FRONTIERGS: EXPLORATION-DRIVEN ACTIVE GAUSSIAN RECONSTRUCTION WITH EFFICIENT FISHER-BASED PLANNING

Anonymous authors
Paper under double-blind review

ABSTRACT

Active Gaussian reconstruction requires robots to explore unknown space while improving the reconstruction quality of observed surfaces. Existing methods may prioritize local reconstruction gains during viewpoint planning, potentially trapping robots in observed regions. Repeated Gaussian map optimization and viewpoint evaluation also require substantial computation. To address these problems, we present FrontierGS, an exploration-driven framework that organizes viewpoint planning around a main path toward unknown space and adopts Fisher information to estimate the geometric benefit of candidate viewpoints. For efficient estimation of reconstruction gains, we reuse the Fisher information computed from past observations, forming a persistent Fisher approximation that avoids recomputing historical information after map updates. To guide viewpoint selection, we accumulate the expected Fisher information of already selected viewpoints to predict which viewpoints near the main path can provide complementary reconstruction benefits. These predictions use previously computed Fisher information between map optimizations, which lead to accumulated prediction errors. We therefore introduce an adaptive refresh strategy that optimizes the map and recomputes candidate Fisher information. Experiments demonstrate that our method improves coverage and rendering quality in large scenes while reducing computational cost.

1 INTRODUCTION

Robots operating in unknown environments must explore and reconstruct their surroundings. Unlike passive reconstruction (Kerbl et al., 2023; Keetha et al., 2024) from predefined image sequences, active reconstruction selects viewpoints to balance the efficiency of data collection and reconstruction quality (Lluvia et al., 2021). However, redundant or insufficient observations can leave surfaces poorly constrained, making viewpoint selection critical for accurate reconstruction.

Early active mapping methods mainly optimize geometric coverage using occupancy or voxel representations (Hornung et al., 2013). Recent scene representations, including NeRFs and 3D Gaussian Splatting (Mildenhall et al., 2021; Kerbl et al., 2023), have enabled active methods (Jiang et al., 2024; Wilson et al., 2025; Feng et al., 2024) to select viewpoints by estimating uncertainty or information gain. Existing methods combine Gaussian maps with spatial representations, travel costs, or hierarchical planning for unknown environments (Jin et al., 2025; Li et al., 2025; Xu et al., 2025). However, reconstruction favors views of partially reconstructed surfaces, whereas exploration requires progress toward unknown space. Balancing these goals can lead to detours or repeated observations. Moreover, recomputing Fisher information over all past observations after each Gaussian map optimization becomes increasingly costly. Reducing the frequency of map optimization and Fisher recomputation saves computation, but information gains of candidate views become less accurate with accumulated observations, since they remain evaluated against the last optimized map.

To address these challenges, we present FrontierGS, an exploration-driven active Gaussian reconstruction framework that preserves the target frontier while selecting complementary viewpoints within a corridor around the exploration path, as shown in Figure 1. This organization prioritizes exploration progress while accommodating local reconstruction needs. To reduce the computational cost of viewpoint selection, we maintain persistent Fisher approximations by linearizing each frame

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

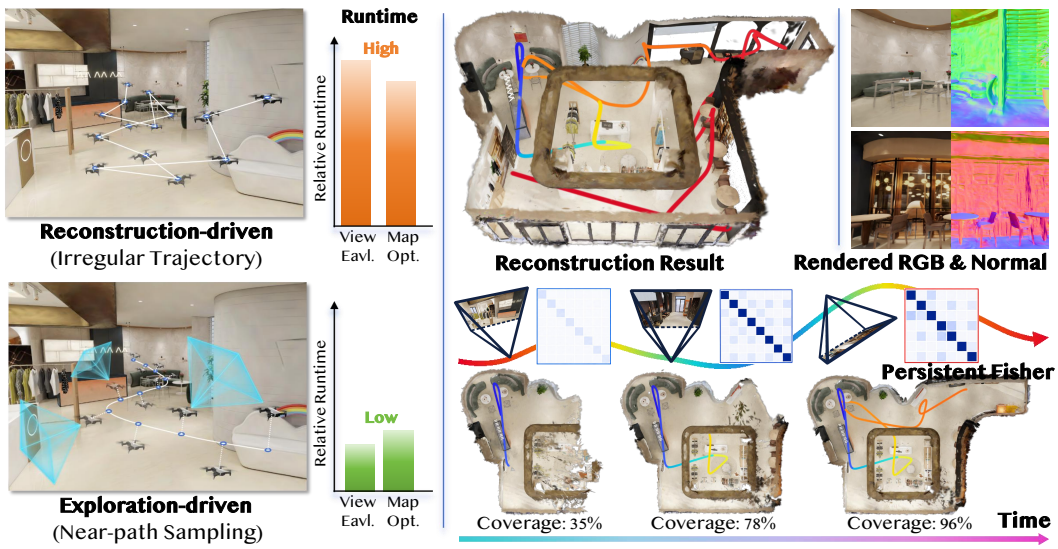

Figure 1: **Exploration-driven active Gaussian reconstruction**. We adopt an exploration-driven active reconstruction method that avoids detours caused by reconstruction-driven planning.

once and reusing its information across subsequent map optimizations, avoiding replay of the observation history. To select complementary views near the exploration path, we use virtual information updates to accumulate the information of selected views in a temporary state and evaluate each candidate's marginal reconstruction gain without re-optimizing the map. However, as these gains are evaluated against the last optimized map, they may become stale as new observations arrive. We therefore use an adaptive refresh strategy to trigger map optimization and Fisher recomputation once the estimated drift is large enough to affect path selection. Our contributions are threefold.

- We present an exploration-driven active Gaussian reconstruction framework that preserves the target frontier and global exploration progress while selecting complementary reconstruction viewpoints near the exploration path.
- We maintain persistent Fisher approximations by linearizing each new observation once for Fisher accumulation and retaining its contribution across map updates, avoiding Jacobian recomputation over historical frames.
- We develop an adaptive map refresh strategy that compares predicted score changes with the score gap between returned paths, while virtual information updates support conditional viewpoint planning between refreshes.

2 Related Work

Exploration-Based Active Mapping. Early active mapping explored unknown scenes with geometric representations by selecting viewpoints that maximize coverage or information criteria. Frontier-based methods (Yamauchi, 1997; Umari & Mukhopadhyay, 2017; Cieslewski et al., 2017) navigate toward the boundary between known and unknown space. Sampling methods select viewpoints using information gain, uncertainty reduction, or mutual information (Connolly, 1985; Hollinger & Sukhatme, 2014; Charrow et al., 2015). These methods typically operate on occupancy grid maps (Hornung et al., 2013), and later extensions optimize surface coverage through Monte Carlo volumetric integration (Guédon et al., 2022) or adopt hierarchical and topological abstractions (Cao et al., 2021) to improve planning efficiency in large environments. However, these methods measure only geometric coverage and completeness, ignoring color and texture, so the recovered maps are insufficient for high-fidelity novel-view rendering.

Active Scene Reconstruction. Radiance fields (Mildenhall et al., 2021) and 3D Gaussian Splatting (Kerbl et al., 2023) enable active reconstruction (Li et al., 2026b) with high-fidelity representa-

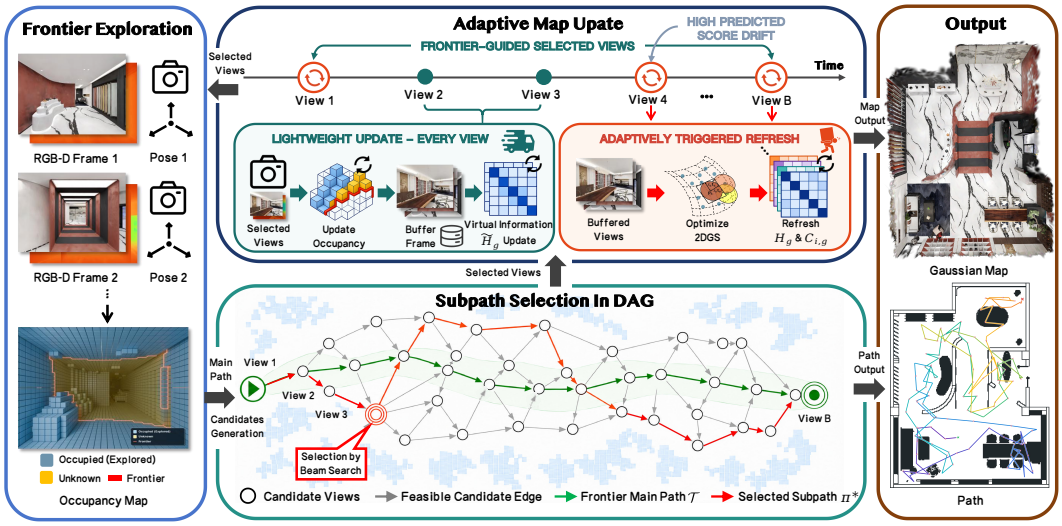

Figure 2: **Overview of FrontierGS.** Frontier exploration plans a route into unknown space. Candidate views along the route form a directed acyclic graph. Beam search selects subpaths by reconstruction gain, exploration coverage, and travel cost. The robot executes the selected view, buffers its RGB-D frame, and updates occupancy and the virtual information state. FrontierGS adaptively refreshes the map and reevaluates candidates when score drift may alter subpath selection.

tions, efficient rendering, and incremental updates. Fisher-information methods (Jiang et al., 2024; Wilson et al., 2025) evaluate viewpoints through parameter uncertainty, while mutual-information (Xie et al., 2025), visibility-based (Xue et al., 2026), and coverage-based methods (Jin et al., 2025; Chen et al., 2025; 2026) use alternative uncertainty or visibility measures. Hierarchical and long-horizon planning frameworks organize the search through topological abstractions (Li et al., 2025), Fisher-information gains (Xu et al., 2025), or imagined coverage maps (Li et al., 2026a). FisherRF (Jiang et al., 2024) and POp-GS (Wilson et al., 2025) reduce redundancy through batch information updates, with FisherRF also planning toward frontiers. However, batch updates alone neither avoid historical recomputation after map changes. FrontierGS reuses past Fisher estimates and decides when to optimize the map based on whether updated scores could change path selection, while virtual information updates support conditional viewpoint planning within the path corridor.

3 Method

At each step k, the robot captures an RGB-D frame $o_k = \{I_k, D_k, T_k, K\}$. Persistent Fisher approximations and virtual information updates support conditional subpath planning, while each adaptive refresh optimizes the Gaussian map with buffered frames and updates persistent and candidate Fisher information, as summarized in Figure 2.

3.1 2DGS Map Representation

Following 2DGS (Huang et al., 2024), we represent the 3D scene with planar Gaussian disks. Given camera pose T_k and intrinsics K, alpha blending renders color and depth at pixel u as follows:

$$\hat{I}_k(u) = \sum_g T_g(u)\alpha_g(u)c_g, \qquad \hat{D}_k(u) = \sum_g T_g(u)\alpha_g(u)d_g(u), \tag{1}$$

where $T_g(u) = \prod_{h<g}(1-\alpha_h(u))$ is the accumulated transmittance, and $\alpha_g(u)$, c_g, and $d_g(u)$ denote projected opacity, color, and depth. The map is optimized from RGB-D observations using

$$\mathcal{L}_k = \mathcal{L}_{\text{rgb}}(\hat{I}_k, I_k) + \lambda_{\text{depth}}\mathcal{L}_{\text{depth}}(\hat{D}_k, D_k), \tag{2}$$

where $\lambda_{\text{rgb}} = 1$ and $\lambda_{\text{depth}} = 0.5$. New Gaussians are added for newly observed regions, while Gaussians with negligible rendering contributions are pruned. Since actively selected viewpoints are sparse, every captured frame is used for map optimization without keyframe selection.

3.2 PERSISTENT FISHER INFORMATION

To capture parameter coupling within each Gaussian, we use geometric Gauss–Newton information as a local proxy for reconstruction uncertainty. We parameterize Gaussian g by $\theta_g = [\mu_g^\top, \log s_g^\top, \rho_g^\top]^\top \in \mathbb{R}^8$, where $\rho_g \in \mathbb{R}^3$ denotes rotation coordinates in a fixed parameterization shared across map states. Following POp-GS (Wilson et al., 2025), each 8×8 block retains within-Gaussian coupling while omitting correlations between Gaussians. For frame k, we compute

$$F_{k,g} = J_{k,g}^\top J_{k,g}, \tag{3}$$

where $J_{k,g}$ is the Jacobian of the RGB-D observation model with respect to θ_g, using the weighting in Eq. 2 while holding appearance and opacity fixed. Computing the accumulated geometric Fisher information $F_{\text{base},g}$ at the current map parameters requires replaying historical frames to recompute their Jacobians, with cost increasing as observations accumulate.

Incremental Fisher Accumulation. We maintain a persistent approximation $\widehat{F}_{\text{base},g}$ to $F_{\text{base},g}$ for each Gaussian. At each refresh, every newly buffered frame k is linearized once at the current map state and accumulated as

$$\widehat{F}_{\text{base},g} \leftarrow \widehat{F}_{\text{base},g} + F_{k,g}. \tag{4}$$

Historical contributions are not relinearized after later map updates and remain associated with Gaussian identities. Retained Gaussians keep their persistent Fisher approximations, newly added Gaussians receive zero blocks, split children receive equal shares of their parent's approximation, and approximations for removed Gaussians are discarded. Candidate Fisher increments affected by topology changes are recomputed at the next refresh. Because active reconstruction progressively shifts observations toward newly discovered regions, well-observed Gaussians typically change little, making their historical linearizations informative for planning.

Uncertainty Estimation. To measure the uncertainty of each Gaussian and the contribution of a candidate viewpoint, we regularize the persistent Fisher approximation to obtain the persistent information base H_g and cache the Fisher increment $C_{i,g}$ of candidate viewpoint i,

$$H_g = \widehat{F}_{\text{base},g} + \lambda I, \qquad C_{i,g} = J_{i,g}^\top J_{i,g}, \tag{5}$$

where $\lambda > 0$ ensures $H_g \succ 0$, and $J_{i,g}$ is evaluated from the current 2DGS rendering, with rendered depth fixed in the depth residual. The marginal reconstruction gain is

$$\Delta_g(i) = \log\det(H_g + C_{i,g}) - \log\det H_g. \tag{6}$$

For a fixed $C_{i,g}$, the gain decreases as information increases, favoring complementary observations.

3.3 SUBPATH SELECTION

Each main path connects consecutive target frontiers. Since map optimization and candidate Fisher recomputation after every observation are costly, we schedule these operations adaptively and use the virtual information state $\widetilde{H}_g$ for conditional planning between refreshes.

Candidate Graph. Candidate viewpoints are sampled at path anchors within a corridor around the main path $\mathcal{T}$, forming a directed acyclic graph ordered along $\mathcal{T}$. A feasible subpath $\pi = (p_{i_1}, \ldots, p_{i_K})$, with $i_1 < \cdots < i_K$, connects the current pose to the target frontier through reachable candidates. Edges carry motion costs c_{ij}, and beam search (Lowerre, 1976) selects a subpath.

After each observation, we update the occupancy map $\mathcal{M}$, remove unreachable candidates, and recompute the unknown voxels visible from the remaining candidates through known free space. If the main path becomes blocked, the subpath is discarded and replanned from the current pose. Otherwise, the main path is retained until the next refresh and candidate graph rebuild.

Conditional Information Gain. At each refresh, we optimize the map, update the persistent Fisher approximations from buffered frames, rebuild the candidate graph when needed, and cache candidate Fisher increments. The virtual information state is then reset as $\widetilde{H}_g = H_g$. Between refreshes, executing viewpoint i updates the state by

$$\widetilde{H}_g \leftarrow \widetilde{H}_g + C_{i,g}. \tag{7}$$

Algorithm 1 FrontierGS active reconstruction

Input: initial pose T_0, Gaussian map G, occupancy map M, persistent Fisher approximation $\hat{F}_{\text{base}}$, intrinsics K, corridor radius r_c, maximum subpath length $K_{\max}$, planning weights, refresh parameters, budget
Output: updated G, $\hat{F}_{\text{base}}$, M

$T \leftarrow T_0$, $\mathcal{T} \leftarrow \emptyset$, Buf $\leftarrow \emptyset$
while exploration is unfinished and budget remains **do**
  if $\mathcal{T} = \emptyset$ **then**
    if Buf is nonempty **then**
      $(G, \hat{F}_{\text{base}}) \leftarrow \text{MapFisherRefresh}(G, \hat{F}_{\text{base}}, \text{Buf})$; clear Buf
    $(V_{\text{main}}, \mathcal{T}) \leftarrow \text{FrontierExploration}(M, T)$
    if no reachable target frontier exists **then break**
    $\text{BuildPlanningState}(\mathcal{T}, M, G, \hat{F}_{\text{base}})$
  $(\pi^\star, m_t) \leftarrow \text{BeamSearch}(T)$ using Eq. 9 and current $\widetilde{H}_g$
  if fewer than two paths are returned **then** $m_t \leftarrow \infty$
  if $\text{AdaptiveRefresh}(\text{Buf}, m_t)$ **then**
    $(G, \hat{F}_{\text{base}}) \leftarrow \text{MapFisherRefresh}(G, \hat{F}_{\text{base}}, \text{Buf})$; clear Buf
    $\text{BuildPlanningState}(\mathcal{T}, M, G, \hat{F}_{\text{base}})$
    rerun $\text{BeamSearch}(T)$ to update $\pi^\star$
  if $\pi^\star = \emptyset$ **then** $T \leftarrow$ next collision-checked pose on $\mathcal{T}$
  else $T \leftarrow$ first viewpoint of $\pi^\star$
  capture $o_k = \{I_k, D_k, T, K\}$, append o_k to Buf, and update M
  if candidate i was executed **then** $\widetilde{H}_g \leftarrow \widetilde{H}_g + C_{i,g}$
  update candidate validity and visible unknown voxels
  if V_{main} is reached or $\mathcal{T}$ is blocked **then** $\mathcal{T} \leftarrow \emptyset$
if Buf is nonempty **then** $(G, \hat{F}_{\text{base}}) \leftarrow \text{MapFisherRefresh}(G, \hat{F}_{\text{base}}, \text{Buf})$
return G, $\hat{F}_{\text{base}}$, M

Unlike batch view selection, the virtual state persists across viewpoint selections, while each beam search branch keeps a separate state to evaluate gains conditioned on its path,

$$\Delta_R(i) = \frac{1}{N_{\text{gs}} D_{\text{geo}}} \sum_{g \in \text{vis}(i)} \left[\log\det\left(\widetilde{H}_g + C_{i,g}\right) - \log\det \widetilde{H}_g \right], \tag{8}$$

where $\text{vis}(i)$ denotes visible Gaussians, N_{gs} is the current Gaussian count, and $D_{\text{geo}} = 8$.

Subpath Selection. Within a corridor, beam search balances reconstruction, discovery, and detour.

$$\pi^\star = \arg\max_{\pi} \left[\bar{I}_R(\pi) + \lambda_{\text{disc}} \frac{I_{\text{disc}}(\pi)}{N_{\text{unknown}}} - \lambda_D \frac{\max\{0, L(\pi) - L_F\}}{L_F} \right]. \tag{9}$$

$\bar{I}_R(\pi)$ sums the marginal reconstruction gains in Eq. 8 along π, while $I_{\text{disc}}(\pi)$ counts distinct unknown voxels visible from its viewpoints and N_{unknown} is the number of unknown voxels. Each partial path tracks covered voxels to avoid double counting. $L(\pi)$ is the motion cost from the current pose through selected viewpoints to the target frontier, and L_F is the remaining main path cost. Search stops after $K_{\max}$ candidates or when no reachable forward candidate remains. Following (Tan et al., 2024), only the first viewpoint of the highest scoring path is executed before replanning.

Adaptive Map Refresh. Because virtual updates reuse cached Fisher increments computed on the last optimized map, their gain estimates become increasingly stale as observations accumulate. We therefore compare the predicted score change with m_t, the score gap between the two best paths returned by beam search. Let n_t be the number of buffered observations since the last refresh and $\widehat{a}_t$ the score change per observation estimated from previous refreshes. We estimate $\widehat{E}_t = n_t \widehat{a}_t$ and refresh when $2\widehat{E}_t \geq \beta m_t$ or $n_t \geq n_{\max}$, where $0 < \beta \leq 1$ controls refresh sensitivity and $n_{\max}$ limits the maximum interval. If no refresh is triggered, the current path is executed. Otherwise, we optimize the map, accumulate Fisher information from buffered frames without relinearizing historical contributions, recompute $C_{i,g}$, reset $\widetilde{H}_g = H_g$, and rerun beam search. Appendix A derives the sufficient condition that motivates this refresh rule.

Table 1: Comparison under fixed RGB-D observation budgets. *Recon. Cov.* and *Expl. Cov.* denote surface coverage on Replica and Matterport3D and exploration coverage on InteriorGS. *Ours (w/o corridor)* expands the sampling region, while *Ours (exploration only)* uses exploration views only. For each dataset, best, second-best, and third-best distinct values in the first five metrics are highlighted in red, orange, and yellow, respectively. All Ours variants share the same backend.

Method	Recon. Cov. / Expl. Cov.↑	Completion (m)↓	PSNR (dB)↑	SSIM↑	LPIPS↓	Path (m)
Replica						
FisherRF	0.658	0.166	17.44	0.645	0.447	6.56
ActiveSplat	0.947	0.026	27.51	0.809	0.230	12.40
ActiveGAMER	0.864	0.060	23.78	0.787	0.324	32.16
ActiveGS	0.966	**0.019**	**32.55**	**0.892**	**0.149**	45.31
Ours	**0.976**	**0.019**	31.68	0.878	0.195	47.92
Matterport3D						
FisherRF	0.807	0.148	16.16	0.613	0.655	59.06†
ActiveSplat	0.366	0.530	12.76	0.571	0.613	19.90†
ActiveGAMER	0.652	0.549	15.84	0.554	0.631	115.02
ActiveGS	0.853	**0.146**	22.02	0.748	0.425	148.50
Ours (w/o corridor)	0.901	0.154	21.31	0.754	0.419	346.19
Ours	**0.910**	0.157	**22.98**	**0.796**	**0.412**	145.22
InteriorGS						
FisherRF	0.552	0.472	12.89	0.602	0.582	38.46†
ActiveSplat	0.102	0.659	12.39	0.473	0.651	10.66†
ActiveGAMER	0.580	0.462	11.23	0.427	0.670	212.03†
ActiveGS	0.665	0.442	16.31	0.647	0.458	202.35†
Ours (exploration only)	0.807	0.198	17.64	0.590	0.510	163.59
Ours	**0.958**	**0.049**	**28.36**	**0.899**	**0.223**	310.26

† The planner frequently stalls or repeatedly visits already explored regions under the fixed observation budget, resulting in shorter traveled paths.

3.4 Frontier-Based Exploration

Frontier exploration (Yamauchi, 1997) supplies the main path $\mathcal{T}$. The robot updates the occupancy map using depth maps $\mathbf{D}_k$ and poses $\mathbf{T}_k$, identifying frontiers at boundaries between free and unknown space. The exploration layer selects a target frontier V_{main} and plans $\mathcal{T}$ toward it. The reconstruction layer preserves this target while selecting viewpoints within the path corridor using conditional information gain. Algorithm 1 summarizes the overall loop.

4 Experiments

4.1 Experimental Setup

Implementation details. We use a 2DGS mapper (Huang et al., 2024) with RGB-D observations generated by Habitat-GS (Xia et al., 2026). Planning uses a corridor radius $r_c = 1$ m, at most 20 candidates, beam width 4, and depth 5. We set $\lambda = 10^{-4}$, $\lambda_{\text{disc}} = 300$, $\lambda_D = 0.1$, $\beta = 0.8$, and $n_{\max} = 6$, with 20 map optimization iterations per selected viewpoint. Occupancy is updated after every observation, while the Gaussian map and candidate Fisher information are refreshed adaptively. Further details are provided in Appendix B.

Datasets and baselines. We evaluate on eight Replica (Straub et al., 2019), five Matterport3D (Chang et al., 2017), and six InteriorGS (Team et al., 2025) scenes, with budgets of 100, 250, and 500 RGB-D observations, respectively. Matterport3D and InteriorGS both represent large indoor environments. We compare with ActiveSplat (Li et al., 2025), FisherRF (Jiang et al., 2024), ActiveGAMER (Chen et al., 2025), and ActiveGS (Jin et al., 2025), retaining their original reconstruction settings under the same observation settings and budgets. HGS-Planner and POp-GS are omitted

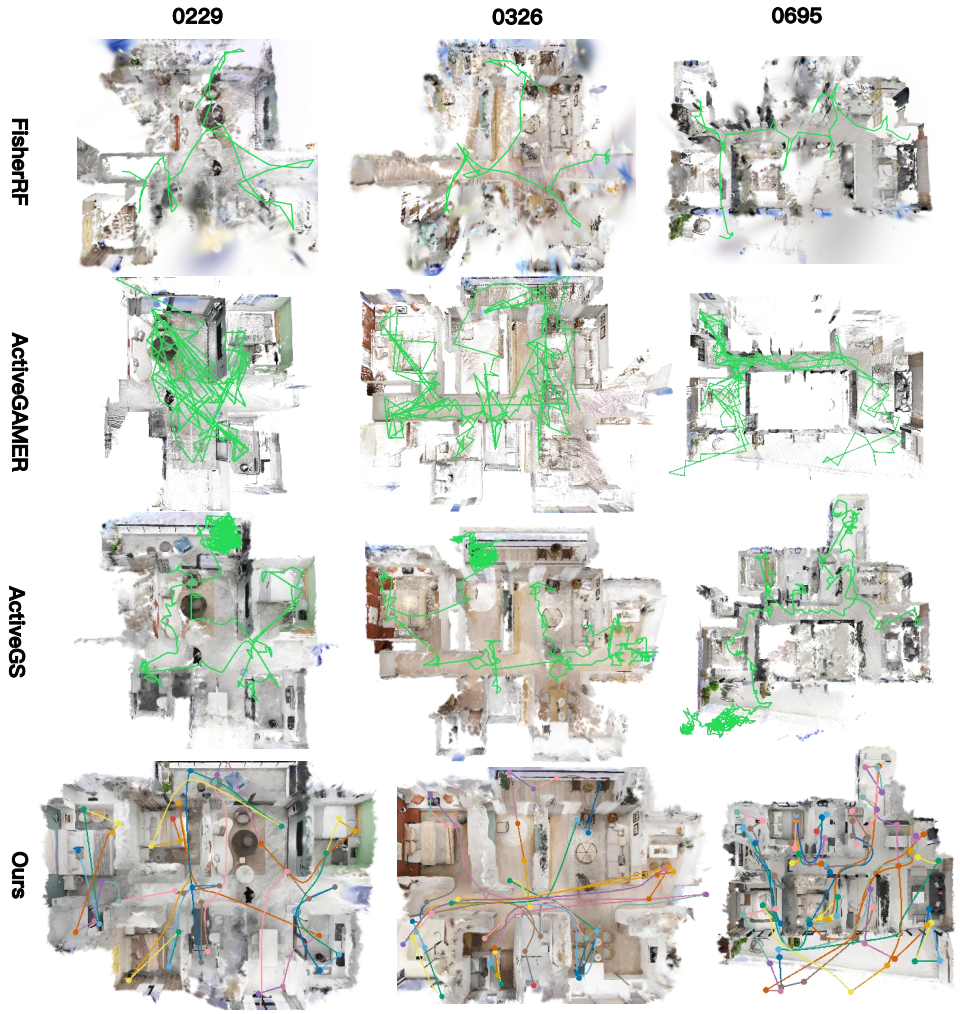

Figure 3: BEV visualization of Gaussian reconstruction and exploration trajectories on InteriorGS scenes. Filled dots mark selected frontiers. FrontierGS follows more regular trajectories and reconstructs more complete scenes than the baselines.

because public implementations are unavailable. All methods use the same scenes and initial poses, with scene lists, aggregation, and initialization details given in Appendix B.

Evaluation metrics. We evaluate geometry using *Completion* (Li et al., 2026a) and *Coverage*. Completion is the mean nearest-neighbor distance from ground-truth surface samples to reconstructed geometry. On Replica and Matterport3D, Coverage is the fraction of ground-truth surface samples within 5 cm of reconstructed geometry. On InteriorGS, it is the fraction of a fixed domain of reachable free space and adjacent surfaces observed as free or occupied, measuring exploration coverage. Novel-view rendering uses *PSNR*, *SSIM* (Wang et al., 2004), and *LPIPS* (Zhang et al., 2018) at identical test viewpoints and ground-truth evaluation regions. Missing reconstructions are not masked. *Path* is cumulative travel distance. See Appendix B for details. Each method retains its original mapper backend, while FrontierGS variants share 2DGS to isolate planning contributions.

4.2 Quantitative Results

As shown in Table 1, FrontierGS achieves reconstruction quality comparable to ActiveGS on compact Replica scenes. Local reconstruction views already cover much of these small scenes, leaving less room for exploration-driven view selection to further improve reconstruction quality.

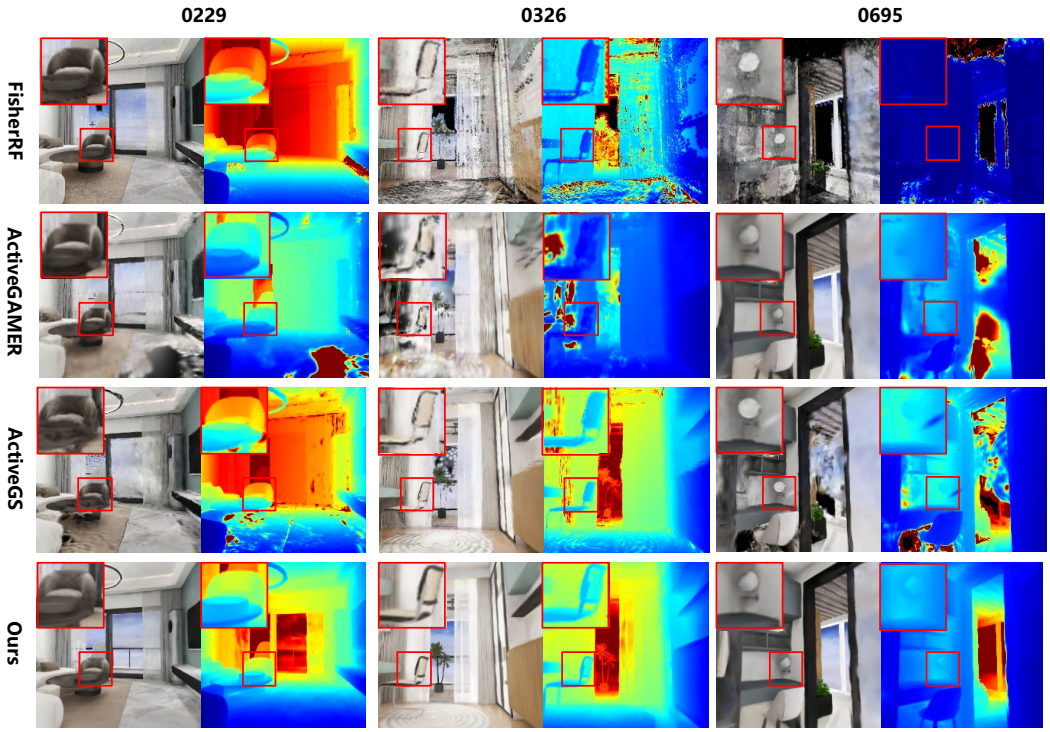

Figure 4: Novel-view rendering comparison on InteriorGS scenes. FrontierGS produces cleaner RGB and depth renderings with more complete scene geometry in large indoor environments.

On larger Matterport3D and InteriorGS scenes, FrontierGS shows clearer gains in coverage and rendering, especially on InteriorGS. At equal budgets, baselines may stay in observed regions, limiting progress. FrontierGS preserves the target frontier and selects complementary reconstruction views along the exploration path. Its longer InteriorGS path reflects sustained progress across successive frontiers. Several baselines stay in observed regions or enter local loops.

4.3 Qualitative Results

We qualitatively evaluate novel-view RGB and depth renderings on InteriorGS in Figure 4. FrontierGS shows fewer missing structures and Gaussian rendering artifacts, with cleaner RGB appearance and more consistent depth geometry than the baselines. Figure 3 compares Gaussian reconstruction and exploration trajectories in BEV. FrontierGS follows more regular trajectories toward successive frontiers and achieves more complete scene reconstruction, while several baselines exhibit irregular motion or become trapped in local exploration loops.

4.4 Ablation Studies

Reconstruction Viewpoint Selection. Table 1 evaluates two variants. On Matterport3D, *Ours (w/o corridor)* expands the sampling region while keeping other settings unchanged, resulting in substantially longer travel. On InteriorGS, *Ours (exploration only)* uses the same reconstruction backend but acquires views only through the exploration module, resulting in lower coverage and reconstruction quality. These results support selecting complementary reconstruction viewpoints near the exploration path.

Persistent Fisher Approximation. Table 2 compares persistent Fisher with full historical relinearization at matched checkpoints on InteriorGS and Matterport3D. At each checkpoint, only the Fisher state changes, while the map, candidate graph, visibility, motion costs, and beam search remain fixed. *Top-1 View Agreement* is the fraction of checkpoints with the same next view. *Top-1*

Table 2: Planning consistency between persistent Fisher and full historical relinearization. Decision Consistency uses a 1% score tolerance.

Dataset	Top-1 View Agree.↑	Top-1 in Full Top-3↑	Decision Cons.↑
InteriorGS	88.6%	96.3%	94.2%
Matterport3D	89.2%	98.8%	96.7%
Average	88.9%	97.6%	95.5%

Table 3: Effect of greedy selection and virtual information updates on subpath planning.

Variant	Recon. gain $\bar{I}_R$ ↑	Detour↓	Path score
Greedy Selection	0.218	**0.203**	0.267
w/o Virtual Update	0.326	0.346	0.392
Ours	**0.474**	0.295	**0.464**

Table 4: Ablation of Fisher maintenance, refresh strategy, and virtual information updates (VU). Runtime is cumulative in seconds, and τ denotes the fixed refresh interval in views.

Configuration			Quality		Updates	Cumulative runtime (s)				
Fisher	Refresh	VU	Coverage↑	PSNR↑ (dB)	# Refresh	Map Opt.↓	Fisher Update↓	Cand. Fisher↓	Beam Search↓	Total↓
Full relinearization	$\tau = 1$	✓	0.96	26.78	498	6109	7168	803	357	14437
Persistent	$\tau = 1$	✓	**0.97**	**27.04**	498	6120	154	963	366	7603
Persistent	$\tau = 5$	×	0.85	22.17	108	3896	66	110	154	4226
Persistent	$\tau = 5$	✓	0.91	23.96	108	3911	**65**	**99**	**151**	4226
Persistent	Adaptive	×	**0.97**	25.13	112	3862	**65**	118	**151**	4196
Persistent	Adaptive	✓	**0.97**	26.63	116	**3850**	**65**	118	155	**4188**

in Full Top-3 measures whether the persistent choice appears among the first views of the three highest-scoring fully relinearized paths. *Decision Consistency* allows 1% relative score loss. We compute this by rescoring both paths with Eq. 9 under full historical relinearization and normalizing by the absolute reference score plus ϵ_0. Persistent Fisher achieves 95.5% consistency. As shown in Table 4, persistent Fisher at $\tau = 1$ retains comparable quality at much lower Fisher update cost.

Virtual Information Updates. Table 3 compares greedy selection, beam search without virtual updates, and our full method over 100 seeded InteriorGS trials with adaptive refresh. Sequential path evaluation uses conditional reconstruction gain $\bar{I}_R$, normalized detour $\max\{0, L(\pi) - L_F\}/L_F$, and the path score in Eq. 9. Compared with beam search without virtual updates, our method improves reconstruction gain and path score while reducing detour. As shown in Table 4, virtual updates at fixed $\tau = 5$ improve PSNR by 1.79 dB and coverage from 0.85 to 0.91.

Adaptive Refresh and Efficiency Analysis. Table 4 compares refresh strategies over 500 RGB-D observations on InteriorGS, using the same exploration, candidate generation, and beam search settings. Without virtual updates, $\widetilde{H}_g = H_g$ stays fixed between refreshes in each branch. With virtual updates, adaptive refresh matches the persistent-Fisher $\tau = 1$ variant in coverage at comparable PSNR, with fewer refreshes and about 45% less cumulative runtime. It also improves coverage and PSNR over fixed $\tau = 5$ at similar runtime. Appendix C reports average Top-1 View Agreement of 94.6% and average Decision Consistency of 96.0%, comparing skipped and immediate refreshes.

5 Conclusion

We present FrontierGS, a framework for active Gaussian reconstruction driven by exploration that selects complementary viewpoints near the frontier exploration path. FrontierGS maintains persistent Fisher approximation without recomputing historical Jacobians, while virtual information updates enable conditional viewpoint selection between map refreshes. Adaptive refresh reduces repeated map optimization and Fisher computation according to their impact on the current planning decision. Experiments on multiple datasets demonstrate improved coverage and reconstruction quality in large scenes with substantially lower computation. A limitation is that the refresh predictor is empirically calibrated, and our evaluation is limited to simulated indoor environments. Future work will investigate more general refresh estimation and validation on real robotic platforms.

AI USE STATEMENT

In this work, we used generative AI tools for assisting with translation. Additionally, we used generative AI tools for editing a research paper to improve readability.

We have reviewed all AI-assisted work. AI-generated translations were reviewed by the authors to ensure that the original meaning and terminology were accurately conveyed. We also use generative AI for editing and polishing our paper to improve readability, which we carefully reviewed to ensure that there was no ambiguity in the wording.

We take responsibility for the final content of this work, including text, claims, or artifacts produced with the aid of generative AI.

REFERENCES

Chao Cao, Hongbiao Zhu, Howie Choset, and Ji Zhang. Tare: A hierarchical framework for efficiently exploring complex 3d environments. In *Robotics: Science and systems*, volume 5, pp. 2, 2021.

Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Niebner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang. Matterport3d: Learning from rgb-d data in indoor environments. In *2017 International Conference on 3D Vision (3DV)*, pp. 667–676, 2017. doi: 10.1109/3DV.2017.00081.

Benjamin Charrow, Gregory Kahn, Sachin Patil, Sikang Liu, Ken Goldberg, Pieter Abbeel, Nathan Michael, and Vijay Kumar. Information-theoretic planning with trajectory optimization for dense 3d mapping. In *Robotics: Science and systems*, volume 11, pp. 3–12. Rome, 2015.

Liyan Chen, Huangying Zhan, Kevin Chen, Xiangyu Xu, Qingan Yan, Changjiang Cai, and Yi Xu. Activegamer: Active gaussian mapping through efficient rendering. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16486–16497. IEEE, 2025.

Yanyu Chen, Jiyue Jiang, Jiahong Liu, Yifei Zhang, Xiao Guo, and Irwin King. Trace: Trajectory-aware comprehensive evaluation for deep research agents. In *Proceedings of the ACM Web Conference 2026*, pp. 2524–2534, 2026.

Titus Cieslewski, Elia Kaufmann, and Davide Scaramuzza. Rapid exploration with multi-rotors: A frontier selection method for high speed flight. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2135–2142. IEEE, 2017.

Cl Connolly. The determination of next best views. In *Proceedings. 1985 IEEE international conference on robotics and automation*, volume 2, pp. 432–435. IEEE, 1985.

Ziyue Feng, Huangying Zhan, Zheng Chen, Qingan Yan, Xiangyu Xu, Changjiang Cai, Bing Li, Qilun Zhu, and Yi Xu. Naruto: Neural active reconstruction from uncertain target observations. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21572–21583. IEEE, 2024.

Antoine Guédon, Pascal Monasse, and Vincent Lepetit. Scone: Surface coverage optimization in unknown environments by volumetric integration. *Advances in Neural Information Processing Systems*, 35:20731–20743, 2022.

Geoffrey A Hollinger and Gaurav S Sukhatme. Sampling-based robotic information gathering algorithms. *The International Journal of Robotics Research*, 33(9):1271–1287, 2014.

Armin Hornung, Kai M Wurm, Maren Bennewitz, Cyrill Stachniss, and Wolfram Burgard. Octomap: An efficient probabilistic 3d mapping framework based on octrees. *Autonomous robots*, 34(3):189–206, 2013.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2d gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH 2024 conference papers*, pp. 1–11, 2024.

Wen Jiang, Boshu Lei, and Kostas Daniilidis. Fisherrf: Active view selection and mapping with radiance fields using fisher information. In *European conference on computer vision*, pp. 422–440. Springer, 2024.

Liren Jin, Xingguang Zhong, Yue Pan, Jens Behley, Cyrill Stachniss, and Marija Popović. Activegs: Active scene reconstruction using gaussian splatting. *IEEE Robotics and Automation Letters*, 10 (5):4866–4873, 2025.

Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21357–21366. IEEE, 2024.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis, et al. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.

Shiyao Li, Antoine Guédon, Shizhe Chen, and Vincent Lepetit. Magician: Efficient long-term planning with imagined gaussians for active mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21606–21615, 2026a.

Yan Li, Yingzhao Li, and Gim Hee Lee. Active3d: Active high-fidelity 3d reconstruction via multi-level uncertainty quantification. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 6557–6565, 2026b.

Yuetao Li, Zijia Kuang, Ting Li, Qun Hao, Zike Yan, Guyue Zhou, and Shaohui Zhang. Activesplat: High-fidelity scene reconstruction through active gaussian splatting. *IEEE Robotics and Automation Letters*, 10(8):8099–8106, 2025.

Iker Lluvia, Elena Lazkano, and Ander Ansuategi. Active mapping and robot exploration: A survey. *Sensors*, 21(7):2445, 2021.

Bruce T Lowerre. *The harpy speech recognition system.* Carnegie Mellon University, 1976.

Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.

Julian Straub, Thomas Whelan, Lingni Ma, Yufan Chen, Erik Wijmans, Simon Green, Jakob J Engel, Raul Mur-Artal, Carl Ren, Shobhit Verma, et al. The replica dataset: A digital replica of indoor spaces. *arXiv preprint arXiv:1906.05797*, 2019.

Rayen Tan, Rohan Ghuge, and Viswanath Nagarajan. Informative path planning with limited adaptivity. In *International Conference on Artificial Intelligence and Statistics*, pp. 4006–4014. PMLR, 2024.

MTI SpatialVerse Research Team et al. Interiorgs: A 3d gaussian splatting dataset of semantically labeled indoor scenes, 2025.

Hassan Umari and Shayok Mukhopadhyay. Autonomous robotic exploration based on multiple rapidly-exploring randomized trees. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1396–1402. IEEE, 2017.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Joey Wilson, Marcelino Almeida, Sachit Mahajan, Martin Labrie, Maani Ghaffari, Omid Ghasemalizadeh, Min Sun, Cheng-Hao Kuo, and Arnab Sen. Pop-gs: Next best view in 3d-gaussian splatting with p-optimality. In *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3646–3655. IEEE, 2025.

Ziyuan Xia, Jingyi Xu, Chong Cui, Yuanhong Yu, Jiazhao Zhang, Qingsong Yan, Tao Ni, Junbo Chen, Xiaowei Zhou, Hujun Bao, et al. Habitat-gs: A high-fidelity navigation simulator with dynamic gaussian splatting. In *European Conference on Computer Vision*, pp. 306–323. Springer, 2026.

Yuhan Xie, Yixi Cai, Yinqiang Zhang, Lei Yang, and Jia Pan. GauSS-MI: Gaussian Splatting Shannon Mutual Information for Active 3D Reconstruction. In *Proceedings of Robotics: Science and Systems*, LosAngeles, CA, USA, June 2025. doi: 10.15607/RSS.2025.XXI.030.

Zijun Xu, Rui Jin, Ke Wu, Yi Zhao, Zhiwei Zhang, Jieru Zhao, Fei Gao, Zhongxue Gan, and Wenchao Ding. Hgs-planner: Hierarchical planning framework for active scene reconstruction using 3d gaussian splatting. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 14161–14167. IEEE, 2025.

S Xue, J Dill, D Ahuja, F Dellaert, P Tsiotras, and D Xu. Uncertainty-driven 3d gaussian splatting active mapping via anisotropic visibility field. 2026.

Brian Yamauchi. 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.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE/CVF conference on computer vision and pattern recognition*, pp. 586–595. IEEE, 2018.

A Conditional Score Stability and Adaptive Refresh

A.1 Conditional score bound.

At decision t, we compare the cached planning state with a reference obtained by immediately optimizing the map and incorporating buffered observations using the same persistent Fisher estimator and observation history. We consider comparisons in which the candidate graph, feasible path family Π_t, occupancy state, motion costs, Gaussian identities, and normalization remain fixed between the two states. Comparisons involving Gaussian creation, splitting, pruning, or normalization changes are excluded.

Let $\widetilde{H}_{t,g}$ denote the virtual information state and $\widetilde{C}_{t,i,g}$ the cached candidate Fisher increments. The reference uses the refreshed persistent information base $H^+_{t,g}$ and candidate Fisher increments $C^+_{t,i,g}$. Assume

$$\|\widetilde{H}_{t,g} - H^+_{t,g}\|_* \le \eta_{H,t}, \qquad \|\widetilde{C}_{t,i,g} - C^+_{t,i,g}\|_* \le \eta_{C,t}, \tag{A1}$$

uniformly over the Gaussians and candidates, where $\|\cdot\|_*$ denotes the nuclear norm and invisible candidate contributions are represented by zero matrices.

Let $\widetilde{Q}_t(\pi)$ and $Q^+_t(\pi)$ denote Eq. 9 evaluated using the cached and reference states under the fixed comparison above. Since any feasible path uses at most $K_{\max}$ candidates,

$$\sup_{\pi\in\Pi_t} |\widetilde{Q}_t(\pi) - Q^+_t(\pi)| \le E_t, \qquad E_t = \frac{2\eta_{H,t} + K_{\max}\eta_{C,t}}{D_{\mathrm{geo}}\lambda}. \tag{A2}$$

All cached and reference information states are bounded below by λI, since the persistent information bases include the same regularization and Fisher increments are positive semidefinite. For $X, Y \succeq \lambda I$,

$$|\log\det X - \log\det Y| \le \lambda^{-1}\|X - Y\|_*. \tag{A3}$$

This follows by integrating $\mathrm{tr}(X(s)^{-1}(X - Y))$ along $X(s) = (1 - s)Y + sX$ and using $\|X(s)^{-1}\|_2 \le \lambda^{-1}$. For a path with k candidates, the sequential marginal gains telescope into a difference of log determinants. The final information states differ in nuclear norm by at most $\eta_{H,t} + k\eta_{C,t}$, while the initial states differ by at most $\eta_{H,t}$. Applying Eq. A3, summing over the fixed N_{gs} Gaussians, and dividing by the normalization $N_{\mathrm{gs}}D_{\mathrm{geo}}$ in Eq. 8 gives Eq. A2. The discovery and detour terms are identical under this comparison and therefore cancel.

A.2 Stability within the returned path set.

Let $\mathcal{A}_t \subseteq \Pi_t$ be the complete paths returned by beam search under the cached state, and let $\widehat{\pi}_t = \arg\max_{\pi\in\mathcal{A}_t} \widetilde{Q}_t(\pi)$. When at least two paths are returned, define

$$m_t = \widetilde{Q}_t(\widehat{\pi}_t) - \max_{\pi\in\mathcal{A}_t\setminus\{\widehat{\pi}_t\}} \widetilde{Q}_t(\pi). \tag{A4}$$

By Eq. A2, for every $\pi \in \mathcal{A}_t \setminus \{\widehat{\pi}_t\}$,

$$Q^+_t(\widehat{\pi}_t) - Q^+_t(\pi) \ge m_t - 2E_t. \tag{A5}$$

Hence $2E_t < m_t$ preserves the selected path relative to all alternatives in $\mathcal{A}_t$. This does not guarantee that rerunning beam search after refreshing returns the same path set, since intermediate pruning decisions may change.

For completeness, define the cached beam search error

$$\delta_t = \max_{\pi\in\Pi_t} \widetilde{Q}_t(\pi) - \widetilde{Q}_t(\widehat{\pi}_t) \ge 0.$$

Applying Eq. A2 to the reference optimum and $\widehat{\pi}_t$ gives

$$\max_{\pi\in\Pi_t} Q^+_t(\pi) - Q^+_t(\widehat{\pi}_t) \le 2E_t + \delta_t. \tag{A6}$$

Thus Eq. A5 characterizes stability within the returned path set, whereas Eq. A6 bounds the score loss over the full feasible family under the same fixed comparison. Both results apply to one decision under the same observation history and do not bound divergence between complete trajectories.

A.3 Empirical adaptive refresh.

The drift quantities in Eq. A1 are not directly available online. We therefore calibrate an empirical predictor of score error from completed refreshes that satisfy the comparison conditions above. At a valid refresh r, let $n_r > 0$ be the number of buffered observations and let $\mathcal{A}_r$ contain the paths returned immediately before the refresh. We record

$$d_r = \max_{\pi \in \mathcal{A}_r} |Q_r^+(\pi) - \widetilde{Q}_r(\pi)|, \qquad a_r = \frac{d_r}{n_r}. \tag{A7}$$

Let $\mathcal{W}_t$ contain the indices of up to the most recent W valid records available at decision t. When $\mathcal{W}_t$ is nonempty, we estimate

$$\widehat{a}_t = \kappa \max_{j \in \mathcal{W}_t} a_j, \qquad \widehat{E}_t = n_t \widehat{a}_t, \qquad \kappa \geq 1, \tag{A8}$$

where n_t is the number of observations since the last refresh. We use $W = 5$ and $\kappa = 1$. Before any valid calibration record is available, we refresh every 3 observations.

When at least two paths are returned, a tentative search provides the margin m_t, and we refresh when

$$2\widehat{E}_t \geq \beta m_t \quad \text{or} \quad n_t \geq n_{\max}, \qquad 0 < \beta \leq 1. \tag{A9}$$

When fewer than two paths are returned, the margin test is omitted and refresh is determined by $n_t \geq n_{\max}$. A triggered refresh is followed by a new search before execution. Refreshes that do not satisfy the comparison conditions above are not used for calibration.

Under these comparison conditions, when the margin test is available and $\widehat{E}_t \geq E_t$, skipping the refresh gives

$$2E_t \leq 2\widehat{E}_t < \beta m_t \leq m_t.$$

Therefore, Eq. A5 guarantees that the selected path remains the highest-scoring path within A_t under the refreshed state, while Eq. A6 bounds its score loss relative to the optimum over Π_t by $\beta m_t + \delta_t$. Finite calibration records in Eqs. A7 and A8 do not guarantee $\widehat{E}_t \geq E_t$. Eq. A9 therefore remains an empirical rule motivated by the sufficient stability condition. Its decision consistency is evaluated in Appendix C.

B Additional Implementation Details

B.1 Scene selection and aggregation.

We evaluate the main comparison on the following scenes.

- Replica: `office0`, `office1`, `office2`, `office3`, `office4`, `room0`, `room1`, `room2`.
- Matterport3D: `17DRP5sb8fy`, `2t7WUuJeko7`, `RPmz2sHmrrY`, `YVUC4YcDtcY`, `x8F5xyUWy9e`.
- InteriorGS: `0759_839979`, `0695_841526`, `0326_840493`, `0229_840306`, `0088_839961`, `0224_840270`.

All methods use the same scenes within each dataset. For each scene, we randomly sample five initial camera poses including position and orientation, and use the same poses for all methods. Each method is run once from each pose, with random seeds controlling candidate sampling. We first average the five runs within each scene and then compute an unweighted average across scenes. The main comparisons use fixed budgets of 100, 250, and 500 RGB-D observations on Replica, Matterport3D, and InteriorGS, respectively. Ablation budgets are specified separately in Section 4.4.

B.2 Hardware, sensor, and viewpoint sampling.

We implement FrontierGS in PyTorch and run all methods on the same NVIDIA RTX 5090 GPU. RGB-D observations are generated with Habitat-GS (Xia et al., 2026), while each method retains

Algorithm 2 Exploration coverage on InteriorGS

Input: Ground-truth occupancy, indoor free space mask, initial pose, online occupancy map
Output: Exploration coverage C_{exp}

1: Construct a 0.2 m voxel grid
2: Downsample the ground-truth surface occupancy to the voxel grid
3: Compute reachable free voxels R from the initial pose using a 6-connected flood fill
4: Compute the adjacent surface shell

$$S = \text{Dilate}_6(R) \cap O$$

5: Define the evaluation domain

$$D' = (R \cap E) \cup S$$

where E is the indoor free space mask
6: Obtain the unknown voxel mask U from the final online occupancy map
7: Compute

$$C_{\text{exp}} = \frac{|D' \cap \neg U|}{|D'|}$$

8: **return** C_{exp}

its original reconstruction mapper. Color and depth images are rendered at 512×512 resolution with horizontal and vertical fields of view of 60° and a depth range of $[0.01, 5]$ m. Reconstruction candidate viewpoints are sampled at 1 m intervals along the main path within the corridor defined in Section 4.1. For Fisher evaluation, each candidate is rendered from the current 2DGS map rather than simulator ground truth. For each candidate viewpoint i and visible Gaussian g, we compute the RGB-D geometry Jacobian $J_{i,g}$ and form the Fisher increment $C_{i,g} = J_{i,g}^\top J_{i,g}$, treating rendered depth as a fixed pseudo-observation for the relative depth residual. Rendering evaluation uses identical test viewpoints, resolution, pixel inclusion rules, and ground-truth evaluation regions across methods, without masking missing reconstructions. Completion and coverage follow the definitions in Section 4.1.

B.3 Frontier exploration.

The exploration module maintains the occupancy map at a voxel resolution of 0.2 m. Frontiers are extracted at boundaries between free and unknown space and clustered within a 2.5 m neighborhood, with at least five voxels per cluster. The ground plane height prior is disabled to retain indoor frontiers near the floor. Viewpoints around each frontier are sampled at radii from 0.5 m to 1.5 m with an angular interval of 15° and a minimum clearance of 0.15 m from occupied voxels. Target selection refines the seven nearest frontiers within 5 m and evaluates up to 15 viewpoints per frontier. Viewpoints with a visibility ratio below 0.8 are discarded. A* search at 0.1 m resolution generates the main path $\mathcal{T}$ while avoiding occupied space.

B.4 Evaluation metrics.

We evaluate reconstruction geometry using *Completion* and *Coverage*. *Completion* is the mean distance in meters from points sampled on the ground-truth mesh to their nearest neighbors in the reconstructed point cloud, so missing surfaces also increase the error. *Coverage* follows two protocols. On Replica and Matterport3D, reconstruction coverage measures the fraction of ground-truth surface samples whose nearest reconstructed geometry lies within 5 cm. On InteriorGS, exploration coverage measures the fraction of a fixed evaluation domain consisting of reachable free space and adjacent surfaces that has been observed as free or occupied in the online occupancy map. This avoids penalizing methods for ground-truth surfaces that cannot be observed from reachable indoor free space. The detailed procedures are given in Algorithms 2 and 3. We evaluate novel-view rendering using PSNR, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018) with identical test viewpoints, 512×512 resolution, pixel inclusion rules, and evaluation regions determined only from ground truth. Missing reconstructions are not masked. Path reports the cumulative travel distance in meters under the fixed observation budget.

Algorithm 3 Surface reconstruction coverage on Replica and Matterport3D

Input: Ground-truth mesh, reconstructed scene representation, captured camera poses
Output: Reconstruction coverage C_{rec}

1: Uniformly sample 5×10^5 points P_{GT} from the ground-truth mesh according to triangle area
2: Render surface depth from the reconstructed map at the captured poses
3: Backproject valid rendered depths to obtain reconstructed points P_{rec}
4: **for** each $x \in P_{\mathrm{GT}}$ **do**
5: $\quad d_x \leftarrow \min_{y \in P_{\mathrm{rec}}} \|x - y\|_2$
6: **end for**
7: Compute

$$C_{\mathrm{rec}} = \frac{1}{|P_{\mathrm{GT}}|} \sum_{x \in P_{\mathrm{GT}}} \mathbf{1}[d_x < 0.05 \text{ m}]$$

8: **return** C_{rec}

Table 5: Decision Consistency at checkpoints where the adaptive rule skips refreshing. The reference performs an immediate refresh using the same observations and candidate graph. Decision Consistency uses a 1% score tolerance.

Dataset	Evaluated checkpoints	Top-1 View Agreement (%) ↑	Decision Consistency (%) ↑
InteriorGS	50	93.7	94.3
Matterport3D	20	95.5	97.7

For Matterport3D, rendered surface depths with opacity at least 0.5 are first fused into a TSDF volume with 2 cm voxels and a truncation distance of three voxels. We extract the resulting surface mesh, uniformly sample 10^6 reconstruction points, and apply the same 5 cm nearest-neighbor criterion from ground-truth samples to reconstructed geometry.

B.5 Fisher information after Gaussian splitting.

When a Gaussian g is replaced by m children $g_1, \ldots, g_m$, we initialize

$$\widehat{F}_{\mathrm{base},g_j} \leftarrow \frac{1}{m} \widehat{F}_{\mathrm{base},g}, \qquad j = 1, \ldots, m.$$

The parent matrix is discarded. This equal allocation avoids duplicating the accumulated information across the children while preserving the sum of the inherited Fisher matrices. Regularization is added separately to each child as $H_{g_j} = \widehat{F}_{\mathrm{base},g_j} + \lambda I$. Subsequent Fisher increments are accumulated independently for each child as in Eq. 3, without recomputing historical contributions. The allocation is an initialization heuristic rather than the historical Fisher evaluated at the child parameters. Gaussian splitting therefore remains outside the stability analysis in Appendix A.

C Decision Consistency of Adaptive Refresh

Top-1 View Agreement measures whether the unrefreshed and refreshed reference states select the same first viewpoint. Decision Consistency also counts decisions with at most 1% relative score loss. Both selected paths are rescored from the same refreshed information base using the subpath objective and sequential virtual updates. The relative loss is the reference score minus the score of the path selected without refreshing, normalized by the absolute reference score plus $\epsilon_0 > 0$. Metrics are computed per run, averaged over the two runs of each scene, and then across scenes. Table 5 shows high agreement on both datasets, indicating that skipped refreshes rarely alter the next viewpoint decision or substantially reduce its refreshed score. The comparison concerns individual decisions under fixed candidates rather than complete trajectory consistency.

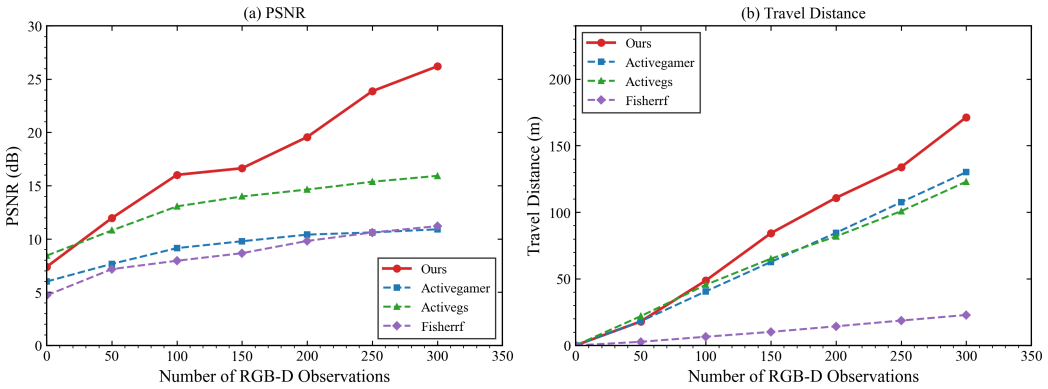

Figure 5: Performance evolution with increasing RGB-D observations on InteriorGS: (a) PSNR and (b) cumulative travel distance. FrontierGS achieves the highest PSNR and the longest cumulative travel distance among the compared methods at 300 observations.

Table 6: Per-scene results on InteriorGS. Values are averaged over the same initial poses, with standard deviations shown in parentheses.

Scene	Coverage↑	Completion (m)↓	PSNR (dB)↑	SSIM↑
0759_839979	0.9735 (±0.0148)	0.041 (±0.018)	29.90 (±1.90)	0.938 (±0.029)
0695_841526	0.9751 (±0.0246)	0.029 (±0.007)	26.76 (±2.10)	0.924 (±0.033)
0326_840493	0.9414 (±0.0189)	0.056 (±0.010)	28.29 (±1.70)	0.941 (±0.035)
0229_840306	0.9449 (±0.0309)	0.059 (±0.013)	28.24 (±2.20)	0.848 (±0.031)
0088_839961	0.9698 (±0.0157)	0.051 (±0.010)	28.01 (±2.50)	0.867 (±0.022)
0224_840270	0.9486 (±0.0344)	0.060 (±0.017)	28.96 (±2.20)	0.879 (±0.035)

D Additional Experimental Results

D.1 Performance Evolution with Increasing Observations

To examine how reconstruction quality and travel distance evolve as observations accumulate, we report PSNR and cumulative travel distance on InteriorGS in Figure 5. At 300 RGB-D observations, FrontierGS achieves the highest PSNR and the longest cumulative travel distance among the compared methods.

D.2 Per-scene Results

To assess performance across different environments, we report per-scene results for FrontierGS on InteriorGS in Table 6. Each value is averaged over the five initial poses used in the main evaluation, with the standard deviation indicating variation across runs. The results indicate that FrontierGS maintains high coverage and reconstruction quality across all six evaluated scenes.

D.3 Sensitivity to the Reconstruction Backend

To assess whether reconstruction backend differences account for the performance gaps in Table 1, we additionally evaluated baseline methods using the same 2DGS backend as FrontierGS on scene 0695 of InteriorGS. Replacing the native backends resulted in only small changes in reconstruction quality. These changes were small relative to the corresponding performance gaps reported in Table 1, suggesting that backend choice alone does not explain the observed gains in the evaluated settings.

D.4 Supplementary Visual Results

Figure 6 provides additional scene-level visual results that complement the examples in the main paper.

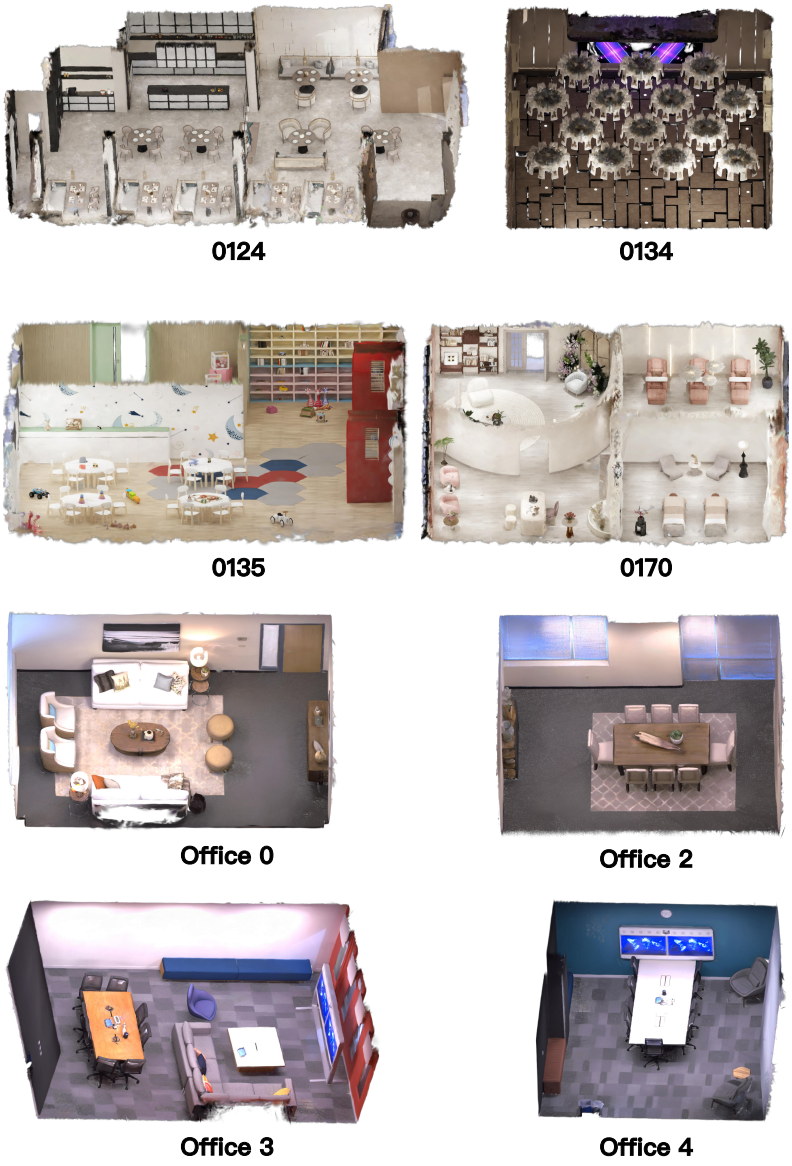

Figure 6: Supplementary reconstruction results on InteriorGS scenes 0124, 0134, 0135, and 0170, and Replica scenes Office 0, Office 2, Office 3, and Office 4.