

Active View Selection with Perturbed Gaussian Ensemble for Tomographic Reconstruction

Yulun Wu¹, Ruyi Zha², Wei Cao¹,
Yingying Li¹, Yuanhao Cai^{3,*}, and Yaoyao Liu¹

¹ University of Illinois Urbana-Champaign

² Australian National University

³ Johns Hopkins University

{yulun5, weicao3, yl101, lyy}@illinois.edu

ruyi.zha@anu.edu.au ycai51@jh.edu

Abstract. Sparse-view computed tomography (CT) is critical for reducing radiation exposure to patients. Recent advances in radiative 3D Gaussian Splatting (3DGS) have enabled fast and accurate sparse-view CT reconstruction. Despite these algorithmic advancements, practical reconstruction fidelity remains fundamentally bounded by the quality of the captured data, raising the crucial yet underexplored problem of X-ray active view selection. Existing active view selection methods are primarily designed for natural-light scenes and fail to capture the unique geometric ambiguities and physical attenuation properties inherent in X-ray imaging. In this paper, we present *Perturbed Gaussian Ensemble*, an active view selection framework that integrates uncertainty modeling with sequential decision-making, tailored for X-ray Gaussian Splatting. Specifically, we identify low-density Gaussian primitives that are likely to be uncertain and apply stochastic density scaling to construct an ensemble of plausible Gaussian density fields. For each candidate projection, we measure the structural variance of the ensemble predictions and select the one with the highest variance as the next best view. Extensive experimental results on arbitrary-trajectory CT benchmarks demonstrate that our density-guided perturbation strategy effectively eliminates geometric artifacts and consistently outperforms existing baselines in progressive tomographic reconstruction under unified view selection protocols.

Keywords: CT Reconstruction · Gaussian Splatting · Active Learning

1 Introduction

X-ray computed tomography (CT) is an indispensable non-invasive imaging modality widely utilized in medical diagnosis, industrial inspection, and scientific research. During a CT scan, an X-ray machine captures multi-angle 2D projections that measure ray attenuation through the material. Tomographic reconstruction can then be performed to recover high-fidelity 3D anatomical

* Corresponding author.




Fig. 1: (a) We compare our approach against the state-of-the-art 3DGS-based active view selection method, FisherRF [25], evaluating both quantitative metrics (3D PSNR $\uparrow$ and SSIM $\uparrow$) and qualitative visual fidelity, including zoomed-in details. Our method achieves the highest reconstruction quality and best preserves fine structural details. (b) We plot the training iterations *vs.* the average 3D PSNR $\uparrow$ (dB) and SSIM $\uparrow$ on the synthetic dataset. As demonstrated, our proposed approach consistently delivers superior reconstruction fidelity throughout the progressive reconstruction process.

structures from the 2D projections. However, the ionizing radiation associated with prolonged X-ray exposure poses significant health risks.

To mitigate these risks, sparse-view CT has emerged as a critical scanning paradigm, aiming to maximize the reconstruction quality with a limited number of projection angles. Mathematically, this reduction in data acquisition transforms the tomographic reconstruction into a highly ill-posed inverse problem. Existing algorithms and solvers often struggle under such sparse conditions.

Recent breakthroughs in neural rendering, particularly 3D Gaussian Splatting (3DGS) [31], have demonstrated a remarkable performance superiority in both sparse-view novel view synthesis [35, 83, 86] and surface reconstruction [7, 23, 24, 71]. Building on this, several methods [5, 78, 80] have successfully adapted 3DGS to X-ray imaging, enabling rapid and highly accurate volumetric reconstructions. However, despite these algorithmic leaps, practical reconstruction quality remains fundamentally bottlenecked by the captured data. Consequently, how to accurately select scanning viewpoints under a limited view-budget, so as to ensure both global silhouette coverage and the capture of local structural details, has emerged as a critical challenge in sparse-view tomography.

While crucial, active view selection for X-ray Gaussian Splatting remains largely unexplored. Existing view selection strategies [25, 55] are predominantly designed for natural-light scenes, relying heavily on surface occlusions and view-dependent specularities to estimate uncertainty or information gain. In stark contrast, X-ray imaging operates on the Beer-Lambert law [29], where the projec-

tion is a purely linear integral of the continuous density field along the ray path, without any occlusion. Furthermore, X-ray attenuation is inherently isotropic, meaning the constituent Gaussian primitives have no spherical harmonic (SH) parameters. Consequently, methods that depend on view-dependent color gradients and assume sparse, occlusion-bounded ray interactions fail to accurately capture the volumetric ambiguity in CT. As a result, these methods struggle to distinguish between the stretched artifacts and true high-density structures, often yielding redundant view selections that fail to provide the orthogonal constraints necessary to eliminate such artifacts.

In this study, we present *Perturbed Gaussian Ensemble*,¹ a novel active view selection framework for progressive reconstruction specifically tailored to X-ray Gaussian Splatting and designed for flexible C-arm CT systems. Our core intuition is that under sparse-view constraints, geometric ambiguities typically manifest as fragile structures, such as uncertain boundaries and needle-like artifacts, whose projections vary dramatically when observed from informative, previously unseen angles. A valid next best view is the one that maximizes the exposure of this underlying structural instability.

To operationalize this insight, we utilize low-density Gaussian primitives as a proxy for uncertainty-prone regions, as they typically correspond to under-constrained boundaries, background noise, or degenerated artifact tails. By applying stochastic perturbations to the densities of these specific primitives, we construct an ensemble of plausible Gaussian density fields. For each candidate viewpoint, we render predictions across this ensemble and quantify their structural disagreement using the variance of the Structural Similarity Index Measure (SSIM) [69]. A high structural variance indicates that minor perturbations in uncertain regions provoke substantial structural discrepancies in that specific projection, marking the viewpoint as highly informative for resolving existing ambiguities. We therefore select the viewpoint that maximizes this structural variance as the optimal next best view.

We comprehensively evaluate our method against 2D-based paradigms, 3D-based paradigms, and conventional rule-based heuristics under two different view selection protocols within a hemispherical scanning search space. Experimental results demonstrate that our density-guided perturbation strategy effectively eliminates geometric artifacts and consistently outperforms existing baselines in progressive tomographic reconstruction, as illustrated in Fig. 1.

Our contribution can be summarized as follows:

- We propose a novel active view selection and progressive reconstruction framework designed for X-ray Gaussian Splatting. Our work bridges the gap between active learning and explicit radiative fields, addressing the unique physical and geometric challenges of sparse-view computed tomography.
- We introduce a novel uncertainty quantification strategy based on the *Perturbed Gaussian Ensemble*. By applying stochastic density perturbations to low-density primitives that are highly susceptible to geometric degradation

¹ Project page: perturbed-gaussian-ensemble.cvmgroup.web.illinois.edu

- and measuring the structural disagreement in projection space, our method accurately localizes epistemic uncertainty and predicts the next best view.
- We establish an active view selection and progressive reconstruction benchmark for radiative Gaussian Splatting by adapting and evaluating previous state-of-the-art baselines. Extensive experiments demonstrate that our approach consistently outperforms existing paradigms, achieving superior volumetric reconstruction and novel view synthesis quality.

2 Related Work

Tomographic Reconstruction aims to recover the 3D internal density field of an object from 2D X-ray projections. Traditional analytical methods [16, 77] suffer from severe streak artifacts and noise amplification under sparse-view conditions. Iterative reconstruction algorithms [2, 53, 61, 64] rely on optimization over iterations; however, they are computationally expensive and tend to oversmooth fine structural details. While supervised deep learning approaches [1, 3, 12, 20, 26, 34, 41–43, 50, 76] leverage semantic priors to enhance image quality, their generalization capabilities often remain constrained. NeRF-based frameworks [6, 59, 63, 79, 81] model the continuous density field using coordinate-based MLPs optimized purely by photometric losses from sparse projections. Despite achieving high-fidelity reconstructions, their training and rendering processes are prohibitively slow. The emergence of 3D Gaussian Splatting (3DGS) [31] offers a compelling alternative due to its explicit parameterization and highly parallelized rasterization. Early adaptations to X-ray imaging [5, 19] focus primarily on novel view synthesis rather than direct volume retrieval. R²-Gaussian [80] achieves a critical breakthrough by identifying and rectifying the omission of a covariance-related scaling factor during 3D-to-2D projection. By introducing tailored radiative Gaussian kernels and a differentiable voxelizer, it enables direct, bias-free, and rapid static tomographic reconstruction. Subsequent works [36, 51, 78, 82] further explore the sparse-view CT reconstruction paradigm or extend the radiative splatting framework into the temporal domain. However, active view selection for CT remains largely unexplored. Therefore, this study undertakes a more in-depth investigation and introduces a targeted solution to this problem.

Active View Selection (AVS) or next best view (NBV) planning aims to incrementally determine the most informative viewpoints for scene reconstruction to minimize data acquisition costs. It is closely related to continual learning and online learning [15, 18, 37–39, 44–49, 52, 84, 87]. It originates from robotics research [13, 62] and sees extensive exploration in 3D reconstruction [10, 11, 14, 17, 27, 57]. With the rise of neural rendering, recent works adapt AVS to neural radiance fields (NeRF) [54]. NeRF-based methods [21, 55, 66, 72, 74] quantify uncertainty via variance estimation to guide view acquisition, but they remain computationally intensive and primarily apply to synthetic settings. The advent of 3D Gaussian Splatting (3DGS) [31] significantly accelerates active view selection and reconstruction due to its explicit representation and highly efficient




Fig. 2: Comparison of active view selection paradigms for X-ray Gaussian Splatting. (a) Gradient-based method [25] estimates the expected information gain (EIG) of candidate views by computing a Fisher Information Matrix (FIM) via backpropagation, but suffers from gradient coupling and the absence of view-dependent parameters. (b) Ensemble-based method quantifies epistemic uncertainty by rendering disagreement across multiple Gaussian representations initialized with different random seeds, resulting in prohibitively high computational burden. (c) Our *Perturbed Gaussian Ensemble* introduces density-guided parameter perturbations to efficiently construct an ensemble, wherein uncertain primitives exhibit pronounced behavioral randomness. By calculating the structural variance of the rendered projections, our approach achieves superior uncertainty modeling and enables more effective active view selection.

rasterization. 3DGS-based methods [25, 28, 40, 70, 73, 85] adopt uncertainty or visibility-driven selection and show great promise for efficiency. However, current gradient-based algorithms assume surface-based occlusion and heavily rely on view-dependent color parameters. In contrast, X-ray imaging operates on a transmission model governed by the Beer-Lambert law [29], where projections serve as linear integrals of the density field without occlusion. This transmissive nature causes high spatial coupling among Gaussians along the ray, severely violating the diagonal approximation assumptions of gradient-based methods. Although uncertainty estimation and active acquisition have also been explored for CT implicit neural representations [67, 88], it remains challenging to radiative Gaussian Splatting. To address this gap, we abandon the gradient-based heuristic and introduce a forward parameter perturbation strategy specifically tailored to the explicit density parameters of radiative Gaussians.

3 Method

Given a sparse initial collection of measured X-ray projections $\mathcal{I}_{\text{init}} = \{\mathbf{I}_{\text{init}}^{(i)} \in \mathbb{R}^{H \times W}\}_{i=1, \dots, N_{\text{init}}}$, the aim is to iteratively select the next best view (NBV) from a pool of candidate scanner poses and incorporate the corresponding measurement into the training-view set until a predefined target number of observations N_{target} is reached. Throughout this process, we progressively reconstruct the 3D volume to maximize the final reconstruction quality. To achieve this, we exploit the physical properties of radiative Gaussian Splatting to introduce a novel active view selection framework. As illustrated in Fig. 2, our approach models viewpoint uncertainty through the use of *Perturbed Gaussian Ensemble*.

In this section, we present our proposed active view selection framework. First, we introduce the radiative Gaussian Splatting model and establish the theoretical basis in Sec. 3.1. Next, we detail the construction of *Perturbed Gaussian Ensemble* through a density-guided parameter perturbation strategy in Sec. 3.2. Finally, Sec. 3.3 formulates our view selection mechanism, which quantifies epistemic uncertainty by evaluating structural variance in the projection space.

3.1 Preliminary: Radiative Gaussian Splatting

To facilitate efficient and physically correct tomographic reconstruction, our active view selection mechanism is built on a radiative Gaussian Splatting framework, which reformulates 3D Gaussian Splatting (3DGS) [31] for transmission imaging. Following [80], we model the object’s 3D density field σ as a mixture of M radiative Gaussian kernels. Each kernel G_i acts as a local volumetric density primitive parameterized by a central density scalar ρ_i , a mean position $\mathbf{p}_i \in \mathbb{R}^3$, and a 3D covariance matrix $\Sigma_i \in \mathbb{R}^{3 \times 3}$. The density field at a spatial location $\mathbf{x} \in \mathbb{R}^3$ is formulated as

$$\sigma(\mathbf{x}) = \sum_{i=1}^M G_i^3(\mathbf{x} \mid \rho_i, \mathbf{p}_i, \Sigma_i) = \sum_{i=1}^M \rho_i \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{p}_i)^\top \Sigma_i^{-1}(\mathbf{x} - \mathbf{p}_i)\right). \quad (1)$$

The rendering process simulates the physical acquisition of X-ray projections. Due to the linearity of the integral operator, the total log-space projection is the sum of the line integrals of individual Gaussians:

$$I_r(\mathbf{r}) = \sum_{i=1}^M G_i^2(\hat{\mathbf{x}} \mid \rho_i \sqrt{\frac{2\pi|\Sigma_i|}{|\hat{\Sigma}_i|}}, \hat{\mathbf{p}}_i, \hat{\Sigma}_i), \quad (2)$$

where $I_r(\mathbf{r})$ is the rendered pixel value; $\hat{\mathbf{x}} \in \mathbb{R}^2$, $\hat{\mathbf{p}} \in \mathbb{R}^2$, $\hat{\Sigma} \in \mathbb{R}^{2 \times 2}$ are obtained by projecting the Gaussians onto the image plane from a viewpoint. The term $\sqrt{2\pi|\Sigma_i|/|\hat{\Sigma}_i|}$ (where $|\cdot|$ denotes the matrix determinant) acts as a normalization factor that ensures conservation of the integrated density mass, rendering the optimization of ρ_i physically meaningful and view-consistent.

Radiative Gaussians are optimized by photometric L1 loss $\mathcal{L}_1$, D-SSIM loss $\mathcal{L}_{\text{ssim}}$ [69], and 3D total variation regularization $\mathcal{L}_{\text{tv}}$ [60]. The overall loss function is defined as

$$\mathcal{L} = \mathcal{L}_1(\mathbf{I}_r, \mathbf{I}_m) + \lambda_1 \mathcal{L}_{\text{ssim}}(\mathbf{I}_r, \mathbf{I}_m) + \lambda_2 \mathcal{L}_{\text{tv}}, \quad (3)$$

where $\mathbf{I}_r, \mathbf{I}_m$ are the rendered and measured projections. λ_1, λ_2 are weight terms.

3.2 Perturbed Gaussian Ensemble

Due to the transmissive and isotropic nature of X-rays, active view selection methods designed for natural-light scenes, such as FisherRF [25], perform poorly in X-ray CT settings. To maintain real-time performance, FisherRF evaluates the expected information gain (EIG) using a diagonal approximation of the Fisher Information Matrix (FIM). While ignoring parameter correlations is valid for natural-light rendering, where pixels are dominated by a few front-most Gaussian primitives, this assumption fails in X-ray imaging. The transmissive nature of X-rays highly couples Gaussians along a ray, causing substantial mathematical bias in the EIG under a diagonal FIM.

To circumvent the computational intractability of evaluating a dense FIM and the inherent inaccuracies of its diagonal approximation, we abandon the gradient-based heuristic. Instead, we propose a forward, sampling-based approach that directly explores the coupled geometric uncertainty.

Uncertainty Quantification via Rendering Disagreement. In sparse-view CT reconstruction, the lack of sufficient multi-view geometry constraints leads to severe ill-posedness, introducing significant uncertainty into the optimization of Gaussian fields. Let $\mathcal{D}$ denote the training data and θ represent the parameters of a radiative Gaussian Splatting model. For an arbitrary scanner pose $\mathbf{v}$, the predicted projection $\mathbf{I}(\mathbf{v})$ is obtained through the rendering process:

$$\mathbf{I}(\mathbf{v}) = f(\mathbf{v}; \theta). \quad (4)$$

During optimization, we minimize the total loss to obtain the optimal model parameters:

$$\hat{\theta} = \arg \min_{\theta} \mathcal{L}(\theta; \mathcal{D}) + \lambda R(\theta), \quad (5)$$

where $R(\theta)$ represents the regularization term. Due to the highly non-convex nature of the objective function, the optimization problem admits multiple minima. Consequently, when initialized with different random seeds $\{k_i\}_{i=1, \dots, N}$, the model converges to distinct local optima:

$$\hat{\theta}_i \sim q(\theta \mid \mathcal{D}), \quad (6)$$

where $q(\theta \mid \mathcal{D})$ represents an empirical approximation of the true posterior $p(\theta \mid \mathcal{D})$. Thus, the posterior predictive distribution at viewpoint $\mathbf{v}$ can be approximated via Monte Carlo integration over the ensemble:

$$p(\mathbf{I}(\mathbf{v}) \mid \mathcal{D}) \approx \frac{1}{N} \sum_{i=1}^N \delta(\mathbf{I}(\mathbf{v}) - \mathbf{I}_i(\mathbf{v})), \quad \mathbf{I}_i(\mathbf{v}) = f(\mathbf{v}; \hat{\theta}_i), \quad (7)$$

where $\delta(\cdot)$ denotes the Dirac delta function. Therefore, the disagreement in the rendered projections across an ensemble of N explicit Gaussian representations $\{\mathcal{G}_i\}_{i=1}^N$ initialized with different random seeds $\{k_i\}_{i=1,\dots,N}$ effectively captures the epistemic uncertainty at a given view. We quantify this uncertainty using the sample variance:

$$\text{Var}[\mathbf{I}(\mathbf{v})] = \frac{1}{N-1} \sum_{i=1}^N (\mathbf{I}_i(\mathbf{v}) - \bar{\mathbf{I}}(\mathbf{v}))^2, \quad \bar{\mathbf{I}}(\mathbf{v}) = \frac{1}{N} \sum_{i=1}^N \mathbf{I}_i(\mathbf{v}). \quad (8)$$

Density-Guided Parameter Perturbation. Training an ensemble of Gaussian models is an intuitive and effective approach for viewpoint uncertainty modeling. However, this strategy requires training and storing N independent Gaussian Splatting models, leading to a linear increase in computational and memory complexity, which is often prohibitive in practical applications. To address this, we propose *Perturbed Gaussian Ensemble*, a simulation-based approach that leverages the intrinsic properties of the radiative Gaussian model.

Specifically, the variance in the rendered outputs across an ensemble at a given viewpoint is fundamentally induced by the randomness in the learned parameters of the individual Gaussian primitives. Therefore, a more efficient strategy is to train only a single Gaussian model $\mathcal{G} = \{G_i\}_{i=1,\dots,M}$ and, during uncertainty evaluation for candidate viewpoints, independently perturb the parameters of the Gaussian representation with repeated noise injections. This efficiently emulates an ensemble that would otherwise be obtained by training with different random seeds:

$$\mathcal{G}_i = \{G_{i,j}\}, \quad G_{i,j} = G_j \oplus \Delta_{i,j}, \quad i = 1, \dots, N, \quad j = 1, \dots, M. \quad (9)$$

To ensure that the rendering disagreement of this perturbed ensemble accurately reflects the epistemic uncertainty arising from the actual training process, the noise injected into each individual Gaussian primitive should be conditioned on its specific uncertainty level.

According to Eqs. (1) and (2), the contribution of a Gaussian primitive during rendering is heavily determined by its density. High-density primitives typically correspond to real, well-defined solid structures (e.g., bones and dense organs). Because these structures strongly attenuate X-rays, they are usually well-constrained even under sparse-view settings. Consequently, the model exhibits high confidence in these regions.

In contrast, low-density primitives often lie near object boundaries (e.g., soft-tissue interfaces) or constitute background noise introduced to fit only a small number of rays, thereby exhibiting higher uncertainty. Moreover, due to insufficient multi-view constraints, the Gaussians tend to overfit the training views during optimization, giving rise to extremely elongated, needle-like artifacts or numerous tiny, low-density Gaussians distributed along the ray paths. These spurious components likewise exhibit low density and high uncertainty.

Motivated by this observation, we isolate a vulnerable Gaussian subset $\mathcal{G}_{\text{low}}$ comprising a specific fraction α (e.g., 10%) of the total Gaussians in $\mathcal{G}$ with the

lowest density values. We then stochastically perturb the density parameters of this subset to probe the stability of the reconstructed scene structure. Specifically, we generate an ensemble of N independently perturbed models. For the i -th ensemble member $\mathcal{G}_i$ ($i = 1, \dots, N$), the density $\rho_{i,j}$ of a perturbed Gaussian primitive $G_{i,j}$ derived from $G_j \in \mathcal{G}$ ($j = 1, \dots, M$) is defined as:

$$\rho_{i,j} = \begin{cases} \rho_j \cdot (1 + \epsilon_{i,j}), & \text{if } G_j \in \mathcal{G}_{\text{low}} \\ \rho_j, & \text{otherwise} \end{cases}, \quad (10)$$

where $\epsilon_{i,j}$ is a random scaling factor sampled from a uniform distribution:

$$\epsilon_{i,j} \sim \text{Uniform}(-\beta, \beta), \quad \beta > 0. \quad (11)$$

This targeted density scaling effectively injects variations into the most ambiguous regions of the volume while preserving the high-density, high-confidence anatomical structures.

3.3 View Selection by Structural Variance

The core objective of active view selection is to identify the candidate viewpoint that provides the most informative structural constraints to resolve existing geometric ambiguities. For an arbitrary candidate viewpoint $\mathbf{v}$, rendered projections are generated using the *Perturbed Gaussian Ensemble* $\{\mathcal{G}, \mathcal{G}_1, \dots, \mathcal{G}_N\}$ via the rendering function:

$$\mathbf{I}(\mathbf{v}) = f(\mathbf{v}; \theta(\mathcal{G})), \quad \mathbf{I}_i(\mathbf{v}) = f(\mathbf{v}; \theta(\mathcal{G}_i)), \quad i = 1, \dots, N. \quad (12)$$

To quantify the epistemic uncertainty at viewpoint $\mathbf{v}$, we leverage the Structural Similarity Index Measure (SSIM) [69] in the projection space. This allows us to evaluate the macroscopic structural disagreement induced by our parameter perturbations:

$$s_i(\mathbf{v}) = \text{SSIM}(\mathbf{I}(\mathbf{v}), \mathbf{I}_i(\mathbf{v})). \quad (13)$$

The uncertainty score $u(\mathbf{v})$ for the candidate viewpoint is formulated as the sample variance of these N structural similarity scores:

$$u(\mathbf{v}) = \frac{1}{N-1} \sum_{i=1}^N (s_i(\mathbf{v}) - \bar{s}(\mathbf{v}))^2, \quad (14)$$

where $\bar{s}(\mathbf{v})$ denotes the mean of the N SSIM scores.

At each view selection iteration t , a new perturbed ensemble is generated from the most recently updated Gaussian representation $\mathcal{G}^{(t)}$. For every viewpoint from the candidate pool $\mathcal{V}_{\text{cand}}^{(t)}$, we calculate the epistemic uncertainty score using this ensemble. Finally, the viewpoint exhibiting the highest uncertainty is selected as the next best view:

$$\mathbf{v}^{*(t)} = \arg \max_{\mathbf{v} \in \mathcal{V}_{\text{cand}}^{(t)}} u(\mathbf{v}). \quad (15)$$

The corresponding physical measurement from $\mathbf{v}^{*(t)}$ is then acquired and appended to the training set for the subsequent phase of progressive optimization.

Table 1: Quantitative comparisons of 3D tomographic reconstruction and novel view synthesis across different view selection strategies under two protocols. Performance is evaluated in terms of 3D PSNR $\uparrow$ (PSNR $_3$, dB), 3D SSIM $\uparrow$ (SSIM $_3$), 2D PSNR $\uparrow$ (PSNR, dB), and 2D SSIM $\uparrow$ (SSIM). The **best** and **second-best** results are highlighted. *Random*, *Uniform*, and *FPS* denote random view selection, uniform sampling, and Farthest Point Sampling, respectively.

Category	Method	24-view				36-view			
		PSNR ₃	SSIM ₃	PSNR	SSIM	PSNR ₃	SSIM ₃	PSNR	SSIM
Synthetic Dataset									
Rule-based	Random	32.629	0.881	41.862	0.954	34.823	0.915	44.552	0.961
	Uniform	<u>33.562</u>	0.890	43.260	0.961	<u>35.877</u>	<u>0.921</u>	46.053	0.968
	FPS	33.508	<u>0.891</u>	43.081	0.957	35.367	0.919	45.386	0.964
2D-based	TOPIQ [9]	33.000	0.882	42.898	0.961	34.951	0.912	45.483	0.967
	MANIQA [75]	32.319	0.882	41.482	0.953	34.543	0.913	44.516	0.962
	MUSIQ [30]	33.083	0.886	42.675	0.957	35.030	0.916	45.133	0.965
3D-based	FisherRF [25]	33.347	0.887	<u>43.289</u>	<u>0.962</u>	35.551	0.919	<u>46.212</u>	<u>0.969</u>
	Ours	34.078	0.896	44.069	0.966	36.226	0.926	46.849	0.971
Real-World Dataset									
Rule-based	Random	36.112	0.903	42.992	0.985	36.765	0.921	44.965	0.989
	FPS	36.134	0.905	43.073	0.985	36.898	0.925	45.171	0.989
2D-based	TOPIQ [9]	36.208	0.906	43.121	0.986	37.167	0.928	45.644	0.990
	MANIQA [75]	36.222	0.906	<u>43.476</u>	0.986	37.187	<u>0.929</u>	45.644	<u>0.990</u>
	MUSIQ [30]	<u>36.257</u>	<u>0.906</u>	43.293	0.985	37.119	0.925	45.494	0.990
3D-based	FisherRF [25]	36.205	0.902	43.179	0.985	<u>37.258</u>	0.926	<u>46.022</u>	0.990
	Ours	36.399	0.909	43.546	<u>0.986</u>	37.480	0.932	46.046	0.991

4 Experiments

4.1 Experimental Settings

Datasets. Following the experimental settings of R²-Gaussian [80], we evaluate our method on both synthetic and real-world datasets. The synthetic benchmark comprises 15 CT volumes aggregated from various public datasets [4, 33, 58, 68]. For real-world evaluation, we utilize three cases from the FIPS dataset [65], reconstructing pseudo-ground-truth (pseudo-GT) volumes via FDK [16]. We adopt an object-centric hemispherical acquisition trajectory rather than a conventional circular orbit. The candidate viewpoint pool for active selection consists of 448 poses uniformly sampled across two concentric hemispheres at different radii, with all optical axes oriented toward the volume center. For novel view synthesis evaluation, 103 test viewpoints are randomly sampled from a single hemisphere per scene. All GT and pseudo-GT projections are simulated using DiffDRR [22].

Protocols. We initialize the Gaussian primitives with 50k random points and train the radiative GS model for 30k iterations per scene under the default configuration. Following FisherRF [25], active view selection is performed at progressively increasing iteration intervals. To ensure stable optimization, we schedule the final view selection prior to the termination of the GS densification




Fig. 3: Visual comparisons of the reconstructed 3D volumes using different view selection strategies. The 3D PSNR↑ (dB) for each scene is displayed at the top-left corner of the corresponding image. Our approach consistently achieves the highest reconstruction quality and better preserves fine structural details.

phase. At each selection step, a single optimal view is acquired and appended to the training set. We evaluate our method and all baselines under two final view-budget protocols: $N_{\text{target}} = 24$ and 36 (inclusive of two initial views).

Baselines. Previous active view selection strategies generally fall into two paradigms: 2D-based approaches that evaluate rendered image quality, and 3D-based approaches that leverage model uncertainty. We benchmark our method against state-of-the-art baselines from both categories, alongside three standard heuristics: random selection, uniform sampling, and Farthest Point Sampling (FPS). **i.** 2D-based: We adopt three no-reference image quality assessment (IQA) metrics, MUSIQ [30], MANIQA [75], and TOPIQ [9], implemented via the Py-IQA [8] toolbox. **ii.** 3D-based: We compare against the prevailing uncertainty-

Table 2: Quantitative comparisons of 3D tomographic reconstruction on sparser-view protocols. Performance is evaluated in terms of 3D PSNR $\uparrow$ (PSNR $_3$, dB) and 3D SSIM $\uparrow$ (SSIM $_3$). The **best** and second-best results are highlighted.

Method	6-view		8-view		12-view		16-view	
	PSNR $_3$	SSIM $_3$	PSNR $_3$	SSIM $_3$	PSNR $_3$	SSIM $_3$	PSNR $_3$	SSIM $_3$
Random	24.58	0.711	25.14	0.729	28.54	0.805	29.98	0.836
FPS	25.29	0.742	26.85	<u>0.780</u>	29.10	0.822	30.50	0.846
Uniform	23.77	0.697	26.84	0.772	28.95	0.811	<u>31.16</u>	<u>0.851</u>
MUSIQ [30]	25.54	0.741	27.04	0.767	29.86	<u>0.824</u>	31.03	0.847
FisherRF [25]	<u>25.80</u>	<u>0.748</u>	<u>27.43</u>	0.779	<u>29.87</u>	0.823	30.79	0.846
Ours	25.93	0.758	27.90	0.786	30.29	0.830	31.30	0.855

driven baseline, FisherRF [25]. All baselines are re-implemented upon the radiative GS framework.

Implementation Details. Our pipeline is implemented in PyTorch [56] and optimized via Adam [32]. We set the ensemble size $N = 10$, perturbation parameters $\alpha = 10\%$ and $\beta = 0.5$, and adopt the learning rates and loss weights from R²-Gaussian [80]. All experiments are conducted on a single NVIDIA A40 GPU. For quantitative evaluation, we report the Peak Signal-to-Noise Ratio (PSNR) computed over the full 3D volume, and the Structural Similarity Index Measure (SSIM) [69] averaged across 2D slices in the axial, coronal, and sagittal planes.

4.2 Results

Tomographic Reconstruction. Table 1 presents the quantitative results of 3D tomographic reconstruction across various active view selection strategies. Our approach demonstrates superior performance in both synthetic and real-world scenarios. In terms of 3D PSNR, our method surpasses the second-best baseline by margins of up to 0.68 dB on synthetic data and 0.22 dB on real-world data. Notably, FisherRF underperforms FPS under the synthetic 24-view setting. This degradation stems from FisherRF’s diagonal approximation of the Fisher Information Matrix. In X-ray imaging, strongly coupled Gaussians along projection rays violate this assumption, yielding severely biased information gain estimates. Qualitatively, as illustrated in Fig. 3, our method yields the most visually accurate reconstructions, faithfully preserving fine structural details, significantly suppressing boundary artifacts and background noise. We also evaluated our approach on sparser-view settings (6/8/12/16 views). As shown in Tab. 2, ours consistently maintains superior performance.

Novel View Synthesis. We also evaluate our method’s performance on novel view synthesis using the synthetic dataset. The qualitative comparisons and quantitative results are reported in Fig. 4 and Tab. 1, respectively. Our approach achieves the best overall rendering quality under both protocols, demonstrating




Fig. 4: Visual comparisons of novel view synthesis results across different view selection strategies. The $\text{PSNR}\uparrow$ (dB) for each scene is displayed at the top-left corner of each image. Our approach achieves the highest rendering quality.

an improvement of up to 0.78 dB in PSNR. Visual comparisons indicate that our method better preserves local structural details, particularly in high-density areas. This improvement stems from our method’s ability to acquire more informative viewpoints, which effectively mitigates the needle-like artifacts and noise that typically appear along rays passing through dense regions, e.g., bones.

4.3 Ablation Study and Analysis

Component Analysis. To demonstrate the impact of our proposed structural variance-based view selection mechanism, we replace SSIM with L1 error and PSNR, respectively, and evaluate the performance under the 24-view protocol. As shown in the first section of Tab. 3, switching to L1 error or PSNR as the uncertainty metric leads to a significant performance drop. There are two primary reasons for this degradation. Firstly, L1 error and PSNR are pixel-wise, independently computed metrics based on absolute or mean-squared errors. They overlook the spatial arrangement and inter-dependencies of adjacent pixels, making them less sensitive to high-frequency topological changes, such as the structural discontinuities caused by perturbing needle-like artifacts. Secondly, due to the linear integral nature of X-ray imaging, perturbing the density of Gaussian primitives inevitably causes overall intensity shifts in the projection space. L1 and PSNR are highly susceptible to these absolute luminance fluctuations, which can easily overwhelm the subtle variance signals originating from genuine geometric

Table 3: Ablation study results on the uncertainty quantification metric, ensemble size (N), perturbed ratio (α), and density scaling amplitude (β). Performance is evaluated in terms of 3D PSNR $\uparrow$ (PSNR $_3$, dB) and 3D SSIM $\uparrow$ (SSIM $_3$). Within each section, the **best** and second-best results are highlighted. $\dagger$ indicates the default setting.

Setting	PSNR $_3 \uparrow$	SSIM $_3 \uparrow$	Setting	PSNR $_3 \uparrow$	SSIM $_3 \uparrow$
L1	<u>33.644</u>	<u>0.894</u>	$\alpha = 5\%$	33.680	0.893
PSNR	33.390	0.889	$\alpha = 10\%^\dagger$	34.078	<u>0.896</u>
SSIM †	34.078	0.896	$\alpha = 15\%$	<u>33.938</u>	0.896
$N = 5$	33.952	0.895	$\alpha = 20\%$	33.589	0.892
$N = 10^\dagger$	34.078	<u>0.896</u>	$\beta = 0.1$	34.040	<u>0.896</u>
$N = 15$	<u>34.034</u>	0.896	$\beta = 0.3$	<u>34.057</u>	0.897
$N = 20$	33.865	0.895	$\beta = 0.5^\dagger$	34.078	0.896
$N = 40$	33.670	0.893	$\beta = 1.0$	33.338	0.890

ambiguities. In contrast, SSIM inherently incorporates luminance and contrast normalization, effectively decoupling absolute intensity shifts from structural information. This enables the SSIM variance to act as a more robust and precise indicator of true epistemic uncertainty and geometric disagreement.

Parameter Analysis. We evaluate the impact of ensemble size (N), ratio of perturbed Gaussian primitives (α), and density scaling amplitude (β) on reconstruction quality, as detailed in the last three sections of Tab. 3. Empirically, we observe that an ensemble size of $N = 10$ or 15 yields the optimal trade-off. Increasing N to 40 leads to a performance degradation. We attribute this to the over-smoothing of the structural disagreement signal: A larger sample size dilutes the impact of extreme catastrophic structural failures caused by perturbations, thereby reducing the discriminative contrast of the SSIM variance across views. A moderate size N successfully preserves the sharpness of the uncertainty landscape, allowing the active agent to accurately penalize geometric degeneracies.

The ratio α controls the proportion of Gaussians subjected to stochastic scaling, with 10% yielding the best results. Both excessively low and high ratios deteriorate the performance. When α is too small, the perturbation is overly restricted to background noise or void space, failing to include the primitives that constitute geometric degeneracies, resulting in a flat structural disagreement landscape. Conversely, a higher ratio indiscriminately extends the perturbations into well-constrained, high-confidence anatomical solids (e.g., bones). Due to the linear integral nature of X-ray projection, modulating these high-density structures induces massive, non-informative global variations that mask the subtle uncertainty signals from ambiguous regions.

For the scaling amplitude β , where the density scaling factor is sampled from $\text{Uniform}(-\beta, \beta)$, both overly conservative and excessively aggressive perturbations degrade the final reconstruction quality. A small β injects insufficient noise to disrupt the fragile geometric equilibrium of overfitting artifacts, providing weak guidance for view selection. Conversely, an excessively large β introduces

Table 4: Ablation study results on different perturbations. Performance is evaluated in terms of 3D PSNR $\uparrow$ (PSNR $_3$, dB) and 3D SSIM $\uparrow$ (SSIM $_3$). The **best** and second-best results are highlighted. † indicates our approach.

Views	Dropout		Position Jittering		PJ + SS		Stochastic Scaling †	
	PSNR $_3$	SSIM $_3$	PSNR $_3$	SSIM $_3$	PSNR $_3$	SSIM $_3$	PSNR $_3$	SSIM $_3$
24	32.718	0.883	<u>34.061</u>	0.895	33.886	<u>0.896</u>	34.078	0.896
36	34.607	0.924	36.112	0.925	<u>36.200</u>	<u>0.925</u>	36.226	0.926

severe out-of-distribution intensity shifts that completely distort local geometry. This structural collapse causes ubiquitous SSIM degradation across all views, diminishing the relative variance contrast and blinding the acquisition function.

Perturbation Strategy Analysis. To demonstrate the effectiveness of our perturbation strategy (Stochastic Scaling), we further compare it with several alternative perturbations. As shown in Tab. 4, our method achieves the best results. Dropout and Position Jittering tend to mislead the estimation of uncertain regions, thus leading to suboptimal view selections.

5 Conclusion

We present a novel active view selection framework for sparse-view CT reconstruction utilizing radiative Gaussian Splatting. To overcome the theoretical flaws of existing active learning in X-ray imaging, we introduce a physics-aware active view selection method based on *Perturbed Gaussian Ensemble*. By injecting stochastic density perturbations into under-constrained, low-density primitives, our approach exposes structural vulnerabilities quantified via projection-space SSIM variance. This mechanism successfully targets and suppresses artifacts. Extensive evaluations confirm that our method yields superior volumetric reconstructions compared to existing baselines. This work bridges the gap between active learning and explicit radiative fields, advancing the practical deployment of 3DGS in dose-sensitive clinical and industrial settings.

Acknowledgements

This research is supported by the National Artificial Intelligence Research Resource (NAIRR) Pilot Awards NAIRR250199, NAIRR260019, and NAIRR260077, the AMD University Program’s AI & HPC Cluster, NVIDIA Academic Grant Program, and Lambda’s Research Grant. Computational resources are also provided by Delta and DeltaAI at the National Center for Supercomputing Applications (NCSA) through ACCESS allocations CIS250012, CIS250816, and CIS251188.

References

1. Adler, J., Öktem, O.: Learned primal-dual reconstruction. *Transactions on Medical Imaging* **37**(6), 1322–1332 (2018)
2. Andersen, A.H., Kak, A.C.: Simultaneous algebraic reconstruction technique (sart): a superior implementation of the art algorithm. *Ultrasonic imaging* **6**(1), 81–94 (1984)
3. Anirudh, R., Kim, H., Thiagarajan, J.J., Mohan, K.A., Champley, K., Bremer, T.: Lose the views: Limited angle ct reconstruction via implicit sinogram completion. In: *CVPR*. pp. 6343–6352 (2018)
4. Armato III, S.G., McLennan, G., Bidaut, L., McNitt-Gray, M.F., Meyer, C.R., Reeves, A.P., Zhao, B., Aberle, D.R., Henschke, C.I., Hoffman, E.A., et al.: The lung image database consortium (lidc) and image database resource initiative (idri): a completed reference database of lung nodules on ct scans. *Medical physics* **38**(2), 915–931 (2011)
5. Cai, Y., Liang, Y., Wang, J., Wang, A., Zhang, Y., Yang, X., Zhou, Z., Yuille, A.: Radiative gaussian splatting for efficient x-ray novel view synthesis. In: *ECCV*. pp. 283–299. Springer (2024)
6. Cai, Y., Wang, J., Yuille, A., Zhou, Z., Wang, A.: Structure-aware sparse-view x-ray 3d reconstruction. In: *CVPR*. pp. 11174–11183 (2024)
7. Cao, W., Zhang, H., Tian, F., Wu, Y., Li, Y., Wang, S., Yu, N., Liu, Y.: Free-Orbit4D: Training-free arbitrary camera redirection for monocular videos via foreground-complete 4D reconstruction. In: *ACM SIGGRAPH Conference Papers* (2026)
8. Chen, C., Mo, J.: IQA-PyTorch: Pytorch toolbox for image quality assessment. [Online]. Available: <https://github.com/chaofengc/IQA-PyTorch> (2022)
9. Chen, C., Mo, J., Hou, J., Wu, H., Liao, L., Sun, W., Yan, Q., Lin, W.: Topiq: A top-down approach from semantics to distortions for image quality assessment. *TIP* **33**, 2404–2418 (2024)
10. Chen, L., Zhan, H., Chen, K., Xu, X., Yan, Q., Cai, C., Xu, Y.: Activegamer: Active gaussian mapping through efficient rendering. In: *CVPR*. pp. 16486–16497 (2025)
11. Chen, X., Li, Q., Wang, T., Xue, T., Pang, J.: Gennbv: Generalizable next-best-view policy for active 3d reconstruction. In: *CVPR*. pp. 16436–16445 (2024)
12. Chung, H., Ryu, D., McCann, M.T., Klasky, M.L., Ye, J.C.: Solving 3d inverse problems using pre-trained 2d diffusion models. In: *CVPR*. pp. 22542–22551 (2023)
13. Connolly, C.: The determination of next best views. In: *ICRA*. vol. 2, pp. 432–435. IEEE (1985)
14. Dhami, H., Sharma, V.D., Tokekar, P.: Pred-nbv: Prediction-guided next-best-view planning for 3d object reconstruction. In: *IROS*. pp. 7149–7154. IEEE (2023)
15. Duan, R., Chen, J., Kortylewski, A., Yuille, A., Liu, Y.: Prompt-based exemplar super-compression and regeneration for class-incremental learning. In: *BMVC* (2025)
16. Feldkamp, L.A., Davis, L.C., Kress, J.W.: Practical cone-beam algorithm. *Journal of the Optical Society of America A* **1**(6), 612–619 (1984)
17. Feng, Z., Zhan, H., Chen, Z., Yan, Q., Xu, X., Cai, C., Li, B., Zhu, Q., Xu, Y.: Naruto: Neural active reconstruction from uncertain target observations. In: *CVPR*. pp. 21572–21583 (2024)
18. Fischer, T., Liu, Y., Jesslen, A., Ahmed, N., Kaushik, P., Wang, A., Yuille, A.L., Kortylewski, A., Ilg, E.: inemo: Incremental neural mesh models for robust class-incremental learning. In: *ECCV*. pp. 357–374 (2024)

19. Gao, Z., Planche, B., Zheng, M., Chen, X., Chen, T., Wu, Z.: Ddgs-ct: Direction-disentangled gaussian splatting for realistic volume rendering. *NeurIPS* **37**, 39281–39302 (2024)
20. Ghani, M.U., Karl, W.C.: Deep learning-based sinogram completion for low-dose ct. In: *IVMSP*. pp. 1–5. IEEE (2018)
21. Goli, L., Reading, C., Sellán, S., Jacobson, A., Tagliasacchi, A.: Bayes’ rays: Uncertainty quantification for neural radiance fields. In: *CVPR*. pp. 20061–20070 (2024)
22. Gopalakrishnan, V., Golland, P.: Fast auto-differentiable digitally reconstructed radiographs for solving inverse problems in intraoperative imaging. In: *Workshop on Clinical Image-Based Procedures*. pp. 1–11. Springer (2022)
23. Huang, H., Wu, Y., Deng, C., Gao, G., Gu, M., Liu, Y.S.: Fatesgs: Fast and accurate sparse-view surface reconstruction using gaussian splatting with depth-feature consistency. In: *AAAI*. pp. 3644–3652 (2025)
24. Huang, H., Wu, Y., Zhou, J., Gao, G., Gu, M., Liu, Y.S.: Neusurf: On-surface priors for neural surface reconstruction from sparse input views. In: *AAAI*. pp. 2312–2320 (2024)
25. Jiang, W., Lei, B., Daniilidis, K.: Fisherrf: Active view selection and mapping with radiance fields using fisher information. In: *ECCV*. pp. 422–440 (2024)
26. Jin, K.H., McCann, M.T., Froustey, E., Unser, M.: Deep convolutional neural network for inverse problems in imaging. *TIP* **26**(9), 4509–4522 (2017)
27. Jin, L., Chen, X., Rückin, J., Popović, M.: Neu-nbv: Next best view planning using uncertainty estimation in image-based neural rendering. In: *IROS*. pp. 11305–11312 (2023)
28. Jin, L., Zhong, X., Pan, Y., Behley, J., Stachniss, C., Popović, M.: Activegs: Active scene reconstruction using gaussian splatting. *IEEE Robotics and Automation Letters* (2025)
29. Kak, A.C., Slaney, M.: Principles of computerized tomographic imaging. SIAM (2001)
30. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: Musiq: Multi-scale image quality transformer. In: *ICCV*. pp. 5148–5157 (2021)
31. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., et al.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
32. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014)
33. Klacansky, P.: Open scivis datasets (December 2017), <https://klacansky.com/open-scivis-datasets/>
34. Lee, S., Chung, H., Park, M., Park, J., Ryu, W.S., Ye, J.C.: Improving 3d imaging with pre-trained perpendicular 2d diffusion models. In: *ICCV*. pp. 10710–10720 (2023)
35. Li, J., Zhang, J., Bai, X., Zheng, J., Ning, X., Zhou, J., Gu, L.: Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In: *CVPR*. pp. 20775–20785 (2024)
36. Li, Y., Fu, X., Li, H., Zhao, S., Jin, R., Zhou, S.K.: 3dgr-ct: Sparse-view ct reconstruction with a 3d gaussian representation. *Medical Image Analysis* **103**, 103585 (2025)
37. Li, Y., Chen, X., Li, N.: Online optimal control with linear dynamics and predictions: Algorithms and regret analysis. *Advances in Neural Information Processing Systems* **32** (2019)
38. Li, Y., Das, S., Li, N.: Online optimal control with affine constraints. In: *AAAI*. pp. 8527–8537 (2021)

39. Li, Y., Qu, G., Li, N.: Online optimization with predictions and switching costs: Fast algorithms and the fundamental limit. *IEEE Transactions on Automatic Control* **66**(10), 4761–4768 (2020)
40. Li, Y., Kuang, Z., Li, T., Hao, Q., Yan, Z., Zhou, G., Zhang, S.: Activesplat: High-fidelity scene reconstruction through active gaussian splatting. *IEEE Robotics and Automation Letters* (2025)
41. Lin, Y., Luo, Z., Zhao, W., Li, X.: Learning deep intensity field for extremely sparse-view cbct reconstruction. In: *MICCAI*. pp. 13–23 (2023)
42. Lin, Y., Yang, J., Wang, H., Ding, X., Zhao, W., Li, X.: C²rv: Cross-regional and cross-view learning for sparse-view cbct reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 11205–11214 (June 2024)
43. Liu, J., Anirudh, R., Thiagarajan, J.J., He, S., Mohan, K.A., Kamilov, U.S., Kim, H.: Dolce: A model-based probabilistic diffusion framework for limited-angle ct reconstruction. In: *ICCV*. pp. 10498–10508 (2023)
44. Liu, Y., Li, Y., Schiele, B., Sun, Q.: Online hyperparameter optimization for class-incremental learning. In: *AAAI*. pp. 8906–8913 (2023)
45. Liu, Y., Li, Y., Schiele, B., Sun, Q.: Wakening past concepts without past data: Class-incremental learning from online placebos. In: *WACV*. pp. 2226–2235 (2024)
46. Liu, Y., Schiele, B., Sun, Q.: Adaptive aggregation networks for class-incremental learning. In: *CVPR*. pp. 2544–2553 (2021)
47. Liu, Y., Schiele, B., Sun, Q.: Rmm: Reinforced memory management for class-incremental learning. *NeurIPS* **34**, 3478–3490 (2021)
48. Liu, Y., Schiele, B., Vedaldi, A., Rupprecht, C.: Continual detection transformer for incremental object detection. In: *CVPR*. pp. 23799–23808 (2023)
49. Liu, Y., Su, Y., Liu, A.A., Schiele, B., Sun, Q.: Mnemonics training: Multi-class incremental learning without forgetting. In: *CVPR*. pp. 12245–12254 (2020)
50. Liu, Z., Bicer, T., Kettimuthu, R., Gursoy, D., De Carlo, F., Foster, I.: Tomogan: low-dose synchrotron x-ray tomography with generative adversarial networks: discussion. *Journal of the Optical Society of America A* **37**(3), 422–434 (2020)
51. Liu, Z., Zha, R., Zhao, H., Li, H., Cui, Z.: 4drgs: 4d radiative gaussian splatting for efficient 3d vessel reconstruction from sparse-view dynamic dsa images. In: *International Conference on Information Processing in Medical Imaging*. pp. 361–374. Springer (2025)
52. Luo, Z., Liu, Y., Schiele, B., Sun, Q.: Class-incremental exemplar compression for class-incremental learning. In: *CVPR*. pp. 11371–11380 (2023)
53. Manglos, S.H., Gagne, G.M., Krol, A., Thomas, F.D., Narayanaswamy, R.: Transmission maximum-likelihood reconstruction with ordered subsets for cone beam ct. *Physics in Medicine & Biology* **40**(7), 1225–1241 (1995)
54. 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)
55. Pan, X., Lai, Z., Song, S., Huang, G.: Activenerf: Learning where to see with uncertainty estimation. In: *ECCV*. pp. 230–246 (2022)
56. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *NeurIPS* **32** (2019)
57. Ran, Y., Zeng, J., He, S., Chen, J., Li, L., Chen, Y., Lee, G., Ye, Q.: Neurar: Neural uncertainty for autonomous 3d reconstruction with implicit neural representations. *IEEE Robotics and Automation Letters* **8**(2), 1125–1132 (2023)

58. Roth, H., Farag, A., Turkbey, E.B., Lu, L., Liu, J., Summers, R.M.: Data from pancreas-ct (2016)
59. Rückert, D., Wang, Y., Li, R., Idoughi, R., Heidrich, W.: Neat: Neural adaptive tomography. *TOG* **41**(4), 1–13 (2022)
60. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: nonlinear phenomena* **60**(1-4), 259–268 (1992)
61. Sauer, K., Bouman, C.: A local update strategy for iterative reconstruction from projections. *IEEE Transactions on Signal Processing* **41**(2), 534–548 (2002)
62. Scott, W.R., Roth, G., Rivest, J.F.: View planning for automated three-dimensional object reconstruction and inspection. *CSUR* **35**(1), 64–96 (2003)
63. Shen, L., Pauly, J., Xing, L.: Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *TNNLS* **35**(1), 770–782 (2022)
64. Sidky, E.Y., Pan, X.: Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology* **53**(17), 4777–4807 (2008)
65. Society, T.F.I.P.: X-ray tomographic datasets (2024), <https://fips.fi/category/open-datasets/x-ray-tomographic-datasets/>
66. Sunderhauf, N., Abou-Chakra, J., Miller, D., et al.: Density-aware nerf ensembles: Quantifying predictive uncertainty in neural radiance fields. In: *Proceedings of the 2023 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 9370–9376. Institute of Electrical and Electronics Engineers Inc. (2023)
67. Vasconcelos, F., He, B., Singh, N., Teh, Y.W.: Uncertainr: Uncertainty quantification of end-to-end implicit neural representations for computed tomography. *arXiv preprint arXiv:2202.10847* (2022)
68. Verboven, P., Dequeker, B., He, J., Pieters, M., Pols, L., Tempelaere, A., Van Doorselaer, L., Van Cauteren, H., Verma, U., Xiao, H., Nicolai, B.: www.x-plant.org - the ct database of plant organs (2022-05-24)
69. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *TIP* **13**(4), 600–612 (2004)
70. Wilson, J., Almeida, M., Mahajan, S., Labrie, M., Ghaffari, M., Ghasemalazadeh, O., Sun, M., Kuo, C.H., Sen, A.: Pop-gs: Next best view in 3d-gaussian splatting with p-optimality. In: *CVPR*. pp. 3646–3655 (2025)
71. Wu, Y., Huang, H., Zhang, W., Deng, C., Gao, G., Gu, M., Liu, Y.S.: Sparis: Neural implicit surface reconstruction of indoor scenes from sparse views. In: *AAAI*. pp. 8514–8522 (2025)
72. Xiao, W., Cruz, R.S., Ahmedt-Aristizabal, D., Salvado, O., Fookes, C., Lebrat, L.: Nerf director: Revisiting view selection in neural volume rendering. In: *CVPR*. pp. 20742–20751 (2024)
73. Xie, Y., Cai, Y., Zhang, Y., Yang, L., Pan, J.: Gauss-mi: Gaussian splatting shannon mutual information for active 3d reconstruction. *arXiv preprint arXiv:2504.21067* (2025)
74. Xue, S., Dill, J., Mathur, P., Dellaert, F., Tsiotra, P., Xu, D.: Neural visibility field for uncertainty-driven active mapping. In: *CVPR*. pp. 18122–18132 (2024)
75. Yang, S., Wu, T., Shi, S., Lao, S., Gong, Y., Cao, M., Wang, J., Yang, Y.: Maniqa: Multi-dimension attention network for no-reference image quality assessment. In: *CVPR*. pp. 1191–1200 (2022)
76. Ying, X., Guo, H., Ma, K., Wu, J., Weng, Z., Zheng, Y.: X2ct-gan: reconstructing ct from biplanar x-rays with generative adversarial networks. In: *CVPR*. pp. 10619–10628 (2019)

77. Yu, L., Zou, Y., Sidky, E.Y., Pelizzari, C.A., Munro, P., Pan, X.: Region of interest reconstruction from truncated data in circular cone-beam ct. *Transactions on Medical Imaging* **25**(7), 869–881 (2006)
78. Yu, W., Cai, Y., Zha, R., Fan, Z., Li, C., Yuan, Y.: X2-gaussian: 4d radiative gaussian splatting for continuous-time tomographic reconstruction. In: *ICCV*. pp. 24728–24738 (2025)
79. Zang, G., Idoughi, R., Li, R., Wonka, P., Heidrich, W.: Intratomo: self-supervised learning-based tomography via sinogram synthesis and prediction. In: *ICCV*. pp. 1960–1970 (2021)
80. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: R²-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. In: *Advances in Neural Information Processing Systems (NeurIPS)* (2024)
81. Zha, R., Zhang, Y., Li, H.: Naf: neural attenuation fields for sparse-view cbct reconstruction. In: *MICCAI*. pp. 442–452 (2022)
82. Zhang, G., Zha, R., He, H., Liang, Y., Yuille, A., Li, H., Cai, Y.: X-lrm: X-ray large reconstruction model for extremely sparse-view computed tomography recovery in one second. *arXiv preprint arXiv:2503.06382* (2025)
83. Zhang, J., Li, J., Yu, X., Huang, L., Gu, L., Zheng, J., Bai, X.: Cor-gs: sparse-view 3d gaussian splatting via co-regularization. In: *ECCV*. pp. 335–352 (2024)
84. Zhang, Y., Li, X., Chen, H., Yuille, A.L., Liu, Y., Zhou, Z.: Continual learning for abdominal multi-organ and tumor segmentation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 35–45. Springer (2023)
85. Zhao, C., Wang, X., Zhang, T., Javed, S., Salzmann, M.: Self-ensembling gaussian splatting for few-shot novel view synthesis. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4940–4950 (2025)
86. Zhu, Z., Fan, Z., Jiang, Y., Wang, Z.: Fsgs: Real-time few-shot view synthesis using gaussian splatting. In: *ECCV*. pp. 145–163. Springer (2024)
87. Zhu, Z., Gong, Y., Xiao, Y., Liu, Y., Hoiem, D.: How to teach large multimodal models new skills? In: *ECCV* (2026)
88. Zidane, A., Shimshoni, I.: Activenaf: A novel nerf-based approach for low-dose ct image reconstruction through active learning. *Physica Medica* **135**, 104997 (2025)