This CVPR paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

ActiveGAMER: Active GAussian Mapping through Efficient Rendering

Liyan Chen^{*†,1,2} Huangying Zhan^{*‡1} Kevin Chen¹ Xiangyu Xu¹ Qingan Yan¹ Changjiang Cai¹ Yi Xu¹

¹ OPPO US Research Center

² Stevens Institute of Technology

Abstract

We introduce ActiveGAMER, an active mapping system that utilizes 3D Gaussian Splatting (3DGS) to achieve highquality scene mapping and efficient exploration. Unlike recent NeRF-based methods, which are computationally demanding and limit mapping performance, our approach leverages the efficient rendering capabilities of 3DGS to enable effective and efficient exploration in complex environments. The core of our system is a rendering-based information gain module that identifies the most informative viewpoints for next-best-view planning, enhancing both geometric and photometric reconstruction accuracy. ActiveG-AMER also integrates a carefully balanced framework, combining coarse-to-fine exploration, post-refinement, and a global-local keyframe selection strategy to maximize reconstruction completeness and fidelity. Our system autonomously explores and reconstructs environments with state-of-the-art geometric and photometric accuracy and completeness, significantly surpassing existing approaches Extensive evaluations on benchmark in both aspects. datasets such as Replica and MP3D highlight ActiveG-AMER's effectiveness in active mapping tasks.

1. Introduction

In computer vision, the ability to generate detailed 3D reconstructions from 2D images or videos has seen tremendous progress, enabling real-time, incremental 3D modeling as new visual data is assimilated. This process, often powered by Simultaneous Localization and Mapping (SLAM), plays a crucial role in robotic applications, where it supports tasks such as planning and navigation. When combined, these functions define Active SLAM, a framework that integrates localization, mapping, planning, and navigation to enable autonomous exploration.

This paper focuses on a subproblem of Active SLAM



Figure 1. *ActiveGAMER* is built on a Gaussian Map backbone and autonomously performs coarse-to-fine exploration to optimize geometric accuracy and photometric fidelity.

known as Active Reconstruction, where localization is assumed to be known, allowing the system to prioritize highquality and complete scene reconstruction. We explore this task by introducing a novel approach based on 3D Gaussian Splatting (3DGS), which provides an efficient radiance field representation optimized for active exploration.

Radiance fields, particularly implicit models like NeRFs, have shown promising results in applications such as 3D object reconstruction [45], novel view synthesis [40, 49, 71, 76], and surface reconstruction [2, 32]. However, NeRF's high computational demands hinder real-time applicability, especially in active vision tasks. To address this, recent efforts have explored hybrid neural representations, which incorporate both implicit and explicit components to enhance rendering efficiency [41, 62].

On the other hand, 3D Gaussian Splatting [26] is a recent and efficient radiance field representation that leverages Gaussian primitives in 3D space. Unlike NeRFs, which rely on dense, computationally intensive sampling, 3DGS uses a sparse set of Gaussian ellipsoids to approximate both geometry and color information. This representation allows for

^{*}Equal contribution

[†]Work done as an intern at OPPO US Research Center

[‡]Corresponding author (zhanhuangying.work@gmail.com)

rapid rendering by projecting each Gaussian onto the image plane, where they are blended using alpha compositing. Due to its efficient design, 3DGS achieves real-time rendering speeds, making it particularly suitable for dynamic applications like SLAM and active vision, where computational resources and time are limited. Recent work [22, 72] has shown that 3DGS maintains high visual fidelity while reducing the computational overhead, enabling its use in scenarios that demand fast, high-quality 3D reconstructions.

Despite these advances, integrating radiance fields into active vision remains challenging. While several studies have explored active reconstruction and path planning with NeRFs [1, 31, 44, 51, 75], these methods often focusing on geometric reconstruction tasks, especially surface reconstruction. Improving photometric reconstruction (rendering) is often ignored with the use of NeRF-based methods due to its slow rendering speeds.

To address existing limitations and enhance photometric reconstruction, we introduce *ActiveGAMER*, an active mapping system that leverages the real-time rendering capabilities of 3D Gaussian Splatting. Our system enables unrestricted 6DoF movement, allowing flexible exploration and high-quality scene reconstruction even in complex environments. The key contributions of our work include:

- An advanced active mapping system based on 3DGS, allowing real-time, unrestricted 6DoF exploration.
- A rendering-based information gain module that efficiently identifies the most informative viewpoints for next-best-view planning.
- A carefully designed system that balances exploration efficiency and reconstruction accuracy for both geometric and photometric reconstruction, through strategies like coarse-to-fine exploration, post-refinement, and globallocal keyframe selection.
- Achieves state-of-the-art performance in active reconstruction, enhancing both geometric accuracy and rendering fidelity over existing approaches.

2. Related Work

Active Reconstruction Autonomous robotics relies on foundational capabilities such as localization, mapping, planning, and motion control [57]. These capabilities have spurred advancements in various areas, including visual odometry [53, 74], monocular depth estimation [3, 16, 73], multi-view stereo [6, 9, 21, 35, 56, 63, 70], structure-frommotion (SfM) [54], and path planning [20, 29, 30]. SLAM systems have also evolved significantly, facilitating simultaneous localization and mapping for autonomous navigation [5, 13, 15, 64, 66]. Active SLAM combines these approaches to enable autonomous localization, mapping, and planning, aiming to reduce uncertainty in the robot's environment representation [12]. Comprehensive reviews of active SLAM can be found in recent survey papers

[5, 36, 48]. Our focus is on active reconstruction, a research area closely related to active SLAM but oriented towards achieving complete and accurate 3D representations. Active reconstruction is often formulated as an exploration problem, where the objective is to determine the most informative viewpoints for capturing detailed scene representations [4, 17, 37, 42, 58, 59, 65]. This approach has been widely applied in scene and object reconstruction from multiple viewpoints, with methods developed to handle occlusions, optimize viewpoint selection, and maximize information gain [11, 14, 23, 27, 39, 46, 47].

Radiance Fields Radiance Field representations have become a cornerstone in 3D scene modeling, enabling continuous and high-quality representations of complex environments. Neural Radiance Fields (NeRFs) [40] pioneered this field by modeling scenes as continuous radiance fields using multi-layer perceptrons (MLPs), achieving impressive results in applications such as novel view rendering [40, 49, 71, 76], object and surface reconstruction [2, 32, 40, 45], and generative modeling [43, 55], and Structure-from-Motion [10, 34, 67]. While NeRFs have set a high standard for rendering quality, they are computationally intensive, making real-time applications challenging.

Recent advancements in 3D Gaussian Splatting [26] provide an alternative approach by representing scenes through a set of Gaussian primitives in 3D space, which can be rendered efficiently with lower computational overhead. 3D Gaussian Splatting has shown promise in achieving real-time rendering, making it suitable for dynamic applications like SLAM [25, 38]. Together, NeRFs and 3D Gaussian Splatting represent complementary approaches to scene representation, each with unique strengths and tradeoffs in terms of rendering quality, computational efficiency, and suitability for active vision tasks.

Active Radiance Fields Building on advancements in NeRF and 3D Gaussian Splatting, recent research has investigated the potential of these representations in active vision, particularly for autonomous scene exploration and mapping. Active radiance fields leverage the strengths of NeRF and 3DGS for tasks requiring continuous exploration and decision-making in 3D environments.

NeRF-based approaches have been applied to path planning [1] and active reconstruction via next-best-view strategies [31, 44, 51], though they are often limited by NeRF's high computational demands, which restricts real-time applications [28]. To overcome these limitations, hybrid models such as ActiveRMAP [75] integrate implicit and explicit representations for improved efficiency. However, many NeRF-based methods still restrict camera motion to constrainted spaces, reducing flexibility in complex 3D environments. NARUTO [18] addresses this by introducing an



Figure 2. ActiveGAMER Framework. At each keyframe step, HabitatSim [52] generates posed RGB-D images, which are stored in a keyframe database, with certain frames designated as Global Keyframes. These observations are used to update a Gaussian Map comprising a collection of 3D Gaussians. Map optimization is achieved by minimizing color and depth rendering losses, based on the rendered RGB-D images and silhouette masks. Using the up-to-date Gaussian Map, rendering-based planning evaluates the information gain across sampled candidate viewpoints and choose the one with highest information gain as the next-best-view.

active neural mapping system with 6DoF movement in unrestricted spaces, while [28] integrates Voronoi planning to scale exploration across larger environments.

3DGS offers a faster alternative, making real-time mapping and exploration more feasible. Concurrent works like ActiveSplat [33] utilize a hybrid map with topological abstractions for efficient planning, and AG-SLAM [24] incorporates 3DGS with Fisher Information to balance exploration and localization in complex environments. In this work, we propose an Active Gaussian Mapping method that leverages the real-time rendering capability of 3DGS to enable effective planning and exploration.

3. ActiveGAMER: Active Gaussian Mapping

In this section, we present ActiveGAMER (Fig. 2), a pioneering 3D Gaussian Splatting framework for active reconstruction that incorporates rendering-based planning. Our approach begins with the Gaussian Splatting mapping module, an efficient representation for real-time, highfidelity geometric and photometric reconstruction. We utilize SplaTAM [25] as the mapping backbone, as detailed in Sec. 3.1, establishing a foundation for dense reconstruction using Gaussian Maps. Building upon this, Sec. 3.2 introduces our rendering-based planning module for goaldirected searching and path planning. To address the limitations of the map update process in SplaTAM, we propose an enhanced update strategy in Sec. 3.3, incorporating a global-local keyframe selection strategy. This module leverages rendering-based information, ensuring seamless integration into existing incremental 3DGS to improve mapping robustness. We then present a post-refinement module that further improves the photometric reconstruction in Sec. 3.4. Finally, we conclude this section with an overview of the Active Gaussian Mapping process in Sec. 3.5.

3.1. Gaussian Mapping

Gaussian Splatting Recent advancements have established 3D Gaussian Splatting as both expressive and efficient. These representations effectively encode a scene's appearance and 3D geometry as a collection of 3D Gaussians, $\mathbf{G} = {\mathbf{G}_i}_{i=1}^N$, enabling real-time rendering into high-fidelity color and depth images. A series of prior works, including [25, 38], have demonstrated the effectiveness of 3DGS in 3D reconstruction. Given a stream of RGB-D images, dense mapping is done by optimizing the 3D Gaussian representation through rendering supervision.

Rather than using the comprehensive 3DGS representation originally proposed in [26], we adopt the simplified Gaussian Mapping suggested in SplaTAM [25]. This simplified approach utilizes only view-independent color and isotropic Gaussians, reducing the number of parameters needed for each Gaussian. The parameters for each Gaussian include color **c**, center position μ , radius *r*, and opacity *o*. A 3D point **x** is influenced by all Gaussians based on a standard Gaussian function weighted by its opacity:

$$f(\mathbf{x}) = o \exp\left(-\frac{|\mathbf{x} - \boldsymbol{\mu}|^2}{2r^2}\right).$$
 (1)

Real-time Rendering A key strength of 3DGS is its realtime rendering capability, enabling high-fidelity color and depth image generation from arbitrary camera poses. Leveraging Gaussian Maps, 3DGS can efficiently render scenes by projecting 3D Gaussians into 2D pixel space. Following the approach proposed in [26], we achieve efficient rendering by transforming all 3D Gaussians to camera space, sorting them front-to-back, and projecting them onto the image plane. Each Gaussian is then splatted in 2D, where color and transparency are composited using alpha-blending. The color at a pixel, $\mathbf{p} = (u, v)$, is defined as:

$$C(\mathbf{p}) = \sum_{i=1}^{n} c_i f_i(\mathbf{p}) \prod_{j=1}^{i-1} (1 - f_j(\mathbf{p})), \qquad (2)$$

where $f_i(\mathbf{p})$ is derived from the Gaussian's position and size in 2D pixel space, computed as:

$$\boldsymbol{\mu}^{\text{2D}} = K \frac{T_t \boldsymbol{\mu}}{d}, \quad r^{\text{2D}} = \frac{fr}{d}, \quad \text{where} \quad d = (T_t \boldsymbol{\mu})_z.$$
 (3)

Here, K represents the camera intrinsics, including focal length f and principal point, while T_t encodes the camera's extrinsics, capturing its rotation and translation in world space at time t. The variable d represents the depth of the Gaussian in camera coordinates.

Similarly, the depth map is rendered as:

$$D(\mathbf{p}) = \sum_{i=1}^{n} d_i f_i(\mathbf{p}) \prod_{j=1}^{i-1} (1 - f_j(\mathbf{p})).$$
(4)

We also generate a silhouette mask to indicate if a pixel contains information from the Gaussian map:

$$S(\mathbf{p}) = \sum_{i=1}^{n} f_i(\mathbf{p}) \prod_{j=1}^{i-1} (1 - f_j(\mathbf{p})).$$
 (5)

Optimization The differentiable nature of this rendering process enables end-to-end optimization, where gradients are calculated directly from the discrepancy between rendered images and RGB-D inputs. These gradients drive updates to each Gaussian's parameters through minimization of the following loss:

$$L = \sum_{\mathbf{p}} \left(S(\mathbf{p}) > 0.99 \right) \left(L_1(D(\mathbf{p})) + 0.5L_1(C(\mathbf{p})) \right).$$
(6)

By thresholding the silhouette mask, we selectively optimize pixels in regions with high-quality visibility, enhancing the stability of the reconstruction.

Gaussian Densification Gaussian Densification is designed to adaptively add new 3D Gaussians in response to incoming data. Using the known camera pose and depth measurements, we compute a densification mask to determine where additional Gaussians are needed, avoiding redundant creation in regions already well-represented by the current model. The densification mask is defined as:

$$M(\mathbf{p}) = (S(\mathbf{p}) < 0.5) + (D_{\text{GT}}(\mathbf{p}) < D(\mathbf{p})) (L_1(D(\mathbf{p})) > \lambda \text{MDE}).$$
(7)

This mask identifies areas where: (1) the density of Gaussians is low (S < 0.5), and (2) new Gaussians are required to refine the geometry, as indicated by a depth error exceeding a threshold of $\lambda = 50$ times the median depth error (MDE).

3.2. Rendering-based Planning

In Sec. 3.1, we introduce a Gaussian Mapping method using known camera parameters and incremental RGB-D observations. Typically, mapping is conducted in a passive manner, where the capture trajectory is controlled manually. However, passive capture does not ensure a complete, highfidelity reconstruction of the scene. In this section, we propose an active mapping system that leverages the Gaussian Map, specifically exploiting its efficient rendering capabilities—advantages that are unattainable with neural radiance fields [18, 69]. We propose a rendering-based information gain to achieve active exploration.

Our approach involves a two-stage active mapping strategy. First, a coarse-to-fine exploration stage reconstructs the scene as comprehensively as possible. Subsequently, a post-refinement stage further enhances the Gaussian Mapping representation for rendering purpose once the exploration phase is complete.

In the exploration stage, our goal is to incrementally determine the next-best-view (NBV) to enhance the completeness of the Gaussian Map. Leveraging the real-time rendering capability of the Gaussian Map, we can extensively generate NBV candidates across the environment and evaluate the information gain for each candidate. We first present our exploration information gain formulation, followed by a mechanism for efficiently maintaining a candidate pool to enable rapid information gain evaluation. After that, we present a coarse-to-fine exploration strategy that further accelerates exploration efficiency. Finally, we introduce a local path planner introduced in [18].

Exploration Information Gain Given a candidate camera pose, we compute its exploration information gain, \mathcal{I} , based on the rendered silhouette mask S with respect to the up-to-date Gaussian Map. The number of missing pixels in the rendered silhouette mask, N_{S_i} , quantifies the information gain for the candidate viewpoint. We further incorporate a motion cost represented by the L_2 distance between the current location $T_{t,x}$ and the candidate location $T_{i,x}$, aiming to reduce overall travel distance during exploration. That is, we prefer a candidate closer to the current camera pose when two candidates yield the same information gain. The distance-weighted information gain is formulated as:

$$\mathcal{I} = (1 - \sigma(l_i)) \cdot \sigma(\log(N_{S_i})), \tag{8}$$

where $\sigma(\cdot)$ is the Softmax function, which normalizes both the distance weighting and N_{S_i} in a relative form; $l_i = ||T_{i,x} - T_{t,x}||_2$ represents the travel distance, and $N_{S_i} = \sum_{\mathbf{p}} \mathbb{I}(S_i(\mathbf{p}) = 0)$ counts the pixels in the silhouette mask that are zero. We select the candidate viewpoint with the highest \mathcal{I} as the goal pose T_g . **Exploration Candidate Pool** While Gaussian Maps provide real-time rendering capabilities, extensively sampling viewpoints within the scene and evaluating their information gain remains computationally expensive. To manage this, we maintain an Exploration Candidate Pool, allowing us to add new sampled candidate viewpoints, update their associated information gain, and remove them from the pool when necessary.

Adding candidates from the current observation requires careful consideration, as redundant candidates with overlapping view frustums can occur. To ensure uniform spatial sampling, we incrementally update an occupancy grid for each incoming observation. We refer to the free space within this grid as the Exploration Map, which we use to sample exploration candidates. Rather than sampling the entire exploration map each time, we identify newly added free space voxels by comparing the latest Exploration Map to its previous state and sample candidates only from these new voxels. Candidate positions $T_{e,x}$ are generated from every v_1 meter, with evenly distributed v_2 viewing directions based on the Fibonacci lattice. Once candidate poses T_e are added to the Exploration Pool, we evaluate all candidates using Eq. (8) and update their information gain, \mathcal{I}_i . For candidates with associated N_{S_i} values below 0.5% of the total pixel count, we consider the viewpoint to be well observed and remove the candidate from the pool to prevent its overpopulation.

Coarse-to-fine Exploration Oversampling candidates enhances scene coverage and optimizes next-best-view selection but increases evaluation time. To improve exploration efficiency, we introduce a coarse-to-fine strategy, which initially covers the environment with minimal candidates and then refines the reconstruction with finer exploration. In the coarse stage, we sample new free voxels at specific heights, such as a single 2D plane, using larger spatial steps $(v_1 = 1)$ and fewer viewing directions $(v_2 = 5)$. The fine stage increases sampling density, employing multiple height levels, smaller steps ($v_1 = 0.5$), and more viewing directions ($v_2 = 15$) to refine the reconstruction. Initially, sampling is restricted to newly added free space; at the start of the fine stage, we sample the entire free space, quickly discarding redundant candidates as most regions have already been observed. This strategy effectively balances exploration speed and comprehensive scene coverage.

Local Path Planning Once the goal pose is identified, our path planning module initiates to generate a viable path from the current pose, T_t , to the goal pose, T_g . For this, we use a sampling-based path planning approach similar to Rapid-exploration Random Tree (RRT) [30], with the Exploration Map as the basis. Executing standard RRT in a large-scale 3D environment can be highly time-intensive.

Algorithm 1 ActiveGAMER

 Initialization Camera Pose T₀; Gaussian Map G₀; Exploration Map M_{e,0}; Observations {O}⁰_{i=0}; Exploration Pool P_e; PLAN_REQUIRED = True; t = 0
 # EXPLORATION

3: while $P_e \neq \emptyset \lor t = 0$ do

- 4: t = t + 1
- 5: # Update Database in keyframe steps
- 6: **Observation**: acquire a new observation O_t
- 7: Update Database: $\{O\}_{i=0}^t \leftarrow \{O\}_{i=0}^{t-1}$
- 8: # Update Mapping Models
- 9: **Map Update:** $\mathbf{G}_t \leftarrow \mathbf{G}_{t-1}, \mathbf{M}_{e,t} \leftarrow \mathbf{M}_{e,t-1},$
- 10: # Update Exploration Pool
- 11: $P_e \leftarrow \text{PoolUpdate}(\mathbf{M}_{e,t}, \mathbf{M}_{e,t-1})$
- 12: **if PLAN_REQUIRED then**
- 13: # Search a new goal from the maps
- 14: **GoalSearch**(\mathbf{G}_t, P_e, T_t) $\rightarrow T_g$
- 15: # Plan a feasible path based on $\mathbf{M}_{e,t}$ towards T_g
- 16: **PathPlanning** $(\mathbf{M}_{e,t}, T_t, T_g) \rightarrow \{T_j\}_{j=t}^g$
- 17: # Set **PLAN_REQUIRED** to False
- 18: **PLAN_REQUIRED** \leftarrow False
- 19: **end if**
- 20: # Update pose from planned path
- 21: Action $T_t \leftarrow \{T_j\}_{j=t}^g$
- 22: # Replanning after reached goal
- 23: CheckPlanRequired: update PLAN_REQUIRED
- 24: end while
- 25: $t_1 = t$ 26: # POST-REFINEMENT
- 27: for $t \leftarrow t_1$ to $t_1 + T$ do
- 28: # Post-refinement based on $\{O\}_{i=0}^{t_1}$
- 29: Map Update: $\mathbf{G}_t \leftarrow \mathbf{G}_{t-1}$
- 30: end for

To address this, we adopt the efficient RRT implementation proposed by [18], and we generate a collision-free path connecting the current and goal locations by sampling within the free space. Additionally, we implement a rotation planning module to ensure a smooth transition from the current orientation to the goal orientation.

3.3. Global-Local Keyframe Selection

Following SplaTAM [25], we design our Gaussian Mapping backbone as detailed in Sec. 3.1. Instead of optimizing with every frame, SplaTAM selects every fifth frame as a keyframe and updates the map using a subset of keyframes that overlaps with the current frame. This approach ensures that the map is optimized with keyframes that influence newly added Gaussians, providing sufficient multiview supervision without processing all prior keyframes.

The k selected local keyframes include the current frame, the last keyframe, and k - 2 keyframes with the highest overlap with the current frame. Overlap is determined by projecting the current frame's depth map into a point cloud and counting points within each keyframe's frustum. This local keyframe strategy yields accurate scene representation, especially in recently updated regions. However, it also causes potential overfitting in local areas. Gaussians that do not contribute directly to rendering may have their opacity minimized, particularly those within the view frustum but behind the primary surface.

To mitigate this issue, we propose a global-local keyframe selection strategy that balances local supervision with global regularization. Global keyframes are selected based on the information they provide, determined by two criteria: (1) Completeness: A keyframe is included if it reveals over 10% new pixels in the silhouette mask. (2) Quality: A keyframe is included if its rendering quality falls below a threshold. Type-1 global keyframes primarily support comprehensive scene reconstruction, while Type-2 keyframes focus on challenging regions. In conclusion, we select half of the keyframes from local overlapping frames and the other half from global supporting keyframes to maintain an effective balance of local detail and global structural accuracy.

3.4. Post-Refinement

The exploration stage focuses on achieving completeness in both geometric (3D) and photometric (rendering) reconstruction. However, rendering quality may be suboptimal due to limited optimization of the Gaussian Map. To address this, we introduce a post-refinement step, leveraging the global keyframes from Sec. 3.3 to enhance photometric reconstruction quality.

We note, however, that post-refinement can degrade geometric reconstruction, as Gaussian Map optimizations for rendering involve operations that are unfavorable for preserving geometric detail. Specifically, optimizing the Gaussian Map for photometric quality may result in Gaussian pruning in low-texture areas and removal of redundant Gaussians, reducing geometric completeness. Despite these trade-offs, we observe overall improvements in rendering quality with post-refinement.

3.5. ActiveGAMER System

After capturing new RGB-D frames, a selection of keyframes is stored in keyframe a database to optimize mapping. Integrating the mapping module from Sec. 3.1 and Sec. 3.3 with the planning module in Sec. 3.2, we establish a comprehensive Active Gaussian Mapping system, namely *ActiveGAMER*, detailed in Algorithm 1 and illustrated in Fig. 2. This system maintains an up-to-date Gaussian Map in an incremental manner, leveraging the planning module for goal searching and path planning.

Methods	Acc. (cm) \downarrow	Comp. (cm) \downarrow	Comp. Ratio (%) ↑
FBE [68]	/	9.78	71.18
UPEN [19]	/	10.60	69.06
OccAnt [50]	/	9.40	71.72
ANM [69]	7.80	9.11	73.15
NARUTO[18]	6.31	3.00	90.18
Ours	1.66	2.30	95.32

Table 1. **MP3D Results.** Our method shows superior performance with better reconstruction quality and completeness.

4. Experiments and Results

4.1. Experimental Setup

Simulator and Dataset Our experiments are conducted in the Habitat simulator [52] and evaluated on two photorealistic datasets: Replica [60] and Matterport3D (MP3D) [7]. We use 8 scenes from Replica [61] and 5 scenes from MP3D [69] for analysis. Each experiment runs for 2000 steps in Replica and 5000 steps in MP3D, with the extended steps in MP3D reflecting its larger scene sizes and need for thorough exploration.

In these experiments, our system processes posed RGB-D images at a resolution of 680×1200 , with vertical and horizontal fields of view set at 60° and 90° , respectively. The voxel size for generating the Exploration Map is fixed at 5 cm across all experiments.

Unlike prior active radiance field methods, which often constrain actions to teleporting between discrete locations [44, 51], moving within hemispheres [75], or navigating limited 2D planes [8, 69], our approach enables 6DoF movement in unrestricted 3D spaces.

Geometric Metrics We evaluate geometric reconstruction using three metrics: *Accuracy* (cm), *Completion* (cm), and *Completion ratio* (%) with a 5 cm threshold. To calculate these metrics, we evenly sample 3D points from the ground-truth meshes and compare them with a point cloud uniformly extracted from the Gaussian Maps.

Rendering Metrics For measuring RGB rendering performance we use PSNR, SSIM and LPIPS. For depth rendering performance we use Depth L1 distance.

4.2. Evaluation

To the best of our knowledge, this work is among the *first* to address active mapping using 3D Gaussian Splatting. Concurrently, recent technical reports [24, 33] explore related concepts, highlighting the growing interest in this area.

As discussed in Sec. 3.4, the Refined Model enhances rendering performance but reduces geometric completeness by removing redundant Gaussians. To accurately assess both aspects, we evaluate geometric performance using the Exploration Model and rendering performance using the Refinement Model. The choice of model depends on the



Figure 3. **3D Reconstruction Results on MP3D.** Shown are two scenes (Left: pLe4; Right: HxpK) with results distinguished by border colors: [Ground Truth, NARUTO [18], Ours]. For NARUTO, black regions indicate neural mapping extrapolation, sometimes with inaccuracies. For our method, colored point clouds are extracted from the Gaussian Map for visualization; note that noisy points here do not reflect actual rendering quality. Unlike neural maps, the explicit Gaussian Map representation avoids inaccurate extrapolation artifacts.

Methods	Metrics	Avg.	Of0	Of1	Of2	Of3	Of4	R0	R1	R2
SplaTAM [25]	PSNR ↑	29.08	30.15	33.60	26.40	23.97	29.87	31.13	29.74	27.79
	SSIM ↑	0.95	0.96	0.96	0.94	0.90	0.96	0.97	0.96	0.96
	LPIPS \downarrow	0.14	0.13	0.16	0.16	0.21	0.13	0.09	0.12	0.11
	L1-D↓	1.38	1.34	0.89	1.75	3.28	1.26	0.64	0.84	1.02
NARUTO [18]	PSNR ↑	26.01	28.88	33.27	24.26	25.32	22.75	24.70	26.17	22.77
	SSIM \uparrow	0.89	0.93	0.96	0.89	0.91	0.88	0.87	0.90	0.82
	LPIPS \downarrow	0.41	0.40	0.27	0.38	0.38	0.44	0.49	0.43	0.52
	L1-D \downarrow	9.54	4.51	1.35	5.16	6.28	7.76	6.07	4.08	41.14
Ours	PSNR ↑	32.02	33.70	35.78	29.98	32.17	33.59	29.64	30.25	31.03
	SSIM ↑	0.97	0.97	0.97	0.97	0.98	0.98	0.96	0.96	0.97
	LPIPS \downarrow	0.11	0.10	0.13	0.10	0.08	0.08	0.14	0.12	0.10
	L1-D \downarrow	1.12	1.38	0.86	1.70	0.99	0.87	0.99	1.01	1.21

Table 2. Novel View Rendering Performance on Replica [60].

specific application requirements, balancing the need for detailed geometry versus high-quality rendering.

Geometric Evaluation Tab. 1 presents a quantitative comparison of our system against prior studies on MP3D. Our approach consistently outperforms previous methods across all evaluation metrics. The accuracy metric reflects the geometric precision of the Gaussian Map, which, given our initialization with observed depths, achieves high accuracy as expected. Additionally, both the Completion and Completion Ratio metrics, which measure the 3D space coverage through active exploration, show that our method attains exceptional completeness. This achievement stems from our rendering-based planning, which efficiently identifies the most informative viewpoints, combined with the agent's unrestricted movement capabilities. We also present qualitative comparisons in Fig. 3.

Rendering Evaluation We evaluate rendering performance on the high-fidelity Replica Dataset, as shown in Tab. 2, using novel view trajectories for Novel View Rendering. Two baselines are used for comparison.

The first baseline is the passive mapping method, SplaTAM [25]. For a fair comparison, we disable its track-

ing thread and simulate handheld scanning with a manually defined capture trajectory. Since manual scanning doesn't cover the entire scene, we exclude uncovered regions from the evaluation, focusing only on captured areas.

The second baseline is NARUTO [18], a state-of-the-art active mapping method that uses a neural radiance field. As NARUTO actively maps the environment, we evaluate all pixels in the novel views to assess its coverage fully.

Our proposed method consistently outperforms both baselines by a significant margin, benefiting from *actively* building a *Gaussian Map* representation capable of high-fidelity rendering, as evidenced in the quantitative results and qualitative examples in Fig. 4.

4.3. Ablation Studies

The Replica dataset features photorealistic 3D indoor scenes, represented by dense meshes with higher completeness than MP3D scenes. Given this level of detail, we primarily conduct our ablation studies on Replica to obtain more representative and robust results. In the following experiments, we use **Full** as the reference system, with ablation studies performed with respect to this configuration. The results are shown in Tab. 3.

Coarse-to-Fine Exploration We first investigate the effectiveness of the coarse-to-fine exploration strategy. While Coarse Exploration alone achieves substantial scene completeness, Fine Exploration—denoted as "w/o Refine"—further enhances both geometric and photometric reconstruction. Fig. 5 illustrates the progression of exploration completeness. As shown, our method quickly reaches high completeness in the Coarse Exploration stage and then refines missing details in the Fine Exploration stage.

Post-Refinement As discussed in Sec. 3.4, the post-refinement step in our full method improves rendering qual-



Figure 4. **Rendering Results on Replica.** Two scenes (office0, office4) are shown in the first and second rows, respectively. The results represent [Ground Truth, NARUTO, SplaTAM w/o tracking, Ours]. Our renders demonstrate higher fidelity across most regions.



Figure 5. Reconstruction Progress in Replica-office0. The distribution of rendering-based information (N_{S_i}) of candidate viewpoints is shown alongside reconstruction performance metrics. As information decreases, completeness and rendering performance improve, with Fine-Exploration yielding an additional boost.

Exp.	Comp. (cm)↓	Comp. Ratio (%) ↑	PSNR (dB) ↑	L1-D (cm)↓
Coarse Exploration	1.77	94.53	29.77	1.80
w/o Global-KF	2.19	94.87	30.73	1.23
w/o Refine.	1.56	96.50	30.67	1.42
Full	1.80	95.45	32.02	1.12

Table 3. Ablation Studies on Replica.

ity by further optimizing the Gaussian Map based on a rendering-oriented losses. However, this step may lead to Gaussian pruning, reducing the completeness of the map. When Post-Refinement is omitted ("w/o Refine"), the geometric reconstruction metrics improve due to the retention of redundant Gaussians, which are beneficial for geometry. Nonetheless, the post-refinement step significantly enhances rendering performance.

Global-Local Keyframe Selection In the experiment "w/o Global-KF," we assess the impact of excluding Global Keyframes. Without the regularization from Global Keyframes, 3D reconstruction performance declines, underscoring the importance of supervising by local and global keyframes to optimize the Gaussian Map effectively.

5. Discussion

In this paper, we introduced *ActiveGAMER*, an active mapping system that leverages 3D Gaussian Splatting (3DGS) for efficient, high-fidelity real-time scene reconstruction and exploration. By harnessing the fast rendering capabilities of 3DGS, our approach overcomes limitations of traditional NeRF-based methods, enabling unrestricted 6DoF movement and enhancing both geometric completeness and photometric fidelity.

Key features of *ActiveGAMER* include a rendering-based information gain module for optimal viewpoint selection and a balanced framework incorporating coarse-to-fine exploration, post-refinement, and global-local keyframe selection. Together, these components enable autonomous, high-quality exploration and reconstruction in complex environments, achieving state-of-the-art accuracy and fidelity. Extensive evaluations on challenging datasets like Replica and MP3D confirm that *ActiveGAMER* outperforms existing methods in both geometric accuracy, reconstruction completeness, and rendering quality, highlighting 3DGS as a valuable tool for real-time active vision applications

While ActiveGAMER demonstrates strong performance, several directions for future work could enhance its applicability. First, real-world scenarios require a robust planning and localization module to address assumptions of known localization and perfect action execution. Second, to increase system versatility, future iterations should account for real-world motion constraints that affect the agent's movement. Additionally, the current rendering-based information gain approach may overlook certain regions, such as double-sided objects. For example, after reconstructing one side of an object, the system may not detect the need to explore the back side if it's not visible in the initial view, leading to incomplete reconstructions. Future work will aim to integrate additional semantic and surface cues to better understand scene complexity and enhance exploration guidance.

References

- [1] Michal Adamkiewicz, Timothy Chen, Adam Caccavale, Rachel Gardner, Preston Culbertson, Jeannette Bohg, and Mac Schwager. Vision-only robot navigation in a neural radiance world. *IEEE Robotics and Automation Letters*, 7(2):4606–4613, 2022. 2
- [2] Dejan Azinović, Ricardo Martin-Brualla, Dan B Goldman, Matthias Nießner, and Justus Thies. Neural rgb-d surface reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6290–6301, 2022. 1, 2
- [3] Jiawang Bian, Zhichao Li, Naiyan Wang, Huangying Zhan, Chunhua Shen, Ming-Ming Cheng, and Ian Reid. Unsupervised scale-consistent depth and egomotion learning from monocular video. *Advances in neural information processing systems*, 32, 2019. 2
- [4] Frederic Bourgault, Alexei A Makarenko, Stefan B Williams, Ben Grocholsky, and Hugh F Durrant-Whyte. Information based adaptive robotic exploration. In *IEEE/RSJ international conference on intelligent robots and systems*, pages 540–545. IEEE, 2002. 2
- [5] Cesar Cadena, Luca Carlone, Henry Carrillo, Yasir Latif, Davide Scaramuzza, José Neira, Ian Reid, and John J Leonard. Past, present, and future of simultaneous localization and mapping: Toward the robustperception age. *IEEE Transactions on robotics*, 32(6): 1309–1332, 2016. 2
- [6] Changjiang Cai, Pan Ji, Qingan Yan, and Yi Xu. Riavmvs: Recurrent-indexing an asymmetric volume for multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 919–928, 2023. 2
- [7] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niessner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. arXiv preprint arXiv:1709.06158, 2017. 6
- [8] Devendra Singh Chaplot, Dhiraj Gandhi, Saurabh Gupta, Abhinav Gupta, and Ruslan Salakhutdinov. Learning to explore using active neural slam. arXiv preprint arXiv:2004.05155, 2020. 6
- [9] Liyan Chen, Weihan Wang, and Philippos Mordohai. Learning the distribution of errors in stereo matching for joint disparity and uncertainty estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17235–17244, 2023. 2
- [10] Shin-Fang Chng, Sameera Ramasinghe, Jamie Sherrah, and Simon Lucey. Gaussian activated neural radiance fields for high fidelity reconstruction and pose estimation. In *European Conference on Computer Vision*, pages 264–280. Springer, 2022. 2

- [11] Cl Connolly. The determination of next best views. In *Proceedings. 1985 IEEE international conference on robotics and automation*, pages 432–435. IEEE, 1985.
 2
- [12] Andrew J Davison and David W. Murray. Simultaneous localization and map-building using active vision. *IEEE transactions on pattern analysis and machine intelligence*, 24(7):865–880, 2002. 2
- [13] Andrew J Davison, Ian D Reid, Nicholas D Molton, and Olivier Stasse. Monoslam: Real-time single camera slam. *IEEE transactions on pattern analysis and machine intelligence*, 29(6):1052–1067, 2007. 2
- [14] Jeffrey Delmerico, Stefan Isler, Reza Sabzevari, and Davide Scaramuzza. A comparison of volumetric information gain metrics for active 3d object reconstruction. *Autonomous Robots*, 42(2):197–208, 2018. 2
- [15] Hugh Durrant-Whyte and Tim Bailey. Simultaneous localization and mapping: part i. *IEEE robotics & automation magazine*, 13(2):99–110, 2006. 2
- [16] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in neural information processing systems*, 27, 2014. 2
- [17] Hans Jacob S Feder, John J Leonard, and Christopher M Smith. Adaptive mobile robot navigation and mapping. *The International Journal of Robotics Research*, 18(7):650–668, 1999. 2
- [18] Ziyue Feng, Huangying Zhan, Zheng Chen, Qingan Yan, Xiangyu Xu, Changjiang Cai, Bing Li, Qilun Zhu, and Yi Xu. Naruto: Neural active reconstruction from uncertain target observations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21572–21583, 2024. 2, 4, 5, 6, 7
- [19] Georgios Georgakis, Bernadette Bucher, Anton Arapin, Karl Schmeckpeper, Nikolai Matni, and Kostas Daniilidis. Uncertainty-driven planner for exploration and navigation. In 2022 International Conference on Robotics and Automation (ICRA), pages 11295– 11302. IEEE, 2022. 6
- [20] Peter Hart, Nils Nilsson, and Bertram Raphael. A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, 4(2):100–107, 1968. 2
- [21] Heiko Hirschmuller. Accurate and efficient stereo processing by semi-global matching and mutual information. In 2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05), pages 807–814. IEEE, 2005. 2
- [22] Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In ACM SIG-

GRAPH 2024 Conference Papers, pages 1–11, 2024. 2

- [23] Stefan Isler, Reza Sabzevari, Jeffrey Delmerico, and Davide Scaramuzza. An information gain formulation for active volumetric 3d reconstruction. In 2016 IEEE International Conference on Robotics and Automation (ICRA), pages 3477–3484. IEEE, 2016. 2
- [24] Wen Jiang, Boshu Lei, Katrina Ashton, and Kostas Daniilidis. Ag-slam: Active gaussian splatting slam. arXiv preprint arXiv:2410.17422, 2024. 3, 6
- [25] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat track & map 3d gaussians for dense rgb-d slam. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 21357–21366, 2024. 2, 3, 5, 7
- [26] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. ACM Trans. Graph., 42(4):139–1, 2023. 1, 2, 3
- [27] Simon Kriegel, Christian Rink, Tim Bodenmüller, and Michael Suppa. Efficient next-best-scan planning for autonomous 3d surface reconstruction of unknown objects. *Journal of Real-Time Image Processing*, 10(4): 611–631, 2015. 2
- [28] Zijia Kuang, Zike Yan, Hao Zhao, Guyue Zhou, and Hongbin Zha. Active neural mapping at scale. arXiv preprint arXiv:2409.20276, 2024. 2, 3
- [29] James J Kuffner and Steven M LaValle. Rrt-connect: An efficient approach to single-query path planning. In Proceedings 2000 ICRA. Millennium Conference. IEEE International Conference on Robotics and Automation. Symposia Proceedings (Cat. No. 00CH37065), pages 995–1001. IEEE, 2000. 2
- [30] Steven M LaValle, James J Kuffner, BR Donald, et al. Rapidly-exploring random trees: Progress and prospects. *Algorithmic and computational robotics: new directions*, 5:293–308, 2001. 2, 5
- [31] Soomin Lee, Le Chen, Jiahao Wang, Alexander Liniger, Suryansh Kumar, and Fisher Yu. Uncertainty guided policy for active robotic 3d reconstruction using neural radiance fields. *IEEE Robotics and Automation Letters*, 2022. 2
- [32] Kejie Li, Yansong Tang, Victor Adrian Prisacariu, and Philip HS Torr. Bnv-fusion: Dense 3d reconstruction using bi-level neural volume fusion. In *Proceedings* of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022. 1, 2
- [33] Yuetao Li, Zijia Kuang, Ting Li, Guyue Zhou, Shaohui Zhang, and Zike Yan. Activesplat: High-fidelity scene reconstruction through active gaussian splatting. *arXiv preprint arXiv:2410.21955*, 2024. 3, 6

- [34] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5741– 5751, 2021. 2
- [35] Jiachen Liu, Pan Ji, Nitin Bansal, Changjiang Cai, Qingan Yan, Xiaolei Huang, and Yi Xu. Planemvs: 3d plane reconstruction from multi-view stereo. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 8665– 8675, 2022. 2
- [36] Iker Lluvia, Elena Lazkano, and Ander Ansuategi. Active mapping and robot exploration: A survey. *Sensors*, 21(7):2445, 2021. 2
- [37] Alexei A Makarenko, Stefan B Williams, Frederic Bourgault, and Hugh F Durrant-Whyte. An experiment in integrated exploration. In *IEEE/RSJ international conference on intelligent robots and systems*, pages 534–539. IEEE, 2002. 2
- [38] Hidenobu Matsuki, Riku Murai, Paul HJ Kelly, and Andrew J Davison. Gaussian splatting slam. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pages 18039–18048, 2024. 2, 3
- [39] Jasna Maver and Ruzena Bajcsy. Occlusions as a guide for planning the next view. *IEEE transactions* on pattern analysis and machine intelligence, 15(5): 417–433, 1993. 2
- [40] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106, 2021. 1, 2
- [41] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. 1
- [42] Paul Newman, Michael Bosse, and John Leonard. Autonomous feature-based exploration. In 2003 IEEE International Conference on Robotics and Automation (Cat. No. 03CH37422), pages 1234–1240. IEEE, 2003. 2
- [43] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11453–11464, 2021. 2
- [44] Xuran Pan, Zihang Lai, Shiji Song, and Gao Huang. Activenerf: Learning where to see with uncertainty estimation. In *European Conference on Computer Vi*sion, pages 230–246. Springer, 2022. 2, 6
- [45] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf:

Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF* conference on computer vision and pattern recognition, pages 165–174, 2019. 1, 2

- [46] Daryl Peralta, Joel Casimiro, Aldrin Michael Nilles, Justine Aletta Aguilar, Rowel Atienza, and Rhandley Cajote. Next-best view policy for 3d reconstruction. In *European Conference on Computer Vision*, pages 558–573. Springer, 2020. 2
- [47] Richard Pito. A solution to the next best view problem for automated surface acquisition. *IEEE Transactions on pattern analysis and machine intelligence*, 21(10): 1016–1030, 1999. 2
- [48] Julio A Placed, Jared Strader, Henry Carrillo, Nikolay Atanasov, Vadim Indelman, Luca Carlone, and José A Castellanos. A survey on active simultaneous localization and mapping: State of the art and new frontiers. arXiv preprint arXiv:2207.00254, 2022. 2
- [49] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10318–10327, 2021. 1, 2
- [50] Santhosh K Ramakrishnan, Ziad Al-Halah, and Kristen Grauman. Occupancy anticipation for efficient exploration and navigation. In *Computer Vision–ECCV* 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16, pages 400– 418. Springer, 2020. 6
- [51] Yunlong Ran, Jing Zeng, Shibo He, Lincheng Li, Yingfeng Chen, Gimhee Lee, Jiming Chen, and Qi Ye. Neurar: Neural uncertainty for autonomous 3d reconstruction. *arXiv preprint arXiv:2207.10985*, 2022. 2, 6
- [52] Manolis Savva, Abhishek Kadian, Oleksandr Maksymets, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, et al. Habitat: A platform for embodied ai research. In Proceedings of the IEEE/CVF international conference on computer vision, pages 9339–9347, 2019. 3, 6
- [53] Davide Scaramuzza and Friedrich Fraundorfer. Visual odometry [tutorial]. *IEEE robotics & automation magazine*, 18(4):80–92, 2011. 2
- [54] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of* the IEEE conference on computer vision and pattern recognition, pages 4104–4113, 2016. 2
- [55] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger. Graf: Generative radiance fields for 3daware image synthesis. Advances in Neural Information Processing Systems, 33:20154–20166, 2020. 2

- [56] Steven M Seitz, Brian Curless, James Diebel, Daniel Scharstein, and Richard Szeliski. A comparison and evaluation of multi-view stereo reconstruction algorithms. In 2006 IEEE computer society conference on computer vision and pattern recognition (CVPR'06), pages 519–528. IEEE, 2006. 2
- [57] Roland Siegwart, Illah Reza Nourbakhsh, and Davide Scaramuzza. *Introduction to autonomous mobile robots*. MIT press, 2011. 2
- [58] Cyrill Stachniss. *Robotic mapping and exploration*. Springer, 2009. 2
- [59] Cyrill Stachniss, Dirk Hahnel, and Wolfram Burgard. Exploration with active loop-closing for fastslam. In 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)(IEEE Cat. No. 04CH37566), pages 1505–1510. IEEE, 2004. 2
- [60] Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. arXiv preprint arXiv:1906.05797, 2019. 6, 7
- [61] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229– 6238, 2021. 6
- [62] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, 2022. 1
- [63] Jian Sun, Nan-Ning Zheng, and Heung-Yeung Shum. Stereo matching using belief propagation. *IEEE Transactions on pattern analysis and machine intel-ligence*, 25(7):787–800, 2003. 2
- [64] Sebastian Thrun. Probabilistic robotics. *Communications of the ACM*, 45(3):52–57, 2002. 2
- [65] Sebastian B Thrun and Knut Möller. Active exploration in dynamic environments. *Advances in neural information processing systems*, 4, 1991. 2
- [66] Weihan Wang, Jiani Li, Yuhang Ming, and Philippos Mordohai. Edi: Eskf-based disjoint initialization for visual-inertial slam systems. In 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), pages 1466–1472. IEEE, 2023. 2
- [67] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 2
- [68] Brian Yamauchi. A frontier-based approach for autonomous exploration. In Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. Towards New Computational Principles for Robotics and Automation', pages 146–151. IEEE, 1997. 6

- [69] Zike Yan, Haoxiang Yang, and Hongbin Zha. Active neural mapping. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10981–10992, 2023. 4, 6
- [70] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, pages 767– 783, 2018. 2
- [71] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1, 2
- [72] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19447–19456, 2024. 2
- [73] Huangying Zhan, Ravi Garg, Chamara Saroj Weerasekera, Kejie Li, Harsh Agarwal, and Ian Reid. Unsupervised learning of monocular depth estimation and visual odometry with deep feature reconstruction. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 340–349, 2018.
 2
- [74] Huangying Zhan, Chamara Saroj Weerasekera, Jia-Wang Bian, and Ian Reid. Visual odometry revisited: What should be learnt? In 2020 IEEE International Conference on Robotics and Automation (ICRA), pages 4203–4210. IEEE, 2020. 2
- [75] Huangying Zhan, Jiyang Zheng, Yi Xu, Ian Reid, and Hamid Rezatofighi. Activermap: Radiance field for active mapping and planning. *arXiv preprint arXiv:2211.12656*, 2022. 2, 6
- [76] Kai Zhang, Gernot Riegler, Noah Snavely, and Vladlen Koltun. Nerf++: Analyzing and improving neural radiance fields. arXiv preprint arXiv:2010.07492, 2020. 1, 2