

Modeling Uncertainty for Gaussian Splatting

Original

Modeling Uncertainty for Gaussian Splatting / SAVANT AIRA, L., Valsesia, D., Magli, E.. - In: IEEE TRANSACTIONS ON NEURAL NETWORKS AND LEARNING SYSTEMS. - ISSN 2162-237X. - (2025), pp. 1-7.
[10.1109/TNNLS.2025.3553582]

Availability:

This version is available at: 11583/3000131 since: 2025-05-14T13:49:15Z

Publisher:

IEEE

Published

DOI:10.1109/TNNLS.2025.3553582

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

IEEE postprint/Author's Accepted Manuscript

©2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collecting works, for resale or lists, or reuse of any copyrighted component of this work in other works.

(Article begins on next page)

Modeling uncertainty for Gaussian Splatting

Luca Savant Aira, *Student Member, IEEE*, Diego Valsesia, *Member, IEEE*, Enrico Magli, *Fellow, IEEE*

Abstract—We present Stochastic Gaussian Splatting (SGS): the first framework for uncertainty estimation using Gaussian Splatting (GS). GS recently advanced the novel-view synthesis field by achieving impressive reconstruction quality at a fraction of the computational cost of Neural Radiance Fields (NeRF). However, contrary to the latter, it still lacks the ability to provide information about the confidence associated with their outputs. To address this limitation, in this paper, we introduce a Variational Inference-based approach that seamlessly integrates uncertainty prediction into the common rendering pipeline of GS. Additionally, we introduce the Area Under Sparsification Error (AUSE) as a new term in the loss function, enabling optimization of uncertainty estimation alongside image reconstruction. Experimental results on the three different datasets demonstrate that our method outperforms existing approaches in terms of both image rendering quality and uncertainty estimation accuracy. Overall, our framework equips practitioners with valuable insights into the reliability of synthesized views, facilitating safer decision-making in real-world applications.

I. INTRODUCTION

Novel-view synthesis, the task of generating images of a scene from viewpoints not observed during data collection, is a fundamental problem in computer vision with numerous applications, including virtual reality, augmented reality, and robotics. Traditionally, this task has been addressed using methods such as Structure from Motion [1] which rely on geometric reconstruction techniques. However, recent advances in deep learning, particularly with the introduction of Neural Radiance Fields (NeRF), have revolutionized the field by enabling high-fidelity synthesis of novel views directly from the underlying scene representation.

NeRF [2] has recently enjoyed great success by representing a scene as a continuous volumetric function that maps 3D spatial coordinates and viewing directions to radiance values. By learning this function from a set of posed images, NeRF can generate photorealistic images from novel viewpoints. However, while NeRF achieves impressive results, its computational complexity and memory requirements limit its practicality for real-time applications. This is the focus of the emerging Gaussian Splatting (GS) technique [3] which offers a more computationally efficient alternative to NeRF while maintaining high-quality novel-view synthesis. GS learns to approximate the radiance field by using a set of Gaussian kernels, enabling real-time rendering with competitive visual fidelity.

The authors are with Politecnico di Torino - Department of Electronics and Telecommunications, Italy. Email: name.surname@polito.it.

This publication is part of the project PNRR-NGEU which has received funding from the MUR - DM 352/2022. This work was partially supported by the European Union under the Italian National Recovery and Resilience Plan (NRRP) of NextGenerationEU, partnership on “Telecommunications of the Future” (PE00000001 - program “RESTART”).

At the same time, research in novel view synthesis has started addressing the problem of estimating the epistemic or aleatoric uncertainty in order to understand the reliability of the generated views. Indeed, any practical downstream task that involves taking actions in the real world (such as robotics and autonomous systems) must consider not only the newly synthesized views but also their corresponding uncertainties, in order to potentially discard too uncertain yet promising actions. As an example, a robot moving in the real world could use a GS representation of the environment around itself, and plans movement according to some goal. It is very important that this goal takes into account also the uncertainty of the representation of the environment, as in application riskier moves (moving too close to a too-much uncertain obstacle) should be penalized. This scenario was first addressed in the seminal work of Shen et al. [4], where they proposed a deep architecture, based on NeRF, called S-NeRF to also estimate meaningful (aleatoric) uncertainty maps for each generated view.

At the moment, GS lacks a mechanism for estimating uncertainty in the synthesized views. In this paper, we seek to address this limitation by proposing a novel framework for uncertainty estimation in GS. We extend the traditional deterministic GS framework to incorporate stochasticity, allowing us to predict uncertainty alongside synthesized views. Our approach leverages Variational Inference (VI) to learn the parameters of the GS radiance field in a Bayesian framework, enabling us to accurately estimate uncertainty without sacrificing computational efficiency. As done in S-NeRF and the subsequent work CF-NeRF, our method SGS primarily focuses on aleatoric uncertainty, not epistemic uncertainty. Aleatoric uncertainty refers to the irreducible uncertainty inherent in the data due to stochastic processes, such as noise in image acquisition or variability in camera pose estimation, while epistemic uncertainty refers to the model’s lack of knowledge, e.g., due to limited training data.

We can summarize our novel contributions as follows:

- we introduce a novel framework for uncertainty estimation in GS, called Stochastic Gaussian Splatting (SGS), enabling real-time synthesis of high-quality images with accurate uncertainty predictions;
- we propose a VI-based approach to learn the parameters of the GS radiance field, allowing us to incorporate uncertainty prediction seamlessly into the rendering pipeline. Moreover, we innovate this learning process by augmenting Empirical Bayes with a loss function dependent on the area under the sparsification curve;
- we demonstrate the effectiveness of our approach through experiments on three different datasets (LLFF, Blender and Mip-NeRF360), showing significant improvements in both rendering quality and uncertainty estimation metrics

compared to state-of-the-art methods.

II. BACKGROUND

A. NeRF and Gaussian Splatting

In recent years, the Structure from Motion and novel view synthesis tasks have been revolutionized by novel techniques based on Neural Radiance Fields (NeRF), introduced in [2] and based on deep neural network architectures. Following the paradigm of implicit neural representation [5], the NeRF network learns a mapping $(\mathbf{x} \in \mathbb{R}^3, \mathbf{d} \in \mathbb{S}^2) \rightarrow (c(\mathbf{x}, \mathbf{d}), \sigma(\mathbf{x}))$ where \mathbf{x} represents a 3D point, \mathbf{d} a view direction, c the color field, and σ the density field. Using the Direct Volume Rendering Integral [6], [7], the physical radiance field \mathbf{I} can be obtained from c and σ , hence an image can be formed. This learning process can be supervised with the pixels from a number of available views, so that the network correctly resynthesizes the original images. After successful training, the network is also able to coherently synthesize novel views.

In an ever-growing literature around NeRF, GS [3] distinguishes itself by achieving state-of-the-art visual quality while maintaining competitive training times and, importantly, allowing high-quality and real-time novel-view synthesis (≥ 30 fps at 1080p resolution). In order to achieve such high inference speed, GS replaces the deep neural network with an elliptical basis function approximation. By carefully choosing the basis functions as elliptical Gaussian kernels, the EWA Volume Splatting technique [8] can be used to achieve real-time view synthesis.

B. Structure from Motion uncertainty estimation

Uncertainty estimation is a long standing problem in the context of machine learning and computer vision, for example semantic segmentation and depth regression tasks, and it is addressed in many works [9]–[11]. Some broader works such as [12], [13] provides an excellent starting point for uncertainty estimation, Bayesian learning and variational inference applied to deep learning.

An emerging research direction for the Structure from Motion field is the quantification of the uncertainty of the synthesized novel-views or of the 3D radiance field itself. This task was first addressed by [4], where it is cast as a Bayesian learning problem and solved with the Variational Inference (VI) framework [13], applied to NeRF. Subsequently, the same authors proposed an evolution of their method [14], which differs from the previous work by dropping the independence assumptions between the color and opacity fields, hence recovering a more complex stochastic dependency graph using the Conditional Normalizing Flows framework [15]. Following these works, [16] drops all independence assumptions by using a generative Flow-GAN model [17]. Another approach is addressed in [18] where the Laplace Approximation framework [13] is used. Finally, two similar works, [19] and [20], focus on developing methods that actively estimate high uncertainty in spatial areas that are not covered by the input views, using VI and Ensemble Learning frameworks, respectively. It is worth noting that all these methods rely on the use of NeRF, while, to the best of our knowledge, there is no work currently

addressing the GS framework. Filling this void in literature is interesting, as GS is both more computational efficient than NeRF and more easily interpretable. Moreover, it is not possible to apply the same techniques used in some previous works on NeRF to GS, as they exploit the presence of a deep neural network, for example [21] or [22].

C. Radiance Field and Direct Volume Rendering

In this section, we recall the notation for volume rendering, which will be useful in the remainder of the paper. Following the optical derivation of [6], let's denote the physical light intensity field with I , which increases or decreases along a ray segment, parameterized by t , by interacting with particles in a particle-filled volume. These interactions are described using the density field σ (also called the extinction coefficient) and the color field c (also called the emission term), with the following Cauchy problem:

$$\begin{cases} \frac{dI}{dt} = -c(t)\sigma(t) + \sigma(t)I(t) & \forall t \in [0, T] \\ I(T) = I_T \end{cases} \quad (1)$$

whose solution value at $t = 0$ is:

$$I(0) = I_T e^{-\int_0^T \sigma(t)dt} + \int_0^T e^{-\int_0^t \sigma(z)dz} c(t)\sigma(t)dt \quad (2)$$

Reparametrizing the segment with $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, setting $I_T = 0$ (in the case of a dark background), and assuming an isotropic density field, the rendering integral is obtained:

$$I(\mathbf{o}, \mathbf{d}) = \int_0^T e^{-\int_0^t \sigma(\mathbf{r}(s))ds} c(\mathbf{r}(t), \mathbf{d}) \sigma(\mathbf{r}(t)) dt \quad (3)$$

From a given set of posed views, a set of corresponding posed pixels and colors can be extracted as $\mathcal{D} = \{(\mathbf{o}_i, \mathbf{d}_i), y_i\}_i$, where i is an index that runs over all the pixels of all the images in the training set, \mathbf{o}_i is the camera center of the i -th pixel, \mathbf{d}_i is the 3d spatial direction of the i -th pixel from its camera center and y_i is the color of the i -th pixel, so that a reconstruction loss can be used to regress the color and density fields:

$$\mathcal{L}_{\text{rec}} = \sum_i \ell(y_i, I(\mathbf{o}_i, \mathbf{d}_i)) \quad (4)$$

To synthesize a novel view, the integral in Eq. (3) must be evaluated by placing \mathbf{o} at the camera origin and letting \mathbf{d} vary for each pixel in the new image plane.

III. METHOD

In this section, we present the proposed method, called Stochastic Gaussian Splatting (SGS), to enable uncertainty quantification in the Gaussian Splatting framework.




Fig. 1. Bayesian Network Graphical Model of SGS. Learnable variables are depicted in blue, while stochastic variables are circled. The “Camera” node represents both the spatial coordinate of the pixel (\mathbf{o}, \mathbf{d}) and the corresponding camera intrinsic and extrinsic parameters. Gray dashed rectangles are used for the plate notation, i.e. variables repetitions.

A. Stochastic Gaussian Splatting

The main difference between NeRF and GS lies in the fact that in the original NeRF formulation [2], the color and density fields are regressed with a multilayer perceptron neural network. Instead, in GS [3], the two fields are learned with an elliptical basis function approximation. Let \mathbf{x} be any point in 3D space, and let \mathbf{d} identify a view direction, then the density and color fields are obtained as:

$$\begin{cases} \sigma(\mathbf{x}) \approx \sum_{k=1}^K \alpha_k \phi(\mathbf{x}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \\ c(\mathbf{x}, \mathbf{d}) \approx c_k(\mathbf{d}) \quad \text{on the support of the } k\text{-th kernel} \end{cases} \quad (5)$$

where:

- ϕ is the elliptical 3D Gaussian kernel, with learnable parameters $\boldsymbol{\mu}_k \in \mathbb{R}^3$ and $\boldsymbol{\Sigma}_k \in \mathbb{R}^{3 \times 3}$ to control the shape and position of the k -th kernel. Moreover, GS assumes that kernel supports overlapping is negligible.

$$\phi(\mathbf{x}; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) = \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu}_k)^T \boldsymbol{\Sigma}_k^{-1}(\mathbf{x} - \boldsymbol{\mu}_k)\right); \quad (6)$$

- $\alpha_k \in [0, 1]$ is a learnable scalar to control the opacity of the k -th kernel;
- $c_k(\mathbf{d})$ is a learned linear combination of spherical harmonics Y_{lm} :

$$c_k(\mathbf{d}) = \sum_{l=0}^L \sum_{m=0}^l c_{klm} Y_{lm}(\mathbf{d}), \quad (7)$$

where l, m are degree and order of the used spherical harmonics.

In GS, the learnable parameters for each Gaussian kernel are $\boldsymbol{\mu}_k$, $\boldsymbol{\Sigma}_k$, c_{klm} , and α_k .

The specific definition of ϕ enables the EWA Volume Splatting technique [8], so that the integral in Eq. (3) can be reduced to the simpler alpha blending technique:

$$I(\mathbf{o}, \mathbf{d}) = \sum_{k=1}^K \alpha_k e^{z_k(\mathbf{o}, \mathbf{d})} c_k(\mathbf{d}) \prod_{j=1}^{k-1} (1 - \alpha_j e^{z_j(\mathbf{o}, \mathbf{d})}) \quad (8)$$

where $z_k(\mathbf{o}, \mathbf{d})$ is the splatting coefficient [8]. This reduction is essentially obtained from Eq. (3) by exploiting the linearity of the integral along the number of gaussian primitives, and leveraging the fact that there exists a closed form solution of these integrals if the opacity falloff follows a gaussian distribution.

Aligning with previous works in the field, the proposed SGS method predicts an uncertainty value for each pixel in the novel-synthesized view. It is natural to define this uncertainty as the standard deviation of the predicted pixel color, requiring the predicted colors to be random variables. Hence, we need to inject stochasticity into the otherwise deterministic process of Eq. (8). We propose to use a Monte Carlo method (sampling the model $S \geq 1$ times) to approximate the variance of pixel colors, so that the expression (8) should be evaluated multiple times, each time by sampling some random variables.

Following the Variational Inference framework, we directly expand upon GS by imposing a prior distribution on each of the following parameters:

- $\boldsymbol{\mu}_k \sim \mathcal{N}(\boldsymbol{\mu}_k^*, \boldsymbol{\Gamma}_k^*)$
- $\ln \frac{\alpha_k}{1 - \alpha_k} \sim \mathcal{N}(\alpha_k^*, \boldsymbol{\Pi}_k^*)$
- $c_{klm} \sim \mathcal{N}(c_{klm}^*, \boldsymbol{\Xi}_{klm}^*)$

so that the original GS parameters are no longer learned but are sampled from the above distributions, whose parameters are the new learning variables. In order to keep the model simple, we decided not to model the uncertainty of the rotation and scaling of the gaussians primitives (hence their covariance matrices), but to model their positions, colors and opacities.

Thanks to the reparameterization trick [23], gradients can flow from the loss in Eq. (4), which is a function of the sam-

ples, to the new distribution parameters $\boldsymbol{\mu}_k^*$, $\boldsymbol{\Gamma}_k^*$, α_k^* , Π_k^* , C_{klm}^* and Ξ_{klm}^* , enabling learning with standard backpropagation techniques.

B. Learning with Variational Inference

The Variational framework introduced in [13] and used in [4] and [14] is required for optimization, since direct optimization of just the loss would otherwise incur in underestimation of the pixel variance as the model could reduce to a deterministic one.

The VI framework stems from an approximation of Bayes' Theorem. Let \mathcal{G} represent a Gaussian Splatting Radiance Field, composed of K Gaussian kernels \mathcal{G}_k , and \mathcal{D} represent the pixels dataset in Eq. (4). Then the Bayes' Theorem reads as:

$$\mathbb{P}(\mathcal{G}|\mathcal{D}) = \frac{\mathbb{P}(\mathcal{D}|\mathcal{G})\mathbb{P}(\mathcal{G})}{\mathbb{P}(\mathcal{D})}. \quad (9)$$

Since this is generally intractable, the VI framework prescribes to approximate the true posterior $\mathbb{P}(\mathcal{G}|\mathcal{D})$ by introducing a parametric distribution q_θ over all GS radiance fields \mathcal{G} and to learn these parameters θ in order to minimize the Kullback-Leibler (KL) divergence between the approximate posterior and the true one:

$$\min_{\theta} \text{KL}(q_\theta(\mathcal{G}) \parallel \mathbb{P}(\mathcal{G}|\mathcal{D})) \quad (10)$$

Now, this problem is further manipulated to get a tractable expression. Let us start with the following manipulation, where only properties of the logarithm and linearity on the expectation are used:

$$\begin{aligned} \text{KL}(q_\theta(\mathcal{G}) \parallel \mathbb{P}(\mathcal{G}|\mathcal{D})) &= \\ &= \mathbb{E}_{q_\theta} \left[\log \frac{q_\theta(\mathcal{G})}{\mathbb{P}(\mathcal{G}|\mathcal{D})} \right] = \mathbb{E}_{q_\theta} \left[\log \frac{q_\theta(\mathcal{G})\mathbb{P}(\mathcal{D})}{\mathbb{P}(\mathcal{D}|\mathcal{G})\mathbb{P}(\mathcal{G})} \right] \\ &= -\mathbb{E}_{q_\theta} [\log \mathbb{P}(\mathcal{D}|\mathcal{G})] + \mathbb{E}_{q_\theta} \left[\log \frac{q_\theta(\mathcal{G})}{\mathbb{P}(\mathcal{G})} \right] + \mathbb{E}_{q_\theta} [\log \mathbb{P}(\mathcal{D})] \\ &= -\mathbb{E}_{q_\theta} [\log \mathbb{P}(\mathcal{D}|\mathcal{G})] + \text{KL}(q_\theta(\mathcal{G}) \parallel \mathbb{P}(\mathcal{G})) + \log \mathbb{P}(\mathcal{D}) \end{aligned} \quad (11)$$

Now, remembering that $\mathcal{D} = \{((\mathbf{o}_i, \mathbf{d}_i), y_i)\}_i$, and that the dataset samples are independently sampled, the first term in the last expression is equivalent to:

$$\mathbb{E}_{q_\theta} [\log \mathbb{P}(\mathcal{D}|\mathcal{G})] = \quad (12)$$

$$= \mathbb{E}_{q_\theta} \left[\log \prod_i \mathbb{P}(\mathcal{D}_i|\mathcal{G}) \right] \quad (13)$$

$$= \sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(\mathcal{D}_i|\mathcal{G})] \quad (14)$$

$$= \sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(((\mathbf{o}_i, \mathbf{d}_i), y_i) | \mathcal{G})] \quad (15)$$

$$= \sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G}) \mathbb{P}((\mathbf{o}_i, \mathbf{d}_i) | \mathcal{G})] \quad (16)$$

$$= \sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G}) \mathbb{P}(\mathbf{o}_i, \mathbf{d}_i)] \quad (17)$$

$$= \sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G})] + \mathbb{E}_{q_\theta} [\log \mathbb{P}(\mathbf{o}_i, \mathbf{d}_i)] \quad (18)$$

$$= \sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G})] + \log \mathbb{P}(\mathbf{o}_i, \mathbf{d}_i) \quad (19)$$

Plugging (12) into (11), we get:

$$\begin{aligned} \text{KL}(q_\theta(\mathcal{G}) \parallel \mathbb{P}(\mathcal{G}|\mathcal{D})) &= \\ &= -\sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G})] + \\ &\quad -\sum_i \log \mathbb{P}(\mathbf{o}_i, \mathbf{d}_i) + \\ &\quad + \text{KL}(q_\theta(\mathcal{G}) \parallel \mathbb{P}(\mathcal{G})) + \log \mathbb{P}(\mathcal{D}) \end{aligned} \quad (20)$$

As the optimization variable is the parameters θ , all the terms that do not depend on θ in the latter equation can be discarded. Finally, we have the optimization problem:

$$\min_{\theta} -\sum_i \mathbb{E}_{q_\theta} [\log \mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G})] + \text{KL}(q_\theta(\mathcal{G}) \parallel \mathbb{P}(\mathcal{G})) \quad (21)$$

The first term in the loss function of problem (21) is the expected negative log-likelihood. This term forces θ to maximize the expected log-likelihood by matching the observations in the dataset \mathcal{D} . So, in spirit, it replaces the standard loss (4). It is estimated with the Monte Carlo method (with S samples) and, for simplicity, by defining the conditional probability distribution $\mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G})$ to be a normal distribution that is pixel-wise independent, i.e.:

$$\log \mathbb{P}(y_i | (\mathbf{o}_i, \mathbf{d}_i), \mathcal{G}) \propto (y_i - I(\mathbf{o}_i, \mathbf{d}_i))^2 \quad (22)$$

Note that this requirement is not too strict as the independence is required in the conditional probability distribution and not for the unconditional probability distribution.

The second term in Eq. (21), instead, limits θ from moving too far away from the prior distribution of the GS radiance field \mathcal{G} . For it to be efficiently tractable, we suppose independence among the K Gaussian kernels, for both the prior and the approximate posterior distributions, i.e.:

$$q_\theta(\mathcal{G}) = \prod_{k=1}^K q_{\theta_k}(\mathcal{G}_k) \quad (23)$$

$$\mathbb{P}(\mathcal{G}) = \prod_{k=1}^K \mathbb{P}(\mathcal{G}_k) \quad (24)$$

This independence assumption between Gaussian kernels is more general with respect to previous work. For instance, Shen et al. [4] prescribed independence between the values of the opacity fields for every pair of points in the radiance field, even if they are very close to each other. Meanwhile, in the proposed method, the independence is prescribed at the Gaussian kernel level and not at the infinitesimal 3D point level.

Thanks to the assumptions in Eqs. (23) and (24), the second term in Eq. (21) can be expanded to become:

$$\text{KL}(q_\theta(\mathcal{G}) \parallel \mathbb{P}(\mathcal{G})) = \sum_{k=1}^K \text{KL}(q_{\theta_k}(\mathcal{G}_k) \parallel \mathbb{P}(\mathcal{G}_k)) \quad (25)$$

If we suppose that the prior distributions are themselves multivariate Normal, each KL term in the right-hand side of Eq. (25) has the following general closed form expression:

$$\begin{aligned} \text{KL}(\mathcal{N}_{\mu_0, \Sigma_0} \parallel \mathcal{N}_{\mu_1, \Sigma_1}) &= -\frac{3}{2} + \frac{1}{2} \text{Tr}(\Sigma_1^{-1} \Sigma_0) \\ &+ \frac{1}{2} (\mu_1 - \mu_0)^T \Sigma_1^{-1} (\mu_1 - \mu_0) \\ &+ \frac{1}{2} \ln \left(\frac{\det \Sigma_1}{\det \Sigma_0} \right) \end{aligned}$$

Hence, we define the total KL contribution to the training loss \mathcal{L}_{KL} as the sum over all Gaussians and over all the learnable parameters of the KL divergence between the prior (hat variables) and the posterior (starred variables) of that parameter:

$$\begin{aligned} \mathcal{L}_{\text{KL}} &= \sum_{k=1}^K \left[\text{KL}(\mathcal{N}_{\hat{\mu}_k, \hat{\Gamma}_k} \parallel \mathcal{N}_{\mu_k^*, \Gamma_k^*}) \right. \\ &+ \text{KL}(\mathcal{N}_{\hat{\alpha}_k, \hat{\Pi}_k} \parallel \mathcal{N}_{\alpha_k^*, \Pi_k^*}) \\ &+ \left. \text{KL}(\mathcal{N}_{\hat{c}_{klm}, \hat{\Xi}_{klm}} \parallel \mathcal{N}_{c_{klm}^*, \Xi_{klm}^*}) \right] \quad (26) \end{aligned}$$

C. Learning with AUSE

In order to assess the accuracy of uncertainty estimation, a quantitative approach involves examining its correlation with the true error map using the Sparsification Curve [24]. First, the predicted values are sorted based on decreasing predicted uncertainty and then progressively removed, starting from those with high predicted uncertainty. By keeping track of a quality metric applied to the remaining values, the Sparsification Curve is generated. The area under this curve is called Area Under Sparsification Curve (AUSC). The Area Under the Sparsification Error (AUSE) metric is defined as the difference between the AUSC of the method and the AUSC

of the oracle, i.e., the curve obtained by sorting the predicted values according to the true error.

The AUSE can be formally described as in the following. We consider a statistical model that predicts a value \hat{y}_i along a predicted uncertainty $\hat{\sigma}_i$. Let $e_i = E(\hat{y}_i, y_i)$ be the error committed by the model, where y_i is the ground truth and $E(\cdot, \cdot)$ is an error function. Consider a collection of predictions of the model, with the associated ground truth $\{\hat{y}_i, \hat{\sigma}_i, y_i\}_{i=1}^N$. Also consider two permutations of $\{1, \dots, N\}$, namely P and Q , such that $\hat{\sigma}_{P(i)} \leq \hat{\sigma}_{P(j)}$ and $e_{Q(i)} \leq e_{Q(j)}, \forall j \geq i$ respectively. These two permutations sort the aforementioned collection according to the predicted uncertainty and the error, respectively. Given a permutation R , the Sparsification Curve is:

$$\mathcal{S}_R(n) = \frac{1}{n} \sum_{i=1}^n e_{R(i)} \quad \forall n \in \{1, \dots, N\}$$

The Area Under the Sparsification Curve (AUSC) is:

$$\text{AUSC}(R) = \sum_{n=1}^N \mathcal{S}_R(n)$$

Finally, the AUSE is defined as the difference of $\text{AUSC}(P)$ and $\text{AUSC}(Q)$:

$$\text{AUSE} = \sum_{n=1}^N \left(\frac{1}{n} \sum_{i=1}^n e_{P(i)} - e_{Q(i)} \right)$$

If the uncertainty prediction was random, the percolation process would also be random, resulting in a sparsification flat curve and a high AUSE. Otherwise, if the uncertainty prediction was positively correlated with the prediction error, improvements in the tracked quality metric would be observed. As the GS technique has significantly lower memory requirements compared to the NeRF used in previous works, it is capable of sampling the whole view multiple times in a single forward pass. This enables us to directly compute the AUSE metric applied to all the pixels in a view, taking the standard deviation of the samples of each pixel as the uncertainty map.

As we will show in the experiments section, using only the standard VI framework is suboptimal. Indeed, GS inherently has low memory requirements and high rendering efficiency, leading to underutilized GPU memory and computational resources, even when using a standard VI framework. Hence, in this work, we propose to augment the VI loss of Eq. (10) with the AUSE metric. Introducing the AUSE loss enhances the framework by providing an additional optimization signal that fully leverages these available resources. Although it introduces computational overhead, the added complexity is outweighed by the benefits of improved uncertainty estimation accuracy and overall performance.

D. End-to-end SGS Training

We now have all the necessary ingredients to define the overall loss used to train the proposed SGS method. In particular, the overall loss is the following combination:

$$\mathcal{L} = \mathcal{L}_{\text{rec}} + \lambda_{\text{SSIM}} \mathcal{L}_{\text{ssim}} + \lambda_{\text{KL}} \mathcal{L}_{\text{KL}} + \lambda_{\text{AUSE}} \mathcal{L}_{\text{AUSE}} \quad (27)$$

where \mathcal{L}_{rec} is defined as in (4) using the ℓ_1 norm, $\mathcal{L}_{\text{ssim}}$ is the SSIM perceptual loss, \mathcal{L}_{KL} is the Kullback–Leibler divergence with the prior from Eq. (26), and $\mathcal{L}_{\text{AUSE}}$ is the loss induced by the AUSE RMSE metric. The ℓ_1 , SSIM and AUSE losses augment the conventional KL loss in order to more explicitly enforce the training tradeoff between distortion, perceptual quality and uncertainty estimation.

Finally, one more aspect in which SGS training differs from previous works is the approach to learning the distribution of the priors. In previous works, stochasticity was introduced in the weights of neural networks, which are typically randomly initialized [25]. However, in GS, the parameters have a more direct physical meaning in the 3D space. For example, minimizing the KL-divergence in GS would tend to fix the center of a Gaussian kernel fixed in a randomly initialized position in 3D space. Instead, we tackle this convergence issue, by taking inspiration from Empirical Bayes. We thus first learn an informative prior with some iterations of classic GS, which serves as an initialization before switching to the SGS formulation.

IV. EXPERIMENTAL RESULTS

A. Experimental Setting

This work addresses the task of synthesizing novel views and associated uncertainty maps from multiple posed views of a static scene. These views consist of RGB photos captured from various camera positions and orientations. The spatial positions of the cameras are referred to as extrinsic parameters, while the intrinsic parameters include the camera’s focal length and central point. Similar to prior research, we assume that both extrinsic and intrinsic camera parameters are known for all input views. The objective is to perform standard novel view synthesis using the provided views while also generating an uncertainty map for each synthesized view. Consequently, for every pixel in the novel view, both its color and uncertainty are predicted. It is crucial that the predicted uncertainty correlates with the true error.

We remark that this is the first paper addressing uncertainty estimation for the GS setting, while current literature focuses on NeRF models. This makes direct comparisons challenging, as both absolute image quality as well as metrics for measuring the effectiveness of uncertainty estimation need to be reported. In particular, we remark that the AUSE metric benchmarks uncertainty estimation against the oracle method, so it is relatively insensitive to the absolute image quality generated by each method.

We compare our proposed method with the current literature on NeRF: the state-of-the-art CF-NeRF [14], the pioneering work of S-NeRF [4], and also with NeRF-W [22], Deep-Ensembles (D.E.) [26] and MC-Dropout [21], as done in previous works. We remark that some methods [4], [14] use extra information in the form of depth maps while we do not, resulting in a setting that is slightly unfair towards SGS. Nevertheless, we report improvements over such methods. We stress that the D.E. method is defined simply as training from scratch 5 vanilla NeRF networks [2], testing them and averaging the results. This method may reach high levels of

uncertainty prediction accuracy at the cost of being 5 time slower of the “backbone” method.

Moreover, we report also the results of the vanilla GS algorithm, even if it does not enable any uncertainty prediction.

As common practice, the experiments are conducted on the LLFF dataset from the original NeRF paper [2]. Moreover, while works in previous related literature [4], [14] report results only for the LLFF dataset, we also report results on the Blender and Mip-NeRF360 datasets in order to validate the robustness and generalization of SGS and to ease the comparison of SGS with other methods. In fact these datasets are a de-facto standard benchmark for general-purpose novel view synthesis methods. Note that the Mip-NeRF360 dataset is much more challenging than the other two datasets, and all the methods based on the original NeRF are not able to correctly represent the scenes therein. This highlights also the added benefit of using GS as the backbone for SGS, as the vanilla GS can correctly reproduce the Mip-NeRF360 dataset.

All the experiments are performed at a fraction of the original resolution, so that all synthesized images and uncertainty maps are composed by approximately 400×400 pixels.

The hyperparameters in the final loss function are: $\lambda_{\text{KL}} = 10^{-3}$, $\lambda_{\text{AUSE}} = 5$, $\lambda_{\text{ssim}} = 0.2$. We use the default GS hyperparameters for each dataset, except for fixing the highest spherical harmonics degree to 1.

At iteration 16000, the current learned GS is fixed and taken as the prior, and the Bayesian regime is introduced. All the prior covariance matrices are initialized as $10^{-2}\mathbb{I}$, where \mathbb{I} is the identity matrix of the correct dimension. The learning rate for all the posterior learnable parameters is set to 10^{-4} . Then we continue the training until iteration 30000. During both training and testing, Monte Carlo sampling from the posterior is performed for $S = 8$ times.

Code for our SGS implementation is publicly available at the following address: <https://github.com/mezzelfo/stochastic-gaussian-splatting>.

B. Main Results

Tables I, II and III report our results on the LLFF, Blender and Mip-NeRF360 dataset respectively, evaluating the quality of the rendered images, as well as the reliability of the associated uncertainty maps. The rendered images are quantitatively evaluated with three metrics: PSNR as a distortion metric and LPIPS and SSIM as perceptual metrics. It can be noticed that our method improves by a large margin all these metrics, hence our method has a much lower test prediction bias with respect to previous work. The uncertainty maps are quantitatively evaluated with two metrics: AUSE RMSE and AUSE MAE. Both metrics are obtained as the area under the sparsification curve, but considering as the error metric the Root Mean Square Error and the Mean Absolute Error, respectively. As shown in Tables I, II and III, our method improves the AUSE RMSE metric, while keeping an AUSE MAE metric comparable with the state of the art, which however also exploits depth information. Fig. 2 shows a qualitative result for a novel view generated by SGS and CF-NeRF together with the predicted uncertainty maps. We can notice that SGS

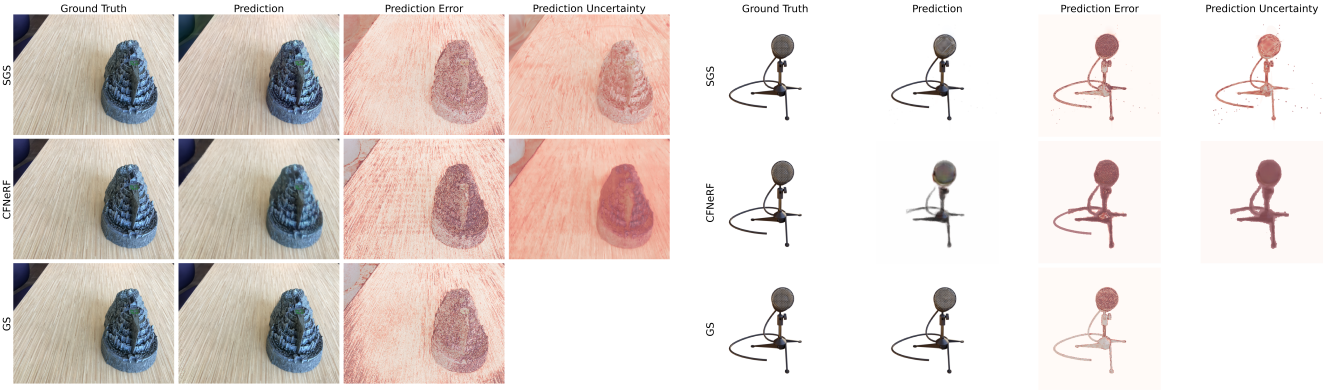


Fig. 2. Two qualitative examples of our method SGS with CF-NeRF [14]. The last column is a visualization of the predicted uncertainty map.

TABLE I
RESULTS ON THE LLFF DATASET. * DENOTES EXTRA DEPTH INFORMATION.

	Rendering Metrics			Uncertainty Metrics	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	AUSE RMSE \downarrow	AUSE MAE \downarrow
GS. (Kerbl et al. 2023)	26.85	0.880	0.055	N.A.	N.A.
D.E. (Lakshminarayanan, Pritzel, and Blundell 2017)	<u>22.32</u>	0.788	0.236	0.0254	0.0122
Drop. (Gal and Ghahramani 2016)	21.90	0.758	0.248	0.0316	0.0162
NeRF-W (Martin-Brualla et al. 2021)	20.19	0.706	0.291	0.0268	0.0113
S-NeRF* (Shen et al. 2021)	20.27	0.738	0.229	0.0248	0.0101
CF-NeRF* (Shen et al. 2022)	21.96	<u>0.790</u>	<u>0.201</u>	<u>0.0177</u>	0.0078
SGS (Ours)	24.20	0.842	0.121	0.0147	<u>0.0092</u>

TABLE II
RESULTS ON THE BLENDER DATASET. THE BASELINE S-NeRF IS MISSING AS NO CODE IS AVAILABLE.

	Rendering Metrics			Uncertainty Metrics	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	AUSE RMSE \downarrow	AUSE MAE \downarrow
GS. (Kerbl et al. 2023)	34.50	0.986	0.011	N.A.	N.A.
D.E. (Lakshminarayanan, Pritzel, and Blundell 2017)	23.72	0.909	<u>0.070</u>	0.0029	0.0007
Drop. (Gal and Ghahramani 2016)	23.59	0.882	0.153	0.0564	0.0250
NeRF-W (Martin-Brualla et al. 2021)	<u>25.90</u>	<u>0.912</u>	0.092	0.0345	0.0122
CF-NeRF (Shen et al. 2022)	23.26	0.860	0.160	0.0085	0.0031
SGS (Ours)	31.13	0.965	0.021	<u>0.0069</u>	<u>0.0027</u>

is capable of producing a sharp view as well as prediction uncertainty which correlates well with the true rendering error.

C. AUSE Loss Ablation

Table IV compares our SGS method with an ablated version, where the proposed AUSE loss term $\mathcal{L}_{\text{AUSE}}$ in equation (27) is removed in order to verify its effectiveness. As reported in the table, the removal of this loss term improves the photometric reconstruction (measured by the three quality metrics: PSNR, SSIM, and LPIPS), while deteriorating the model’s ability to predict accurate uncertainty maps. Hence, this ablation study proves that one of our key contribution, that is to incorporate the AUSE loss term, improves the quality of the predicted uncertainty maps. Moreover, the hyperparameter λ_{AUSE} provides a natural way to control the impact of this loss term, so that an application-specific trade-off between reconstruction quality and accurate uncertainty prediction can be found for downstream tasks.

D. Components Ablation

SGS imposes a prior distribution on the position, color, and opacity of the Gaussian primitives. Table V compares the full proposed SGS method with versions of it where only one attribute at a time is modeled as a stochastic variable, showing that the best uncertainty maps are obtained when all the parameters considered are stochastic.

E. Runtime complexity

We analyzed the runtime performance of our SGS method to compare it with baselines. From a theoretical perspective, the proposed method is S times slower than conventional GS, where S is the number of Monte Carlo samples used in the variational inference process. Hence, there is a trade-off: a higher S leads to better uncertainty estimation accuracy, but also to a linearly higher computational overhead. Note that S is a dynamic parameter, i.e., it can be changed even after the training phase is completed, depending on the application needs. In the limit case $S = 1$, one recovers the classic GS, losing the ability to estimate uncertainty.

TABLE III
RESULTS ON THE MIP-NeRF360 DATASET. THE BASELINE S-NeRF IS MISSING AS NO CODE IS AVAILABLE.

	Rendering Metrics			Uncertainty Metrics	
	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	AUSE RMSE \downarrow	AUSE MAE \downarrow
GS. (Kerbl et al. 2023)	30.66	0.923	0.033	N.A.	N.A.
D.E. (Lakshminarayanan, Pritzel, and Blundell 2017)	20.72	0.487	0.533	0.0379	0.0231
Drop. (Gal and Ghahramani 2016)	19.21	0.402	0.628	0.0536	0.0338
NeRF-W (Martin-Brualla et al. 2021)	19.40	0.457	0.503	0.0649	0.0421
CF-NeRF (Shen et al. 2022)	19.51	0.422	0.633	0.0292	0.0171
SGS (Ours)	26.30	0.873	0.091	0.0139	0.0084

TABLE IV
AUSE LOSS TERM ABLATION OF LLFF DATASET.

	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow	AUSE RMSE \downarrow
No AUSE Loss	26.65	0.869	0.082	0.0291
SGS (Ours)	24.20	0.842	0.121	0.0147

TABLE V
COMPONENTS ABLATION ON *kitchen* SCENE OF THE MIP-NeRF360 DATASET.

	PSNR (dB) \uparrow	SSIM \uparrow	LPIPS \downarrow	AUSE RMSE \downarrow
Position	27.93	0.950	0.043	0.0105
Color	28.94	0.947	0.040	0.0121
Opacity	29.05	0.954	0.043	0.0112
SGS (Ours)	28.61	0.949	0.034	0.0090

We experimentally compared the runtime of the proposed SGS method with CF-NeRF. In order to have a comparison as fair as possible, we run the original evaluation code for both methods on the test split of the *hotdog* scene of Blender dataset ($N = 200$ views to be rendered) but all the disk accesses are disabled, GPU-CPU transfers are reduced as much as possible, and the same Nvidia Quadro P6000 GPU is used. Moreover, we compare both methods at an equal value of the parameter $S = 8$. We report the results measures as seconds per view per sample ($\frac{sec}{NM}$). SGS achieves $\approx 0.002 \frac{sec}{NM}$, while CF-NeRF is substantially slower at approximately $4.7 \frac{sec}{NM}$.

F. Effect of Sample Size on AUSE Performance

In this section, we evaluate the quality of uncertainty estimation in SGS as a function of the number of Monte Carlo samples, S , utilized in the variational inference process. We selected the *kitchen* scene from the Mip-NeRF360 dataset as our test case. For each value of S in $\{2, 4, 8, 16, 32\}$, we ran the model and computed the AUSE metric across eight trials, averaging the results for each value of S . The average AUSE values are plotted in Figure 3. The results reveal an almost linear decrease in AUSE as S doubles, indicating that increasing the number of Monte Carlo samples improves the quality of uncertainty estimation. However, this improvement comes with a tradeoff: higher values of S lead to more accurate uncertainty estimates but also result in longer runtimes due to the increased computational load required for additional sampling. Therefore, SGS allows for a balance between achieving better uncertainty estimation and managing computational efficiency, depending on the choice of S .




Fig. 3. SGS trend between the number of samples S used in the Monte Carlo estimation and the resulting AUSE, in the *kitchen* scene of the Mip-NeRF360 dataset.

G. Limitations and Future Work

While Stochastic Gaussian Splatting (SGS) represents a significant advancement in uncertainty estimation for novel view synthesis, our approach inherits the limitations of the vanilla GS framework it is based on. Firstly, SGS, like GS, is designed for static scenes. This limitation means that the framework may not perform optimally in dynamic or evolving environments where scene content changes over time. Additionally, as GS (but also as all the previous related works and baselines), our method requires re-training for each new scene, reducing its practicality for scenarios involving frequent scene changes.

Moreover, our method addresses only the aleatoric uncertainty and not the epistemic one because we are modeling uncertainty in the 3D points that are covered by the Gaussian primitives. For example, SGS cannot predict any uncertainty for a novel view that is completely outside the scene extent, as no Gaussians are placed there. This kind of high uncertainty is epistemic as can be reduced by obtaining training views that effectively enlarge the scene extent.

An interesting direction for future work is how to use the uncertainty map predicted by SGS to aid and inform the reconstruction process. Even if the uncertainty is represented as a variance on the image plane, making it non-trivial to

directly map this back to the 3D space for adjusting the Gaussian parameters, it can potentially be used as a more principled densification criterion for vanilla GS. Using uncertainty information, future research could focus on optimizing adaptive density control methods (densification/pruning strategies) in GS to improve reconstruction quality and computational efficiency.

V. CONCLUSIONS

In this paper, we proposed a novel approach for uncertainty estimation in GS-based novel-view synthesis tasks. Leveraging the efficiency and real-time capabilities of GS, we introduced a stochastic extension to the traditional deterministic GS framework. Our method incorporates uncertainty prediction through a Bayesian framework, specifically using VI to learn the parameters of the GS radiance field. The training is further augmented with direct optimization of the AUSE metric to control the tradeoff between reconstruction quality and accuracy of uncertainty estimation.

Experimental results on three different datasets demonstrated the effectiveness of our approach. We outperformed state-of-the-art methods in terms of rendering quality metrics, while also improving upon uncertainty estimation metrics. Notably, our work advances the state of the art by being the first to introduce uncertainty estimation for GS-based novel-view synthesis tasks.

REFERENCES

- [1] U. S., “The interpretation of structure from motion,” *Proceedings of the Royal Society of London. Series B. Biological Sciences*, vol. 203, no. 1153, pp. 405–426, 1979.
- [2] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: Representing scenes as neural radiance fields for view synthesis,” *Communications of the ACM*, vol. 65, no. 1, pp. 99–106, 2021.
- [3] B. Kerbl, G. Kopanas, T. Leimkühler, and G. Drettakis, “3d gaussian splatting for real-time radiance field rendering,” *ACM Transactions on Graphics (ToG)*, vol. 42, no. 4, pp. 1–14, 2023.
- [4] J. Shen, A. Ruiz, A. Agudo, and F. Moreno-Noguer, “Stochastic neural radiance fields: Quantifying uncertainty in implicit 3d representations,” in *2021 International Conference on 3D Vision (3DV)*. IEEE, 2021, pp. 972–981.
- [5] Z. Chen and H. Zhang, “Learning implicit fields for generative shape modeling,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019, pp. 5939–5948.
- [6] N. Max, “Optical models for direct volume rendering,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 1, no. 2, pp. 99–108, 1995.
- [7] S. Chandrasekhar, *Radiative transfer*. Courier Corporation, 2013.
- [8] M. Zwicker, H. Pfister, J. Van Baar, and M. Gross, “Ewa volume splatting,” in *Visualization, 2001. VIS 01. Proceedings*. IEEE, 2001, pp. 29–538.
- [9] A. Kendall and Y. Gal, “What uncertainties do we need in bayesian deep learning for computer vision?” *Advances in neural information processing systems*, vol. 30, 2017.
- [10] A. G. Kendall, “Geometry and uncertainty in deep learning for computer vision,” Ph.D. dissertation.
- [11] Y. Lu, X. Ma, L. Yang, T. Zhang, Y. Liu, Q. Chu, J. Yan, and W. Ouyang, “Geometry uncertainty projection network for monocular 3d object detection,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021, pp. 3111–3121.
- [12] Y. Gal *et al.*, “Uncertainty in deep learning.”
- [13] L. V. Jospin, H. Laga, F. Boussaid, W. Buntine, and M. Bennamoun, “Hands-on bayesian neural networks—a tutorial for deep learning users,” *IEEE Computational Intelligence Magazine*, vol. 17, no. 2, pp. 29–48, 2022.
- [14] J. Shen, A. Agudo, F. Moreno-Noguer, and A. Ruiz, “Conditional-flow nerf: Accurate 3d modelling with reliable uncertainty quantification,” in *European Conference on Computer Vision*. Springer, 2022, pp. 540–557.
- [15] G. Yang, X. Huang, Z. Hao, M.-Y. Liu, S. Belongie, and B. Hariharan, “Pointflow: 3d point cloud generation with continuous normalizing flows,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2019, pp. 4541–4550.
- [16] S. Wei, J. Zhang, Y. Wang, F. Xiang, H. Su, and H. Wang, “Fg-nerf: Flow-gan based probabilistic neural radiance field for independence-assumption-free uncertainty estimation,” *arXiv preprint arXiv:2309.16364*, 2023.
- [17] A. Grover, M. Dhar, and S. Ermon, “Flow-gan: Combining maximum likelihood and adversarial learning in generative models,” *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, 04 2018.
- [18] L. Goli, C. Reading, S. Sellán, A. Jacobson, and A. Tagliasacchi, “Bayes’ rays: Uncertainty quantification for neural radiance fields,” *arXiv preprint arXiv:2309.03185*, 2023.
- [19] J. Shen, R. Ren, A. Ruiz, and F. Moreno-Noguer, “Estimating 3d uncertainty field: Quantifying uncertainty for neural radiance fields,” *arXiv preprint arXiv:2311.01815*, 2023.
- [20] N. Sünderhauf, J. Abou-Chakra, and D. Miller, “Density-aware nerf ensembles: Quantifying predictive uncertainty in neural radiance fields,” in *2023 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2023, pp. 9370–9376.
- [21] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: Representing model uncertainty in deep learning,” in *international conference on machine learning*. PMLR, 2016, pp. 1050–1059.
- [22] R. Martin-Brualla, N. Radwan, M. S. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth, “Nerf in the wild: Neural radiance fields for unconstrained photo collections,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021, pp. 7210–7219.
- [23] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [24] F. K. Gustafsson, M. Danelljan, and T. B. Schon, “Evaluating scalable bayesian deep learning methods for robust computer vision,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, 2020, pp. 318–319.
- [25] M. V. Narkhede, P. P. Bartakke, and M. S. Sutaone, “A review on weight initialization strategies for neural networks,” *Artificial intelligence review*, vol. 55, no. 1, pp. 291–322, 2022.
- [26] B. Lakshminarayanan, A. Pritzel, and C. Blundell, “Simple and scalable predictive uncertainty estimation using deep ensembles,” *Advances in neural information processing systems*, vol. 30, 2017.