, 16, 17
- [3] Krishna Murthy Jatavallabhula, Alihusein Kuwajerwala, Qiao Gu, Mohd Omama, Tao Chen, Shuang Li, Ganesh Iyer, Soroush Saryazdi, Nikhil Keetha, Ayush Tewari, et al. Conceptfusion: Open-set multimodal 3d mapping. *arXiv preprint arXiv:2302.07241*, 2023. 1, 3
- [4] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. *arXiv preprint arXiv:2205.15585*, 2022. 1, 3, 4, 7, 17
- [5] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. 1, 3, 7
- [6] 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. 1, 3, 4
- [7] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. *arXiv preprint arXiv:2211.15654*, 2022. 1, 3
- [8] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and René Ranftl. Language-driven semantic segmentation. *arXiv preprint arXiv:2201.03546*, 2022. 1, 3, 7, 8
- [9] Yiwu Zhong, Jianwei Yang, Pengchuan Zhang, Chunyuan Li, Noel Codella, Liunian Harold Li, Luwei Zhou, Xiyang Dai, Lu Yuan, Yin Li, et al. Regionclip: Region-based language-image pretraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16793–16803, 2022. 1, 3
- [10] 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, 3
- [11] Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis, Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. *arXiv preprint arXiv:2111.02114*, 2021. 2, 3
- [12] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9630–9640, 2021. 2, 3, 6
- [13] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 3, 6
- [14] Huaishao Luo, Junwei Bao, Youzheng Wu, Xiaodong He, and Tianrui Li. Segclip: Patch aggregation with learnable centers for open-vocabulary semantic segmentation. *arXiv preprint arXiv:2211.14813*, 2022. 3
- [15] Jiarui Xu, Shalini De Mello, Sifei Liu, Wonmin Byeon, Thomas Breuel, Jan Kautz, and Xiaolong Wang. Groupvit: Semantic segmentation emerges from text supervision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18134–18144, 2022.
- [16] 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, 8
- [17] Xueyan Zou, Zi-Yi Dou, Jianwei Yang, Zhe Gan, Linjie Li, Chunyuan Li, Xiyang Dai, Harkirat Behl, Jianfeng Wang, Lu Yuan, et al. Generalized decoding for pixel, image, and language. *arXiv preprint arXiv:2212.11270*, 2022.
- [18] Jian Ding, Nan Xue, Gui-Song Xia, and Dengxin Dai. Decoupling zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11583–11592, 2022.
- [19] Feng Liang, Bichen Wu, Xiaoliang Dai, Kunpeng Li, Yanan Zhao, Hang Zhang, Peizhao Zhang, Peter Vajda, and Diana Marculescu. Open-vocabulary semantic segmentation with mask-adapted clip. *arXiv preprint arXiv:2210.04150*, 2022. 3, 7, 8
- [20] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning*, pages 4904–4916. PMLR, 2021. 3
- [21] Yangguang Li, Feng Liang, Lichen Zhao, Yufeng Cui, Wanli Ouyang, Jing Shao, Fengwei Yu, and Junjie Yan. Supervision exists everywhere: A data efficient contrastive language-image pre-training paradigm. *arXiv preprint arXiv:2110.05208*, 2021.
- [22] Norman Mu, Alexander Kirillov, David Wagner, and Saining Xie. Slip: Self-supervision meets language-image pre-training. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXVI*, pages 529–544. Springer, 2022. 3
- [23] 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, 8
- [24] Golnaz Ghiasi, Xiuye Gu, Yin Cui, and Tsung-Yi Lin. Scaling open-vocabulary image segmentation with image-level labels. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXVI*, pages 540–557. Springer, 2022. 3
- [25] Iro Armeni, Sasha Sax, Amir R Zamir, and Silvio Savarese. Joint 2d-3d-semantic data for indoor scene understanding. *arXiv preprint arXiv:1702.01105*, 2017. 3
- [26] Jens Behley, Martin Garbade, Andres Milioto, Jan Quenzel, Sven Behnke, Cyrill Stachniss, and Jurgen Gall. Semantickitti: A dataset for semantic scene understanding of lidar sequences. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9297–9307, 2019.
- [27] 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. 3
- [28] 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. 3
- [29] Dave Zhenyu Chen, Angel X Chang, and Matthias Nießner. Scanrefer: 3d object localization in rgb-d scans using natural language. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XX*, pages 202–221. Springer, 2020.

- [30] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE conference on computer vision and pattern recognition*, pages 3354–3361. IEEE, 2012.
- [31] Binh-Son Hua, Quang-Hieu Pham, Duc Thanh Nguyen, Minh-Khoi Tran, Lap-Fai Craig Yu, and Sai-Kit Yeung. Scenenn: A scene meshes dataset with annotations. *2016 Fourth International Conference on 3D Vision (3DV)*, pages 92–101, 2016.
- [32] Yiyi Liao, Jun Xie, and Andreas Geiger. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45: 3292–3310, 2021.
- [33] Kaichun Mo, Shilin Zhu, Angel X Chang, Li Yi, Subarna Tripathi, Leonidas J Guibas, and Hao Su. Partnet: A large-scale benchmark for fine-grained and hierarchical part-level 3d object understanding. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 909–918, 2019.
- [34] Pei Sun, Henrik Kretzschmar, Xerxes Dotiwalla, Aurelien Chouard, Vijaysai Patnaik, Paul Tsui, James Guo, Yin Zhou, Yuning Chai, Benjamin Caine, et al. Scalability in perception for autonomous driving: Waymo open dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2446–2454, 2020.
- [35] Lyne P. Tchapmi, Christopher Bongsoo Choy, Iro Armeni, JunYoung Gwak, and Silvio Savarese. Segcloud: Semantic segmentation of 3d point clouds. *2017 International Conference on 3D Vision (3DV)*, pages 537–547, 2017.
- [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. 3
- [37] Runyu Ding, Jihan Yang, Chuhui Xue, Wenqing Zhang, Song Bai, and Xiaojuan Qi. Language-driven open-vocabulary 3d scene understanding. *arXiv preprint arXiv:2211.16312*, 2022. 3
- [38] Kirill Mazur, Edgar Sucar, and Andrew J Davison. Feature-realistic neural fusion for real-time, open set scene understanding. *arXiv preprint arXiv:2210.03043*, 2022.
- [39] Huy Ha and Shuran Song. Semantic abstraction: Open-world 3d scene understanding from 2d vision-language models. In *Conference on Robot Learning*, 2022.
- [40] 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. *arXiv preprint arXiv:2301.04926*, 2023. 3
- [41] Iro Armeni, Ozan Sener, Amir Roshan Zamir, Helen Jiang, Ioannis K. Brilakis, Martin Fischer, and Silvio Savarese. 3d semantic parsing of large-scale indoor spaces. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1534–1543, 2016. 3
- [42] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European Conference on Computer Vision*, 2022. 3, 7, 15
- [43] Rahul Goel, Dhawal Sirikonda, Saurabh Saini, and PJ Narayanan. Interactive segmentation of radiance fields. *arXiv preprint arXiv:2212.13545*, 2022. 3
- [44] Vadim Tschernezki, Iro Laina, Diane Larlus, and Andrea Vedaldi. Neural feature fusion fields: 3d distillation of self-supervised 2d image representations. *2022 International Conference on 3D Vision (3DV)*, pages 443–453, 2022. 3
- [45] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J. Davison. In-place scene labelling and understanding with implicit scene representation. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 15818–15827, 2021. 3, 4, 7
- [46] Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Buló, Norman Müller, Matthias Nießner, Angela Dai, and Peter Kotschieder. Panoptic lifting for 3d scene understanding with neural fields. *arXiv preprint arXiv:2212.09802*, 2022. 3
- [47] Vadim Tschernezki, Diane Larlus, and Andrea Vedaldi. Neurdiff: Segmenting 3d objects that move in egocentric videos. *2021 International Conference on 3D Vision (3DV)*, pages 910–919, 2021. 3
- [48] Jesus Zarzar, Sara Rojas, Silvio Giancola, and Bernard Ghanem. Segnerf: 3d part segmentation with neural radiance fields. *arXiv preprint arXiv:2211.11215*, 2022. 3
- [49] Zhiwen Fan, Peihao Wang, Yifan Jiang, Xinyu Gong, DeJia Xu, and Zhangyang Wang. Nerf-sos: Any-view self-supervised object segmentation on complex scenes. *arXiv preprint arXiv:2209.08776*, 2022. 3, 4, 6
- [50] Shengnan Liang, Yichen Liu, Shangzhe Wu, Yu-Wing Tai, and Chi-Keung Tang. Onerf: Unsupervised 3d object segmentation from multiple views. *arXiv preprint arXiv:2211.12038*, 2022. 3

- [51] Karl Stelzner, Kristian Kersting, and Adam R Kosiorek. Decomposing 3d scenes into objects via unsupervised volume segmentation. *arXiv preprint arXiv:2104.01148*, 2021. 3
- [52] Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021. 3
- [53] Ce Zhou, Qian Li, Chen Li, Jun Yu, Yixin Liu, Guangjing Wang, Kai Zhang, Cheng Ji, Qiben Yan, Lifang He, et al. A comprehensive survey on pretrained foundation models: A history from bert to chatgpt. *arXiv preprint arXiv:2302.09419*, 2023. 3
- [54] Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI blog*, 2019. 3
- [55] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [56] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [57] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. 3
- [58] Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S Weld, Luke Zettlemoyer, and Omer Levy. Spanbert: Improving pre-training by representing and predicting spans. *Transactions of the Association for Computational Linguistics*, 8:64–77, 2020. 3
- [59] Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*, 2019. 3
- [60] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*, 2020. 3
- [61] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019. 3
- [62] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 3
- [63] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in Neural Information Processing Systems*, 35:36479–36494, 2022. 3
- [64] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 3
- [65] Zhaoqing Wang, Yu Lu, Qiang Li, Xunqiang Tao, Yandong Guo, Mingming Gong, and Tongliang Liu. Cris: Clip-driven referring image segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11686–11695, 2022. 3
- [66] Haoyu Song, Li Dong, Weinan Zhang, Ting Liu, and Furu Wei. Clip models are few-shot learners: Empirical studies on vqa and visual entailment. In *Annual Meeting of the Association for Computational Linguistics*, 2022. 3
- [67] Narek Tumanyan, Omer Bar-Tal, Shai Bagon, and Tali Dekel. Splicing vit features for semantic appearance transfer. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10738–10747, 2022. 3
- [68] Mark Hamilton, Zhoutong Zhang, Bharath Hariharan, Noah Snaveley, and William T Freeman. Unsupervised semantic segmentation by distilling feature correspondences. *arXiv preprint arXiv:2203.08414*, 2022. 4, 6, 9
- [69] Shir Amir, Yossi Gandelsman, Shai Bagon, and Tali Dekel. Deep vit features as dense visual descriptors. *arXiv preprint arXiv:2112.05814*, 2(3):4, 2021. 3
- [70] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019. 7, 10, 17
- [71] Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of annotated 3d objects. *arXiv preprint arXiv:2212.08051*, 2022. 7

- [72] Tong Wu, Jiarui Zhang, Xiao Fu, Yuxin Wang, Jiawei Ren, Liang Pan, Wayne Wu, Lei Yang, Jiaqi Wang, Chen Qian, et al. Omniobject3d: Large-vocabulary 3d object dataset for realistic perception, reconstruction and generation. *arXiv preprint arXiv:2301.07525*, 2023.
- [73] Jeremy Reizenstein, Roman Shapovalov, Philipp Henzler, Luca Sbordone, Patrick Labatut, and David Novotny. Common objects in 3d: Large-scale learning and evaluation of real-life 3d category reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10901–10911, 2021. [7](#)
- [74] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021. [7](#), [10](#), [17](#)
- [75] Johannes Lutz Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [16](#)
- [76] Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 2019. [16](#)

Appendix

This document provides supplementary materials for *Weakly Supervised 3D Open-vocabulary Segmentation* in implementation details (Appendix A), more ablations (Appendix B), more evaluations (Appendix C), and more results (Appendix D).

A Implementation Details

A.1 Model Architecture

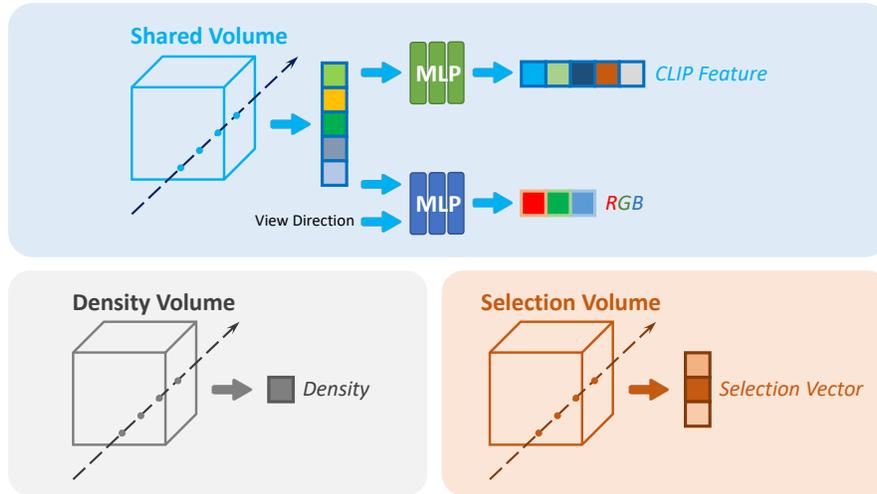


Figure 6: **Model architecture.**

We use TensorRF [42] as our base NeRF architecture for efficiency, the plane size is also the same as the default setting of TensorRF. The RGB and CLIP feature branches share the same volume and use the same intermediate features. The selection volume and density volume are two other independent volumes. We directly use the features extracted from the selection volume and density volume as the selection vector and the density value, as they have low dimensions and are view-independent. We use the original MLP architecture in TensorRF to extract the view-dependent RGB value and use another MLP which discards view direction input to extract the rendered CLIP feature. The architecture is illustrated in Fig. 6.

A.2 Hyperparameters

We set $\tau = 0.2$ to get the shaper segmentation probability distribution \hat{P} . The offset b is set to 0.7 to measure the similarities of the DINO features, meaning that two DINO features are considered similar if their cosine similarity is larger than 0.7, and different if less than 0.7. We use 3 scales of CLIP features, and the patch sizes of each scale are set as $s/5$, $s/7$, and $s/10$, where s is the smaller value in the width and height of the input image I . In the ablation studies, we use $s/7$ as the patch size of the single-scale CLIP feature input. The weights associated with similar and dissimilar DINO features in \mathcal{L}_{FDA} are set as $\lambda_{pos} = 200$ and $\lambda_{neg} = 0.2$ by default. In certain scenes, we find that setting λ_{neg} to 0.22 or 0.18 can produce better results. We use ViT-B/16 CLIP model to extract the image and text features and use version 1 dino_vitb8 model to extract the DINO features because it employs the smallest downsampling factor of 8 which is advantageous for high-precision segmentation.

A.3 Training

To reconstruct a NeRF from multiview images of a scene, we follow the same training settings as TensorRF. For segmentation training, we train the model for 15k iterations. In the first 5k iterations,

Table 3: **Dataset.** We list the collected 10 scenes and the corresponding text labels. The background labels are in *Italic font*.

Scene	Text Labels
bed	red bag, black leather shoe, banana, hand, camera, <i>white sheet</i>
sofa	a stack of UNO cards, a red Nintendo Switch joy-con controller, Pikachu, Gundam, Xbox wireless controller, <i>grey sofa</i>
lawn	red apple, New York Yankees cap, stapler, black headphone, hand soap, <i>green lawn</i>
room	shrilling chicken, weaving basket, rabbit, dinosaur, baseball, <i>wood wall</i>
bench	Portuguese egg tart, orange cat, green grape, mini offroad car, dressing doll, <i>pebbled concrete wall, wood</i>
table	a wooden ukulele, a beige mug, a GPU card with fans, a black Nike shoe, a Hatsune Miku statue, <i>lime wall</i>
office desk	the book of The Unbearable Lightness of Being, a can of red bull drink, a white keyboard, a pack of pocket tissues, <i>desktop, blue partition</i>
blue sofa	a bottle of perfume, sunglasses, a squirrel pig doll, a JBL bluetooth speaker, an aircon controller, <i>blue-grey sofa</i>
snacks	Coke Cola, orange juice drink, calculator, pitaya, Glico Pocky chocolate, biscuits sticks box, <i>desktop</i>
covered desk	Winnie-the-Pooh, Dove body wash, gerbera, electric shaver, canned chili sauce, <i>desktop</i>

we freeze the shared volume and density volume, and train the selection volume and the CLIP feature branch. For the rest 10k iterations, we further finetune the shared volume and the RGB branch. We use Adam optimizer with $\text{betas} = (0.9, 0.99)$. The learning rates for training the volume and MLP branch are respectively set to 0.02 and $1e - 4$. For finetuning the volume and the MLP, the learning rates are set to $5e - 3$ and $5e - 5$. We also employ a learning rate decay with a factor of 0.1.

The multi-scale pixel-level CLIP features of training views are pre-computed before training and the DINO features are computed with sampled patches on the fly during training. When computing $\mathcal{L}_{\text{supervision}}$ and \mathcal{L}_{RDA} , we randomly sample rays with a batch size of 4096. When computing \mathcal{L}_{FDA} we randomly sample patches of size 256×256 with a batch size of 8. We use a downsampling factor of 8 when sampling rays and a factor of 5 when sampling patches. The model is trained on an NVIDIA A5000 GPU with 24G memory for $\sim 1\text{h}30\text{min}$ for each scene.

A.4 Dataset

We capture 10 scenes using smartphones and use Colmap [75] to extract camera parameters for each image. We capture 20 \sim 30 images for each scene and the resolution of each image is 4032×3024 . We follow the data structure of LLFF [76]. We manually annotate the segmentation maps of 5 views for each scene as the ground truth for evaluation.

We list the text labels used in our experiments in Tab. 3. Note that we sometimes choose to add a general word to describe the whole background, such as *wall, desktop*, following LERF [2]. The text labels in our dataset contain many long-tail classes, which can be used to fully assess the open-vocabulary capability.

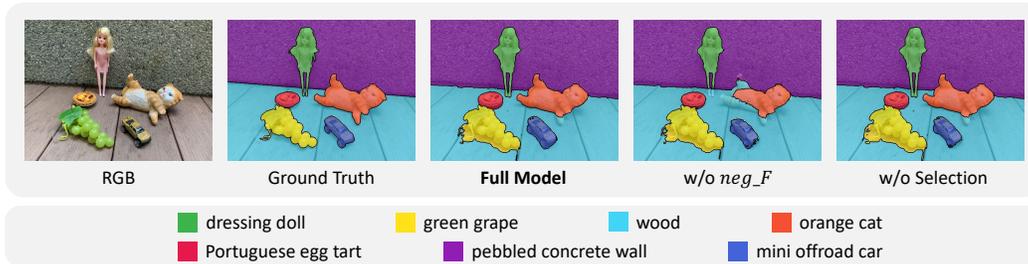


Figure 7: **More ablations.**

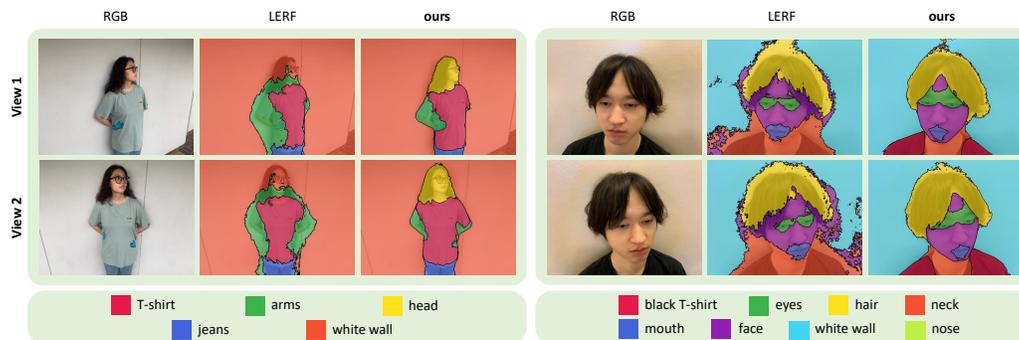


Figure 8: Evaluation on human body dataset (left). Evaluation on human head dataset (right).

B More Ablations

We perform two more ablation studies on the Selection Volume and the FDA loss, as shown in Tab. 4 and Fig. 7. Without the Selection Volume, we simply average the multi-scale CLIP features rather than learning the appropriate scale. We can see that both the mIoU score and the Accuracy score are inferior to the full model. We could discard the dissimilar part neg_F since dissimilar DINO features often impair the stability of the correlation loss. However, neg_F encourages different segmentation probabilities for different semantic regions and it plays a crucial role for precise object boundary extraction.

Table 4: **More Ablations.**

	mIoU	Accuracy
w/o neg_F	76.9	92.4
w/o Selection	84.8	95.3
full model	86.2	95.8

C More Evaluations

We additionally perform evaluations on human body, human head, indoor datasets with low-quality images [70, 74], and a complex scene from LERF datasets [2]. We compare with the concurrent work LERF qualitatively due to the lack of labels or the defective annotations as pointed out in [4]. We also perform experiments with different text prompts. We use the same scale level number and patch sizes in all comparisons.

Human body and head. As shown in Fig. 8, our method segments more precise parts than LERF. Specifically, LERF fails to segment the "head" in the human body and the "black T-shirt" in the human head. In contrast, our method can recognize and segment these parts correctly because our designed RDA loss addresses the ambiguity of the CLIP features effectively.

Indoor scenes with low-quality images. Fig. 9 shows experiments on the indoor datasets [70, 74], where many images are unrealistically rendered with less photorealistic appearances (as indicated in [4]) and have limited spatial resolution (640×480 or 1024×768). Due to these data constraints,

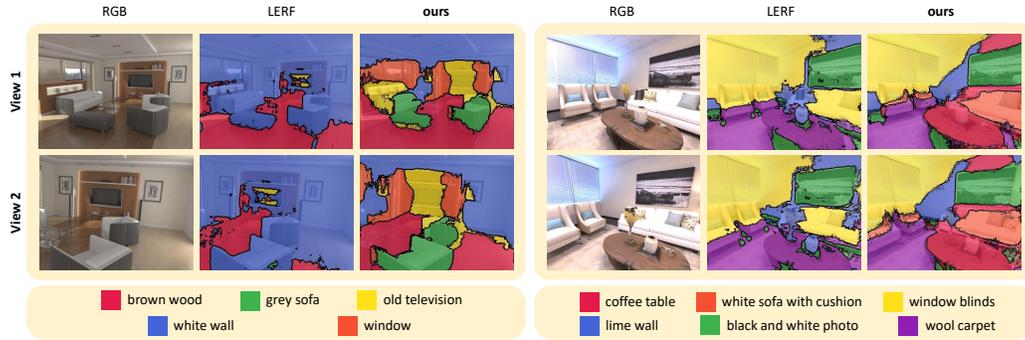


Figure 9: Evaluation on indoor datasets with lower quality images.

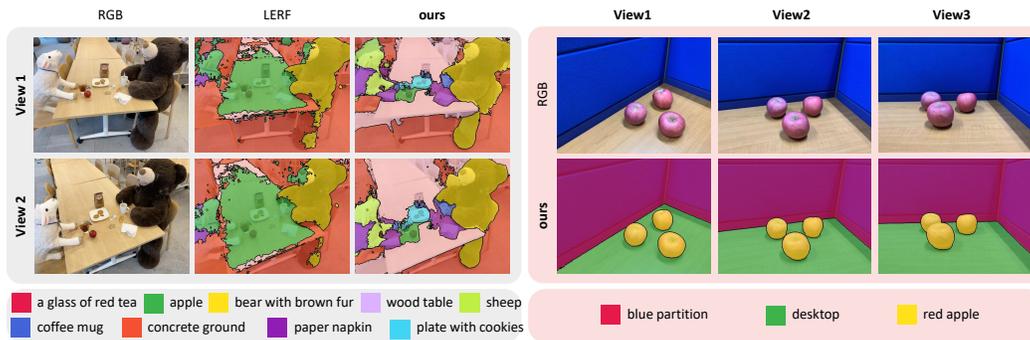


Figure 10: Evaluation on a complex scene from LERF datasets (left). Evaluation on a scene with multiple instances of a same class (right).

our method sometimes confuses labels with similar appearances. However, we can see that our method still outperforms LERF by successfully segmenting more labels.

Complex scenes. Fig. 10 (left) shows the segmentation of one challenging sample from the LERF dataset, where the scene has complex geometry as well as many objects of varying sizes. It can be observed that LERF cannot segment most objects due to the ambiguities of CLIP features while our method can segment more objects correctly with more precise boundaries. Fig. 10 (right) shows a scene with multiple instances of a same class. Since instances of the same class often share similar appearance, texture, etc., they also have similar DINO features. As a result, FDA will not mistakenly segment them into different classes. The RDA loss will further help by assigning all these instances to the same text label. In the experiment, we observed that our method successfully segments all three apples into the same class with accurate boundaries.

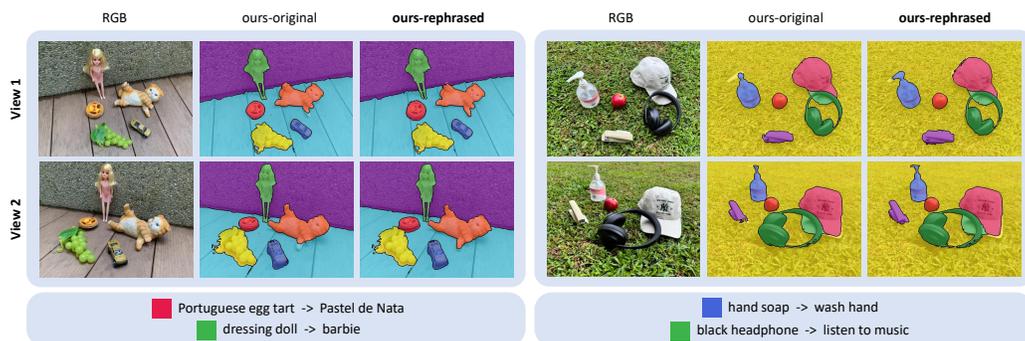


Figure 11: Evaluation on rephrased texts.

Segmentation with different text prompts. We conduct experiments to segment scenes with different text prompts. In the experiments, we replaced the original texts with different languages (e.g., Portuguese egg tart -> Pastel de Nata), names (e.g., dressing doll -> Barbie), and actions (e.g., hand soap -> wash hand, black headphone -> listen to music). As Fig. 11 shows, with the rephrased text prompts, our method can still segment the scenes reliably. The experiments are well aligned with the quantitative experiments as shown in Tab. 5.

Table 5: Evaluation on rephrased texts.

	mIoU	Accuracy	mIoU	Accuracy
original	88.2	97.3	89.3	96.3
rephrased	89.3	97.2	88.4	96.6

D More Results

We show more segmentation visualizations of our method in Fig. 12 (bed), Fig. 13 (sofa), Fig. 14 (lawn), Fig. 15 (room), Fig. 16 (bench), Fig. 17 (table), Fig. 18 (office desk), Fig. 19 (blue sofa), Fig. 20 (snacks), and Fig. 21 (covered desk). The quantitative results are listed in Tab. 6.

Table 6: **Quantitative results.**

<i>bed</i>		<i>sofa</i>		<i>lawn</i>		<i>room</i>		<i>bench</i>	
mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy
89.5	96.7	74.0	91.6	88.2	97.3	92.8	98.9	89.3	96.3
<i>table</i>		<i>office desk</i>		<i>blue sofa</i>		<i>snacks</i>		<i>covered desk</i>	
mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy	mIoU	Accuracy
88.8	96.5	91.7	96.2	82.8	97.7	95.8	99.1	88.6	97.2

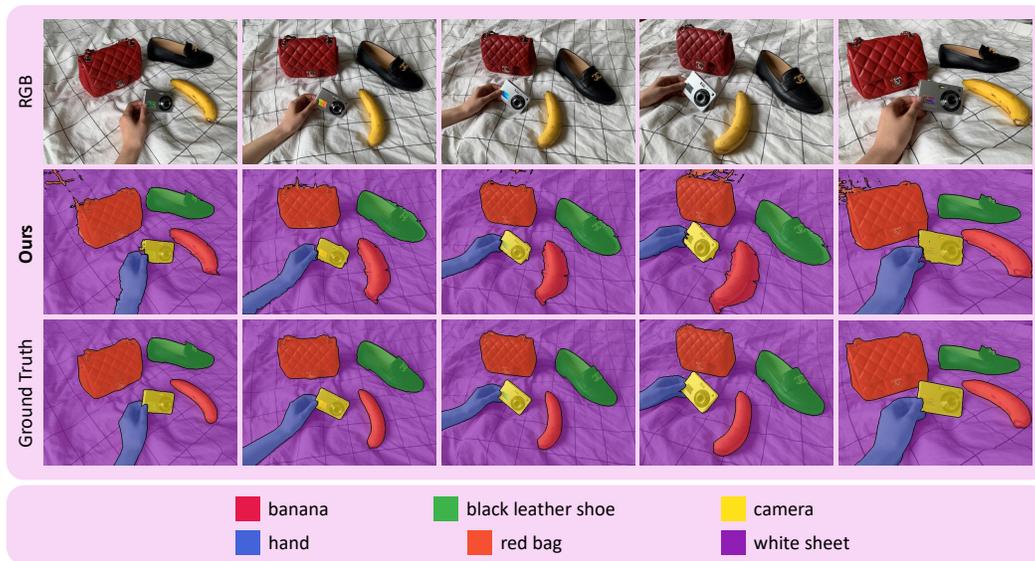


Figure 12: bed.

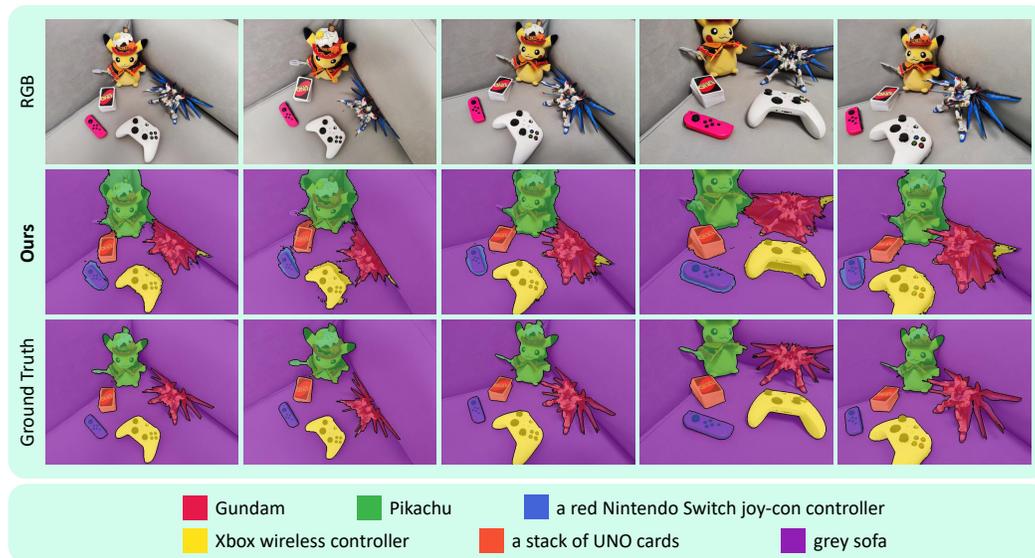


Figure 13: sofa.

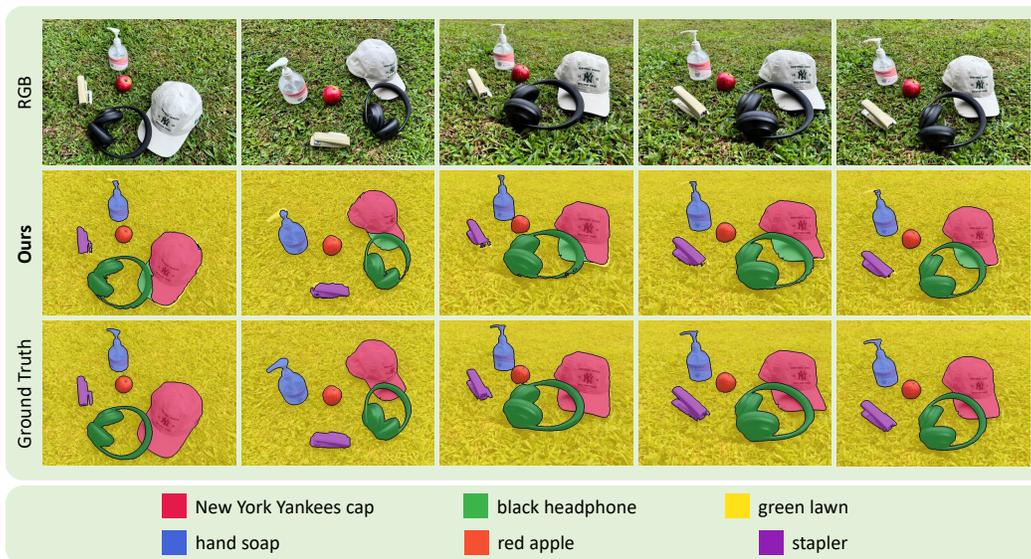


Figure 14: lawn.

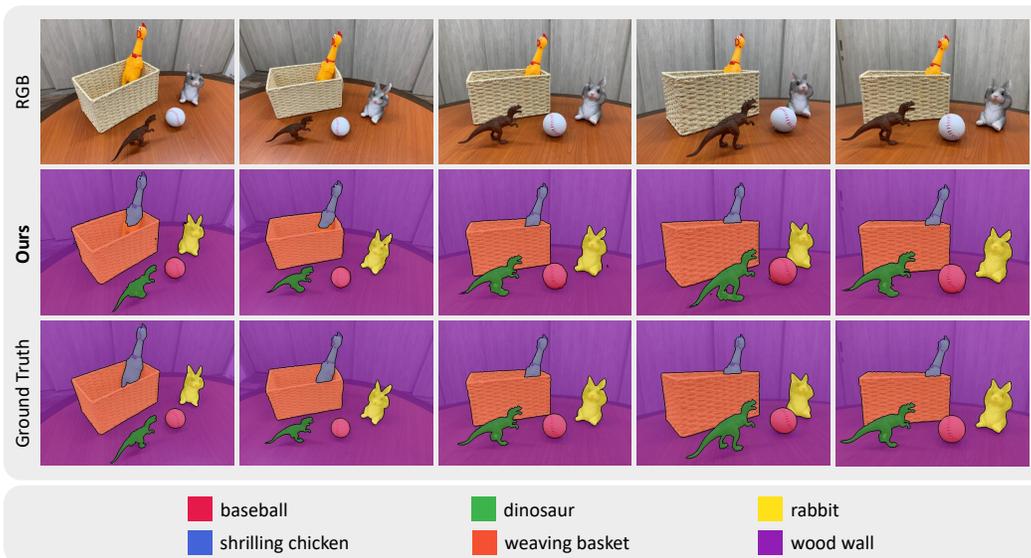


Figure 15: room.

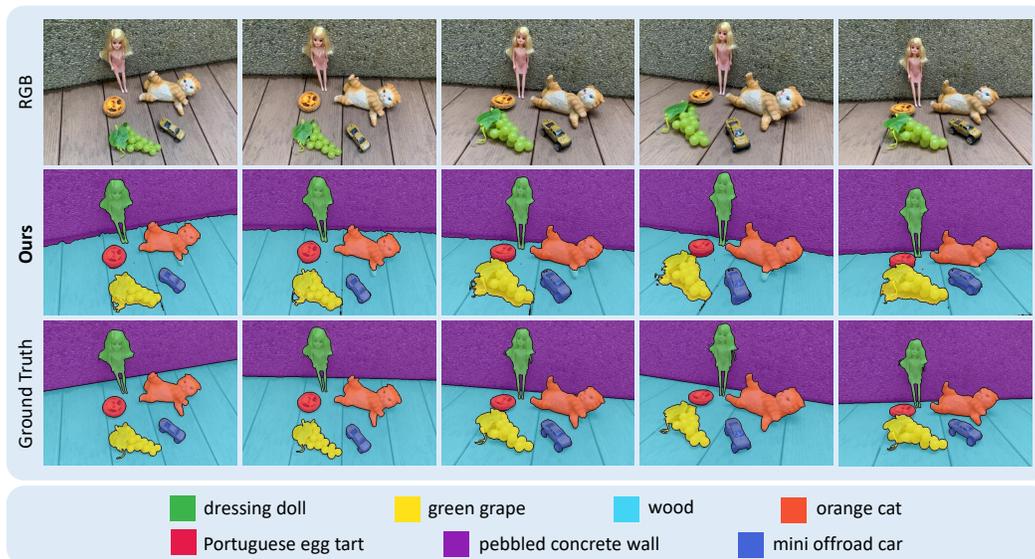


Figure 16: **bench.**

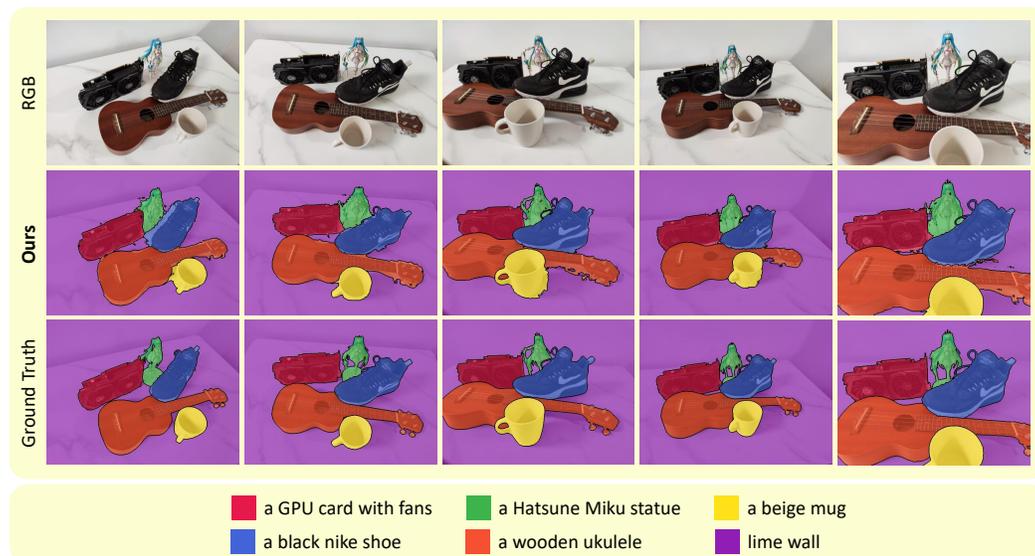


Figure 17: **table.**

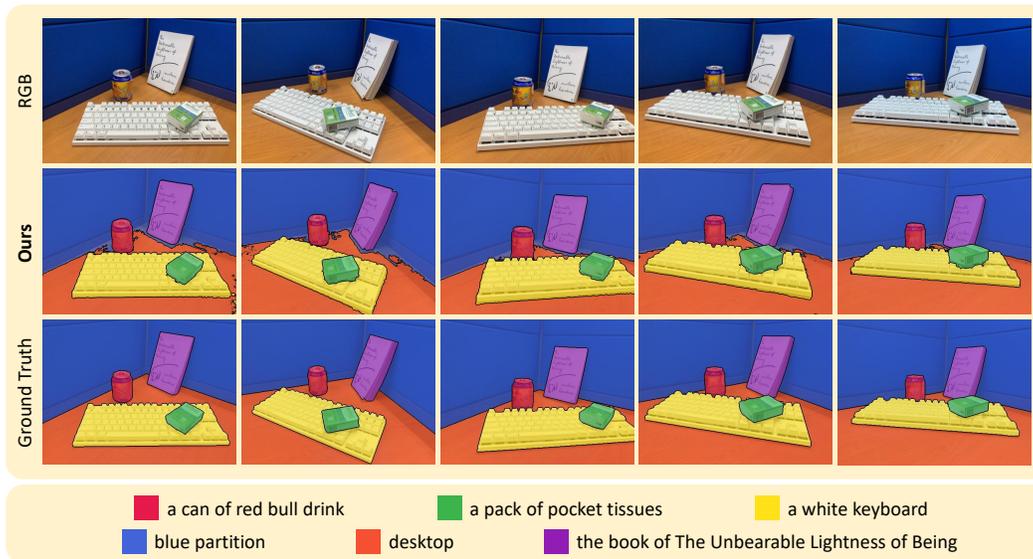


Figure 18: office desk.

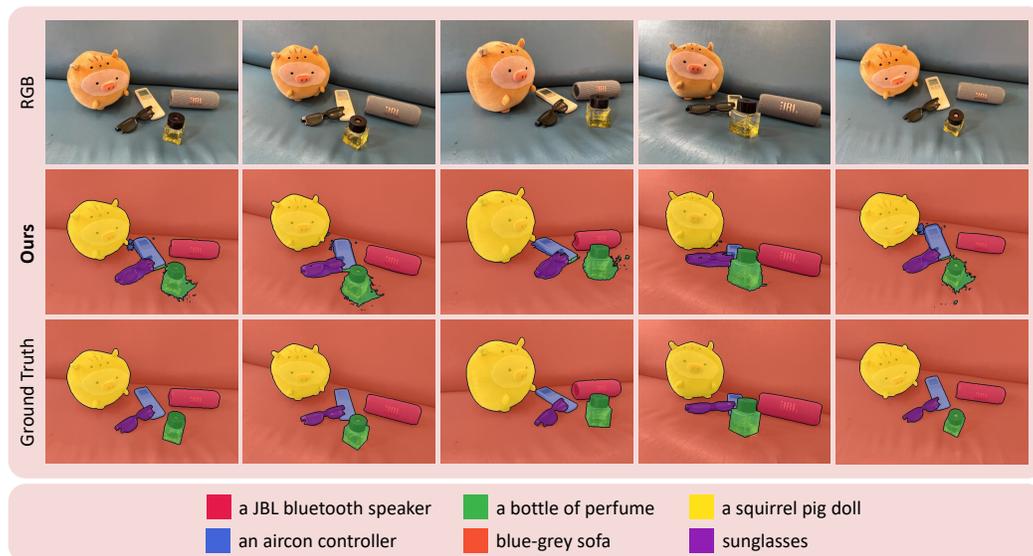


Figure 19: blue sofa.

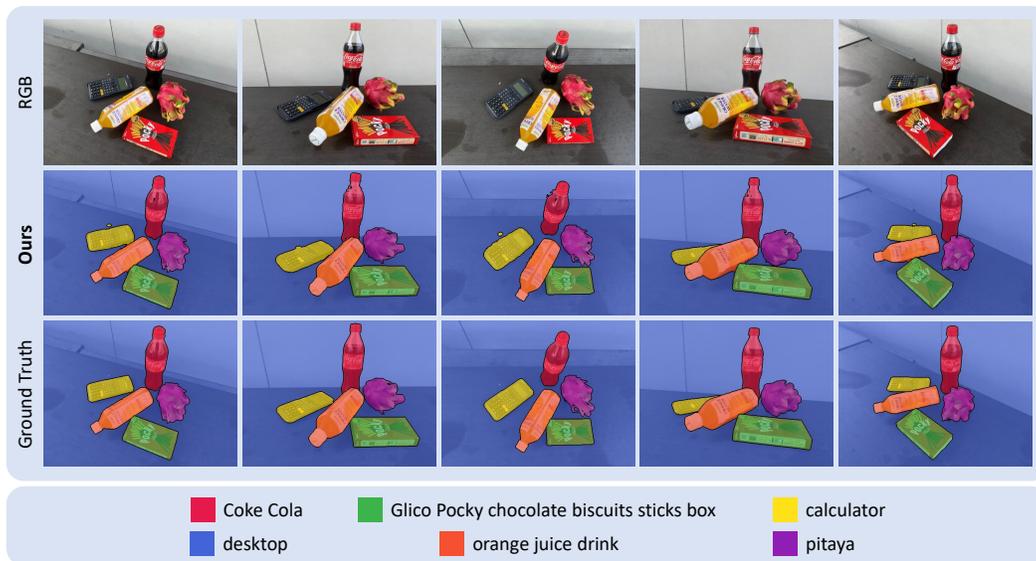


Figure 20: snacks.

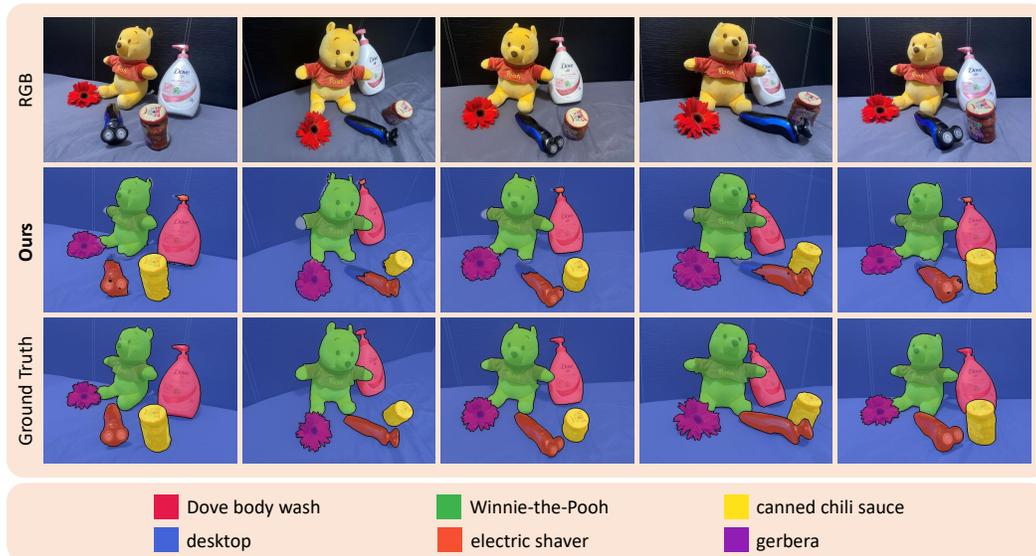


Figure 21: covered desk.