

Uncertainty-driven 3D Gaussian Splatting Active Mapping via Anisotropic Visibility Field

Shangjie Xue^{*†} Jesse Dill^{*} Dhruv Ahuja^{*} Frank Dellaert Panagiotis Tsiotras Danfei Xu
Georgia Institute of Technology

Abstract

We present *Gaussian Splatting Anisotropic Visibility Field (GAVIS)*, a novel framework for uncertainty quantification and active mapping in 3DGS. Our key insight is that regions unseen from the training views yield unreliable predictions from the 3DGS. To address this, we introduce a principled and efficient method for quantifying the visibility field in 3DGS, defined as the anisotropic visibility of each particle with respect to the training views, and represented using spherical harmonics. The resulting visibility field is integrated into a Bayesian Network-based uncertainty-aware 3DGS rasterizer, enabling real-time (200 FPS) uncertainty quantification for synthesized views. Active mapping is further performed within a maximum information gain framework building on this formulation. Extensive experiments across diverse environments demonstrate that GAVIS consistently and significantly outperforms prior approaches in both accuracy and efficiency. Moreover, beyond standalone use, our method can be applied post-hoc to improve the performance of existing approaches.

1. Introduction

When an autonomous robot is deployed in a new environment, it should explore and understand its surroundings. This process, known as active mapping [7, 51, 105], is essential for applications such as household robots, space robots, and search-and-rescue robots. One common goal of active mapping is to achieve comprehensive visual coverage of the scene, formulated as a planning objective that reduces uncertainty in the reconstructed map [8, 104]. A suitable representation is critical for active mapping, as it must not only enable high-quality and fast reconstruction but also enable accurate uncertainty quantification [66].

Recent advancements in radiance fields [57] have demonstrated remarkable capabilities in high-quality scene reconstruction. Notably, the emergence of 3D Gaussian Splatting (3DGS) [39] has shown superior reconstruction quality and rendering speed. However, due to the large




Figure 1. **GAVIS overview.** Gaussian Splatting Anisotropic Visibility Field (GAVIS) quantifies uncertainty in 3DGS by modeling visibility, i.e., whether a region is observed by the training views. Observed regions have low uncertainty (left room), whereas unobserved regions have high uncertainty (right room).

number of trainable parameters in 3DGS, accurately and efficiently quantifying its uncertainty remains challenging. Recent approaches to uncertainty quantification in 3DGS adopted established epistemic uncertainty quantification methods from machine learning, including variational inference [54] and Laplace approximation [32], to tackle this challenge. However, a gap exists between the objective of machine learning uncertainty quantification and active mapping. In active mapping, the robot aims to achieve high-quality reconstruction by providing visual coverage of the entire scene, where predictions for regions never visible in the training views are always considered unreliable. However, as an approximation method, existing learning-based uncertainty quantification techniques are unable to guarantee this, and often underestimate the uncertainty in these unseen, out-of-distribution regions [22, 63, 95].

We propose 3D Gaussian Splatting Anisotropic Visibility Field (GAVIS), a principled method that reliably assigns high uncertainty to unobserved regions for active mapping, while incurring minimal computational overhead. This method can be used independently or integrated seamlessly, as a post-hoc module, into existing learning-based uncertainty quantification framework that estimates uncertainty in 3DGS parameters or volumetric radiance field outputs, thereby improving its performance.

^{*}Equal contribution. [†]xsj@gatech.edu




Figure 2. **GAVIS framework.** (Left) Given a trained 3DGS, GAVIS constructs a visibility field (VF CONST) to represent regions invisible to the training views. It then quantifies uncertainty over sampled candidate views using an uncertainty-aware 3DGS rasterizer (UA 3DGS Rasterizer) that queries the visibility field (VF QUERY). Finally, the maximum-uncertainty view is selected as the next observation. (Right) GAVIS achieves top performance across all evaluation metrics (see Sec. 5 for details).

GAVIS extends the 3D Gaussian Splatting framework to quantify the extent to which each region is covered by the training views, inspired by NVF [104], which shows that visibility is closely related to regional uncertainty. However, NVF relies on training a neural network to approximate a view-independent (isotropic) visibility field for NeRF, making it computationally expensive and impractical for real-world applications, while also introducing approximation errors and neglecting the inherently view-dependent nature of visibility. In contrast, GAVIS explicitly models the view-dependent visibility of each Gaussian particle relative to all training views (see Sec. 4.1 for details), with analytical, efficient, gradient-free computation. The uncertainty-aware 3DGS rasterization is then modeled as a Bayesian network, which integrates both field-based uncertainty and visibility into the ray-based camera observations. Within this framework, the color distribution along a ray is modeled as a Gaussian Mixture Model (GMM). The entropy of the GMM serves as an objective to guide the agent in selecting the optimal poses for active mapping.

Extensive evaluations are conducted across diverse environments, including standard NeRF Blender assets, indoor scenes, and space robot scenarios. The results demonstrate that our method outperforms all existing approaches in both accuracy and efficiency. Moreover, we also demonstrate that our method can be used independently or as a post-hoc method to enhance the performance of existing methods. To summarize, our main contributions are:

- We propose a principled uncertainty quantification method for 3DGS that incorporates anisotropic visibility.
- The proposed method can be used independently or integrated seamlessly as a post-hoc enhancement to improve the performance of existing frameworks.
- Through extensive evaluations across diverse environments, we demonstrate that our method is both effective for active mapping tasks and computationally more efficient than existing approaches.

2. Related Work

Uncertainty Quantification in Radiance Fields. With the introduction of NeRF [57], subsequent work has improved

efficiency [20, 39, 59] and enabled online SLAM [36, 55] with explicit representations such as 3DGS [39]. Uncertainty quantification in radiance fields [41] has seen growing interest for artifact removal [23, 96] and active mapping [32, 46, 104]. Parameter uncertainty is estimated via ensembles [89], variational inference [48, 54, 76], conditional flows [77], Fisher information via Laplace approximation [23, 32, 49], or Monte Carlo dropout [21]. Alternatively, radiance field outputs can be modeled as probability distributions [33, 46, 47, 64, 70, 100, 104, 107], including as Gaussian-distributed RGB [64, 70], occupancy from NeRF density [19, 46, 78, 107] or 3DGS opacity [12, 35, 50], or as a volume rendering weighted GMM to incorporate visibility into ray-wise uncertainty [104].

Active Mapping. Active mapping is typically posed as an optimization over information gain [5, 6, 9, 10, 14, 43, 56, 58, 65, 71, 82, 85, 86, 94, 98], and augmented with frontier-based [15, 18, 24, 26, 37, 38, 52, 68, 71, 74, 75, 93, 105, 106] or sampling-based [7, 8, 16, 44, 65, 68, 74, 79, 90, 93, 113] heuristics to reduce computational cost [66]. Recent radiance field active mapping methods pair voxel grids [12, 19, 34, 35, 103, 108] or Voronoi graphs [45, 50], while a complementary line of work learns planning policies end-to-end [11, 42, 61, 97, 114].

Radiance Fields Visibility. Visibility has received growing attention in radiance fields for uncertainty quantification and active mapping, providing a measure of where the scene is poorly observed. Matsuki et al. [55] define single-view visibility via an occupancy threshold; in 3DGS this reduces to particle opacity [12, 50]; however, single-view metrics ignore prior observations. Multi-view methods [60, 104] aggregate visibility across views, yet face limitations, specifically, NVF [104] requires neural-network training to approximate positional visibility, and ProvNeRF [60] models per-point provenance via additional post-hoc optimization, both incur minutes to hours of runtime, limiting real-time applicability. Beyond active mapping, visibility also serves as a cue for inpainting by identifying high-uncertainty regions [81], and has been incorporated into radiance-field pipelines for relighting [84], material decomposition [111], occlusion-aware view synthesis [13, 27, 80, 109], and scene reconstruction [83, 112].

3. Background

This section reviews the background on radiance fields, uncertainty-aware volume rendering, and active mapping, with further details available in appendix Sec. 8, 9, and [39, 57, 104].

Radiance field. A radiance field [57] F_Θ maps a 3D location $\mathbf{x} \in \mathbb{R}^3$ and a view direction $\mathbf{d} \in \mathbb{S}^2$ to color $\mathbf{c}(\mathbf{x}, \mathbf{d}) \in \mathbb{R}^3$ and volume density $\sigma(\mathbf{x}) \in \mathbb{R}_{\geq 0}$. Given a camera ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$ with near/far bounds t_n, t_f , the pixel color is computed via volume rendering:

$$\mathbf{C}(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt, \quad (1)$$

where $T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right)$ represents the (accumulated) transmittance from the near bounds t_n to point t . In practice, the integral is discretized and computed numerically. In NeRF [57], F_Θ is a neural network with parameters Θ ; in 3DGS [39], the scene is explicitly represented by a trainable mixture of ellipsoidal Gaussian primitives parameterized by Θ .

Active Mapping. Active mapping, or next-best-view (NBV) planning, aims to determine the robot’s next action that most effectively reduces the uncertainty of the reconstructed map. A practical formulation is to select the action with the highest predicted entropy in observation:

$$\tau^* = \arg \max_{\tau} \mathcal{H}(\mathbf{Z}_\tau), \quad (2)$$

where τ denotes the action (i.e., a candidate camera pose) and $\mathbf{Z}_\tau$ is the random variable which represents the observation obtained after executing action τ , and $\mathcal{H}$ is the entropy. This objective has been widely adopted in methods that represent scenes using radiance fields [46, 54, 104, 107], and can be interpreted as an approximation of the fundamental objective of maximizing the expected information gain of the map obtained from new observations (see Sec. 9 for details). However, a key challenge remains: accurately and efficiently quantifying the uncertainty of synthesized views generated by a radiance field.

Uncertainty-aware Volume Rendering. To quantify synthesized-view uncertainty, NVF [104] extends deterministic volume rendering of NeRF to a stochastic formulation. By assuming the emitted color $\mathbf{c}$ is Gaussian, the probability density functions (PDF) of the pixel colors are modeled as Gaussian mixture models (GMM) induced by a Bayesian-network view of uncertainty-aware volume rendering, specifically $p(\mathbf{z}) = \sum_i w_i \mathcal{N}(\boldsymbol{\mu}_{\mathbf{c}_i}, \mathbf{Q}_{\mathbf{c}_i})$, where $p(\mathbf{z})$ is the PDF of observed pixel color, $\boldsymbol{\mu}_{\mathbf{c}_i}$ and $\mathbf{Q}_{\mathbf{c}_i}$ are the mean and variance of emission color at point i , w_i is the mixture weight, which can be obtained from (1) (details in Sec. 8). NVF [104] highlighted the crucial role of visibility in uncertainty quantification and active mapping, as predictions for regions unseen by training views are unreliable.

Accordingly, a visibility correction term is incorporated into the GMM of the pixel-color PDF:

$$p(\mathbf{z}_0) = \sum_i w_i^* v_i \mathcal{N}(\boldsymbol{\mu}_{\mathbf{c}_i}, \mathbf{Q}_{\mathbf{c}_i}) + \mathcal{N}(\boldsymbol{\mu}_0, \mathbf{Q}_0) \sum_i w_i^* (1 - v_i). \quad (3)$$

where w_i^* denotes the visibility-corrected mixture weight (see Sec. 8 for details), $\mathcal{N}(\boldsymbol{\mu}_0, \mathbf{Q}_0)$ is a prior distribution with large variance representing emitted color of invisible regions, v_i represents the probability that point $\mathbf{x}_i$ is visible in at least one training view, quantified as $v_i = V(\mathbf{x}_i)$, where $V : \mathbb{R}^3 \rightarrow [0, 1]$ is the (isotropic) **visibility field** as a function of position. In NVF [104], a neural network is used to both learn the radiance field and approximate the visibility field $V_\theta(\mathbf{x}_i)$. Moreover, its accuracy is limited since it relies on a neural approximation and models visibility solely as a function of position, ignoring directional anisotropy. Details are discussed in the next section.

4. Method

The main challenge in applying radiance-field-based uncertainty quantification to active mapping lies in accurately estimating uncertainty that enables effective exploration while maintaining a practical runtime. Visibility-field-based approach NVF [104], which assigns high uncertainty to unseen regions, could potentially facilitate effective exploration. However, NVF typically require retraining a neural network for the visibility field at every planning step, which takes **several minutes** and severely limits their applicability to real-world robotic systems.

To address this challenge, we propose 3D Gaussian Splatting Anisotropic Visibility Field (GAVIS), which explicitly models the visibility field using the 3DGS representation. This formulation enables analytical construction of the visibility field without training, achieving construction **within 1 second**, which is **500× faster** than NVF [104], while providing more accurate uncertainty quantification that leads to more effective exploration.

In this section, we first formulate the visibility field within the 3DGS representation (Sec. 4.1). In particular, we introduce a direction-aware visibility field that can account for visibility change given a specific viewing direction. We then introduce an efficient method to construct and query the visibility field (Sec. 4.2) by representing it with spherical harmonics. Finally, we summarize the overall active mapping pipeline that integrates the visibility field with uncertainty-aware 3DGS rasterization (Sec. 4.3).

4.1. Formulation

In this subsection, we introduce a novel formulation of *visibility field* in 3DGS. Notably, a straightforward extension of the Neural Visibility Field [104] (as a function of position only) from NeRF to 3DGS (i.e. assigning a scalar visibility value to each Gaussian) is insufficient for reliable

uncertainty quantification. Since a 3DGS particle can occlude itself (i.e. self-occlusion), observing a particle from one direction provides little information about its appearance from the opposite side. As a simple example, viewing only one side of a wall reveals nothing about the opposite side (Fig. 7). A reliable uncertainty quantification method should assign high uncertainty there to drive exploration.

To account for this, a visibility field should depend on the viewing direction. We therefore define the **Anisotropic Visibility Field** $V^{(i)}(\mathbf{d})$ for 3DGS, where for each Gaussian particle i , the visibility is formulated as a function of the rendering direction $\mathbf{d}$, with $\mathbf{d}$ being the unit vector pointing toward the synthesized view.

To obtain $V^{(i)}(\mathbf{d})$, which is the visibility with respect to the entire training set (views observed so far), we first compute the visibility with respect to a single training view. Let $V_p^{(i)}(\mathbf{d})$ denote the visibility from a specific camera pose $\mathbf{p} \in \mathcal{P}$, where $\mathcal{P} = \{\mathbf{p}_1, \mathbf{p}_2, \dots\}$ is the set of camera poses in the training set. We have

$$V_p^{(i)}(\mathbf{d}) = \underbrace{\Phi_{i,\mathbf{p}}}_{\text{FOV}} \underbrace{T_p(t_i^{\mathbf{p}})}_{\text{Transmittance}} \underbrace{\nu(\mathbf{d}; \mathbf{d}_p)}_{\text{Directional Visibility}} \quad (4)$$

where

$$\nu(\mathbf{d}; \mathbf{d}_p) := \zeta \exp(\kappa \mathbf{d} \cdot \mathbf{d}_p), \quad \zeta = \exp(-\kappa)$$

This equation consists of three components:

FOV $\Phi_{i,\mathbf{p}} \in \{0, 1\}$ is a binary indicator denoting whether the particle i lies within the field of view of camera $\mathbf{p}$, where $\Phi_{i,\mathbf{p}} = 1$ if it is within the field of view and 0 otherwise.

Transmittance $T_p(t_i^{\mathbf{p}})$ denotes the probability that the ray travels along $\mathbf{d}_p$ from $\mathbf{x}_p$ to particle i without occlusion, where $\mathbf{d}_p = \frac{\mathbf{x}_i - \mathbf{x}_p}{\|\mathbf{x}_i - \mathbf{x}_p\|}$, $\mathbf{x}_p$ is the position of camera view $\mathbf{p}$. This term can be directly obtained from the radiance field output. The first two terms $\Phi_{i,\mathbf{p}} T_p(t_i^{\mathbf{p}})$ define isotropic visibility, identical to the formulation of NVF [104].

Directional Visibility Function $\nu(\mathbf{d}; \mathbf{d}_p)$ captures how visibility $V_p^{(i)}(\mathbf{d})$ changes as the rendering direction $\mathbf{d}$ of a novel view deviates from the training view direction $\mathbf{d}_p$. Viewing an object from one direction does not imply knowledge of its appearance from the opposite side; the larger the angular difference between $\mathbf{d}$ and $\mathbf{d}_p$, the lower the visibility and the higher the uncertainty. We use the spherical function $\zeta \exp(\kappa \mathbf{d} \cdot \mathbf{d}_p)$, proportional to the widely used von Mises-Fisher distribution, an analogue of the Gaussian distribution for spherical data. Here, κ is analogous to the reciprocal of variance in the Gaussian distribution, controlling the concentration around the mean direction, and $\zeta = \exp(-\kappa)$ is a constant factor ensuring $\nu(\mathbf{d}_p; \mathbf{d}_p) = 1$ when the rendering direction matches the training direction.

Therefore, the anisotropic visibility field for the entire training set can be computed from the per-training-view anisotropic visibility $V_p^{(i)}(\mathbf{d})$ (i.e., the probability of being




Figure 3. **Schematic of directional visibility.** (Top) A 2D illustration of how visibility varies with viewing direction. Anisotropic visibility is plotted in polar coordinates for (1) a single training view, (2) a single training view with occlusion, and (3) multiple training views. (Bottom) A 3D illustration of spherical harmonics expansion of $\hat{V}(\mathbf{d})$. See Secs. 4.1 and 4.2 for details.

visible from a single view). Specifically, the probability that particle i is visible to at least one camera in $\mathcal{P}$ is given by

$$V^{(i)}(\mathbf{d}) = 1 - \prod_{\mathbf{p} \in \mathcal{P}} (1 - V_p^{(i)}(\mathbf{d})). \quad (5)$$

An illustrative example of the anisotropic visibility field under different occlusion conditions is provided in Fig. 3. In the following subsection, we present an accurate and efficient algorithm for visibility field construction and querying, built upon the above formulation.

4.2. Efficient Construction and Querying of Anisotropic Visibility Field

So far, we have introduced the formulation of the anisotropic visibility field. However, directly applying the analytical expression obtained by combining (4) and (5) for uncertainty quantification and active mapping is impractical in real-world scenarios. Querying visibility from a single direction $\mathbf{d}$ requires accessing all training view directions $\mathbf{d}_p$, where $\mathbf{p} \in \mathcal{P}$, causing both runtime and memory cost to scale with the length of the history trajectory. For real-world robotic applications, it is crucial to ensure fast construction and querying of the visibility field. Following the same philosophy as radiance fields, we seek a representation of the anisotropic visibility field $V^{(i)}(\mathbf{d})$ that can be stored in constant memory and queried in constant time, independent of the trajectory length. Unlike NVF [104], which relies on a neural network to approximate the (isotropic) vis-

ibility field, we propose an analytical representation of the anisotropic visibility field based on spherical harmonics that is both more accurate and computationally efficient. In this subsection, we present the core components of our method that enable accurate and efficient construction and querying of $V^{(i)}(\mathbf{d})$ for uncertainty-driven active mapping.

Overview. Inspired by volumetric radiance field [20, 59], and 3DGS, where the directional appearance component (color) is represented using spherical harmonics [39], we express the **anisotropic visibility field** in the orthonormal basis of **spherical harmonics (SH)** $Y_\ell^m(\mathbf{d})$, as in Fig. 3

$$\tilde{V}^{(i)}(\mathbf{d}) = \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} \gamma_{\ell m}^{\mathcal{P}} Y_\ell^m(\mathbf{d}), \quad (6)$$

where $\tilde{V}^{(i)}(\mathbf{d})$ is an auxiliary function related to visibility field $V^{(i)}(\mathbf{d})$, and $\gamma_{\ell m}^{\mathcal{P}}$ denotes the spherical harmonics coefficients, L denotes the maximum degree of the spherical harmonics. Next, we describe how auxiliary function $\tilde{V}^{(i)}(\mathbf{d})$ is used to compute the visibility field $V^{(i)}(\mathbf{d})$, and show how the coefficients $\gamma_{\ell m}^{\mathcal{P}}$ are computed analytically during visibility-field construction.

SH Representation of $\nu(\mathbf{d}; \mathbf{d}_p)$. We notice that the directional visibility function $\nu(\mathbf{d}; \mathbf{d}_p)$ can be decomposed into spherical harmonics as

$$\nu(\mathbf{d}; \mathbf{d}_p) = \zeta \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_\ell^m(\mathbf{d}), \quad (7)$$

where the coefficients $a_{\ell m}$ can be computed analytically as

$$a_{\ell m} = 4\pi i_\ell(\kappa) Y_\ell^{m*}(\mathbf{d}_p), \quad (8)$$

and $Y_\ell^m(\mathbf{d})$ represents the spherical harmonic basis function, and $Y_\ell^{m*}(\mathbf{d}_p)$ is its complex conjugate, $i_\ell(\kappa)$ denotes the modified spherical Bessel function of the first kind [62]. We include the full derivation in appendix Sec. 7.1.

SH Representation of $V^{(i)}(\mathbf{d})$. Since spherical harmonics form a linear basis, the sum of two functions expressed in spherical harmonics can be obtained by simply adding their corresponding coefficients. However, the multiplication of two spherical harmonics [99] is computationally expensive [53]. Directly substituting (7) and (4) into (5) to compute the SH coefficients of $V^{(i)}(\mathbf{d})$ is impractical, as it requires a computational complexity of $O(|\mathcal{P}|^4 L^3)$ and a memory cost of $O(|\mathcal{P}|^2 L^2)$ to store the parameters of $V^{(i)}(\mathbf{d})$ (see Sec. 7.2 for details). Consequently, the query-time complexity scales quadratically with the number of training views $|\mathcal{P}|$, which is not suitable for real-time applications. To address this, we propose an efficient method that bypasses the direct multiplication of spherical harmonics, achieving computational complexity that is linear in $|\mathcal{P}|$ during construction and constant during query time. Specifically, we can obtain a lower bound of $V^{(i)}(\mathbf{d})$ by applying the arithmetic-geometric mean (AM-GM) inequality to (5):

Algorithm 1 Query Anisotropic Visibility Field

Input: SH coefficients $\{\gamma_{\ell m}^{(g)}\}$ for particle g ; query direction $\mathbf{d}$; SH degree L ; number of training views $|\mathcal{P}|$

- 1: $\tilde{V} \leftarrow \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} \gamma_{\ell m}^{(g)} Y_\ell^m(\mathbf{d})$
- 2: $V^{(g)}(\mathbf{d}) \leftarrow 1 - (1 - \frac{\tilde{V}}{|\mathcal{P}|})^{|\mathcal{P}|} \triangleright$ AM-GM lower-bound estimator
- 3: **return** $V^{(g)}(\mathbf{d})$

$$V^{(i)}(\mathbf{d}) \geq 1 - \left(1 - \frac{\tilde{V}^{(i)}(\mathbf{d})}{|\mathcal{P}|}\right)^{|\mathcal{P}|}. \quad (9)$$

where

$$\tilde{V}^{(i)}(\mathbf{d}) := \sum_{\mathbf{p}} V_{\mathbf{p}}^{(i)}(\mathbf{d}) \quad (10)$$

By substituting (4) into (10) and comparing it with (6), we can obtain the SH coefficients of $\tilde{V}^{(i)}(\mathbf{d})$ as

$$\gamma_{\ell m}^{\mathcal{P}} = 4\pi \zeta i_\ell(\kappa) \sum_{\mathbf{p} \in \mathcal{P}} \Phi_{i,\mathbf{p}} T_{\mathbf{p}}(t_i^{\mathbf{p}}) Y_\ell^{m*}(\mathbf{d}_{\mathbf{p}}). \quad (11)$$

More details are provided in Sec. 7.3.

Visibility Field Construction. During the construction stage of the visibility field, given a trained 3DGS, we first use (11) to compute the SH coefficients of $\tilde{V}^{(i)}(\mathbf{d})$, where the single-view particle visibility $\Phi_{i,\mathbf{p}} T_{\mathbf{p}}(t_i^{\mathbf{p}})$ is efficiently computed using a modified 3DGS rasterizer (see appendix Alg. 2 for details). For each particle, the set of coefficients $\{\gamma_{\ell m}^{\mathcal{P}}\}$ is stored as the parameters of the visibility field, requiring $(L+1)^2$ parameters per particle. In practice, we find that $L=2$ is sufficient to accurately capture visibility anisotropy. This SH degree is also commonly used in 3DGS to represent view-dependent color [39].

Constant-time Visibility Field Query. So far, we have introduced an efficient algorithm that constructs the visibility field using spherical harmonics and stores its parameters in the coefficients $\{\gamma_{\ell m}^{\mathcal{P}}\}$. When querying the visibility field in a given direction $\mathbf{d}$, we first compute $\tilde{V}^{(i)}(\mathbf{d})$ from (6) using the visibility field parameters $\{\gamma_{\ell m}^{\mathcal{P}}\}$. We then apply the lower bound in (9) to estimate the anisotropic visibility (see Alg. 1 for details). Importantly, this formulation enables **constant-time querying of the visibility field**.

Visibility Field Density Control. A core component of 3DGS is its adaptive density control [39]: particles are densified in regions of interest and pruned in free space, yielding high-quality constructions without substantial memory or runtime overhead. However, this mechanism complicates uncertainty estimation. Empty regions may correspond either to true free space (pruned) or to underexplored areas with too few initialized particles (e.g., from sparse SfM/depth initialization). Because particles in unseen areas contribute no loss, they receive no gradient and thus are not

densified, making both cases appear similarly empty. As a result, particle-centric uncertainty measures often overlook such regions and falsely assign uniformly low uncertainty to both, even though distinguishing them is crucial for active mapping, where the uncertainty quantification method should assign high uncertainty to underexplored areas.

To tackle this challenge, we propose using *virtual particles* to distinguish free space from underexplored regions. Given a trained 3DGS, we uniformly sample particles across the scene with zero opacity and compute their visibility as $1 - \prod_{p \in \mathcal{P}} (1 - \Phi_{i,p} T_p(t_i^p))$ with respect to all training views. Low visibility indicates that a particle lies in an underexplored (unseen) region, while high visibility suggests it resides in free space. Therefore, we prune virtual particles in free space using a visibility threshold and concatenate the remaining particles with the trained 3DGS for uncertainty quantification, where the visibility of virtual particles is set to $V^{(i)}(\mathbf{d}) = 0$ for all directions. We find that using a relatively small number of virtual particles (approximately 5–10% of the total particles) is sufficient to distinguish underexplored regions from free space. Further implementation details are provided in appendix Sec. 7.5.

4.3. Active Mapping Pipeline

So far, we have introduced our core contribution: an accurate and efficient method to compute the 3DGS visibility field $V^{(i)}(\mathbf{d})$. We then discuss how to apply it to uncertainty quantification and active mapping.

Our approach extends the NeRF-based uncertainty-aware volume rendering pipeline in NVF [104] to 3DGS rasterization, while providing significant improvements in visibility estimation and computational efficiency. Given the observed images and corresponding camera poses, a 3DGS is first trained. We then construct the visibility field by computing the visibility field parameters $\{\gamma_{\ell m}^p\}$ using (11), and integrating virtual particles that identify the underexplored regions. Once the visibility field is constructed, uncertainty quantification is performed by evaluating the entropy of synthesized views within a candidate view set sampled from a prior distribution (see Sec. 5.1). For each candidate view, uncertainty-aware 3DGS rasterization (see Alg. 6) is applied following (3), where v_i is obtained by querying the visibility field $V^{(i)}(\mathbf{d})$ along ray direction $\mathbf{d}$ using (9). Finally, the next best view is selected as the one that maximizes entropy, as defined in (2). An overview of this pipeline is illustrated in Fig. 2, and further details are provided in appendix Sec. 8 & 9.

5. Experiments

In this section, we design experiments to address the following key questions: (Q1) Can GAVIS accurately quantify uncertainty to enable effective active mapping? (Q2) Can GAVIS operate efficiently in practice? (Q3) How ef-

fective is each component of GAVIS? (Q4) Can GAVIS be applied as a post-hoc method to improve the performance of existing uncertainty quantification approaches?

5.1. Experiment Setups

Environments. To demonstrate the advantages of our method across diverse scenes and robotic applications, we conduct experiments on three types of scenarios spanning 3 domains and 4 datasets: (1) the standard NeRF Synthetic dataset [57] for object reconstruction; (2) a space dataset [104] consisting of the Hubble Space Telescope (HST) and the International Space Station (ISS) for space robotics; and (3) indoor scenes, including 8 environments from the Habitat-Matterport 3D (HM3D) [69] and 8 from Gibson [101] datasets, for household robotics. We set the planning steps to 10 (NeRF-Synthetic), 40 (Gibson), and 80 (HM3D), determined as the minimal steps at which the strongest methods achieve reasonable reconstruction quality. Further details are provided in appendix Sec. 11.1.

Baselines. We evaluate GAVIS against state-of-the-art (SOTA) uncertainty quantification approaches for radiance fields in active mapping, including 3DGS-based FisherRF [32], VIMC [54], and NeRF-based NVF [104]. To ensure a fair comparison of uncertainty quantification methods, we use the same active mapping pipeline for all methods by following the setup in NVF [104]. All methods use the same view sampler that uniformly draws collision-free poses from the scene without any additional heuristics to fairly evaluate the performance of the uncertainty quantification method. Further details on baseline and training setup are provided in Sec. 11.3.

Metrics. For each method, we evaluate performance after mapping is complete using PSNR, SSIM, and LPIPS [110], as well as mesh-based metrics including completion ratio (CR) [88] and scene visual coverage (VIS) [104], to assess the effectiveness of active mapping. We additionally report mesh metrics, including completion (Comp) and accuracy (Acc), in the appendix Sec. 12.1. However, we note that Acc is not well aligned with the objective of active mapping under 3DGS, further discussion is provided in Appendix Sec. 11.4. To evaluate the computational efficiency of each method, we measure the runtime during active mapping. Since planning time depends on the number of candidate views used for uncertainty quantification, we decompose the runtime into two components: (a) **Uncertainty Preparation Time** (T_{UP}), the additional time required after training the radiance field, independent of the number of candidate views. For GAVIS and NVF, this corresponds to visibility field construction; for FisherRF, to estimate the Hessian matrix of model parameters; and for VIMC, to the extra time compared to standard 3DGS training. (b) **Uncertainty Quantification Time** (UQ FPS), the frame rate (FPS) for evaluating uncertainty for each candi-

Method	NeRF Synthetic							Space						
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CR $\uparrow$	VIS $\uparrow$	UQ FPS $\uparrow$	T_{UP} $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CR $\uparrow$	VIS $\uparrow$	UQ FPS $\uparrow$	T_{UP} $\downarrow$
FisherRF	22.34	0.870	0.119	0.626	0.376	<u>146.3</u>	<u>0.42</u>	24.17	0.834	0.158	0.547	0.474	<u>141.5</u>	<u>0.37</u>
VIMC	<u>23.14</u>	<u>0.880</u>	<u>0.107</u>	<u>0.651</u>	0.407	144.5	9.48	<u>24.56</u>	<u>0.841</u>	<u>0.150</u>	<u>0.612</u>	0.510	127.6	17.75
NVF	22.59	0.859	0.147	0.549	<u>0.431</u>	11.9	149.13	23.76	0.796	0.202	0.499	<u>0.564</u>	10.7	140.54
GAVIS (ours)	24.26	0.894	0.097	0.711	0.437	251.5	0.17	26.14	0.857	0.140	0.630	0.582	232.4	0.17

Method	Gibson							HM3D						
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CR $\uparrow$	VIS $\uparrow$	UQ FPS $\uparrow$	T_{UP} $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CR $\uparrow$	VIS $\uparrow$	UQ FPS $\uparrow$	T_{UP} $\downarrow$
FisherRF	18.11	0.720	0.419	0.431	0.469	39.8	<u>0.90</u>	18.32	0.693	0.446	0.447	0.558	37.7	<u>1.59</u>
VIMC	15.70	0.668	0.465	0.337	0.366	<u>57.0</u>	90.47	17.15	0.645	0.477	0.476	0.618	<u>50.1</u>	114.51
NVF	<u>23.29</u>	<u>0.798</u>	<u>0.402</u>	0.880	0.915	4.2	219.87	<u>22.69</u>	<u>0.760</u>	<u>0.434</u>	<u>0.819</u>	<u>0.873</u>	4.2	285.26
GAVIS (ours)	24.42	0.812	0.323	<u>0.831</u>	<u>0.890</u>	207.3	0.42	23.97	0.791	0.338	0.820	0.876	192.6	0.79

Table 1. **Quantitative results.** Active mapping performance on all datasets across all baselines and our method. Best results are in **bold**; second-best are underlined. Here, T_{UP} denotes the uncertainty preparation time, and UQ FPS is the frame rate (FPS) for quantifying uncertainty for each candidate view. See Sec. 5 for details.

date view. We also report the Area Under the Sparsification Error curve (AUSE) [29, 67] for uncertainty quality. Noting that the standard depth-based variant (AUSE-D) [23, 32, 54] can be misaligned with active mapping goals, we additionally introduce a visibility-based variant (AUSE-V) that better correlates with active mapping performance. Details and analysis are provided in Sec. 11.4.

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CR $\uparrow$	VIS $\uparrow$
GAVIS	24.70	0.839	0.224	0.748	0.697
Isotropic	23.97	0.827	0.231	0.741	0.671
w/o DC	24.18	0.830	0.234	0.712	0.668
Iso. w/o DC	23.38	0.819	0.240	0.691	0.625

Table 2. **Ablation study.** We evaluate GAVIS for active mapping by isolating the effects of (i) anisotropic visibility ($\nu(\mathbf{d}; \mathbf{d}_p) = 1$) and (ii) the Visibility-Field Density Control module.

	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	CR $\uparrow$	VIS $\uparrow$
Fisher	20.73	0.779	0.285	0.513	0.469
VIMC	20.14	0.758	0.300	0.519	0.475
Fisher+GAVIS	24.70	0.842	0.220	0.748	0.699
VIMC+GAVIS	24.21	0.833	0.227	0.719	0.672

Table 3. **GAVIS as post-hoc module.** Active mapping results showing that applying GAVIS post hoc improves the performance of baseline 3DGS uncertainty-quantification methods.

	GAVIS (Ours)	NVF	FisherRF	VIMC
AUSE-D $\downarrow$	0.224	<u>0.381</u>	0.463	0.504
AUSE-V $\downarrow$	0.176	<u>0.231</u>	0.496	0.447

Table 4. **Uncertainty quantification.** Quality of uncertainty is evaluated using the depth-based AUSE (AUSE-D) and visibility-based variant (AUSE-V). See Sec. 5 for details.

5.2. Core Results

GAVIS accurately quantifies uncertainty and enables effective active mapping. As shown in the quantitative results (Tab. 1) and qualitative results (Fig. 5, and Fig. 4), our method consistently outperforms all 3DGS-based baselines (FisherRF and VIMC) across all datasets and evaluation metrics. The improvement is particularly pronounced in challenging indoor scenes such as Gibson and HM3D, where GAVIS significantly surpasses FisherRF and VIMC. This is because these methods do not model visibility, highlighting that visibility modeling is crucial for active mapping in complex scenes, which is consistent with the findings in NVF [104]. GAVIS outperforms NVF across all image-based metrics. This highlights the importance of

analytical anisotropic visibility modeling for high-quality novel view synthesis. For visibility and mesh metrics, the gains are less pronounced because these metrics are inherently isotropic and direction-agnostic: a mesh face is marked visible once it is observed from any single direction, and a single depth observation may be sufficient to reconstruct the local geometry. Consequently, the benefits of anisotropic visibility, which encourages revisiting the same region from diverse directions, are not fully captured. To further demonstrate that GAVIS accurately quantifies uncertainty and enables effective exploration, particularly by assigning higher uncertainty to unseen regions, we present qualitative results of the quantified uncertainty maps. We train all methods with the same set of views covering only part of a room, leaving remaining areas underexplored to mimic a robot exploration process. As shown in Fig. 6, GAVIS successfully identifies unseen regions, such as areas behind doors or walls, consistent with the ground-truth visibility derived from the mesh (brighter indicates unseen during training, darker indicates seen). In contrast, other 3DGS-based baselines fail to detect these regions, assigning falsely low uncertainty and thus reducing their incentive to explore them. While NVF captures some underexplored regions (e.g., behind walls), its neural network-based approximation struggles to accurately detect smaller unseen areas (e.g., behind doors). Moreover, we emphasize that a reliable UQ method should assign high uncertainty to regions with low GT visibility; however, high uncertainty does not necessarily imply low GT visibility (e.g., when the test view direction significantly deviates from training views direction, see Sec. 11.4 for details). We additionally evaluate uncertainty quantification on 5 representative scenes across all datasets using AUSE-D and AUSE-V, as shown in Tab. 4. GAVIS achieves the best results on both metrics, demonstrating the superior quality of its uncertainty estimates.

GAVIS is significantly more computationally efficient than existing methods. We conduct extensive runtime evaluations to demonstrate the speed advantage of GAVIS over existing baselines. To ensure fair comparison, all methods are trained with the same views obtained from a representative trajectory that fully explores each scene, and

runtime is measured at five evenly spaced steps along the trajectory, mimicking different levels of scene exploration. Experiments were run on a single NVIDIA A40 GPU. As shown in Tab. 1, our method significantly outperforms all baselines in both uncertainty preparation and quantification time. Compared to our closest counterpart, NVF, which is the strongest baseline in active mapping performance, GAVIS achieves about $500\times$ speedup in visibility field construction, thanks to our efficient spherical-harmonics-based construction and fast visibility query algorithm. It is also $30\times$ faster in uncertainty quantification. GAVIS further surpasses other 3DGS-based baselines: FisherRF requires gradient backpropagation through the 3DGS rasterizer, leading to high computational cost for uncertainty preparation and quantification; while VIMC must co-train a stochastic radiance field and, during inference, perform multiple rasterizations from each sample of the field parameters, resulting in substantially higher runtime.

Ablation studies show that each major component of GAVIS plays a crucial role. We ablate two key components of our method: (1) anisotropic visibility and (2) visibility field density control. Results averaged over all datasets are reported in Tab. 2. Removing either component causes a noticeable drop across all metrics, while jointly removing both reduces the model to a direct 3DGS extension of NVF, yielding the worst performance. We include additional qualitative results in the appendix Fig. 7 to further illustrate the importance of each component.

GAVIS can be applied post-hoc to improve the performance of existing active mapping methods. We integrate GAVIS into the existing baselines FisherRF and VIMC as a post-hoc module to refine their field uncertainty estimation, assigning higher uncertainty to low-visibility regions. Details of the derivation and implementation are provided in appendix Sec. 10. Results averaged over all datasets are reported in Tab. 3. GAVIS significantly improves the performance of both baseline methods, highlighting the crucial role of visibility in active mapping. Notably, FisherRF+GAVIS only marginally outperforms GAVIS alone, consistent with the observation in [104] that visibility modeling is the dominant factor in effective active mapping. Meanwhile, VIMC+GAVIS still underperforms GAVIS, as VIMC relies on sampling-based uncertainty estimation, where the noise introduced by sampling overshadows the benefit of visibility modeling.

6. Conclusion

In this work, we present Gaussian Splatting Anisotropic Visibility Field (GAVIS), a principled approach to quantifying uncertainty in 3D Gaussian Splatting through accurate visibility modeling. We propose an analytical method for computing the anisotropic visibility field in 3DGS using spherical harmonics, representing a substantial improve-




Figure 4. **Qualitative active mapping.** Reconstruction results and camera-view distributions (green frustums) from different methods’ active-mapping trajectories on Gibson scene (top) and HM3D scene (bottom). Full results are provided in Sec. 12.




Figure 5. **Qualitative active mapping.** Reconstruction results and camera-view distributions (green frustums) from different methods’ active-mapping trajectories on HST scene (top) and Lego scene (bottom). Full results are provided in Sec. 12.




Figure 6. **Qualitative uncertainty estimation.** All methods are trained on the same set of views that only partially cover the scene, leaving some regions underexplored. GT Vis. indicates rasterized ground-truth mesh visibility (binary face labels), where brighter denotes higher uncertainty (invisible faces) and darker denotes lower uncertainty (visible faces). Our method accurately assigns high uncertainty to invisible regions, aligning with GT Vis.

ment over prior work that relies on training neural networks, and achieving significant gains in both efficiency and accuracy. Empirically, GAVIS consistently outperforms all baselines across diverse evaluation scenes. The proposed anisotropic visibility field is not limited to active mapping and can potentially be extended to other tasks that require visibility estimation in 3DGS.

Acknowledgments

This work is supported by NSF grant 2101250. We thank the anonymous reviewers for their valuable suggestions.

References

- [1] Milton Abramowitz and Irene A Stegun. *Handbook of mathematical functions with formulas, graphs, and mathematical tables*. US Government printing office, 1948. 14
- [2] Lewis AG Stuart, Andrew Morton, Ian Stavness, and Michael P Pound. 3dgs-to-pc: 3d gaussian splatting to dense point clouds. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3730–3739, 2025. 22
- [3] NASA Visualization Technology Applications and Development (VTAD). Hubble space telescope 3d model. <https://science.nasa.gov/resource/hubble-space-telescope-3d-model-2/>, 2019. 21
- [4] NASA Visualization Technology Applications and Development (VTAD). International space station 3d model. <https://science.nasa.gov/resource/international-space-station-3d-model/>, 2019. 21
- [5] Nikolay Atanasov, Jerome Le Ny, Kostas Daniilidis, and George J Pappas. Decentralized active information acquisition: Theory and application to multi-robot slam. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4775–4782. IEEE, 2015. 2
- [6] Ana Batinovic, Antun Ivanovic, Tamara Petrovic, and Stjepan Bogdan. A shadowcasting-based next-best-view planner for autonomous 3d exploration. *IEEE Robotics and Automation Letters*, 7(2):2969–2976, 2022. 2
- [7] Andreas Bircher, Mina Kamel, Kostas Alexis, Helen Oleynikova, and Roland Siegwart. Receding horizon” next-best-view” planner for 3d exploration. In *2016 IEEE international conference on robotics and automation (ICRA)*, pages 1462–1468. IEEE, 2016. 1, 2
- [8] Andreas Bircher, Mina Kamel, Kostas Alexis, Helen Oleynikova, and Roland Siegwart. Receding horizon path planning for 3d exploration and surface inspection. *Autonomous Robots*, 42(2):291–306, 2018. 1, 2
- [9] Luca Carlone, Jingjing Du, Miguel Kaouk Ng, Basilio Bona, and Marina Indri. Active slam and exploration with particle filters using kullback-leibler divergence. *Journal of Intelligent & Robotic Systems*, 75(2):291–311, 2014. 2
- [10] Henry Carrillo, Ian Reid, and José A Castellanos. On the comparison of uncertainty criteria for active slam. In *2012 IEEE International Conference on Robotics and Automation*, pages 2080–2087. IEEE, 2012. 2
- [11] Fanfei Chen, John D Martin, Yewei Huang, Jinkun Wang, and Brendan Englot. Autonomous exploration under uncertainty via deep reinforcement learning on graphs. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6140–6147. IEEE, 2020. 2
- [12] Liyan Chen, Huangying Zhan, Kevin Chen, Xiangyu Xu, Qingan Yan, Changjiang Cai, and Yi Xu. Activegamer: Active gaussian mapping through efficient rendering. *arXiv preprint arXiv:2501.06897*, 2025. 2
- [13] Xingyu Chen, Qi Zhang, Xiaoyu Li, Yue Chen, Ying Feng, Xuan Wang, and Jue Wang. Hallucinated neural radiance fields in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12943–12952, 2022. 2
- [14] Cl Connolly. The determination of next best views. In *Proceedings. 1985 IEEE international conference on robotics and automation*, pages 432–435. IEEE, 1985. 2
- [15] Anna Dai, Sotiris Papatheodorou, Nils Funk, Dimos Tzoumanikas, and Stefan Leutenegger. Fast frontier-based information-driven autonomous exploration with an mav. In *2020 IEEE international conference on robotics and automation (ICRA)*, pages 9570–9576. IEEE, 2020. 2
- [16] Tung Dang, Frank Mascarich, Shehryar Khattak, Christos Papachristos, and Kostas Alexis. Graph-based path planning for autonomous robotic exploration in subterranean environments. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 3105–3112. IEEE, 2019. 2
- [17] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12882–12891, 2022. 21
- [18] Christian Dornhege and Alexander Kleiner. A frontier-void-based approach for autonomous exploration in 3d. *Advanced Robotics*, 27(6):459–468, 2013. 2
- [19] 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
- [20] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5501–5510, 2022. 2, 5
- [21] Yarin Gal and Zoubin Ghahramani. Dropout as a bayesian approximation: Representing model uncertainty in deep learning. In *international conference on machine learning*, pages 1050–1059. PMLR, 2016. 2
- [22] Jakob Gawlikowski, Cedrique Rovile Njiteucheu Tassi, Mohsin Ali, Jongseok Lee, Matthias Humt, Jianxiang Feng, Anna Kruspe, Rudolph Triebel, Peter Jung, Ribana Roscher, et al. A survey of uncertainty in deep neural networks. *Artificial Intelligence Review*, 56(Suppl 1):1513–1589, 2023. 1
- [23] Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. Bayes’ rays: Uncertainty quantification for neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20061–20070, 2024. 2, 7, 20, 22

- [24] Clara Gomez, Alejandra C Hernandez, and Ramon Barber. Topological frontier-based exploration and map-building using semantic information. *Sensors*, 19(20):4595, 2019. [2](#)
- [25] Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q Weinberger. On calibration of modern neural networks. In *International conference on machine learning*, pages 1321–1330. PMLR, 2017. [22](#)
- [26] Dirk Holz, Nicola Basilico, Francesco Amigoni, and Sven Behnke. Evaluating the efficiency of frontier-based exploration strategies. In *ISR 2010 (41st International Symposium on Robotics) and ROBOTIK 2010 (6th German Conference on Robotics)*, pages 1–8. VDE, 2010. [2](#)
- [27] Xuan Huang, Hanhui Li, Zejun Yang, Zhisheng Wang, and Xiaodan Liang. 3d visibility-aware generalizable neural radiance fields for interacting hands. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 2400–2408, 2024. [2](#)
- [28] Marco F Huber, Tim Bailey, Hugh Durrant-Whyte, and Uwe D Hanebeck. On entropy approximation for gaussian mixture random vectors. In *2008 IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pages 181–188. IEEE, 2008. [18](#)
- [29] Eddy Ilg, Ozgun Cicek, Silvio Galesso, Aaron Klein, Osama Makansi, Frank Hutter, and Thomas Brox. Uncertainty estimates and multi-hypotheses networks for optical flow. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 652–667, 2018. [7](#), [22](#)
- [30] John David Jackson. *Classical Electrodynamics*. John Wiley & Sons, Inc., Hoboken, NJ, 3rd edition, 1999. [14](#)
- [31] Wen Jiang, Boshu Lei, Katrina Ashton, and Kostas Daniilidis. Ag-slam: Active gaussian splatting slam. 2024. [19](#)
- [32] Wen Jiang, Boshu Lei, and Kostas Daniilidis. Fisherrf: Active view selection and uncertainty quantification for radiance fields using fisher information. *ECCV*, 2024. [1](#), [2](#), [6](#), [7](#), [16](#), [18](#), [20](#), [22](#)
- [33] Liren Jin, Xieyuanli Chen, Julius Rücker, and Marija Popović. Neu-nbv: Next best view planning using uncertainty estimation in image-based neural rendering. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11305–11312. IEEE, 2023. [2](#), [22](#)
- [34] Liren Jin, Xingguang Zhong, Yue Pan, Jens Behley, Cyrill Stachniss, and Marija Popović. Activegs: Active scene reconstruction using gaussian splatting. *IEEE Robotics and Automation Letters*, 2025. [2](#)
- [35] Rui Jin, Yuman Gao, Yingjian Wang, Yuze Wu, Haojian Lu, Chao Xu, and Fei Gao. Gs-planner: A gaussian-splatting-based planning framework for active high-fidelity reconstruction. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 11202–11209. IEEE, 2024. [2](#)
- [36] 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](#)
- [37] Matan Keidar and Gal A Kaminka. Robot exploration with fast frontier detection: Theory and experiments. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 1*, pages 113–120, 2012. [2](#)
- [38] Matan Keidar and Gal A Kaminka. Efficient frontier detection for robot exploration. *The International Journal of Robotics Research*, 33(2):215–236, 2014. [2](#)
- [39] 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](#), [5](#), [15](#), [17](#), [21](#)
- [40] Andreas Kirsch and Yarin Gal. Unifying approaches in active learning and active sampling via fisher information and information-theoretic quantities. *arXiv preprint arXiv:2208.00549*, 2022. [19](#), [20](#)
- [41] Marcus Klasson, Riccardo Mereu, Juho Kannala, and Arno Solin. Sources of uncertainty in 3d scene reconstruction. In *European Conference on Computer Vision*, pages 271–289. Springer, 2025. [2](#)
- [42] Thomas Kollar and Nicholas Roy. Trajectory optimization using reinforcement learning for map exploration. *The International Journal of Robotics Research*, 27(2):175–196, 2008. [2](#)
- [43] Yves Kompis, Luca Bartolomei, Ruben Mascaro, Lucas Teixeira, and Margarita Chli. Informed sampling exploration path planner for 3d reconstruction of large scenes. *IEEE Robotics and Automation Letters*, 6(4):7893–7900, 2021. [2](#)
- [44] Georgios Kopanas and George Drettakis. Improving nerf quality by progressive camera placement for free-viewpoint navigation. 2023. [2](#)
- [45] Zijia Kuang, Zike Yan, Hao Zhao, Guyue Zhou, and Hongbin Zha. Active neural mapping at scale. In *2024 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 7152–7159. IEEE, 2024. [2](#)
- [46] 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*, 7(4):12070–12077, 2022. [2](#), [3](#), [19](#)
- [47] Sibeak Lee, Kyeongsu Kang, Seongbo Ha, and Hyeonwoo Yu. Bayesian nerf: Quantifying uncertainty with volume density for neural implicit fields. *arXiv preprint arXiv:2404.06727*, 2024. [2](#)
- [48] Ruiqi Li and Yiu-ming Cheung. Variational multi-scale representation for estimating uncertainty in 3d gaussian splatting. *Advances in Neural Information Processing Systems*, 37:87934–87958, 2024. [2](#)
- [49] You Li, Rui Li, Shengjun Tang, and Yanjie Gu. Bavs-nerf: Batch active view selection for neural radiance field using scene uncertainty. In *The First Workshop on Populating Empty Cities–Virtual Humans for Robotics and Autonomous Driving at CVPR 2024*. [2](#)
- [50] Yuetao Li, Zijia Kuang, Ting Li, Qun Hao, Zike Yan, Guyue Zhou, and Shaohui Zhang. Activesplat: High-fidelity scene reconstruction through active gaussian splatting. *IEEE Robotics and Automation Letters*, 2025. [2](#), [22](#)

- [51] Iker Lluvía, Elena Lazkano, and Ander Ansuategi. Active mapping and robot exploration: A survey. *Sensors*, 21(7): 2445, 2021. 1
- [52] Liang Lu, Carlos Redondo, and Pascual Campoy. Optimal frontier-based autonomous exploration in unconstructed environment using rgb-d sensor. *Sensors*, 20(22):6507, 2020. 2
- [53] Shengjie Luo, Tianlang Chen, and Aditi S Krishnapriyan. Enabling efficient equivariant operations in the fourier basis via gaunt tensor products. *arXiv preprint arXiv:2401.10216*, 2024. 5, 15
- [54] Linjie Lyu, Ayush Tewari, Marc Habermann, Shunsuke Saito, Michael Zollhöfer, Thomas Leimkühler, and Christian Theobalt. Manifold sampling for differentiable uncertainty in radiance fields. In *SIGGRAPH Asia 2024 Conference Papers*, pages 1–11, 2024. 1, 2, 3, 6, 7, 16, 18, 19, 20, 22, 23
- [55] 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, 21
- [56] Zehui Meng, Hailong Qin, Ziyue Chen, Xudong Chen, Hao Sun, Feng Lin, and Marcelo H Ang. A two-stage optimized next-view planning framework for 3-d unknown environment exploration, and structural reconstruction. *IEEE Robotics and Automation Letters*, 2(3):1680–1687, 2017. 2
- [57] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 1, 2, 3, 6, 17, 21
- [58] Beipeng Mu, Matthew Giamou, Liam Paull, Ali-akbar Agha-mohammadi, John Leonard, and Jonathan How. Information-based active slam via topological feature graphs. In *2016 IEEE 55th Conference on decision and control (Cdc)*, pages 5583–5590. IEEE, 2016. 2
- [59] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022. 2, 5
- [60] Kiyohiro Nakayama, Mikaela Angelina Uy, Yang You, Ke Li, and Leonidas J Guibas. Provernerf: Modeling per point provenance in nerfs as a stochastic field. *Advances in Neural Information Processing Systems*, 37:99145–99160, 2024. 2
- [61] Farzad Niroui, Kaicheng Zhang, Zendai Kashino, and Goldie Nejat. Deep reinforcement learning robot for search and rescue applications: Exploration in unknown cluttered environments. *IEEE Robotics and Automation Letters*, 4(2):610–617, 2019. 2
- [62] Frank WJ Olver. *NIST handbook of mathematical functions hardback and CD-ROM*. Cambridge university press, 2010. 5, 14
- [63] Yaniv Ovadia, Emily Fertig, Jie Ren, Zachary Nado, David Sculley, Sebastian Nowozin, Joshua Dillon, Balaji Lakshminarayanan, and Jasper Snoek. Can you trust your model’s uncertainty? evaluating predictive uncertainty under dataset shift. *Advances in neural information processing systems*, 32, 2019. 1
- [64] Xuran Pan, Zihang Lai, Shiji Song, and Gao Huang. Activenerf: Learning where to see with uncertainty estimation. In *European Conference on Computer Vision*, pages 230–246. Springer, 2022. 2
- [65] Christos Papachristos, Shehryar Khattak, and Kostas Alexis. Uncertainty-aware receding horizon exploration and mapping using aerial robots. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 4568–4575. IEEE, 2017. 2
- [66] 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. *IEEE Transactions on Robotics*, 2023. 1, 2, 19
- [67] Matteo Poggi, Filippo Aleotti, Fabio Tosi, and Stefano Mattoccia. On the uncertainty of self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3227–3237, 2020. 7, 22
- [68] Wenchuan Qiao, Zheng Fang, and Bailu Si. Sample-based frontier detection for autonomous robot exploration. In *2018 IEEE International Conference on Robotics and Biomimetics (ROBIO)*, pages 1165–1170. IEEE, 2018. 2
- [69] Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. 6, 21
- [70] Yunlong Ran, Jing Zeng, Shibo He, Jiming Chen, Lincheng Li, Yingfeng Chen, Gimhee Lee, and Qi Ye. Neurar: Neural uncertainty for autonomous 3d reconstruction with implicit neural representations. *IEEE Robotics and Automation Letters*, 8(2):1125–1132, 2023. 2
- [71] Victor Massagué Respal, Dmitry Devitt, Roman Fedorenko, and Alexandr Klimchik. Fast sampling-based next-best-view exploration algorithm for a mav. In *2021 IEEE International Conference on Robotics and Automation (ICRA)*, pages 89–95. IEEE, 2021. 2
- [72] Jun John Sakurai and Jim Napolitano. *Modern quantum mechanics*. Cambridge University Press, 2020. 15
- [73] 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. 21
- [74] Magnus Selin, Mattias Tiger, Daniel Duberg, Fredrik Heintz, and Patric Jensfelt. Efficient autonomous exploration planning of large-scale 3-d environments. *IEEE Robotics and Automation Letters*, 4(2):1699–1706, 2019. 2
- [75] PGCN Senarathne and Danwei Wang. Towards autonomous 3d exploration using surface frontiers. In *2016*

- IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)*, pages 34–41. IEEE, 2016. 2
- [76] Jianxiong Shen, Adria Ruiz, Antonio Agudo, and Francesc Moreno-Noguer. Stochastic neural radiance fields: Quantifying uncertainty in implicit 3d representations. In *2021 International Conference on 3D Vision (3DV)*, pages 972–981. IEEE, 2021. 2
- [77] Jianxiong Shen, Antonio Agudo, Francesc Moreno-Noguer, and Adria Ruiz. Conditional-flow nerf: Accurate 3d modelling with reliable uncertainty quantification. In *European Conference on Computer Vision*, pages 540–557. Springer, 2022. 2
- [78] Jianxiong Shen, Ruijie Ren, Adria Ruiz, and Francesc Moreno-Noguer. Estimating 3d uncertainty field: Quantifying uncertainty for neural radiance fields. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2375–2381. IEEE, 2024. 2
- [79] Shaojie Shen, Nathan Michael, and Vijay Kumar. Stochastic differential equation-based exploration algorithm for autonomous indoor 3d exploration with a micro-aerial vehicle. *The International Journal of Robotics Research*, 31(12):1431–1444, 2012. 2
- [80] Yujiao Shi, Hongdong Li, and Xin Yu. Self-supervised visibility learning for novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9675–9684, 2021. 2
- [81] Meng-Li Shih, Wei-Chiu Ma, Lorenzo Boyce, Aleksander Holynski, Forrester Cole, Brian Curless, and Janne Kontkanen. Extranerf: Visibility-aware view extrapolation of neural radiance fields with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20385–20395, 2024. 2
- [82] Robert Sim and Nicholas Roy. Global a-optimal robot exploration in slam. In *Proceedings of the 2005 IEEE international conference on robotics and automation*, pages 661–666. IEEE, 2005. 2
- [83] Nagabhushan Somraj and Rajiv Soundararajan. Vip-nerf: Visibility prior for sparse input neural radiance fields. In *ACM SIGGRAPH 2023 conference proceedings*, pages 1–11, 2023. 2
- [84] Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7495–7504, 2021. 2
- [85] 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
- [86] Cyrill Stachniss, Giorgio Grisetti, and Wolfram Burgard. Information gain-based exploration using rao-blackwellized particle filters. In *Robotics: Science and systems*, pages 65–72, 2005. 2
- [87] 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. 22
- [88] 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
- [89] Niko Sünderhauf, Jad Abou-Chakra, and Dmitry Miller. Density-aware nerf ensembles: Quantifying predictive uncertainty in neural radiance fields. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pages 9370–9376. IEEE, 2023. 2
- [90] Donny Sutantyo, Paul Levi, Christoph Möslinger, and Mark Read. Collective-adaptive lévy flight for underwater multi-robot exploration. In *2013 IEEE International Conference on Mechatronics and Automation*, pages 456–462. IEEE, 2013. 2
- [91] Yuezhan Tao, Dexter Ong, Varun Murali, Igor Spasojevic, Pratik Chaudhari, and Vijay Kumar. Rt-guide: Real-time gaussian splatting for information-driven exploration. *IEEE Robotics and Automation Letters*, 2025. 19
- [92] Barry N Taylor, Chris E Kuyatt, et al. *Guidelines for evaluating and expressing the uncertainty of NIST measurement results*. US Department of Commerce, Technology Administration, National Institute of Standards and Technology, 1994. 21
- [93] Hassan Umari and Shayok Mukhopadhyay. Autonomous robotic exploration based on multiple rapidly-exploring randomized trees. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1396–1402. IEEE, 2017. 2
- [94] Rafael Valencia and Juan Andrade-Cetto. Active pose slam. In *Mapping, planning and exploration with pose SLAM*, pages 89–108. Springer, 2017. 2
- [95] Deng-Bao Wang, Lei Feng, and Min-Ling Zhang. Rethinking calibration of deep neural networks: Do not be afraid of overconfidence. *Advances in Neural Information Processing Systems*, 34:11809–11820, 2021. 1
- [96] Frederik Warburg, Ethan Weber, Matthew Tancik, Aleksander Holynski, and Angjoo Kanazawa. Nerfbusters: Removing ghostly artifacts from casually captured nerfs. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 18120–18130, 2023. 2
- [97] Shuhuan Wen, Yanfang Zhao, Xiao Yuan, Zongtao Wang, Dan Zhang, and Luigi Manfredi. Path planning for active slam based on deep reinforcement learning under unknown environments. *Intelligent Service Robotics*, 13(2):263–272, 2020. 2
- [98] Peter Whaite and Frank P Ferrie. Autonomous exploration: Driven by uncertainty. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 19(3):193–205, 1997. 2
- [99] Eugene Wigner. *Group theory: and its application to the quantum mechanics of atomic spectra*. Elsevier, 2012. 5, 15
- [100] Joey Wilson, Marcelino Almeida, Min Sun, Sachit Mahajan, Maani Ghaffari, Parker Ewen, Omid Ghasemalizadeh, Cheng-Hao Kuo, and Arnie Sen. Modeling uncertainty in 3d gaussian splatting through continuous semantic splat-

- ting. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3284–3290. IEEE, 2025. [2](#)
- [101] Fei Xia, Amir R Zamir, Zhiyang He, Alexander Sax, Jitendra Malik, and Silvio Savarese. Gibson env: Real-world perception for embodied agents. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9068–9079, 2018. [6](#), [21](#)
- [102] Hangao Xin, Zhiqian Zhou, Di An, Ling-Qi Yan, Kun Xu, Shi-Min Hu, and Shing-Tung Yau. Fast and accurate spherical harmonics products. *ACM Trans. Graph.*, 40(6):280–1, 2021. [15](#)
- [103] Zijun Xu, Rui Jin, Ke Wu, Yi Zhao, Zhiwei Zhang, Jieru Zhao, Fei Gao, Zhongxue Gan, and Wenchao Ding. Hgs-planner: Hierarchical planning framework for active scene reconstruction using 3d gaussian splatting. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 14161–14167. IEEE, 2025. [2](#)
- [104] Shangjie Xue, Jesse Dill, Pranay Mathur, Frank Dellaert, Panagiotis Tsiotras, and Danfei Xu. Neural visibility field for uncertainty-driven active mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18122–18132, 2024. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#), [17](#), [18](#), [19](#), [21](#), [22](#), [23](#)
- [105] 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. [1](#), [2](#)
- [106] Brian Yamauchi. Frontier-based exploration using multiple robots. In *Proceedings of the second international conference on Autonomous agents*, pages 47–53, 1998. [2](#)
- [107] Dongyu Yan, Jianheng Liu, Fengyu Quan, Haoyao Chen, and Mengmeng Fu. Active implicit object reconstruction using uncertainty-guided next-best-view optimization. *IEEE Robotics and Automation Letters*, 2023. [2](#), [3](#), [19](#)
- [108] Jing Zeng, Qi Ye, Tianle Liu, Yang Xu, Jin Li, Jinming Xu, Liang Li, and Jiming Chen. Multi-robot autonomous 3d reconstruction using gaussian splatting with semantic guidance. *IEEE Robotics and Automation Letters*, 2025. [2](#)
- [109] Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d gaussian splatting for unconstrained image collections. In *European Conference on Computer Vision*, pages 341–359. Springer, 2024. [2](#)
- [110] Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 586–595, 2018. [6](#)
- [111] Xiuming Zhang, Pratul P Srinivasan, Boyang Deng, Paul Debevec, William T Freeman, and Jonathan T Barron. Nerfactor: Neural factorization of shape and reflectance under an unknown illumination. *ACM Transactions on Graphics (ToG)*, 40(6):1–18, 2021. [2](#)
- [112] Yikang Zhang and Rui Fan. Visibility-aware densification for 3d gaussian splatting in dynamic urban scenes. *arXiv preprint arXiv:2510.09364*, 2025. [2](#)
- [113] Zichao Zhang and Davide Scaramuzza. Perception-aware receding horizon navigation for mavs. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2534–2541. IEEE, 2018. [2](#)
- [114] Yuke Zhu, Roozbeh Mottaghi, Eric Kolve, Joseph J Lim, Abhinav Gupta, Li Fei-Fei, and Ali Farhadi. Target-driven visual navigation in indoor scenes using deep reinforcement learning. In *2017 IEEE international conference on robotics and automation (ICRA)*, pages 3357–3364. IEEE, 2017. [2](#)

Uncertainty-driven 3D Gaussian Splatting Active Mapping via Anisotropic Visibility Field

Supplementary Material

7. Anisotropic Visibility Field Details

In this section, we provide details on anisotropic visibility field $V^{(i)}(\mathbf{d})$. We first derive the spherical harmonics (SH) representation of the directional visibility function $\nu(\mathbf{d}; \mathbf{d}_p)$ (Sec. 7.1), then discuss the SH representation of the visibility field $V^{(i)}(\mathbf{d})$ (Sec. 7.2, 7.3), and finally describe how $V^{(i)}(\mathbf{d})$ is efficiently constructed and queried during uncertainty-aware rasterization (Sec. 7.4).

7.1. Derivation of (8)

We represent $\nu(\mathbf{d}; \mathbf{d}_p)$ in the orthonormal basis of spherical harmonics $Y_\ell^m(\mathbf{d})$,

$$\begin{aligned} \nu(\mathbf{d}; \mathbf{d}_p) &:= \zeta \exp(\kappa \mathbf{d} \cdot \mathbf{d}_p) \\ &= \zeta \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} a_{\ell m} Y_\ell^m(\mathbf{d}), \end{aligned} \quad (12)$$

and derive an analytical expression for the coefficients $a_{\ell m}$.

Since $\exp(\kappa \mathbf{d} \cdot \mathbf{d}_p)$ depends only on the dot-product $x = \mathbf{d} \cdot \mathbf{d}_p$, we first expand in the one-dimensional basis of Legendre polynomials $P_\ell(\mathbf{d} \cdot \mathbf{d}_p)$ [1]:

$$\exp(\kappa \mathbf{d} \cdot \mathbf{d}_p) = \sum_{\ell=0}^{\infty} b_\ell P_\ell(\mathbf{d} \cdot \mathbf{d}_p), \quad (13)$$

By comparing (13) with the Legendre expansion of the exponential [62]:

$$e^{\kappa x} = \sum_{\ell=0}^{\infty} (2\ell+1) i_\ell(\kappa) P_\ell(x).$$

Thus, we can identify the coefficients b_ℓ as

$$b_\ell = (2\ell+1) i_\ell(\kappa). \quad (14)$$

where i_ℓ is the modified spherical Bessel function of the first kind [62].

Then we apply the addition theorem for spherical harmonics [30]:

$$P_\ell(\mathbf{d} \cdot \mathbf{d}_p) = \frac{4\pi}{2\ell+1} \sum_{m=-\ell}^{\ell} Y_\ell^m(\mathbf{d}) Y_\ell^{m*}(\mathbf{d}_p). \quad (15)$$

where $Y_\ell^{m*}(\mathbf{d}_p)$ is the complex conjugate of $Y_\ell^m(\mathbf{d}_p)$.

Substituting (15) and (14) into (13) gives

$$\nu(\mathbf{d}; \mathbf{d}_p) = 4\pi\zeta \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} i_\ell(\kappa) Y_\ell^m(\mathbf{d}) Y_\ell^{m*}(\mathbf{d}_p). \quad (16)$$

Comparing with (7), we identify the coefficients

$$a_{\ell m} = 4\pi i_\ell(\kappa) Y_\ell^{m*}(\mathbf{d}_p). \quad (17)$$

where the modified spherical Bessel function of the first kind $i_\ell(\kappa)$ can be precomputed since κ is a fixed hyperparameter (set to $\kappa = 1$).

Note that $\nu(\mathbf{d}; \mathbf{d}_p) : \mathbb{S}^2 \rightarrow [0, 1]$ is real-valued, so in practice we expand it in a real spherical harmonics basis for implementation.

$$\nu(\mathbf{d}; \mathbf{d}_p) = 4\pi\zeta \sum_{\ell=0}^{\infty} \sum_{m=-\ell}^{\ell} i_\ell(\kappa) Y_{\ell m}(\mathbf{d}) Y_{\ell m}(\mathbf{d}_p). \quad (18)$$

where $Y_{\ell m}(\mathbf{d})$ is the real spherical harmonics function defined as

$$Y_{\ell m}(\mathbf{d}) = \begin{cases} \sqrt{2} (-1)^m \text{Im} [Y_\ell^{|m|}(\mathbf{d})], & m < 0 \\ Y_\ell^0(\mathbf{d}), & m = 0 \\ \sqrt{2} (-1)^m \text{Re} [Y_\ell^m(\mathbf{d})], & m > 0 \end{cases}$$

and it could be verified that (18) is equivalent to (16).

Therefore, we obtain a closed-form expression for the coefficients representing the directional visibility function $\nu(\mathbf{d}; \mathbf{d}_p)$ in spherical harmonics.

7.2. Details of SH Representation of $V^{(i)}(\mathbf{d})$

In this subsection, we provide further details on why computing the anisotropic visibility field $V^{(i)}(\mathbf{d})$ by directly combining (7) and (5) is impractical due to computational complexity.

We first rewrite (5) in a recursive form:

$$V_{\mathcal{P} \cup \{\mathbf{p}\}}^{(i)}(\mathbf{d}) = V_{\mathcal{P}}^{(i)}(\mathbf{d}) + V_{\mathbf{p}}^{(i)}(\mathbf{d}) - V_{\mathcal{P}}^{(i)}(\mathbf{d}) V_{\mathbf{p}}^{(i)}(\mathbf{d}) \quad (19)$$

starting from $V_{\emptyset}^{(i)}(\mathbf{d}) = 0$, and then recursively computing the anisotropic visibility field over the entire training set. Since spherical harmonics form a linear basis, the sum of two functions expressed in spherical harmonics can be obtained by simply adding their corresponding coefficients $a_{\ell m}$. However, their product poses significant challenges, particularly, the product of two spherical harmonics can be re-expanded as

$$\begin{aligned} Y_{\ell_1}^{m_1}(\mathbf{d}) Y_{\ell_2}^{m_2}(\mathbf{d}) &= \sum_{L=|\ell_1-\ell_2|}^{\ell_1+\ell_2} \sum_{M=-L}^L \sqrt{\frac{(2\ell_1+1)(2\ell_2+1)}{4\pi(2L+1)}} \\ &\times C_{\ell_1 0 \ell_2 0}^{L 0} C_{\ell_1 m_1 \ell_2 m_2}^{L M} Y_L^M(\mathbf{d}). \end{aligned} \quad (20)$$




Figure 7. Illustration of the two major components of GAVIS: (1) anisotropic visibility (effects highlighted with gray boxes) and (2) visibility-field density control (effects highlighted with red boxes). From left to right: (i) top-down view of the scene with training views (blue frustums) covering only the right room and the queried view (green frustum) for uncertainty quantification; (ii) synthesized RGB image from a 3DGS model trained only on the partially covered region, leaving the rest of the scene unexplored; (iii) uncertainty estimated by full GAVIS; (iv) GAVIS with isotropic visibility; (v) GAVIS without visibility field density control; (vi) GAVIS with isotropic visibility and without density control. All methods are applied to the same learned 3DGS model.

where $C_{\ell_1 m_1 \ell_2 m_2}^{LM}$ denote the Clebsch–Gordan coefficients [72, 99]. As is shown in the above equation, multiplying two spherical harmonics of degree L produces a spherical harmonics expansion of degree up to $2L$. Since the multi-view aggregation in (5) requires performing the multiplication $|\mathcal{P}|$ times, the resulting degree can grow to $|\mathcal{P}|L$, requiring $O(|\mathcal{P}|^2 L^2)$ parameters to store. Since multiplying two spherical harmonics of degree L can be done in $O(L^3)$ time [53, 102], computing the anisotropic visibility field over the entire training set $\mathcal{P}$ results in an overall complexity of $O(|\mathcal{P}|^4 L^3)$. Therefore, obtaining a closed-form expression for the spherical harmonics coefficients of the anisotropic visibility field is impractical.

7.3. Derivation of (9)

Starting from (5), we have

$$1 - V^{(i)}(\mathbf{d}) = \prod_{\mathbf{p} \in \mathcal{P}} (1 - V_{\mathbf{p}}^{(i)}(\mathbf{d})) \quad (21)$$

Applying the arithmetic-geometric mean (AM-GM) inequality to the right-hand side of (21) yields

$$\sqrt[|\mathcal{P}|]{\prod_{\mathbf{p} \in \mathcal{P}} (1 - V_{\mathbf{p}}^{(i)}(\mathbf{d}))} \leq \frac{1}{|\mathcal{P}|} \sum_{\mathbf{p} \in \mathcal{P}} (1 - V_{\mathbf{p}}^{(i)}(\mathbf{d})) \quad (22)$$

Therefore, combining (21) and (22) yields the lower bound on $V^{(i)}(\mathbf{d})$.

$$V^{(i)}(\mathbf{d}) \geq 1 - \left(1 - \frac{\sum_{\mathbf{p} \in \mathcal{P}} V_{\mathbf{p}}^{(i)}(\mathbf{d})}{|\mathcal{P}|}\right)^{|\mathcal{P}|} \quad (23)$$

which is equivalent to (9) since $\tilde{V}(\mathbf{d}) = \sum_{\mathbf{p} \in \mathcal{P}} V_{\mathbf{p}}^{(i)}(\mathbf{d})$.

This bound is tight when all $V_{\mathbf{p}}^{(i)}(\mathbf{d})$ are equal, i.e., when the visibility contributions from all camera views $\mathbf{p}$ are the same. In practice, this bound provides a computationally efficient way to estimate the overall visibility field $V^{(i)}(\mathbf{d})$ without having to compute the product of spherical harmonics of $V_{\mathbf{p}}^{(i)}(\mathbf{d})$ over all camera poses $\mathbf{p} \in \mathcal{P}$.

7.4. Details of Visibility Field Construction and Query

We provide additional details on constructing the anisotropic visibility field and on efficient querying. Firstly, we describe our efficient algorithm for computing the single-view (isotropic) particle visibility $\Phi_{i,\mathbf{p}} T_{\mathbf{p}}(t_i^{\mathbf{p}})$ for each particle i in training view $\mathbf{p}$ in Alg. 2, which builds on a modified 3DGS rasterizer.

Note that the averaging step at L15 is necessary when 3D Gaussian particles span multiple pixels during rasterization. Since anisotropic visibility is modeled solely as a function of viewing direction, similar to view-dependent color in 3DGS [39], different rays from the same camera can interact with different points on the same particle, while sharing the same viewing direction. Therefore, averaging their transmittance contributions provides a consistent way to aggregate them into a single per-view value.

Next, we construct the anisotropic visibility field $V^{(i)}(\mathbf{d})$ by computing the SH coefficients $\gamma_{\ell m}^{(i)}$ for each particle i , accumulating contributions from all training views $\mathcal{P}$ via (11) as outlined in Alg. 3. The single-view visibility terms $\Phi_{i,\mathbf{p}}, T_{\mathbf{p}}(t_i^{\mathbf{p}})$ are obtained using Alg. 2.

Finally, we describe the algorithm for querying the anisotropic visibility field at a direction $\mathbf{d}$ in Alg. 4. Given

Algorithm 2 Single-view Particle Visibility

Input: $\mathbf{u}_g \in \mathbb{R}^2$: image-space projection of the Gaussian center to training view $\mathbf{p}$; $\Sigma_g \in \mathbb{R}^{2 \times 2}$: image-space conic (precision) matrix; $z_g \in \mathbb{R}$: depth (for front-to-back sorting); $o_g \in [0, 1]$: base opacity; ϵ_T : transmittance cutoff (default 10^{-4});

Output: $\{v_g\}_{g=1}^N$: isotropic particle visibility with respect to training view $\mathbf{p}$

```

1: Initialize  $v_g \leftarrow 0$ ,  $c_g \leftarrow 0$  for all  $g$ 
2: for each pixel  $\mathbf{x}$  in the image do
3:    $T \leftarrow 1$ 
4:    $\mathcal{G} \leftarrow$  Gaussians overlapping  $\mathbf{x}$  (sorted by increasing  $z$ )
5:   for each  $g \in \mathcal{G}$  do
6:      $\alpha \leftarrow o_g \exp(-(\mathbf{u}_g - \mathbf{x})^\top \Sigma_g (\mathbf{u}_g - \mathbf{x}))$ 
7:      $v_g \leftarrow v_g + T$ ;  $c_g \leftarrow c_g + 1$ 
8:      $T \leftarrow T(1 - \alpha)$ 
9:     if  $T < \epsilon_T$  then break
10:    end if
11:  end for
12: end for
13: for  $g = 1$  to  $N$  do
14:   if  $c_g > 0$  then
15:     $v_g \leftarrow v_g / c_g$  ▷ Average when a particle overlaps multiple pixels.
16:   end if
17: end for
18: Return  $\{v_g\}_{g=1}^N$ 

```

Algorithm 3 Anisotropic Visibility Field Construction

Input: Training views $\mathcal{P}$ with camera centers $\{\mathbf{x}_p\}$; 3DGS particles $\mathcal{G}$ with centers $\{\mathbf{x}_g\}$; SH degree L ; concentration κ ; scale $\zeta = \exp(-\kappa)$

Output: Coefficients $\gamma \in \mathbb{C}^{|\mathcal{G}| \times (L+1)^2}$

```

1: Precompute  $\{i_\ell(\kappa)\}_{\ell=0}^L$  ▷ Modified spherical Bessel function of the first kind
2:  $a_\ell \leftarrow 4\pi \zeta i_\ell(\kappa)$  for  $\ell = 0, \dots, L$ 
3: Initialize  $\gamma_{\ell m}^{(g)} \leftarrow 0 \quad \forall g \in \mathcal{G}, 0 \leq \ell \leq L, -\ell \leq m \leq \ell$ 
4: for  $\mathbf{p} \in \mathcal{P}$  do ▷ Accumulate per-view contributions
5:    $\mathbf{v} \leftarrow \text{SINGLEVIEWPARTICLEVISIBILITY}(\mathbf{p}, \mathcal{G})$  ▷
6:    $\mathbf{v}[g] = \Phi_{g, \mathbf{p}} T_{\mathbf{p}}(t_g^{\mathbf{p}})$  via Alg. 2
7:   for all  $g \in \mathcal{G}$  do
8:      $\mathbf{d}_p^{(g)} \leftarrow \frac{\mathbf{x}_g - \mathbf{x}_p}{\|\mathbf{x}_g - \mathbf{x}_p\|}$ 
9:      $\gamma_{\ell m}^{(g)} \leftarrow \gamma_{\ell m}^{(g)} + \mathbf{v}[g] \cdot a_\ell \cdot Y_{\ell m}(\mathbf{d}_p^{(g)})$  for  $\ell = 0:L$ ,  $m = -\ell:\ell$ 
10:  end for
11: end for
12: return  $\{\gamma_{\ell m}^{(g)}\}$ 

```

Algorithm 4 Query Anisotropic Visibility Field

Input: SH coefficients $\{\gamma_{\ell m}^{(g)}\}$ for particle g ; query direction $\mathbf{d}$; SH degree L ; number of training views $|\mathcal{P}|$

Output: Visibility $V^{(g)}(\mathbf{d})$ for particle g in direction $\mathbf{d}$

```

1:  $y[\ell, m] \leftarrow \text{SPHERICALHARMONICS}(\ell, m, \mathbf{d})$  for  $\ell = 0:L, m = -\ell:\ell$ 
2:  $\tilde{V} \leftarrow \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} \gamma_{\ell m}^{(g)} y[\ell, m]$ 
3:  $V^{(g)} \leftarrow 1 - (1 - \frac{\tilde{V}}{|\mathcal{P}|})^{|\mathcal{P}|}$  ▷ AM-GM lower-bound estimator
4: return  $V^{(g)}$ 

```

the stored SH coefficients $\gamma_{\ell m}^{(g)}$ for each particle g , we evaluate the SH basis at the query direction $\mathbf{d}$ to obtain $\tilde{V}(\mathbf{d})$, and then compute the visibility $V^{(g)}(\mathbf{d})$ using the AM-GM lower-bound estimator in (9). Because we use $L = 2$ in practice, each particle requires only 9 SH coefficients, and the query cost is independent of the number of training views $|\mathcal{P}|$, enabling (per particle) constant-time querying.

Moreover, our method is significantly memory efficient. The anisotropic visibility field requires $\sim 28\%$ additional parameters over standard 3DGS, including SH coefficients and virtual particles (see next subsection). Consequently, our uncertainty quantification method incurs substantially lower memory overhead compared to other 3DGS-based approaches. For example, FisherRF [32] introduces 100% additional parameters to store per-parameter uncertainty, while VIMC [54] requires at least 200% overhead to represent the manifold of radiance field parameters.

7.5. Details of Visibility Field Density Control

We use *virtual particles* to distinguish free space from underexplored regions. Given a trained 3DGS, we uniformly sample virtual particles within the scene. Each particle is initialized with zero opacity, identity rotation, and an isotropic scale s . The total number of virtual particles is

$$N_{vp} = \rho \cdot \mathbf{V}_{\text{scene}},$$

where $\mathbf{V}_{\text{scene}}$ is the scene volume and $\rho = 100$ is a density hyperparameter specifying the number of virtual particles per unit volume. The scale s is chosen such that the expected total volume of virtual particles per unit volume of the scene slightly exceeds 1, i.e.,

$$\frac{4}{3}\pi s^3 \rho = 1 + \eta,$$

with $\eta = 0.5$ ensuring that the total particle volume slightly exceeds the scene volume, which encourages broad coverage of the space. Note that s is substantially larger than the scale of the original 3DGS particles. This formulation allows a sparse set of virtual particles to approximate coverage of the entire scene, making it significantly more efficient than using a dense grid of particles.

We then combine virtual particles and original particles to compute the visibility of each virtual particle using the same single-view visibility algorithm as for the original particles, with virtual particles treated as transparent during visibility computation. The multi-view visibility of each virtual particle is then computed as

$$1 - \prod_{p \in \mathcal{P}} (1 - \Phi_{i,p} T_p(t_i^p))$$

Finally, we prune virtual particles using a conservative visibility threshold $\epsilon_v = 0.95$. Particles with visibility greater than ϵ_v are treated as free space and discarded, while those with visibility below ϵ_v are retained as indicators of underexplored regions.

When querying the anisotropic visibility field for virtual particles, we treat all retained virtual particles as invisible in every direction, which is equivalent to setting $\gamma_{\ell m}^{(i)} = 0$ for all ℓ, m . For computational efficiency, however, we simply return $V^{(i)}(\mathbf{d}) = 0$ for all query directions $\mathbf{d}$.

Note that, unlike density control in 3DGS training [39], which is applied repeatedly throughout optimization, the visibility-field density control is a post-hoc procedure and needs to be executed only once after training. This makes it substantially faster in practice. Details of the algorithm are provided in Alg. 5.

Finally, an illustration of the two major components of GAVIS: (1) anisotropic visibility and (2) visibility-field density control, is provided in Fig. 7.

8. Details of Uncertainty-aware 3DGS Rasterization

In this section, we provide additional details on uncertainty-aware 3DGS rasterization, specifically how to quantify the uncertainty of the observed color along a ray direction $\mathbf{d}$ using the anisotropic visibility field $\{V^{(g)}(\mathbf{d})\}_{g \in \mathcal{G}}$. This section extends the uncertainty-aware volume rendering formulation introduced in NVF [104], originally developed for NeRF, to 3DGS.

The rasterization process for 3DGS (along a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$) is expressed as:

$$\hat{C}(\mathbf{r}) = \sum_i w_i c_i(\mathbf{d}), \quad (24)$$

where

$$w_i = \alpha_i \prod_{j=0}^{i-1} (1 - \alpha_j),$$

and $\alpha_i = 1 - \exp(-\sigma_i \delta_i)$ is the opacity of the i -th particle, with σ_i its density and δ_i the interval length. In 3DGS, effective α_i is computed by evaluating a 2D Gaussian with covariance Σ_i at the pixel location, multiplied by the learned per-particle opacity o_i , where Σ_i is obtained by projecting

Algorithm 5 Visibility Field Density Control

Input: Training views $\mathcal{P}$; original 3DGS particles $\mathcal{G}$ with centers $\{\mathbf{x}_g\}$, rotations $\{\mathbf{R}_g\}$, scales $\{s_g\}$, opacities $\{o_g\}$; scene volume $\mathbf{V}_{\text{scene}}$; virtual particle density ρ ; visibility threshold ϵ_v ; hyperparameter η (default 0.5)

Output: Augmented 3DGS $\tilde{\mathcal{G}} = \mathcal{G} \cup \mathcal{G}_{\text{virt}}^{\text{keep}}$ for uncertainty quantification

```

1:  $N_{\text{virt}} \leftarrow \lfloor \rho \cdot |\mathbf{V}_{\text{scene}}| \rfloor$   $\triangleright$  Number of virtual particles
2:  $s \leftarrow \sqrt[3]{\frac{3(1+\eta)}{4\pi\rho}}$   $\triangleright$  Scale of virtual particles
3: Sample virtual centers  $\{\mathbf{x}_j\}_{j=1}^{N_{\text{virt}}}$  uniformly in  $\mathbf{V}_{\text{scene}}$ 
4:  $\mathcal{G}_{\text{virt}} \leftarrow \{(\mathbf{x}_j, \mathbf{I}_3, s, 0)\}_{j=1}^{N_{\text{virt}}}$   $\triangleright$  Virtual particles: identity rotation, scale  $s$ , zero opacity
5:  $\mathcal{G}_{\text{all}} \leftarrow \mathcal{G} \cup \mathcal{G}_{\text{virt}}$ 
6:  $p_g \leftarrow 1 \quad \forall g \in \mathcal{G}_{\text{virt}}$   $\triangleright$  Accumulator for  $\prod_p (1 - \Phi_{g,p} T_p(t_g^p))$ 
7: for all  $p \in \mathcal{P}$  do  $\triangleright$  Accumulate per-view contributions
8:    $\mathbf{v}_p \leftarrow \text{SINGLEVIEWPARTICLEVISIBILITY}(p, \mathcal{G}_{\text{all}})$ 
    $\triangleright \mathbf{v}_p[g] = \Phi_{g,p} T_p(t_g^p)$  via Alg. 2
9:   for all  $g \in \mathcal{G}_{\text{virt}}$  do
10:      $p_g \leftarrow p_g \cdot (1 - \mathbf{v}_p[g])$ 
11:   end for
12: end for
13:  $\mathcal{G}_{\text{virt}}^{\text{keep}} \leftarrow \emptyset$ 
14: for all  $g \in \mathcal{G}_{\text{virt}}$  do
15:    $\tilde{v}_g \leftarrow 1 - p_g$   $\triangleright$  Multi-view visibility
    $1 - \prod_{p \in \mathcal{P}} (1 - \Phi_{g,p} T_p(t_g^p))$ 
16:   if  $\tilde{v}_g \leq \epsilon_v$  then  $\triangleright$  Low visibility  $\Rightarrow$  underexplored region
17:      $\mathcal{G}_{\text{virt}}^{\text{keep}} \leftarrow \mathcal{G}_{\text{virt}}^{\text{keep}} \cup \{(\mathbf{x}_g, \mathbf{I}_3, s, 0)\}$ 
18:   else  $\triangleright$  High visibility  $\Rightarrow$  free space, discard
19:     continue
20:   end if
21: end for
22:  $\tilde{\mathcal{G}} \leftarrow \mathcal{G} \cup \mathcal{G}_{\text{virt}}^{\text{keep}}$ 
23: return  $\tilde{\mathcal{G}}$ 

```

the 3D Gaussian onto the image plane [39]. The emitted color $c_i(\mathbf{d})$ is obtained from the spherical harmonics coefficients of particle i , evaluated at direction $\mathbf{d}$.

To extend (24) to uncertainty-aware rasterization, consider computing the posterior distribution of the observed color $\hat{C}(\mathbf{r})$ along direction $\mathbf{d}$, denoted as $p(\mathbf{z}_0)$. Each emitted color is modeled as a Gaussian random variable $c_i(\mathbf{d}) \sim \mathcal{N}(\mu_{c_i}, \mathbf{Q}_{c_i})$, where μ_{c_i} and $\mathbf{Q}_{c_i}$ are respectively the mean and covariance. Since density σ_i can be interpreted as the differential probability that a ray terminates at position i [57], the effective opacity α_i represents the probability of ray terminating at that position. The rasterization process can thus be expressed as a Bayesian Network [104]

via the recursion

$$p(\mathbf{z}_i) = \alpha_i p(\mathbf{c}_i) + (1 - \alpha_i) p(\mathbf{z}_{i+1}), \quad (25)$$

where $p(\mathbf{z}_i)$ denotes the distribution of the observed color if the camera were placed at position i along the ray. Equation (25) states that with probability α_i a ray terminates at particle i , in which case the observed color is $\mathbf{c}_i$; otherwise, with probability $1 - \alpha_i$, the observed color is inherited from the next position along the ray. The recursion terminates at the last particle, whose emitted color distribution is Gaussian.

By applying the recursive expression in (25) from the last particle to the first, we can compute the posterior of the observed color at the camera position $\mathbf{z}_0$, which is the output of the rasterization process.

$$p(\mathbf{z}_0) = \sum_i w_i \mathcal{N}(\mu_{\mathbf{c}_i}, \mathbf{Q}_{\mathbf{c}_i}), \quad (26)$$

which is a Gaussian mixture model (GMM) with N (number of particles along the ray) components.

As noted in NVF [104], directly quantifying uncertainty from the GMM in (26) is limited because the predicted color variance $\mathbf{Q}_{\mathbf{c}_i}$ is often inaccurate due to approximations in the uncertainty estimation process. To address this issue, NVF [104] introduces a (isotropic) visibility-modulated variance correction which ensures particles with lower visibility exhibit higher uncertainty.

We extend this correction to anisotropic visibility and 3DGS. Let V_i be a binary random variable indicating whether particle i is visible along the ray, with visibility probability given by the visibility field, $P(V_i = 1) = V^{(i)}(\mathbf{d})$. The distribution of the emitted color is therefore modified to

$$p(\mathbf{c}_i | V_i) = \begin{cases} \mathcal{N}(\mu_{\mathbf{c}_i}, \mathbf{Q}_{\mathbf{c}_i}), & \text{if } V_i = 1, \\ \mathcal{N}(\mu_0, \mathbf{Q}_0), & \text{otherwise,} \end{cases} \quad (27)$$

where $\mathcal{N}(\mu_0, \mathbf{Q}_0)$ is a prior with large variance $\mathbf{Q}_0$ to reflect the high uncertainty of invisible particles.

Opacity is also visibility-modulated, since opacity predictions are unreliable for particles that are rarely or never observed. Similar to [104], the compensated opacity for 3DGS is

$$\alpha_i^* = (v_i + \beta(1 - v_i)) \alpha_i + o_0(1 - \beta)(1 - v_i), \quad (28)$$

where $v_i := V^{(i)}(\mathbf{d})$ for shorthand, o_0 is the prior opacity for invisible particles, and β is a hyperparameter representing the reliability of the opacity prediction for invisible regions. The first term in (28) captures the case where the particle is visible or the opacity prediction is reliable. The second term introduces a correction when the particle

is invisible and the predicted opacity is unreliable, by substituting the predicted opacity in place of a prior opacity o_0 . Note that here we use the per-particle opacity o_0 instead of the effective opacity that incorporates the Gaussian decay as a function of the distance between the particle center and the pixel location.

Therefore, the visibility-modulated uncertainty-aware 3DGS rasterization process can be expressed as the following Gaussian mixture model:

$$p(\mathbf{z}_0) = \sum_i w_i^* v_i \mathcal{N}(\mu_{\mathbf{c}_i}, \mathbf{Q}_{\mathbf{c}_i}) + \mathcal{N}(\mu_0, \mathbf{Q}_0) \sum_i w_i^* (1 - v_i), \quad (29)$$

where $w_i^* = \alpha_i^* \prod_{j=1}^{i-1} (1 - \alpha_j^*)$ denotes the visibility-compensated weights.

The entropy of the GMM in (29) is then used to quantify the uncertainty of the observed color $\mathbf{z}_0$. Because the entropy of a Gaussian mixture lacks a closed-form expression, the following upper bound [28], also adopted in NVF [104], is used as an accurate estimator of the entropy:

$$\mathcal{H}(\mathbf{z}_0) \leq \sum_{i=0}^N \bar{w}_i \left(-\log(\bar{w}_i) + \frac{1}{2} \log((2\pi e)^D |\mathbf{Q}_i|) \right), \quad (30)$$

where we denote $\bar{w}_i := w_i^* v_i$ for $i \geq 1$, and $\bar{w}_0 := \sum_{i=1}^N w_i^* (1 - v_i)$, and $\mathbf{Q}_i := \mathbf{Q}_{\mathbf{c}_i}$ for $i \geq 1$ and $\mathbf{Q}_0$ for $i = 0$. Here, $D = 3$ corresponds to the RGB color channels.

It is worth noting that the visibility field v_i is estimated using the lower bound in (23). This approximation underestimates visibility and therefore overestimates how much of the mixture should be attributed to invisible regions. Consequently, the RHS in (30) remains a valid upper bound on the entropy computed with exact visibility, since the inequality is preserved under the assumption $\mathbf{Q}_0 \succeq \mathbf{Q}_{\mathbf{c}_i}$ for all i . This condition is enforced by construction, as $\mathbf{Q}_0$ is chosen to be a large-variance prior.

As demonstrated in [104], when using visibility-modulated uncertainty-aware volume rendering, setting $\mathbf{Q}_{\mathbf{c}_i}$ to a constant already yields strong performance in active mapping, since the visibility field effectively captures uncertainty in unseen regions, which is the dominant factor for active mapping. Incorporating predicted variances $\mathbf{Q}_{\mathbf{c}_i}$ yields only marginal performance gains.

Motivated by this observation, our base GAVIS model adopts a constant variance $\mathbf{Q}_{\mathbf{c}_i} = \sigma_c^2 \mathbf{I}$, which offers a favorable trade-off between performance and computational cost. In addition, existing uncertainty quantification methods for 3DGS, including FisherRF [32] and VIMC [54], could be used to predict $\mathbf{Q}_{\mathbf{c}_i}$, leading to two variants of the base GAVIS model: (1) FisherRF+GAVIS and (2) VIMC+GAVIS. Further details on how these methods are used to obtain $\mathbf{Q}_{\mathbf{c}_i}$ are provided in the Sec. 10, and quantitative experiments of the two variants are included in the

Sec. 12, demonstrating that GAVIS can serve as a post-hoc module that significantly improves the performance of existing methods.

Finally, we describe the algorithm for 3DGS uncertainty quantification using the visibility field at a synthesis view $\mathbf{p}$ in Alg. 6. This procedure is efficiently implemented using a modified 3DGS rasterizer, enabling real-time uncertainty quantification at around 200 FPS.

Algorithm 6 3DGS Uncertainty Quantification with Visibility Field

Input: $\mathbf{u}_g \in \mathbb{R}^2$: image-space projection of the Gaussian center to synthesis view $\mathbf{p}$; $\Sigma_g \in \mathbb{R}^{2 \times 2}$: image-space conic (precision) matrix; $z_g \in \mathbb{R}$: depth (for front-to-back sorting); $o_g \in [0, 1]$: base opacity; $v_g \in [0, 1]$: queried visibility field value for Gaussian g along direction $\mathbf{d}_p$; $\mathbf{Q}_{c_g}$: predicted color covariance for Gaussian g along direction $\mathbf{d}_p$; ϵ_T : transmittance cutoff (default 10^{-4}); $\beta, o_0, \mathbf{Q}_0$: hyperparameters for visibility compensation (default $\beta = 0.5, o_0 = 0.15, \mathbf{Q}_0 = \mathbf{I}_3$).

Output: Entropy map $\mathcal{H} \in \mathbb{R}^{H \times W}$ for synthesis view $\mathbf{p}$.

```

1: Initialize  $\mathbf{H}[\mathbf{x}] \leftarrow 0$  for all pixels  $\mathbf{x}$  in the image
2: for each pixel  $\mathbf{x}$  in the image do
3:    $T \leftarrow 1$  ▷ current transmittance
4:    $\bar{w}_0 \leftarrow 0$  ▷ accumulated weight of the invisible component
5:    $\mathcal{G} \leftarrow$  Gaussians overlapping  $\mathbf{x}$  (sorted by increasing  $z$ )
6:   for each  $g \in \mathcal{G}$  do
7:      $\alpha \leftarrow o_g \exp(-(\mathbf{u}_g - \mathbf{x})^\top \Sigma_g (\mathbf{u}_g - \mathbf{x}))$ 
8:      $\alpha^* \leftarrow (v_g + \beta(1 - v_g)) \alpha + o_0(1 - \beta)(1 - v_g)$ 
9:      $w^* \leftarrow \alpha^* T$ 
10:     $h \leftarrow -\log(w^* v_g) + \frac{1}{2} \log |\mathbf{Q}_{c_g}| + \frac{3}{2} \log(2\pi e)$ 
11:     $\mathbf{H}[\mathbf{x}] \leftarrow \mathbf{H}[\mathbf{x}] + w^* v_g h$  ▷ Eq. (30)
12:     $\bar{w}_0 \leftarrow \bar{w}_0 + w^*(1 - v_g)$ 
13:     $T \leftarrow T(1 - \alpha^*)$ 
14:    if  $T < \epsilon_T$  then break
15:    end if
16:  end for
17:   $h \leftarrow -\log \bar{w}_0 + \frac{1}{2} \log |\mathbf{Q}_0| + \frac{3}{2} \log(2\pi e)$ 
18:   $\mathbf{H}[\mathbf{x}] \leftarrow \mathbf{H}[\mathbf{x}] + \bar{w}_0 h$  ▷ Add the contribution of the invisible component
19: end for
20: Return  $\mathbf{H}$ 

```

9. Active Mapping Details

In this section, we provide additional details on our active mapping method, which follows a formulation similar to that in [104].

The goal of active mapping or next-best-view planning is

to select the camera pose $\mathbf{p}$ that maximizes the expected information gain (EIG) about the scene, which can be written as:

$$I(\mathbf{Z}; \boldsymbol{\theta} | \mathbf{p}) = \mathcal{H}[\mathbf{Z} | \mathbf{p}] - \mathcal{H}[\mathbf{Z}_p | \boldsymbol{\theta}, \mathbf{p}], \quad (31)$$

where $I(\cdot)$ denotes mutual information, $\mathcal{H}[\cdot]$ is entropy, $\mathbf{Z}$ is the random variable corresponding to the rendered observation, and $\boldsymbol{\theta}$ represents the radiance-field parameters.

We assume a Gaussian observation model, $p(\mathbf{Z} | \mathbf{p}, \boldsymbol{\theta}) = \mathcal{N}(f(\mathbf{p}; \boldsymbol{\theta}), \Sigma_z)$, where $f(\mathbf{p}; \boldsymbol{\theta})$ is the rendering function and Σ_z is the observation noise covariance. Assuming Σ_z is constant and isotropic, i.e., $\Sigma_z = \sigma_z^2 \mathbf{I}$ then the entropy term $\mathcal{H}[\mathbf{Z} | \boldsymbol{\theta}, \mathbf{p}]$ becomes constant and can be omitted during optimization. Hence, maximizing EIG reduces to maximizing the entropy:

$$\mathbf{p}^* = \arg \max_{\mathbf{p}} \mathcal{H}(\mathbf{Z}_p), \quad (32)$$

where we use $\mathcal{H}(\mathbf{Z}_p)$ as shorthand for $\mathcal{H}[\mathbf{Z}_p | \mathbf{p}]$, which could be obtained from the posterior of the rendered observation. This objective is commonly adopted in radiance-field-based active mapping [46, 54, 104, 107]. Note that uncertainties arising from robot localization or dynamics are not modeled here, as the primary focus of these methods is the uncertainty of the radiance field itself. Moreover, the action τ is restricted to the next camera pose, and the EIG of an entire trajectory is not considered, since efficiently computing the joint entropy over a sequence of observations remains an open challenge [40, 66]. A simplified workaround assumes independence across views [31, 91], though this approximation is not the focus of this work. Extending the GAVIS framework to support trajectory-level EIG is a promising direction for future research, but it lies beyond the scope of this paper.

We adopt the spatial correlation correction term from NVF [104] to recover the joint image-level entropy from per-pixel entropies:

$$\mathcal{H}(\mathbf{Z}_p) = \sum_{m,n} \left(\mathcal{H}(\mathbf{Z}_p^{mn}) - f_{\text{corr}}(\mathcal{H}(\mathbf{Z}_p^{mn}); d_p^{mn}) \right), \quad (33)$$

where $\mathbf{Z}_p^{mn}$ denotes the observed color associated with pixel (m, n) , d_p^{mn} is its expected depth, and $f_{\text{corr}}(\cdot)$ is a depth-dependent correction term that compensates for spatial correlations between neighboring pixels. Additional details can be found in [104].

Finally, we describe the active mapping algorithm with GAVIS in Alg. 7. The algorithm iteratively trains the 3DGS model, constructs the visibility field, and selects the next-best view based on the expected information gain computed from the visibility-aware 3DGS uncertainty quantification.

10. Applying GAVIS as a Post-hoc Module

In this section, we provide additional details on applying GAVIS as a post-hoc module to enhance existing uncer-

Algorithm 7 Active Mapping with GAVIS

Input: $\mathcal{P}$: initial poses; $\mathcal{Z}$: initial images; T : number of active mapping planning steps.

Output: $\mathcal{G}$: the trained 3DGS model after active mapping.

```

1: for  $i = 1$  to  $T$  do
2:    $\mathcal{G} \leftarrow \text{train3DGS}(\mathcal{P}, \mathcal{Z})$  ▷ train 3DGS
3:    $\{\gamma_{\ell m}^{(g)}\} \leftarrow \text{VFConstruction}(\mathcal{G}, \mathcal{P})$  ▷ Alg. 3
4:    $\tilde{\mathcal{G}} \leftarrow \text{VFDensityControl}(\mathcal{G}, \mathcal{P})$  ▷ Alg. 5
5:    $\mathcal{P}_c \leftarrow \text{samplePoses}(\tilde{\mathcal{G}})$  ▷ sample candidate poses
6:   for  $p$  in  $\mathcal{P}_c$  do
7:      $v_g \leftarrow \text{VFQuery}(\tilde{\mathcal{G}}, \{\gamma_{\ell m}^{(g)}\}, p)$  ▷ Alg. 4
8:      $\mathcal{Z}_p \leftarrow \text{UQVF}(\tilde{\mathcal{G}}, v_g, p)$  ▷ Alg. 6
9:   end for
10:   $p_i \leftarrow \arg \max_{p \in \mathcal{P}_c} \mathcal{H}(\mathcal{Z}_p)$ 
11:   $\mathcal{P} \leftarrow \{p_i\} \cup \mathcal{P}$ 
12:   $\mathcal{Z} \leftarrow \text{takeImageAt}(\{p_i\}) \cup \mathcal{Z}$  ▷ update training set
13: end for
14: return  $\mathcal{G}$ 

```

tainty quantification methods, including FisherRF [32] and VIMC [54]. The key idea is to use the uncertainty estimated by FisherRF or VIMC to refine the variance of emitted color $\mathbf{Q}_{c_i}$ in (3) along ray direction $\mathbf{d}$, whereas in base GAVIS this quantity is treated as a constant hyperparameter for all particles. We first describe how to integrate GAVIS with FisherRF and VIMC, and then discuss the relationship between GAVIS and these methods.

10.1. Integration with FisherRF

FisherRF [32] approximates the expected information gain of a candidate camera pose by maximizing

$$\arg \max_p \text{tr}(\mathbf{H}''[\mathcal{Z}_p | p, \theta^*] \mathbf{H}''[\theta^* | D^{\text{train}}]^{-1}), \quad (34)$$

where the observation $\mathcal{Z}_p$ is the rendered image for pose p , θ^* are the estimated 3DGS parameters, and the training set is $D^{\text{train}} := \{(p, \mathcal{Z}_p)\}_{p \in \mathcal{P}}$. The observed information is given by:

$$\mathbf{H}''[\theta^* | D^{\text{train}}] = \sum_{(p, \mathcal{Z}_p) \in D^{\text{train}}} \mathbf{H}''[\mathcal{Z}_p | p, \theta^*], \quad (35)$$

where $\mathbf{H}''[\mathcal{Z}_p | p, \theta^*]$ is the Hessian of the negative log-likelihood under the Gaussian observation model

$$-\log p(\mathcal{Z}_p | p, \theta) = \frac{1}{2\sigma_z^2} (\mathcal{Z}_p - f(p; \theta))^{\top} (\mathcal{Z}_p - f(p; \theta)) + C, \quad (36)$$

with rendering model $f(p; \theta)$ and observation noise variance σ_z^2 . FisherRF adopts the Laplace approximation [40], leading to:

$$\mathbf{H}''[\mathcal{Z} | p, \theta^*] \approx \frac{1}{\sigma_z^2} \text{diag}(\nabla_{\theta} f(p; \theta)^{\top} \nabla_{\theta} f(p; \theta)) + \lambda \mathbf{I} \quad (37)$$

Thus the posterior covariance of θ is given by:

$$\Sigma_{\theta} \approx \mathbf{H}''[\theta^* | D^{\text{train}}]^{-1} \quad (38)$$

Further details on these derivations can be found in [23, 32, 40].

We then propagate this parameter uncertainty to the emitted color of each particle. Let $\theta_{\ell, m, k}^{(i)}$ denote the color coefficient of the i -th particle in channel k , parameterized in the spherical harmonics basis. The emitted color in channel k along direction $\mathbf{d}$ is given by

$$c_{i, k}(\mathbf{d}) = \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} \theta_{\ell, m, k}^{(i)} Y_{\ell, m}(\mathbf{d}), \quad (39)$$

which is linear in the coefficients $\theta_{\ell, m, k}^{(i)}$. The variance of the emitted color in channel k is obtained by linear uncertainty propagation:

$$Q_{c_{i, k}}(\mathbf{d}) = \sum_{\ell=0}^L \sum_{m=-\ell}^{\ell} Y_{\ell, m}(\mathbf{d})^2 \Sigma_{\ell, m, k}^{(i)}, \quad (40)$$

where $\Sigma_{\ell, m, k}^{(i)}$ is the variance of the coefficient $\theta_{\ell, m, k}^{(i)}$ extracted from the diagonal of Σ_{θ} , as in FisherRF [32] all parameters are assumed independent and the covariance matrix Σ_{θ} is therefore diagonal. Since each color channel is also modeled independently, the covariance of the emitted color vector $\mathbf{Q}_{c_i}$ is diagonal with diagonal entries $Q_{c_{i, k}}$.

10.2. Integration with VIMC

VIMC [54] estimates the distribution of θ from a learned low-dimensional manifold of the radiance field parameters. We draw J parameter samples (with $J = 2$ by default in VIMC). For each sampled parameter θ_j , we compute the emitted color $c_{i, k, j}(\mathbf{d})$ following (39), and estimate the per-channel uncertainty by taking the sample variance across the J samples:

$$Q_{c_{i, k}}(\mathbf{d}) = \frac{1}{J-1} \sum_{j=1}^J (c_{i, k, j}(\mathbf{d}) - \bar{c}_{i, k}(\mathbf{d}))^2, \quad (41)$$

where $\bar{c}_{i, k}(\mathbf{d}) = \frac{1}{J} \sum_{j=1}^J c_{i, k, j}(\mathbf{d})$ is the sample mean. Under the assumption that the three color channels are independent, the resulting emitted-color covariance $\mathbf{Q}_{c_i}$ similarly reduces to a diagonal matrix whose diagonal entries are given by $Q_{c_{i, k}}(\mathbf{d})$.

10.3. Discussion

In principle, an accurate radiance-field uncertainty estimation method should assign high uncertainty to unseen regions, as these areas do not contribute to the training loss. However, existing approaches struggle to capture this behavior.

In FisherRF, the diagonal Laplace approximation considers only the diagonal entries of the Hessian. This ignores the strong correlations among parameters of the same Gaussian particle (e.g., RGB spherical-harmonics coefficients within a color channel), causing the approximation to underestimate uncertainty, especially for viewing directions far from the training views.

In VIMC, uncertainty is approximated from Monte Carlo samples drawn from a learned low-dimensional manifold. With only a small number of samples, variance estimates are noisy and unreliable, while increasing the sample count significantly raises computational cost, making it unsuitable for real-time applications.

Integrating GAVIS provides an efficient way to compensate for these limitations. Since GAVIS reliably assigns higher uncertainty to unseen regions, it improves the accuracy of uncertainty quantification and leads to better active-mapping performance (see Sec. 12).

11. Experimental Details

11.1. Environment Setup

We evaluate our method on four datasets: (1) the standard NeRF Synthetic dataset [57]; (2) a space dataset consisting of the Hubble Space Telescope (HST) [3] and the International Space Station (ISS) [4] for space robotics scenarios; (3) eight indoor environments from the Habitat-Matterport 3D (HM3D) [69] dataset

- 00208-SQqGpSHzfSr
- 00299-bdp1XNEdvmW
- 00321-JWWJBQWHv64
- 00323-yHLr6bvWsVm
- 00415-rBmEe6ab5VP
- 00441-4MRLulyET6a
- 00446-tL6i2PtktSh
- 00670-mDdyQ6azhVD

and (4) eight scenes from the Gibson [101] dataset (annawan, bremerton, creede, eagerville, eastville, helix, hometown, quantico). NeRF Synthetic and the space dataset are simulated in Blender, whereas HM3D and Gibson scenes are simulated in Habitat-Sim [73]. All images are rendered at a resolution of 512×512 pixels, with both horizontal and vertical fields of view set to 90° . For a fair comparison, all methods are initialized from the same set of closely spaced initial views that cover only a portion of the scene, mimicking the starting condition of a robotic active mapping process. The agent

begins with 3 initial views for NeRF Synthetic and the space dataset, and with 5 spherically sampled views from a fixed starting position for HM3D and Gibson. We use 10 active-mapping steps for NeRF Synthetic and the space dataset, 40 steps for Gibson, and 80 steps for HM3D, chosen as the minimal steps at which the strongest methods achieve acceptable reconstruction quality ($\text{PSNR} > 24$). Each scene is evaluated over three runs with different random seeds, and the mean and propagated standard error due to seed stochasticity [92] of all metrics are reported. All experiments for each scene and each method are performed on a single NVIDIA A40 GPU.

11.2. Active Mapping Setup

To ensure a fair comparison of uncertainty quantification methods for active mapping, all methods are evaluated using the same active mapping pipeline, differing only in the uncertainty quantification module, following the setup in NVF [104]. In particular, each method uses the same minimally biased candidate view sampler, which samples only collision-free candidate poses without additional heuristics. This design reflects the principle that an optimal active mapping policy should, in principle, be achievable by relying solely on accurate uncertainty quantification, without the need for heuristic guidance. For the NeRF Synthetic and space datasets, candidate poses are sampled from $\text{SE}(3)$ within the scene bounding box. For the Gibson and HM3D datasets, candidate poses are sampled from $\text{SE}(2)$ on a plane at positions known to be collision-free, with a fixed height of 1.5 m, mimicking the active mapping process of a mobile robot.

11.3. Training Details

For a fair comparison, we adopt the standard 3DGS training setup from [39] with minimal differences across all 3DGS-based methods. Since FisherRF and GAVIS are post-hoc methods, they share the exact same 3DGS training setup. VIMC uses the same setup as well, with the only difference being an additional loss term used to learn the low-dimensional manifold in the space of 3DGS parameters for uncertainty quantification. NVF is trained following the same configuration as in [104]. For the NeRF Synthetic and space datasets, 3DGS is trained for 3500 iterations at each planning step. For the Gibson and HM3D indoor datasets, 3DGS is trained for 7500 iterations at each planning step, and depth supervision from [55] is incorporated to accelerate convergence for all 3DGS-based methods, using simulator-rendered depth images. For a fair comparison, a similar depth loss [17] is also applied when training NVF. For shorter-horizon datasets (NeRF Synthetic, space, and Gibson), 3DGS is initialized using a subsampled sparse point cloud constructed from all previously observed depth images, to reduce training time. For the longer-horizon

HM3D dataset, 3DGS at each planning step is initialized by merging the previous 3DGS model with a subsampled sparse point cloud derived from the most recent depth image.

11.4. Metrics Details

Visual Coverage. For visual coverage evaluation, we adopt the ground-truth visibility (VIS) metric from [104]. This metric assigns each face in the ground-truth mesh a binary visibility value, all faces are initialized with value 0, and then the value of a face is set to 1 (visible) if it is observed by the training views without occlusion. The Vis score is computed as the ratio of the total area of visible faces to the total mesh surface area.

We also visualize the GT visibility map (used in Figs. 6, 9, and 8) by rendering the visibility-annotated mesh from each candidate query pose. Similar to the uncertainty map, darker regions indicate visible areas (value 1), while brighter regions indicate invisible areas (value 0). Ideally, an uncertainty quantification method tailored for active mapping should produce uncertainty maps that correlate with the GT visibility map, assigning high uncertainty (bright) to invisible regions and low uncertainty (dark) to visible regions.

We note that ground-truth (GT) visibility is inherently an isotropic metric that does not account for the view-dependent nature of visibility. In particular, observing a face from a single direction, regardless of viewing distance, may already be considered fully visible. As a result, such metrics do not favor uncertainty quantification methods that encourage revisiting regions from diverse viewpoints. Designing evaluation metrics for visual coverage that account for viewpoint diversity and better align with the active mapping’s objective of acquiring observations from diverse directions remains an open problem.

Mesh Metrics. To assess reconstruction completeness, we also include the completion ratio defined in [87]. Ground-truth 3D points are sampled from the original scene meshes. Predicted points are extracted from a trained 3DGS model using the approach described in [2]. For consistency with prior work [104], we use a completion ratio threshold of 0.01 for the NeRF Synthetic and space datasets. For HM3D and Gibson environments, we adopt a threshold of 0.05.

We additionally report mesh-based metrics, including completion (Comp) and accuracy (Acc). However, we adopt completion ratio (CR) as the primary mesh metric, as it is widely used, easier to interpret, and more robust to outliers compared to Comp [33, 50]. Moreover, Acc is not well aligned with the objective of active mapping under 3DGS, despite its use in NeRF-based active mapping [104]. This is because 3DGS cannot distinguish empty space from under-explored regions, causing reconstructed point clouds to reflect only explored areas. As a result, Acc is biased toward

policies that oversample a limited region while penalizing broader exploration in 3DGS active mapping. Empirically, Acc shows minimal improvement with additional views and exhibits no statistically significant differences across methods; accordingly, it is rarely used in 3DGS-based active mapping [32, 54]. As shown in Tab. 6, Acc varies little across methods and shows weak correlation with other metrics.

Uncertainty Quantification. To quantitatively evaluate the quality of uncertainty quantification, we adopt the Area Under the Sparsification Error curve (AUSE) metric [29, 67], which measures the correlation between predicted uncertainty and actual error. However, the depth-based AUSE metric (AUSE-D) used in [23, 32, 54], which evaluates the correlation between predicted uncertainty and depth error computed from the predicted depth and the ground-truth mesh, can be misaligned with the objective of active mapping. In particular, radiance fields can produce accurate depth predictions in unobserved regions via plausible imagination. As a result, generic uncertainty quantification methods, which are not designed for active mapping and tend to assign low uncertainty to these regions, can still achieve strong AUSE-D scores. In contrast, effective active mapping requires high uncertainty in unobserved regions to encourage exploration. Therefore, a strong AUSE-D score does not necessarily indicate good active mapping performance.

We note that alternative metrics for evaluating uncertainty quantification exist, including negative log-likelihood (NLL) and expected calibration error (ECE) [25]. However, these metrics are less commonly used in the context of radiance fields, since they emphasize the quality of a fully specified and calibrated predictive distribution, which is not necessary for downstream tasks such as active mapping and artifact removal [23]. Moreover, it is unclear how to fairly compare these metrics across different methods, for example, GMM-based approaches (GAVIS&NVF), methods assuming a Gaussian distribution (VIMC), and methods that do not explicitly predict a distribution (FisherRF). Furthermore, similar to AUSE-D, these metrics can also be misaligned with the objective of active mapping.

Therefore, designing uncertainty quantification metrics tailored to active mapping remains an open problem. Toward this direction, we introduce a ground-truth visibility-based variant, denoted as AUSE-V, which measures the correlation between the predicted uncertainty map and the ground-truth visibility map. The ground-truth visibility map is obtained by assigning binary face-level values on the ground-truth mesh and rendering them using a rasterizer. AUSE-V better captures the requirement of active mapping for high uncertainty in unobserved regions, and thus exhibits stronger correlation with both active mapping performance and qualitative behavior. Further results are provided

in Sec. 12.1.

However, as discussed above, GT visibility is inherently isotropic and does not account for the view-dependent nature of visibility, so it may not fully capture the benefits of anisotropic visibility modeling. Moreover, a reliable uncertainty quantification method for active mapping should assign high uncertainty to regions with low GT visibility, but need not always assign low uncertainty to regions with high GT visibility. High uncertainty in observed regions may arise when the test view direction deviates substantially from the training views direction or when model capacity is limited. As a result, predictions can still be inaccurate even in observed regions. Therefore, although AUSE-V is more informative than AUSE-D, it may not fully capture uncertainty quality for active mapping.

12. Additional Results

12.1. Additional Quantitative Results

We present the complete active mapping results for all methods, including the baselines FisherRF, VIMC, and NVF, as well as our proposed methods GAVIS, FisherRF+GAVIS, and VIMC+GAVIS, with all evaluation metrics reported for each dataset in Tab. 6, with additional Comp and Acc metrics included.

As shown in Tab. 6, our method significantly outperforms all 3DGS baselines in all metrics excluding Acc. Moreover, integrating GAVIS as a post-hoc module consistently yields substantial improvements over each corresponding baseline across all datasets and all metrics excluding Acc. Our method also surpasses NVF in active mapping performance while requiring substantially less computation time. GAVIS outperforms NVF on all image-based metrics (PSNR, SSIM, and LPIPS) across all datasets. On the primary mesh metric (CR) and visibility metric (VIS), GAVIS outperforms NVF on all datasets except Gibson. Since these metrics are inherently isotropic and direction agnostic, e.g., observing a region from a single direction may already provide high GT mesh visibility and low point-cloud error (especially when depth loss is used), the performance gains are less pronounced on these metrics.

Among the variants of our method (GAVIS, FisherRF+GAVIS, and VIMC+GAVIS), FisherRF+GAVIS overall outperforms the base GAVIS model on most datasets except the space dataset. This suggests that more accurately modeling Q_{c_i} can further improve performance. The advantage is less pronounced in the space dataset, since color variance tends to be more monotonic across these scenes. VIMC+GAVIS outperforms the base GAVIS model only on the NeRF Synthetic dataset and not on the others. This is because VIMC relies on sampling-based uncertainty estimation, and with only 2 samples (the default setting in [54]), the sampling noise tends to overshadow the benefits intro-

duced by visibility modeling.

Although performance gains can be achieved by more accurately modeling Q_{c_i} , using a constant covariance already yields superior performance, as visibility is the dominant factor in active mapping, which is consistent with the findings in [104]. Since obtaining more accurate Q_{c_i} with FisherRF+GAVIS or VIMC+GAVIS requires additional processing with substantially higher computational cost, the base GAVIS model provides a more favorable trade-off between performance and efficiency in practice.

We include additional uncertainty quantification results using AUSE-D and AUSE-V in Tab. 5, covering all baseline methods and ablations. GAVIS achieves the best performance in AUSE-V among all baselines and ablations that disable individual components, which is consistent with its superior active mapping performance and qualitative results. GAVIS also attains the best AUSE-D score among all baseline methods. However, it does not outperform the isotropic ablation that disables anisotropic visibility by setting $\nu(d; d_p) = 1$. This discrepancy arises from the misalignment between AUSE-D and the objective of active mapping, as discussed in Sec. 11. A simple example can be constructed from a scenario similar to Fig. 7, where the 3DGS model produces plausible imaginations behind a wall, leading to relatively accurate depth predictions in those regions and consequently improving the overall AUSE-D score, even though the model incorrectly assigns low uncertainty to unobserved areas without considering anisotropic visibility.

	GAVIS	NVF	FisherRF	VIMC	Iso.	w/o DC	Iso. w/o DC
AUSE-D ↓	0.224	0.381	0.463	0.504	0.205	0.515	0.413
AUSE-V ↓	0.176	0.231	0.496	0.447	0.292	0.480	0.514

Table 5. Additional uncertainty quantification results

We conduct additional ablation studies on the SH degree L , anisotropy parameter κ , and the density of virtual particles ρ in visibility field density control, compared against the default GAVIS settings ($L = 2$, $\kappa = 1$, $\rho = 100$). The results in Tab. 7, averaged over all datasets, indicate that GAVIS is relatively robust to these hyperparameters. Specifically, for κ , the default $\kappa = 1$ performs best overall. Smaller values modestly reduce rendering quality, while overly large values (e.g., $\kappa = 10$, a sharper $\nu(d; d_p)$) degrade visual coverage. For SH degree L , $L = 3$ yields a slight improvement over $L = 2$, but gains plateau at higher degrees; we retain $L = 2$ as the default for efficiency. For virtual particle density ρ , higher values consistently improve performance with diminishing returns, and $\rho = 100$ provides a practical balance between performance and computational cost.

12.2. Additional Qualitative Results

In this subsection, we provide additional comprehensive qualitative results across a wide variety of scenes, demonstrating the generality of GAVIS compared to baseline methods for uncertainty quantification and active mapping.

Fig. 8–9 visualize uncertainty quantification results. For each scene, all methods are trained on the same set of views, which only partially observe the underlying scene. To enable direct evaluation of uncertainty quantification, we generate ground-truth visibility (GT VIS) maps by marking mesh faces that are invisible from the training views with a bright color (and visible faces with a dark color), and rendering this visibility-annotated mesh from the same queried viewpoints. An accurate uncertainty quantification algorithm should assign high uncertainty to invisible regions, consistent with the GT visibility.

As shown in Figs. 8–9, GAVIS provides uncertainty estimation that aligns most closely with the GT VIS maps, thanks to its analytical anisotropic visibility modeling. In contrast, NVF relies on a black-box neural network approximation of isotropic visibility, leading to a noisy and less accurate uncertainty map. All visibility-aware methods significantly outperform baselines that ignore visibility, which fail to assign high uncertainty to invisible regions. Qualitative results for FisherRF+GAVIS and VIMC+GAVIS are also presented, demonstrating significant improvements over the original method by reliably assigning high uncertainty to invisible regions.

Finally, additional active-mapping reconstruction results for all Gibson and HM3D scenes are shown in Figs. 10 and 11, demonstrating that our method leverages accurate uncertainty quantification to guide more effective exploration, resulting in broader scene coverage and higher-quality reconstructions.

Dataset	Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Acc $\downarrow$	Comp $\downarrow$	CR $\uparrow$	VIS $\uparrow$
NeRF Synth.	FisherRF	22.34 $\pm$ 0.31	0.870 $\pm$ 0.004	0.119 $\pm$ 0.003	0.014 $\pm$ 2e-4	0.014 $\pm$ 0.001	0.626 $\pm$ 0.011	0.376 $\pm$ 0.010
	VIMC	23.14 $\pm$ 0.25	0.880 $\pm$ 0.003	0.107 $\pm$ 0.003	0.013 $\pm$ 4e-4	0.015 $\pm$ 0.001	0.651 $\pm$ 0.010	0.407 $\pm$ 0.011
	NVF	22.59 $\pm$ 0.26	0.859 $\pm$ 0.005	0.147 $\pm$ 0.004	0.025 $\pm$ 0.002	0.020 $\pm$ 0.001	0.549 $\pm$ 0.014	0.431 $\pm$ 0.005
	GAVIS (ours)	24.26 $\pm$ 0.25	<u>0.894 $\pm$ 0.002</u>	<u>0.097 $\pm$ 0.002</u>	<u>0.014 $\pm$ 2e-4</u>	0.011 $\pm$ 0.001	<u>0.711 $\pm$ 0.009</u>	<u>0.437 $\pm$ 0.006</u>
	FisherRF+GAVIS (ours)	24.55 $\pm$ 0.17	0.898 $\pm$ 0.002	0.092 $\pm$ 0.002	<u>0.013 $\pm$ 3e-4</u>	<u>0.012 $\pm$ 4e-4</u>	<u>0.703 $\pm$ 0.008</u>	0.440 $\pm$ 0.005
	VIMC+GAVIS (ours)	<u>24.30 $\pm$ 0.18</u>	<u>0.894 $\pm$ 0.002</u>	<u>0.096 $\pm$ 0.001</u>	0.014 $\pm$ 2e-4	<u>0.012 $\pm$ 3e-4</u>	0.713 $\pm$ 0.006	<u>0.433 $\pm$ 0.005</u>
Space	FisherRF	24.17 $\pm$ 0.08	0.834 $\pm$ 0.004	0.158 $\pm$ 0.002	0.015 $\pm$ 3e-4	0.020 $\pm$ 0.001	0.547 $\pm$ 0.017	0.474 $\pm$ 0.016
	VIMC	24.56 $\pm$ 0.55	0.841 $\pm$ 0.004	0.150 $\pm$ 0.002	<u>0.015 $\pm$ 4e-4</u>	0.017 $\pm$ 0.003	0.612 $\pm$ 0.008	0.510 $\pm$ 0.009
	NVF	23.76 $\pm$ 0.46	0.796 $\pm$ 0.012	0.202 $\pm$ 0.010	0.037 $\pm$ 0.003	0.015 $\pm$ 0.002	0.499 $\pm$ 0.019	0.564 $\pm$ 0.017
	GAVIS (ours)	26.14 $\pm$ 0.10	0.857 $\pm$ 0.003	<u>0.140 $\pm$ 0.002</u>	<u>0.014 $\pm$ 4e-4</u>	0.013 $\pm$ 0.002	<u>0.630 $\pm$ 0.019</u>	<u>0.582 $\pm$ 0.017</u>
	FisherRF+GAVIS (ours)	<u>25.44 $\pm$ 0.31</u>	<u>0.857 $\pm$ 0.003</u>	0.139 $\pm$ 0.002	0.013 $\pm$ 0.001	<u>0.014 $\pm$ 0.001</u>	<u>0.631 $\pm$ 0.014</u>	<u>0.588 $\pm$ 0.023</u>
	VIMC+GAVIS (ours)	<u>25.47 $\pm$ 0.34</u>	<u>0.857 $\pm$ 0.002</u>	<u>0.141 $\pm$ 0.001</u>	0.016 $\pm$ 0.001	<u>0.013 $\pm$ 2e-4</u>	0.640 $\pm$ 0.008	0.597 $\pm$ 0.009
Gibson	FisherRF	18.11 $\pm$ 0.50	0.720 $\pm$ 0.007	0.419 $\pm$ 0.006	<u>0.038 $\pm$ 0.001</u>	0.891 $\pm$ 0.139	0.431 $\pm$ 0.031	0.469 $\pm$ 0.035
	VIMC	15.70 $\pm$ 0.20	0.668 $\pm$ 0.003	0.465 $\pm$ 0.003	<u>0.046 $\pm$ 0.002</u>	1.006 $\pm$ 0.067	0.337 $\pm$ 0.013	0.366 $\pm$ 0.014
	NVF	23.29 $\pm$ 0.12	0.798 $\pm$ 0.001	0.402 $\pm$ 0.002	0.029 $\pm$ 3e-4	0.033 $\pm$ 0.001	0.880 $\pm$ 0.002	0.915 $\pm$ 0.002
	GAVIS (ours)	<u>24.42 $\pm$ 0.14</u>	<u>0.812 $\pm$ 0.003</u>	<u>0.323 $\pm$ 0.004</u>	<u>0.040 $\pm$ 0.001</u>	<u>0.044 $\pm$ 0.002</u>	<u>0.831 $\pm$ 0.005</u>	<u>0.890 $\pm$ 0.003</u>
	FisherRF+GAVIS (ours)	24.58 $\pm$ 0.12	0.815 $\pm$ 0.002	0.319 $\pm$ 0.003	0.041 $\pm$ 0.001	<u>0.044 $\pm$ 0.002</u>	<u>0.836 $\pm$ 0.006</u>	<u>0.890 $\pm$ 0.007</u>
	VIMC+GAVIS (ours)	<u>24.04 $\pm$ 0.19</u>	<u>0.807 $\pm$ 0.002</u>	<u>0.324 $\pm$ 0.003</u>	0.041 $\pm$ 3e-4	0.075 $\pm$ 0.006	0.772 $\pm$ 0.009	0.834 $\pm$ 0.010
HM3D	FisherRF	18.32 $\pm$ 0.44	0.693 $\pm$ 0.009	0.446 $\pm$ 0.009	<u>0.040 $\pm$ 0.001</u>	0.503 $\pm$ 0.097	0.447 $\pm$ 0.025	0.558 $\pm$ 0.030
	VIMC	17.15 $\pm$ 0.26	0.645 $\pm$ 0.006	0.477 $\pm$ 0.006	<u>0.052 $\pm$ 0.001</u>	0.178 $\pm$ 0.038	0.476 $\pm$ 0.017	0.618 $\pm$ 0.020
	NVF	22.69 $\pm$ 0.14	0.760 $\pm$ 0.002	0.434 $\pm$ 0.002	0.036 $\pm$ 0.001	0.040 $\pm$ 0.001	<u>0.819 $\pm$ 0.003</u>	<u>0.873 $\pm$ 0.001</u>
	GAVIS (ours)	<u>23.97 $\pm$ 0.07</u>	<u>0.791 $\pm$ 0.001</u>	<u>0.338 $\pm$ 0.002</u>	0.043 $\pm$ 4e-4	<u>0.040 $\pm$ 3e-4</u>	<u>0.820 $\pm$ 0.003</u>	<u>0.876 $\pm$ 0.001</u>
	FisherRF+GAVIS (ours)	24.23 $\pm$ 0.07	0.797 $\pm$ 0.001	0.330 $\pm$ 0.002	<u>0.043 $\pm$ 3e-4</u>	<u>0.040 $\pm$ 3e-4</u>	0.821 $\pm$ 0.002	0.877 $\pm$ 0.001
	VIMC+GAVIS (ours)	<u>23.04 $\pm$ 0.13</u>	<u>0.774 $\pm$ 0.002</u>	<u>0.347 $\pm$ 0.002</u>	0.043 $\pm$ 5e-4	0.057 $\pm$ 0.002	0.750 $\pm$ 0.008	0.826 $\pm$ 0.005

Table 6. **Quantitative results.** Active mapping performance across all datasets and methods. Best results are shown in **bold**; second and third best are underlined.

Setting	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	Acc $\downarrow$	Comp $\downarrow$	CR $\uparrow$	VIS $\uparrow$
GAVIS (default)	24.70 $\pm$ 0.08	0.839 $\pm$ 0.001	0.224 $\pm$ 0.001	0.028 $\pm$ 2e-4	0.027 $\pm$ 0.001	0.748 $\pm$ 0.006	0.697 $\pm$ 0.005
$\kappa = 0$	24.22 $\pm$ 0.14	0.832 $\pm$ 0.001	0.223 $\pm$ 0.001	0.027 $\pm$ 3e-4	0.027 $\pm$ 0.001	0.744 $\pm$ 0.005	0.666 $\pm$ 0.005
$\kappa = 0.1$	24.26 $\pm$ 0.18	0.834 $\pm$ 0.001	0.222 $\pm$ 0.001	0.027 $\pm$ 4e-4	0.026 $\pm$ 0.001	0.747 $\pm$ 0.005	0.671 $\pm$ 0.005
$\kappa = 0.3$	24.39 $\pm$ 0.13	0.836 $\pm$ 0.002	0.222 $\pm$ 0.001	0.028 $\pm$ 3e-4	0.026 $\pm$ 4e-4	0.753 $\pm$ 0.006	0.686 $\pm$ 0.003
$\kappa = 3$	24.71 $\pm$ 0.11	0.843 $\pm$ 0.001	0.216 $\pm$ 0.001	0.026 $\pm$ 2e-4	0.043 $\pm$ 0.004	0.725 $\pm$ 0.007	0.676 $\pm$ 0.006
$\kappa = 10$	24.04 $\pm$ 0.16	0.835 $\pm$ 0.002	0.223 $\pm$ 0.002	0.024 $\pm$ 4e-4	0.112 $\pm$ 0.017	0.662 $\pm$ 0.009	0.610 $\pm$ 0.013
$L = 1$	24.61 $\pm$ 0.10	0.837 $\pm$ 0.002	0.224 $\pm$ 0.002	0.027 $\pm$ 3e-4	0.028 $\pm$ 0.001	0.749 $\pm$ 0.005	0.694 $\pm$ 0.004
$L = 3$	24.90 $\pm$ 0.09	0.844 $\pm$ 0.002	0.217 $\pm$ 0.001	0.027 $\pm$ 2e-4	0.027 $\pm$ 0.001	0.761 $\pm$ 0.003	0.702 $\pm$ 0.006
$L = 4$	24.73 $\pm$ 0.10	0.841 $\pm$ 0.001	0.220 $\pm$ 0.001	0.027 $\pm$ 2e-4	0.027 $\pm$ 0.001	0.758 $\pm$ 0.004	0.700 $\pm$ 0.006
$L = 5$	24.74 $\pm$ 0.15	0.840 $\pm$ 0.003	0.220 $\pm$ 0.002	0.028 $\pm$ 3e-4	0.027 $\pm$ 0.001	0.755 $\pm$ 0.008	0.697 $\pm$ 0.008
$\rho = 10$	24.53 $\pm$ 0.08	0.837 $\pm$ 0.001	0.225 $\pm$ 0.001	0.028 $\pm$ 3e-4	0.033 $\pm$ 0.002	0.738 $\pm$ 0.005	0.685 $\pm$ 0.006
$\rho = 25$	24.54 $\pm$ 0.12	0.836 $\pm$ 0.001	0.224 $\pm$ 0.001	0.028 $\pm$ 2e-4	0.031 $\pm$ 0.001	0.734 $\pm$ 0.003	0.683 $\pm$ 0.003
$\rho = 50$	24.54 $\pm$ 0.11	0.839 $\pm$ 0.001	0.222 $\pm$ 0.001	0.028 $\pm$ 3e-4	0.029 $\pm$ 0.001	0.750 $\pm$ 0.003	0.693 $\pm$ 0.003
$\rho = 200$	24.99 $\pm$ 0.08	0.844 $\pm$ 0.001	0.216 $\pm$ 0.001	0.027 $\pm$ 2e-4	0.026 $\pm$ 2e-4	0.765 $\pm$ 0.003	0.703 $\pm$ 0.004
$\rho = 400$	24.95 $\pm$ 0.09	0.844 $\pm$ 0.001	0.215 $\pm$ 0.001	0.027 $\pm$ 2e-4	0.025 $\pm$ 3e-4	0.766 $\pm$ 0.005	0.701 $\pm$ 0.004
$\rho = 1000$	25.07 $\pm$ 0.10	0.847 $\pm$ 0.001	0.212 $\pm$ 0.001	0.026 $\pm$ 2e-4	0.030 $\pm$ 0.004	0.763 $\pm$ 0.004	0.700 $\pm$ 0.005

Table 7. **Additional ablation studies.** Active mapping performance of GAVIS with different hyperparameter settings. GAVIS uses $\kappa = 1$, $L = 2$ and $\rho = 100$ by default.




Figure 8. **Qualitative results for uncertainty quantification** in indoor scenes (Gibson and HM3D). From left to right: (1) Top-down view of the scene with training views (blue frustums) that cover only part of the room, and the queried view for uncertainty evaluation (green frustum); (2) Ground-truth RGB image rendered from the scene mesh; (3) Synthesized RGB image from a 3DGS model trained on the partial-view dataset; Uncertainty maps produced by (4) FisherRF, (5) VIMC, (6) NVF, (7) GAVIS, (8) FisherRF+GAVIS, (9) VIMC+GAVIS; (10) Ground-truth visibility map used as reference, where bright regions indicate invisible areas and dark regions indicate visible areas. An accurate uncertainty quantification method should closely align with the GT visibility by reliably assigning high uncertainty to invisible regions. All methods are trained using the same set of training views.




Figure 9. **Qualitative results for uncertainty quantification** in HST (from space dataset) and Lego (from NeRF Synthetic dataset) scenes. From left to right: (1) Isometric view of the object with training views (blue frustums) that cover only part of the object geometry, and the queried view for uncertainty evaluation (green frustum); (2) Ground-truth RGB image rendered from the scene mesh; (3) Synthesized RGB image from a 3DGS model trained on the partial-view dataset; Uncertainty maps produced by (4) FisherRF, (5) VIMC, (6) NVF, (7) GAVIS, (8) FisherRF+GAVIS, (9) VIMC+GAVIS; (10) Ground-truth visibility map used as reference, where bright regions indicate invisible areas and dark regions indicate visible areas. An accurate uncertainty quantification method should closely align with the GT visibility by reliably assigning high uncertainty to invisible regions. All methods are trained using the same set of training views.




Figure 10. **Qualitative results for active mapping** in all HM3D scenes From left to right: (1) Ground-truth top-down view; Reconstruction results from (2) FisherRF, (3) VIMC, (4) NVF, (5) GAVIS, (6) FisherRF+GAVIS, (7) VIMC+GAVIS. Planned camera poses are shown as green frustums.



Figure 11. **Qualitative results for active mapping** in all Gibson scenes From left to right: (1) Ground-truth top-down view; Reconstruction results from (2) FisherRF, (3) VIMC, (4) NVF, (5) GAVIS, (6) FisherRF+GAVIS, (7) VIMC+GAVIS. Planned camera poses are shown as green frustums.