

FisherRF: Active View Selection and Mapping with Radiance Fields using Fisher Information

Wen Jiang¹, Boshu Lei¹, and Kostas Daniilidis^{1,2}

¹University of Pennsylvania

²Archimedes, Athena RC

Abstract. This study addresses the challenging problem of active view selection and uncertainty quantification within the domain of Radiance Fields. Neural Radiance Fields (NeRF) have greatly advanced image rendering and reconstruction, but the cost of acquiring images poses the need to select the most informative viewpoints efficiently. Existing approaches depend on modifying the model architecture or hypothetical perturbation field to indirectly approximate the model uncertainty. However, selecting views from indirect approximation does not guarantee optimal information gain for the model. By leveraging Fisher Information, we directly quantify observed information on the parameters of Radiance Fields and select candidate views by maximizing the Expected Information Gain (EIG). Our method achieves state-of-the-art results on multiple tasks, including view selection, active mapping, and uncertainty quantification, demonstrating its potential to advance the field of Radiance Fields.

1 Introduction

Neural Radiance Fields brought back image rendering and reconstruction from multiple views to the center of attention in the field of computer vision. Novel volumetric representations of radiance fields and differentiable volumetric rendering enabled unprecedented advances in image-based rendering of complex scenes both in terms of perceptual quality and speed.

Recently, 3D Gaussian Splatting [25] has demonstrated distinct advantages in real-time rendering and explicit point-based parameterizations without neural representations. However, to achieve satisfactory rendering quality, numerous viewpoints are required to train a radiance field, especially in large-scale scenarios. It is crucial to establish a criterion for selecting information-maximizing views before obtaining image observations at those locations. Quantifying the observed information of a Radiance Field model is challenging, given that Radiance Field models are typically regression-based and scene-specific. The challenge intensifies when we aim to leverage quantified observed information for active view selection and mapping, especially when the selection candidates are only $SE(3)$ camera poses for acquiring new observations, a.k.a capturing new images.

Previous approaches select viewpoints by quantifying the uncertainty in Radiance Fields. They can be broadly categorized into two groups: variational

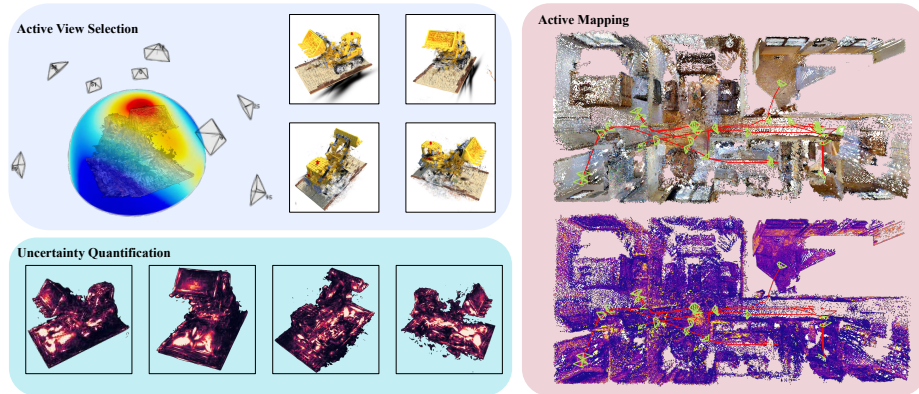


Fig. 1: A brief overview of our method Given a Radiance Field that was trained with a limited number of views, our method could find the next best view that could maximize information gain by computing the Fisher Information of the radiance field. We illustrate the Information Gain as a heat map on the viewing sphere and show four of the candidate views. Our model can quantify pixel-wise uncertainty, visualized at the bottom left, by examining the Fisher Information on the related parameters of the ray. Our algorithm can be also adapted into an active mapping system, as showcased at the right, that could actively explore and reconstruct the environment.

white-box models and black-box models. White-box models integrate conventional NeRF architectures with Bayesian models, such as reparameterization [36, 43, 48] and Normalizing Flows [55]. Black-box methods, on the other hand, do not modify the existing model architecture but seek to quantify predictive uncertainty by examining the distribution of predicted outcomes [31, 43, 67, 71]. White-box models depend on specific model architectures and are often characterized by slower training times due to the challenges associated with probabilistic learning. Conversely, existing black-box models either focus solely on studying the uncertainty on NeRF-style query points through network prediction [43] or hypothetical perturbation field [16], or rely on Monte Carlo sampling techniques [67] to quantify the uncertainty at the candidate views.

In this study, our primary objective is to quantify the observed information of a Radiance Field model and utilize it to select the optimal view with the highest information gain for downstream tasks, Fig. 1. To achieve this, we propose using Fisher Information, which represents the expectation of observation information. This quantity is directly linked to the second-order derivatives or Hessian matrix of the loss function involved in optimizing Radiance Field models. Importantly, the Hessian of the objective function in volumetric rendering is independent of ground truth data or the actual image measurement. This property allows us to compute the information gain between the training dataset and the candidate viewpoint pool using only the camera parameters of the candidate views. This capability facilitates an efficient next-best-view selection algorithm. In addition to the formulation of Information Gain, the Hessian of the loss function has

an intuitive interpretation: the perturbation at flat minima of the loss function. From an intuitive standpoint, one can perceive the Fisher Information matrix as a metric of the curvature of the log-likelihood function at specific parameter instantiations denoted as \mathbf{w}^* . Lower Fisher Information suggests that the log-likelihood function exhibits a flatter profile around \mathbf{w}^* , implying that the loss is less prone to changes when \mathbf{w}^* is perturbed. The flat minimum interpretation has attracted substantial attention and research within various domains of machine learning [19,20,28,34,58]. Moreover, as we have estimated the Fisher information for the model’s parameters, which correspond to specific 3D locations in recent Radiance Field models such as 3D Gaussian Splatting [25] and Plenoxels [52], we can derive pixel-wise uncertainty in the model’s predictions by examining the Fisher Information on parameters that contribute to the prediction for each pixel.

We implement the computation of Fisher Information on top of two types of Radiance Field models: point-based 3D Gaussian Splatting [25] and Plenoxels [52]. In 3D Gaussian Splatting, we compute the Hessians of the parameters on each 3D Gaussian with respect to the negative log-likelihood. As Fisher Information is additive, we further extend our view selection algorithm into batch selection and path selection, which are closer to real-world applications. To make the computation of information gain tractable, we use an approximation of the decrease in entropy by its upper bound, which is the trace of the product of the Hessian of the candidate view times the Hessian with respect to all previous views. These matrices are highly sparse due to the disentangling of scene parameters with respect to disparate rays. The sparsity allows us a computation of the matrix trace above that is as cost-effective as back-propagation, enabling us to evaluate views at 70 fps when we use 3D Gaussian Splatting. We carried out extensive evaluations in different benchmarks, including active view selection, active mapping, and uncertainty quantification. The quantitative and qualitative results unequivocally demonstrate that our approach surpasses previous methods and heuristic baseline by a significant margin. To summarize, our main contributions are as follows:

- A novel formulation to quantify the observed information and Information Gain for Radiance Fields.
- An effective and efficient view selection method exploiting the sparse structure of the scene rendering problem.
- An efficient pixel-wise uncertainty quantification and visualization method for 3D Gaussian Splatting.
- Extensive comparative study showing that our method outperforms existing active approaches on view selection, active mapping, and uncertainty quantification.

2 Related Works

Active Learning and Radiance Fields are prosperous research fields with various research directions. In this section, we limit our literature review to works that

directly relate to our tasks. We refer the readers to literature reviews if they are interested in Radiance Fields [10, 13] or Deep Active Learning [50, 72]. Besides, Fisher Information has been extensively studied in deep active learning [1, 2, 27, 29]. Notably, Kirsch *et al.* [28] unified existing works in active learning for deep learning problems from the perspective of Fisher Information, which shares many insights with our method.

Uncertainty Quantifications for Radiance Fields In the Neural Radiance Fields, Lee *et al.* [31], Yan *et al.* [66] and Zhan *et al.* [71] attempted to quantify the uncertainty in a scene by the distribution of densities on a casted ray. Shen *et al.* [55, 56] designed Bayesian models by re-parametrizing the NeRF model. However, they only tackled the predictive uncertainty of the model that did not relate to the observed information of the parameters. Sunderhauf *et al.* [61] propose an additional uncertainty measure in uncharted regions, determined by the ray termination probability on the learned geometry. Regarding concurrent work, Goli *et al.* [16] introduces a hypothetical perturbation field and interpolates uncertainty over NeRF-style query points for uncertainty quantification and NeRF clean-up task. The perturbation field with respect to query points is also not compatible with recent radiance field models like 3D Gaussian Splatting [25], which do not take query points as inputs. Besides, their algorithm is not designed for an active learning problem since selecting candidates with the largest per-pixel uncertainty may not be optimal, as the information gain for the model parameters is not concerned. In contrast, we focus on active learning problems. We derive our theory from the expected information gain for the model and compute the Hessian for the model parameters, allowing us to apply our algorithm to recent advancements in radiance fields such as 3D Gaussian splatting. Our paper focuses on active view selection and mapping based on our novel expected information gain objective. Moreover, their implementation relies on multiple function calls to PyTorch backward engine whereas our efficient CUDA implementation for the diagonal Hessian computation consumes much less time ($11.3 \text{ ms} \pm 33.3 \text{ } \mu\text{s}$ v.s. $1.1 \text{ s} \pm 2.5 \text{ ms}$). To the best of our knowledge, we are the first method that quantifies the uncertainty directly on model parameters for Radiance Fields, thanks to the powerful framework of 3D Gaussian and our efficient formulation.

Active View Selection and Mapping Although the next best view selection problem was extensively studied before radiance fields gained popularity, there has not been much literature on active “training” view selection for volumetric rendering. Jin *et al.* [22] trains a neural network to predict per-pixel uncertainty for view selection. However, their method requires an image capture as the input of the network, which is closer to dataset subsampling rather than active learning. ActiveNeRF [43] studied the next best view selection by directly adding variances as the output of the vanilla NeRF, and it is the closest approach to the view selection problem we are working on. In a larger scope, Active Perception [3, 4] has been studied along with the advancement of robotics and computer vision. Traditionally, frontier-based approach [64] and Rapidly exploring

Random Tree (RRT) [30] are the most representative and have been extended and enhanced for different data representations [12, 23, 57, 74] and complex conditions [44, 68]. For the larger scope of Active Simultaneous Localization and Mapping (SLAM), please refer to literature reviews [33, 45]. In recent years, Guedon *et al.* [17, 18] attempted to improve the coverage for mapping systems using point cloud representation with neural networks. Dhami *et al.* [11] studied the next best view selection in view planning by first completing the point cloud from partial observation. Georgakis *et al.* [14] measured uncertainty through deep ensemble. Chaplot *et al.* [8] and Ramakrishnan *et al.* [46] developed active exploration methods with deep reinforcement learning. The mapping system that uses implicit field [42, 51, 59, 75] especially 3D Gaussians [24, 37, 65, 70] as representation has been widely studied. This presents the need for active mapping algorithms that could maximize the mapping efficiency in an embodied robot system. Similar to ActiveNeRF [43], Ran *et al.* [48] performed trajectory planning for object-centric autonomous reconstruction systems by predicting variance as an extra output of the network. Yan *et al.* [67] discussed the intuition of flat minimum and quantified the predictive uncertainty of neural mapping models through the lens of loss landscape. However, they only approximate the uncertainty of the variances from the output of neighboring points. Our method directly quantifies the observed information of the parameters by computing the Hessian matrix of the log-likelihood function of the model’s objective. Besides, we introduce the Expected Information Gain for mapping systems from the first principle as a complementary method to previous heuristics.

3 Technical Approach

In this section, we first introduce how we use Fisher Information in Sec. 3.1. By leveraging Fisher Information, we demonstrate how to select the next best view and next best batch of views in Sec. 3.2 and Sec. 3.3. We then extend our method to large-scale scenes with active mapping. Finally, we showcase that our method could quantify the pixel-wise uncertainties in Sec. 3.5.

3.1 Fisher Information in Volumetric Rendering

Fisher Information is a measurement of information that an observation (\mathbf{x}, \mathbf{y}) carries about the unknown parameters \mathbf{w} that model $p(\mathbf{y}|\mathbf{x}; \mathbf{w})$. In the problem of novel view synthesis, (\mathbf{x}, \mathbf{y}) are the camera pose \mathbf{x} and image observation \mathbf{y} at pose \mathbf{x} , respectively, whereas \mathbf{w} are the parameters of the radiance field. The objective of neural rendering is equivalent to minimizing the negative log likelihood (NLL) between rendered images and ground truth images on the holdout set, which inherently represents the quality of scene reconstruction

$$-\log p(\mathbf{y}|\mathbf{x}, \mathbf{w}) = (\mathbf{y} - f(\mathbf{x}, \mathbf{w}))^T (\mathbf{y} - f(\mathbf{x}, \mathbf{w})) \quad (1)$$

where $f(\mathbf{x}, \mathbf{w})$ is our rendering model. Under regularity conditions [54], the Fisher Information of the model $\log p(\mathbf{y}|\mathbf{x}; \mathbf{w})$ is the Hessian of the log-likelihood

function with respect to the model parameters \mathbf{w} :

$$\mathcal{I}(\mathbf{w}) = -\mathbb{E}_{\log p(\mathbf{y}|\mathbf{x}, \mathbf{w})} \left[\frac{\partial^2 \log p(\mathbf{y}|\mathbf{x}, \mathbf{w})}{\partial \mathbf{w}^2} \middle| \mathbf{w} \right] = \mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}] \quad (2)$$

where $\mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}]$ is the Hessian matrix of Eq. (1).

3.2 Next Best View Selection Using Fisher Information

In the active view selection problem, we start with a training set D^{train} and have an initial estimation of parameters \mathbf{w}^* using D^{train} . The aim is to select the next best view that maximizes the Information Gain [21, 26, 32] between candidates views $\mathbf{x}_i^{acq} \in D^{pool}$ and D^{train} , where D^{pool} is the pool of candidate views:

$$\begin{aligned} \mathcal{I}[\mathbf{w}^*; \{\mathbf{y}_i^{acq}\} | \{\mathbf{x}_i^{acq}\}] \\ = H[\mathbf{w}^* | D^{train}] - H[\mathbf{w}^* | \{\mathbf{y}_i^{acq}\}, \{\mathbf{x}_i^{acq}\}, D^{train}] \end{aligned} \quad (3)$$

where $H[\cdot]$ is the entropy [28].

When the log-likelihood has the form of Eq. (1), in our case, the rendering error, the difference of the entropies in the R.H.S. of Eq. (3) can be approximated as [28]:

$$\begin{aligned} \frac{1}{2} \log \det (\mathbf{H}''[\{\mathbf{y}_i^{acq}\} | \{\mathbf{x}_i^{acq}\}, \mathbf{w}^*] \mathbf{H}''[\mathbf{w}^* | D^{train}]^{-1} + I) \\ \leq \frac{1}{2} \text{tr} (\mathbf{H}''[\{\mathbf{y}_i^{acq}\} | \{\mathbf{x}_i^{acq}\}, \mathbf{w}^*] \mathbf{H}''[\mathbf{w}^* | D^{train}]^{-1}). \end{aligned} \quad (4)$$

As Fisher Information is additive, $\mathbf{H}''[\mathbf{w}^* | D^{train}]^{-1}$ can be computed by summing the Hessians of model parameters across all different views in $\{D^{train}\}$ before inverting. We can choose the next best view \mathbf{x}_i^{acq} by optimizing

$$\arg \max_{\mathbf{x}_i^{acq}} \text{tr} (\mathbf{H}''[\mathbf{y}_i^{acq} | \mathbf{x}_i^{acq}, \mathbf{w}^*] \mathbf{H}''[\mathbf{w}^* | D^{train}]^{-1}). \quad (5)$$

The Hessian $\mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}^*]$ of our model can be computed as:

$$\mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}^*] = \nabla_{\mathbf{w}} f(\mathbf{x}; \mathbf{w}^*)^T \nabla_{f(\mathbf{x}; \mathbf{w}^*)}^2 H[\mathbf{y} | f(\mathbf{x}; \mathbf{w}^*)] \nabla_{\mathbf{w}} f(\mathbf{x}; \mathbf{w}^*) \quad (6)$$

where $\mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}^*]$ in our case is equal to the covariance of the RGB measurement that we set equal to one. Hence, the Hessian matrix can be computed just from the Jacobian matrix of $f(\mathbf{x}, \mathbf{w})$

$$\mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}^*] = \nabla_{\mathbf{w}} f(\mathbf{x}; \mathbf{w}^*)^T \nabla_{\mathbf{w}} f(\mathbf{x}; \mathbf{w}^*). \quad (7)$$

We can optimize the objective in Eq. (5) without knowing the ground truth of candidate views $\{\mathbf{y}_i^{acq}\}$, which was expected since the Fisher Information never depends on the observations themselves. The Hessian in (7) has only a limited number of off-diagonal elements because each pixel is considered independent

in $-\log p(\mathbf{y}|\mathbf{x}, \mathbf{w})$. Furthermore, recent NeRF models [40, 49, 52, 60] typically employ structured local parameters that each parameter only contributes to the radiance and density of a limited spatial region for faster convergence and rendering. Therefore, only parameters that contribute to the color of the pixels would share non-zero values in the Hessian matrix $\mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}^*]$. However, the number of optimizable parameters is typically more than 200 million, which means it is impossible to compute without sparsification or approximation. In practice, we apply Laplace approximation [9, 35] that approximates the Hessian matrix with its diagonal values plus a log-prior regularizer λI

$$\mathbf{H}''[\mathbf{y}|\mathbf{x}, \mathbf{w}^*] \simeq \text{diag}(\nabla_{\mathbf{w}} f(\mathbf{x}, \mathbf{w}^*)^T \nabla_{\mathbf{w}} f(\mathbf{x}, \mathbf{w}^*)) + \lambda I. \quad (8)$$

3.3 Batch Active View Selection

Selecting multiple views to capture new images is useful for its possible applications, such as view planning and scene reconstruction. If we simply use Eq. (5) to select a batch of acquisition samples $\{\mathbf{x}_i^{acq}\}$, we could possibly select very similar views inside the acquisition set $\{\mathbf{x}_i^{acq}\}$ as we do not consider the mutual information between our selections. However, we would face a combinatorial explosion if we directly attempt to maximize the expected information gain between training and acquisition samples and simultaneously minimize the mutual information across acquisition samples. Therefore, we employ a greedy optimization algorithm as illustrated in Alg. 1, which is $1/e$ -approximate in Fisher Information [41]. When batch size B is 1, our algorithm is equivalent to Eq. (5). Please note that we focus on batch active view selection instead of dataset sub-sampling as it is more related to real-world scenarios where we wish to plan a trajectory for an agent to acquire more training views.

Algorithm 1: Batch Active Views Selection

Input: $\{\mathbf{H}''[\mathbf{w}^*|\mathbf{x}_i^{acq}]\}$, $\mathbf{H}''[\mathbf{w}^*|D^{train}]$, number of views to select B
Output: Selected Views S_B

- 1 $S_0 \leftarrow \emptyset$;
- 2 $H_0 \leftarrow \mathbf{H}''[\mathbf{w}^*|D^{train}]$;
- 3 **for** $b \leftarrow 1$ **to** B **do**
- 4 $i \leftarrow \arg \max_{\mathbf{x}_i^{acq} \in D^{pool} \setminus S_{b-1}} \text{tr}(\mathbf{H}''[\mathbf{y}_i^{acq}|\mathbf{x}_i^{acq}, \mathbf{w}^*] H_{b-1}^{-1})$;
- 5 $S_b \leftarrow S_{b-1} \cup \{i\}$;
- 6 $H_b \leftarrow H_{b-1} + \mathbf{H}''[\mathbf{y}_i^{acq}|\mathbf{x}_i^{acq}, \mathbf{w}^*]$;
- 7 **end**

3.4 Active Mapping with 3D Gaussian Splatting

With the batch selection algorithm discussed in Sec. 3.3, our method can be extended as an active mapping algorithm. By maximizing the EIG from

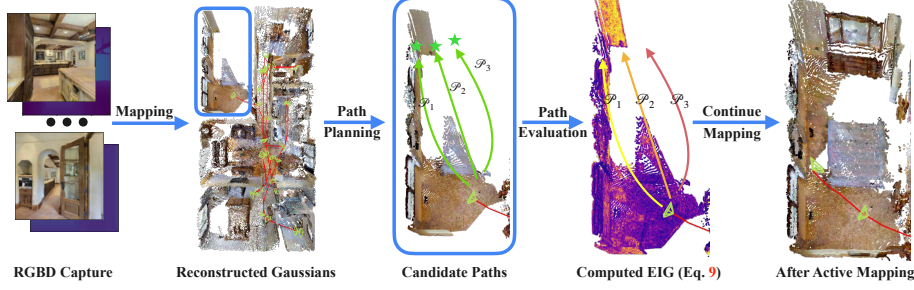


Fig. 2: An Illustration of Our Active Mapping System Given RGBD captures, we first reconstruct the environment using 3D Gaussians as representation. Afterward, we select a set of goal points from map frontiers and plan the shortest path to each goal. The EIG is computed for each path, and we choose the one with the highest EIG to continue exploration.

candidate trajectories $\{\mathcal{P}_j\}$, we can determine the most informative path for exploration:

$$\arg \max_{\mathcal{P}_j} \text{tr} \left(\left(\sum_{\mathbf{x}_i^{acq} \in \mathcal{P}_j} \mathbf{H}''[\mathbf{y}_i^{acq} | \mathbf{x}_i^{acq}, \mathbf{w}^*] \right) \mathbf{H}''[\mathbf{w}^* | D^{train}]^{-1} \right). \quad (9)$$

We build our active mapping system on top of existing passive mapping architecture SplatAM [24] that uses 3D Gaussians as basic data representation. We employ frontier-based exploration [64] to propose path candidates and sample cameras along the path concerning the mobility of the agent. The pipeline is illustrated in Fig. 2. Our path selection algorithm is compatible with other active mapping systems and can be used as an extension to select the most informative path.

3.5 Pixel-wise Uncertainty with Volumetric Rendering

Previously, we discussed quantifying the Information Gain for any camera views. Our model can also be adapted to obtain pixel-wise uncertainties. As we have discussed in Sec. 3.2, we can approximate the uncertainty on each parameter with the diagonal elements of the hessian matrix $\mathbf{H}''[\mathbf{y}_i^{acq} | \mathbf{x}_i^{acq}, \mathbf{w}^*]$. In recent Radiance Field models [25, 52], the parameters directly correspond to a spatial location in the scene. Therefore, we can compute the uncertainty in our rendered pixels by examining the diagonal Hessians along the casted ray for volumetric rendering

$$\mathbf{U}(\mathbf{r}) = \sum_{n=1}^{N_s} T_i (1 - \exp(-\sigma_n \delta_n)) \text{tr}(\mathbf{G}_n) \quad (10)$$

where \mathbf{G}_n is the submatrix of $\mathbf{H}''[\mathbf{w}^* | D^{train}]$ containing the rows and columns that correspond to parameters at location n . Please note that this is a relative

uncertainty derived from the observed information on the model parameters, which does not involve metric scales. If we wish to estimate absolute uncertainties on predictions like depth maps, we need to denormalize the uncertainty in each term by its depth d_n .

4 Experiments

In this Section, we present the empirical evaluation of our approach. The running time analysis for our customized CUDA kernel and other implementation details can be found in the supplementary. We first focus on the experiments of active view selection and compare our method quantitatively and qualitatively against the previous method (Subsection 4.1). We further present our results on active mapping (Subsection 4.2). Finally, we present more results of our method on uncertainty quantification results (Subsection 4.3).

4.1 Active View Selection

We conducted extensive experiments on view selections to demonstrate that our expected information gain could help the model find the next best views. Here, we first introduce the dataset we use and detailed experimental settings. Then, we compare our method with random baselines and previous state-of-the-art ActiveNeRF [43] quantitatively and qualitatively.

Datasets Our approach is extensively evaluated on two common benchmark datasets: Blender Dataset [38] and the challenging real-world Mip-NeRF360 dataset [5]. The Blender dataset comprises eight synthetic objects with intricate geometry and realistic non-Lambertian materials. Each scene in this dataset includes 100 training and 200 test views, all with a resolution of 800×800 . Our method uses the 100 training views as a candidate pool to select training views, and we evaluate all the models on the full 200-view test set. We use the default training configuration as in 3D Gaussian and Plenoxels for this dataset. Mip-NeRF360 [5] is a real-world 360° dataset captured for nine different scenes. It has been widely used as a quality benchmark for novel view synthesis models [6, 39]. We train our models at the resolution of 1066×1600 following 3D Gaussian Splatting [25].

Metrics Our evaluations utilize image quality metrics such as peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) [62]. Additionally, we incorporate LPIPS [73], which provides a more accurate reflection of human perception.

Baselines We quantitatively and qualitatively compare our method against the current state-of-the-art ActiveNeRF [43] and random selection baseline. To make a fair comparison with ActiveNeRF [43], We re-implemented a similar variance estimation algorithm in 3D Gaussian Splatting and Plenoxels in CUDA. We assign

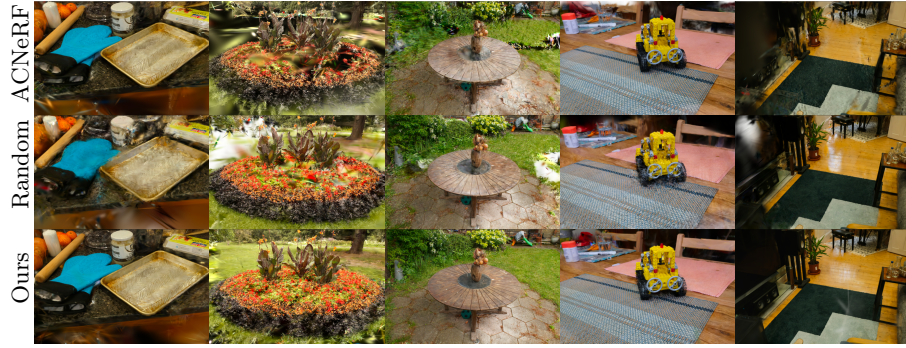


Fig. 3: Qualitative Study of our method on Mip360 Dataset From the top to bottom are results from ActiveNeRF, random baseline, our method, and the ground truth. All the models in this figure are implemented on top of 3D Gaussian Splatting [25] for better performance on this challenging dataset. We could see baseline models exhibited artifacts in some renderings due to their lack of constraints from nearby training views.

each 3D Gaussian (or grid cell in Plenoxels) a variance parameter and use volumetric rendering to render a variance map. The supplementary materials provide more details about our re-implementation. To show the difference between EIG and uncertainty, we adapt the uncertainty estimation method BayesRays [16] into the view selection algorithm by directly using the rendered uncertainty value on each candidate view as a metric for view selection.

Experiment Settings We experiment with our model with 3D Gaussian Splatting backend on both the next view selection and the next batch of view selections across both the Blender and Mip360 Dataset. Each model is initialized with the same random seed and was trained on the same four uniform views. Each model is trained for 30,000 iterations following the default configurations of 3D Gaussian Splatting [25]. Similar to the training program of Gaussian Splatting, we reset the opacity every time we select new views to avoid degeneration of the training procedure. All the external settings in the experiment are kept the same except for the view selection algorithms.

Similarly, we also showcase the implementations of our active view selection algorithm on Plenoxels in the Blender Dataset. The experimental settings for initial views and view selection schedules are the same, except view selection was made every four epochs.

- Sequential Active View Selection: 1 new view is selected every 100 epochs till the model has 20 training views.
- Batch Active View Selection: 4 new views are selected every 300 epochs till the model has 20 training views.

The quantitative results of active view selections can be found in Table. 1, Table. 2 and Table. 4. As can be seen, our method achieved better results across different metrics and datasets. We also compare our method qualitatively on the

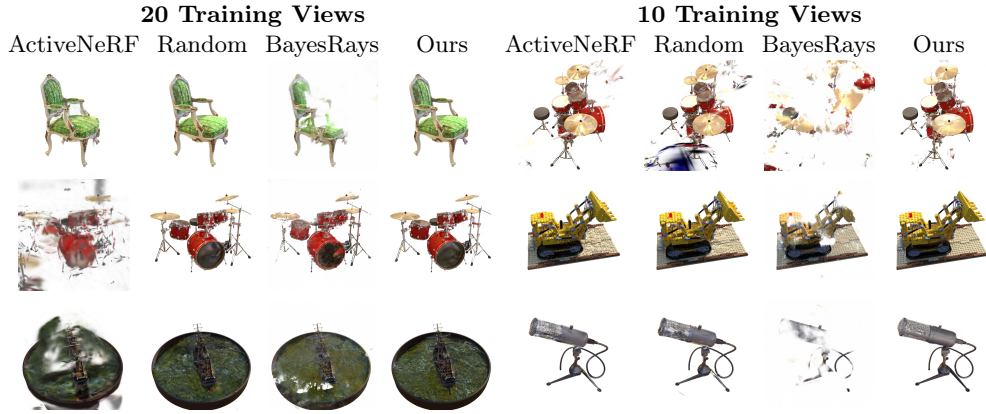


Fig. 4: Qualitative Results on Blender Dataset with 20 and 10 Training Views. All the methods are implemented on 3D Gaussian Splatting and compared in the same training configuration except for different training views selected by different methods.

Blender and Mip-NeRF360 datasets in Fig. 3 and Fig. 4. Our method demonstrates superior results in both single-view selection and batch-view selection. Moreover, our method achieved better performance improvements compared to the random backbone compared to BayesRays. This further renders the importance of using EIG rather than uncertainty as the metric for view selection. Furthermore, we experimented with our model with the challenging ten-view selection task on the Blender Dataset. Each model is initialized with the same random seed and two uniform initial views. Each new view is added after every 100 epochs till the model has ten training views. The quantitative and qualitative results are in Table. 3 and Fig. 4. Again, our method selects necessary views given the extremely limited observations and preserves fine details of reconstructed objects.

4.2 Active Mapping

We experimented with our active mapping system on room-scale scene dataset Matterport3D [7] and Gibson [63] through the Habitat Simulator [47, 53]. Our active mapping system takes posed RGB-D images as input and uses the same configurations for the simulator as previous approaches [15, 67]. We turn our 3D Gaussian representation into a point cloud using only the mean $\mu \in R^3$ of each 3D Gaussian as the previous evaluation benchmark is designed for point cloud or 3D mesh. We use the following evaluation metric to evaluate the quality of our active mapping system.

Comp. The completeness metric quantifies the completeness of the active exploration algorithm in 3D space by computing the per-point nearest distance between points sampled from the ground truth mesh and the final reconstruction

Table 1: Active View Selections on Blender Dataset The best, second, and third results are highlighted in red, orange, and yellow, respectively. *: Numbers are taken directly from ActiveNeRF [43] as reference.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
ActiveNeRF*	26.240	0.856	0.124
Plenoxels + Random	23.906	0.863	0.174
Plenoxels + ActiveNeRF	23.522	0.857	0.150
Nerfacto + Random	22.614	0.890	0.111
Nerfacto + BayesRays	22.504	0.887	0.115
3D Gaussian + Random	28.732	0.939	0.053
3D Gaussian + ActiveNeRF	26.610	0.905	0.081
Plenoxel + Ours	24.513	0.876	0.157
3D Gaussian + Ours	29.525	0.944	0.043

Table 2: Active Batch View Selections on Blender Dataset Each time, 4 views are selected based on different view selection methods. *: Numbers are taken directly from ActiveNeRF [43] as reference.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
ActiveNeRF*	26.240	0.856	0.124
Plenoxels + Random	23.242	0.862	0.158
Plenoxels + ActiveNeRF	23.147	0.857	0.148
Nerfacto + Random	21.622	0.875	0.137
Nerfacto + BayesRays	22.160	0.878	0.127
3D Gaussian + Random	27.135	0.927	0.065
3D Gaussian + ActiveNeRF	27.326	0.912	0.076
Plenoxels + Ours	24.212	0.878	0.139
3D Gaussian + Ours	29.094	0.938	0.053

Table 3: Active View Selections on Blender Dataset with only 10 views Our view selection algorithm could select necessary views when the number of training views is extremely limited.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Nerfacto + Random	16.430	0.809	0.254
Nerfacto + BayesRays	14.950	0.798	0.282
Plenoxels + Random	19.950	0.812	0.233
Plenoxels + ActiveNeRF	19.770	0.804	0.210
3D Gaussian + Random	22.493	0.873	0.112
3D Gaussian + ActiveNeRF	22.979	0.876	0.111
Plenoxels + Ours	20.670	0.824	0.205
3D Gaussian + Ours	23.681	0.883	0.102

of the scene. We calculate both the percentage of points within a 5cm threshold (Comp. (%)) and the average nearest distance (Comp. (cm)).

PSNR & Depth MAE We uniformly sample 1000 navigable poses in the scene and render color and depth images at the sampled poses and compute average PSNR and Mean Absolute Error (MAE) for rendered color and depth images, respectively.

We compare with previous state-of-the-art in Table. 5 and Figure. 5. Our active mapping system outperformed previous methods by a large margin on both geometry and rendering quality. Notably, we are the first active mapping system to capture fine-grained details and high-quality textures, thanks to our view selection objectives, which aim to improve rendering quality when measuring the EIG. In Figure 5, we demonstrate the qualitative comparison between our method and Active-INR [67]. In Table 5, our method outperforms Active-INR by a large margin in Comp. metric, showing the superiority of the active mapping algorithm. To compare the rendering quality with previous methods, we train a 3D Gaussian Splatting model using the trajectory from UPEN [15] and Active-INR [67] and evaluate the render quality at the same holdout set. The result is summarized in Table 6. Our method efficiently selects the most informative path and outperforms other methods when using the same backbone model. We also notice that UPEN is worse than Active Neural Mapping in

Table 4: Active View Selections on Mip-NeRF360 Dataset We compare our method on real-world dataset Mip360. †: batch active view selection setting. Our method outperformed previous state-of-the-art with a clear margin.

Method	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
3D Gaussian + Random	17.914	0.564	0.430
3D Gaussian + Random†	19.542	0.568	0.376
3D Gaussian + ActiveNeRF	17.889	0.533	0.414
3D Gaussian + ActiveNeRF†	18.303	0.539	0.406
3D Gaussian + Ours	20.351	0.601	0.361
3D Gaussian + Ours†	20.568	0.608	0.365

Table 5: Scene Coverage on Gibson and MP3D dataset Our method outperformed previous methods by a large margin when evaluating the coverage metric. Numbers for other methods are taken from Active Neural Mapping [67].

Method	Gibson		MP3D	
	Comp. (%) \uparrow	Comp. (cm) \downarrow	Comp. (%) \uparrow	Comp. (cm) \downarrow
Random	45.80	34.48	45.67	26.53
FBE	68.91	14.42	71.18	9.78
UPEN	63.30	21.09	69.06	10.60
OccAnt	61.88	23.25	71.72	9.40
Active Neural Mapping	80.45	7.44	73.15	9.11
Ours	92.89	5.64	89.41	2.91

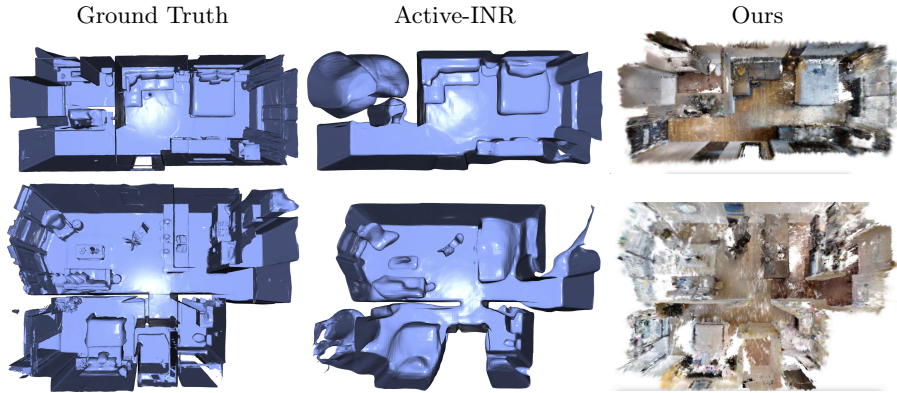


Fig. 5: Qualitative Comparisons on Active Mapping We compare our method against the ground truth mesh and Active Neural Mapping. Our method exhibited better details and coverage. Our method can successfully produce detailed and high-fidelity reconstruction for the rooms compared to the previous method.

Table 5 but better than the other in 6 when evaluated using the 3D Gaussian Splatting model. We ascribe this gap in UPEN to the difference in the scene representations and the evaluation protocol.

4.3 Uncertainty Quantification

As discussed in Sec. 3.5, our model can be extended to compute pixel-wise uncertainties on training views. Following previous methods on uncertainty estimation [55, 56], we evaluate our method on the Light Field (LF) Dataset [69] using the Area Under Sparsification Error (AUSE) metric. The pixels are filtered twice, once using the absolute error with ground truth depth and once using the uncertainty. The difference in the mean absolute error on the remaining pixels between the two sparsification processes produces two different error curvatures, where the area between those two curvatures is the AUSE, which evaluates the correlation between uncertainties and the predicted error. A low AUSE indi-

Table 6: Evaluation for Render Quality on Gibson and MP3D We also compare our method with previous method on the rendering quality for both color and depth reconstruction.

Method	Gibson		MP3D	
	PSNR \uparrow	Dpth MAE (m) \downarrow	PSNR \uparrow	Depth MAE (m) \downarrow
FBE	19.35	0.1751	16.68	0.3627
UPEN	21.13	0.1893	18.40	0.2690
Active Neural Mapping	19.64	0.2943	17.16	0.4621
Ours	22.58	0.0924	19.96	0.1667

Table 7: Uncertainty Estimation on LF Dataset Numbers are AUSE, the lower the better. The best results are highlighted in red, and the second-best results are in orange color.

Method	Statue \downarrow	Africa \downarrow	Torch \downarrow	Basket \downarrow	Average \downarrow
CF-NeRF	0.54	0.34	0.50	0.14	0.38
ActiveNeRF	0.40	0.39	0.21	0.34	0.33
BayesRays	0.17	0.27	0.22	0.28	0.23
Ours	0.21	0.26	0.24	0.18	0.22

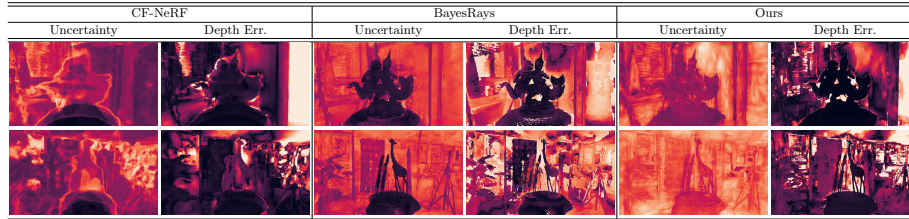


Fig. 6: Visualizations for Uncertainty on Depth Prediction the LF Dataset We compare the per-pixel uncertainty maps from CF-NeRF [55], BayesRays [16], and ours. Our per-pixel uncertainty maps demonstrate a strong correlation with the actual depth error.

cates our model is confident in the correctly estimated depths and could predict a high uncertainty in the regions where we are likely to have larger errors. As seen in Table 7 and Fig. 6, our model exhibited a better correlation between our uncertainty estimation and depth error than the previous state-of-the-art CF-NeRF [55]. More quantitative results and visualizations on uncertainty estimation can be found in the supplementary materials.

5 Conclusion and Limitations

We presented FisherRF, a novel method for active view selection, active mapping, and uncertainty quantification in Radiance Fields. Leveraging Fisher Information, our method provides an efficient and effective means to quantify the observed information of Radiance Field models. The flexibility of our approach allows it to be applied to various model parametrizations, including 3D Gaussian Splatting and Plenoxels. Our extensive evaluation of active view selection, active mapping, and uncertainty quantification has consistently shown superior performance compared to existing methods and heuristic baselines. These results highlight the potential of our approach to significantly enhance the quality and efficiency of image rendering and reconstruction tasks with limited viewpoints. However, our method is limited to static scenes in a confined scenario. Reconstructing dynamically changing the Radiance Field and quantifying its Fisher Information is still an open problem. More work could be done to overcome the limitations and extend the proposed method to more challenging settings.

Acknowledgements The authors gratefully appreciate support through the following grants: NSF FRR 2220868, NSF IIS-RI 2212433, NSF TRIPODS 1934960, NSF CPS 2038873. The authors thank Pratik Chaudhari for the insightful discussion and Yinshuang Xu for proofreading the drafts.

References

1. Ash, J.T., Goel, S., Krishnamurthy, A., Kakade, S.M.: Gone fishing: Neural active learning with fisher embeddings. In: NeurIPS. vol. abs/2106.09675 (2021)
2. Ash, J.T., Zhang, C., Krishnamurthy, A., Langford, J., Agarwal, A.: Deep batch active learning by diverse, uncertain gradient lower bounds. In: ICLR (2020), <https://openreview.net/forum?id=ryghZJBKPS>
3. Bajcsy, R.: Active perception. *Proceedings of the IEEE* **76**(8), 966–1005 (1988)
4. Bajcsy, R., Aloimonos, Y., Tsotsos, J.K.: Revisiting active perception. *Autonomous Robots* **42**, 177–196 (2018)
5. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. *CVPR* (2022)
6. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. In: ICCV (2023)
7. Chang, A., Dai, A., Funkhouser, T., Halber, M., Niessner, M., Savva, M., Song, S., Zeng, A., Zhang, Y.: Matterport3d: Learning from rgb-d data in indoor environments. *International Conference on 3D Vision* (2017)
8. Chaplot, D.S., Gandhi, D., Gupta, S., Gupta, A., Salakhutdinov, R.: Learning to explore using active neural slam. *arXiv preprint arXiv:2004.05155* (2020)
9. Daxberger, E., Kristiadi, A., Immer, A., Eschenhagen, R., Bauer, M., Hennig, P.: Laplace redux—effortless Bayesian deep learning. In: NeurIPS (2021)
10. Dellaert, F., Yen-Chen, L.: Neural volume rendering: Nerf and beyond (2021)
11. Dhami, H., Sharma, V.D., Tokekar, P.: Pred-nbv: Prediction-guided next-best-view planning for 3d object reconstruction. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 7149–7154. IEEE (2023)
12. Dornhege, C., Kleiner, A.: A frontier-void-based approach for autonomous exploration in 3d. *Advanced Robotics* **27**(6), 459–468 (2013)
13. Gao, K., Gao, Y., He, H., Lu, D., Xu, L., Li, J.: Nerf: Neural radiance field in 3d vision, a comprehensive review (2023)
14. Georgakis, G., Bucher, B., Arapin, A., Schmeckpeper, K., Matni, N., Daniilidis, K.: Uncertainty-driven planner for exploration and navigation. In: 2022 International Conference on Robotics and Automation (ICRA). pp. 11295–11302. IEEE (2022)
15. Georgakis, G., Bucher, B., Arapin, A., Schmeckpeper, K., Matni, N., Daniilidis, K.: Uncertainty-driven planner for exploration and navigation. In: ICRA (2022)
16. Goli, L., Reading, C., Sellán, S., Jacobson, A., Tagliasacchi, A.: Bayes’ Rays: Uncertainty quantification in neural radiance fields. *arXiv* (2023)
17. Guédon, A., Monasse, P., Lepetit, V.: Scone: Surface coverage optimization in unknown environments by volumetric integration. *Advances in Neural Information Processing Systems* **35**, 20731–20743 (2022)
18. Guédon, A., Monnier, T., Monasse, P., Lepetit, V.: Macarons: Mapping and coverage anticipation with rgb online self-supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 940–951 (2023)

19. Hinton, G.E., van Camp, D.: Keeping the neural networks simple by minimizing the description length of the weights. In: Proceedings of the Sixth Annual Conference on Computational Learning Theory. p. 5–13. COLT '93, Association for Computing Machinery, New York, NY, USA (1993). <https://doi.org/10.1145/168304.168306>, <https://doi.org/10.1145/168304.168306>
20. Hochreiter, S., Schmidhuber, J.: Simplifying neural nets by discovering flat minima. In: NeurIPS (1994)
21. Houlaby, N., Huszar, F., Ghahramani, Z., Lengyel, M.: Bayesian active learning for classification and preference learning. CoRR **abs/1112.5745** (2011), <http://dblp.uni-trier.de/db/journals/corr/corr1112.html#abs-1112-5745>
22. Jin, L., Chen, X., Rückin, J., Popović, M.: Neu-nbv: Next best view planning using uncertainty estimation in image-based neural rendering. arXiv preprint arXiv:2303.01284 (2023)
23. Karaman, S., Frazzoli, E.: Incremental sampling-based algorithms for optimal motion planning. Robotics Science and Systems VI **104**(2), 267–274 (2010)
24. Keetha, N., Karhade, J., Jatavallabhula, K.M., Yang, G., Scherer, S., Ramanan, D., Luiten, J.: Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. arXiv preprint arXiv:2312.02126 (2023)
25. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics **42**(4) (July 2023), <https://repo-sam.inria.fr/fungraph/3d-gaussian-splatting/>
26. Kirsch, A., Amersfoort, J.v., Gal, Y.: Batchbald: Efficient and diverse batch acquisition for deep bayesian active learning. In: NeurIPS (2019)
27. Kirsch, A., Farquhar, S., Atighehchian, P., Jesson, A., Branchaud-Charron, F., Gal, Y.: Stochastic batch acquisition for deep active learning. arXiv preprint arXiv:2106.12059 (2021)
28. Kirsch, A., Gal, Y.: Unifying approaches in active learning and active sampling via fisher information and information-theoretic quantities. Transactions on Machine Learning Research (2022), <https://openreview.net/forum?id=UVDKQANOW>, expert Certification
29. Kothawade, S.N., Beck, N.A., Killamsetty, K., Iyer, R.K.: SIMILAR: Submodular information measures based active learning in realistic scenarios. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems (2021), <https://openreview.net/forum?id=VGDFaLNFFk>
30. LaValle, S.: Rapidly-exploring random trees: A new tool for path planning. Research Report 9811 (1998)
31. Lee, S., Chen, L., Wang, J., Liniger, A., Kumar, S., Yu, F.: Uncertainty guided policy for active robotic 3d reconstruction using neural radiance fields. IEEE Robotics and Automation Letters **7**(4), 12070–12077 (2022)
32. Lindley, D.V.: On a Measure of the Information Provided by an Experiment. The Annals of Mathematical Statistics **27**(4), 986 – 1005 (1956). <https://doi.org/10.1214/aoms/1177728069>, <https://doi.org/10.1214/aoms/1177728069>
33. Lluvia, I., Lazkano, E., Ansuategi, A.: Active mapping and robot exploration: A survey. Sensors **21**(7), 2445 (2021)
34. MacDonald, L.E., Valmadre, J., Lucey, S.: On progressive sharpening, flat minima and generalisation (2023)
35. MacKay, D.J.C.: Bayesian Interpolation. Neural Computation **4**(3), 415–447 (05 1992). <https://doi.org/10.1162/neco.1992.4.3.415>, <https://doi.org/10.1162/neco.1992.4.3.415>

36. Martin-Brualla, R., Radwan, N., Sajjadi, M.S.M., Barron, J.T., Dosovitskiy, A., Duckworth, D.: NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In: CVPR (2021)
37. Matsuki, H., Murai, R., Kelly, P.H.J., Davison, A.J.: Gaussian Splatting SLAM (2024)
38. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
39. Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P., Martin-Brualla, R., Barron, J.T.: MultiNeRF: A Code Release for Mip-NeRF 360, Ref-NeRF, and RawNeRF (2022), <https://github.com/google-research/multinerf>
40. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.* **41**(4), 102:1–102:15 (Jul 2022). <https://doi.org/10.1145/3528223.3530127>, <https://doi.org/10.1145/3528223.3530127>
41. Nemhauser, G.L., Wolsey, L.A., Fisher, M.L.: An analysis of approximations for maximizing submodular set functions—i. *Mathematical programming* **14**, 265–294 (1978)
42. Ortiz, J., Clegg, A., Dong, J., Sucar, E., Novotny, D., Zollhoefer, M., Mukadam, M.: isdf: Real-time neural signed distance fields for robot perception. In: *Robotics: Science and Systems* (2022)
43. Pan, X., Lai, Z., Song, S., Huang, G.: Activenerf: Learning where to see with uncertainty estimation. In: ECCV. pp. 230–246. Springer (2022)
44. Papachristos, C., Khattak, S., Alexis, K.: Uncertainty-aware receding horizon exploration and mapping using aerial robots. In: 2017 IEEE international conference on robotics and automation (ICRA). pp. 4568–4575. IEEE (2017)
45. Placed, J.A., Strader, J., Carrillo, H., Atanasov, N., Indelman, V., Carlone, L., Castellanos, J.A.: A survey on active simultaneous localization and mapping: State of the art and new frontiers. *IEEE Transactions on Robotics* (2023)
46. Ramakrishnan, S.K., Al-Halah, Z., Grauman, K.: Occupancy anticipation for efficient exploration and navigation. In: *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V* 16. pp. 400–418. Springer (2020)
47. Ramakrishnan, S.K., Gokaslan, A., Wijmans, E., Maksymets, O., Clegg, A., Turner, J.M., Undersander, E., Galuba, W., Westbury, A., Chang, A.X., Savva, M., Zhao, Y., Batra, D.: 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)
48. Ran, Y., Zeng, J., He, S., Chen, J., Li, L., Chen, Y., Lee, G., Ye, Q.: Neurar: Neural uncertainty for autonomous 3d reconstruction with implicit neural representations. *IEEE Robotics and Automation Letters* **8**(2), 1125–1132 (Feb 2023). <https://doi.org/10.1109/lra.2023.3235686>, <http://dx.doi.org/10.1109/LRA.2023.3235686>
49. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: ICCV (2021)
50. Ren, P., Xiao, Y., Chang, X., Huang, P.Y., Li, Z., Gupta, B.B., Chen, X., Wang, X.: A survey of deep active learning. *ACM computing surveys (CSUR)* **54**(9), 1–40 (2021)
51. Sandström, E., Li, Y., Van Gool, L., Oswald, M.R.: Point-slam: Dense neural point cloud-based slam. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 18433–18444 (2023)

52. Sara Fridovich-Keil and Alex Yu, Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: CVPR (2022)
53. Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J., Parikh, D., Batra, D.: Habitat: A Platform for Embodied AI Research. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
54. Schervish, M.: Theory of Statistics. Springer Series in Statistics, Springer New York (2012), <https://books.google.com/books?id=s5LHBgAAQBAJ>
55. Shen, J., Agudo, A., Moreno-Noguer, F., Ruiz, A.: Conditional-flow nerf: Accurate 3d modelling with reliable uncertainty quantification. In: ECCV (2022)
56. Shen, J., Ruiz, A., Agudo, A., Moreno-Noguer, F.: Stochastic neural radiance fields: Quantifying uncertainty in implicit 3d representations. CoRR **abs/2109.02123** (2021), <https://arxiv.org/abs/2109.02123>
57. Shen, S., Michael, N., Kumar, V.: Autonomous indoor 3d exploration with a micro-aerial vehicle. In: 2012 IEEE international conference on robotics and automation. pp. 9–15. IEEE (2012)
58. Smith, S., Le, Q.V.: A bayesian perspective on generalization and stochastic gradient descent. In: ICLR (2018), <https://openreview.net/pdf?id=BJij4yg0Z>
59. Sucar, E., Liu, S., Ortiz, J., Davison, A.: iMAP: Implicit mapping and positioning in real-time. In: Proceedings of the International Conference on Computer Vision (ICCV) (2021)
60. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In: CVPR (2022)
61. Sünderhauf, N., Abou-Chakra, J., Miller, D.: Density-aware nerf ensembles: Quantifying predictive uncertainty in neural radiance fields. In: ICRA (2023)
62. Wang, Z., Simoncelli, E.P., Bovik, A.C.: Multiscale structural similarity for image quality assessment. In: NeurIPS. vol. 2, pp. 1398–1402. IEEE (2003)
63. Xia, F., Zamir, A.R., He, Z., Sax, A., Malik, J., Savarese, S.: Gibson env: Real-world perception for embodied agents. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 9068–9079 (2018)
64. Yamauchi, B.: 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’. pp. 146–151. IEEE (1997)
65. Yan, C., Qu, D., Wang, D., Xu, D., Wang, Z., Zhao, B., Li, X.: Gs-slam: Dense visual slam with 3d gaussian splatting. arXiv preprint arXiv:2311.11700 (2023)
66. Yan, D., Liu, J., Quan, F., Chen, H., Fu, M.: Active implicit object reconstruction using uncertainty-guided next-best-view optimization (2023)
67. Yan, Z., Yang, H., Zha, H.: Active neural mapping. In: ICCV (2023)
68. Ye, K., Dong, S., Fan, Q., Wang, H., Yi, L., Xia, F., Wang, J., Chen, B.: Multi-robot active mapping via neural bipartite graph matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14839–14848 (2022)
69. Yücer, K., Sorkine-Hornung, A., Wang, O., Sorkine-Hornung, O.: Efficient 3d object segmentation from densely sampled light fields with applications to 3d reconstruction. ACM Trans. Graph. **35**(3) (mar 2016). <https://doi.org/10.1145/2876504>, <https://doi.org/10.1145/2876504>
70. Yugay, V., Li, Y., Gevers, T., Oswald, M.R.: Gaussian-slam: Photo-realistic dense slam with gaussian splatting (2023)
71. Zhan, H., Zheng, J., Xu, Y., Reid, I., Rezatofighi, H.: Activermap: Radiance field for active mapping and planning (2022)

- 72. Zhan, X., Wang, Q., hao Huang, K., Xiong, H., Dou, D., Chan, A.B.: A comparative survey of deep active learning (2022)
- 73. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR. pp. 586–595 (2018)
- 74. Zhu, C., Ding, R., Lin, M., Wu, Y.: A 3d frontier-based exploration tool for mavs. In: 2015 IEEE 27th International Conference on Tools with Artificial Intelligence (ICTAI). pp. 348–352. IEEE (2015)
- 75. Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z., Oswald, M.R., Pollefeys, M.: Nice-slam: Neural implicit scalable encoding for slam (2021)