

R3CON: Radiance-Field-Free Active Reconstruction via Renderability

Xiaofeng Jin¹, Matteo Frosi¹, Yiran Guo¹, and Matteo Matteucci¹

Politecnico di Milano, Milan, Italy

{xiaofeng.jin, matteo.frosi, yiran.guo, matteo.matteucci}@polimi.it

Abstract. In active reconstruction, an embodied agent must decide where to look next to efficiently acquire views that support high-quality novel-view rendering. Recent work on active view planning for neural rendering largely derives next-best-view (NBV) criteria by backpropagating through radiance fields or estimating information entropy over 3D Gaussian primitives. While effective, these strategies tightly couple view selection to heavy, representation-specific mechanisms and fail to account for the computational and resource constraints required for lightweight online deployment. In this paper, we revisit active reconstruction from a renderability-centric perspective. We propose R3CON, a radiance-fields-free active reconstruction framework that induces an implicit, pose-conditioned renderability field over $SE(3)$ from a lightweight voxel map. Our formulation aggregates per-voxel online observation statistics into a unified scalar renderability score that is cheap to update and can be queried in closed form at arbitrary candidate viewpoints in milliseconds, without requiring gradients or radiance-field training. This renderability field is strongly correlated with image-space reconstruction error, naturally guiding NBV selection. We further introduce a panoramic extension that estimates omnidirectional (360°) view utility to accelerate candidate evaluation. In the standard indoor Replica dataset, R3CON achieves more uniform novel-view quality and higher 3D Gaussian splatting (3DGS) reconstruction accuracy than recent active GS baselines with matched view and time budgets. .

Keywords: Active view planning · Information gain · 3D Gaussian splatting

1 Introduction

Recent advances in 3D Gaussian splatting (3DGS) [17] enable high-fidelity novel view synthesis and dense reconstruction at real-time frame rates [17, 24], making Gaussian-based maps attractive for online robotic perception. However, 3DGS quality depends strongly on view density and coverage: sparse or poorly distributed inputs leave regions weakly constrained, leading to artifacts in novel

Project page: <https://github.com/jkff00/R3CON>




Fig. 1: Overview of renderability. Left: green cameras denote training views and black dots denote novel test views; we visualize RGB renderings, per-view renderability, and panoramic renderability. Colored boxes highlight local regions whose novel-view artifacts vary with viewpoint and are well captured by renderability. Right: distributions over 227 novel views of LPIPS, SSIM, PSNR, and our renderability, showing strong alignment with image-space reconstruction quality.

views. Therefore, active view selection during exploration is essential to reduce redundant data while still achieving high-quality reconstruction [1, 4, 39]. We study online next-best-view (NBV) selection for 3DGS with an on-board RGB-D camera, where an agent incrementally collects observations and must decide where to look next to maximize downstream rendering quality.

A key challenge in online active reconstruction is to score candidate views under strict latency and memory budgets while providing reliable improvements in reconstruction quality. Most existing methods estimate information gain from the internal state of an online-optimized radiance representation, including NeRF-style fields and 3DGS: they compute uncertainty- or information-theoretic scores from model parameters, gradients, or training losses, and select the next view by maximizing these scores [12, 13, 38]. This model-state dependence is problematic in online settings. Maintaining and repeatedly querying a large radiance representation is expensive on resource-limited platforms, and optimization-state signals can be unreliable before convergence. Recent sensitivity analysis [7] on 3DGS shows that loss-based second-order signals become stable only after the representation has sufficiently converged, making early-stage view utility estimates noisy and potentially leading to myopic selection.

We decouple view utility from transient optimization state and express information gain in terms of how well view-dependent radiance is constrained by past observations. As Fig. 1 shows, we represent this constraint with an implicit renderability field derived directly from lightweight directional observation statistics aggregated online. Motivated by this, we propose R3CON, a radiance-fields-free active reconstruction framework that maintains a voxel-statistics map and queries a closed-form renderability score at arbitrary candidate viewpoints in milliseconds. We further introduce a panoramic, pose-level formulation that estimates omnidirectional (360°) view utility to accelerate candidate evaluation.

On the Replica dataset, the proposed method improves active 3DGS reconstruction over recent baselines under matched view and time budgets.

Our contributions are: (1) A radiance-fields-free, closed-form view-utility formulation based on observation-statistics-driven renderability, enabling stable online scoring without relying on transient model optimization state. (2) A lightweight voxel-statistics map that supports millisecond-level NBV evaluation at arbitrary candidate viewpoints. (3) A 360° panoramic, pose-level utility and extensive Replica experiments demonstrating improved active 3DGS reconstruction under matched budgets.

2 Related Works

2.1 Geometry-driven Active Reconstruction.

NBV planning is a classic topic in online exploration and active mapping, where a robot incrementally builds a probabilistic occupancy map and selects viewpoints to reduce geometric uncertainty [27]. Representative strategies include frontier-based exploration [40] and information-gain objectives that score actions by expected entropy reduction [2, 31]. These ideas have also been extended to active volumetric 3D reconstruction by computing information gain on voxel or octree maps [8, 9]. While effective for geometric coverage and unknown-space discovery, such criteria are only indirectly tied to photometric fidelity, motivating rendering-aware NBV when neural rendering becomes the reconstruction target.

2.2 Online Photometric Mapping.

Neural rendering makes novel-view quality depend on the density and directional distribution of captured views [23]. Recent online reconstruction backbones make it practical to optimize radiance representations incrementally from streaming frames: rather than retraining on all past data, they refine the model by updating a sliding window and/or a small active parameter set, often with keyframe buffering or replay [15, 29, 34, 35, 44]. In parallel, 3DGS provides an explicit primitive-based radiance representation for real-time, high-fidelity rendering [17], enabling online 3DGS mapping pipelines such as SplaTAM and MonoGS [16, 22] as well as further variants for challenging settings [26, 33, 43].

2.3 Rendering-driven Active Reconstruction.

Building on online photometric mapping, rendering-driven NBV replaces geometric maps with an online photometric reconstruction state and scores candidate poses by information gain or uncertainty reduction on the evolving model. Accordingly, most rendering-driven NBV methods derive view utility from internal optimization signals of the evolving photometric model. For radiance fields, representative approaches estimate information gain or uncertainty in parameter space, like Fisher information or related criteria [5, 12, 25]; for 3DGS, utilities




Fig. 2: Renderability in IBR and radiance-field rendering. IBR estimates view feasibility from a coarse geometric proxy and source-view overlap, while 3DGS depends on how well primitive appearance is constrained across viewing directions.

are often computed from primitive-level confidence, probabilistic uncertainty, or mutual-information-style objectives [3, 13, 20, 38]. Other pipelines couple view selection with online mapping and planning [10, 11, 21] or use learned ranking signals for fast candidate filtering [37]. Despite differences in formulation, a common thread is the reliance on repeatedly maintaining and querying a nontrivial optimization state (parameters, gradients, or losses), which is costly under online latency and memory constraints and can be unreliable before convergence; notably, PUP-3DGS [7] shows that loss-based second-order sensitivity becomes meaningful mainly after sufficient optimization.

2.4 Renderability-Centric Active Reconstruction.

To mitigate the limitations of model-state-driven information gain, we adopt a renderability-centric view, characterizing view utility by how reliably a candidate viewpoint can be rendered given past observations. The previous work introduced renderability in image-based rendering using a coarse geometric proxy and observation statistics to estimate novel-view synthesis reliability for viewpoint or trajectory selection [41]. The following work adapted this metric to 3DGS to guide pseudo-view sampling under sparse inputs [14]. Existing renderability models, however, are largely grounded in image-based rendering and do not directly characterize uncertainty in radiance representations. Motivated by analyses of uncertainty sources in 3DGS [18], we reinterpret renderability through an uncertainty lens and extend it to an online formulation enabling efficient view selection, yielding lightweight active photometric reconstruction.

3 Technical Approach

3.1 Overview

We propose an observation-statistics-driven renderability utility for millisecond online NBV scoring without relying on transient reconstruction optimization state. First, we recap the notion of renderability from the image-based rendering (IBR) perspective and contrast it with renderability in view-dependent radiance representations. We then derive an analytic, primitive-level renderability

formulation and show how it can be computed online from lightweight observation statistics. Finally, we integrate the resulting renderability score into an active reconstruction pipeline for efficient NBV selection, including a panoramic pose-level evaluation for fast view-direction planning.

3.2 Preliminaries

Renderability was introduced by [41] as a pre-evaluation metric for IBR. By comparing a target viewpoint against past observations, as illustrated in Fig. 2, renderability essentially measures how well the source pixels used for synthesis are aligned with the target view. When the scene is purely Lambertian, IBR becomes less sensitive to viewpoint changes, and the dependence on viewpoint differences is reduced.

For radiance representations, the meaning of renderability changes. Radiance fields are fundamentally view-dependent, so sufficient directional observations are still required even in weakly reflective regions. From this perspective, renderability is determined by how well local primitives can fit view-dependent appearance at the query viewing direction, given the history of sampled rays.

3.3 Renderability in Radiance Representations

We analyze renderability for a single primitive via its spherical-harmonic (SH) fitting behavior. Following uncertainty sources in 3DGS, we characterize SH fitting quality with two factors: observation bias and noise (Fig. 3). Observation bias measures the angular discrepancy between the query direction and historical viewing directions; zero bias indicates identical directions and minimal fitting loss. Observation noise summarizes the dispersion of historical radiance observations and reflects material-dependent view variability, which is typically higher for non-Lambertian regions. Accordingly, for each primitive p_i , the observation acquired from a viewpoint s (camera pose $T^{(s)}$) is modeled as:

$$\underbrace{\mathbf{z}_i^{(s)} - \mathbf{z}_i^*}_{\text{rendering deviation}} = \mathcal{F}\left(b_i^{(s)}, \varepsilon_i^{(s)}\right), \quad (1)$$

where $\mathbf{z}_i^*$ denotes the true (unknown) sample value, $b_i^{(s)}$ represents viewpoint-dependent bias, and $\varepsilon_i^{(s)}$ is the noise from all historical observations of p_i .

For spherical-harmonic fitting, the same sampling bias can have markedly different effects depending on where the query direction lies relative to the historical samples. From least-squares error propagation and analyses of noise amplification under ill-conditioned sampling [28], fitting error can grow rapidly when the query falls in an extrapolation regime. To capture this effect, we follow standard spherical geometry that treat the spherical convex hull of historical viewing directions as the valid interpolation domain.

As shown in Fig. 3, we project historical directions to the tangent plane at the query direction and classify extrapolation if the query falls outside the 2D bounding box of projected samples (a practical local convex-hull approximation).




Fig. 3: Renderability factors. Renderability combines directional bias, appearance noise, and a resolution term. v denotes viewing direction and z the observed radiance; orange points indicate the (unknown) radiance at a novel direction. Noise captures view-dependent variability, and bias measures how far the query direction deviates from the directional support of past observations. Extrapolation is detected by projecting past directions to the query tangent plane and checking whether the query lies outside their support region.

For bias term, We denote the extrapolation test result by a binary factor $\kappa_i^{(s)} \in \{1, 2\}$, where $\kappa_i^{(s)} = 1$ for interpolation and $\kappa_i^{(s)} = 2$ for extrapolation. let $\{\mathbf{r}_i^{(t)}\}_{t=0}^{n_i}$ denote the unit viewing directions of all past observations, and let $\mathbf{r}_i^{(s)}$ be the unit viewing direction of the current candidate view. Computing the angular deviation between the current and the closest historical direction:

$$\cos \theta_i^{(s)} = \max_{t \in \{0, \dots, n_i\}} \langle \mathbf{r}_i^{(t)}, \mathbf{r}_i^{(s)} \rangle \quad (2)$$

$$b_i^{(s)} = \left(\cos \theta_i^{(s)} \right)^{\kappa_i^{(s)}}. \quad (3)$$

For the noise term, it is quantified by the mutual differences among its historical observations $\{\mathbf{z}_i^{(t)}\}_{t=1}^{n_i}$. This captures how much the appearance of p_i varies across views. Concretely, the average pairwise discrepancy as

$$\delta_i = 1 - \frac{1}{n_i(n_i - 1)} \sum_{\substack{t, t'=1 \\ t \neq t'}}^{n_i} \|\mathbf{z}_i^{(t)} - \mathbf{z}_i^{(t')}\|^2, \quad (4)$$

where a smaller δ_i indicates higher observation noise.

Then convert this discrepancy into a renderability-related term

$$\varepsilon_i^{(s)} = \delta_i^{\kappa_i^{(s)}(1-b_i^{(s)})}. \quad (5)$$

This exponential form has two useful properties. First, the base $\delta_i \in [0, 1]$ reflects noise consistency: larger δ_i yields a higher value, meaning more reliable historical observations. Second, the exponent controls the decay rate. When the angular deviation is small ($b_i^{(s)} \approx 1$), the exponent approaches zero and $\delta_i^{1-b_i^{(s)}} \approx$

1, indicating almost no rendering loss. Larger deviations increase the exponent and accelerate the decay, penalizing biased viewpoints more strongly.

As illustrated in Fig. 3, observations are discretized in the pixel domain, so changes in image resolution introduce additional differences in rendering quality. We therefore introduce an auxiliary resolution term to account for range-dependent variations in observation quality. Imaging resolution is inversely proportional to the viewing distance. For a primitive at unbiased depth D , the projected pixel footprint satisfies

$$\text{pixel size} \propto \frac{1}{D}. \quad (6)$$

We define a resolution-related auxiliary term using the depth-based resolution factor

$$\rho_i^{(s)} = \frac{1}{D_i^{(s)}}, \quad \rho_i^{\max} = \max_t \frac{1}{D_i^{(t)}}, \quad (7)$$

which reflects the relative pixel footprint of primitive p_i under different viewpoints. The corresponding resolution gain is then

$$\gamma_i^{(s)} = \min \left(1, \frac{\rho_i^{\max}}{\rho_i^{(s)}} \right)^{1 - \delta_i b_i^{(s)}}. \quad (8)$$

indicating whether the current view demands higher effective resolution than the best historical observation. If $\rho_i^{(s)} \leq \rho_i^{\max}$, then $\gamma_i^{(s)} = 1$; otherwise, $\gamma_i^{(s)} < 1$ indicates lower reconstruction reliability due to insufficient historical resolution.

The exponent $1 - \delta_i b_i^{(s)}$ controls the decay rate of this reliability score: primitives with stable appearance (large δ_i) and well-supported viewpoints (large $b_i^{(s)}$) require less additional resolution, yielding weaker attenuation, whereas primitives with high appearance variation or biased viewpoints demand higher resolution and therefore induce stronger decay.

We finally define renderability as the product of three normalized factors: viewpoint bias $b_i^{(s)}$, appearance stability $\varepsilon_i^{(s)}$, and resolution gain $\gamma_i^{(s)}$,

$$R_i^{(s)} = \mathcal{F}(b_i^{(s)}, \varepsilon_i^{(s)}) \gamma_i^{(s)} = b_i^{(s)} \varepsilon_i^{(s)} \gamma_i^{(s)}, \quad (9)$$

where $b_i^{(s)}, \varepsilon_i^{(s)}, \gamma_i^{(s)} \in [0, 1]$. The multiplicative form yields $R_i^{(s)} \in [0, 1]$ and assigns high renderability only when the primitive is well-supported directionally, exhibits consistent appearance, and is observed at sufficient resolution.

3.4 Online Renderability

Computing renderability naively requires aggregating statistics over all past observations. This is impractical online for two reasons: (i) the bias term would have to retain all historical viewing directions on the continuous sphere $\mathbb{S}^2$, incurring

Algorithm 1 Online bias term with Fibonacci-binned directions

Global (once): choose N bins by θ_{res} and precompute $\{\mathbf{q}_k\}_{k=1}^N$
 $A \leftarrow 2\pi(1 - \cos(\theta_{\text{res}}/2)); \quad N \leftarrow \lceil 4\pi/A \rceil; \quad \varphi_g \leftarrow (1 + \sqrt{5})/2$
for $k = 1..N$ **do**
 $z_k \leftarrow 1 - 2^{\frac{k+0.5}{N}}; \quad r_k \leftarrow \sqrt{1 - z_k^2}; \quad \phi_k \leftarrow \frac{2\pi k}{\varphi_g}; \quad \mathbf{q}_k \leftarrow (r_k \cos \phi_k, r_k \sin \phi_k, z_k)$
end for
State (per primitive p_i): visited-bin mask $\mathcal{I}_i \in \{0, 1\}^N$
Update with a novel view $\mathbf{r}_i^{(t)}$: $k^* \leftarrow \arg \max_k \langle \mathbf{q}_k, \mathbf{r}_i^{(t)} \rangle; \quad \mathcal{I}_i[k^*] \leftarrow 1.$
Query for direction $\mathbf{r}_i^{(s)}$: $b_i^{(s)} \leftarrow \max_{k: \mathcal{I}_i[k]=1} \langle \mathbf{q}_k, \mathbf{r}_i^{(s)} \rangle; \quad \text{return } b_i^{(s)}.$

Algorithm 2 Online noise estimation with Welford updates

State (per-voxel p_i): $n_i, \boldsymbol{\mu}_i \in \mathbb{R}^3, \mathbf{M}_i \in \mathbb{R}^{3 \times 3}$ (init: $n_i=0, \boldsymbol{\mu}_i=\mathbf{0}, \mathbf{M}_i=\mathbf{0}$)
Update($\mathbf{z}_i^{(t)} \in \mathbb{R}^3$):
 $n_i \leftarrow n_i + 1; \quad \Delta \leftarrow \mathbf{z}_i^{(t)} - \boldsymbol{\mu}_i; \quad \boldsymbol{\mu}_i \leftarrow \boldsymbol{\mu}_i + \Delta/n_i; \quad \mathbf{M}_i \leftarrow \mathbf{M}_i + (\mathbf{z}_i^{(t)} - \boldsymbol{\mu}_i)\Delta^\top.$
Query (assume $n_i \geq 2$): $\text{trs}_i \leftarrow \text{tr}(\mathbf{M}_i)/(n_i - 1); \quad \delta_i \leftarrow 1 - \alpha\sqrt{\text{trs}_i}, \quad \alpha = \sqrt{2/1.5}.$

prohibitive memory and increasing query cost; and (ii) different primitives accumulate different numbers of observations, which prevents efficient vectorization and further amplifies computation. The noise term suffers from the same history dependence, so in incremental exploration the update cost grows quickly with the number of frames. We therefore adopt a principled approximation that makes renderability computation independent of the number of stored observations.

In Algorithm 1, we discretize the unit sphere with a Fibonacci lattice [6] $\{\mathbf{q}_k\}_{k=1}^N$. Each primitive maintains a fixed-size visited-bin mask over these uniformly distributed bins, enabling efficient bias tracking and evaluation with constant memory cost and observation count-independent, fully vectorizable computation.

For the noise term, as shown in Algorithm 2, each voxel maintains fixed-size Welford moments to track color covariance online. We map the covariance trace to a normalized consistency score $\delta_i \in [0, 1]$ using α .

3.5 Active Reconstruction via Renderability

We integrate the proposed renderability into online active reconstruction by building on Active-GS [13] and replacing its view utility with our renderability-based information gain. As shown in Fig. 4, given a posed RGB-D stream, we maintain two maps. The first is a coarse occupancy map used for exploration and safety, following Active-GS. The second is a lightweight voxel-statistics map used for renderability computation. The voxel-statistics map is downsampled to a resolution of 5 cm, and each voxel corresponds to one local primitive that stores the online statistics. To evaluate candidate views efficiently, we extend the rasterizer in ODGS [19] for any queried pose, which can quickly compute the visible voxel set and a panoramic depth proxy.




Fig. 4: System overview. Posed RGB-D updates a coarse occupancy grid for exploration and a fine voxel-statistics map for renderability. Candidate views are scored by exploration gain and renderability, with panoramic direction selection accelerated by sphere discretization and FoV aggregation.

For a candidate viewpoint s with position $t^{(s)}$, we compute the visible unexplored grid cells $\mathcal{G}^{(s)}$, and the visible set on the voxel-statistics map for renderability evaluation, $\mathcal{V}^{(s)}$. The exploration utility $U_G(s) = |\mathcal{G}^{(s)}|$ is simply the count of grid cells. While rendering utility as:

$$U_R(s) = \sum_{v \in \mathcal{V}^{(s)}} (1 - R_v^{(s)}). \quad (10)$$

The overall pose score is a weighted combination, as in Active-GS, to prioritize exploration at early stages:

$$U_{\text{view}}(s) = \lambda_1 U_G(s) + U_R(s) - \lambda_2 U_{\text{path}}(s), \quad (11)$$

where λ_1 and λ_2 are weights for the exploration term and the path-cost term in Active-GS, respectively, and $U_{\text{path}}(s)$ is the travel cost to reach s computed on the exploration grid.

We further select the viewing direction using the Fibonacci-sphere discretization (Sec. 3.4). We form a unified 3D point set by combining centers of $\mathcal{G}^{(s)}$ and $\mathcal{V}^{(s)}$, project each point to a unit ray direction on $\mathbb{S}^2$, and assign it to the nearest Fibonacci bin k . For each bin k , we precompute the FoV-covered bin set $\mathcal{N}(k)$ when the optical axis aligns with k , and compute a direction score by summing per-bin scores over $\mathcal{N}(k)$. The bin with the maximum aggregated score is selected as the final viewing direction.

4 Experiments

4.1 Experimental Setup

Across all scenes, the Fibonacci sphere used for voxel renderability updates contains 64 samples. For panoramic direction selection, the Fibonacci discretization




Fig. 5: Evaluation pipeline. Each method runs online to generate a trajectory and collect posed RGB and scene pointcloud. A 3DGS model is then trained from the recorded data under matched budgets. For evaluation, test views are sampled from free-space grid locations, each rendered with six canonical viewing directions, to measure scene-level generalization on held-out novel views and reconstruction metrics.

uses a target central angle of 10° . The voxel size is set to 5 cm to represent local primitives. All other hyperparameters follow Active-GS.

All methods are evaluated in Habitat-Sim [30]. The RGB image resolution is 512×512 , and the camera FoV is $[60^\circ, 60^\circ]$. Experiments are run on a workstation with an NVIDIA Quadro RTX 6000 GPU (24 GB) under Ubuntu 20.04.

Active-GS [13] and ActiveGAMER [3] are used as strong 3DGS-based baselines. Both support full 6-DoF camera motion and achieve competitive performance on Replica. ActiveSplat [21] tailored to ground navigation is not included, as their motion assumptions do not match our free-flying setting. In our tables, Ours replaces the scoring module of Active-GS with our proposed information gain, and Ours-Pano further adds the panoramic extension.

4.2 Replica-Dense Testset

The goal of reconstruction is high-fidelity rendering from arbitrary free-space viewpoints. To this end, we construct a dense test set, Replica-Dense, on nine Replica [32] indoor scenes. For each scene, collision-free free space is approximated and sampled by a 3D grid with spacing $1.0\text{ m} \times 1.0\text{ m} \times 0.75\text{ m}$ in (x, y, z) . At each valid grid location, six test cameras are placed with optical axes aligned to $\pm x$, $\pm y$, and $\pm z$. All methods are evaluated on the same Replica-Dense set.

As shown in Fig. 5, all methods are treated purely as data collectors. Each method starts from the same initial pose, runs its policy until a fixed budget, and records RGB images, depth-derived point clouds, and camera poses. After data collection, the planner state is discarded, and a vanilla 3DGS model is trained from scratch on the collected views using the same 3DGS implementation and hyperparameters for all methods, without method-specific regularization or augmentation. This decouples planning from reconstruction, so differences on Replica-Dense can be attributed to view selection.

Each method is evaluated under two budget regimes: a time budget of 300 s and a view budget of 300 keyframes. In each regime, we subsample the collected trajectory at $\{60, 120, 180, 240, 300\}$ (seconds or views), yielding 2×5 acquisition

Method	Metric	Hotel0	Off0	Off1	Off2	Off3	Off4	Room0	Room1	Room2
ActiveGAMER	SSIM	0.8338	0.8859	0.8032	0.7468	0.7339	0.7396	0.6846	0.8008	0.7169
	PSNR	20.51	26.16	25.74	15.95	16.45	16.02	14.86	19.75	14.55
	LPIPS	0.2738	0.1791	0.2549	0.3658	0.3727	0.3530	0.4369	0.3397	0.3995
Active-GS	SSIM	0.8867	0.9138	0.8957	0.9006	0.8792	0.8982	0.8672	0.8778	0.9060
	PSNR	24.60	29.87	31.39	26.91	25.38	26.45	23.01	24.51	25.07
	LPIPS	0.2134	0.1513	<u>0.1849</u>	0.1936	0.2209	0.2012	0.2597	0.2508	0.2049
Ours	SSIM	0.9124	0.9493	0.9338	0.9162	0.9140	0.9269	<u>0.9053</u>	<u>0.9149</u>	0.9426
	PSNR	26.93	<u>33.95</u>	37.14	<u>29.34</u>	<u>28.93</u>	31.16	<u>27.85</u>	<u>28.85</u>	31.24
	LPIPS	0.1824	<u>0.1058</u>	0.1581	<u>0.1720</u>	<u>0.1826</u>	0.1549	<u>0.2109</u>	<u>0.2053</u>	0.1569
Ours-Pano	SSIM	<u>0.8975</u>	<u>0.9480</u>	<u>0.9298</u>	0.9305	0.9225	<u>0.9184</u>	0.9159	0.9201	<u>0.9314</u>
	PSNR	<u>25.95</u>	34.17	<u>36.70</u>	30.74	30.91	<u>30.34</u>	29.41	29.43	<u>29.23</u>
	LPIPS	<u>0.1992</u>	0.1017	0.1581	0.1553	0.1668	<u>0.1590</u>	0.1964	0.1963	<u>0.1711</u>

Table 1: Per-scene results on Replica-Dense under a fixed time budget. Higher is better for PSNR/SSIM, and lower is better for LPIPS. Best values are highlighted in bold with color, and second-best values are underlined.




Fig. 6: Quantitative view quality and exploration efficiency. Each group contains scenes Hotel0 and Office0. The left two pictures is the final novel-view rendering quality versus time. The right two images shows the number of collected images versus time. Under the same time budget, Ours achieves the highest data acquisition efficiency.

sets per scene. For each set, we train a vanilla 3DGS model and render all Replica-Dense test views. We report PSNR, SSIM [36], and LPIPS [42] together with the corresponding acquisition cost, and also record peak GPU memory usage per method and scene to quantify the trade-off between quality and resource consumption.

4.3 Rendering Quality under Budgets

With a fixed exploration time of 300s, Tab. 1 shows that Ours and Ours-Pano achieve the best or second-best PSNR/SSIM and the lowest LPIPS on nearly all scenes, improving PSNR by about 3–5 dB over Active-GS (and by a larger margin over ActiveGAMER). Fig. 6 further shows that our PSNR rises faster and saturates higher on Hotel0 and Office0; for the same wall-clock budget, we also collect more keyframes, indicating that renderability-based scoring is both more informative and cheaper to evaluate.

With a fixed budget of 300 keyframes, all methods are trained on the same number of images, so quality depends solely on viewpoint selection. Under this

Method	Metric	Hotel0	Off0	Off1	Off2	Off3	Off4	Room0	Room1	Room2
ActiveGAMER	SSIM	0.8831	0.8921	0.8816	0.8735	0.8752	0.8726	0.8796	0.8807	0.8796
	PSNR	23.63	26.93	31.32	22.79	24.61	25.20	24.42	24.42	23.26
	LPIPS	0.2102	0.1639	0.1851	0.2250	0.2226	0.2122	0.2389	0.2428	0.2313
Active-GS	SSIM	<u>0.9085</u>	0.9314	0.8782	0.9142	<u>0.9220</u>	0.9179	0.9037	0.9139	0.9298
	PSNR	26.36	31.65	29.65	27.06	<u>29.57</u>	28.92	26.68	27.64	27.40
	LPIPS	0.1854	0.1244	0.1925	0.1781	<u>0.1743</u>	0.1680	0.2173	0.2080	0.1752
Ours	SSIM	0.9084	0.9532	0.9263	<u>0.9243</u>	0.9054	0.9282	<u>0.9200</u>	<u>0.9227</u>	0.9357
	PSNR	<u>26.72</u>	34.69	<u>36.01</u>	<u>29.93</u>	27.69	30.95	<u>29.60</u>	<u>29.84</u>	<u>29.69</u>
	LPIPS	<u>0.1847</u>	0.0987	0.1599	<u>0.1627</u>	0.1934	0.1549	<u>0.1948</u>	<u>0.1946</u>	0.1636
Ours-Pano	SSIM	0.9119	<u>0.9510</u>	<u>0.9248</u>	0.9277	0.9240	<u>0.9239</u>	0.9251	0.9249	<u>0.9345</u>
	PSNR	27.18	<u>34.34</u>	36.50	30.70	31.26	<u>30.77</u>	30.41	30.39	29.92
	LPIPS	0.1825	<u>0.1021</u>	<u>0.1603</u>	0.1596	0.1685	<u>0.1553</u>	0.1877	0.1913	<u>0.1654</u>

Table 2: Per-scene results on Replica-Dense under a fixed view budget. Higher is better for PSNR/SSIM, and lower is better for LPIPS. Best values are highlighted in bold with color, and second-best values are underlined.




Fig. 7: Quantitative results under a view budget. Top row is PSNR versus view budget on four scenes (Office2/3/4 and Room0). Bottom row is PSNR distribution shifts on Office1 across methods (ActiveGAMER, Active-GS, Ours-Pano, and Ours) as the view budget increases.

setting, Tab. 2 shows that Ours consistently outperforms Active-GS and ActiveGAMER in PSNR/SSIM and LPIPS, validating the effectiveness of renderability-based information gain. Ours-Pano typically improves further, while in small scenes such as Office0 it matches the base version due to saturated coverage.

Fig. 7 summarizes performance across budgets: Ours and Ours-Pano remain above both baselines throughout, reflecting higher sample efficiency at low budgets and fewer redundant views at high budgets. The PSNR distribution is also noticeably tighter, with reduced low-quality tails compared to baselines. Qualitative results in Fig. 8 corroborate these trends, showing sharper textures, cleaner edges, and fewer floaters/ghosting; Ours-Pano further stabilizes quality across scenes.



Fig. 8: Qualitative comparison under a fixed view budget. The novel view shows that our method produces sharper details and fewer artifacts than prior methods, and the panoramic variant yields more stable reconstruction quality across all scenes. Red boxes highlight regions where differences are subtle.




Fig. 9: GPU memory usage across four scenes (Office2/3/4 and Room0). Each plot shows the peak GPU memory consumption of different methods as the number of collected views increases.




Fig. 10: Time cost of Ours-Pano. The figure shows, in Office2/3/4 and Room0, how the renderability update/query time changes as the number of keyframes increases.

4.4 Online Performance Analysis

Fig. 9 reports peak GPU memory versus the number of views on four scenes. Active-GS shows the steepest growth as it accumulates a heavy radiance field, whereas Ours and Ours-Pano remain below 300 MB across all scenes since the planner only maintains a lightweight voxel-statistics map.




Fig. 11: Degeneracy analysis of Ours. We visualize 3D heatmaps of camera viewpoints colored by acquisition ID for four methods. The top row is the state in Hotel0, while for Office3, Ours collapses into a local region with repeated, degenerate viewpoints.

As shown in Fig. 10, per-frame renderability update/query stays in the millisecond range and remains nearly constant as keyframes grow to hundreds, demonstrating observation-count-independent scalability that is critical for long-horizon exploration.

Fig. 11 visualizes camera distributions, comparing with the performance in Hotel0, our non-panoramic variant may degenerate into locally repeated viewpoints due to the trade-off between renderability gain and path cost under random local sampling. By scoring omnidirectional utility at each pose, Ours-Pano mitigates this local trapping and yields more uniform coverage, leading to more stable performance across scenes and budgets.

5 Conclusion and Limitations

In this work, we presented R3CON, a radiance-field-free active reconstruction framework that scores candidate views with an observation-statistics-driven renderability field. By decomposing renderability into directional bias, appearance noise, and resolution, and computing these terms online from a lightweight voxel-statistics map, our method enables closed-form, millisecond-level NBV evaluation without maintaining an online radiance field. A panoramic pose-level extension aggregates omnidirectional utility per pose and accelerates view-direction selection. On the Replica-Dense benchmark, R3CON consistently improves 3DGS reconstruction quality over recent active Gaussian-splatting baselines under matched time and view budgets, while using less GPU memory and compute, making it a promising plug-and-play module to pair with existing SLAM systems in real-world autonomous data acquisition.

A current limitation is that candidate viewpoints are sampled locally at random, which can be suboptimal in large or cluttered environments; future work

will couple renderability-based scoring with learned or longer-horizon planning strategies to generate more informative candidates.

References

1. Alsadik, B., Mahdi, H.A., Abdulateef, N.A.: A structured review and taxonomy of next-best-view strategies for 3d reconstruction. *ISPRS Open Journal of Photogrammetry and Remote Sensing* **17**, 100098 (2025). <https://doi.org/10.1016/j.ophoto.2025.100098>
2. Bourgault, F., Makarenko, A.A., Williams, S.B., Grocholsky, B., Durrant-Whyte, H.F.: Information based adaptive robotic exploration. In: *IEEE/RSJ international conference on intelligent robots and systems*. vol. 1, pp. 540–545. IEEE (2002)
3. Chen, L., Zhan, H., Chen, K., Xu, X., Yan, Q., Cai, C., Xu, Y.: Activegamer: Active gaussian mapping through efficient rendering. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 16486–16497 (2025)
4. Chen, X., Li, Q., Wang, T., Xue, T., Pang, J.: Gennbv: Generalizable next-best-view policy for active 3d reconstruction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16436–16445 (2024)
5. 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: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 21572–21583 (2024)
6. Álvaro González: Measurement of areas on a sphere using fibonacci and latitude–longitude lattices. *Mathematical Geosciences* **42**, 49–64 (2010)
7. 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*. pp. 5949–5958 (2025)
8. Hornung, A., Wurm, K.M., Bennewitz, M., Stachniss, C., Burgard, W.: Octomap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous robots* **34**(3), 189–206 (2013)
9. Isler, S., Sabzevari, R., Delmerico, J., Scaramuzza, D.: An information gain formulation for active volumetric 3d reconstruction. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*. pp. 3477–3484. IEEE (2016)
10. Jiang, W., Lei, B., Ashton, K., Daniilidis, K.: Ag-slam: Active gaussian splatting slam. *arXiv preprint arXiv:2410.17422* (2024). <https://doi.org/10.48550/arXiv.2410.17422>
11. Jiang, W., Lei, B., Ashton, K., Daniilidis, K.: Multimodal llm guided exploration and active mapping using fisher information. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 5392–5404 (2025)
12. Jiang, W., Lei, B., Daniilidis, K.: Fisherrf: Active view selection and mapping with radiance fields using fisher information. In: *European Conference on Computer Vision (ECCV)*. pp. 422–440 (2024). https://doi.org/10.1007/978-3-031-72624-8_24
13. Jin, L., Zhong, X., Pan, Y., Behley, J., Stachniss, C., Popović, M.: Activegs: Active scene reconstruction using gaussian splatting. *IEEE Robotics and Automation Letters* **10**(5), 4866–4873 (2025). <https://doi.org/10.1109/LRA.2025.3555149>
14. Jin, X., Fang, Y., Frosi, M., Ge, J., Xiao, J., Matteucci, M.: Rendering anywhere you see: Renderability field-guided gaussian splatting. In: *2025 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 7850–7857. IEEE (2025)

15. Johari, M.M., Carta, C., Fleuret, F.: Eslam: Efficient dense slam system based on hybrid representation of signed distance fields. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 17408–17419 (2023)
16. Keetha, N., Karhade, J., Jatavallabhula, K.M., Yang, G., Scherer, S., Ramanan, D., Luiten, J.: Splatam: Splat track & map 3d gaussians for dense rgb-d slam. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 21357–21366 (2024)
17. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
18. Klasson, M., Mereu, R., Kannala, J., Solin, A.: Sources of uncertainty in 3d scene reconstruction. In: *European Conference on Computer Vision*. pp. 271–289. Springer (2024)
19. Lee, S., Chung, J., Huh, J., Lee, K.M.: ODGS: 3D scene reconstruction from omnidirectional images with 3D Gaussian splatting. In: *Advances in Neural Information Processing Systems*. vol. 37 (2024). <https://doi.org/10.52202/079017-1818>
20. Li, Y., Li, Y., Lee, G.H.: Active3d: Active high-fidelity 3d reconstruction via hierarchical uncertainty quantification. *arXiv preprint arXiv:2511.20050* (2025). <https://doi.org/10.48550/arXiv.2511.20050>
21. 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* **10**(8), 8099–8106 (2025). <https://doi.org/10.1109/LRA.2025.3580331>
22. 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). <https://doi.org/10.1109/CVPR52733.2024.01708>
23. 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)
24. Niedermayr, S., Stumpfegger, J., Westermann, R.: Compressed 3d gaussian splatting for accelerated novel view synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10349–10358 (2024)
25. Pan, X., Lai, Z., Song, S., Huang, G.: Activenerf: Learning where to see with uncertainty estimation. In: *European Conference on Computer Vision*. pp. 230–246. Springer (2022)
26. Pan, Y., Zhong, X., Jin, L., Wiesmann, L., Popovic, M., Behley, J., Stachniss, C.: Pings: Gaussian splatting meets distance fields within a point-based implicit neural map. In: *Proceedings of Robotics: Science and Systems*. Los Angeles, CA, USA (Jun 2025). <https://doi.org/10.15607/RSS.2025.XXI.040>
27. Placed, J.A., Strader, J., Carrillo, H., Atanasov, N., Indelman, V., Carlone, L., Castellanos, J.A.: A survey on active simultaneous localization and mapping: State of the art and new frontiers. *IEEE Transactions on Robotics* **39**(3), 1686–1705 (2023)
28. Rauhut, H., Ward, R.: Sparse recovery for spherical harmonic expansions. *arXiv preprint arXiv:1102.4097* (2011). <https://doi.org/10.48550/arXiv.1102.4097>
29. Rosinol, A., Leonard, J.J., Carlone, L.: Nerf-slam: Real-time dense monocular slam with neural radiance fields. In: *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. pp. 3437–3444. IEEE (2023)
30. Savva, M., Kadian, A., Maksymets, O., Zhao, Y., Wijmans, E., Jain, B., Straub, J., Liu, J., Koltun, V., Malik, J., Parikh, D., Batra, D.: Habitat: A platform for

- embodied ai research. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9338–9346 (2019)
31. Stachniss, C., Grisetti, G., Burgard, W.: Information gain-based exploration using rao-blackwellized particle filters. In: Robotics: Science and systems. vol. 2, pp. 65–72 (2005)
 32. Straub, J., Whelan, T., Ma, L., Chen, Y., Wijmans, E., Green, S., Engel, J.J., Mur-Artal, R., Ren, C.Y., Verma, S., Clarkson, A., Yan, M., Budge, B., Yan, Y., Pan, X., Yon, J., Zou, Y., Leon, K., Carter, N., Briales, J., Gillingham, T., Mueggler, E., Pesqueira, L., Savva, M., Batra, D., Strasdat, H.M., De Nardi, R., Goesele, M., Lovegrove, S., Newcombe, R.A.: The Replica dataset: A digital replica of indoor spaces. arXiv preprint arXiv:1906.05797 (2019). <https://doi.org/10.48550/arXiv.1906.05797>
 33. Su, Y., Lin, C., Zhang, K., Yu, Z., Hou, C., Zhao, Z.: Gaus-slam: Dense rgb-d slam with gaussian surfels. IEEE Robotics and Automation Letters **11**(3), 3390–3397 (2026). <https://doi.org/10.1109/LRA.2026.3655276>
 34. Sucar, E., Liu, S., Ortiz, J., Davison, A.J.: imap: Implicit mapping and positioning in real-time. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6229–6238 (2021)
 35. Wang, H., Wang, J., Agapito, L.: Co-slam: Joint coordinate and sparse parametric encodings for neural real-time slam. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13293–13302 (2023)
 36. 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)
 37. Wang, Z., Bhalgat, Y., Li, R., Prisacariu, V.A.: Active view selector: Fast and accurate active view selection with cross reference image quality assessment. arXiv preprint arXiv:2506.19844 (2025). <https://doi.org/10.48550/arXiv.2506.19844>
 38. Xie, Y., Cai, Y., Zhang, Y., Yang, L., Pan, J.: GauSS-MI: Gaussian Splatting Shannon Mutual Information for Active 3D Reconstruction. In: Proceedings of Robotics: Science and Systems. Los Angeles, CA, USA (June 2025). <https://doi.org/10.15607/RSS.2025.XXI.030>
 39. Xue, S., Dill, J., Mathur, P., Dellaert, F., Tsiotra, P., Xu, D.: Neural visibility field for uncertainty-driven active mapping. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18122–18132 (2024)
 40. Yamauchi, B.: A frontier-based approach for autonomous exploration. In: Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'. pp. 146–151. IEEE (1997)
 41. Yi, Z., Xie, K., Lyu, J., Gong, M., Huang, H.: Where to render: Studying renderability for ibr of large-scale scenes. In: 2023 IEEE Conference Virtual Reality and 3D User Interfaces (VR). pp. 356–366. IEEE (2023)
 42. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 586–595 (2018)
 43. Zheng, J., Zhu, Z., Bieri, V., Pollefeys, M., Peng, S., Armeni, I.: Wildgs-slam: Monocular gaussian splatting slam in dynamic environments. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 11461–11471 (2025)
 44. Zhu, Z., Peng, S., Larsson, V., Xu, W., Bao, H., Cui, Z., Oswald, M.R., Pollefeys, M.: Nice-slam: Neural implicit scalable encoding for slam. In: Proceedings of the

IEEE/CVF conference on computer vision and pattern recognition. pp. 12786–12796 (2022)