# ActiveGS: Active Scene Reconstruction Using Gaussian Splatting

Liren Jin<sup>1</sup>, Xingguang Zhong<sup>1</sup>, Yue Pan<sup>1</sup>, Jens Behley<sup>1</sup>, Cyrill Stachniss<sup>1,3</sup> and Marija Popović<sup>2</sup>



Fig 1: Our approach actively reconstructs an unknown scene. We illustrate the reconstruction progress over mission time, displaying planned camera viewpoints (green pyramids) and paths (red lines). We present examples of RGB and confidence maps (redder colour indicates lower confidence) rendered at the same viewpoint at different mission times (magenta and cyan arrows) and at two distinct viewpoints at the same mission time (cyan and blue arrows). By integrating confidence modelling into the Gaussian splatting pipeline, our approach enables targeted view planning to build a high-fidelity Gaussian splatting map. The complete camera path and final reconstruction results, including RGB rendering and surface mesh, are visualised on the right.

Abstract—Robotics applications often rely on scene reconstructions to enable downstream tasks. In this work, we tackle the challenge of actively building an accurate map of an unknown scene using an RGB-D camera on a mobile platform. We propose a hybrid map representation that combines a Gaussian splatting map with a coarse voxel map, leveraging the strengths of both representations: the high-fidelity scene reconstruction capabilities of Gaussian splatting and the spatial modelling strengths of the voxel map. At the core of our framework is an effective confidence modelling technique for the Gaussian splatting map to identify under-reconstructed areas, while utilising spatial information from the voxel map to target unexplored areas and assist in collision-free path planning. By actively collecting scene information in under-reconstructed and unexplored areas for map updates, our approach achieves superior Gaussian splatting reconstruction results compared to state-of-the-art approaches. Additionally, we demonstrate the real-world applicability of our framework using an unmanned aerial vehicle.

#### Index Terms—Mapping; RGB-D Perception

Manuscript received: December 20, 2024; Revised: February 23, 2025; Accepted: March 20, 2025.

This paper was recommended for publication by Editor Javier Civera upon evaluation of the Associate Editor and Reviewers' comments. This work has been fully funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy, EXC-2070 – 390732324 (PhenoRob).

<sup>1</sup>L. Jin, X. Zhong, Y. Pan, J. Behley, C. Stachniss are with the Center for Robotics, University of Bonn, Germany. <sup>2</sup>M. Popović is with the MAVLab, TU Delft, the Netherlands. <sup>3</sup>C. Stachniss is also with the Lamarr Institute for Machine Learning and Artificial Intelligence, Germany.

Digital Object Identifier (DOI): see top of this page.

# I. INTRODUCTION

CTIVE exploration and reconstruction of unknown scenes are relevant capabilities for developing fully autonomous robots [3]. For scenarios, including search and rescue, agricultural robotics, and industrial inspection, online active reconstruction using mobile robots demands both mission efficiency and reconstruction quality. Achieving this requires two key components: high-fidelity map representations for modelling fine-grained geometric and textural details of the scenes and adaptive view planning strategies for effective sensor data acquisition.

In this work, we tackle the problem of actively reconstructing unknown scenes using posed RGB-D camera data. Given a limited budget, e.g. mission time, our goal is to obtain an accurate 3D representation of the scene by actively positioning the robot with its camera online during a mission. Existing active scene reconstruction frameworks mainly rely on conventional map representations such as voxel grids, meshes, or point clouds [2, 16, 29, 30, 31, 38, 39]. However, these methods often do not deliver high-fidelity reconstruction results due to their sparse representations. Recent advances in implicit neural representations, e.g. neural radiance fields (NeRFs) [20], have attracted significant research interest for their accurate dense scene reconstruction capabilities and low memory footprints. In the context of active reconstruction, several emerging works [5, 12, 13, 24, 26, 36] incorporate uncertainty estimation in NeRFs and exploit it to guide view

planning. While these approaches demonstrate promising results, the rather costly volumetric rendering procedure during online incremental mapping poses limitations for NeRF-based active scene reconstruction.

Dense reconstruction using Gaussian splatting (GS) [15] offers a promising alternative to NeRFs, addressing rendering inefficiencies while preserving representation capabilities. GS explicitly models scene properties through Gaussian primitives and utilises efficient differentiable rasterisation to achieve dense reconstruction. Its fast map updates and explicit structure make it well-suited for online incremental mapping. Building on these strengths, we adopt GS for active scene reconstruction in this work.

Incorporating GS into an active scene reconstruction pipeline presents significant challenges. First, active reconstruction often requires evaluating the reconstruction quality to guide view planning. However, this is difficult without ground truth information at novel viewpoints. Second, the Gaussian primitives represent only occupied space, making it hard to distinguish between unknown and free space, which are important for exploration and path planning.

The key contribution of this work is addressing these challenges through our novel GS-based active scene reconstruction framework, denoted as ActiveGS. To tackle the first challenge, we propose a simple yet effective confidence modelling technique for Gaussian primitives based on viewpoint distribution, enabling view planning for inspecting under-reconstructed surfaces. For the second challenge, we combine the GS map with a conventional coarse voxel map, exploiting the strong representation capabilities of GS for scene reconstruction with the spatial modelling strengths of voxel maps for exploration and path planning.

We make the following three claims: (i) our ActiveGS framework achieves superior reconstruction performance compared to state-of-the-art NeRF-based approach and GS-based baselines; (ii) our explicit confidence modelling for Gaussian primitives enables informative viewpoint evaluation and targeted inspection around under-reconstructed surfaces, further improving mission efficiency and reconstruction quality; (iii) we validate our approach in different synthetic indoor scenes and in a real-world scenario using an unmanned aerial vehicle (UAV). To support reproducibility and future research, we open-source our implementation code at: https://github.com/dmar-bonn/active-gs.

### II. RELATED WORK

Our work uses Gaussian splatting as the map representation for active scene reconstruction. In this section, we overview state-of-the-art high-fidelity map representations, focusing on GS and active scene reconstruction methods.

#### A. Gaussian Splatting as Map Representation

Conventional map representations such as voxel grids [8], meshes [1, 33], and point clouds [38] often only capture coarse scene structures, struggling to provide fine-grained geometric and textural information crucial for many robotics applications [3]. Recent advances in implicit neural representations, such as NeRFs [20, 21], show promising results in highfidelity scene reconstruction by modelling scene attributes continuously. Although they achieve impressive reconstruction results, NeRFs require dense sampling along rays for view synthesis, a computationally intensive process that limits their online applicability.

3D GS [15] offers an efficient alternative for high-fidelity scene reconstruction by combing explicit map structures with volumetric rendering. Unlike NeRFs, 3D GS stores scene information using explicit 3D Gaussian primitives, eliminating the need for inefficient dense sampling during volumetric rendering. This explicit nature also makes it well-suited for online incremental mapping, which requires frequent data fusion and scene attributes modification. Follow-up works further enhance geometric quality by regularising 3D GS training [6] or directly adopting 2D GS for improved surface alignment [4, 9]. 2D GS collapses the 3D primitive volume into 2D oriented planar Gaussian primitives, enabling more accurate depth estimation and allowing the integration of normal information into the optimisation process. Motivated by its strong performance, we utilise 2D GS, specifically Gaussian surfel [4], as our GS map representation for active scene reconstruction.

#### B. Active Scene Reconstruction

Active scene reconstruction using autonomous mobile robots is an area of active research [3]. Given an unknown scene, the goal is to explore and map the scene by actively planning the robot's next viewpoints for effective data acquisition. Existing active scene reconstructions utilise map representations such as voxel maps [2, 10, 16, 29, 39], meshes [30, 31], or point clouds [38?]. These approaches primarily focus on fully covering the unknown space, rather than achieving high-fidelity reconstruction. However, highfidelity scene reconstruction is crucial for downstream robotic tasks that rely on accurate map information.

To address this, recent research explores implicit neural representations for active reconstruction applications. In an object-centric setup, several methods incorporate uncertainty estimation into NeRFs [12, 13, 24, 26, 36] and use this information to select next best viewpoints. For scene-level reconstruction, Yan et al. [37] and Kuang et al. [17] investigate the loss landscape of implicit neural representations during training to identify under-reconstructed areas. NARUTO [5] learns an uncertainty grid map alongside a hybrid neural scene representation, guiding data acquisition in uncertain regions. These implicit neural representations often face challenges such as inefficient map updates and catastrophic forgetting during incremental mapping.

Several concurrent works propose using GS for active scene reconstruction. GS-Planner [14] and HGS-planner [34] incorporate unknown voxels into the GS rendering pipeline and detect unseen regions for exploration. Li et al. [18] use a Voronoi graph to extract a traversable topological map from the GS representation for path planning. The approach is designed for a 2D planning space, reducing its effectiveness in cluttered environments. FisherRF [11] evaluates the information content



Fig 2: An overview of our framework. Our hybrid map representation consists of a 2D GS map for high-fidelity scene reconstruction with a coarse voxel map for exploration and path planning. Our view planner leverages unexplored regions in the voxel map for exploration and low-confidence Gaussian primitives for targeted inspection, collecting informative measurements at planned viewpoints for map updates. We iterate between the map update and view planning steps until the pre-allocated mission time is reached.

of a novel view by measuring the Fisher information value in the GS parameters. This procedure requires computationally expensive gradient calculations at each previously visited and candidate viewpoint, making view planning inefficient for online missions. We build upon the idea of using GS for active scene reconstruction, while introducing a key innovation: we explicitly model the confidence of each Gaussian primitive, enabling viewpoint sampling around low-confidence Gaussian primitives for targeted inspection and fast feed-forward confidence rendering for efficient viewpoint evaluation.

# III. OUR APPROACH

We introduce ActiveGS, a novel framework for active scene reconstruction using GS for autonomous robotic tasks. An overview of our framework is shown in Fig. 2. Our goal is to reconstruct an unknown scene using a mobile robot, e.g. a UAV, equipped with an on-board RGB-D camera. Given posed RGB-D measurements as input, we update a coarse voxel map to model the spatial occupancy and incrementally train a GS map for high-fidelity scene reconstruction. To actively guide view planning to reconstruct the scene in a targeted manner, we propose using our confidence modelling technique in the GS map and information about unexplored regions in the voxel map as the basis for planning. The framework alternates between mapping and planning steps until a predefined mission time is reached.

#### A. Hybrid Map Representation

Given the bounding box of the scene to be reconstructed, we uniformly divide the enclosed space into voxels, forming our voxel map  $\mathcal{V}$ , where each voxel  $v_i \in \mathcal{V}$  represents the volume occupancy probability in the range [0, 1].

Our GS map is based on Gaussian surfel [4], a state-ofthe-art 2D GS representation. The GS map  $\mathcal{G}$  comprises a collection of Gaussian primitives. Each primitive  $g_i \in \mathcal{G}$  is defined by its parameters  $g_i = (x_i, q_i, s_i, c_i, o_i, k_i)$ , where  $x_i \in \mathbb{R}^3$  denotes the position of the primitive centre;  $q_i \in \mathbb{R}^4$ is its rotation in the form of a quaternion;  $s_i = [s_i^x, s_i^y] \in \mathbb{R}^2_+$  are the scaling factors along the two axes of the primitive;  $c_i \in [0,1]^3$  represents the RGB colour;  $o_i \in [0,1]$  is the opacity; and  $k_i \in \mathbb{R}_+$  is the confidence score introduced later in Sec. III-C. The distribution of the Gaussian primitive  $g_i$  in world coordinate is represented as:

$$\mathcal{N}(\boldsymbol{x}; \boldsymbol{x}_i, \boldsymbol{\Sigma}_i) = \exp\left(-\frac{1}{2}(\boldsymbol{x} - \boldsymbol{x}_i)^\top \boldsymbol{\Sigma}_i^{-1}(\boldsymbol{x} - \boldsymbol{x}_i)\right),$$
 (1)

where  $\Sigma_i = \mathbf{R}(\mathbf{q}_i) \operatorname{diag}((s_i^x)^2, (s_i^y)^2, 0) \mathbf{R}(\mathbf{q}_i)^\top$  is the covariance matrix, with  $\mathbf{R}(\mathbf{q}_i) \in SO(3)$  as the rotation matrix derived from the corresponding quaternion  $\mathbf{q}_i$  and  $\operatorname{diag}(\cdot)$ indicating a diagonal matrix with the specified diagonal elements. The normal of the Gaussian primitive can be directly obtained from the last column of the rotation matrix as  $\mathbf{n}_i = \mathbf{R}(\mathbf{q}_i)_{:,3}$ .

Given the GS map, we can render the colour map I, depth map D, normal map N, opacity map O, and confidence map K at a viewpoint using the differentiable rasterisation pipeline [4]. Without loss of generality, the rendering function for a pixel u on the view is formulated as:

$$\mathbf{O}(\boldsymbol{u}) = \sum_{i=1}^{n} w_i, \mathbf{M}(\boldsymbol{u}) = \sum_{i=1}^{n} w_i m_i, \qquad (2)$$

$$w_i = T_i \alpha_i, T_i = \prod_{j < i} (1 - \alpha_j), \alpha_i = \mathcal{N}(\boldsymbol{u}; \boldsymbol{u}_i, \boldsymbol{\Sigma}'_i) o_i, \quad (3)$$

where  $\mathbf{M} \in {\{\mathbf{I}, \mathbf{D}, \mathbf{N}, \mathbf{K}\}}$  and  $m_i \in {\{c_i, d_i, n_i, k_i\}}$  is the corresponding modality feature, with  $d_i$  being the distance from the viewpoint centre to the intersection point of the camera ray and the Gaussian primitive  $g_i$ ;  $w_i$  indicates the rendering contribution of  $g_i$  to pixel u; and  $\Sigma'_i$  and  $u_i$  are the primitive's covariance matrix and centre projected onto the image space [40]; For more technical details about the rendering process, please refer to Gaussian surfel [4].

# B. Incremental Mapping

We collect measurements captured at planned viewpoints and incrementally update our map representation. Given the current RGB image I<sup>\*</sup> and depth image D<sup>\*</sup> measurements, we generate a per-pixel point cloud using known camera parameters. We then update our voxel map  $\mathcal{V}$  probabilistically based on the new point cloud, following OctopMap [8].

For GS map update, we first add Gaussian primitives to  $\mathcal{G}$  where needed. To this end, we render the colour map I, depth map D, and opacity map O at the current camera viewpoint. We calculate a binary mask to identify the pixels that should be considered for densifying the GS map:

$$\mathbf{B}(\boldsymbol{u}) = (\mathbf{O}(\boldsymbol{u}) < 0.5) \lor (\operatorname{avg}(|\mathbf{I}(\boldsymbol{u}) - \mathbf{I}^{\star}(\boldsymbol{u})|) > 0.5) \lor (\mathbf{D}(\boldsymbol{u}) - \mathbf{D}^{\star}(\boldsymbol{u}) > \lambda \mathbf{D}^{\star}(\boldsymbol{u})), \qquad (4)$$

where  $avg(\cdot)$  is the channel-wise averaging operation to calculate per-pixel colour error and  $\lambda$  is a constant accounting for depth sensing noise, set to 0.05 in our pipeline. This mask indicates areas where opacity is low, colour rendering is inaccurate, or new geometry appears in front of the current depth estimate, signalling the need for new Gaussian primitives. We spawn new Gaussian primitives by unprojecting pixels on these

areas into 3D space, with initial parameters defined by the corresponding point cloud position, pixel colour, and normal estimated by applying central differencing on the bilateral-filtered depth image [22], which helps mitigate noise contained in the depth sensing. We also set scale values to 1 cm, opacity value to 0.5, and confidence value to 0.

At each mapping step, we train our GS map  $\mathcal{G}$  using all collected RGB-D measurements for 10 iterations. Specifically, for each iteration, we select the 3 most recent frames and 5 random frames from the measurement history. The loss for a frame  $\{\hat{\mathbf{I}}, \hat{\mathbf{D}}\}$  in the training batch is formulated as the weighted sum of individual loss terms:

$$\mathcal{L} = w_c \mathcal{L}_c + w_d \mathcal{L}_d + w_n \mathcal{L}_n \,, \tag{5}$$

where the photometric loss  $\mathcal{L}_c = L_1(\mathbf{I}, \hat{\mathbf{I}})$  and the depth loss  $\mathcal{L}_d = L_1(\mathbf{D}, \hat{\mathbf{D}})$  are both calculated using the  $L_1$  distance. The normal loss  $\mathcal{L}_n = D_{cos}(\mathbf{N}, \widetilde{\mathbf{N}}) + TV(\mathbf{N})$  consists of the cosine distance  $D_{cos}$  between the rendered normal map and the normal map  $\widetilde{\mathbf{N}}$  derived from the rendered depth map [4], along with the total variation TV loss [27] to enforce smooth normal rendering between neighbouring pixels. Note that the training process involves only a subset of the Gaussian primitive parameters  $(\boldsymbol{x}_i, \boldsymbol{q}_i, \boldsymbol{s}_i, \boldsymbol{c}_i, o_i)$ , while the modelling of non-trainable  $k_i$  is introduced in Sec. III-C.

After every 5 mapping steps, we perform a visibility check on all Gaussian primitives and delete those invisible to all history views to compact the GS map. We consider a Gaussian primitive visible from a viewpoint if at least one pixel rendered in that view receives its rendering contribution greater than a threshold, as defined in Eq. (3). Unlike previous works utilising density control during offline training [4, 9, 15], our approach adds necessary primitives and removes invisible ones during online missions, achieving computationally efficient scene reconstruction.

#### C. Confidence Modelling for Gaussian Primitives

A Gaussian primitive can be effectively optimised if observed from different viewpoints. Based on this insight, we derive the confidence of a Gaussian primitive from the spatial distribution of viewpoints in the measurement history. Specifically, we connect the Gaussian centre  $x_i$  to the viewpoint centre  $x_{p_j}$ , denoted as  $d_{ij} = x_{p_j} - x_i = d_{ij}v_{ij}$ , where  $d_{ij}$ is the distance and  $v_{ij}$  is the normalised view direction, with  $j \in S(g_i)$  and S being the index set of viewpoints from which the Gaussian primitive  $g_i$  has been observed. The confidence  $k_i$  is finally formulated as:

$$k_i = \gamma_i \exp(\beta_i), \tag{6}$$

$$\gamma_i = \sum_{j \in \mathcal{S}(\boldsymbol{g}_i)} \left( 1 - \frac{d_{ij}}{d_{\text{far}}} \right) \boldsymbol{n}_i \cdot \boldsymbol{v}_{ij} \,, \tag{7}$$

$$\beta_i = 1 - \|\boldsymbol{\mu}_i\|, \boldsymbol{\mu}_i = \frac{1}{|\mathcal{S}(\boldsymbol{g}_i)|} \sum_{j \in \mathcal{S}(\boldsymbol{g}_i)} \boldsymbol{v}_{ij}, \qquad (8)$$

where  $\gamma_i$  accounts for distance-weighted cosine similarity between the Gaussian primitive's normal  $n_i$  and view direction  $v_{ij}$ , with  $d_{\text{far}}$  as the maximum depth sensing range. Note that we increase the impact of viewpoints that are closer to the Gaussian primitive's centre or provide view directions similar to the primitive's normal.  $\beta_i$  measures the dispersion of directions from which  $g_i$  is observed, with  $\beta_i$  closer to 0 indicating similar view directions. Our confidence formulation assigns higher confidence to Gaussian primitives densely observed from viewpoints with varying angles, whereas lower confidence to those with sparse and similar observations.

#### D. Viewpoint Utility Formulation

Active scene reconstruction requires both exploration, to cover unknown areas, and exploitation, to closely inspect under-reconstructed surfaces. In this work, we combine utility derived from the voxel map for exploration and the GS map for exploitation, enabling these behaviours effectively.

A candidate viewpoint  $p_i^c \in \mathbb{R}^5$  is defined by its 3D position, yaw, and pitch angles in our framework. To simplify path planning, we constrain the positions to a discrete lattice placed at the centres of all free voxels. We follow existing active scene exploration frameworks [2, 10, 23, 29, 39] and define the exploration utility based on the number of unexplored voxels visible from a candidate viewpoint. Without relying on time-consuming ray-casting operations, our framework leverages efficient rendering of our GS map to identify visible voxels. We achieve this by checking whether the projected depth of the in-view voxel centres in the camera coordinate is smaller than the corresponding depth value in the rendered depth from the GS map.

Combining unexplored region information in the voxel map and confidence rendering from the GS map, we define the utility of a candidate viewpoint  $p_i^c$  as:

$$U_{\text{view}}(\boldsymbol{p}_{i}^{c}) = \phi \, U_{\mathcal{V}}(\boldsymbol{p}_{i}^{c}) + U_{\mathcal{G}}(\boldsymbol{p}_{i}^{c}) \,, \tag{9}$$

where  $U_{\mathcal{V}}(\boldsymbol{p}_i^c) = \frac{N_u(\boldsymbol{p}_i^c)}{|\mathcal{V}|}$  is the exploration utility, defined by the ratio of the number of visible unexplored voxels  $N_u(\boldsymbol{p}_i^c)$  to the total number of voxels in the voxel map;  $U_{\mathcal{G}}(\boldsymbol{p}_i^c) = -\text{mean}(\mathbf{K}_i)$  is the exploitation utility, calculated as the negative mean of the confidence map  $\mathbf{K}_i$  rendered at  $\boldsymbol{p}_i^c$  following Eq. (2); and  $\phi$  is the exploration weight.

### E. Viewpoint Sampling and Evaluation

Our viewpoint sampling strategy involves two types of candidate viewpoints. First, we randomly sample  $N_{random}$ candidate viewpoints within a specified range around the current viewpoint. However, relying solely on random local sampling can lead to local minima. To address this, we introduce additional candidate viewpoints based on regions of interest (ROI) defined in the voxel map. We begin by identifying frontier voxels [35] and add them to our ROI set  $\mathcal{R}$ . By explicitly modelling the confidence of each Gaussian primitive, we can identify and also include voxels containing low-confidence Gaussian primitives in  $\mathcal{R}$ . Inspired by previous work [16], we define normals for each voxel in  $\mathcal{R}$  to indicate the most informative viewing direction. For voxels with lowconfidence Gaussian primitives, this is simply the average normal of these Gaussian primitives. The normal of frontier voxels is determined by finding their neighbouring free voxels



Fig 3: We show a 2D case of our ROI-based candidate viewpoint generation. We identify ROI voxels by retrieving voxels containing low-confidence Gaussian primitives and frontier voxels. Normals for low-confidence voxels are generated by averaging the normals of low-confidence primitives, while frontier voxel normals are calculated using average vectors to neighbouring free voxels. Given the voxel centres and directional normals, we generate candidate viewpoints within the sampling region, as illustrated on the right.

and calculating the average directional vector from the frontier voxel to these neighbours. To generate ROI-based candidate viewpoints, we create a fixed number of candidate viewpoints within the cone defined by the minimum and maximum sampling distances to each ROI's centre, and the maximum angular difference relative to its normal. Starting from the closest ROI voxel, we continue outward until we have collected up to  $N_{\rm ROI}$  viewpoints in free space. We illustrate the sampling process in Fig. 3.

We evaluate the utility of all candidate viewpoints following Eq. (9). We use the  $A^*$  algorithm [7] to find the shortest traversable path from the current viewpoint position to all candidate viewpoint positions. Taking travel distance into account, we select the next best viewpoint  $p^*$  by:

$$\boldsymbol{p}^{\star} = \operatorname*{arg\,max}_{\boldsymbol{p}_{i}^{c}} \left( \frac{U_{\operatorname{view}}(\boldsymbol{p}_{i}^{c})}{\sum_{i=1}^{N_{\operatorname{total}}} U_{\operatorname{view}}(\boldsymbol{p}_{i}^{c})} - \delta \frac{U_{\operatorname{path}}(\boldsymbol{p}_{i}^{c})}{\sum_{i=1}^{N_{\operatorname{total}}} U_{\operatorname{path}}(\boldsymbol{p}_{i}^{c})} \right),$$
(10)

where  $N_{\text{total}} = N_{\text{random}} + N_{\text{ROI}}$ ;  $U_{\text{path}}$  is the travel distance to the candidate viewpoint positions; and  $\delta$  is a weighting factor for the travel cost.

# IV. EXPERIMENTAL EVALUATION

Our experimental results support our three claims: (i) we show that our ActiveGS framework outperforms state-of-theart NeRF-based and GS-based active scene reconstruction methods; (ii) we show that our confidence modelling of Gaussian primitives enables informative viewpoint evaluation and targeted candidate viewpoint generation, improving reconstruction performance; (iii) we validate our framework in different simulation scenes and in a real-world scenario to show its applicability.

# A. Implementation Details

**Mapping.** We use a voxel size of  $20 \text{ cm} \times 20 \text{ cm} \times 20 \text{ cm}$  for the voxel map. We set the loss weights in Eq. (5) as:  $w_c = 1.0, w_d = 0.8$ , and  $w_n = 0.1$ . For visibility checks, a minimum rendering contribution threshold of 0.3 is applied.

**Planning**. We set the exploration weight  $\phi = 1000$ in Eq. (9) to encourage exploratory behaviour during the initial phase of an online mission. We set  $\delta = 0.5$  for weighting travel costs in Eq. (10). We consider  $N_{\text{total}} = 100$  candidate viewpoints, including up to  $N_{\text{ROI}} = 30$  ROI-based samples and  $N_{\text{random}} = N_{\text{total}} - N_{\text{ROI}}$  random samples generated within 0.5 m of the current viewpoint position.

We test our implementation on a desktop PC with an Intel Core i9-10940X CPU and an NVIDIA RTX A5000 GPU. In this setup, one mapping and planning steps take approximately 1 s and 0.5 s, respectively. The whole framework consumes 4-5 GB GPU RAM during an online mission, with approximately 10% allocated to the voxel map update.

#### **B.** Simulation Experiments

Setup. We conduct our simulation experiments using the Habitat simulator [28] and the Replica dataset [32]. The experiments utilise an RGB-D camera with a field of view of  $[60^\circ, 60^\circ]$  and a resolution of [512, 512] pixels. The camera has a depth sensing range of [0.1, 5.0] m and Gaussian noise in the depth measurements with linearly increased standard deviation  $\sigma = 0.01d$ , where d represents the depth value.

**Evaluation Metrics.** We report the reconstruction performance over total mission time, defined as the summation of mapping time, planning time, and action time, assuming a constant robot velocity of 1 m/s. The reconstruction performance is evaluated on both rendering and mesh quality. For the rendering evaluation, we generate ground truth RGB images captured from 1000 uniformly distributed test viewpoints in the scene's free space. We report PSNR [20] of RGB images rendered from our GS map at test viewpoints as the rendering quality metric. For the mesh evaluation, we run TSDF fusion [22] on depth images rendered at training viewpoints and extract the scene mesh using Marching Cubes [19]. We use the completeness ratio [5] as the mesh quality metric with a distance threshold set to 2 cm.

We consider the following methods:

- Ours: Our full ActiveGS framework utilising both exploration and exploitation utility measures. We consider ROI-based sampling to achieve targeted candidate viewpoint generation as described in Sec. III-E.
- Ours (w/o ROI): A variant of our ActiveGS that leverages only local random sampling, with N<sub>ROI</sub> = 0.
- Ours<sup>†</sup>: A variant of our ActiveGS with an alternative confidence formulation, assigning higher confidence to Gaussian primitives with more visible viewpoints, without considering their spatial distribution.
- *FBE* [35]: Frontier-based exploration framework that solely focuses on covering unexplored regions, without accounting for the quality of the GS map. We use the collected RGB-D data to update the GS map, similar to our framework.



Fig 4: We report the reconstruction performance evaluated in rendering and mesh quality over online mission time. Our ActiveGS outperforms baselines in all test scenes. Our view planner considers unexplored regions for exploration, while exploiting low-confidence Gaussian primitives for further inspection. Compared to GS-based approaches, our approach proposes explicit confidence modelling of Gaussian primitives, enabling targeted candidate viewpoint generation and fast viewpoint evaluation. Our approaches demonstrate a large performance gain to the state-of-the-art NeRF-based approach, *NARUTO*, motivating the use of GS in active scene reconstruction.

- *FisherRF* [11]: GS-based active scene reconstruction using frontier voxels for ROI-based candidate viewpoint generation and Fisher information for viewpoint evaluation. We replace its 3D GS map with our 2D GS.
- *NARUTO* [5]: A state-of-the-art NeRF-based active scene reconstruction pipeline.

We run 5 trials for all methods across 8 test scenes. We set the maximum mission time to 300 s and evaluate reconstruction performance every 60 s. We report the mean and standard deviation for PSNR and completeness ratio.

We present the results of simulation experiments in Fig. 4. Our approach achieves the best performance in both rendering and mesh quality across all test scenes, supporting our first claim that it outperforms state-of-the-art NeRF and GSbased methods. The NeRF-based active scene reconstruction framework, *NARUTO*, exhibits a significant performance gap compared to our approach, particularly in RGB rendering. This disparity arises because NeRF-based methods often compromise model capacity for faster map updates, limiting their representation quality in scene-level reconstruction. *FisherRF* evaluates viewpoint utility by calculating the Fisher information in the parameters of the Gaussian primitives within its field of view. This requires computationally expensive gradient calculation for all candidate and training viewpoints, leading to prolonged planning times and incomplete reconstruction under limited mission time. Additionally, since Fisher information is conditioned on the candidate viewpoint, the viewpoint must be selected before its utility can be evaluated, preventing direct viewpoint sampling informed by Fisher information. In contrast, our approach models the confidence of each Gaussian



Fig 5: Visual comparison of reconstruction results using different approaches. We show RGB rendering and surface meshes for two scenes, with red circles highlighting areas of low-quality reconstruction from baseline approaches. Our ActiveGS considers both unexplored regions in voxel map and confidence value of GS map to enable targeted view planning, achieving complete and high-fidelity scene reconstruction.

primitive, enabling fast feed-forward confidence rendering for viewpoint evaluation and identification of low-confidence surfaces for targeted candidate viewpoint generation, significantly enhancing reconstruction quality and efficiency. FBE focuses solely on exploration and ignores surface reconstruction quality, limiting its performance. While our approach balances exploration and exploitation by accounting for both unexplored regions and low-confidence Gaussian primitives. The ablation study comparing Ours and Ours (w/o ROI) demonstrates the benefits of ROI-based sampling for targeted inspection, reflected by higher means and smaller standard deviations in both evaluation metrics. Our confidence formulation also outperforms the variant in Ours<sup>†</sup> by considering viewpoint distribution. These results confirm that our confidence modelling is effective in achieving efficient and high-fidelity active scene reconstruction, validating our second claim. We visually compare the reconstruction results in Fig. 5.

# C. Real-World Experiments

We demonstrate the applicability our framework in a realworld experiment using a UAV equipped with an Intel RealSense 455 RGB-D camera to reconstruct a scene of size  $6 \text{ m} \times 6 \text{ m} \times 3 \text{ m}$ . Unlike simulation experiments, we do not account for the pitch angle of viewpoints in this experiment due to control limitations. The UAV pose is tracked by an OptiTrack motion capture system. Given the limited on-board resources, we run ActiveGS on our desktop PC, where it receives RGB-D and pose data from the UAV for map updates and sends planned collision-free waypoints to guide the UAV. All communication is handled via ROS [25].

Our real-world experiments indicate that our approach is effective for actively reconstructing unknown scenes by



Fig 6: Our real-world experiments using a UAV equipped with an RGB-D camera. We show the experimental setup (left) and the RGB rendering from our GS map (right).

considering both unexplored regions in the voxel map and under-reconstructed surfaces in the GS map. We show the experimental setup in Fig. 6 and the online active scene reconstruction in the supplementary video.

# V. CONCLUSIONS AND FUTURE WORK

In this paper, we propose ActiveGS, a GS-based active scene reconstruction framework. Our approach employs a hybrid map representation, combining the high-fidelity scene reconstruction capabilities of Gaussian splatting with the spatial modelling strengths of the voxel map. We propose an effective method for confidence modelling of Gaussian primitives, enabling targeted viewpoint generation and informative viewpoint evaluation. Our view planning strategy leverages the confidence information of Gaussian primitives to inspect under-reconstructed areas, while also considering unexplored regions in the voxel map for exploration. Experimental results demonstrate that ActiveGS outperforms baseline approaches in both rendering and mesh quality. A limitation of our current framework is the assumption of perfect localisation. Future work will incorporate localisation uncertainty into confidence modelling of Gaussian primitives. We also plan to better integrate the voxel map with GS map for more efficient mapping and leverage optimisation-based approaches to enhance view planning quality.

#### VI. ACKNOWLEDGMENT

We thank Hang Yu, Jakub Plonka, and Moji Shi for their assistance in conducting the real-world experiments.

# REFERENCES

- J. Behley and C. Stachniss, "Efficient surfel-based slam using 3d laser range data in urban environments," in *Proc. of Robotics: Science and Systems*, 2018.
- [2] A. Bircher, M. Kamel, K. Alexis, H. Oleynikova, and R. Siegwart, "Receding horizon "next-best-view" planner for 3d exploration," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation*, 2016.
- [3] S. Chen, Y. Li, and N. M. Kwok, "Active vision in robotic systems: A survey of recent developments," *Intl. Journal of Robotics Research*, vol. 30, no. 11, pp. 1343–1377, 2011.
- [4] P. Dai, J. Xu, W. Xie, X. Liu, H. Wang, and W. Xu, "High-quality surface reconstruction using gaussian surfels," in *Proc. of the Intl. Conf. on Computer Graphics and Interactive Techniques*, 2024.
- [5] Z. Feng, H. Zhan, Z. Chen, Q. Yan, X. Xu, C. Cai, B. Li, Q. Zhu, and Y. Xu, "Naruto: Neural active reconstruction from uncertain target observations," in *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, 2024.
- [6] A. Guédon and V. Lepetit, "Sugar: Surface-aligned gaussian splatting for efficient 3d mesh reconstruction and high-quality mesh rendering," in Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition, 2024.
- [7] P. E. Hart, N. J. Nilsson, and B. Raphael, "A formal basis for the heuristic determination of minimum cost paths," *IEEE Trans. on Systems Science* and Cybernetics, vol. 4, no. 2, pp. 100–107, 1968.
- [8] A. Hornung, K. M. Wurm, M. Bennewitz, C. Stachniss, and W. Burgard, "Octomap: An efficient probabilistic 3d mapping framework based on octrees," *Autonomous Robots*, vol. 34, pp. 189–206, 2013.
- [9] B. Huang, Z. Yu, A. Chen, A. Geiger, and S. Gao, "2d gaussian splatting for geometrically accurate radiance fields," in *Proc. of the Intl. Conf. on Computer Graphics and Interactive Techniques*, 2024.
- [10] S. Isler, R. Sabzevari, J. Delmerico, and D. Scaramuzza, "An information gain formulation for active volumetric 3d reconstruction," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation*, 2016.
- [11] W. Jiang, B. Lei, and K. Daniilidis, "Fisherrf: Active view selection and mapping with radiance fields using fisher information," in *Proc. of the Europ. Conf. on Computer Vision*, 2024.
- [12] L. Jin, X. Chen, J. Rückin, and M. Popović, "Neu-nbv: Next best view planning using uncertainty estimation in image-based neural rendering," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2023.
- [13] L. Jin, H. Kuang, Y. Pan, C. Stachniss, and M. Popović, "Stair: Semantictargeted active implicit reconstruction," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2024.
- [14] R. Jin, Y. Gao, H. Lu, and F. Gao, "Gs-planner: A gaussian-splattingbased planning framework for active high-fidelity reconstruction," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2024.
- [15] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, "3d gaussian splatting for real-time radiance field rendering," ACM Trans. on Graphics, vol. 42, no. 4, pp. 139–1, 2023.
- [16] Y. Kompis, L. Bartolomei, R. Mascaro, L. Teixeira, and M. Chli, "Informed sampling exploration path planner for 3d reconstruction of large scenes," *IEEE Robotics and Automation Letters*, vol. 6, no. 4, pp. 7893–7900, 2021.
- [17] Z. Kuang, Z. Yan, H. Zhao, G. Zhou, and H. Zha, "Active neural mapping at scale," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2024.
- [18] Y. Li, Z. Kuang, T. Li, G. Zhou, S. Zhang, and Z. Yan, "Activesplat: High-fidelity scene reconstruction through active gaussian splatting," *arXiv preprint arXiv:2410.21955*, 2024.

- [19] W. E. Lorensen and H. E. Cline, "Marching cubes: A high resolution 3d surface construction algorithm," in *Seminal Graphics: Pioneering Efforts that Shaped the Field*, 1998, pp. 347–353.
- [20] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," in *Proc. of the Europ. Conf. on Computer Vision*, 2020.
- [21] T. Müller, A. Evans, C. Schied, and A. Keller, "Instant neural graphics primitives with a multiresolution hash encoding," ACM Trans. on Graphics, vol. 41, no. 4, pp. 1–15, 2022.
- [22] R. A. Newcombe, S. Izadi, O. Hilliges, D. Molyneaux, D. Kim, A. J. Davison, P. Kohli, J. Shotton, S. Hodges, and A. Fitzgibbon, "KinectFusion: Real-Time Dense Surface Mapping and Tracking," in *Proc. of the Intl. Symp. on Mixed and Augmented Reality*, 2011.
- [23] E. Palazzolo and C. Stachniss, "Effective Exploration for MAVs Based on the Expected Information Gain," *Drones*, vol. 2, no. 1, pp. 59–66, 2018.
- [24] X. Pan, Z. Lai, S. Song, and G. Huang, "Activenerf: Learning where to see with uncertainty estimation," in *Proc. of the Europ. Conf. on Computer Vision*, 2022.
- [25] M. Quigley, K. Conley, B. Gerkey, J. Faust, T. Foote, J. Leibs, R. Wheeler, and A. Y. Ng, "Ros: an open-source robot operating system," in *ICRA Workshop on Open Source Software*, 2009.
- [26] Y. Ran, J. Zeng, S. He, J. Chen, L. Li, Y. Chen, G. Lee, and Q. Ye, "Neurar: Neural uncertainty for autonomous 3d reconstruction with implicit neural representations," *IEEE Robotics and Automation Letters*, vol. 8, no. 2, pp. 1125–1132, 2023.
- [27] L. Rudin and S. Osher, "Total variation based image restoration with free local constraints," in *Proc. of the IEEE Intl. Conf. on Image Processing*, 1994.
- [28] M. Savva, A. Kadian, O. Maksymets, Y. Zhao, E. Wijmans, B. Jain, J. Straub, J. Liu, V. Koltun, J. Malik *et al.*, "Habitat: A platform for embodied ai research," in *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision*, 2019.
- [29] L. Schmid, M. Pantic, R. Khanna, L. Ott, R. Siegwart, and J. Nieto, "An efficient sampling-based method for online informative path planning in unknown environments," *IEEE Robotics and Automation Letters*, vol. 5, no. 2, pp. 1500–1507, 2020.
- [30] S. Song and S. Jo, "Online inspection path planning for autonomous 3d modeling using a micro-aerial vehicle," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation*, 2017.
- [31] —, "Surface-based exploration for autonomous 3d modeling," in *Proc. of the IEEE Intl. Conf. on Robotics & Automation*, 2018.
- [32] J. Straub, T. Whelan, L. Ma, Y. Chen, E. Wijmans, S. Green, J. J. Engel, R. Mur-Artal, C. Ren, S. Verma, A. Clarkson, M. Yan, B. Budge, Y. Yan, X. Pan, J. Yon, Y. Zou, K. Leon, N. Carter, J. Briales, T. Gillingham, E. Mueggler, L. Pesqueira, M. Savva, D. Batra, H. M. Strasdat, R. D. Nardi, M. Goesele, S. Lovegrove, and R. Newcombe, "The Replica dataset: A digital replica of indoor spaces," *arXiv preprint arXiv:1906.05797*, 2019.
- [33] T. Whelan, S. Leutenegger, R. F. Salas-Moreno, B. Glocker, and A. J. Davison, "Elasticfusion: Dense slam without a pose graph." in *Proc. of Robotics: Science and Systems*, 2015.
- [34] Z. Xu, R. Jin, K. Wu, Y. Zhao, Z. Zhang, J. Zhao, F. Gao, Z. Gan, and W. Ding, "Hgs-planner: Hierarchical planning framework for active scene reconstruction using 3d gaussian splatting," *arXiv preprint* arXiv:2409.17624, 2024.
- [35] B. Yamauchi, "A frontier-based approach for autonomous exploration," in Proc. of the IEEE Intl. Symp. on Computational Intelligence in Robotics and Automation, 1997.
- [36] D. Yan, J. Liu, F. Quan, H. Chen, and M. Fu, "Active implicit object reconstruction using uncertainty-guided next-best-view optimization," *IEEE Robotics and Automation Letters*, vol. 8, no. 10, pp. 6395–6402, 2023.
- [37] Z. Yan, H. Yang, and H. Zha, "Active neural mapping," in *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision*, 2023.
- [38] R. Zeng, W. Zhao, and Y.-J. Liu, "Pc-nbv: A point cloud based deep network for efficient next best view planning," in *Proc. of the IEEE/RSJ Intl. Conf. on Intelligent Robots and Systems*, 2020.
- [39] B. Zhou, Y. Zhang, X. Chen, and S. Shen, "Fuel: Fast uav exploration using incremental frontier structure and hierarchical planning," *IEEE Robotics and Automation Letters*, vol. 6, no. 2, pp. 779–786, 2021.
- [40] M. Zwicker, H. Pfister, J. Van Baar, and M. Gross, "Ewa volume splatting," in *Proc. of Visualization*, 2001.

# CERTIFICATE OF REPRODUCIBILITY

The authors of this publication declare that:

- 1) The software related to this publication is distributed in the hope that it will be useful, support open research, and simplify the reproducability of the results but it comes without any warranty and without even the implied warranty of merchantability or fitness for a particular purpose.
- 2) *Liren Jin* primarily developed the implementation related to this paper. This was done on Ubuntu20.04.
- 3) *Yue Pan* verified that the code can be executed on a machine that follows the software specification given in the Git repository available at:

https://github.com/dmar-bonn/active-gs

4) *Yue Pan* verified that the experimental results presented in this publication can be reproduced using the implementation used at submission, which is labeled with a tag in the Git repository and can be retrieved using the command:

git checkout ral2025