

Online Segment 3D Gaussians via Launching Virtual Drones

Liwei Liao^{1,2,3}, Rongjie Wang², and Ronggang Wang^{1,2,3}✉

¹ Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology, Shenzhen Graduate School, Peking University, Shenzhen, China

² Pengcheng Laboratory, Shenzhen, China

³ Peking University, Beijing, China

lewio@pku.edu.cn, rgwang@pkusz.edu.cn

Abstract. Interactive segmentation of 3D Gaussians offers a compelling opportunity for real-time manipulation of 3D scenes, thanks to the real-time rendering capability of 3D Gaussian Splatting (3DGS). However, existing methods require a time-consuming per-scene setup—typically tens of seconds or even minutes—before interactive segmentation can begin on a raw 3DGS scene. This setup involves multi-view mask preparation, mask lifting, and feature distillation, creating a major bottleneck for online applications. To address this limitation, we aim to completely eliminate the setup stage for interactive 3DGS segmentation while keeping the segmentation time practical (under 1 second). In this work, we present **SAGO** (Segment Any Gaussians One), a novel setup-free framework for interactive 3DGS segmentation. By introducing **virtual drones**, our method reframes the 3D segmentation problem as an online Next-Best-View (NBV) planning task formulated within a Markov process. Extensive experiments demonstrate that **SAGO** can extract clean 3D assets directly from 3D Gaussians with sub-second latency, thereby enabling a broad range of downstream applications such as object manipulation and scene editing. Moreover, our method achieves over a **50×** speedup compared to the previous setup-free 3DGS segmentation frameworks.

Keywords: Interactive 3D Segmentation · 3D Gaussian Splatting · Next Best View · Online Planning · 3D Segmentation · Active Vision

1 Introduction

In recent years, 3D Gaussian Splatting (3DGS) [14] has seen remarkable progress in 3D representation and has shown great potential as a medium for 3D scene understanding in downstream applications such as embodied AI [8, 27], robotics [28, 58], data synthesis [19, 45] and immersive media [7, 20, 21, 31, 39–42, 44, 46, 47]. To equip 3DGS with capabilities of interaction, editing, and scene understanding, a series of 3D segmentation methods [4, 6, 11, 12, 18, 22–24, 32, 37, 43, 48, 51, 53–55] have achieved remarkable progress. Most of them follow a pipeline as shown in Fig. 1 that requires a scene-specific setup stage to pre-process the pretrained




Fig. 1: The common pipeline of interactive 3DGS segmentation. Existing methods require a setup stage of tens of seconds or even minutes before interactive segmentation can be performed on a given raw 3DGS scene.

3DGS scene to make it interactively segmentable. The setup stage typically involves pre-segmenting training views with 2D segmentation models, such as the Segment Anything (SAM) series [3, 16, 34], to generate 2D masks. These masks are subsequently utilized for mask-lifting, feature distillation, or other forms of optimization. Due to the setup stage, existing methods require an additional time of tens of seconds or even minutes to process a new 3DGS scene before interactive segmentation can be performed, which significantly limits their efficiency and user experience in practical applications.

To eliminate the setup stage, we explore an online and inference-only approach for 3DGS segmentation. Our goal is to enable interactive segmentation of a given raw 3DGS scene within a user-tolerable time frame after the user provides a prompt describing the target object (e.g., clicks, boxes, or text), while achieving comparable segmentation quality to existing methods requiring a per-scene setup stage (i.e., offline methods). To achieve this, we revisit the problem of online view planning within the context of active robotic vision [13, 50, 52], where autonomous drones are deployed to explore unknown regions of physical scenes for high-fidelity digital reconstruction. Each drone is equipped with a camera that can be maneuvered to capture images from various viewpoints, and the Next-Best-View (NBV) planning module is responsible for determining the optimal camera parameters to maximize the coverage of the target object. This NBV planning problem shares a significant similarity with our 3D segmentation task: *each additional effective 2D view can provide substantial cues for the final 3D task*. In our 3D segmentation framework, 2D segmentation masks from novel viewpoints serve as spatial constraints to effectively “prune” background Gaussians that project outside the mask contours.

Building on this insight, we introduce **virtual drones** and formulate the 3D segmentation as an NBV planning problem for virtual drones. Starting from the initial view (i.e., the view for user interaction), the virtual drone iteratively plans its NBV based on the current state in a Markov Process framework. At each step, the virtual drone captures a new view of the scene, and the 2D mask obtained from SAM is used to update the state by pruning background Gaussians that project outside the mask contours. By iteratively updating the state and

planning the NBV, our method can effectively segment the target object in 3D without requiring a time-consuming setup stage, while still maintaining high segmentation quality. In summary, our key contributions are as follows:

- We propose **SAGO**, an effective and efficient online 3DGS segmentation framework. SAGO completely eliminates the time-consuming scene-specific setup stage required by most existing methods, while delivering segmentation performance comparable (or superior) to state-of-the-art offline approaches.
- We introduce the concept of **virtual drones** and present a principled formulation that recasts 3D online segmentation as a Next-Best-View (NBV) planning problem for a virtual drone navigating and observing the scene.
- In contrast to previous setup-free methods, ours achieves over **50×** inference speedup while simultaneously delivering superior segmentation quality.

2 Related Works

Optimization-based 3DGS Segmentation. Most existing approaches [4, 6, 32, 48, 51, 57] rely on a preparatory setup stage in which 2D segmentation masks are first generated for a set of training views, followed by a dedicated optimization procedure—commonly involving mask lifting, feature distillation, or contrastive learning. For example, SAGA [4] employs scale-aware contrastive learning to extract scale-gated affinity features, while Click-Gaussian [6] aggregates global feature candidates to provide consistent supervision for Gaussian feature learning. COB-GS [51] introduces boundary-adaptive Gaussian splitting to better resolve ambiguous regions near object boundaries. LangSplat [32] distills semantic knowledge from CLIP [33] into 3D Gaussians using an auto-encoder architecture. These optimization-based methods typically require iterative gradient-based optimization, resulting in substantial computational time and memory costs during the setup phase.

Online 3DGS Segmentation. Recent methods [11, 12, 23, 36, 38, 53] have explored optimization-free approaches to 3D Gaussian Splatting (3DGS) segmentation, a paradigm we refer to as *online 3DGS segmentation*. These methods generally bypass time-consuming per-scene optimization and gradient-based back-propagation. For example, GaussianCut [12] constructs a graph over Gaussians and reformulates 3D segmentation as a graph partitioning problem. iSegMan [53] introduces epipolar-guided interaction propagation combined with visibility-based Gaussian voting to achieve fully training-free segmentation. SAGD [11] employs Gaussian decomposition to alleviate jagged boundary artifacts. Trace3D [36] proposes Gaussian Instance Tracing (GIT), which explicitly associates each Gaussian with its corresponding instance labels across multiple views through a learned weight matrix.

Although these online 3DGS segmentation methods eliminate per-scene optimization, most still rely on non-trivial preprocessing steps and are not truly free of any setup phase. More critically, their per-interaction latency remains significantly higher than that of optimization-based approaches, which benefit from precomputed features. For example, GaussianCut [12] performs online




Fig. 2: The pipeline of SAGO. Following user input, SAGO launches virtual drones in both clockwise and counter-clockwise directions from the interaction viewpoint to explore Next-Best-Views. Each newly captured view provides a segmentation mask that effectively filters out background Gaussians via Mask-shaped Frustum Filtering (MFF), with the process repeating until full coverage is achieved.

graph construction, while iSegMan [53] requires online epipolar-guided interaction propagation and visibility-based Gaussian voting—both of which incur considerable computational cost during user interactions. In contrast, our proposed method introduces the virtual drones and leverages the NBV planning theory to achieve a completely setup-free pipeline with highly efficient interaction times, typically under 1 second per interaction.

3 Method

We propose SAGO, an online 3DGS segmentation method based on virtual drone NBV planning, enabling interactive segmentation without the need for a time-consuming setup stage. In this section, we first introduce the relevant knowledge of 3D Gaussian Splatting (3DGS), particularly the geometric projection of 3D Gaussians onto the 2D plane, which allows us to use 2D masks for background filtering, forming the theoretical basis for our online segmentation method. Next, we will detail the state initialization process and the online view planning and state update mechanism, which are core components of our method.

3.1 Preliminary: 3D Gaussian Splatting

3D Gaussian Splatting (3DGS) [14] has emerged as a prominent approach for real-time novel view synthesis, representing scenes using a collection of explicit 3D Gaussians, achieving an excellent trade-off between rendering quality and speed. A set of pretrained 3D Gaussians $\mathcal{G} = \{G_i\}_{i=1}^N$ is obtained by optimizing the parameters of the Gaussians through differentiable rendering techniques. Each Gaussian G_i is parameterized as $G_i = \{x_i, q_i, s_i, c_i, o_i\}$, where $x_i \in \mathbb{R}^3$

is the 3D position, $q_i \in \mathbb{R}^4$ is the quaternion representing the rotation, $s_i \in \mathbb{R}^3$ is the scale, $c_i \in \mathbb{R}^{48}$ is the view-dependent color represented by three-order spherical harmonics [9], and $o_i \in \mathbb{R}$ is the opacity. Based on their role in 3D segmentation, Gaussian attributes can be categorized into two groups: **geometric attributes** (position, covariance (including rotation and scaling)) and **appearance attributes** (color and opacity). Geometric attributes can serve as a discriminative basis for distinguishing the foreground from the background, while appearance attributes provide extra semantic information to distinguish different instances. In the rendering stage, the 2D covariance matrix Σ_i^{2D} of a Gaussian G_i can be computed as follows:

$$\Sigma_i^{2D} = JW(R_i S_i S_i^T R_i^T) W^T J^T, \quad (1)$$

where R_i is the rotation matrix derived from the quaternion q_i , S_i is the diagonal matrix of the scale s_i , W is the projection matrix, and J is the Jacobian matrix of the projection. The 2D covariance Σ_i^{2D} represents the i -th Gaussian's shape in the 2D plane, and the projected center of the Gaussian is given by $x_i^{2D} = \text{project}(x_i)$. Thus, the i -th Gaussian can be projected onto the 2D plane as an ellipse (referred to as a "splat"), defined by $\{x_i^{2D}, \Sigma_i^{2D}\}$. This projection allows us to leverage 2D masks to inform the segmentation of the 3D Gaussians (see Fig. 3), forming the basis for our online segmentation pipeline.

3.2 Overview

Problem Definition. Given a set of 3D Gaussians $\mathcal{G}$ representing a scene and a user-provided prompt (e.g., point clicks, scribbles, or text) on a 2D view, interactive 3D segmentation aims to partition $\mathcal{G}$ into two disjoint, non-empty subsets $\mathcal{A}$ and $\mathcal{B}$. Here, $\mathcal{A}$ represents the Gaussians corresponding to the target object (the foreground), while $\mathcal{B}$ represents the background, such that $\mathcal{A} \cup \mathcal{B} = \mathcal{G}$ and $\mathcal{A} \cap \mathcal{B} = \emptyset$. This can be formulated as:


$$\mathcal{A}, \mathcal{B} = \text{segment}(\mathcal{G}, \text{prompt}). \quad (2)$$

Overall pipeline. As shown in Fig. 2, we model the interactive 3D segmentation process as an NBV planning problem for virtual drones, where the state at each time step t is represented as S_t . The virtual drones explore the pre-defined search space to find optimal views for segmentation. At each new view, a segmentation mask is generated via SAM, which is then used to filter out background Gaussians. With the updated segmentation results, the state is updated and the next NBV is planned, iterating until full coverage of the target object is achieved. The core components of this pipeline include state initialization, NBV planning, and state update, which will be detailed in the following sections.

The key notations used in this paper are summarized in Tab. 1 for clarity and reference throughout the subsequent sections.

Table 1: A list of key notations used in this paper.

Symbol	Description at Step t
$\mathcal{A}^{(t)}, \mathcal{B}^{(t)}$	Foreground/background subsets at step t
S_t	Segmentation state at step t
$\mathcal{M}_t$	Segmentation memory bank at step t
τ_t	Coverage indicator at step t
$\mathbf{v}_t$	Camera pose of virtual drone at step t
$\mathbf{c}_t$	Target center/the point virtual drones look at
r_t	The distance between $\mathbf{c}_t$ and the virtual drone
ϕ_t	The yaw angle of the virtual drone at step t
θ_t	The pitch angle of the virtual drone at step t

**Fig. 3:** Mask-shaped Frustum Filtering (MFF) illustration.

3.3 State Initialization

We leverage virtual drone Next-Best-View (NBV) planning to address the online 3DGS segmentation problem, which can be modeled as a Markov process where the next state depends only on the current state. Starting from an initial 2D user-prompted view, which serves as the virtual drone’s origin, we iteratively plan a sequence of NBVs. At each new viewpoint, a segmentation mask is generated via SAM with memories saved from previous states to filter out background Gaussians, continuing until full coverage is achieved. We define the state at each step using a four-tuple of variables: the current set of foreground Gaussians $\mathcal{A}^{(t)}$, the virtual drone’s pose $\mathbf{v}_t$, a memory bank $\mathcal{M}_t$ for SAM tracking, and a coverage indicator τ_t . The drone’s pose is represented in an object-centric spherical coordinate system, consisting of the target center $\mathbf{c}_t$, radius r_t , yaw ϕ_t , and pitch θ_t . This can be formulated as:

$$\begin{aligned} \mathcal{S}_t &= \{\mathcal{A}^{(t)}, \mathbf{v}_t, \mathcal{M}_t, \tau_t\}, \\ \mathbf{v}_t &= \{\mathbf{c}_t, r_t, \phi_t, \theta_t\}. \end{aligned} \quad (3)$$

According to Eq. (3), we can initialize the state $\mathcal{S}_0 = \{\mathcal{A}^{(0)}, \mathbf{v}_0, \mathcal{M}_0, \tau_0\}$. The initial virtual drone pose $\mathbf{v}_0$ is derived from the camera pose of the interaction viewpoint with yaw ϕ_0 set to 0, and τ_0 is set to 0 initially. The memory bank $\mathcal{M}_0$ is initialized via a SAM memory encoder [34] based on the user’s prompt. The initial foreground Gaussian set $\mathcal{A}^{(0)}$ is a subset of $\mathcal{G}$, with a part of background Gaussians pruned by mask-shaped frustum filtering based on the initial mask.

Mask-shaped Frustum Filtering (MFF). The 2D mask on each view can be used to filter out certain background Gaussians, and we call this process mask-shaped frustum filtering. Specifically, as shown in Fig. 3, we project each Gaussian onto the 2D plane and check its spatial relationship with the mask region. If the Gaussian is determined to overlap with the mask, we consider it as a potential foreground candidate and include it in $\mathcal{A}$. Otherwise, we classify it as background set $\mathcal{B}$ and exclude it from the foreground set. This filtering process

helps to reduce the number of Gaussians that need to be considered in subsequent iterations, improving the efficiency of the online segmentation process. The MFF in each view can exclude a group of Gaussians $\Delta\mathcal{A}$ from the foreground set $\mathcal{A}$ to the background set $\mathcal{B}$, which can be formulated as:

$$\begin{aligned}\mathcal{A}^{(t+1)} &= \mathcal{A}^{(t)} - \Delta\mathcal{A}^{(t)}, \\ \mathcal{B}^{(t+1)} &= \mathcal{B}^{(t)} + \Delta\mathcal{A}^{(t)}.\end{aligned}\tag{4}$$

MFF can be implemented in two ways: center-based MFF and splat-based MFF. Center-based MFF directly projects the center of each Gaussian onto the 2D mask, classifying it as background if the center falls outside the mask. Splat-based MFF, on the other hand, projects each Gaussian as a splat and classifies it as foreground if there is any overlap between the splat and the mask; otherwise, it is classified as background. We adopt the center-based MFF in our implementation for its simplicity and efficiency, and the performance is even better (see Sec. 4.5).

3.4 Online NBV Planning and State Update

As illustrated in Fig. 2, we launch two virtual drones simultaneously from the interaction viewpoint, exploring NBVs in both clockwise and counter-clockwise directions. Each drone is responsible for exploring a 180-degree yaw range. Due to the symmetry of the two trajectories, we demonstrate the update process of one trajectory, which will be explained in the following using the clockwise direction as an example (see Fig. 4).

To reach a full coverage of the target object, we leverage an object-centric spherical coordinate system for NBVs planning. The camera pose of the virtual drone at step t is represented as $\mathbf{v}_t = (\mathbf{c}_t, r_t, \phi_t, \theta_t)$, where $\mathbf{c}_t$ is the target center that the virtual drone looks at, r_t is the distance between $\mathbf{c}_t$ and the virtual drone, ϕ_t is the yaw angle of the virtual drone at step t , and θ_t is the pitch angle. The target center $\mathbf{c}_t$ is initialized as the average of the 3D positions of the Gaussians that are projected inside the initial mask m_0 at step 0, and is updated as the average of the 3D positions of the currently segmented foreground Gaussians $\mathcal{A}^{(t-1)}$ at step t (if $t \geq 2$), as follows:

$$\mathbf{c}_t = \begin{cases} \frac{1}{|X_{\text{reproj}(m_0)}|} \sum_{x \in X_{\text{reproj}(m_0)}} x, & \text{if } t = 0, 1, \\ \frac{1}{|X_{\mathcal{A}^{(t-1)}}|} \sum_{x \in X_{\mathcal{A}^{(t-1)}}} x, & \text{if } t \geq 2, \end{cases}\tag{5}$$

where X denotes the 3D positions of Gaussians or 3D points. Eq. (5) actually determines the geometric centroid or 3D location of the target object, and the centroid $\mathbf{c}_t$ acts as the origin of the spherical coordinate system for NBV planning. In the initial stage, our method can estimate the centroid $\mathbf{c}_t$ roughly through the 3D points obtained by depth back-projection of the 2D mask from




Fig. 4: The illustration of NBV-based online 3D segmentation (clockwise direction).

the initial view, which usually has a certain offset and lies near the surface of the object. The offset always leads to the object being not well centered in the next viewpoint when the yaw span is too large. But when the foreground set $\mathcal{A}$ is updated closer to the ground truth, the centroid $\mathbf{c}_t$ can be updated to be more accurate using the mean of the 3D positions of the current $\mathcal{A}^{(t)}$.

The radius r_t is initialized to a fixed value that is the distance between the target center $\mathbf{c}_t$ and the start point of the virtual drones, and remains constant during the state update process to make the virtual drone maintain a consistent distance from the target object.

The yaw angle ϕ_t controls the horizontal circular motion of the virtual drone, starting at an initial value of 0° . When ϕ_t reaches 360° , the virtual drone completes a full revolution around the object and returns to its starting position. In our setup, the clockwise direction is defined as ϕ increasing from 0° to 180° .

The pitch angle θ_t governs the vertical movement of the virtual drone, enabling it to explore different elevations around the target object and thereby achieve more comprehensive segmentation coverage. This angle is constrained to the interval $[0^\circ, 90^\circ]$.

Since r_t remains constant and $\mathbf{c}_t$ is deterministically computed from the current foreground set, our NBV planning primarily focuses on determining the optimal yaw angle ϕ_t and pitch angle θ_t at each time step. We employ an **Exploration-Evaluation (EE)** strategy to determine the next angles ϕ_{t+1} and θ_{t+1} for the virtual drone, and set the search space for NBV planning as candidate sets $\Phi = \{90^\circ, 60^\circ, 30^\circ\}$ and $\Theta = \{\theta_0, 0^\circ, 30^\circ, 60^\circ\}$ for yaw and pitch angles, respectively. At time step t , the virtual drone searches over candidate combinations of $\Delta\phi \in \Phi$ and $\theta \in \Theta$ to select the optimal pose as the next-best view (NBV) $\mathbf{v}_{t+1}$ and navigates to this position. We use $\mathbf{v} = \{\mathbf{c}_t, r_t, \phi_t + \Delta\phi, \theta\}$ to denote the candidate pose that is selected as the NBV, which can be formulated as:

$$\mathbf{v}^* = \arg \max_{\mathbf{v}(\Delta\phi, \theta)} |\Delta\mathcal{A}'(\mathbf{v})| \quad \text{s.t.} \quad \text{mIoU}(I(\mathbf{v}_0, \mathcal{A}^{(t)} - \Delta\mathcal{A}'(\mathbf{v})), m_0) \geq \sigma, \quad (6)$$

where

$$\Delta\mathcal{A}' = \text{MFF}(\text{SAM}(\mathcal{I}(\mathbf{v}, \mathcal{A}^{(t)}), \mathcal{M}_t), \mathcal{A}^{(t)}). \quad (7)$$

Eq. (7) denotes the action of segmentation that finds the change in the foreground Gaussian set $\mathcal{A}$ after performing Mask-shaped Frustum Filtering (MFF) based on the segmentation mask obtained from SAM with memory $\mathcal{M}_t$ at the candidate view $\mathcal{I}(\mathbf{v}, \mathcal{A}^{(t)})$. Eq. (6) formulates the NBV planning as an optimization problem that aims to find the candidate view $\mathbf{v}$ that maximizes the number of Gaussians that can be filtered out as background (i.e., $|\Delta\mathcal{A}'|$) while ensuring that the segmentation result of the candidate view maintains a certain level of consistency with the initial mask m_0 as measured by mIoU (σ denotes the mIoU threshold), which helps to avoid cascading over-segmentation errors along the Markov chain. As illustrated in Fig. 4, starting from $\phi = 0^\circ$, the virtual drone first attempts a large yaw increment of 90° . If SAM successfully tracks the target, it flies directly to $\phi = 90^\circ$ for state update. Otherwise, it falls back to a smaller increment of 60° and finally to 30° . In the case of successful tracking at each step, the virtual drone can achieve full coverage with only two exploration steps, which makes our NBV-based online segmentation efficient.

State Update. After the NBV $\mathbf{v}_{t+1}$ is selected ($\mathbf{v}_{t+1} = \mathbf{v}^*$), the virtual drone navigates to this new position and captures a new view of the scene. The segmentation mask obtained from SAM at this new viewpoint is then used to update the foreground Gaussian set $\mathcal{A}$ and the background Gaussian set $\mathcal{B}$ through MFF. This process effectively refines the segmentation by filtering out additional background Gaussians based on the new view, and updates the state of the system for the next iteration of NBV planning. We refer to this iterative process of updating the segmentation state based on new views as **State Update**, which is crucial for achieving full coverage of the target object through successive NBV explorations. The core variables in state $\mathcal{S}_{t+1}$ can be updated as follows:

$$\mathcal{A}^{(t+1)} = \mathcal{A}^{(t)} - \Delta\mathcal{A}(\mathbf{v}^*), \quad \mathbf{v}_{t+1} = \mathbf{v}^*, \quad \mathcal{M}_{t+1} = \mathcal{M}_t + \mathcal{M}(\mathbf{v}^*, \mathcal{M}_t), \quad (8)$$

where $\mathcal{M}(\mathbf{v}^*, \mathcal{M}_t)$ means the new memory generated by SAM [34] at the new view $\mathbf{v}^*$, and $\Delta\mathcal{A}(\mathbf{v}^*)$ is the change in the foreground Gaussian set after MFF based on the new mask obtained from SAM at the NBV $\mathbf{v}^*$. The coverage indicator $\tau_{t+1} = (\phi_t + \Delta\phi)/180^\circ$ for a single direction virtual drone.

4 Experiments

4.1 Experimental settings

Datasets. To evaluate the effectiveness of our approach, we collect 3D scenes from multiple sources that encompass the majority of existing 3DGS segmentation datasets, including SPIn-NeRF [30], NVOS [35], LLFF [29], 3D-OVS [25], LERF-Mask [15], MipNerf-360 [1], T & T [17], and PKU-DyMVHumans [56]. These datasets serve as the basis for our qualitative analysis. For quantitative comparison, we utilize SPIn-NeRF, NVOS, 3D-OVS, and LERF-Mask.



Fig. 5: Online segmentation results of SAGO on the Figurines scene. The granularity is simply determined by the initial mask obtained via SAM prompted by the user, and the segmentation is iteratively refined through NBV-based online updates.

Implementation Details. We employ SAM2 [34] as our segmentation and tracking agent. The rendering resolution of each virtual drone viewpoint is set to 512×512 , which matches the input resolution used by SAM2. The search space consists of yaw increments $\Delta\phi \in \{90^\circ, 60^\circ, 30^\circ\}$ and pitch angles $\theta \in \{\theta_0, 0^\circ, 30^\circ, 60^\circ\}$. We use Grounding-DINO [26] to transform the text prompt to a bounding box when using text prompts. All experiments are conducted on a single NVIDIA RTX 4090 GPU.

4.2 Qualitative Results

A key advantage of our method is that it does not need a time-consuming setup stage before performing online segmentation. We provide sufficient qualitative results in Fig. 5 and Fig. 6 to demonstrate this advantage and generalization capacity of our method across various scenes and datasets.

Fig. 5 illustrates the online segmentation results of SAGO on the Figurines scene. The granularity of segmentation is determined by the initial mask obtained via SAM prompted by the user, and the segmentation is iteratively refined through NBV-based online updates. Fig. 6 further tests the generalization capacity of SAGO by applying it to ten scenes from five different datasets, performing online segmentation of the main objects without any scene-specific setup. The average segmentation latency per interaction is 500 ms, demonstrating the efficiency and adaptability of our method across diverse scenarios.




Fig. 6: Generalization test of SAGO. We select 10 scenes from 5 datasets [1, 15, 17, 29, 35] and apply SAGO to segment the object described by the user via prompts like clicks, a box and text. The average segmentation latency per interaction is 500 ms.

Table 3: Quantitative comparison on LERF-Mask (mIoU $\uparrow$).

Method	Figurines	Ramen	Teatime	Average
Lerf [15]	33.5	28.3	49.7	37.1
Gau-Group [48]	69.7	77.0	71.7	66.1
OmniSeg3D [49]	87.1	77.3	73.8	79.4
ClickGaussian [6]	93.2	90.9	83.2	89.1
Ours	93.9	90.6	88.6	91.0

Table 4: Quantitative comparison on 3D-OVS dataset (mIoU $\uparrow$).

Method	bed	bench	room	sofa	lawn	mean
LERF [15]	73.5	53.2	46.6	27.0	73.7	54.8
3D-OVS [25]	89.5	89.3	92.8	74.0	88.2	86.8
LangSplat [32]	92.5	94.2	94.1	90.0	96.1	93.4
N2F2 [2]	93.8	92.6	93.5	92.1	96.3	93.9
SAGA [4]	97.4	95.4	96.8	93.5	96.6	96.0
Ours	98.5	96.6	95.7	94.3	95.9	96.2

4.3 Quantitative Comparison

We perform quantitative comparisons on four mainstream datasets (LERF-Mask [15], 3D-OVS [25], SPIn-NeRF [30], and NVOS [35]), with additional efficiency analysis on NVOS and SPIn-NeRF. As shown in Tabs. 2 to 4, by comparing with current state-of-the-art (SOTA) methods, our approach achieves comparable—or even superior—performance in both mIoU and mAcc, without requiring any scene-specific setup.




Fig. 7: Comparison to SAGA [4], the milestone of offline methods.

Method	Type	SPIn-NeRF		NVOS		Time	
		mIoU	mAcc	mIoU	mAcc	Setup	Seg.
MVSeg [30]	Offline	90.4	98.8	-	-	-	-
NVOS [35]	Offline	-	-	70.1	92.0	-	-
ISRF [10]	Offline	71.5	95.5	83.8	96.4	-	-
SA3D [5]	Offline	91.9	98.8	90.3	98.2	2~5m	15~30s
LangSplat [32]	Offline	69.5	94.5	74.0	94.0	~2.5h	-
SAGA [4]	Offline	93.4	99.2	<u>92.6</u>	<u>98.6</u>	~1h	10ms
Flashsplat [37]	Online	-	-	91.8	<u>98.6</u>	30~60s	1~2s
iSegMan [53]	Online	92.4	99.1	92.0	98.4	50s~90s	4~6s
GaussianCut [12]	Online	<u>92.9</u>	<u>99.2</u>	92.5	98.4	N/A	40~60s
Ours	Online	92.5	99.3	92.7	98.7	N/A	0.4~0.9s

Table 2: Comparison of interactive 3D segmentation on SPIn-NeRF [30] and NVOS [35]. First, second, and third best results are highlighted.

Results on SPIn-NeRF and NVOS. As shown in Tab. 2, on the NVOS dataset, our method achieves the best mIoU of 92.7% and mAcc of 98.7%, outperforming all compared methods including both offline and online approaches. Compared to the offline method SAGA [4], which requires approximately 1 hour of setup time for feature distillation, our method improves mIoU by 0.1% and mAcc by 0.1% on NVOS without any setup overhead. On SPIn-NeRF, our method achieves 92.5% mIoU and the highest mAcc of 99.3%, ranking second and first respectively, closely matching the best offline method SAGA (93.4% mIoU). Compared to online methods, our approach surpasses iSegMan [53] by 0.1% mIoU and 0.2% mAcc on SPIn-NeRF, and by 0.7% mIoU and 0.3% mAcc on NVOS, while being significantly faster. Notably, in contrast to the existing setup-free SOTA method GaussianCut [12], our method improves mIoU by 0.2% and mAcc by 0.3% on NVOS, and achieves comparable mIoU with 0.1% higher mAcc on SPIn-NeRF, while reducing segmentation time from 40~60 seconds to 0.4~0.9 seconds, achieving a speedup of over **50**×

Results on LERF-Mask. As shown in Tab. 3, our method achieves the highest average mIoU of 91.0% across the three scenes, surpassing ClickGaussian [6] by 1.9%. In particular, our method obtains the best performance on the Figurines (93.9%) and Teatime (88.6%) scenes, with significant improvements of 5.4% over ClickGaussian on Teatime. On the Ramen scene, our method achieves a competitive 90.6%, closely matching the best result of 90.9%.

Results on 3D-OVS. As shown in Tab. 4, our method achieves the highest mean mIoU of 96.2%, outperforming the previous best method SAGA [4] by




Fig. 8: Online Editing 3DGS via SAGO. SAGO enables editing on a raw 3DGS scene in test-time.

0.2%. Specifically, our approach achieves the best results on three out of five scenes (*bed*: 98.5%, *bench*: 96.6%, *sofa*: 94.3%), and remains highly competitive on the remaining two scenes.

4.4 3D Assets Extraction and Online Editing

The core advantage of our method is its ability to extract 3D assets and perform online editing. Our method extracts user-specified 3D assets from a raw 3DGS scene in the shortest time (less than 1 second), and after the extraction is completed, users can directly perform online editing on these 3D assets. As shown in Fig. 8, we demonstrate the editing results of SAGO on the Figurines scene. After extracting the 3D asset of the leftmost figurine, we perform online editing operations such as relocation, removing, combination, copying on this asset, showcasing our method as a powerful tool for interactive 3D content creation and manipulation in 3DGS scenes.

4.5 Ablation Study

NBV strategy vs. Real views. The motivation behind our virtual drone strategy is to explore the Next-Best-View (NBV) for online segmentation, while the real views are the actual camera views in the original 3DGS scene. To validate the effectiveness of our NBV strategy, we compare it against using only real

NBV	Center-based MFF	mIOU			
		LERF-Mask	NVOS	SPIIn-NeRF	3D-OVS
✓	✓	91.0	92.7	92.5	96.2
✗	✓	36.7 (-54.3↓)	89.1 (-3.6↓)	90.1 (-2.4↓)	91.6 (-4.6↓)
✓	✗	90.4 (-0.6↓)	92.5 (-0.2↓)	91.1 (-1.4↓)	96.0 (-0.2↓)

Table 5: Ablation study on multiple datasets.

views from the original scene. As shown in Tab. 5, removing the NBV strategy (NBV ✗) causes a dramatic performance drop across all datasets, with LERF-Mask (complex and heavy occlusion) dropping by 54.3% mIoU to 36.7%, NVOS by 3.6%, SPIIn-NeRF by 2.4%, and 3D-OVS by 4.6%. This demonstrates that virtual drone viewpoints generated by our NBV strategy are crucial for exploring occluded regions and improving segmentation accuracy beyond what real views alone can provide. The more complex the scene and heavier the occlusion, the more significant the performance drop without NBV, highlighting the importance of our NBV strategy in challenging scenarios.

Splat-based MFF vs. Center-based MFF. Mask-shaped Frustum Filtering (MFF) filters out the background Gaussians based on the 2D mask at each state. We compare our center-based MFF with a splat-based MFF that retains Gaussians whose splats have any overlap with the 2D mask. As shown in Tab. 5, removing the center-based MFF (Center-based MFF ✗) results in a performance drop across all datasets, with LERF-Mask dropping by 0.6% mIoU, NVOS by 0.2%, SPIIn-NeRF by 1.4%, and 3D-OVS by 0.2%. Although the performance drop is not significant, the center-based MFF is computationally simpler, which is why we adopt it in our implementation.

5 Conclusion

In this paper, we present SAGO, a novel setup-free framework for interactive segmentation of 3D Gaussians. Central to our approach is the introduction of **virtual drones**, a key innovation that reformulates the 3D segmentation problem as a multi-step online Next-Best-View (NBV) planning task. By leveraging virtual drones and NBV planning, SAGO effectively eliminates the time-consuming per-scene setup required by existing methods, while maintaining practical interaction times (typically within 1 second).

Limitations. (1) The center-based MFF takes each Gaussian’s center as a representative point. Since Gaussians typically have volume and cannot be treated as ideal points, SAGO currently cannot address jagged boundaries. However, this does not pose a barrier to subsequent work based on SAGO, as online boundary trimming methods such as GaussianTrimmer [24] can be introduced to improve boundary segmentation performance. (2) SAGO requires a high-quality initial

view for effective state initialization, which may restrict its applicability when the initial view is suboptimal or unavailable. Future work could investigate strategies for using v-drones to actively explore an optimal initial view.

Acknowledgements

This work is financially supported by Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology (Grant No. 2024B1212010006). This work is also financially supported for Fundamental and Interdisciplinary Disciplines Breakthrough Plan of the Ministry of Education of China (Grant No. JYB2025XDXM413), Outstanding Talents Training Fund in Shenzhen, Shenzhen Science and Technology Program (Grant No. SYSPG20241211173440004).

References

1. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: CVPR. pp. 5470–5479 (2022) [9](#), [11](#)
2. Bhalgat, Y., Laina, I., Henriques, J.F., Zisserman, A., Vedaldi, A.: N2f2: Hierarchical scene understanding with nested neural feature fields. In: European Conference on Computer Vision. pp. 197–214. Springer (2024) [11](#)
3. Carion, N., Gustafson, L., Hu, Y.T., Debnath, S., Hu, R., Suris, D., Ryali, C., Alwala, K.V., Khedr, H., Huang, A., et al.: Sam 3: Segment anything with concepts. arXiv preprint arXiv:2511.16719 (2025) [2](#)
4. Cen, J., Fang, J., Yang, C., Xie, L., Zhang, X., Shen, W., Tian, Q.: Segment any 3d gaussians. In: Proceedings of the AAAI conference on artificial intelligence. vol. 39, pp. 1971–1979 (2025) [1](#), [3](#), [11](#), [12](#)
5. Cen, J., Zhou, Z., Fang, J., Shen, W., Xie, L., Jiang, D., Zhang, X., Tian, Q., et al.: Segment anything in 3d with nerfs. *Advances in Neural Information Processing Systems* **36**, 25971–25990 (2023) [12](#)
6. Choi, S., Song, H., Kim, J., Kim, T., Do, H.: Click-gaussian: Interactive segmentation to any 3d gaussians. In: ECCV. pp. 289–305. Springer (2024) [1](#), [3](#), [11](#), [12](#)
7. Deng, Y., Wang, Z., Wu, J., Liang, J., Ma, J., Hu, Y., Wang, R.: Pano-gs: Perception-aware gaussian optimization with gradient consistency and multi-criteria densification for high-quality rendering. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 3560–3568 (2026) [1](#)
8. Duan, J., Yu, S., Tan, H.L., Zhu, H., Tan, C.: A survey of embodied ai: From simulators to research tasks. *IEEE Transactions on Emerging Topics in Computational Intelligence* **6**(2), 230–244 (2022) [1](#)
9. Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5501–5510 (2022) [5](#)
10. Goel, R., Sirikonda, D., Saini, S., Narayanan, P.: Interactive segmentation of radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4201–4211 (2023) [12](#)

11. Hu, X., Wang, Y., Fan, L., Fan, J., Peng, J., Lei, Z., Li, Q., Zhang, Z.: Sagd: Boundary-enhanced segment anything in 3d gaussian via gaussian decomposition. arXiv preprint arXiv:2401.17857 (2024) [1](#), [3](#)
12. Jain, U., Mirzaei, A., Gilitschenski, I.: GaussianCut: Interactive segmentation via graph cut for 3d gaussian splatting. *Advances in Neural Information Processing Systems* **37**, 89184–89212 (2024) [1](#), [3](#), [12](#)
13. Jin, R., Gao, Y., Wang, Y., Wu, Y., Lu, H., Xu, C., Gao, F.: Gs-planner: A gaussian-splatting-based planning framework for active high-fidelity reconstruction. In: 2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 11202–11209. IEEE (2024) [2](#)
14. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM TOG* **42**(4), 1–14 (2023) [1](#), [4](#)
15. Kerr, J., Kim, C.M., Goldberg, K., Kanazawa, A., Tancik, M.: Lerf: Language embedded radiance fields. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 19729–19739 (2023) [9](#), [11](#)
16. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: *ICCV*. pp. 4015–4026 (2023) [2](#)
17. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM TOG* **36**(4), 1–13 (2017) [9](#), [11](#)
18. Li, H., Wu, Y., Meng, J., Gao, Q., Zhang, Z., Wang, R., Zhang, J.: Instance-gaussian: Appearance-semantic joint gaussian representation for 3d instance-level perception. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 14078–14088 (2025) [1](#)
19. Li, X., Li, J., Zhang, Z., Zhang, R., Jia, F., Wang, T., Fan, H., Tseng, K.K., Wang, R.: Robogsim: A real2sim2real robotic gaussian splatting simulator. arXiv preprint arXiv:2411.11839 (2024) [1](#)
20. Liang, J., Wang, R., Peng, R., Zhang, Z., Xiong, K., Wang, R.: High fidelity aggregated planar prior assisted patchmatch multi-view stereo. In: *Proceedings of the 32nd ACM International Conference on Multimedia*. pp. 3141–3150 (2024) [1](#)
21. Liang, J., Wu, J., Wang, C., Yang, J., Zheng, X., Xiong, K., Wang, Z., Yan, J., Gao, F., Wang, R.: Clipstream: Clip-stream gaussian splatting for any length and any motion multi-view dynamic scene reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 41022–41032 (June 2026) [1](#)
22. Liao, G., Li, J., Bao, Z., Ye, X., Wang, J., Li, Q., Liu, K.: Clip-gs: Clip-informed gaussian splatting for real-time and view-consistent 3d semantic understanding. arXiv preprint arXiv:2404.14249 (2024) [1](#)
23. Liao, L., Li, X., Zheng, X., Liu, B., Gao, F., Wang, R.: Zero-shot visual grounding in 3d gaussians via view retrieval. In: *ICASSP 2026-2026 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 12692–12696. IEEE (2026) [1](#), [3](#)
24. Liao, L., Wang, R.: Gaussiantrimmer: Online trimming boundaries for 3dgs segmentation. arXiv preprint arXiv:2601.12683 (2026) [1](#), [14](#)
25. Liu, K., Zhan, F., Zhang, J., Xu, M., Yu, Y., Saddik, A.E., Theobalt, C., Xing, E., Lu, S.: Weakly supervised 3d open-vocabulary segmentation. arXiv preprint arXiv:2305.14093 (2023) [9](#), [11](#)
26. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Jiang, Q., Li, C., Yang, J., Su, H., et al.: Grounding dino: Marrying dino with grounded pre-training for open-set object detection. In: *ECCV*. pp. 38–55. Springer (2024) [10](#)

27. Liu, Y., Jia, B., Lu, R., Ni, J., Zhu, S.C., Huang, S.: Building interactable replicas of complex articulated objects via gaussian splatting. In: ICLR (2025) **1**
28. Lu, G., Zhang, S., Wang, Z., Liu, C., Lu, J., Tang, Y.: Manigaussian: Dynamic gaussian splatting for multi-task robotic manipulation. In: European Conference on Computer Vision. pp. 349–366. Springer (2024) **1**
29. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (ToG)* **38**(4), 1–14 (2019) **9, 11**
30. Mirzaei, A., Aumentado-Armstrong, T., Derpanis, K.G., Kelly, J., Brubaker, M.A., Gilitschenski, I., Levinstein, A.: Spin-nerf: Multiview segmentation and perceptual inpainting with neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20669–20679 (2023) **9, 11, 12**
31. Peng, R., Xu, W., Tang, L., Liao, L., Jiao, J., Wang, R.: Structure consistent gaussian splatting with matching prior for few-shot novel view synthesis. *Advances in Neural Information Processing Systems* **37**, 97328–97352 (2024) **1**
32. Qin, M., Li, W., Zhou, J., Wang, H., Pfister, H.: Langsplat: 3d language gaussian splatting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20051–20060 (2024) **1, 3, 11, 12**
33. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PmLR (2021) **3**
34. Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., Mintun, E., Pan, J., Alwala, K.V., Carion, N., Wu, C.Y., Girshick, R., Dollár, P., Feichtenhofer, C.: Sam 2: Segment anything in images and videos. *arXiv preprint arXiv:2408.00714* (2024), <https://arxiv.org/abs/2408.00714> **2, 6, 9, 10**
35. Ren, Z., Agarwala, A., Russell, B., Schwing, A.G., Wang, O.: Neural volumetric object selection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6133–6142 (2022) **9, 11, 12**
36. Shen, H., Ni, J., Chen, Y., Li, W., Pei, M., Huang, S.: Trace3d: Consistent segmentation lifting via gaussian instance tracing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6656–6666 (2025) **3**
37. Shen, Q., Yang, X., Wang, X.: Flashsplat: 2d to 3d gaussian splatting segmentation solved optimally. In: ECCV. pp. 456–472. Springer (2024) **1, 12**
38. Sun, W., Wu, Q., Xu, H., Gao, K., Xu, Z., Chen, Y., Zhang, D., Ma, L., Zelek, J.S., Li, J.: Sagonline: Segment any gaussians online. *arXiv preprint arXiv:2508.08219* (2025) **3**
39. Wang, Z., Wang, Z., Xiong, K., Jiahao, W., Deng, Y., Wang, R.: Sap: Exact sorting in splatting via screen-aligned primitives. *Advances in Neural Information Processing Systems* **38**, 48975–49001 (2026) **1**
40. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20310–20320 (2024) **1**
41. Wu, J., Peng, R., Jiao, J., Yang, J., Tang, L., Xiong, K., Liang, J., Yan, J., Liu, R., Wang, R.: Localdygs: Multi-view global dynamic scene modeling via adaptive local implicit feature decoupling. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 9519–9529 (2025) **1**

42. Wu, J., Peng, R., Wang, Z., Xiao, L., Tang, L., Yan, J., Xiong, K., Wang, R.: Swift4d: Adaptive divide-and-conquer gaussian splatting for compact and efficient reconstruction of dynamic scene. arXiv preprint arXiv:2503.12307 (2025) [1](#)
43. Wu, Y., Meng, J., Li, H., Wu, C., Shi, Y., Cheng, X., Zhao, C., Feng, H., Ding, E., Wang, J., et al.: Opengaussian: Towards point-level 3d gaussian-based open vocabulary understanding. arXiv preprint arXiv:2406.02058 (2024) [1](#)
44. Xiong, K., Peng, R., Wu, J., Wang, Z., Liang, J., Zheng, X., Gao, F., Wang, R.: Intrinsic geometry-appearance consistency optimization for sparse-view gaussian splatting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 40918–40928 (2026) [1](#)
45. Xu, Y., Liao, L., Wang, R.: Nvpose: Novel view data augmentation for human pose estimation. In: 2025 IEEE International Conference on Multimedia and Expo (ICME). pp. 1–6. IEEE (2025) [1](#)
46. Yan, J., Peng, R., Tang, L., Wang, R.: 4d gaussian splatting with scale-aware residual field and adaptive optimization for real-time rendering of temporally complex dynamic scenes. In: Proceedings of the 32nd ACM International Conference on Multimedia. pp. 7871–7880 (2024) [1](#)
47. Yang, J., Tang, L., Wu, J., Liang, J., Gao, F., Wang, R.: i3dv: Intelligent 3d volumetric video coding standard and platform. IEEE Transactions on Circuits and Systems for Video Technology (2026) [1](#)
48. Ye, M., Danelljan, M., Yu, F., Ke, L.: Gaussian grouping: Segment and edit anything in 3d scenes. In: ECCV. pp. 162–179. Springer (2024) [1](#), [3](#), [11](#)
49. Ying, H., Yin, Y., Zhang, J., Wang, F., Yu, T., Huang, R., Fang, L.: Omniseg3d: Omniversal 3d segmentation via hierarchical contrastive learning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20612–20622 (2024) [11](#)
50. Zeng, R., Wen, Y., Zhao, W., Liu, Y.J.: View planning in robot active vision: A survey of systems, algorithms, and applications. Computational Visual Media **6**(3), 225–245 (2020) [2](#)
51. Zhang, J., Jiang, J., Chen, Y., Jiang, K., Liu, X.: Cob-gs: Clear object boundaries in 3dgs segmentation based on boundary-adaptive gaussian splitting. arXiv preprint arXiv:2503.19443 (2025) [1](#), [3](#)
52. Zhang, T., Liu, G., Tian, G.: A novel view planning with joint optimization for efficient 3d building inspection. IEEE Robotics and Automation Letters **11**(2), 1162–1169 (2025) [2](#)
53. Zhao, Y., Xu, W., Zheng, R., Qiao, P., Liu, C., Chen, J.: isegman: Interactive segment-and-manipulate 3d gaussians. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 661–670 (2025) [1](#), [3](#), [4](#), [12](#)
54. Zheng, X., Li, X., Liao, L., Gao, F., Wang, S., Wang, R.: Space-time gaussian surfels for high-fidelity dynamic objects segmentation and representation. IEEE Transactions on Circuits and Systems for Video Technology (2025) [1](#)
55. Zheng, X., Liao, L., Jiao, J., Gao, F., Wang, R.: Surface-sos: Self-supervised object segmentation via neural surface representation. IEEE Transactions on Image Processing **33**, 2018–2031 (2024) [1](#)
56. Zheng, X., Liao, L., Li, X., Jiao, J., Wang, R., Gao, F., Wang, S., Wang, R.: Pku-dymvhumans: A multi-view video benchmark for high-fidelity dynamic human modeling. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22530–22540 (2024) [9](#)
57. Zhu, R., Qiu, S., Liu, Z., Hui, K.H., Wu, Q., Heng, P.A., Fu, C.W.: Rethinking end-to-end 2d to 3d scene segmentation in gaussian splatting. In: Proceedings

- of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3656–3665 (2025) [3](#)
58. Zhu, S., Wang, G., Kong, X., Kong, D., Wang, H.: 3d gaussian splatting in robotics: A survey. arXiv preprint arXiv:2410.12262 (2024) [1](#)