

# WildSeg3D: Segment Any 3D Objects in the Wild from 2D Images

Yansong Guo<sup>1</sup>, Jie Hu<sup>2</sup>, Yansong Qu<sup>1</sup>, Liujuan Cao<sup>1†</sup>

<sup>1</sup>Key Laboratory of Multimedia Trusted Perception and Efficient Computing,  
Ministry of Education of China, Xiamen University

<sup>2</sup>National University of Singapore, Singapore

## Abstract

*Recent advances in intuitive 3D segmentation from 2D images have demonstrated impressive performance. However, current models typically require extensive scene-specific training to accurately reconstruct and segment objects, which limits their applicability in real-time scenarios. In this paper, we introduce WildSeg3D, an efficient approach that enables the segmentation of arbitrary 3D objects across diverse environments using a feed-forward mechanism. A key challenge of this feed-forward approach lies in the accumulation of 3D alignment errors across multiple 2D views, which can lead to inaccurate 3D segmentation results. To address this issue, we propose Dynamic Global Aligning (DGA), a technique that improves the accuracy of global multi-view alignment by focusing on difficult-to-match 3D points across images, using a dynamic adjustment function. Additionally, for real-time intuitive segmentation, we introduce Multi-view Group Mapping (MGM), a method that utilizes an object mask cache to integrate multi-view segmentations and respond rapidly to user prompts. WildSeg3D demonstrates robust generalization across arbitrary scenes, thereby eliminating the need for scene-specific training. Specifically, WildSeg3D not only attains the accuracy of state-of-the-art (SOTA) methods but also achieves a 40× speedup compared to existing SOTA models. Code will be released at <https://github.com/Ethan16162/WildSeg3D>.*

## 1. Introduction

Intuitive 3D segmentation from 2D images plays a critical role in 3D scene understanding and remains a fundamental challenge in computer vision, attracting significant attention from the research community [11, 18, 21, 38, 47–49]. This technology has broad applications across various fields, including virtual and augmented reality, real-time intuitive systems, and automatic labeling. Recent advancements in intuitive 3D segmentation have demonstrated remarkable

performance, particularly based on Neural Radiance Fields (NeRF) [33] and 3D Gaussian Splatting (3DGS) [17]. For instance, models such as SA3D [4] and SANeRF-HQ [27] integrate NeRF with foundational segmentation models like Segment Anything Model (SAM) [19], aligning semantic information with 3D representations to enable effective 3D object segmentation. Similarly, 3DGS-based approaches [5, 15, 40, 59, 67] address the high training demands of NeRF by constructing Gaussian feature fields in combination with SAM, facilitating faster model training. However, both NeRF-based and 3DGS-based methods typically rely on extensive scene-specific training to obtain accurate 3D priors, which significantly hinders their applicability in real-time scenarios.

To overcome this limitation, we propose WildSeg3D, a novel approach with a feed-forward manner, eliminating the need for scene-specific training. Inspired by Dust3r [51] and Mast3r [24], our feed-forward approach represents 3D scenes as pointmaps and performs scene reconstruction via global alignment. A key challenge in this process is the influence of redundant background points and the difficulty of matching 3D points across different views, which accumulates 3D alignment errors across multiple 2D views. This can lead to inaccurate 3D segmentation with confusion between background and target objects. To address this, we introduce Dynamic Global Alignment (DGA), a method that dynamically adjusts attention on view-specific points during matching, minimizing alignment errors during the global registration of pointmaps across multiple views. For real-time intuitive segmentation, we also propose a mask cache constructed during preprocessing, leveraging the multi-frame segmentation capabilities of SAM2 [43]. This cache stores consistent object masks across all view-points, providing offline data that can be used for subsequent real-time segmentation. We further introduce Multi-view Group Mapping (MGM), which combines the mask cache with the DGA strategy to integrate multi-view segmentation results into an aligned global coordinate system. The MGM module retrieves target feature points from the mask cache based on user inputs and generates a unified 3D

<sup>†</sup>Corresponding Author.

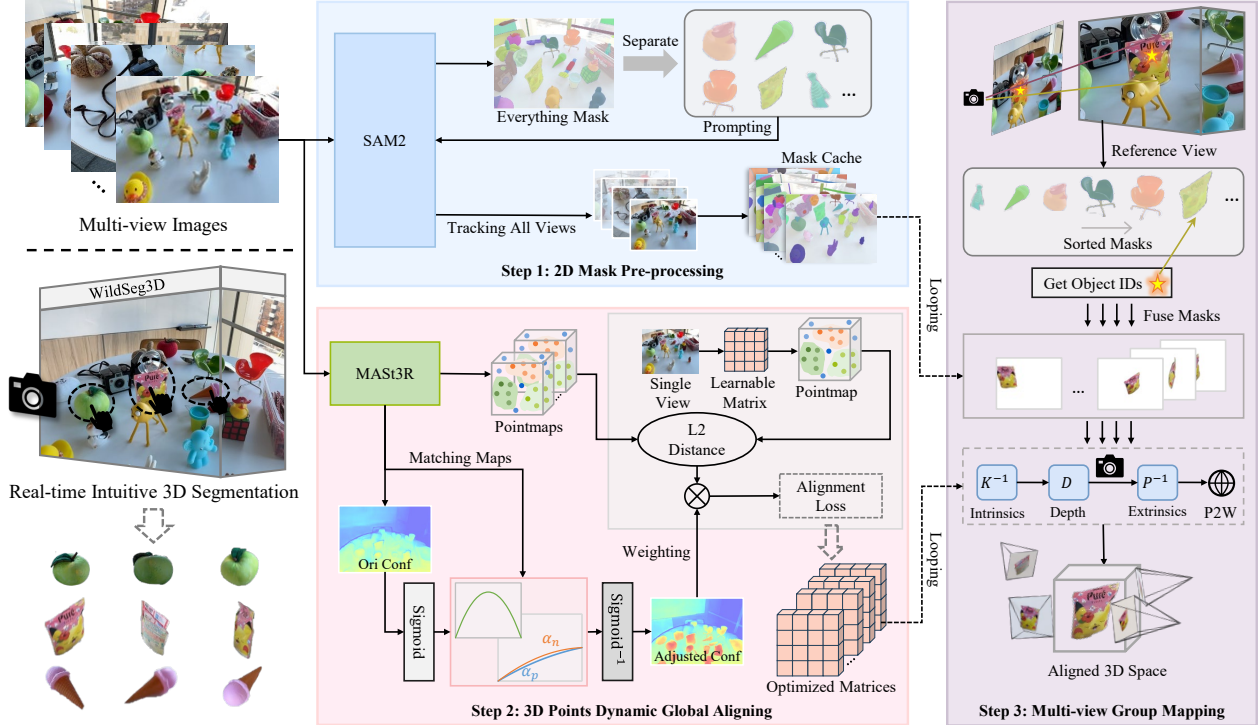


Figure 1. **Framework of WildSeg3D.** WildSeg3D operates in three stages. First, during the pre-processing stage, 2D feature masks are constructed from multi-view images using SAM2, providing offline support for intuitive segmentation. Second, in the Dynamic Global Alignment (DGA) stage, a dynamic weight adjustment strategy is applied to achieve global alignment of the pointmaps generated by MAST3r, thereby improving the accuracy of object reconstruction. Finally, in the Multi-view Group Mapping (MGM) stage, multi-view masks for the target are retrieved from the mask cache based on user input, and these masks are transformed into an aligned 3D space in real time, where “P2W” refers to the transformation from pixel coordinates to the aligned world coordinates.

mask for the target across all viewpoints.

WildSeg3D demonstrates robust generalization across diverse scenes, transforming 2D target masks into an aligned 3D space, thus enabling accurate 3D segmentation without the need for scene-specific training. We conduct extensive experiments on WildSeg3D using complex real-world scenes to evaluate both segmentation accuracy and efficiency. Our results show that WildSeg3D not only matches the accuracy of state-of-the-art (SOTA) methods [4, 5, 11, 35, 44, 45, 48, 61] but also achieves a 40× speedup compared to the SOTA models. Specifically, WildSeg3D completes scene reconstruction in under 30 seconds, a significant reduction compared to the fastest model, SA3D, which requires 780 seconds. Additionally, real-time interaction results are delivered in just 5-20 milliseconds.

Our main contributions can be summarized as follows:

- We introduce WildSeg3D, the first feed-forward 3D segmentation model that operates directly from 2D views, eliminating the need for scene-specific training and enabling immediate segmentation of arbitrary 3D objects in diverse environments.
- we propose Dynamic Global Alignment (DGA), a novel method for addressing 3D alignment errors and enhanc-

ing segmentation accuracy. We also propose Multi-view Group Mapping (MGM) that enables real-time intuitive 3D segmentation with robust generalization across diverse scenes.

- Extensive experiments on complex real-world scenes demonstrate that WildSeg3D achieves both high segmentation accuracy and significant efficiency, processing 2D views to 3D segmentation over 40× faster than existing SOTA models.

## 2. Related Work

### 2.1. 2D Image-based 3D Reconstruction

Recent advancements in image-based 3D reconstruction using neural networks have led to significant progress. Innovations like Neural Radiance Fields (NeRF) [33] have shown strong performance in generating realistic novel viewpoints for view synthesis. However, their reliance on neural networks leads to long training and rendering times. To improve surface reconstruction, [10, 13, 28, 50, 55] leverage the signed distance function (SDF) for surface representation and introduce an innovative volume rendering technique to learn an SDF model. Kerbl *et al.* intro-

duced 3D Gaussian Splatting (3DGS) [17], which provides an explicit representation of 3D scene information, bypassing the time-consuming implicit reconstruction process of NeRF via MLP. Some studies [16, 18, 40, 41, 59, 67], incorporate semantic information into 3DGS, equipping it with semantic awareness through training. Other researches [30, 52, 56–58] extend 3DGS by incorporating deformation fields to track the positions of 3D Gaussians at each timestamp, capturing dynamic 3D environments. DUST3R [51] presents a novel approach for dense, unconstrained 3D reconstruction from arbitrary image collections without requiring camera calibration or viewpoint poses, unlike 3DGS and NeRF, which depend on dense viewpoints for scene construction. Building on DUST3R, MAST3R [24] reframes image matching as a 3D task, achieving superior 3D reconstruction performance. We extended MAST3R’s feed-forward mechanism to enable 3D scene perception and address alignment errors.

## 2.2. 2D Foundation Models

Foundation models (FMs) have emerged as a transformative paradigm in AI. These models are typically trained on extensive datasets, possess a large number of parameters, and demonstrate adaptability across a broad spectrum of downstream tasks. Specifically, 2D visual foundation models (VFMs) [12, 19, 25, 37, 42, 43, 54] have gained significant attention due to their ability to process and understand visual data. Kirillov *et al.* proposed the Segment Anything Model (SAM) [19], a 2D segmentation foundation model for prompt-based segmentation. SAM generates segmentation masks based on prompts that identify target objects in an image, allowing it to generalize across unseen categories. As the successor to SAM, Segment Anything 2 (SAM2) [43] unifies video and image segmentation by utilizing a larger training dataset and incorporating architectural enhancements to improve performance across a wide range of tasks. Our method leverages SAM2 to segment 2D images, maintaining consistent masks across multiple viewpoints, thereby addressing the challenges of real-time 3D intuitive segmentation.

## 2.3. Intuitive 3D Segmentation

Few approaches [11, 16, 22, 23, 39, 63] support intuitive segmentation directly in 3D space. For example, [22] directly segment 3D point clouds based on user clicks. Meanwhile, representations in 2D image-based 3D reconstruction have driven progress in intuitive 3D segmentation from 2D images. Inspired by the advancements in 3D Neural Scene Representation [17, 33], several studies [2, 4, 7, 26, 36, 46, 62, 64] have explored 3D segmentation within these frameworks. With SAM [19], SA3D [4] proposed an automatic strategy for cross-view prompt collection, leveraging SAM to obtain 2D masks and guide 3D

feature training. Within the context of 3DGS, Feature3DGS [67] converts features from SAM’s encoder into 3D space and uses SAM’s decoder to generate masks. Mask-lifting-based methods directly map 2D segmentation masks from SAM into 3D space. Notable examples include SAGA [5], Gaussian Grouping [59], SAGS [15], Click-Gaussian [8], and FlashSplat [45]. These approaches computational costs of feature-to-mask conversion but still suffer from slow reconstructing. Our method distinguishes itself by achieving feed-forward segmentation, without scene-specific training.

## 3. Methods

### 3.1. Preliminary: Feed-Forward Mechanism

The feed-forward mechanism for 3D scene reconstruction was initially proposed in DUST3R [51] and further refined in MAST3R [24]. Unlike methods such as NeRF [33] and 3DGS [17], which require scene-specific pre-training, the feed-forward approach enables general 3D scene reconstruction through two key steps: pointmap prediction and global alignment.

**Pointmap Prediction.** The process of pointmap prediction can be described as a network function  $\mathcal{F} : (I^n, I^m) \rightarrow (X^{n,e}, C^{n,e}, F^{n,e}, X^{m,e}, C^{m,e}, F^{m,e})$ , the inputs are two RGB images  $I^n, I^m \in \mathbb{R}^{W \times H \times 3}$  from different views of the scene, the outputs include two corresponding pointmaps,  $X^{n,e}, X^{m,e} \in \mathbb{R}^{W \times H \times 3}$ , confidence maps  $C^{n,e}, C^{m,e} \in \mathbb{R}^{W \times H}$ , and dense local features  $F^{n,e}, F^{m,e} \in \mathbb{R}^{W \times H \times d}$ . Note that  $e = (n, m)$  refers to the image pair formed by  $I^n$  and  $I^m$ , and both pointmaps are positioned in the camera coordinate system of  $I^n$ . The predicted pointmaps locate the 3D positions for every pixel of the input 2D images.

**Global Alignment.** Global alignment is used as a post-process that optimizes the pointmaps from multiple views into an aligned 3D coordinate system. Given a set of images  $\{I^1, I^2, \dots, I^N\}$  from a scene, a connectivity graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  is constructed, where the vertices  $\mathcal{V}$  represent the  $N$  images, and each edge  $e = (n, m) \in \mathcal{E}$  connects an image pair  $I^n$  and  $I^m$ . By traversing the connected graph  $\mathcal{G}$ , globally aligned pointmaps  $\{\chi^n \in \mathbb{R}^{W \times H \times 3}\}$  are recovered for all pixel coordinates  $(i, j) \in \{1 \dots W\} \times \{1 \dots H\}$  and all cameras for different views  $n = 1, \dots, N$ . The global optimization process is formulated as follows:

$$\chi^* = \arg \min_{\chi, P, \sigma} \sum_{e \in \mathcal{E}} \sum_{v \in e} \sum_{i=1}^{HW} C_i^{v,e} \|\chi_i^v - \sigma_e P_e X_i^{v,e}\|, \quad (1)$$

where  $P_e \in \mathbb{R}^{3 \times 4}$  represents the pairwise pose, which is a rigid transformation used to align the pointmaps  $X^{n,e}, X^{m,e}$  with the world-coordinate pointmaps  $\chi^n, \chi^m$ . Additionally,  $\sigma_e$  is a scale factor, subject to the constraint that  $\prod_e \sigma_e = 1$  for all  $e \in \mathcal{E}$ .

### 3.2. Task Formulation: Segment with Pointmaps

Although the feed-forward approach offers the advantage of generalizing across various scenes without the need for scene-specific pre-training, its application to 3D segmentation often results in lower accuracy. This is primarily due to the accumulation of 3D alignment errors during the global alignment stage. Given multiple 2D images from different viewpoints and a user-provided prompt specifying a target object in 2D, our task is to optimize the alignment of 3D object segmentation across these views. The goal is to transform the 2D object masks, related to the provided prompt, into an aligned 3D space. Two key challenges need to be addressed in this task: (1) obtaining real-time segmentation masks of the target across all viewpoints based on user prompts, and (2) refining the alignment loss function to enhance segmentation accuracy, which requires a strategy to minimize discrepancies between predicted pointmaps from multiple views and globally aligned pointmaps.

### 3.3. WildSeg3D

As illustrated in Figure 1, WildSeg3D is a feed-forward framework designed for real-time intuitive 3D segmentation from 2D views. The framework operates in three stages: 2D mask pre-processing, 3D point dynamic global alignment, and multi-view group mapping. In the 2D mask pre-processing stage, segmentation masks are generated by SAM2 from the input multi-view images and stored in a mask cache for efficient access during real-time segmentation. In the 3D point dynamic global aligning stage, the proposed DGA refines the alignment of 3D pointmaps by focusing on challenging pixel correspondences across different views. This approach ensures more accurate 3D scene reconstruction by addressing misalignments caused by complex or occluded regions. In the multi-view group mapping stage, the stored masks are retrieved from the mask cache and the transform matrix learned through DGA is applied to map the multi-view pointmaps into an aligned 3D coordinate system in real time. This process allows for the accurate 3D segmentation results based on user prompts.

#### 3.3.1. 2D Mask Pre-processing

In the 2D mask pre-processing stage, we eliminate the need for online computation during 2D segmentation by introducing a mask cache, enabling rapid, accurate segmentation across multiple viewpoints. First, by leveraging the video tracking capabilities of SAM2, we perform panoptic segmentation on a single viewpoint, generating precise masks for each target object within the scene. These object masks are then stored offline in the mask cache, creating a repository of segmentation data that can be efficiently accessed during real-time processing. In subsequent stages, these pre-generated masks serve as accurate prompts for SAM2’s tracking functionality, allowing consistent tracking

and segmentation of each object across multiple viewpoints in later frames. This approach ensures continuity and coherence in the segmentation process, even as the viewpoint changes. By pre-generating and storing the masks offline, the mask cache reduces the computational burden at runtime. Furthermore, this offline caching mechanism not only accelerates the segmentation process but also enhances its robustness.

#### 3.3.2. 3D Point Dynamic Global Aligning

While the segmentation masks from the first step can be globally aligned to unify the 3D pointmaps into a single world coordinate system, challenges remain in achieving precise 3D segmentation across different viewpoints. Specifically, misalignment and loss of detail can occur due to cluttered and inconsistent backgrounds in images from various viewpoints, weakening the alignment of target objects during the global alignment process. To address these issues, we propose the Dynamic Global Aligning (DGA) method, which introduces soft masks to minimize background interference and dynamically adjusts aligning weights. This adjustment enhances the focus on challenging sample points, thus improving the overall alignment effectiveness for the target objects.

**Soft-mask and Confidence Aggregation.** As described in Eq. 1, the original global alignment approach considers all pixels from all viewpoints for alignment. However, due to significant background differences between images from different viewpoints, the alignment process suffers from two issues: (1) background features are difficult to be aligned, and (2) the large proportion of the background in the images weakens the alignment focus on the objects to be segmented. To mitigate this, we propose to soften the masks generated by SAM2 and store them in the mask cache. Concretely, begin by multiplying the confidence of the pointmaps,  $C \in \mathbb{R}^{W \times H}$ , with the corresponding soft masks,  $S \in \mathbb{R}^{W \times H}$ , to obtain a weighted confidence. A sigmoid function,  $\sigma(\cdot)$ , is then applied to the weighted confidence to map it to a confidence score in the range of 0 to 1. For all views  $v = 1, \dots, N$ , this process is expressed as follows:

$$F_i^{v,e} = \sigma(S_i^v \times C_i^{v,e}) = \frac{1}{1 + e^{-S_i^v \times C_i^{v,e}}}, \quad (2)$$

where  $F \in \mathbb{R}^{W \times H}$  represents the adjusted confidence after applying the soft mask.

**Transform Matrix.** We initialize a transform matrix for each view, composed of camera intrinsics, extrinsics, and depth information, enabling projection from pixel coordinates to world coordinates. Given a 2D point  $(x, y)$  from a single view, with depth map  $\mathbf{D} \in \mathbb{R}^{W \times H}$ , camera intrinsic matrix  $\mathbf{K} \in \mathbb{R}^{3 \times 3}$ , and extrinsic matrix  $\mathbf{P}$ , its corresponding 3D point  $\mathbf{p} = (x, y, z)^T$  in world coordinates can be

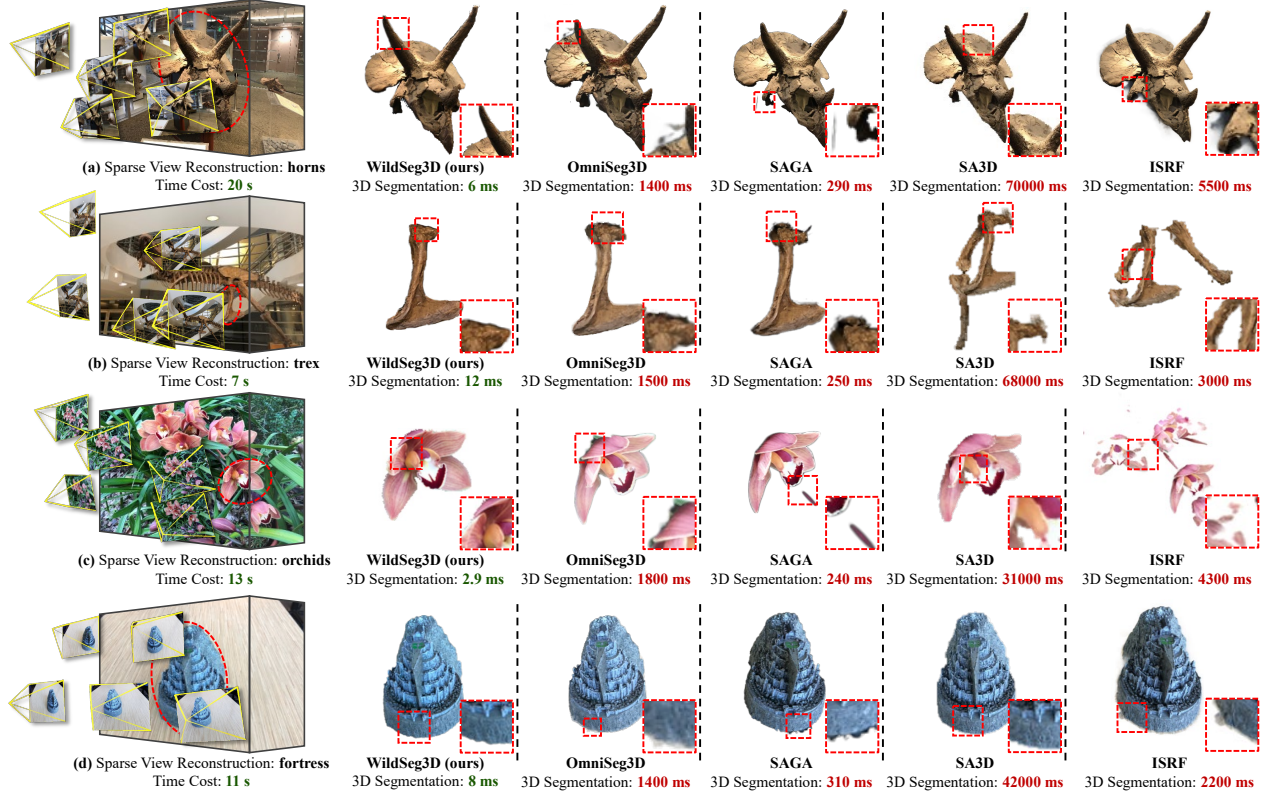


Figure 2. **Visualization on the NVOS dataset.** (a)-(d) show the sparse view reconstruction and timing results on horns, trex, orchids, and fortress scenes, including preprocessing and DGA-based scene reconstruction. Target objects for segmentation are marked with red dashed lines in the first column. From left to right: segmentation results and elapsed time from prompt input to segmentation acquisition across models.

estimated as:

$$\mathbf{p} = \mathbf{P}^{-1} \mathbf{K}^{-1} \begin{pmatrix} x \cdot \mathbf{D}(x, y) \\ y \cdot \mathbf{D}(x, y) \\ \mathbf{D}(x, y) \end{pmatrix}. \quad (3)$$

For all views  $v = 1, \dots, N$ , each 3D point  $\chi_i^v$  in the world coordinate, where  $i \in \mathbb{R}^{HW}$ , can be derived using the above transform matrix.

**Dynamic Aligning Loss.** To handle challenging target sample points, we perform point pre-matching on all image pairs in the connectivity graph  $\mathcal{G}$ . For each image pair, we generate a matching map  $\Phi^{v,e} \in \mathbb{R}^{W \times H}$ , which indicates whether each sample point  $i \in \Phi^{v,e}$  has a matched point. Based on the confidence scores derived from Eq. 2, we define a dynamic adjustment function as follows:

$$A_i^{v,e} = \frac{F_i^{v,e} + \alpha_i^{v,e} \cdot F_i^{v,e} \cdot (1 - F_i^{v,e})}{1 + |\alpha_i^{v,e}| \cdot F_i^{v,e} \cdot (1 - F_i^{v,e}) + \epsilon}, \quad (4)$$

where  $F_i^{v,e} \cdot (1 - F_i^{v,e})$  enhances attention on difficult-to-match points, with confidence scores approaching 0.5, the axis of symmetry of the quadratic function. The denominator ensures that the adjusted confidence score  $A_i^{v,e}$  remains

within an effective range. The adjustment factor  $\alpha$  is defined as:

$$\alpha_i^{v,e} = \begin{cases} \alpha_p & \text{if } \Phi_i^{v,e} = 1 \\ -\alpha_n & \text{if } \Phi_i^{v,e} = 0 \end{cases}, \quad (5)$$

where  $\alpha_p$  is the positive adjustment factor for matching points, and  $\alpha_n$  is the negative adjustment factor for non-matching points. Both are used to differentiate the weights between matching and non-matching points.

To map the dynamically adjusted confidence score back to its original value range, we apply the inverse of the sigmoid function,  $\sigma^{-1}(\cdot)$ , as follows:

$$W_i^{v,e} = \sigma^{-1}(A_i^{v,e}) = -\ln\left(\frac{1}{A_i^{v,e}} - 1\right). \quad (6)$$

Finally, we compute the L2 distance between pointmaps of all image pairs and their corresponding pointmaps  $\chi^v$  in world coordinates. The final Dynamic Global Aligning is defined as:

$$\chi^* = \arg \min_{\chi, P, \sigma} \sum_{e \in \mathcal{E}} \sum_{v \in e} \sum_{i=1}^{HW} W_i^{v,e} \|\chi_i^v - \sigma_e P_e X_i^{v,e}\|, \quad (7)$$

Method	Scene-specific training	mIoU (%)	mAcc (%)	Total
NVOS [44]†	need	70.1	92.0	-
ISRF [11]	need	83.8	96.4	840 s
SGISRF [48]†	need	84.5	97.2	-
SA3D [4]	need	90.3	98.2	<u>780 s</u>
SAGA [5]	need	90.9	98.3	2280 s
OmniSeg3D [61]	need	91.7	98.4	8220 s
FlashSplat [45]†	need	<u>91.8</u>	<u>98.6</u>	1500 s
WildSeg3D (ours)	<b>no need</b>	<b>94.1</b>	<b>99.0</b>	<b>30 s</b>

Table 1. **Quantitative results** on NVOS dataset. **Boldface** highlights the best results and underline the second-best. “Total” denotes the overall time from scene reconstruction to completing an intuitive 3D segmentation. The symbol † denotes data taken from the respective references, as the code was not fully released.

Method	Scene-specific training	mIoU (%)	mAcc (%)	Time
Single view[4]	need	74.6	95.5	-
MVSeg [35]	need	90.9	98.9	180-360 s
ISRF [11]	need	77.4	93.46	2-3 s
SA3D [4]	need	92.4	<u>98.9</u>	120-600 s
SAGA [5]	need	88.0	98.5	<u>0.08-0.9 s</u>
OmniSeg3D [61]	need	<b>94.3</b>	<b>99.3</b>	1-2 s
WildSeg3D (ours)	<b>no need</b>	<u>94.0</u>	98.6	<b>0.005-0.02 s</b>

Table 2. **Quantitative results** on SPIn-NeRF dataset.

where  $\chi_i^v$  is estimated from the transform matrices as described in Eq. 3. After the above training, we obtain an optimized transform matrix for each viewpoint, which can effectively minimize the aligning errors introduced by the feed-forward mechanism.

### 3.3.3. Multi-view Group Mapping

WildSeg3D enables real-time 3D segmentation of target objects by using single-view prompts as input. To efficiently map multi-view object masks into an aligned 3D space based on user prompts, we propose a multi-view group mapping method designed to optimize mask retrieval and integration. The process begins by sorting all object masks within the mask cache for the current viewpoint in ascending order based on their area, with priority given to smaller, fine-grained objects. Based on the user’s prompts, we then sequentially filter the relevant masks from the cache, appending the corresponding object IDs to the result set  $O$ . Once the relevant masks are identified, we compute the union of these masks across all viewpoints in the dataset. Specifically, the unified mask for each viewpoint is defined as:

$$M = \left\{ M^v \mid M^v = \bigcup_{o \in O} m_o^v, v \in V \right\}, \quad (8)$$

where  $V$  represents the set of viewpoints,  $m_o^v$  denotes the mask of object  $o$  in viewpoint  $v$ , and  $M$  is the collection of masks for each viewpoint after prompt-based retrieval.

We apply the transform matrices learned by DGA to convert the segmentation masks from all viewpoints in  $M$  from pixel coordinates to aligned world coordinates. The result-

ing 3D segmentation is denoted as  $\mathcal{P}$ :

$$\mathcal{P} = \bigcup_{v \in V} \left\{ \mathbf{P}_v^{-1} \mathbf{K}_v^{-1} \begin{pmatrix} x \mathbf{D}_v(x, y) \\ y \mathbf{D}_v(x, y) \\ \mathbf{D}_v(x, y) \end{pmatrix} \mid (x, y) \in M_v, M_v(x, y) = 1 \right\}, \quad (9)$$

where  $\mathcal{P}$  represents a set of 3D point cloud coordinates for the object. Through this framework, WildSeg3D efficiently aggregates 2D masks from multiple viewpoints, enabling real-time intuitive segmentation via the P2W (pixel to world coordinates) strategy.

## 4. Experiments

### 4.1. Datasets

To evaluate the effectiveness of our method, we conducted experiments on multiple benchmark datasets [1, 20, 35, 44]. The NVOS dataset provides a reference view with scribble annotations for the segmented targets, as well as a target view with the corresponding segmentation mask, both captured from frontal perspectives. The SPIn-NeRF dataset, a 3D scene dataset, is annotated using the widely-adopted NeRF datasets [9, 20, 31, 32, 60]. It is used to assess the performance of intuitive 3D segmentation methods, including the evaluation of segmentation quality in more complex 3D environments. For qualitative experiments, we used the NVOS[44], SPIn-NeRF [35], T&T[20], and Mip-NeRF360 [1] datasets. These datasets were chosen to compare our method with existing approaches and to showcase the segmentation results produced by our method across different 3D scenes and viewpoints. To evaluate segmentation accuracy and facilitate comparisons, we use mean Intersection over Union (mIoU) and mean Accuracy (mAcc) as primary metrics. Additionally, we assess both the training duration and the time required for intuitive 3D segmentation on a single NVIDIA RTX 3090 GPU to evaluate the models’ efficiency and real-time performance.

### 4.2. Quantitative Results

**NVOS Dataset.** To ensure experimental fairness, we adopted the evaluation approach used by models such as SAGA [5], utilizing the scribble annotations provided by the NVOS dataset [44] as input to generate 2D masks for the SAM model [19]. Additionally, we performed random point sampling on the scribble annotations of the reference view to acquire point prompts for segmentation. The results of our experiments on the NVOS dataset are presented in Table 1, where we compare the performance of our WildSeg3D framework against other state-of-the-art methods. As shown in the table, WildSeg3D outperforms existing approaches in both mIoU and mAcc. In terms of runtime efficiency, NeRF-based methods, including NVOS [44], ISRF [11], SGISRF [48], and SA3D [4], as well as 3DGS-based methods such as SAGA [5], OmniSeg3D [61], and FlashSplat [45], typically require scene-specific training.

Datasets	Scene	D-GA	D-DGA	M-GA	M-SGA	M-DGA
NVOS	fern	82.5%	<b>85.1%</b>	94.1%	71.7%	<b>94.2%</b>
	flower	90.3%	<b>90.9%</b>	73.8%	91.3%	<b>94.3%</b>
	fortress	95.5%	<b>96.3%</b>	95.5%	96.0%	<b>96.8%</b>
	horns (center)	88.4%	<b>93.6%</b>	93.0%	92.9%	<b>95.9%</b>
	horns (left)	89.9%	<b>95.2%</b>	94.8%	<b>95.2%</b>	95.0%
	leaves	<b>65.8%</b>	61.9%	88.0%	93.5%	<b>96.7%</b>
	orchids	78.2%	<b>82.6%</b>	84.1%	84.1%	<b>93.5%</b>
	trex	79.8%	<b>80.2%</b>	85.3%	61.0%	<b>86.4%</b>
	average	83.8%	<b>85.7%</b>	88.6%	85.7%	<b>94.1%</b>
SPIn-NeRF	fern	82.5%	<b>85.1%</b>	94.1%	71.7%	<b>94.2%</b>
	fork	85.4%	<b>88.7%</b>	<b>89.8%</b>	88.3%	88.9%
	fortress	95.5%	<b>96.3%</b>	95.5%	96.0%	<b>96.8%</b>
	horns	88.4%	<b>93.6%</b>	93.0%	92.9%	<b>95.9%</b>
	leaves	65.8%	61.9%	88.0%	93.5%	<b>96.7%</b>
	lego	77.5%	<b>79.4%</b>	84.3%	81.5%	<b>92.8%</b>
	orchids	78.2%	<b>82.6%</b>	84.1%	84.1%	<b>93.5%</b>
	pinecone	84.2%	<b>88.6%</b>	85.2%	89.2%	<b>95.7%</b>
	room	<b>90.4%</b>	90.3%	88.7%	74.1%	<b>91.5%</b>
	truck	82.5%	<b>82.6%</b>	90.9%	90.5%	<b>93.7%</b>
average	83.1%	<b>84.9%</b>	89.4%	86.2%	<b>94.0%</b>	

Table 3. **Effect of different alignment strategies** for DUST3R-based and MAST3R-based 3D segmentation models on NVOS and SPIn-NeRF datasets. “D” and “M” represent the pointmaps from DUST3R and MAST3R, respectively. “GA” refers to global alignment, “SGA” denotes sparse global alignment in MAST3R, and “DGA” stands for our proposed dynamic global alignment.

In contrast, our approach demonstrates strong generalization across diverse scenes, eliminating the need for scene-specific training. This advantage leads to significant efficiency improvements, with our model requiring only 3.8% of the computation time of SA3D (NeRF-based) and 2% of FlashSplat (3DGS-based), while maintaining high segmentation accuracy.

**SPIn-NeRF Dataset.** The quantitative results of our experiments on the SPIn-NeRF dataset are presented in Table 2. We evaluate the accuracy of our method by projecting the 3D segmentation masks onto the reference views and comparing them against the ground truth. In the MVSeg method [14], an intuitive segmentation model is first employed to acquire object masks from a single view. These masks are then used to segment training views, which are treated as a video sequence and processed through a video instance segmentation model [3, 53] to generate 3D masks. These masks are further refined using a semantic NeRF model [34, 65, 66]. For SA3D [4], intuitive segmentation requires traversing all training views. Cross-view self-prompting is used to guide SAM in generating 2D masks, which are subsequently projected into 3D space using Mask Inverse Rendering. This process incurs significant time costs due to the need to process each view individually. SAGA [5] incorporates low-dimensional 3D features for each Gaussian, which are rendered into 2D feature maps through differentiable rasterization during training. This approach allows for contrastive training with the segmentation results generated by SAM. In comparison to MVSeg [35] and SA3D [4], our method achieves comparable accuracy




Figure 3. **Performance of WildSeg3D on indoor and outdoor scenes.** For each scene, the left side shows the sparse views for reconstruction, with the segmentation target indicated by red dashed lines in the first view as prompts.

while requiring only a fraction—approximately one thousandth—of the processing time. Moreover, when compared to SAGA, our approach not only demonstrates superior accuracy but also outperforms SAGA in terms of processing time. Overall, our method delivers exceptional performance, making real-time intuitive 3D segmentation feasible in practical applications.

### 4.3. Qualitative Results

We conduct a comparative analysis of our method against existing approaches and present qualitative experimental results. As shown in Figure 2, the first column illustrates the reconstruction time of our model for various scenes based on sparse views. The subsequent columns compare the segmentation results produced by our method with those from other approaches, along with the time taken to obtain 3D segmentation results after user input. The final column displays the segmentation results of ISRF [11] across different scenes. ISRF utilizes the TensorRF representation [6] for scene training and rendering, incorporating DINO features [3] within each voxel to enable 2D-to-3D semantic matching. Segmentation is performed via nearest neighbor feature matching (NNFM). However, ISRF struggles with distinguishing semantically similar objects, as demonstrated in the Trex and Orchids scenes, where the method faces challenges in accurate segmentation. SA3D [4], upon receiving user prompts, employs a cross-view self-

Number of Views	3	4	5	7	9
NVOS	89.8%	92.9%	<b>94.1%</b>	94.0%	93.8%
SPIn-NeRF	87.8%	92.2%	94.0%	<b>94.9%</b>	94.7%
Total Time	15 s	18 s	23 s	35 s	50 s

Table 4. **Ablation Experiments** on the Impact of Viewpoint Quantity on mIoU.

prompting mechanism and SAM [19] to generate 2D masks for each view. These masks are then mapped to 3D space via Mask Inverse Rendering, with the process repeated for each user interaction. This approach incurs high computational costs due to the need for repeated processing. SAGA [5] and OmniSeg3D [61] transform 2D masks into 3D features and require additional training for each 3D Gaussian. SAGA relies on a collection of loss functions during training, whereas OmniSeg3D uses hierarchical contrastive learning to map discontinuous multi-view segmentations to consistent 3D features. In contrast, our method achieves superior performance with high-quality transform matrices learned through DGA, which map 2D views into an aligned 3D coordinate system. The MGM module then enables fast retrieval of 2D masks from multiple views, which are subsequently transformed into the unified 3D space. This design results in significantly improved computational efficiency, allowing our method to achieve faster intuitive segmentation compared to existing approaches.

As demonstrated in Figure 3, our method does not require scene-specific training, enabling near real-time 3D intuitive segmentation for arbitrary scenes. We showcase this capability using indoor and outdoor scenes from the MipNeRF360 [1], T&T [20], SPIn-NeRF [35] datasets, and two additional scenes in the wild. By leveraging only sparse views, our method is able to perform both scene reconstruction and near real-time 3D segmentation, underscoring its versatility and ability to segment objects in diverse, uncontrolled environments.

#### 4.4. Ablation Studies

**Comparison of Different Viewpoint Counts.** Table 4 presents an ablation study on the number of viewpoints input into WildSeg3D. Considering overall performance, we randomly select five viewpoints in prior experiments. On the SPIn-NeRF dataset, our method outperforms OmniSeg3D when seven viewpoints are selected, despite the results in Table 2.

**Effect of Dynamic Global Alignment.** To further validate the effectiveness of DGA, we also integrate it into DUS3R, replacing the matching maps with SIFT [29]. As shown in Table 3, incorporating DGA significantly improves performance in both DUS3R and MAST3R models compared to traditional GA training methods.

Furthermore, as shown in Figure 4, comparative visualizations on the NVOS and LERF [18] datasets demonstrate




Figure 4. **Visualization** of ablation experiments on DGA.

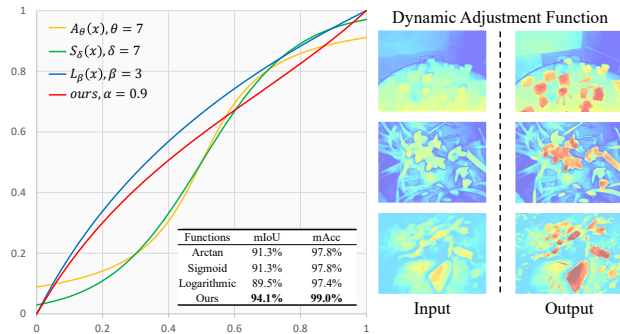


Figure 5. **Ablation Experiments** on NVOS Dataset. Left: visualization of different confidence adjustment functions for DGA. Right: Results of these functions.

that DGA effectively reduces 3D alignment errors, minimizing blurred details and background confusion.

**Different Confidence adjustment functions.** Figure 5 compares the confidence adjustment function in DGA with other adjustment functions, defined as:  $A_\theta(x) = \arctan(\frac{\theta \cdot (x-0.5)}{\pi}) + 0.5$ ,  $S_\delta(x) = \frac{1}{1+e^{-\delta \cdot (x-0.5)}}$ ,  $L_\beta(x) = \frac{\log(1+\beta \cdot x)}{\log(1+\beta)}$ , where  $\theta$ ,  $\delta$ , and  $\beta$  control the scaling of the Arctan, Sigmoid, and Logarithmic functions, respectively.

Our function performs adaptive adjustments centered at 0.5 confidence, improving the weight of hard-to-match points while ensuring smooth transitions across both low and high confidence levels, achieving the best performance.

## 5. Conclusion

We introduce WildSeg3D, a feed-forward method that enables real-time intuitive 3D segmentation from 2D images without scene-specific pre-training. With Dynamic Global Aligning, WildSeg3D learns high-quality transform matrix for each view, aligning 2D images to an aligned coordinate system. Additionally, mask cache stores masks from multiple views consistently. For real-time intuitive segmentation, MGM uses a search strategy to retrieve target masks and map them into an aligned 3D space, responding promptly to user inputs. Through extensive experiments, WildSeg3D demonstrates a significant speedup over existing methods while maintaining high accuracy.

## References

- [1] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5470–5479, 2022. 6, 8
- [2] Yash Bhargat, Iro Laina, Joao F Henriques, Andrew Zisserman, and Andrea Vedaldi. Contrastive lift: 3d object instance segmentation by slow-fast contrastive fusion. *arXiv preprint arXiv:2306.04633*, 2023. 3
- [3] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 9650–9660, 2021. 7
- [4] 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. 1, 2, 3, 6, 7
- [5] Jiazhong Cen, Jiemin Fang, Chen Yang, Lingxi Xie, Xiaopeng Zhang, Wei Shen, and Qi Tian. Segment any 3d gaussians, 2024. 1, 2, 3, 6, 7, 8
- [6] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *European conference on computer vision*, pages 333–350. Springer, 2022. 7
- [7] Hanlin Chen, Chen Li, Mengqi Guo, Zhiwen Yan, and Gim Hee Lee. Gnesf: Generalizable neural semantic fields. *Advances in Neural Information Processing Systems*, 36:36553–36565, 2023. 3
- [8] Seokhun Choi, Hyeonseop Song, Jaechul Kim, Taehyeong Kim, and Hoseok Do. Click-gaussian: Interactive segmentation to any 3d gaussians. *arXiv preprint arXiv:2407.11793*, 2024. 3
- [9] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5501–5510, 2022. 6
- [10] Qiancheng Fu, Qingshan Xu, Yew Soon Ong, and Wenbing Tao. Geo-neus: Geometry-consistent neural implicit surfaces learning for multi-view reconstruction. *Advances in Neural Information Processing Systems*, 35:3403–3416, 2022. 2
- [11] Rahul Goel, Dhawal Sirikonda, Saurabh Saini, and PJ Narayanan. Interactive segmentation of radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4201–4211, 2023. 1, 2, 3, 6, 7
- [12] Yubin Gu, Honghui Xu, Yueqian Quan, Wanjun Chen, and Jianwei Zheng. Orsi salient object detection via bidimensional attention and full-stage semantic guidance. *IEEE Transactions on Geoscience and Remote Sensing*, 61:1–13, 2023. 3
- [13] Jianfei Guo, Nianchen Deng, Xinyang Li, Yeqi Bai, Botian Shi, Chiyu Wang, Chenjing Ding, Dongliang Wang, and Yikang Li. Streetsurf: Extending multi-view implicit surface reconstruction to street views. *arXiv preprint arXiv:2306.04988*, 2023. 2
- [14] Yuying Hao, Yi Liu, Zewu Wu, Lin Han, Yizhou Chen, Guowei Chen, Lutao Chu, Shiyu Tang, Zhiliang Yu, Zeyu Chen, et al. Edgeflow: Achieving practical interactive segmentation with edge-guided flow. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1551–1560, 2021. 7
- [15] Xu Hu, Yuxi Wang, Lue Fan, Junsong Fan, Junran Peng, Zhen Lei, Qing Li, and Zhaoxiang Zhang. Semantic anything in 3d gaussians. *arXiv preprint arXiv:2401.17857*, 2024. 1, 3
- [16] Chi Huang, Xinyang Li, Shengchuan Zhang, Liujuan Cao, and Rongrong Ji. Nerf-dets: Enhancing multi-view 3d object detection with sampling-adaptive network of continuous nerf-based representation. *arXiv e-prints*, pages arXiv–2404, 2024. 3
- [17] 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. 1, 3
- [18] 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. 1, 3, 8
- [19] 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. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4015–4026, 2023. 1, 3, 6, 8
- [20] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017. 6, 8
- [21] Sosuke Kobayashi, Eiichi Matsumoto, and Vincent Sitzmann. Decomposing nerf for editing via feature field distillation. *Advances in Neural Information Processing Systems*, 35:23311–23330, 2022. 1
- [22] Theodora Kontogianni, Ekin Celikkan, Siyu Tang, and Konrad Schindler. Interactive object segmentation in 3d point clouds. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2891–2897. IEEE, 2023. 3
- [23] Itai Lang, Fei Xu, Dale Decatur, Sudarshan Babu, and Rana Hanocka. iseg: Interactive 3d segmentation via interactive attention. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024. 3
- [24] Vincent Leroy, Yohann Cabon, and Jérôme Revaud. Grounding image matching in 3d with mast3r. *arXiv preprint arXiv:2406.09756*, 2024. 1, 3
- [25] Jiangang Lin, Yunhang Shen, Bingquan Wang, Shaohui Lin, Ke Li, and Liujuan Cao. Weakly supervised open-vocabulary object detection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 3404–3412, 2024. 3

- [26] Yichen Liu, Benran Hu, Junkai Huang, Yu-Wing Tai, and Chi-Keung Tang. Instance neural radiance field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 787–796, 2023. 3
- [27] Yichen Liu, Benran Hu, Chi-Keung Tang, and Yu-Wing Tai. Sanerf-hq: Segment anything for nerf in high quality. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3216–3226, 2024. 1
- [28] Xiaoxiao Long, Cheng Lin, Peng Wang, Taku Komura, and Wenping Wang. Sparseneus: Fast generalizable neural surface reconstruction from sparse views. In *European Conference on Computer Vision*, pages 210–227. Springer, 2022. 2
- [29] G Lowe. Sift-the scale invariant feature transform. *Int. J.* 2 (91-110):2, 2004. 8
- [30] Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. *arXiv preprint arXiv:2308.09713*, 2023. 3
- [31] 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)*, 38(4):1–14, 2019. 6
- [32] 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. 6
- [33] 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, 2, 3
- [34] Ashkan Mirzaei, Yash Kant, Jonathan Kelly, and Igor Gilitschenski. Laterf: Label and text driven object radiance fields. In *European Conference on Computer Vision*, pages 20–36. Springer, 2022. 7
- [35] Ashkan Mirzaei, Tristan Aumentado-Armstrong, Konstantinos G Derpanis, Jonathan Kelly, Marcus A Brubaker, Igor Gilitschenski, and Alex Levinshtein. Spin-nerf: Multiview segmentation and perceptual inpainting with neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20669–20679, 2023. 2, 6, 7, 8
- [36] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021. 3
- [37] 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. 3
- [38] 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. 1
- [39] Zhipeng Qian, Yiwei Ma, Jiayi Ji, and Xiaoshuai Sun. X-refseg3d: Enhancing referring 3d instance segmentation via structured cross-modal graph neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 4551–4559, 2024. 3
- [40] 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. 1, 3
- [41] Yansong Qu, Shaohui Dai, Xinyang Li, Jianghang Lin, Lijuan Cao, Shengchuan Zhang, and Rongrong Ji. Goi: Find 3d gaussians of interest with an optimizable open-vocabulary semantic-space hyperplane. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pages 5328–5337, 2024. 3
- [42] 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. 3
- [43] Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, et al. Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714*, 2024. 1, 3
- [44] Zhongzheng Ren, Aseem Agarwala, Bryan Russell, Alexander G Schwing, and Oliver Wang. Neural volumetric object selection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6133–6142, 2022. 2, 6
- [45] Qihong Shen, Xingyi Yang, and Xinchao Wang. Flashsplat: 2d to 3d gaussian splatting segmentation solved optimally. *arXiv preprint arXiv:2409.08270*, 2024. 2, 3, 6
- [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. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9043–9052, 2023. 3
- [47] Ayça Takmaz, Elisabetta Fedele, Robert W Sumner, Marc Pollefeys, Federico Tombari, and Francis Engelmann. Openmask3d: Open-vocabulary 3d instance segmentation. *arXiv preprint arXiv:2306.13631*, 2023. 1
- [48] Songlin Tang, Wenjie Pei, Xin Tao, Tanghui Jia, Guangming Lu, and Yu-Wing Tai. Scene-generalizable interactive segmentation of radiance fields. In *Proceedings of the 31st ACM International Conference on Multimedia*, pages 6744–6755, 2023. 2, 6
- [49] Vadim Tschernezki, Iro Laina, Diane Larlus, and Andrea Vedaldi. Neural feature fusion fields: 3d distillation of self-supervised 2d image representations. In *2022 International Conference on 3D Vision (3DV)*, pages 443–453. IEEE, 2022. 1

- [50] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 2
- [51] Shuzhe Wang, Vincent Leroy, Yohann Cabon, Boris Chidlovskii, and Jerome Revaud. Dust3r: Geometric 3d vision made easy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20697–20709, 2024. 1, 3
- [52] Shizun Wang, Xingyi Yang, Qihong Shen, Zhenxiang Jiang, and Xinchao Wang. Gflow: Recovering 4d world from monocular video. *arXiv preprint arXiv:2405.18426*, 2024. 3
- [53] Wenguan Wang, Tianfei Zhou, Fatih Porikli, David Crandall, and Luc Van Gool. A survey on deep learning technique for video segmentation. *arXiv e-prints*, pages arXiv–2107, 2021. 7
- [54] Xinlong Wang, Xiaosong Zhang, Yue Cao, Wen Wang, Chunhua Shen, and Tiejun Huang. Seggpt: Towards segmenting everything in context. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1130–1140, 2023. 3
- [55] Yiqun Wang, Ivan Skorokhodov, and Peter Wonka. Hf-neus: Improved surface reconstruction using high-frequency details. *Advances in Neural Information Processing Systems*, 35:1966–1978, 2022. 2
- [56] Yuze Wang, Junyi Wang, Ruicheng Gao, Yansong Qu, Wantong Duan, Shuo Yang, and Yue Qi. Look at the sky: Sky-aware efficient 3d gaussian splatting in the wild. *IEEE Transactions on Visualization and Computer Graphics*, 2025. 3
- [57] Zeyu Yang, Hongye Yang, Zijie Pan, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. *arXiv preprint arXiv:2310.10642*, 2023.
- [58] Ziyi Yang, Xinyu Gao, Wen Zhou, Shaohui Jiao, Yuqing Zhang, and Xiaogang Jin. Deformable 3d gaussians for high-fidelity monocular dynamic scene reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20331–20341, 2024. 3
- [59] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. *arXiv preprint arXiv:2312.00732*, 2023. 1, 3
- [60] Lin Yen-Chen, Pete Florence, Jonathan T Barron, Tsung-Yi Lin, Alberto Rodriguez, and Phillip Isola. Nerf-supervision: Learning dense object descriptors from neural radiance fields. In *2022 international conference on robotics and automation (ICRA)*, pages 6496–6503. IEEE, 2022. 6
- [61] Haiyang Ying, Yixuan Yin, Jinzhi Zhang, Fan Wang, Tao Yu, Ruqi Huang, and Lu Fang. Omnise3d: Omniversal 3d segmentation via hierarchical contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20612–20622, 2024. 2, 6, 8
- [62] Hong-Xing Yu, Leonidas J Guibas, and Jiajun Wu. Unsupervised discovery of object radiance fields. *arXiv preprint arXiv:2107.07905*, 2021. 3
- [63] Yuanwen Yue, Sabarinath Mahadevan, Jonas Schult, Francis Engelmann, Bastian Leibe, Konrad Schindler, and Theodora Kontogianni. Agile3d: Attention guided interactive multi-object 3d segmentation. *arXiv preprint arXiv:2306.00977*, 2023. 3
- [64] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15838–15847, 2021. 3
- [65] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15838–15847, 2021. 7
- [66] Shuaifeng Zhi, Edgar Sucar, Andre Mouton, Iain Haughton, Tristan Laidlow, and Andrew J Davison. ilabel: Interactive neural scene labelling. *arXiv preprint arXiv:2111.14637*, 2021. 7
- [67] 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. 1, 3