

Predictive Photometric Uncertainty in Gaussian Splatting for Novel View Synthesis

Chamuditha Jayanga Galappaththige^{1,2*}, Thomas Gottwald^{3*}, Peter Stehr^{3*},
Edgar Heinert⁴, Niko Suenderhauf^{1,2}, Dimity Miller^{1,2}, and Matthias
Rottmann⁴

¹ QUT Centre for Robotics, Australia

² ARIAM Hub, Australia

³ University of Wuppertal, Germany

⁴ University of Osnabrück, Germany

Abstract. Recent advances in 3D Gaussian Splatting have enabled impressive photorealistic novel view synthesis. However, to transition from a pure rendering engine to a reliable spatial map for autonomous agents and safety-critical applications, knowing where the representation is uncertain is as important as the rendering fidelity itself. We bridge this critical gap by introducing a lightweight, plug-and-play framework for pixel-wise, view-dependent predictive uncertainty estimation. Our post-hoc method formulates uncertainty as a Bayesian-regularized linear least-squares optimization over reconstruction residuals. This architecture-agnostic approach extracts a per-primitive uncertainty channel without modifying the underlying scene representation or degrading baseline visual fidelity. Crucially, we demonstrate that providing this actionable reliability signal successfully translates 3D Gaussian splatting into a trustworthy spatial map, further improving state-of-the-art performance across three critical downstream perception tasks: active view selection, pose-agnostic scene change detection, and pose-agnostic anomaly detection. Code is available at github.io/3DGS-Uncertainty.

Keywords: Gaussian Splatting · Novel View Synthesis · Uncertainty Estimation

1 Introduction

Radiance fields [23, 40] have evolved from novel view synthesis (NVS) engines into foundational spatial maps for autonomous agents [11, 39, 63]. However, constructing a radiance field from 2D images is an inherently ill-posed inverse problem [55]. To function reliably in real-world deployments, where severe occlusions, unobserved regions, and geometric ambiguities are inevitable, these systems must be capable of rigorously quantifying their predictive uncertainty [53]. 3D Gaussian Splatting (3DGS) [23] has rapidly emerged as the leading representation for these tasks, combining the expressiveness of volumetric rendering with the

* Equally contributed.




Fig. 1: Per-pixel UE for NVS. Our post-hoc method generates a view-dependent uncertainty map that closely mirrors regions of error within the RGB render.

efficiency of rasterization. While recent advancements have drastically improved 3DGS visual fidelity [25, 61], geometric consistency [9, 24], and efficiency [16, 38], equipping these models with robust, system-level uncertainty estimation (UE) remains a critical, underexplored challenge.

Existing UE methods for 3DGS fall into two complementary families, depending on whether they quantify uncertainty in the *learned representation*, *i.e.*, the Gaussian parameters, or in the *rendered radiance field*, *i.e.*, the rendered pixels. Most prior work targets the learned representation. Stochastic formulations [1, 33, 37, 52] model distributions over Gaussian parameters, but rely on sampling-based optimization or complex inference that introduces prohibitive latency, requires architectural modifications degrading baseline visual fidelity, and lacks modularity with the rapidly expanding ecosystem of 3DGS variants. Post-hoc alternatives instead estimate epistemic uncertainty in parameter space via Hessian approximations [21, 58]. However, as shown in our experiments (Sec. 4.1), these parameter-centric methods capture view-dependent uncertainty poorly, hindering their use in downstream perception tasks.

In contrast, we model uncertainty directly in the rendered radiance field [15], introducing an efficient, plug-and-play system that estimates pixel-wise *predictive* uncertainty for NVS. Our approach generates view-dependent uncertainty maps alongside RGB images, with areas of high uncertainty reflecting regions of low rendering fidelity in the novel view (see Fig. 1). This yields a single, actionable reliability signal that is particularly valuable for NVS, where downstream tasks consume and act directly upon the rendered pixels. We demonstrate this utility across three perception tasks, leveraging our uncertainty estimates to improve performance for next-best-view planning [43], pose-agnostic scene change detection [13], and pose-agnostic anomaly detection [31].

Our approach exploits a simple but powerful observation: regions that exhibit high reconstruction residual often correspond to areas where 3DGS fails to accurately represent the scene (see Fig. 2). We leverage this signal to guide the learning of a dedicated uncertainty channel within each 3DGS primitive, correlating representation error with high uncertainty. Yet reconstruction residuals alone are not a reliable proxy for uncertainty – in sparse-view settings, the highly parameterized nature of 3DGS enables the model to overfit to the training views, driving residuals toward zero (thus signaling zero predictive uncertainty). To mitigate this, we introduce a Bayesian prior within our linear least-squares formulation of uncertainty estimation. This prior encourages uncertainty esti-

mates to revert toward a maximal uncertainty level when insufficient multi-view evidence is available. This process is entirely post-hoc: it requires zero modifications to the base architecture or its optimization pipeline, ensuring preservation of the original rendering fidelity and compatibility with any 3DGS variant.

In summary, we make the following claims:

- We introduce a lightweight, plug-and-play system for view-dependent UE for novel view synthesis in 3DGS. We achieve state-of-the-art performance for UE while strictly preserving rendering fidelity (see Sec. 4.1).
- We formulate this estimation as a (Bayesian-inspired) regularized linear least-squares problem, which ensures robust uncertainty predictions across both dense and sparse capture settings (see Sec. 4.2).
- Beyond state-of-the-art UE in 3DGS, our uncertainty maps can be practically utilized in downstream tasks to improve state-of-the-art performance in active view selection (Sec. 4.3), pose-agnostic scene change detection (Sec. 4.4), and pose-agnostic anomaly detection (Sec. 4.5).

2 Related Work

2.1 Uncertainty Estimation in Gaussian Splatting

UE in NeRFs [40] has been extensively explored through stochastic sampling [50], ensemble models [54], conditional normalizing flows [51], and provenance modeling [41]. Bayes’ Rays [14] interprets uncertainty as allowable volumetric variation to identify under-constrained regions. While foundational, these designs remain tightly coupled to neural volumetric representations, making direct application to 3DGS [23] non-trivial due to its explicit, non-neural parameterization.

Existing UE techniques for 3DGS primarily rely on stochastic formulations. Continuous Semantic Splatting [59] learns semantic distributions over Gaussian primitives, interpreting their variance as a measure of semantic uncertainty. Several stochastic extensions [1, 33, 37] derive pixel-level uncertainty from multi-sample variance. Stochastic-GS [1] learns probabilistic distributions over Gaussian parameters, while Variational-3DGS [33] leverages multi-scale Gaussian hierarchies to construct a diversified sampling space. Lyu *et al.* [37] propose a



Fig. 2: Visual example of reconstruction residual on training views. (Left) ground truth, (Center) rendering, (Right) residual map. The residual maps capture errors caused by both imperfect geometry (*e.g.*, under-reconstructed regions in vegetation) and the limited capacity of view-dependent appearance (*e.g.*, reflections on the table).

low-dimensional manifold constraint to reduce the heavy sampling cost inherent in stochastic rendering. Despite their advances, these stochastic methods are computationally expensive and demand substantial architectural or optimization modifications. This lacks plug-and-play applicability across 3DGS variants and often compromises rendering fidelity.

Current post-hoc approaches for uncertainty estimation in 3DGS predominantly focus on the parameter space of the model. FisherRF [21] uses Fisher information to quantify confidence in the learned representation. POp-GS [58] extends this via optimal experimental design while also modeling inter-parameter correlations. Similar to PRIMU [15], our approach operates directly in image space. However, rather than relying on hand-crafted features and a regressor trained on hold-out views, we introduce a lightweight linear least-squares formulation. This allows us to directly learn primitive-level, view-dependent uncertainty that efficiently aggregates into pixel-wise maps from arbitrary viewpoints, entirely bypassing the need for a regressor.

A parallel line of work assigns error-based scores to primitives to guide pruning and densification. PUP 3D-GS [17] and Speedy-Splat [16] compute pruning scores based on L^2 loss sensitivity. Rota Bulò *et al.* [44] derive a scalar densification score directly from the rendering error. However, their per-pixel formulation ignores interactions between overlapping Gaussians, assigning elevated scalar scores to all contributing primitives even if only one is responsible for an artifact. While this heuristic suits densification, where over-densification is corrected in subsequent optimization, it falsely penalizes reliable primitives in frozen models. By analyzing all views simultaneously through our least-squares formulation, we disentangle these interactions to consistently identify correct Gaussians, even when they contribute to localized artifacts.

2.2 Active View Selection in Gaussian Splatting

Active view selection (AVS) seeks the optimal next viewpoint to maximize reconstruction quality under a constrained acquisition budget [21, 29, 37, 43, 54]. Standard approaches select a candidate view from a predefined pool to either minimize the overall uncertainty [37, 54] or maximize the expected information gain [21, 43]. Within 3DGS specifically, FisherRF [21] optimizes information gain between candidate and training views, whereas Lyu *et al.* [37] reduce reconstruction uncertainty via low-dimensional manifold sampling. In contrast, we adopt a deterministic strategy that selects the next view to maximally reduce the per-primitive uncertainty estimated from our least-squares residual formulation, achieving state-of-the-art results in AVS.

2.3 Pose-Agnostic Scene Change Detection

Scene change detection (SCD) has been widely studied as a bi-temporal comparison problem, where changes are identified between two aligned images [2, 7, 10, 32, 34, 46–48, 56]. Such methods rely on strict viewpoint alignment and extensive supervision, limiting scalability in real-world applications such as robotic

inspection [13]. Recent studies [3, 8, 22, 26] explore zero-shot SCD using large visual foundation models [27, 42]; however, they still assume paired captures with minimal viewpoint disparity.

The emergence of high-fidelity 3D scene representations [23, 60] has enabled *pose-agnostic* SCD, where pre-change scenes are reconstructed and rendered into post-change viewpoints for comparison [12, 13, 20, 36]. However, they remain highly sensitive to reconstruction inaccuracies and view-dependent artifacts, which are frequently misclassified as true scene changes. We address this critical flaw by leveraging our predicted per-pixel uncertainty to suppress unreliable regions, effectively filtering out rendering-induced false positives. To the best of our knowledge, this is the first work to apply uncertainty estimation to pose-agnostic SCD. We further demonstrate that incorporating our uncertainty maps improves state-of-the-art pose-agnostic SCD methods [13], as our UE accurately captures erroneous regions in novel-view renderings.

2.4 Pose-Agnostic Anomaly Detection

Traditional visual anomaly detection (AD) methods [5, 45] excel at localizing defects but typically require strictly controlled, fixed camera viewpoints. To relax this constraint, recent approaches leverage 3D scene representations [23, 40] to enable pose-agnostic AD [31, 35, 62], allowing query captures from arbitrary viewpoints to be directly compared against novel-view renderings of a reconstructed reference model of an object. However, much like in pose-agnostic SCD, this direct comparison introduces a critical vulnerability: systemic rendering artifacts and geometric ambiguities are frequently misclassified as physical defects. By explicitly quantifying the reliability of the reference model, our view-dependent uncertainty maps provide a principled mechanism to mask these inherently ambiguous regions. This effectively suppresses rendering-induced false positives, significantly enhancing the robustness of the AD pipeline.

3 Method

An overview of our approach is illustrated in Fig. 3. The goal of this work is to predict *pixel-wise, view-dependent uncertainty maps* for novel views while preserving the rendering fidelity of 3DGS. We begin by revisiting the fundamentals of volumetric rendering in 3DGS (Sec. 3.1) and reformulate the rendering process to predict uncertainty rather than color (Sec. 3.2). Subsequently, we describe our learning formulation (Sec. 3.3) and introduce a Bayesian-inspired regularization (Sec. 3.4) that improves uncertainty consistency in highly novel viewpoints.

3.1 3D Gaussian Splatting

In 3DGS [23], a radiance field is represented using a set of 3D Gaussian primitives, each defined by a mean μ , covariance matrix Σ , scalar opacity factor a , and view-dependent color $c(d)$ modeled via spherical harmonics (SH). Given a




Fig. 3: Overview of our approach: We learn a primitive-level uncertainty channel u_k on a *well-trained* 3DGS, leveraging a (Bayesian-inspired) regularized linear least-squares formulation of training residuals. The learned view-dependent uncertainty can be rasterized to any given (novel) viewpoint P_j . Our method is post-hoc and therefore can be seamlessly integrated with any 3DGS variant without loss of fidelity.

camera with position, viewing direction d , and intrinsic parameters, a view can be rendered by projecting the 3D Gaussians onto the image plane and α -blending their contributions to obtain the final pixel color. This projection, also referred to as splatting [64], employs a local affine approximation of the projective transformation, such that 3D Gaussians are mapped to 2D Gaussians on the image plane. The opacity of a Gaussian at pixel x is given by $\alpha(x) = a \cdot \Phi(x, \mu^{2D}, \Sigma^{2D})$, where μ^{2D} and Σ^{2D} denote the mean and covariance of the projected 2D Gaussian, and $\Phi(x, \mu^{2D}, \Sigma^{2D})$ is its density at pixel x . The pixel-wise color is computed as a weighted sum over the primitives’ color values along the viewing ray:

$$C(x) = \sum_{k=1}^K c_k(d) \alpha_k(x) T_k(x), \quad (1)$$

where $c_k(d)$ and $\alpha_k(x)$ denote the color and opacity of the k -th Gaussian, and $T_k(x) = \prod_{l=1}^{k-1} (1 - \alpha_l(x))$ is the transmittance up to the $(k-1)$ -th Gaussian.

3.2 Per-primitive Uncertainty

We estimate pixel-wise, view-dependent uncertainty by learning per-primitive uncertainty channels from training-view reconstruction residuals. Since uncertainty may vary with the viewing direction, it is modeled using SH, analogously to color. For any view, a pixel-wise uncertainty map can be rendered in the same manner (and as fast) as an RGB image:

$$U(x) = \sum_{k=1}^K u_k(d) \alpha_k(x) T_k(x), \quad (2)$$

where $u_k(d)$ denotes the uncertainty of Gaussian k from direction d .

We hypothesize that (in dense capture settings**) reconstruction residuals serve as a reliable proxy for *predictive* uncertainty: just as 3DGS generalizes color

** As in typical 3DGS datasets, *e.g.* Mip-NeRF360.

to novel views, a well-fit residual channel approximates the expected photometric error there. This proxy aggregates aleatoric, epistemic, and optimization-related effects into a single, actionable signal.

While any discrepancy metric L could serve as the target residual, we explicitly employ the same combination of L^1 loss and DSSIM used to train the base 3DGS model [23]. To learn the primitive-level uncertainty channel u_k , we freeze all other representation parameters. This ensures the original reconstruction fidelity is perfectly preserved while the new uncertainty parameters are optimized to best explain the observed training errors.

We formulate the post-hoc estimation of the uncertainty channel as a least-squares problem. Given a well-reconstructed scene with K primitives $G_1, \dots, G_K$, we enumerate all pixels across all training views as $x_1, \dots, x_M$, each corresponding to a viewing ray determined by its camera pose and pixel location. For simplicity, we initially assume u_k to be view-independent and later describe how directional dependence can be incorporated. For each pixel x_j , depth-sorting the primitives front-to-back along its viewing ray defines a permutation σ_j , where $\sigma_j(k)$ denotes the index of the k -th primitive along the ray. In this setting, Eq. (2) can be expressed as

$$U(x_j) = \sum_{k=1}^K u_k \alpha_k(x_j) \hat{T}_k(x_j), \quad (3)$$

where $\hat{T}_k(x_j) = \prod_{l=1}^{\sigma_j(k)-1} (1 - \alpha_{\sigma_j(l)}(x_j))$ and $u_k \in \mathbb{R}$ is the uncertainty of primitive G_k . By gathering these volumetric blending weights into a matrix $A \in \mathbb{R}^{M \times K}$ such that $A_{jk} = \alpha_k(x_j) \hat{T}_k(x_j)$, the rendered uncertainties for all pixels become the matrix-vector product Au , where $u = (u_1, \dots, u_K)^\top$. Let $y \in \mathbb{R}^M$ be the vector of corresponding pixel-wise reconstruction losses, such that $y_j = L_{x_j}$. Estimating the primitive uncertainties that best explain these residuals amounts to solving the linear least-squares problem:

$$\arg \min_u \|y - Au\|_2^2. \quad (4)$$

Incorporating view-dependency for u via SH preserves this linearity; the SH basis functions are simply absorbed into A , and u expands to contain the concatenated SH coefficients across all primitives (see App. A for the explicit formulation).

3.3 Optimization

Due to the massive scale of typical scenes, a direct solution to this problem is computationally intractable. Specifically, for the normal equations $A^T A u = A^T y$, the dimension of the matrix $A^T A$ scales with the number of primitives (often in the millions) times the 16 SH coefficients per Gaussian. Despite the sparsity of this matrix, its size renders direct solvers impractical. Instead, we approximate the solution using stochastic gradient descent (SGD), following the same optimization procedure employed for the other Gaussian parameters. The convexity of the linear least-square objective ensures efficient and stable convergence.

After reconstruction is completed, for each pixel x in the view, we first calculate the residual error L_x for all training viewpoints. We adopt 3DGS’s photometric error and keep $\lambda = 0.2$ following 3DGS [23]: $L_x = (1 - \lambda)L_x^1 + \lambda L_x^{\text{DSSIM}}$. The overall training objective is then minimizing the L^2 loss of the residual L_x and the rendered uncertainty U_x ,

$$L = \sum_x (L_x - U_x)^2. \quad (5)$$

Gradients are backpropagated exclusively to the uncertainty channels u_k , leaving the matrix A fixed, thus ensuring linearity.

3.4 Bayesian-Inspired Regularization

The least-squares formulation in Eq. (4) learns directional uncertainty only from viewing directions that are present in the training data. Consequently, directions absent from the training data provide no supervisory signal. Furthermore, in sparse-view scenarios, a trivially low rendering error on the training views often masks a poorly reconstructed underlying 3D geometry. As a result, uncertainty estimates for highly novel viewpoints can become unstable or unintuitive.

To this end, we introduce an L^2 regularization that imposes a prior over the directional uncertainty functions. In directions unsupported by training views, the learned uncertainty $u_k(d)$ is encouraged to default to a maximal uncertainty level $b \in \mathbb{R}$. This approach is grounded in the well-known interpretation of L^2 -regularized linear regression as Bayesian inference with a Gaussian prior [6] centered at the maximal uncertainty value b .

For a Gaussian G_k , the SH coefficients of its uncertainty feature define a function $u_k : S^2 \rightarrow \mathbb{R}$ on the unit sphere that encodes directional uncertainty. The difference of u_k and the constant b (which encodes maximal and isotropic uncertainty) is quantified by their L^2 distance:

$$l_k = \int_{S^2} (b - u_k(r))^2 dr. \quad (6)$$

This integral is approximated by evaluating $u_k(d)$ at discrete points on the sphere using Gauss-Legendre sampling [30]. Summing over all K Gaussian primitives yields the global regularization loss: $L_{\text{reg}} = \sum_{k=1}^K l_k$. The final training objective regularizes the base loss L from Eq. (5):

$$L' = L + \lambda_{\text{reg}} L_{\text{reg}}, \quad (7)$$

where $\lambda_{\text{reg}} \geq 0$ controls the influence of the prior. The choice of λ_{reg} and conditions under which this regularization is beneficial are discussed in Sec. 4.2.

4 Experiments

Datasets: We follow the exact evaluation protocol established in 3DGS [23] and assess our UE performance on three standard benchmarks: Mip-NeRF360 [4], Tanks & Temples [28], and Deep Blending [18].




Fig. 4: Qualitative comparison of predicted uncertainty maps on the *bonsai* and *stump* scenes of Mip-NeRF360 [4]. Our UE more closely reflects the true error when compared to the baselines. Additional qualitative examples can be found in App. F.

Table 1: Quantitative results for UE for NVS. We consistently outperform all baselines across all metrics and datasets while achieving minimal training overhead (OH).

Method	Mip-NeRF360 [4]						Tanks&Temples [28]						Deep Blending [18]					
	AUSE(↓)			Pearson(↑)			AUSE(↓)			Pearson(↑)			AUSE(↓)			Pearson(↑)		
	L^1	DSSIM		L^1	DSSIM	OH(↓) (%)	L^1	DSSIM		L^1	DSSIM	OH(↓) (%)	L^1	DSSIM		L^1	DSSIM	OH(↓) (%)
FisherRF [21]	0.708	0.606	-0.055	0.009	14.2	0.691	0.709	-0.087	-0.145	19.3	0.751	0.853	-0.116	-0.190	19.7			
Manifold [37]	0.520	0.559	0.070	-0.005	30.2	0.574	0.654	0.053	0.008	23.6	0.503	0.548	0.095	0.074	48.1			
Var3DGS [33]	0.558	0.495	0.118	0.160	>100	0.539	0.567	0.161	0.176	>100	0.588	0.671	0.106	0.006	>100			
Ours	0.328	0.214	0.369	0.547	12.8	0.299	0.233	0.427	0.571	14.0	0.376	0.356	0.243	0.244	13.0			

Baselines: We evaluate against representative state-of-the-art 3DGS UE approaches, including the post-hoc method FisherRF [21] and the variational inference-based methods of Lyu et al. [37] (Manifold) and Li et al. [33] (Var3DGS). The latter two modify the 3DGS framework by learning distributions over Gaussian parameters, which slightly reduces rendering fidelity. A detailed comparison of their rendering quality against standard 3DGS is provided in App. D.

Metrics: To evaluate how well predicted uncertainty maps correspond to actual reconstruction errors in NVS, we follow prior work [21, 33, 37] and report Area Under the Sparsification Error curve (AUSE) [19] and Pearson correlation.

We report both metrics using two error formulations: the standard intensity difference L^1 , and DSSIM. DSSIM is particularly important for evaluating errors in NVS, as these often arise from structural artifacts such as geometric misalignment, blur, or floaters rather than simple pixel-wise deviations. Since DSSIM is sensitive to luminance, contrast, and structural inconsistencies, it provides a more perceptually meaningful measure of NVS error than L^1 .

4.1 Uncertainty Estimation in NVS

Analysis of results: Tab. 1 reports UE comparison on the hold-out (novel) views; here we set $\lambda_{\text{reg}} = 0$, as our Bayesian-inspired regularization yields only marginal gains in dense capture and is instead designed for sparse-view regimes (analyzed in Sec. 4.2). Our method consistently outperforms all baselines by a substantial margin. Notably, our method demonstrates particularly strong performance when evaluated using DSSIM error maps, indicating superior align-

Table 2: Results for (a) spherical harmonics degree ablation and (b) AVS.

(a) Analysis of view-dependence on Mip-NeRF360 [4]. Lowering the capacity of the spherical harmonic degrades the performance of our view-dependent uncertainty maps.

Method	AUSE ($\downarrow$)		Pearson ($\uparrow$)	
	L1	DSSIM	L1	DSSIM
No View-Dependence	0.391	0.310	0.217	0.356
SH Degree = 1	0.352	0.245	0.327	0.498
SH Degree = 2	0.331	0.218	0.353	0.535
SH Degree = 3	0.328	0.214	0.369	0.547

(b) Quantitative results for AVS on the Mip-NeRF360 dataset [4]. Our predicted uncertainty guided AVS achieves the best performance across all metrics.


Method	PSNR	SSIM	LPIPS
	($\uparrow$)	($\uparrow$)	($\downarrow$)
FisherRF [21]	20.266	0.593	0.363
Manifold [37]	19.732	0.595	0.373
Manifold [†]	20.088	0.611	0.350
Ours	20.676	0.615	0.344

ment with human perceptual error relative to existing methods. For instance, we achieve more than $3\times$ *higher Pearson correlation* and *less than half the AUSE* (DSSIM) compared to our best competitor Var3DGS [33] on Mip-NeRF360. Additionally, we achieve minimal training overhead (measured as percentage of additional time required for UE compared to vanilla 3DGS [23] training time). Memory overhead is likewise minimal: we add only $(\text{SH degree} + 1)^2$ scalars per Gaussian (16 at degree 3), and freezing the base parameters keeps training memory below standard 3DGS. App. B further shows that our method is fully plug-and-play across different 3DGS variants, yielding similar or improved UE performance. In Fig. 4, our method is visually compared to the baselines, demonstrating a much higher correlation with the true error.


Analysis on View-Dependence: We ablate the view-dependence of our uncertainty channel u_k by constraining the SH degree used to model it (Tab. 2a). Similar to view-dependent color, the best performances are obtained with a degree of 3, while reducing the capacity degrades UE.

4.2 UE on Highly Novel Views

Experiment Setup: We investigate the contribution of our Bayesian-inspired regularization on highly novel views by adopting a sparse setup using four well-



(a) Visualization of the training and test split for UE on highly novel views. Red: training views; Blue: test views.



(b) Quantitative results for UE on highly novel views on Mip-NeRF360 [4] where our Bayesian-inspired regularization improves performance.

Fig. 5: Highly novel view experiments: (a) example for view setup (b) results




Fig. 6: Qualitative examples of uncertainty maps with Bayesian-inspired regularization for different weights λ_{reg} (see Eq. (7)) in the setting of highly novel views. The given AUSE values are with respect to DSSIM.

distributed training views. These four views are selected by maximizing pairwise camera center distances. We retain the test views in the standard setting (as in Sec. 4.1 and 3DGS [23]) for fair comparison. The train-test split is visualized in Fig. 5a. Due to the sparse coverage of the scene, most test views observe the scene from highly novel viewing directions.

We evaluate on Mip-NeRF360 [4], training the base 3DGS model for 4,000 iterations and our uncertainty channel for 400 iterations to account for the sparse view count. For experiments with Bayesian-inspired regularization, we set the maximal uncertainty level to $b = 1$. This choice follows naturally from our training objective, as uncertainty is regressed toward the photometric reconstruction residual, whose magnitude is bounded by the normalized image intensity range. Thus, a value of 1 corresponds to maximal reconstruction error and provides a principled prior. In these experiments, background regions are treated as maximally uncertain and assigned an uncertainty value of 1. No explicit background prior is imposed when Bayesian regularization is disabled. As a baseline, we compare against FisherRF [21], the only publicly available post-hoc method that can be directly applied to the same splatting outputs as ours.

Analysis of Results: In Fig. 5b, we study the effect of the regularization weight λ_{reg} by evaluating values from 0 (no regularization) to 20.48, doubling the weight from 0.02 onward. UE is assessed on both L^1 and DSSIM error maps. Our approach consistently outperforms FisherRF across all λ_{reg} values in this challenging setting. Compared to the standard UE study, AUSE scores are generally higher, likely due to reduced splatting quality and fewer UE training views, which can make the residual error maps inconsistent and harder to learn. Overall, the Bayesian-inspired regularization improves performance in this regime, though its effect differs between error metrics. For L^1 error, the best AUSE scores are achieved for $0.16 \leq \lambda_{\text{reg}} \leq 0.32$, while for DSSIM, optimal performance occurs at $\lambda_{\text{reg}} = 10.24$, with marginal changes beyond. Notably, these optimal ranges are largely scene-independent; scene-wise AUSE plots are provided in App. C. A qualitative comparison of predicted uncertainty maps for different λ_{reg} values is shown in Fig. 6.

4.3 UE in Active View Selection

Experiment Setup: We follow the AVS setup proposed by FisherRF [21] and evaluate on the Mip-NeRF360 dataset [4]. Following prior work [21,33,37], we re-

port PSNR, SSIM, and LPIPS on the holdout views after reconstruction. Strictly adhering to the 3DGS evaluation protocol [23], we evaluate at full resolution across all methods; while this naturally lowers metrics compared to downsampled versions in prior work [21], it guarantees a fair comparison.

We benchmark our approach against state-of-the-art AVS baselines FisherRF [21] and the method by Lyu *et al.* [37] (Manifold), which employs a stochastic variant of 3DGS. Because Manifold’s variational optimization may degrade the visual fidelity of the underlying 3DGS representation, direct comparison is unfair. Therefore, we introduce Manifold[†], a variant that uses Manifold’s view predictions to guide a vanilla 3DGS [23].

We initialize the AVS procedure with four training views selected by maximizing pairwise camera distances, consistent with our sparse-view setup. Each AVS method then sequentially selects 16 further views, resulting in a total of 20 training views. Between selections, the base 3DGS is trained for $100 \times N_{\text{views}}$ iterations, where N_{views} is the current training view count. For ours, we train the uncertainty channel for $50 \times N_{\text{views}}$ iterations (ablated in App. E) prior to each selection, then selecting the candidate view with highest total uncertainty.

Quantitative results: As reported in Tab. 2b, our method consistently outperforms all baselines across all metrics. This demonstrates that our predicted view-dependent uncertainty maps provide a highly reliable signal for identifying informative viewpoints, substantially improving downstream AVS performance.

4.4 UE in Pose-Agnostic Scene Change Detection

Experiment Setup: Given multi-view captures of a pre-change (*reference*) scene and a post-change (*inference*) scene, pose-agnostic SCD first reconstructs a 3DGS representation of the reference scene $\mathcal{R}_{\text{ref}}$ using reference images $\mathcal{I}_{\text{ref}}$ and their SfM poses $\mathcal{P}_{\text{ref}}$ [49]. For each inference image I_{inf}^k with pose P_{inf}^k , a corresponding rendered view I_{ren}^k is generated by querying $\mathcal{R}_{\text{ref}}$, producing aligned image pairs $(I_{\text{ren}}^k, I_{\text{inf}}^k)$ for change detection. Per-pixel change maps M^k are obtained in a label-free manner, either by comparing dense DINOv2 features [42], direct pixel-level similarity [57], or a combination of both. A threshold is applied to produce binary masks M_{bin}^k [13]. The state-of-the-art MV3DCD [13] further aggregates the predictions $\{M^k\}$ into a 3D change representation, improving robustness via multi-view consistency.

However, the rendered views I_{ren}^k often contain artifacts (Fig. 7), caused by incomplete scene coverage, large capture pose discrepancies, or inherent 3DGS limitations. These inconsistencies lead to false positives, as differences between I_{ren}^k and I_{inf}^k may arise from rendering errors rather than actual scene changes.

Leveraging uncertainty for pose-agnostic SCD: We first apply our UE method from Sec. 3 to estimate uncertainties for the reference scene $\mathcal{R}_{\text{ref}}$. Using these per-primitive values, we render an uncertainty map M_{unc}^k alongside the RGB image I_{ren}^k for each queried inference pose P_{inf}^k . The standard change detection pipeline remains unchanged; however, before thresholding, we attenuate high-uncertainty pixels in the predicted change map M^k , obtaining an uncertainty-guided change map $\hat{M}^k = M^k \odot (1 - M_{\text{unc}}^k)$, where $\odot$ denotes




Fig. 7: Qualitative examples of uncertainty-guided scene change detection. Our predicted uncertainty maps effectively capture rendering artifacts, allowing the system to suppress false positives caused by these artifacts rather than actual scene changes.

Hadamard multiplication. Following standard practice [12, 13], thresholding this map at the midpoint to obtain the final binary mask M_{bin}^k effectively filters out false positives induced by rendering artifacts rather than genuine scene changes.

Quantitative Results: We evaluate our approach on the PASLCD [13] benchmark following MV3DCD [13] and adopt four baselines for comparison. First, we include the *Feature Diff.* baseline, which infers a change mask from differences between dense DINOv2 [42] features. Second, we use MV3DCD’s zero-shot variant (MV3DCD-ZS), which fuses cues from feature- and structure-aware masks [13]. We further incorporate our uncertainty guidance into MV3DCD’s multi-view fusion framework and finally couple it with the existing state-of-the-art online pose-agnostic SCD approach, Online-SCD [12]. As shown in Tab. 3, uncertainty-guidance consistently improves pose-agnostic SCD performance, yielding higher F1 scores and mean Intersection-over-Union (mIoU) across all baselines.

4.5 UE in Pose-Agnostic Anomaly Detection

Experiment Setup: Given multi-view captures of a fault-free object, pose-agnostic AD first reconstructs a 3DGS reference model, denoted as $\mathcal{R}_{\text{ref}}$. During inference, for a given query image I_{inf}^k , its camera pose P_{inf}^k is estimated on-the-fly [31, 35]. Rendering $\mathcal{R}_{\text{ref}}$ at P_{inf}^k yields a corresponding fault-free view, I_{ren}^k . The resulting aligned image pair $(I_{\text{ren}}^k, I_{\text{inf}}^k)$ is compared within the feature space of a pre-trained vision encoder [62] to generate a spatial anomaly score map S^k .

Leveraging Uncertainty for Pose-Agnostic AD: We apply our UE framework (described in Sec. 3) to the optimized reference model $\mathcal{R}_{\text{ref}}$, rendering a per-pixel uncertainty map M_{unc}^k along with I_{ren}^k . Crucially, we leave the underlying anomaly scoring pipeline untouched. Instead, we perform a post-hoc attenuation of the raw score map S^k using our UE to produce an uncertainty-guided anomaly map: $\hat{S}^k = S^k \odot (1 - M_{\text{unc}}^k)$. This suppresses anomalous score activations in regions where the reference 3DGS model lacks confidence.

Quantitative Results: We evaluate our approach on the MAD-Real benchmark [62], reporting AUROC and AUPRO metrics following prior work [31, 35, 62]. By seamlessly integrating our uncertainty-guided attenuation into existing

Table 3: Numerical results for pose-agnostic SCD with uncertainty on PASLCD [13]. Our method consistently improves both state-of-the-art offline and online baselines. Percentage increase over the baseline is given in the $\Delta(\%)$ column.

Metric	Feature Diff.			MV3DCD-ZS [13]			MV3DCD [13]			Online-SCD [12]		
	Base	+ Ours	$\Delta(\%)$	Base	+ Ours	$\Delta(\%)$	Base	+ Ours	$\Delta(\%)$	Base	+ Ours	$\Delta(\%)$
mIoU ($\uparrow$)	0.278	0.359	29.1	0.382	0.439	14.9	0.470	0.498	6.0	0.486	0.498	2.5
F1 ($\uparrow$)	0.402	0.502	24.9	0.526	0.593	12.7	0.621	0.649	4.5	0.638	0.651	2.1

Table 4: Quantitative results for pose-agnostic AD on MAD-Real [62]. Our UE guidance enables SplatPose [31] to match or surpass its successor, SplatPosePlus [35]. Furthermore, applying our framework directly to SplatPosePlus advances the current state-of-the-art. Percentage increase over the baseline is given in the $\Delta(\%)$ column.

Metric	SplatPose [31]			SplatPosePlus [35]		
	Base	+ Ours	$\Delta(\%)$	Base	+ Ours	$\Delta(\%)$
AUROC ($\uparrow$)	0.929	0.939	1.1	0.940	0.956	0.7
AUPRO ($\uparrow$)	0.700	0.765	9.3	0.761	0.798	4.9

3DGS-based AD baselines SplatPose [31] and SplatPosePlus [35], we demonstrate a consistent improvement in AD across both architectures (see Tab. 4).

5 Limitations

Our approach has two main limitations. First, due to its strict post-hoc nature, the spatial granularity of our uncertainty maps is bottlenecked by the density and scale of the underlying Gaussians: in coarsely reconstructed regions (few large primitives), the predicted uncertainty is equally coarse. We do not alter primitives, as our goal is a lightweight estimator preserving model fidelity and structure. Second, we estimate a single *predictive* uncertainty by treating residuals as a proxy for the expected photometric error; this conflates aleatoric, epistemic, and optimization-related effects rather than disentangling them, and does not yield a probabilistic distribution. These limitations are not unique; existing approaches also lack disentanglement, while only stochastic formulations [33, 37] provide a probabilistic form. We further study controlled aleatoric and epistemic perturbations in App. G and find our UE behaves as expected.

6 Conclusion

We introduced a lightweight, plug-and-play framework for predictive photometric uncertainty estimation in 3DGS-based NVS. By formulating post-hoc uncertainty estimation as a linear least-squares problem with Bayesian-inspired regularization, our method effectively predicts pixel-level uncertainty maps under both dense and sparse captures. This enables state-of-the-art results in active view selection, pose-agnostic scene change detection, and anomaly detection.

Acknowledgment

This work was supported by the Australian Research Council Research Hub in Intelligent Robotic Systems for Real-Time Asset Management (IH210100030) (ARIAM) and Abyss Solutions. C.J., N.S., and D.M. also acknowledge ongoing support from the QUT Centre for Robotics. T.G. P.S., and M.R. acknowledge support by the state of North Rhine-Westphalia and the European Union within the EFRE/JTF project “Just scan it 3D”, grant no. EFRE-20800529. E.H. and M.R. acknowledge support through the junior research group project “UnREAL” by the German Federal Ministry of Education and Research (BMBF), grant no. 01IS22069. M.R. also acknowledges mobility support by the German Academic Exchange Service (DAAD PPP), grant no. 57700453.

References

1. Aira, L.S., Valsesia, D., Magli, E.: Modeling uncertainty for gaussian splatting. *IEEE Transactions on Neural Networks and Learning Systems* **36**(6), 11657–11663 (2025)
2. Alcantarilla, P.F., Stent, S., Ros, G., Arroyo, R., Gherardi, R.: Street-view change detection with deconvolutional networks. *Autonomous Robots* **42**(7), 1301–1322 (2018)
3. Alpherts, T., Ghebreab, S., van Noord, N.: Emplace: Self-supervised urban scene change detection. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 39 (2), pp. 1737–1745 (2025)
4. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5470–5479 (2022)
5. Bergmann, P., Fauser, M., Sattlegger, D., Steger, C.: Mvtec ad — a comprehensive real-world dataset for unsupervised anomaly detection. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 9584–9592 (2019)
6. Bishop, C.M., Nasrabadi, N.M.: *Pattern recognition and machine learning*, vol. 4. Springer (2006)
7. Chen, S., Yang, K., Stiefelhagen, R.: Dr-tanet: Dynamic receptive temporal attention network for street scene change detection. In: *2021 IEEE Intelligent Vehicles Symposium (IV)*. pp. 502–509. IEEE (2021)
8. Cho, K., Kim, D.Y., Kim, E.: Zero-shot scene change detection. In: *Proceedings of the AAAI Conference on Artificial Intelligence*. vol. 39 (3), pp. 2509–2517 (2025)
9. Chung, J., Oh, J., Lee, K.M.: Depth-regularized optimization for 3d gaussian splatting in few-shot images. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 811–820 (2024)
10. Daudt, R.C., Le Saux, B., Boulch, A.: Fully convolutional siamese networks for change detection. In: *2018 25th IEEE international conference on image processing (ICIP)*. pp. 4063–4067. IEEE (2018)
11. Fei, B., Xu, J., Zhang, R., Zhou, Q., Yang, W., He, Y.: 3d gaussian splatting as new era: A survey. *IEEE Transactions on Visualization and Computer Graphics* (2024)

12. Galappaththige, C.J., Lai, J., Windrim, L., Dansereau, D., Suenderhauf, N., Miller, D.: Changes in real time: Online scene change detection with multi-view fusion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 32246–32256 (2026)
13. Galappaththige, C.J., Lai, J., Windrim, L., Dansereau, D., Sunderhauf, N., Miller, D.: Multi-view pose-agnostic change localization with zero labels. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 11600–11610 (2025)
14. Goli, L., Reading, C., Sellán, S., Jacobson, A., Tagliasacchi, A.: Bayes’ rays: Uncertainty quantification for neural radiance fields. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 20061–20070 (2024)
15. Gottwald, T., Heinert, E., Stehr, P., Galappaththige, C.J., Rottmann, M.: Primu: Uncertainty estimation for novel views in gaussian splatting from primitive-based representations of error and coverage. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 11871–11880 (2026)
16. Hanson, A., Tu, A., Lin, G., Singla, V., Zwicker, M., Goldstein, T.: Speedy-splat: Fast 3d gaussian splatting with sparse pixels and sparse primitives. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 21537–21546 (2025)
17. Hanson, A., Tu, A., Singla, V., Jayawardhana, M., Zwicker, M., Goldstein, T.: PUP 3D-GS: Principled uncertainty pruning for 3d gaussian splatting. In: *Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR)*. pp. 5949–5958 (2025)
18. Hedman, P., Philip, J., Price, T., Frahm, J.M., Drettakis, G., Brostow, G.: Deep blending for free-viewpoint image-based rendering. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* **37**(6), 257:1–257:15 (2018)
19. Ilg, E., Cicek, O., Galesso, S., Klein, A., Makansi, O., Hutter, F., Brox, T.: Uncertainty estimates and multi-hypotheses networks for optical flow. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 652–667 (2018)
20. Jiang, B., Huang, R., Zhao, Q., Zhang, Y.: Gaussian difference: Find any change instance in 3d scenes. In: *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. pp. 1–5. IEEE (2025)
21. Jiang, W., Lei, B., Daniilidis, K.: Fisherrf: Active view selection and mapping with radiance fields using fisher information. In: *European Conference on Computer Vision*. pp. 422–440. Springer (2024)
22. Kamman, S.S., Min, B.C.: Zeroscd: Zero-shot street scene change detection. In: *2025 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 4665–4671 (2025)
23. Kerbl, B., Kopanas, G., Leimkuehler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics (TOG)* **42**(4), 1–14 (2023)
24. Kerbl, B., Meuleman, A., Kopanas, G., Wimmer, M., Lanvin, A., Drettakis, G.: A hierarchical 3d gaussian representation for real-time rendering of very large datasets. *ACM Transactions on Graphics (TOG)* **43**(4), 1–15 (2024)
25. Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.C., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: 3d gaussian splatting as markov chain monte carlo. *Advances in Neural Information Processing Systems* **37**, 80965–80986 (2024)
26. Kim, J.W., Kim, U.H.: Towards generalizable scene change detection. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 24463–24473 (2025)

27. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 4015–4026 (2023)
28. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics* **36**(4) (2017)
29. Kopanas, G., Drettakis, G.: Improving NeRF Quality by Progressive Camera Placement for Free-Viewpoint Navigation. In: Guthe, M., Grosch, T. (eds.) *Vision, Modeling, and Visualization*. The Eurographics Association (2023)
30. Kowsky, W.: A quadrature formula over the sphere with application to high resolution spherical harmonic analysis. *Bull. Gdod* **60**, 1–14 (1986)
31. Kruse, M., Rudolph, M., Woiwode, D., Rosenhahn, B.: Splatpose & detect: Pose-agnostic 3d anomaly detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3950–3960 (2024)
32. Lei, Y., Peng, D., Zhang, P., Ke, Q., Li, H.: Hierarchical paired channel fusion network for street scene change detection. *IEEE Transactions on Image Processing* **30**, 55–67 (2020)
33. Li, R., Cheung, Y.m.: Variational multi-scale representation for estimating uncertainty in 3d gaussian splatting. *Advances in Neural Information Processing Systems* **37**, 87934–87958 (2024)
34. Lin, C.J., Garg, S., Chin, T.J., Dayoub, F.: Robust scene change detection using visual foundation models and cross-attention mechanisms. In: 2025 IEEE International Conference on Robotics and Automation (ICRA). pp. 8337–8343. IEEE (2025)
35. Liu, Y., Hu, Y.S., Chen, Y., Zelek, J.: Splatpose+: Real-time image-based pose-agnostic 3d anomaly detection. In: European Conference on Computer Vision Workshops. pp. 378–391. Springer (2024)
36. Lu, Z., Ye, J., Leonard, J.: 3dgs-cd: 3d gaussian splatting-based change detection for physical object rearrangement. *IEEE Robotics and Automation Letters* (2025)
37. Lyu, L., Tewari, A., Habermann, M., Saito, S., Zollhöfer, M., Leimkühler, T., Theobalt, C.: Manifold sampling for differentiable uncertainty in radiance fields. In: *SIGGRAPH Asia 2024 Conference Papers*. pp. 1–11 (2024)
38. Mallick, S.S., Goel, R., Kerbl, B., Steinberger, M., Carrasco, F.V., De La Torre, F.: Taming 3dgs: High-quality radiance fields with limited resources. In: *SIGGRAPH Asia 2024 Conference Papers*. pp. 1–11 (2024)
39. Matsuki, H., Murai, R., Kelly, P.H., Davison, A.J.: Gaussian splatting slam. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18039–18048 (2024)
40. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
41. Nakayama, K., Uy, M.A., You, Y., Li, K., Guibas, L.J.: Provnerf: Modeling per point provenance in nerfs as a stochastic field. *Advances in Neural Information Processing Systems* **37**, 99145–99160 (2024)
42. Oquab, M., Darcet, T., Moutakanni, T., Vo, H., Szafraniec, M., Khalidov, V., Fernandez, P., Haziza, D., Massa, F., El-Nouby, A., et al.: Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193* (2023)
43. Pan, X., Lai, Z., Song, S., Huang, G.: Activenerf: Learning where to see with uncertainty estimation. In: Avidan, S., Brostow, G., Cissé, M., Farinella, G.M., Hassner, T. (eds.) *Computer Vision – ECCV 2022*. pp. 230–246. Springer Nature Switzerland, Cham (2022)

44. Rota Bulò, S., Porzi, L., Kotschieder, P.: Revising densification in gaussian splatting. In: Proceedings of the European conference on computer vision (ECCV). p. 347–362 (2024)
45. Roth, K., Pemula, L., Zepeda, J., Scholkopf, B., Brox, T., Gehler, P.: Towards total recall in industrial anomaly detection. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 14298–14308 (2021)
46. Sachdeva, R., Zisserman, A.: The change you want to see. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3993–4002 (2023)
47. Sachdeva, R., Zisserman, A.: The change you want to see (now in 3d). In: Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops. pp. 2060–2069 (2023)
48. Sakurada, K., Okatani, T.: Change detection from a street image pair using cnn features and superpixel segmentation. In: Proceedings of the British Machine Vision Conference (BMVC). pp. 61.1–61.12 (2015)
49. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4104–4113 (2016)
50. Seo, S., Chang, Y., Kwak, N.: Flipnerf: Flipped reflection rays for few-shot novel view synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 22883–22893 (2023)
51. Shen, J., Agudo, A., Moreno-Noguer, F., Ruiz, A.: Conditional-flow nerf: Accurate 3d modelling with reliable uncertainty quantification. In: European Conference on Computer Vision. pp. 540–557. Springer (2022)
52. Shen, J., Ruiz, A., Agudo, A., Moreno-Noguer, F.: Stochastic neural radiance fields: Quantifying uncertainty in implicit 3d representations. In: 2021 International Conference on 3D Vision (3DV). pp. 972–981. IEEE (2021)
53. Sünderhauf, N., Brock, O., Scheirer, W., Hadsell, R., Fox, D., Leitner, J., Upcroft, B., Abbeel, P., Burgard, W., Milford, M., et al.: The limits and potentials of deep learning for robotics. The International journal of robotics research **37**(4-5), 405–420 (2018)
54. Sünderhauf, N., Abou-Chakra, J., Miller, D.: Density-aware nerf ensembles: Quantifying predictive uncertainty in neural radiance fields. In: 2023 IEEE International Conference on Robotics and Automation (ICRA). pp. 9370–9376 (2023)
55. Tarantola, A.: Inverse problem theory and methods for model parameter estimation. SIAM (2005)
56. Varghese, A., Gubbi, J., Ramaswamy, A., Balamuralidhar, P.: Changenet: A deep learning architecture for visual change detection. In: Proceedings of the European conference on computer vision (ECCV) workshops. pp. 0–0 (2018)
57. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing **13**(4), 600–612 (2004)
58. Wilson, J., Almeida, M., Mahajan, S., Labrie, M., Ghaffari, M., Ghasemalizadeh, O., Sun, M., Kuo, C.H., Sen, A.: POp-GS: Next Best View in 3D-Gaussian Splatting with P-Optimality. In: Proceedings of the Computer Vision and Pattern Recognition Conference (CVPR). pp. 3646–3655 (2025)
59. Wilson, J., Almeida, M., Sun, M., Mahajan, S., Ghaffari, M., Ewen, P., Ghasemalizadeh, O., Kuo, C.H., Sen, A.: Modeling uncertainty in 3d gaussian splatting through continuous semantic splatting. In: 2025 IEEE International Conference on Robotics and Automation (ICRA). pp. 3284–3290 (2025)

60. Wu, G., Yi, T., Fang, J., Xie, L., Zhang, X., Wei, W., Liu, W., Tian, Q., Wang, X.: 4d gaussian splatting for real-time dynamic scene rendering. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 20310–20320 (2024)
61. Ye, Z., Li, W., Liu, S., Qiao, P., Dou, Y.: Absgs: Recovering fine details in 3d gaussian splatting. In: Proceedings of the 32nd ACM International Conference on Multimedia. pp. 1053–1061 (2024)
62. Zhou, Q., Li, W., Jiang, L., Wang, G., Zhou, G., Zhang, S., Zhao, H.: Pad: A dataset and benchmark for pose-agnostic anomaly detection. *Advances in Neural Information Processing Systems* **36**, 44558–44571 (2023)
63. Zhu, L., Li, Y., Sandström, E., Huang, S., Schindler, K., Armeni, I.: Loopsplat: Loop closure by registering 3d gaussian splats. In: 2025 International Conference on 3D Vision (3DV). pp. 156–167. IEEE (2025)
64. Zwicker, M., Pfister, H., Van Baar, J., Gross, M.: Ewa volume splatting. In: Proceedings Visualization, 2001. VIS'01. pp. 29–36. IEEE (2001)