

SA-RESGS: SELF-AUGMENTED RESIDUAL 3D GAUSSIAN SPLATTING FOR NEXT BEST VIEW SELECTION

Anonymous authors

Paper under double-blind review

ABSTRACT

We propose Self-Augmented Residual 3D Gaussian Splatting (SA-ResGS), a novel framework for stabilizing uncertainty quantification and enhancing uncertainty-aware supervision in next-best-view (NBV) selection for active scene reconstruction. SA-ResGS improves both the reliability of uncertainty estimates and their effectiveness for supervision by generating Self-Augmented point clouds (SA-Points) via triangulation between a training view and a rasterized extrapolated view, enabling efficient scene coverage estimation. While improving scene coverage through physically guided view selection, SA-ResGS also addresses the challenge of under-supervised Gaussians, exacerbated by sparse and wide-baseline views, by introducing the first residual learning strategy tailored for 3D Gaussian Splatting. This targeted supervision enhances gradient flow in high-uncertainty Gaussians by combining uncertainty-driven filtering with dropout- and hard-negative-mining-inspired sampling. Our contributions are threefold: (1) a physically grounded view selection strategy that promotes efficient and uniform scene coverage; (2) an uncertainty-aware residual supervision scheme that amplifies learning signals for weakly contributing Gaussians, improving training stability and uncertainty estimation across scenes with diverse camera distributions; (3) an implicit unbiasing of uncertainty quantification as a consequence of constrained view selection and residual supervision, which together mitigate conflicting effects of wide-baseline exploration and sparse-view ambiguity in NBV planning. Experiments on active view selection demonstrate that SA-ResGS outperforms state-of-the-art baselines in both reconstruction quality and view selection robustness.

1 INTRODUCTION

Recent advances in neural rendering—particularly Neural Radiance Fields (NeRFs) (Mildenhall et al., 2020) and 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023)—have significantly advanced photorealistic scene reconstruction (Yu et al., 2024; Niedermayr et al., 2024; Kulhanek et al., 2024), enabling high-fidelity, real-time applications across diverse environments. Beyond static scene capture, these methods have spurred broader interest in tackling complex challenges, such as active view selection (Wang et al., 2023; Xiao et al., 2024; Chen et al., 2024), uncertainty quantification for next-best-view (NBV) selection (Wen et al., 2024). While pre-captured, dense-view training methods can achieve impressive reconstruction quality, in-situ (active) reconstruction—where views are selected and added progressively—remains challenging due to artifacts caused by shape-radiance ambiguity and further exacerbated by limited training views and dynamics of view addition strategy. Despite the inherent difficulty of estimating reliable uncertainty under such conditions, recent post-hoc approaches—such as Laplacian approximation-based, model-agnostic methods (Wen et al., 2024; Goli et al., 2024)—have shown promise by providing uncertainty signals without altering the rendering pipeline. However, several critical challenges remain unaddressed:

- **Disregarded physical constraints:** Computational uncertainty is often misaligned with physical plausibility of the reconstructed geometry.
- **Underutilized supervision:** Existing methods rarely convert uncertainty cues into actionable learning signals, leaving weakly contributing Gaussians under-supervised throughout training.
- **Performance dependency:** The reliability of uncertainty estimation remains tightly coupled with training dynamics, particularly in the early stages when scene coverage is incomplete.

In response to these challenges, we propose SA-ResGS, a Self-Augmented Residual 3D Gaussian Splatting framework that stabilizes uncertainty quantification and enhances uncertainty-aware supervision in next-best-view selection for progressive scene reconstruction, as shown in Fig. 1. SA-ResGS strategically decouples view selection from strong dependence on uncertainty estimates driven by internal learning dynamics, promoting more robust and geometry-aware surface coverage. Concretely, we physically prefilter a subset of candidate views based on their geometric dissimilarity and then apply uncertainty-based scoring within this subset—effectively implementing a physically grounded, uncertainty-informed selection strategy. To guide this process, we construct SA-Points by triangulating dense correspondences between a given training view and its rasterized extrapolated views at each view-selection step, after a fixed number of initial views have been trained. These SA-Points are then encoded using a hash-based scene representation, enabling efficient similarity measurement between candidate views and pre-selected views. We select the most dissimilar candidate views in the encoded space, encouraging coverage of previously unseen regions.

While the physically grounded and uncertainty-informed selection strategy enhances overall scene coverage, it inadvertently increases sparsity in multiview overlap—since more dissimilar views are less likely to observe shared regions, thereby weakening multi-view geometric constraints. To counterbalance this contradiction without sacrificing the benefits of diverse view selection, we introduce a residual supervision mechanism that reinforces training using ground truth RGB images. The proposed SA-ResGS—particularly its residual supervision module—is fully compatible with conventional 3DGS pipelines that rely on point-cloud-based rasterization and gradient-based optimization. Beyond the original supervision, SA-ResGS additionally rasterizes color images using a selected Gaussian subset: a small fraction of the most uncertain Gaussians is combined with a majority subset (e.g., 90%) of the original visible Gaussians. This strategy mirrors the principles of residual learning and draws conceptual inspiration from Dropout (Park et al., 2025; Srivastava et al., 2014) and Hard Negative Mining (Xuan et al., 2020; Jang et al., 2019), by amplifying supervision for Gaussians that typically receive weak gradients during backpropagation—similar to the role of skip connections in ResNet (He et al., 2016) in mitigating vanishing gradients. Our method thus provides additional learning signals specifically targeted at under-optimized Gaussians which are often overlooked due to their low contribution in 3DGS’s rendering mechanism, receive repeated 2D photometric supervision, resulting in more stable optimization and improved reconstruction in sparse or ambiguous regions.

The main contributions of the SA-ResGS are threefold:

- **Physically grounded view selection:** We propose a geometry-aware strategy using Self-Augmented Points (SA-Points) to guide next-best-view selection, enforcing physical plausibility and promoting more balanced, coverage-oriented exploration.
- **Residual learning for 3DGS:** We introduce the first residual supervision framework specifically designed for 3D Gaussian Splatting, addressing the vanishing gradient problem by reinforcing weakly supervised Gaussians and improving both optimization stability and reconstruction quality.
- **Unbiased uncertainty quantification:** By jointly improving view distribution and supervising under-optimized Gaussians, SA-ResGS mitigates geometric sparsity and density bias, leading to fairer and more reliable uncertainty estimates throughout training.

2 RELATED WORKS

Next-best-view selection. NBV selection originated from the robotics community as a strategy to efficiently guide in-situ scene capture, where the goal was to incrementally select viewpoints that maximally reduce reconstruction ambiguity with minimal sensor movement (Connolly, 1985; Scott et al., 2003; Delmerico et al., 2018). Classical NBV methods typically followed rule-based paradigms, selecting views based on geometric coverage (Dunn and Frahm, 2009; Guédon et al., 2022; Guédon et al., 2023), viewpoint entropy (Vázquez et al., 2001), or visibility heuristics (Bircher et al., 2016; Sun et al., 2021). As NBV became increasingly relevant to 3D computer vision, especially under sparse-view constraints, learning-based approaches emerged to model scene-specific view policies via reinforcement learning or active learning frameworks (Wang et al., 2024). While these data-driven methods demonstrate improved adaptability over hand-crafted rules, they often rely on task-specific reward definitions and struggle to generalize across scene types.

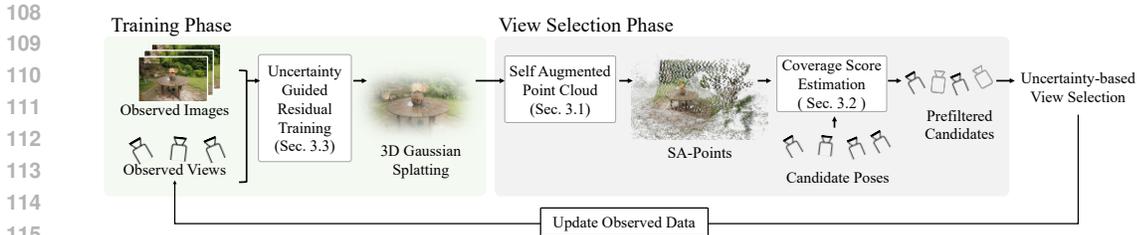


Figure 1: **Overview of SA-ResGS.** The framework alternates between view selection and training. At each NBV step, Self-Augmented Points are generated via triangulation from dense correspondences between a training view and its extrapolated render, enabling surface-aware coverage estimation (Sec.4.1). Candidate views are first physically filtered using hash-encoded feature dissimilarity, then ranked by uncertainty quantification scores for final selection (Sec.4.2). During training, residual supervision (Sec. 4.4) combines full and uncertainty-intensified renders to reinforce gradients toward weakly contributing Gaussians, improving training stability and reconstruction quality under sparse-view conditions.

More recent work explore representation-aware NBV strategies based on NeRF and 3DGS. Approaches such as NARUTO Feng et al. (2024), which learns a grid-based uncertainty field, and ActiveGAMER Chen et al. (2025), which derives uncertainty from 3DGS visibility, show how representation-specific cues can guide informative view selection. Building on this direction, information-theoretic models such as FisherRF Wen et al. (2024) offer a principled formulation for uncertainty-based NBV in neural fields. In our work, we build on this line (Wen et al., 2024; Goli et al., 2024; Hanson et al., 2024; Wilson et al., 2025) by integrating physically grounded geometry priors to further stabilize early-stage view planning, particularly when minimal visual input available.

Uncertainty quantification for 3DGS and neural rendering. Uncertainty quantification plays a pivotal role in active reconstruction, particularly for guiding view selection and supervision. In 3D Gaussian Splatting (3DGS), however, the high dynamic nature of primitive splitting and sensitivity to initialization often leads to unstable training, degrading the reliability of intermediate uncertainty signals. Earlier methods based on variational inference (Shen et al., 2021; Lee et al., 2025; Shen et al., 2022; Lyu et al., 2024) and ensemble-based estimates (Sünderhauf et al., 2023) enable stochastic or distributional reasoning, but they require costly model retraining or parallel inference and are typically incompatible with standard rendering pipelines. Recent post-hoc approaches (Wilson et al., 2025; Hanson et al., 2024) such as FisherRF (Wen et al., 2024) and BayesRays (Goli et al., 2024) estimate uncertainty using Laplacian approximations without altering model structure, and have demonstrated promising results on NeRF and 3DGS variants. Complementary to geometry-based methods, image-level approach (Wang et al., 2025b) leverage perceptual quality on current rendering result as a proxy for uncertainty. However, these methods remain strongly coupled with the density of underlying Gaussians—leading to biased uncertainty estimates in early training stages when geometry is sparse or unevenly distributed, often misinterpreting under-observed regions as confident. This overlooked bias limits the reliability of NBV guidance when it’s most needed. To mitigate limitation, we introduce residual learning with the help of using physically grounded view selection, enabling robust and a little loose-coupled uncertainty estimation in an early-stage of view selection while emphasizing the effect of targeting high-uncertainty Gaussian focused supervision.

Residual supervision in 3DGS. While accurate uncertainty estimation helps localize regions requiring stronger supervision, it alone does not guarantee that gradients effectively reach under-optimized Gaussians in the 3DGS pipeline. Residual learning, as popularized by ResNet (He et al., 2016), has proven effective in mitigating vanishing gradients and improving training stability through skip connections and additive refinement, yet it remains underexplored in the context of 3D Gaussian Splatting. Existing 3DGS methods mainly rely on direct photometric losses (Kerbl et al., 2023) or external depth priors (Li et al., 2024; Xu et al., 2024), which often fail to sufficiently supervise Gaussians with low opacity or minimal rendering contributions. Recent studies such as pixelSplat (Charatan et al., 2024), PAPR (Zhang et al., 2023), and PAPR-in-Motion (Peng et al., 2024) explicitly discuss the vanishing gradient issue and propose solutions including differentiable parameterization of Gaussians, proximity attention-based differentiable renderer, adaptive updates, and activation tuning. Despite

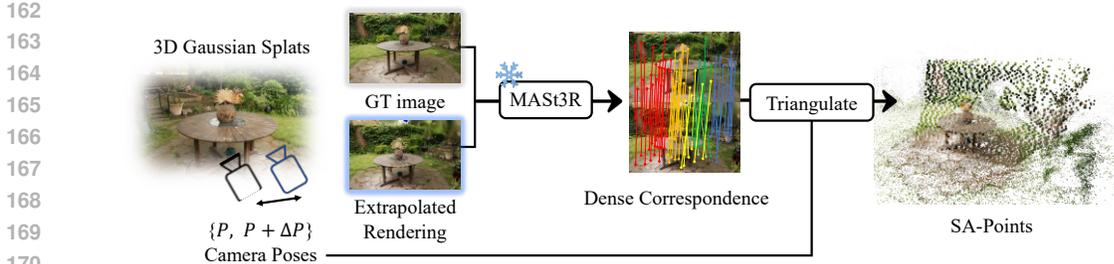


Figure 2: **SA-Points Generation.** An extrapolated image is rendered from a perturbed camera pose. Dense correspondences with the reference image are predicted using MAST3R, and triangulated to produce SA-Points, which are filtered by reprojection error for reliable surface geometry.

various strategies proposed to mitigate the vanishing gradient problem, prior approaches lack an explicit mechanism for correcting weakly supervised Gaussians, which remain largely unresolved due to insufficient gradient signals. Although dropout-based approaches (Park et al., 2025) help increase gradient diversity, they operate stochastically and do not target supervision to the most uncertain or least updated Gaussians. Our method addresses these limitations by introducing the first residual supervision strategy for 3DGS, applying uncertainty-guided rendering to intentionally amplify gradients for under-supervised Gaussians—without altering underlying rasterization process.

3 SELF-AUGMENTED RESIDUAL 3D GAUSSIAN SPLATTING

The proposed SA-ResGS framework is illustrated in Fig. 1. SA-ResGS builds upon the state-of-the-art next-best-view selection method, FisherRF (Wen et al., 2024), extending it with SA-Points to support two core ideas: (1) guiding physically grounded view selection with reduced reliance on uncertainty estimation, and (2) applying residual supervision to uncertain Gaussians, mitigating the vanishing gradient problem, wherein weakly contributing Gaussians—those with minimal impact on rasterized pixel values—receive insufficient gradients during backpropagation. The physically grounded view selection is enabled by a geometrically encoded surface representation, constructed using SA-Points derived from a single training view. To construct SA-Points, we employ the 3D vision foundational model, MAST3R (Leroy et al., 2024), to predict dense correspondences between a given training image and a rasterized extrapolated view rendered from a nearby camera pose. The resulting 2D correspondences are then triangulated to produce 3D SA-Points.

3.1 SELF-AUGMENTED POINTS GENERATION

Given a reference image I_r with camera pose $\mathbf{T}_r = [\mathbf{R}_r \mid \mathbf{t}_r]$, we render an extrapolated image I_e from a perturbed pose $\mathbf{T}_e = [\mathbf{R}_e \mid \mathbf{t}_e + \Delta\mathbf{t}]$ using 3D Gaussian Splatting (Kerbl et al., 2023). Dense correspondences $\{(\mathbf{p}_r^i, \mathbf{p}_e^i)\}$ between I and I_e are predicted using the pretrained MAST3R model (Leroy et al., 2024), which is robust to moderate viewpoint changes and capable of producing contextually meaningful matches even in the presence of minor geometric distortions. Each SA-Point \mathbf{X}^i is triangulated from a 2D correspondence pair using the projection matrices \mathbf{P}_r and \mathbf{P}_e derived from COLMAP, including intrinsic matrices over each extrinsic pose \mathbf{T} . However, because triangulation is performed repeatedly during training—while the model is still fitting to a sparse and incomplete geometry—rasterized extrapolated images may occasionally contain rendering noise due to inaccurately placed Gaussians. To ensure reliable geometry while fully leveraging the generalization capability of MAST3R, we apply reprojection error-based filtering:

$$\varepsilon^i = \frac{1}{2} (\|\mathbf{p}_r^i - \pi(\mathbf{P}_r \mathbf{X}^i)\| + \|\mathbf{p}_e^i - \pi(\mathbf{P}_e \mathbf{X}^i)\|), \quad \text{retain if } \varepsilon^i < \tau. \quad (1)$$

This filtering step discards geometrically inconsistent points while preserving accurate SA-Points from dense, context-aware matches—even when the extrapolated image is noisier than the original training view. Compared to prior methods such as CoMapGS (Jang and Pérez-Pellitero, 2025) or MP-SfM (Pataki et al., 2025), our triangulation pipeline produces scale-consistent, surface-aware geometry from a single image by leveraging extrapolated viewpoints rather than requiring multiview input or monocular depth estimates. The overall steps are visualized in Fig. 2.

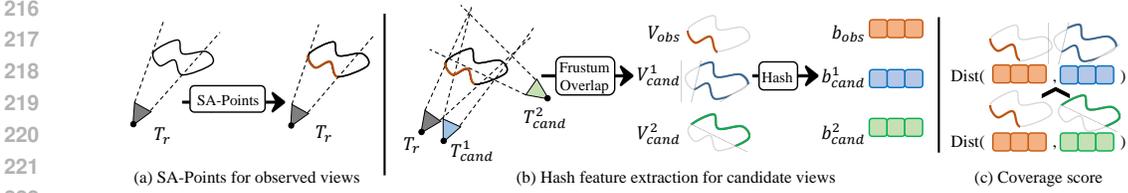


Figure 3: **Physically grounded candidate view selection via surface coverage.** (a) SA-Points from training views define observed voxels \mathcal{V}_{obs} . (b) Each candidate view generates a binary hash-encoded feature \mathbf{b} , via frustum-based visibility estimation. (c) Normalized Hamming distance between hash-encoded features quantifies coverage dissimilarity, enabling efficient selection of geometrically complementary views without rendered images or uncertainty scores.

3.2 PHYSICALLY GROUNDED VIEW SELECTION ALGORITHM

We present our physically grounded view selection algorithm for next-best-view (NBV) selection, illustrated in Fig. 3. As discussed in Sec. 2, NBV selection in 3D Gaussian Splatting (3DGS) is particularly challenging due to the tight coupling between uncertainty estimation and the quality of reconstructed geometry—both of which are highly sensitive to the sparsity and distribution of Gaussian splats. Under sparse-view settings, where reconstruction begins with as few as four images and new views are incrementally added every 100 training iterations, uncertainty-based NBV strategies often become unreliable. This is because uncertainty signals are inherently biased or unstable when the geometry is incomplete or under-constrained. To address this, we introduce a surface-aware guidance mechanism based on SA-Points, enabling view selection to operate independently of the computed uncertainty quantification. By decoupling view selection from the internal training dynamics of 3DGS, our method provides more stable and physically meaningful candidate views during the early reconstruction phase—even before the model has accumulated sufficient confidence to produce reliable uncertainty maps.

We begin by discretizing the 3D scene into a voxel grid $\mathcal{V} = \{v_k\}_{k=1}^K$, where each voxel represents a unit volume in the scene. The bounding volume of \mathcal{V} is defined by the sparse point cloud obtained from structure-from-motion (SfM). A voxel $v_k \in \mathcal{V}$ is marked as observed if it intersects with any SA-Point \mathbf{X}^i (Sec. 4.1), forming the subset $\mathcal{V}_{\text{obs}} \subset \mathcal{V}$. To account for potential localization errors and promote coverage continuity, we dilate each occupied voxel using a 3D kernel \mathcal{K}_r of radius r :

$$\tilde{\mathcal{V}}_{\text{obs}} = \bigcup_{v_k \in \mathcal{V}_{\text{obs}}} \mathcal{K}_r(v_k), \quad (2)$$

where $\tilde{\mathcal{V}}_{\text{obs}}$ denotes the dilated observed region for the current set of training views.

For each candidate view j , we compute a frustum $\mathcal{F}_j \subset \mathcal{V}$, derived from camera intrinsics (field of view) and near/far planes estimated from the SfM point distribution. A voxel is considered potentially visible from view j if its center lies within the frustum:

$$\mathcal{V}_{\text{cand}}^{(j)} = \{v_k \in \mathcal{V} \mid v_k \in \mathcal{F}_j\}. \quad (3)$$

To estimate the geometric dissimilarity between the current coverage and a candidate view (see Fig. 3), we compute the normalized Hamming distance:

$$d_j = \frac{1}{K} \left\| \mathbf{b}_{\text{obs}} \oplus \mathbf{b}_{\text{cand}}^{(j)} \right\|_1, \quad (4)$$

where \oplus denotes the element-wise XOR operation between binary vectors, and $\|\cdot\|_1$ is the ℓ_1 norm (i.e., the number of differing entries). Here, \mathbf{b} denotes a binary occupancy vector obtained by mapping voxel coordinates through a fixed random hashing function, following the spatial hashing strategy introduced in Instant-NGP (Müller et al., 2022). The resulting value $d_j \in [0, 1]$ measures the proportion of voxels with inconsistent occupancy status between the currently observed volume and the candidate view. Candidate views are then ranked in descending order of their normalized Hamming distances d_j , and the top $N\%$ (e.g., $N = 20$) are retained to form the physically filtered candidate set \mathcal{C}' .

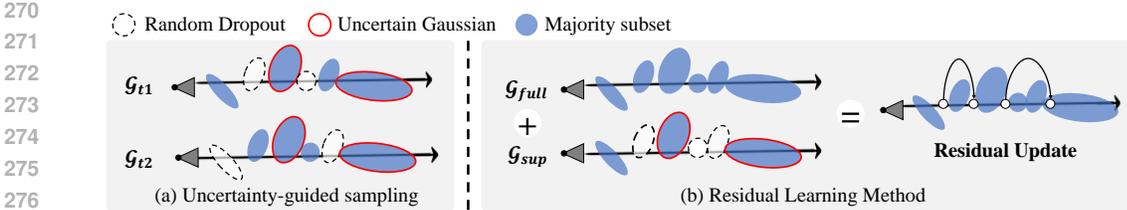


Figure 4: **Residual supervision in 3DGS.** (a) At each iteration (t_1, t_2), G_{sup} combines random and top-uncertain Gaussians; (b) residual supervision in 3DGS mimics ResNet-style skip connections.

We apply uncertainty quantification only within \mathcal{C}' , finalizing the view selection with finer-level scoring. This two-stage pipeline follows a coarse-to-fine strategy: it first expands the observed surface area using explicit geometric cues from SA-Points, then refines the choice using uncertainty-aware reasoning. By restricting uncertainty estimation to a smaller candidate pool, this approach improves computational efficiency while maintaining scene-aware diversity in the selected views. This strategy enables balanced, physically grounded exploration of unobserved regions.

3.3 UNCERTAINTY-GUIDED RESIDUAL LEARNING IN 3DGS

We propose the first residual learning for 3DGS to address the vanishing gradient issue affecting weakly contributing Gaussians, as shown in Fig. 4. These Gaussians often receive insufficient supervision during training due to their limited impact on rasterized pixels—particularly in sparse or ambiguous regions. While ResNet (He et al., 2016) mitigates similar issues via skip connections, such mechanisms are infeasible in 3DGS due to the dynamic and view-dependent nature of Gaussian properties. Instead, we propose a rasterizer-agnostic strategy that enhances gradient flow by generating auxiliary renders that emphasize high-uncertainty Gaussians. These images are supervised with ground-truth RGB images, forming the basis for a residual supervision scheme detailed below.

Residual supervision. To reinforce under-supervised Gaussians, we introduce a residual supervision scheme that leverages two rasterized images from the same training viewpoint: one using the full set of Gaussians \mathcal{G} and another from a guided subset \mathcal{G}_{sup} , as shown in Fig. 4(a). We define this subset as:

$$\mathcal{G}_{sup} = \mathcal{G}_{rand} \cup \mathcal{G}_{uncertain}, \quad (5)$$

where \mathcal{G}_{rand} is a random sample comprising $\alpha\%$ of \mathcal{G} (e.g., $\alpha=90$), and $\mathcal{G}_{uncertain}$ contains the top- β most uncertain Gaussians (e.g., $\beta=10$). To estimate uncertainty, we analyze two per-Gaussian attributes: opacity and scale. Gaussians with low opacity contribute minimally to alpha blending during rasterization, while those with large scale blur across pixels and tend to dominate ambiguous or low-texture regions. This rank identifies Gaussians that are both visually suppressed and spatially diffuse—making them key targets for correction. The combination \mathcal{G}_{sup} ensures that \mathcal{G}_{rand} maintains overall scene fidelity, while $\mathcal{G}_{uncertain}$ provides targeted supervision to gradient-deficient areas.

We compute two rendered images: I_{full} using the full set of Gaussians \mathcal{G} , and I_{sup} using the uncertainty-intensified subset \mathcal{G}_{sup} . Each image is supervised independently against the ground-truth image I_{gt} using ℓ_1 and SSIM losses:

$$\mathcal{L} = \sum_{i \in \{full, sup\}} \lambda_i [\mathcal{L}_{rgb}(I_i, I_{gt}) + \mathcal{L}_{ssim}(I_i, I_{gt})], \quad (6)$$

where $\lambda_{full} + \lambda_{sup} = 1$, and we set both to 0.5 in practice. We denote the losses for the full set render I_{full} and the uncertainty-intensified subset render I_{sup} as full loss (\mathcal{L}_{full}) and subset loss (\mathcal{L}_{sup}), respectively. The uncertainty-intensified rasterization strategy is conceptually inspired by Dropout (Park et al., 2025; Srivastava et al., 2014) and Hard Negative Mining (Xuan et al., 2020; Jang et al., 2019). Randomly sampling \mathcal{G}_{rand} provides stochastic diversity, allowing weakly contributing Gaussians to be supervised when dominant ones are excluded. Meanwhile, the deterministic inclusion of $\mathcal{G}_{uncertain}$ ensures consistent gradient flow to Gaussians that are persistently under-optimized. This dual mechanism reinforces learning in uncertain or ambiguous regions without modifying the rasterization process, and complements full-image supervision to maintain global photometric fidelity.

Category	Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow	Category	Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Rule-based	Random	19.969	0.584	0.456	Rule-based	Random	24.847	0.893	0.117
	ACP	20.325	0.596	0.449		ACP	22.718	0.855	0.138
2D-based	MUSIQ	19.850	0.575	0.466	2D-based	MUSIQ	25.237	0.889	0.119
	CrossScore	21.076	0.612	0.448		CrossScore	23.746	0.868	0.130
3D-based	FisherRF	20.642	0.595	0.450	3D-based	FisherRF	25.190	0.892	0.116
	Ours	21.410	0.613	0.451		Ours	26.580	0.907	0.110

(a) Mip-NeRF 360 dataset

(b) NeRF Synthetic dataset

Category	Methods	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
Rule-based	Random	18.918	0.694	0.390
	ACP	19.604	0.711	0.377
2D-based	MUSIQ	18.541	0.686	0.403
	CrossScore	19.709	0.725	0.366
3D-based	FisherRF	19.455	0.710	0.381
	Ours	20.060	0.722	0.377

(c) Deep Blending & Tank and Temples dataset

Table 1: **Quantitative results for the Active View Selection.** We compare our model with (1) Rule-based models (Random, ACP (Kopanas and Drettakis, 2023)), (2) 2D-based models (MUSIQ, CrossScore (Chen et al., 2024)), and 3D-based models (FisherRF (Wen et al., 2024)). Results are averaged over 9 scenes from the Mip-NeRF 360, 7 scenes from NeRF-synthetic dataset and 7 additional scenes from the Deep Blending and Tanks and Temples. We conduct four trials for each scene and report average scores. For statistics for all trial please refer to Appendix E

By supervising both the full and uncertainty-intensified images, we promote stronger gradient flow toward uncertain or low-opacity Gaussians without compromising photometric quality. This strategy mirrors the role of residual skip connections in ResNet (He et al., 2016) (Fig. 4(b)), supporting more stable convergence and mitigating overfitting in sparse or wide-baseline training settings. It is particularly effective in the early stages of next-best-view selection, where reconstruction is sensitive to both sparsely initialized regions and supervision bias caused by overfitting to limited views.

4 EXPERIMENTAL RESULTS

Dataset. We evaluate our approach on two benchmark dataset: NeRF-Synthetic Mildenhall et al. (2020), and Mip-NeRF 360 (Barron et al., 2022). While both datasets comprises from synthetic object-scale to real-world outdoor scenes with full 360-degree coverage, its uniform and curated camera trajectories provide limited challenge for active view selection, since even simple heuristics (e.g., furthest-distance selection) already perform reliably under balanced coverage (Xiao et al., 2024). To address this limitation, we carefully curate an extended benchmark dataset including seven diverse scenes from Deep Blending (Hedman et al., 2018) and Tanks and Temples (Knapitsch et al., 2017), which introduce unbalanced view distributions and varied scene scales that better reflect practical conditions. All experiments are conducted using images at their original resolutions, and further details on dataset curation are provided in Sec. A of the Appendix.

Counterparts. We compare our method quantitatively and qualitatively against several active 3DGS baselines which operate solely on RGB images: FisherRF (Wen et al., 2024), ACP (Kopanas and Drettakis, 2023), and random view selection. We also included 2D based view selection methods adopted from Active View Selector framework (Wang et al., 2025b). Following this framework, we incorporated two image quality assessment (IQA) models (MUSIQ and CrossScore) to evaluate perceptual quality. Both models were re-implemented according to the authors’ official instructions and publicly available code.

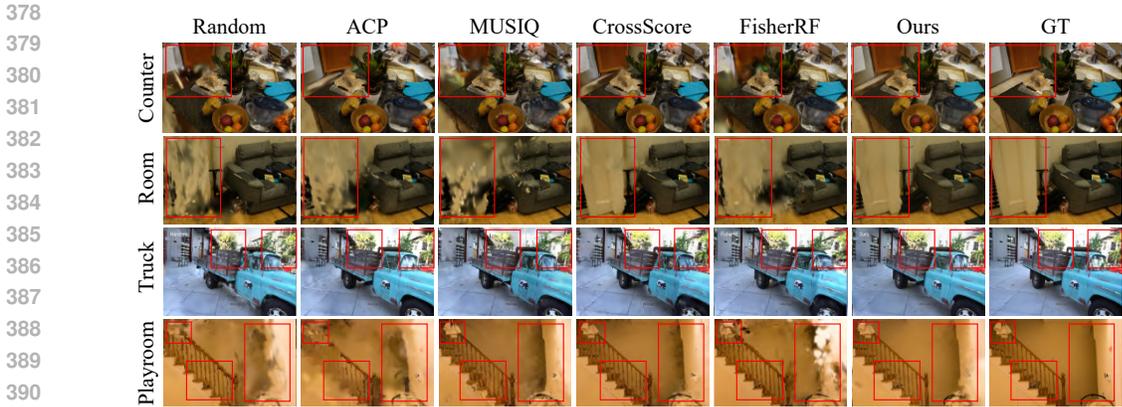


Figure 5: **Qualitative Comparison of Active View Selection.** Reconstruction from 20 selected views per scene. Our method shows improved completeness and fewer artifacts compared to baselines. For multi-view visualization, please refer the Appendix. Sec. C and supplementary video.

4.1 ACTIVE VIEW SELECTION

Experimental settings. Following the experimental protocols outlined by Wen et al. (2024), we adopt their prescribed initial view configurations and view selection schedules. Specifically, our experiments initiate with four uniformly distributed views, subsequently selecting an additional view every 100 epochs until reaching a total of 20 training views (For fewer training views, see Appendix Sec. C). We apply the same active selection strategy consistently across all datasets. For consistency, each model is initialized with the same random seed and trained for 20,000 iterations. All other settings remain unchanged across experiments, with exception of the view selection algorithms.

Results. Quantitative and qualitative results for the Mip-NeRF 360 dataset and additional scenes from the Deep Blending and Tank and Temples datasets are summarized in Table 1 and Fig. 5. The counterparts exhibit limited performance in 3D reconstruction, particularly in regions with sparse observations, due primarily to biased view selection and overfitting caused by vanishing gradients. This results in incomplete reconstructions, characterized by floating artifacts and missing geometry, *i.e.* holes and missing objects. In contrast, our method consistently delivers improved reconstruction quality, ranking first in PSNR and SSIM metrics and second in LPIPS for the Mip-NeRF 360 dataset. Our uncertainty-guided residual learning approach, based on dropout effects, produces smoother reconstructions in uncertain regions without compromising comparative performance. For camera view distribution, see Appendix Sec. B.

Additionally, experimental outcomes on the additional datasets validate the generalizability of our method. The Deep Blending dataset demonstrates our method’s effectiveness in handling real-world-like diverse camera distributions. Likewise, the extensive outdoor scenes of the Tank and Temples dataset further illustrate our enhanced coverage, exemplified by the Truck scene in Figure 5.

4.2 COMPARISON ON UNCERTAINTY ESTIMATION

Experimental settings. We evaluate the effectiveness of [our method](#) in improving uncertainty estimation accuracy. Specifically, we examine whether incorporating residual loss (\dagger +ResGS) and [self-augmented prefiltering](#) (\dagger +SA-ResGS) enhances the alignment between depth errors and predicted uncertainties under otherwise identical conditions. To measure this, we utilize the Area Under the Sparsification Error (AUSE) metric—a standard for evaluating uncertainty calibration adopted in (Wen et al., 2024; Goli et al., 2024; Shen et al., 2022)—where lower scores indicate better alignment between uncertainty and actual error. Thus, a lower AUSE score indicates better alignment between uncertainty predictions and actual errors, reflecting superior uncertainty calibration.

Following the approach used in CF-NeRF Shen et al. (2022), we employ depth maps from the NerfingMVS Wei et al. (2021) network, optimized using stereo depth from COLMAP at test time. Experiments are conducted on the [all nine scenes from Mip-NeRF 360](#) under identical view selection and evaluated using [all test views](#).

Methods	Our proposed methods		Metrics		
	Sec. 4.2	Sec. 4.4	PSNR	SSIM	LPIPS
FisherRF [†]	-	-	20.814	0.603	0.452
[‡] +ResGS	-	✓	21.022	0.596	0.459
[†] +ResGS	-	✓	20.740	0.610	0.442
[†] +SA-HashGS	✓	-	21.051	0.608	0.450
[†] +SA-ResGS	✓	✓	21.325	0.610	0.450

Table 2: **Ablation using Mip-NeRF 360 dataset.** [‡] denotes fixed-order view selection, replicating the original selection sequence from FisherRF[†], whereas [†] indicates dynamically updated view selection based on the model’s progressive training status.

Results. We compare the average AUSE among 9 scenes in Mip-NeRF 360. By incorporating residual learning applied in ResGS, we observed a reduction in AUSE from 0.327 to 0.323, and subsequently to 0.297, when progressing from baseline (FisherRF[†]) to [‡]+ResGS (adding residual loss), and [†]+ResGS (adding self-augmented prefiltering) respectively. These reductions demonstrate that both residual learning and self-augmented prefiltering enhance uncertainty calibration. The re-

sults signify enhanced alignment between predicted uncertainty and actual depth errors (see Fig. 6). This improvement is attributed to the ResGS approach, which increases certainty in regions with low prediction error while enabling continued learning through skip connections in uncertain regions. Consequently, our model achieves both structurally and quantitatively superior uncertainty calibration, leading to improved accuracy in uncertainty-driven tasks such as active mapping.

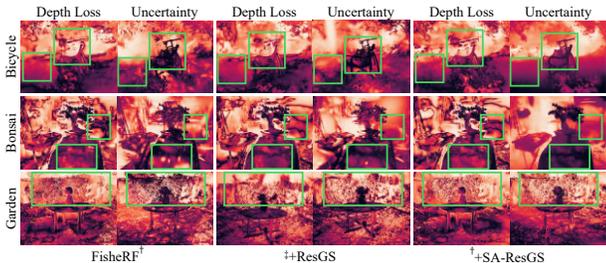


Figure 6: **Uncertainty Comparison.** We visualize depth loss with their corresponding uncertainty. Highlighted green boxes show that +SA-ResGS yields better alignment between large depth errors and high uncertainty

4.3 ABLATION STUDIES

Effect of individual components. The ablation study in Table 2 highlights the contributions of each component in our proposed SA-ResGS framework. Incorporating residual learning (ResGS) alone improves reconstruction stability, particularly in ambiguous or sparse regions ([‡]+ResGS). However, it shows limitations when view selection remains uncontrolled ([†]+ResGS), indicating that training improvements alone are insufficient, especially under high computational uncertainty quantification errors. As shown in Sec. 4.2, our training module effectively aligns predicted uncertainty with actual losses, yet purely uncertainty-driven view selection strategies remain vulnerable to bias from internal learning dynamics.

In contrast, self-augmented prefiltering (SA-HashGS) independently improves geometric coverage through physically grounded view selection. Notably, combining SA-HashGS with ResGS yields substantial synergistic improvements, confirming the complementarity between robust view selection and residual supervision in optimizing both uncertainty quantification and reconstruction quality. For visual comparison please see Appendix Sec. D.

Effect of the full loss. We compare our method with a variant that updates Gaussians only through guided subset (Fig. 7). Without the full loss, the model often over-smooths ambiguous or low-confidence regions, causing high-frequency detail loss. This occurs because rule-based updates in 3DGS, such as pruning, remove Gaussians that lack sufficient gradients. Without the main reconstruction loss reinforcing these regions, the subset-only variant lacks a mechanism to preserve or re-activate Gaussians that require continued refinement. The full ResGS, by integrating residual and full losses, prevents under-updated Gaussians from collapsing and preserves both global structure and fine details. For ablation on uncertainty-guided sampling please refer Appendix D, Fig. S8.

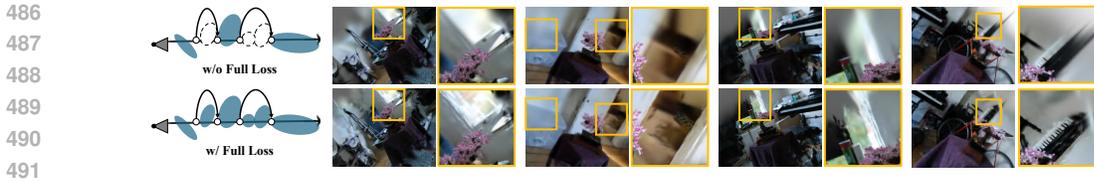


Figure 7: **Comparison of novel-view renderings with and without the full loss in ResGS.** The w/o full loss (\mathcal{L}_{sup}) variant exhibits smoothing artifacts, whereas incorporating the full loss ($\mathcal{L}_{full} + \mathcal{L}_{sup}$) preserves scene structure and high-frequency details. Note that both employ fixed-order view selection, using pre-selected view-order from FisherRF[†] to ignore view selection effect.

Robustness to correspondence noise. SA-Points are obtained through triangulation of dense correspondences, and their accuracy can influence reconstruction and view selection. To examine robustness, we conducted experiments with synthetic correspondence noise (0.0–10.0 pixels). The system remains stable up to moderate noise levels (≤ 5.0 pixels) when combined with a 1-pixel reprojection filter, as shown in Table. 3. Larger perturbations cause noticeable degradation, indicating that the system tolerates moderate correspondence errors without significant performance drop.

Noise	0.0	0.5	1.0	5.0	10.0
PSNR	24.441	24.199	24.117	24.311	23.121

Table 3: **Robustness to synthetic correspondence noise.** PSNR remains stable under moderate perturbations (≤ 5 pixels), showing SA-Points are resilient to realistic levels of noise.

4.4 COMPUTATION EFFICIENCY ANALYSIS

A key challenge in active view selection is computational cost, as FisherRF computes per-Gaussian Fisher information via backpropagation across all candidate views, creating bottlenecks in large-scale datasets. To evaluate our prefiltering strategy, we conducted a runtime analysis on the Bonsai scene using a mid-range GPU (38 TFLOPS fp32), summarized in Table 4. SA-ResGS replaces exhaustive Fisher evaluation with a four-step process: dense correspondence prediction (MASt3R), triangulation for SA-Points, voxel-based prefiltering, and Fisher computation on only 20% of views. Despite these extra steps, view selection is 55% faster (28.0s \rightarrow 12.5s), with only a modest increase in per-iteration cost (0.005s \rightarrow 0.027s). Unlike FisherRF, whose cost grows with candidate views and becomes impractical for large datasets such as ScanNet (about 1600 images), SA-ResGS maintains fixed training cost once the schedule is set. GPU memory usage rises slightly but remains within standard limits, making our approach both scalable and practical. We further report the end-to-end active-reconstruction runtime under same setting. In five repeated run, SA-ResGS reduced runtime by 55%. This aligns with Table 4: the reduced view-selection cost outweighs the extra operations.

Process	MASt3R	Triangulation	prefilter	Fisher	Total	Raster.	GPU	End-to-end
FisherRF	—	—	—	28.00 s	28.00 s	0.005 s/iter	~8K	32 m 59 s
Ours	0.40s	1.49 s	5.00s	5.60 s	12.50 s	0.027 s/iter	~10.5K	19 m 45 s

Table 4: **Runtime comparison of FisherRF and SA-ResGS on the Bonsai scene.** Breakdown of view selection and training costs. SA-ResGS introduces additional prefiltering steps but reduces total runtime by 55%, with only modest increases in per-iteration cost and GPU memory usage.

5 CONCLUSION

This paper presents SA-ResGS, a Self-Augmented Residual 3D Gaussian Splatting framework for stabilize uncertainty quantification and enhance uncertainty-aware supervision in next-best-view (NBV) selection for active scene reconstruction. We introduce Self-Augmented Points, triangulated from a training view and a rasterized extrapolated view. These points enable physically grounded view selection and help mitigate erroneous uncertainty bias, while supporting targeted 3D supervision in high-uncertainty regions. Furthermore, we propose the first residual learning strategy tailored to 3D Gaussian Splatting, enabling effective supervision for both uncertain image regions and weakly contributing Gaussian splats. This leads to improved photometric reconstruction in novel view synthesis. Experimental results demonstrate the effectiveness of SA-ResGS across a range of realistic scenes, encompassing both indoor and outdoor environments and varying scene scales.

LIMITATIONS AND FUTURE WORK

SA-ResGS has several limitations. The method also relies on hyperparameter choices, potentially affecting performance to specific scenarios. Method to probabilistically choosing parameter would have potential to handle such limitations. [We believe that applying residual supervision intermittently or adopting a tighter prefiltering ratio—without compromising reconstruction quality can further reduce computational cost.](#) Current method relies on 2D feedforward matching models, which show robust correspondence matching in most cases, application with more stronger model would further enhance the performance of the proposed methods. Extending SA-ResGS to dynamic environments involving moving objects or illumination changes remains an open challenge.

ETHIC STATEMENT

Our work focuses on advancing active 3D scene reconstruction through residual supervision and physically guided view selection in 3D Gaussian Splatting. The proposed method is entirely built upon publicly available datasets such as MipNeRF360 (Barron et al., 2022), Deep Blending (Hedman et al., 2018), and Tank and Temples dataset (Knapitsch et al., 2017), which are widely used in the vision community. No personally identifiable or sensitive information is included in our experiments. However, as with all vision datasets, there may exist inherent biases in scene distributions (e.g., indoor environments being more common than outdoor scenes). These biases could affect generalization in certain deployment scenarios. We encourage users to be aware of such limitations when applying our method to broader real-world applications.

RREPRODUCIBILITY STATEMENT

We provide detailed descriptions of our algorithmic components in Sec. 3, including uncertainty-guided supervision, view selection strategies, and residual learning for 3DGS. The datasets used (MipNeRF360 (Barron et al., 2022), DeepBlending (Hedman et al., 2018), Tanks and Temples datasets (Knapitsch et al., 2017)) and the evaluation protocols are specified in Sec. 4. To facilitate reproducibility, the hyperparameters, training recipes, and implementation details are documented in Appendix B. All experiments were conducted with standard PyTorch libraries and commonly available GPUs, ensuring that the results can be reproduced with reasonable computational resources.

REFERENCES

- Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2022.
- Andreas Bircher, Mina Kamel, Kostas Alexis, Helen Oleynikova, and Roland Siegwart. Receding horizon "next-best-view" planner for 3d exploration. In *ICRA*, 2016.
- David Charatan, Sizhe Lester Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024.
- Liyan Chen, Huangying Zhan, Kevin Chen, Xiangyu Xu, Qingan Yan, Changjiang Cai, and Yi Xu. Activegamer: Active gaussian mapping through efficient rendering. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2025.
- Xiao Chen, Quanyi Li, Tai Wang, Tianfan Xue, and Jiangmiao Pang. Gennbv: Generalizable next-best-view policy for active 3d reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition (IEEE Conf. Comput. Vis. Pattern Recog.)*, 2024.

- 594 C. I. Connolly. The determination of next best views. In *Proceedings. 1985 IEEE International*
595 *Conference on Robotics and Automation*, volume 2, 1985.
- 596 Jeffrey Delmerico, Stefan Isler, Reza Sabzevari, and Davide Scaramuzza. A comparison of volumetric
597 information gain metrics for active 3d object reconstruction. *Autonomous Robots*, 42(2), 2018.
- 599 Enrique Dunn and Jan-Michael Frahm. Next best view planning for active model improvement. In
600 *British Machine Vision Conference*, 2009. URL [https://api.semanticscholar.org/](https://api.semanticscholar.org/CorpusID:1345419)
601 [CorpusID:1345419](https://api.semanticscholar.org/CorpusID:1345419).
- 602 Ziyue Feng, Huangying Zhan, Zheng Chen, Qingan Yan, Xiangyu Xu, Changjiang Cai, Bing Li,
603 Qilun Zhu, and Yi Xu. Naruto: Neural active reconstruction from uncertain target observations. In
604 *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024.
- 606 Lily Goli, Cody Reading, Silvia Sellán, Alec Jacobson, and Andrea Tagliasacchi. Bayes’ Rays:
607 Uncertainty quantification in neural radiance fields. *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024.
- 608 Antoine Guédon, Pascal Monasse, and Vincent Lepetit. Scone: surface coverage optimization
609 in unknown environments by volumetric integration. In *Proceedings of the 36th International*
610 *Conference on Neural Information Processing Systems*, 2022.
- 612 Antoine Guédon, Tom Monnier, Pascal Monasse, and Vincent Lepetit. Macarons: Mapping and
613 coverage anticipation with rgb online self-supervision. In *IEEE Conf. Comput. Vis. Pattern Recog.*,
614 2023.
- 615 Alex Hanson, Allen Tu, Vasu Singla, Mayuka Jayawardhana, Matthias Zwicker, and Tom Goldstein.
616 Pup 3d-gs: Principled uncertainty pruning for 3d gaussian splatting. *arXiv*, 2024.
- 617 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image
618 recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2016.
- 620 Peter Hedman, Julien Philip, True Price, Jan-Michael Frahm, George Drettakis, and Gabriel Brostow.
621 Deep blending for free-viewpoint image-based rendering. 37(6), 2018.
- 622 Youngkyoon Jang and Eduardo Pérez-Pellitero. Comapgs: Covisibility map-based gaussian splatting
623 for sparse novel view synthesis, 2025.
- 625 Youngkyoon Jang, Hatice Gunes, and Ioannis Patras. Registration-free face-ssd: Single shot analysis
626 of smiles, facial attributes, and affect in the wild. *Computer Vision and Image Understanding*, 182,
627 2019.
- 628 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting
629 for real-time radiance field rendering. *ACM Trans. Graph.*, 2023.
- 631 Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking
632 large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4), 2017.
- 633 Georgios Kopanas and George Drettakis. Improving NeRF Quality by Progressive Camera Placement
634 for Free-Viewpoint Navigation. In Michael Guthe and Thorsten Grosch, editors, *Vision, Modeling,*
635 *and Visualization*. The Eurographics Association, 2023. ISBN 978-3-03868-232-5. doi: 10.2312/
636 vmv.20231222.
- 637 Jonas Kulhanek, Songyou Peng, Zuzana Kukelova, Marc Pollefeys, and Torsten Sattler. WildGaus-
638 sians: 3D gaussian splatting in the wild. In *Advances in Neural Information Processing Systems*
639 *(NeurIPS)*, 2024.
- 641 Sibaek Lee, Kyeongsu Kang, Seongbo Ha, and Hyeonwoo Yu. Bayesian nerf: Quantifying uncertainty
642 with volume density for neural implicit fields. *IEEE Robotics Autom. Lett.*, 10(3), 2025.
- 643 Vincent Leroy, Yohann Cabon, and Jerome Revaud. Grounding image matching in 3d with mast3r,
644 2024.
- 645 Jiahe Li, Xiao Bai Jiawei Zhang, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. Dngaussian: Depth-
646 regularized neural gaussian splatting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024.

- 648 Linjie Lyu, Ayush Tewari, Marc Habermann, Shunsuke Saito, Michael Zollhöfer, Thomas
649 Leimkühler, and Christian Theobalt. Manifold sampling for differentiable uncertainty in ra-
650 diance fields. In *SIGGRAPH Asia Conference Proceedings*, 2024. doi: 10.1145/3680528.3687655.
651
- 652 Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and
653 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *Eur. Conf.*
654 *Comput. Vis.*, 2020.
- 655 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics
656 primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4), July 2022. doi:
657 10.1145/3528223.3530127. URL <https://doi.org/10.1145/3528223.3530127>.
658
- 659 Simon Niedermayr, Josef Stumpfegger, and Rüdiger Westermann. Compressed 3d gaussian splatting
660 for accelerated novel view synthesis. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2024.
- 661 Hyunwoo Park, Gun Ryu, and Wonjun Kim. Dropgaussian: Structural regularization for sparse-view
662 gaussian splatting, 2025. URL <https://arxiv.org/abs/2504.00773>.
663
- 664 Zador Pataki, Paul-Edouard Sarlin, Johannes L. Schönberger, and Marc Pollefeys. MP-SfM: Monoc-
665 ular Surface Priors for Robust Structure-from-Motion. In *IEEE Conf. Comput. Vis. Pattern Recog.*,
666 2025.
- 667 Shichong Peng, Yanshu Zhang, and Ke Li. Papr in motion: Seamless point-level 3d scene interpolation.
668 In *IEEE Conf. Comput. Vis. Pattern Recog.*, 2024.
669
- 670 Johannes L. Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *IEEE Conf.*
671 *Comput. Vis. Pattern Recog.*, 2016.
- 672 Johannes L. Schonberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view
673 selection for unstructured multi-view stereo. In *Eur. Conf. Comput. Vis.*, 2016.
674
- 675 William R. Scott, Gerhard Roth, and Jean-François Rivest. View planning for automated three-
676 dimensional object reconstruction and inspection. *ACM Computing Surveys (CSUR)*, 35(1), 2003.
677
- 678 Jianxiong Shen, Adria Ruiz, Antonio Agudo, and Francesc Moreno-Noguer. Stochastic neural
679 radiance fields: Quantifying uncertainty in implicit 3d representations. In *International Conference*
680 *on 3D Vision*, 2021.
- 681 Jianxiong Shen, Antonio Agudo, Francesc Moreno-Noguer, and Adria Ruiz. Conditional-flow nerf:
682 Accurate 3d modelling with reliable uncertainty quantification. In *Eur. Conf. Comput. Vis.*, 2022.
683
- 684 Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov.
685 Dropout: a simple way to prevent neural networks from overfitting. *Journal of Machine Learning*
686 *Research*, 15(1), 2014. ISSN 1532-4435.
- 687 Yifan Sun, Qixing Huang, Dun-Yu Hsiao, Li Guan, and Gang Hua. Learning view selection for 3d
688 scenes. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2021.
689
- 690 Niko Sünderhauf, Jad Abou-Chakra, and Dimity Miller. Density-aware nerf ensembles: Quantifying
691 predictive uncertainty in neural radiance fields. In *2023 IEEE International Conference on Robotics*
692 *and Automation (ICRA)*, 2023.
- 693 Pere-Pau Vázquez, Miquel Feixas, Mateu Sbert, and Wolfgang Heidrich. Viewpoint selection using
694 viewpoint entropy. In *Proceedings of the Vision Modeling and Visualization Conference 2001*,
695 2001.
- 696 Tao Wang, Weibin Xi, Yong Cheng, Hao Han, and Yang Yang. Rl-nbv: A deep reinforcement learning
697 based next-best-view method for unknown object reconstruction. *Pattern Recognition Letters*, 184,
698 2024.
699
- 700 Yifan Wang, Jianjun Zhou, Haoyi Zhu, Wenzheng Chang, Yang Zhou, Zizun Li, Junyi Chen,
701 Jiangmiao Pang, Chunhua Shen, and Tong He. π^3 : Scalable permutation-equivariant visual
geometry learning, 2025a. URL <https://arxiv.org/abs/2507.13347>.

702 Zihang Wang et al. Uncertainty-aware ensemble nerf for active view planning. In *IEEE Conf. Comput.*
703 *Vis. Pattern Recog.*, 2023.

704

705 Zirui Wang, Yash Bhargat, Ruining Li, and Victor Adrian Prisacariu. Active view selector: Fast and
706 accurate active view selection with cross reference image quality assessment. 2025b.

707

708 Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfingmvs: Guided
709 optimization of neural radiance fields for indoor multi-view stereo. In *Int. Conf. Comput. Vis.*,
710 2021.

711

712 Jiang Wen, Lei Boshu, and Daniilidis Kostas. Fisherrf: Active view selection and uncertainty
713 quantification for radiance fields using fisher informations. In *Eur. Conf. Comput. Vis.*, 2024.

714

715 Joey Wilson, Marcelino Almeida, Sachit Mahajan, Martin Labrie, Maani Ghaffari, Omid Ghase-
716 malizadeh, Min Sun, Cheng-Hao Kuo, and Arnab Sen. Pop-gs: Next best view in 3d-gaussian
717 splatting with p-optimality. 2025.

718

719 Wenhui Xiao, Rodrigo Santa Cruz, David Ahmedt-Aristizabal, Olivier Salvado, Clinton Fookes, and
720 Leo Lebrat. Nerf director: Revisiting view selection in neural volume rendering. In *IEEE Conf.*
721 *Comput. Vis. Pattern Recog.*, June 2024.

722

723 Wangze Xu, Huachen Gao, Shihe Shen, Jianbo Jiao Rui Peng, and Ronggang Wang. Mvpgs:
724 Excavating multi-view priors for gaussian splatting from sparse input views. In *Eur. Conf. Comput.*
725 *Vis.*, 2024.

726

727 Hong Xuan, Abby Stylianou, Xiaotong Liu, and Robert Pless. Hard negative examples are hard, but
728 useful. In *Eur. Conf. Comput. Vis.*, 2020.

729

730 Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free
731 3d gaussian splatting. In *IEEE Conf. Comput. Vis. Pattern Recog.*, June 2024.

732

733

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

APPENDIX OVERVIEW

This supplementary document provides additional implementation details and results that support and extend the main paper. It is organized as follows:

- Sec. A. Implementation Details.
- Sec. B. Camera Distribution Analysis.
- Sec. C. Extended Qualitative Comparisons.
- Sec. D. Qualitative Results for Ablation Studies.
- Sec. E. Additional Ablation Analysis.
- Sec. F. Supplementary Video Overview.

A IMPLEMENTATION DETAILS

A.1 COVERAGE ESTIMATION

We follow the baseline setup of FisherRF (Wen et al., 2024), initializing our experiments with a sparse point cloud reconstructed via COLMAP (Schonberger and Frahm, 2016; Schonberger et al., 2016). This SfM output defines the axis-aligned bounding volume for discretizing the scene into voxel grids. To enable physically grounded view selection, we construct Self-Augmented Points (SA-Points) via triangulation between a training view and its extrapolated view. SA-Points are used to estimate surface occupancy and to encode observed geometry into binary voxel features. Candidate views are scored based on their voxel-level dissimilarity to the encoded training views, supporting robust coverage estimation without relying on early uncertainty signals.

In the following, we detail (1) Self-Augmented Points (SA-Points) Generation, and (2) Observed Surface Coverage Estimation and View Frustum Construction.

A.1.1 SELF-AUGMENTED POINTS (SA-POINTS) GENERATION

Extrapolated View Generation. To synthesize a novel rasterized view while maintaining sufficient scene overlap, we perturb the original camera center by ± 0.25 units along the x and y axes and translate it backward by 0.5 units along the z -axis. This backward-only perturbation ensures that the extrapolated view retains a high degree of visibility overlap with the original training view, keeping most scene content within the shared frustums. As a result, the extrapolated view covers a large portion of the original training image while still providing a novel perspective of the same surfaces, enabling reliable correspondence estimation for SA-Point generation.

Dense Correspondence Matching. We compute dense correspondences between the ground truth image of original training view and the rendering from extrapolated view using the pretrained MAST3R model (Leroy et al., 2024) (MASt3R_ViTLarge_BaseDecoder_512_catmlpdpt_metric), which is robust to moderate viewpoint perturbations. As discussed in the main paper, this 3D vision foundation model produces context-aware dense correspondences by capturing structured scene semantics over 16×16 local patches, even when the extrapolated view includes geometric artifacts. This robustness allows us to extract reliable matches despite distortions in the rasterized extrapolated images, enabling consistent 3D reconstruction from single-view observations. Consequently, we can leverage extrapolated viewpoints to augment sparse training views without requiring multiview supervision.

Triangulation and Reprojection Filtering. To ensure geometric consistency, we triangulate 3D points from matched correspondences and filter them based on reprojection error, discarding points with a bidirectional reprojection error exceeding 0.5 pixels. For computational efficiency, we parallelize the triangulation process across multiple CPU threads and apply a spatial stride of 5 pixels in both x and y directions to subsample correspondence pairs. These filtering strategies significantly accelerate processing while preserving high-fidelity geometric structure.

810 A.1.2 OBSERVED SURFACE COVERAGE ESTIMATION AND VIEW FRUSTUM CONSTRUCTION

811
812 **Voxel Grid Construction.** Given the sparse SfM points, we define an axis-aligned bounding box
813 (AABB) that encapsulates the entire 3D point cloud. The AABB is computed from the minimum and
814 maximum bounds of the reconstructed points.

815 **Initial Occupancy Estimation.** We discretize the scene into a voxel grid and mark a voxel as
816 occupied if it contains a minimum number of SfM points. To better represent scene geometry,
817 we apply N -fold upsampling to the occupied voxels. For the minimum number for determining
818 occupancy, we use 2 for outdoor scenes, and 5 for indoor cases.

819 **Observed Region Calculation.** SA-Points are mapped to their nearest voxels to define the observed
820 surface region. To account for possible triangulation errors and improve spatial robustness, we apply
821 a 3D dilation operation to the occupied voxels. For all cases, a dilation radius of 2 is applied.
822

823 **View Frustum Determination.** To evaluate candidate views, we define view frustums using camera
824 intrinsics and the global maximum bounds computed from the SfM point cloud. These frustums ignore
825 visibility constraints but serve as a conservative estimate of potential scene coverage. Coverage scores
826 for physically grounded view selection are then computed by measuring voxel-level intersections
827 between the candidate frustums and the observed surface, represented via hash-encoded voxel
828 occupancy representation.

829 A.2 EXPERIMENTAL SETUP

830
831 All models presented in the main manuscript, including the ablation variants of our proposed
832 method, were trained and evaluated on a single NVIDIA V100 GPU with 32 GB of memory. CPU-
833 based components—such as SA-Points generation, voxel grid processing, and triangulation—were
834 parallelized across 8 threads for efficiency. Each experiment was run once with a fixed random seed
835 of 0 to ensure reproducibility.
836

837 A.2.1 DATASET DETAIL

838
839 We evaluate our method and baselines on [three](#) types of datasets: (1) Mip-NeRF 360, (2) NeRF-
840 Synthetic and (3) Extended NBV benchmark datasets. Mip-NeRF 360 consists of nine real-world
841 scenes with dense 360-degree camera coverage, captured across both indoor and outdoor scenarios.
842 [NeRF-Synthetic dataset consists of 8 object-centric, multi-view images with dense 360-degree
843 camera view distribution, rendered via Blender. To ensure consistency across datasets, we generated
844 sparse point clouds for all NeRF-Synthetic scenes using COLMAP. We note that Ficus scene fails to
845 reconstruct reliably under COLMAP, therefore we excluded it and evaluated on the remaining scenes.](#)

846 While this datasets provides a controlled and well-curated benchmark, its uniform view distribution
847 limits the difficulty of next-best-view (NBV) evaluation. Such settings often make next-best-view
848 (NBV) strategies appear less critical, since even simple heuristics (e.g., furthest-distance selection)
849 already perform reliably under balanced coverage (Xiao et al., 2024)

850 To address this, we additionally construct an Extended NBV benchmark by selecting seven challeng-
851 ing scenes from Deep Blending and Tanks and Temples, characterized by irregular camera trajectories
852 and diverse scene scales. This curated set introduces more realistic and unbalanced conditions,
853 offering a complementary testbed for assessing robustness in active view selection.

854 These datasets contain large-scale and geometrically complex scenes with irregular camera distri-
855 butions, but their difficulty also means that many methods fail outright. To filter such degenerate
856 cases, we trained the FisherRF baseline under the standard scheme and retained only scenes where it
857 achieved at least 17 dB PSNR, ensuring that the comparisons remained fair and informative.

858 This procedure produced seven representative scenes: Horse, Truck, Francis, Ballroom, Barn, Ponche,
859 and Playroom. The selected set spans both indoor and outdoor environments, and includes highly
860 complex camera distributions (e.g., Ballroom, Ponche, and Playroom) that deviate substantially
861 from the curated coverage of Mip-NeRF 360. These characteristics create more realistic stress
862 tests for NBV strategies by introducing occlusions, scale variations, and unbalanced observations.
863 A brief visualization of scene configurations and camera trajectories is provided in Sec.2 of this
supplementary material.

B CAMERA DISTRIBUTION ANALYSIS

To evaluate the effectiveness of our view selection strategy, we compare the camera distributions produced by different methods. Fig. S1, S3, and S1 visualize these distributions from both bird’s-eye and side perspectives. As discussed in the main paper, FisherRF tends to produce clustered view selections due to its tight coupling with the internal 3DGS learning dynamics—visibly highlighted in the semi-transparent yellow regions. In contrast, our method yields a more spatially uniform and well-dispersed distribution of viewpoints.

This distinction is further illustrated in Fig. S4, where in the Room scene from the Mip-NeRF 360 dataset, FisherRF often selects redundant or near-parallel views. Our method instead promotes angular diversity, leading to broader scene exploration. A similar trend is observed in the Deep Blending dataset, where our approach selects viewpoints across a wider vertical range. These comparisons reinforce our claim that SA-ResGS facilitates physically grounded and geometrically diverse view selection, contributing to improved scene coverage.

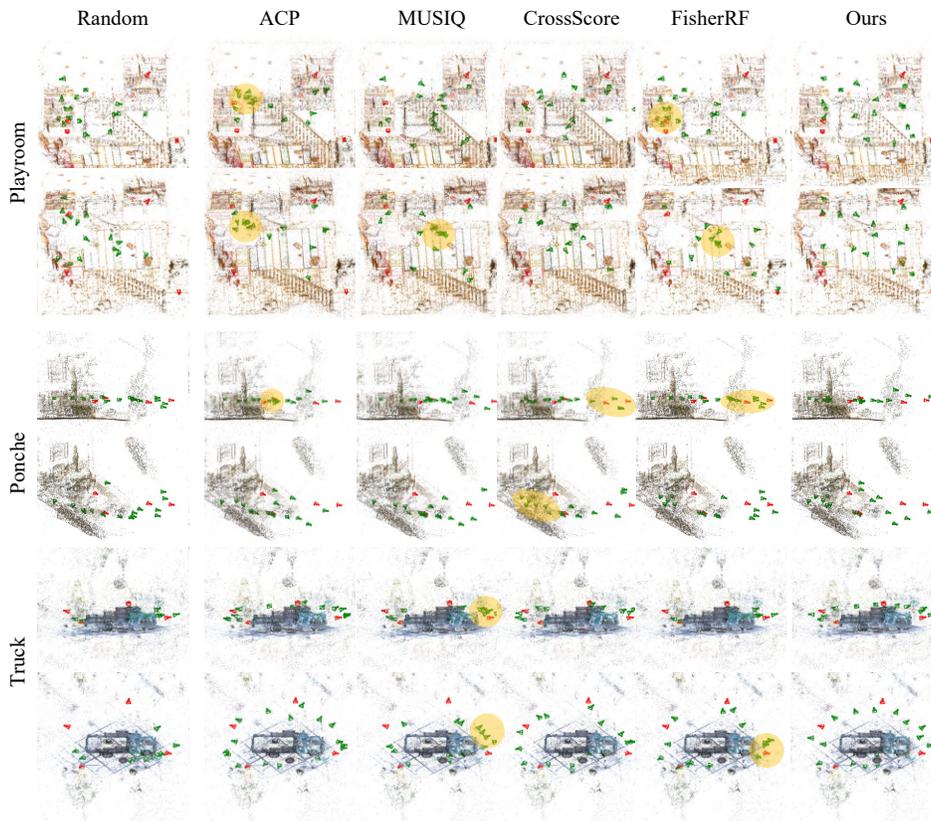
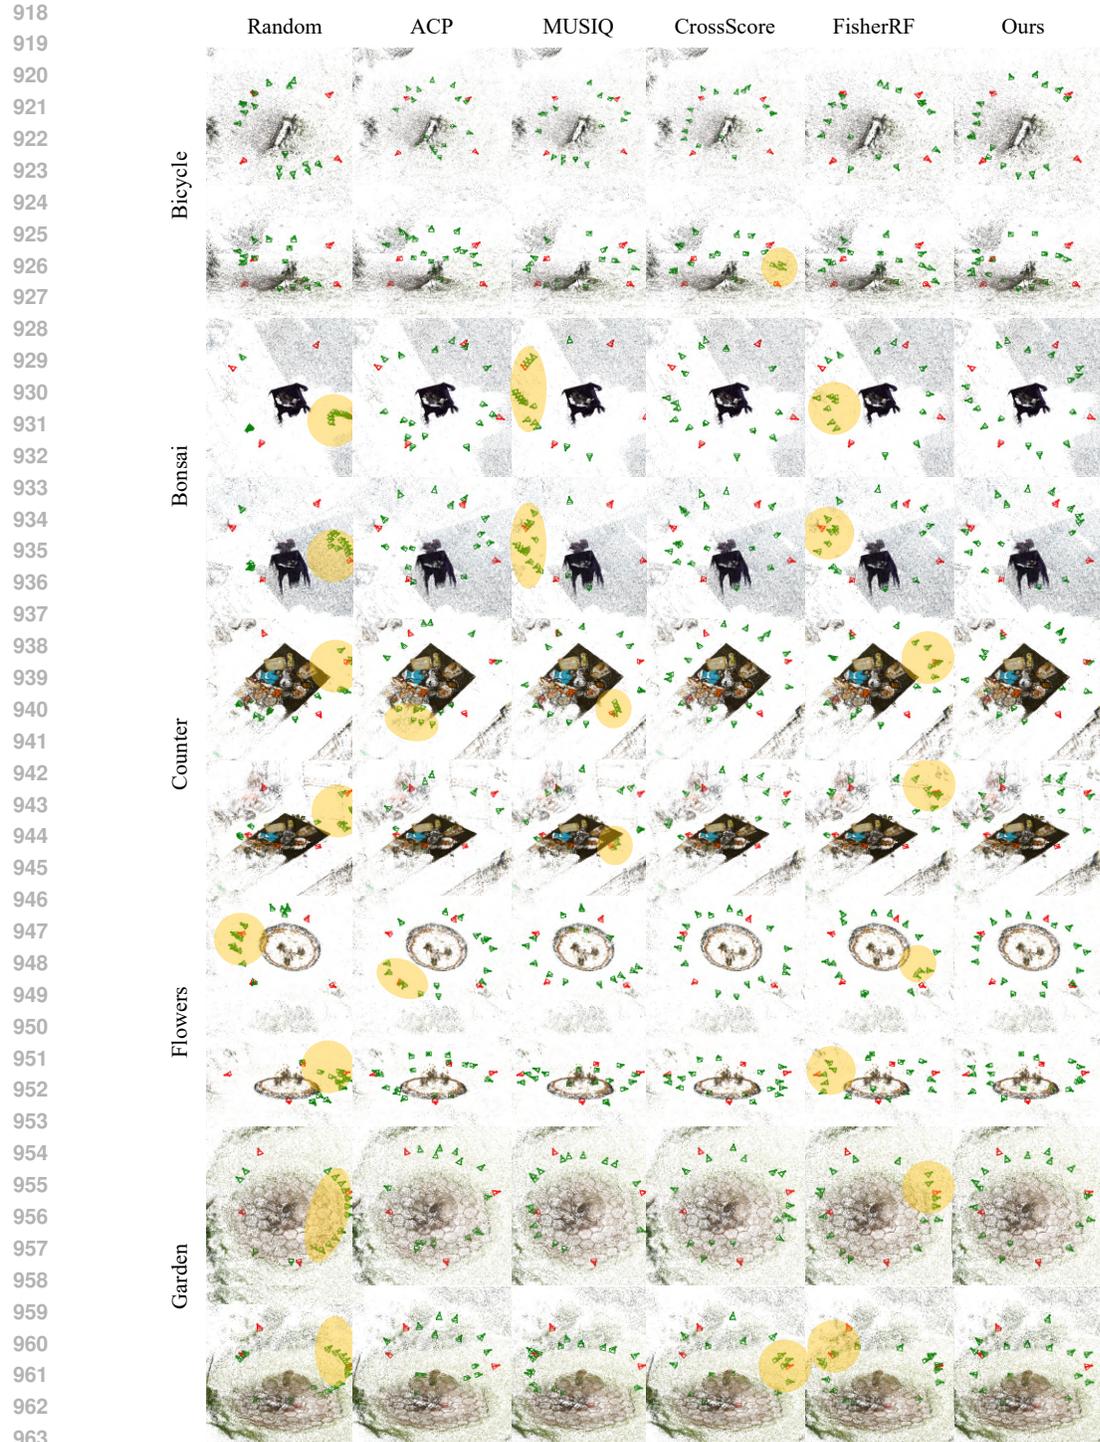


Figure S1: **Camera View Distribution on Deep Blending and Tanks & Temples** View distributions on additional datasets, following the same color and annotation scheme as Fig. S2. Our method maintains broader spatial coverage, while baseline methods often exhibit clustering (yellow circles), consistent with the biases observed in Mip-NeRF 360.



964
965
966
967
968
969
970
971

Figure S2: **Camera View Distribution on Mip-NeRF 360.** Visualization of camera poses selected by each method on the Mip-NeRF 360 dataset. Red frustums indicate the initial views, while green frustums denote views added during active selection. Our method produces a more uniformly distributed set of viewpoints, while both baselines exhibit clustered view selections (highlighted in yellow circle), particularly in FisherRF due to its reliance on uncertainty signals entangled with 3DGS training dynamics.

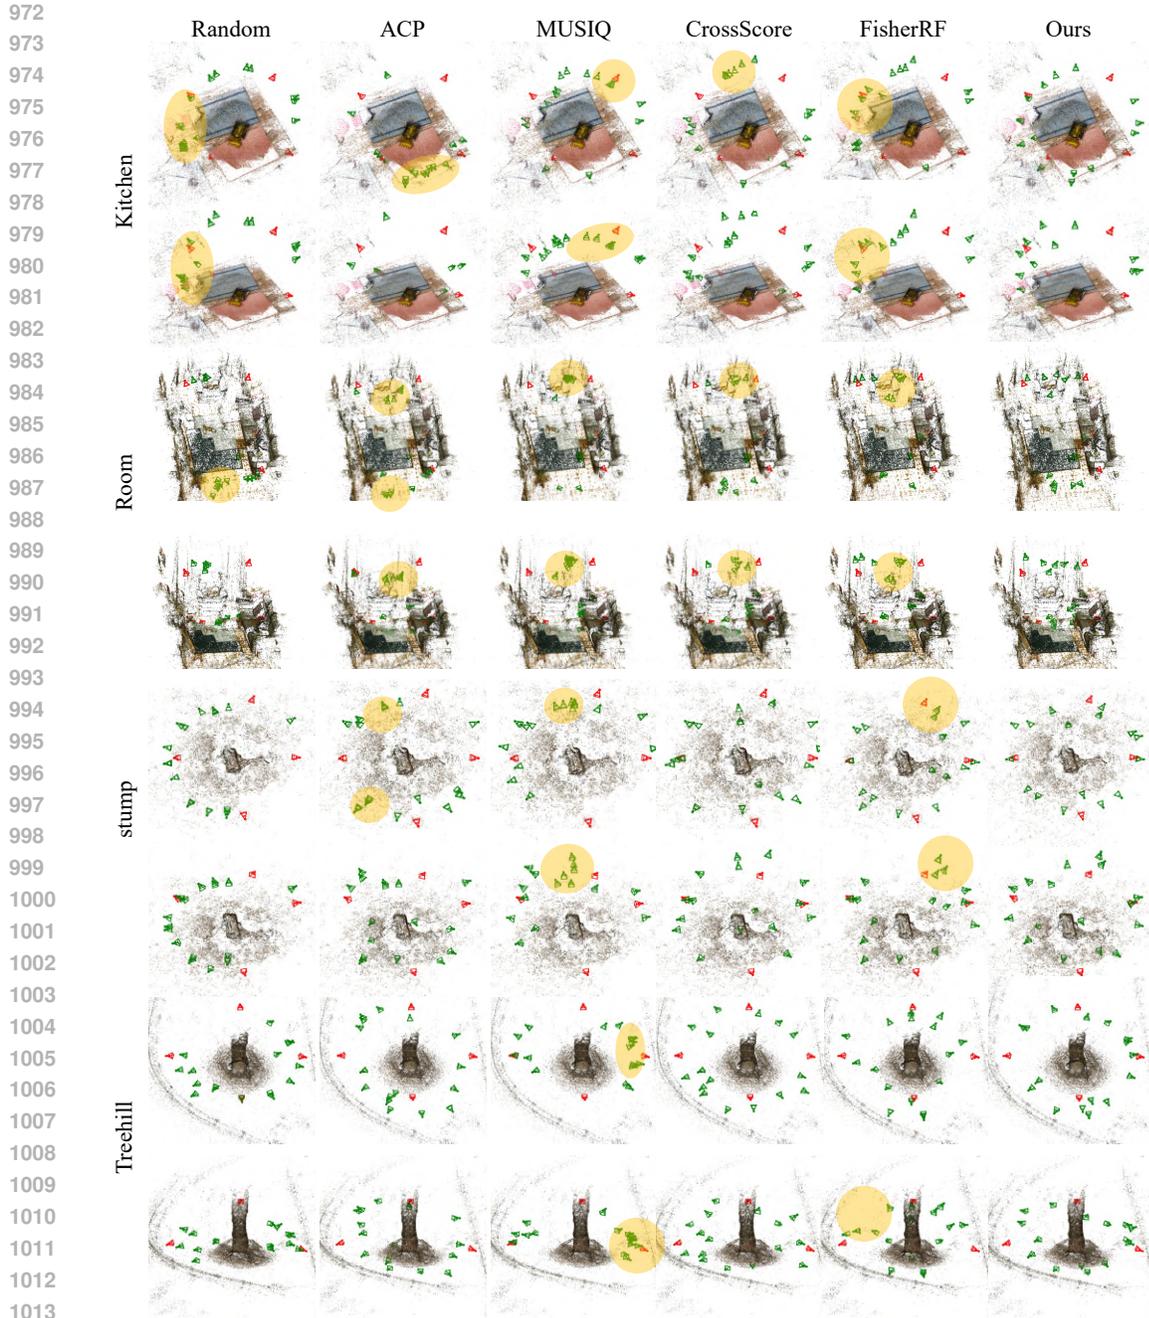


Figure S3: **Camera View Distribution on Mip-NeRF 360.** Visualization of camera poses selected by each method on the Mip-NeRF 360 dataset. Red frustums indicate the initial views, while green frustums denote views added during active selection. Our method produces a more uniformly distributed set of viewpoints, while both baselines exhibit clustered view selections (highlighted in yellow circle), particularly in FisherRF due to its reliance on uncertainty signals entangled with 3DGS training dynamics.

C EXTENDED QUALITATIVE COMPARISONS

The qualitative results presented in the main paper are limited by space constraints, which may obscure the full advantages of our method. To address this, we provide extended visualizations from

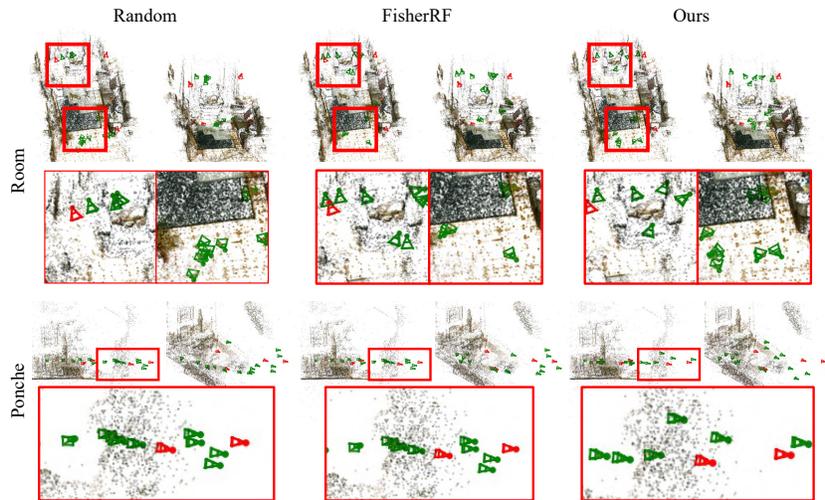
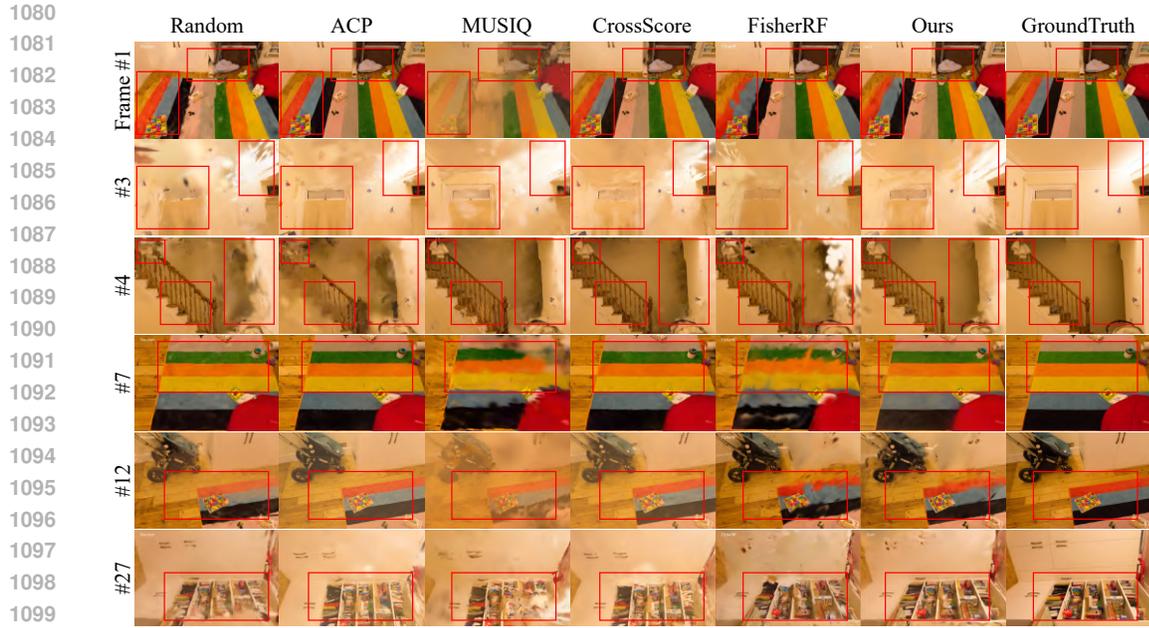


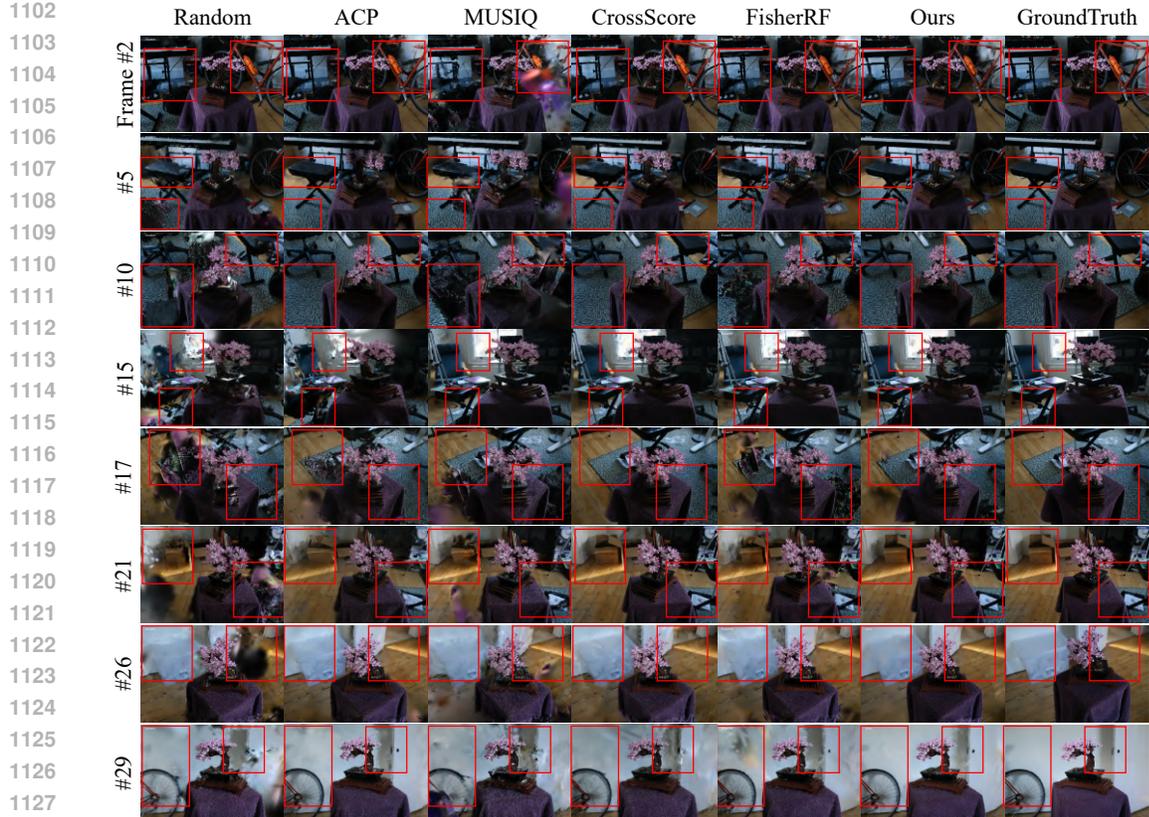
Figure S4: **Diversity of Selected Views (Zoomed-in Analysis)** We highlight the spatial diversity of selected camera poses by presenting zoomed-in regions corresponding to the boxed areas above. Compared to the baselines, our method selects views from a broader range of angles (Room, top) and elevations (Ponche, bottom), resulting in more comprehensive scene coverage.

multiple test viewpoints. Fig. S5a to S7 display results from five scenes across the Mip-NeRF 360, Deep Blending, and Tanks and Temples datasets, with six to eight novel test views per scene. Our method consistently achieves broader and more complete scene coverage compared to the FisherRF.

As discussed in the main paper, our residual supervision strategy further improves geometric consistency and reconstruction robustness, particularly in sparse or limited-view scenarios. This is especially beneficial under the standard protocol of active or next-best-view selection, where training begins with a small number of views (e.g., 4) and progressively adds new views (typically one at a time). The synergy between physically grounded view selection and residual learning enables high-fidelity reconstruction even from limited initial observations. As mentioned in Sec. F, we also include a supplementary video with 360-degree novel view renderings. We encourage reviewers to view this video to better appreciate the improvements in coverage and structural accuracy provided by our proposed SA-ResGS.

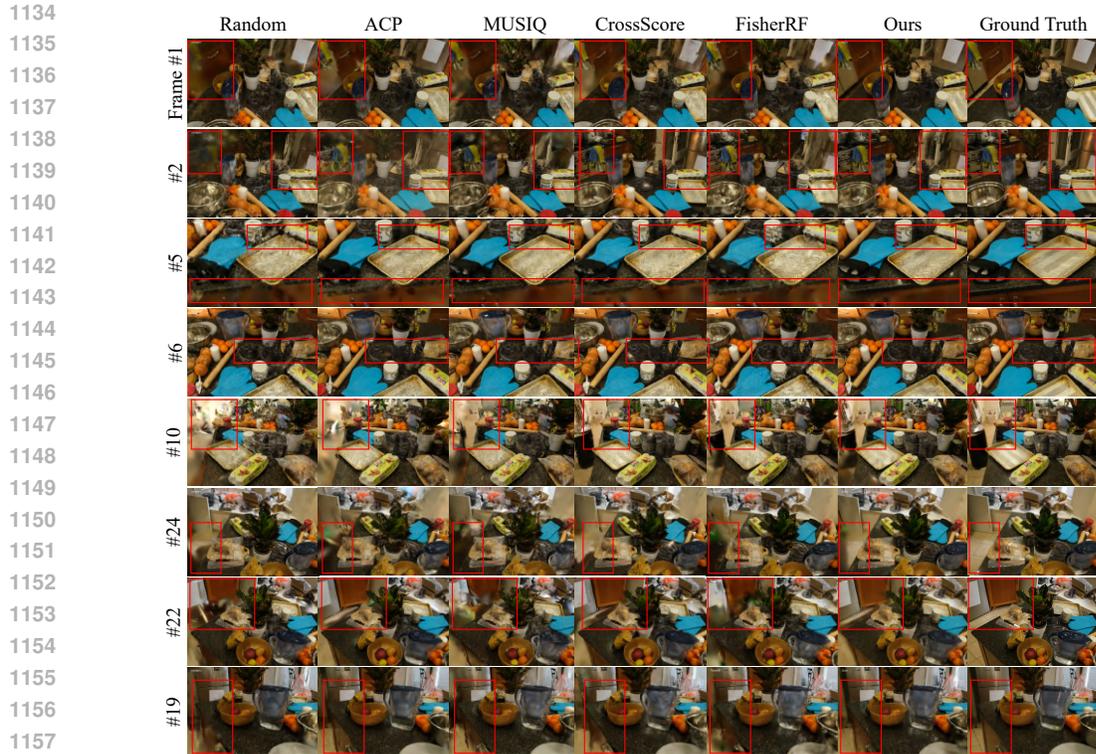


(a) Playroom scene. Six test-time novel views.

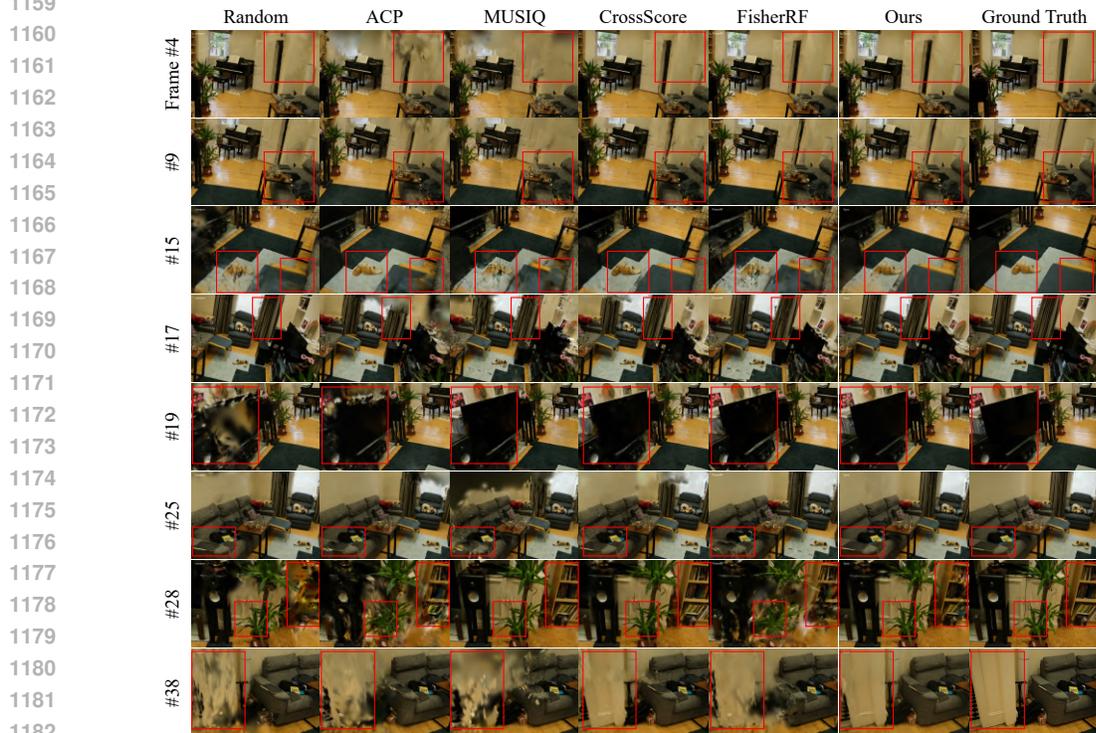


(b) Bonsai scene. Eight test-time novel views.

1130 **Figure S5: Qualitative comparison across multiple test-time novel views.** Across all scenes, our
 1131 method produces more complete and consistent reconstructions, particularly in **occluded or sparsely**
 1132 **observed regions** (red boxes). In contrast, baseline methods (Random, FisherRF) often exhibit
 1133 **missing geometry, blurring, or structural artifacts** due to biased or clustered view selection.



(a) Counter scene. Eight test-time novel views.



(b) Room scene. Eight test-time novel views.

1184
1185
1186
1187

Figure S6: **Qualitative comparison across multiple test-time novel views.** Across all scenes, our method produces more complete and consistent reconstructions, particularly in occluded or sparsely observed regions (red boxes). In contrast, baseline methods (Random, FisherRF) often exhibit missing geometry, blurring, or structural artifacts due to biased or clustered view selection.

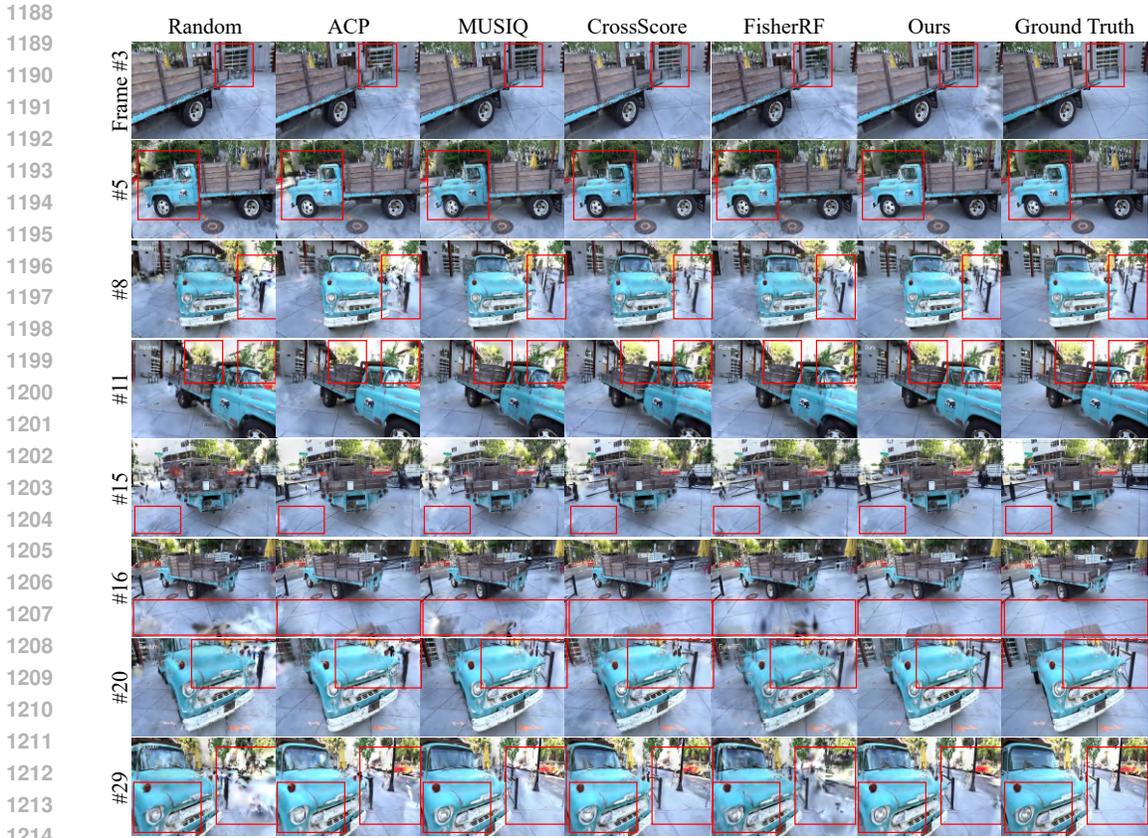


Figure S7: **Qualitative Comparison Across Multiple Test-Time Views (Truck Scene)**. Across all scenes, our method produces more complete and consistent reconstructions, particularly in occluded or sparsely observed regions (red boxes). In contrast, baseline methods (Random, FisherRF) often exhibit missing geometry, blurring, or structural artifacts due to biased or clustered view selection.

D ADDITIONAL RESULTS FOR ABLATION STUDIES

Qualitative Results for Ablation studies. To supplement the quantitative results presented in the ablation study, we provide extended qualitative comparisons in Fig. S9, evaluating different model variants on the **Room** and **Counter** scenes. The variants include the baseline FisherRF, residual supervision with fixed-order view selection (\ddagger +ResGS), our physically grounded view selection method (SA-HashGS), and the full SA-ResGS model combining both components. In the **Room** scene, the baseline FisherRF exhibits missing geometry and artifacts near occluded regions and object boundaries (orange boxes). SA-HashGS mitigates these issues by selecting geometrically diverse viewpoints, leading to improved surface coverage. Residual supervision (\ddagger +ResGS) further refines local details—even in regions already observed by the baseline—demonstrating enhanced reconstruction without additional view coverage, as evidenced by improved toy geometry and floor textures (red boxes). The full SA-ResGS model combines both benefits, producing reconstructions that closely align with ground-truth in terms of both structural completeness and fine detail.

In the **Counter** scene, residual supervision (\ddagger +ResGS) reduces floating artifacts and improves under-optimized regions caused by occlusion, such as the shadow-like area on the left side of the countertop (first row). SA-HashGS enhances global coverage, alleviating severe blurring and recovering geometry in previously unobserved areas (red boxes). The full SA-ResGS model integrates both benefits, yielding sharper and more complete reconstructions by jointly leveraging physically grounded view selection and residual learning. These results validate the complementary roles of physically grounded view selection and uncertainty-guided residual learning. Their integration in SA-ResGS consistently improves reconstruction fidelity under sparse-view settings.

Additional 360-degree renderings for these ablation results are included in the supplementary video and are recommended for further comparison.

Ablation study on Residual Learning. To further analyze the effect of the residual loss in ResGS, we conduct an expanded ablation study in Sec 4.3, as presented in Fig. S8. We compare three configurations: (a) w/ full loss, and w/ subset loss only (b) w/ full Loss, and w/ subset loss using random sampling only, and (c) w/ full Loss, and w/ subset loss using both uncertainty-guided sampling and random sampling. For fair comparison, the sampling ratio in (b) and (c) is matched by increasing the amount of random sampling in (b). To separate the effect of view sampling, we employ fixed-order view selection, replicating the original selection sequence from FisherRF^T, to isolate and examine the effect of residual loss.

Consistent with the observations in the main paper, the absence of the full loss (a) results in pronounced over-smoothing, particularly in ambiguous or low-confidence regions. Moreover, when the supervised loss relies solely on random sampling (b), structural degradation becomes evident in areas that frequently undergo occlusions—such as the staircase behind the door or the window-frame region. This occurs because random sampling alone cannot reliably propagate gradients to Gaussians that require sustained refinement.

In contrast, the proposed full configuration (c), which incorporates both uncertainty-guided and random sampling, consistently preserves fine details and maintains structural coherence. By enforcing continuous gradient flow to under-updated or geometrically unstable Gaussians—particularly those that are bulky, floating, or insufficiently activated—our approach effectively corrects misaligned geometry and enhances robustness in challenging regions. These results highlight the advantage of our residual-learning-based update mechanism, which yields qualitatively more faithful and geometrically consistent renderings across diverse novel viewpoints.

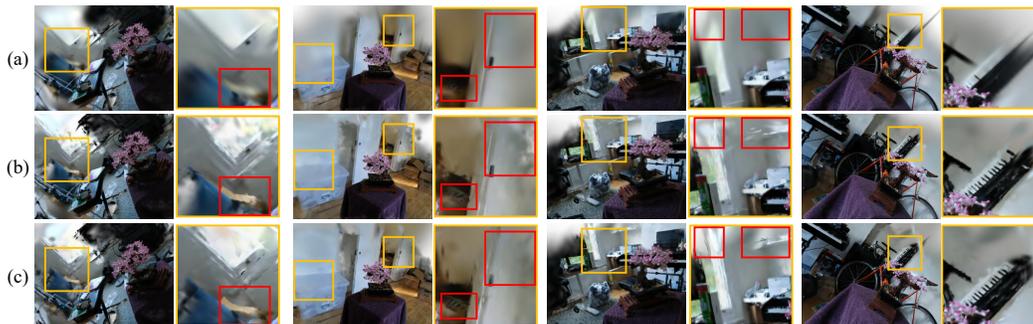


Figure S8: **Qualitative comparison across multiple test-time novel views.** (a) w/ subset loss only, (b) w/ full loss and subset loss using random sampling only, and (c) w/ full loss and subset loss using both uncertainty-guided and random sampling. For a fair comparison between (b) and (c), the overall sampling ratio is matched by increasing the random sampling rate in (b). Highlighted regions (red boxes) denote areas where differences between the variants become especially pronounced.

Ablation study on extended datasets. We additionally evaluated all ablation configurations on the Deep Blending and Tanks & Temples datasets to verify their generality, as shown in Table S1. Across scenes, ResGS consistently improves reconstruction under an identical view sequence, whereas combining ResGS with dynamic selection remains sensitive to early uncertainty noise. SA-HashGS stabilizes early decisions via surface-aware filtering, and the full configuration (SA-HashGS + ResGS) achieves the strongest and most stable improvements across datasets. Note that, the Barn scene was excluded due to repeated convergence failures across some ablation methods.

E ADDITIONAL ABLATION ANALYSIS

Robustness to Correspondence Noise. SA-Points are obtained through triangulation of dense correspondences, so their accuracy can influence reconstruction and view selection. We adopt

1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349

Methods	Our proposed methods		Metrics		
	Sec. 4.2	Sec. 4.4	PSNR	SSIM	LPIPS
FisherRF [†]	-	-	19.293	0.715	0.378
‡+ResGS	-	✓	19.462	0.719	0.381
†+ResGS	-	✓	19.887	0.730	0.374
†+SA-HashGS	✓	-	19.890	0.730	0.364
†+SA-ResGS	✓	✓	20.305	0.740	0.362

Table S1: **Ablation using Deep Blending & Tank and Temples dataset.** ‡ denotes fixed-order view selection, replicating the original selection sequence from FisherRF[†], whereas † indicates dynamically updated view selection based on the model’s progressive training status.

Reprojection error (pixels)	PSNR	SSIM	LPIPS	Coverage
0.5	22.634	0.817	0.332	57.74%
1.0	24.441	0.838	0.317	94.11%
2.0	24.244	0.834	0.318	94.89%

Table S2: **Effect of reprojection filtering thresholds.** A 1-pixel threshold offers the best trade-off between accuracy and coverage, confirming that accurate filtering is essential for stable SA-Points.

reprojection error to filter out inconsistent points, to examine robustness, we conducted experiments with different reprojection thresholds.

Varying the reprojection error threshold further confirmed this robustness. As shown in Table. S2, very strict filtering (0.5 pixels) reduced coverage, while very loose filtering (2.0 pixels) introduced slight inaccuracies. A threshold of 1.0 pixel provided the best trade-off, ensuring sufficient coverage and reliable geometry.

These results demonstrate that SA-Points are resilient to moderate correspondence errors. Our combination of reprojection filtering and voxel dilation effectively balances coverage and accuracy, enabling stable and reliable NBV selection in practice.

Effects of number of selected views. The number of available input views is a critical factor in reconstruction quality, particularly under sparse-view settings. To analyze this effect, we conducted an ablation on the Bonsai scene, starting from 4 fixed initial views and incrementally adding views using our selection strategy. For clarity, we uniformly subsampled training views and report PSNR across different totals.

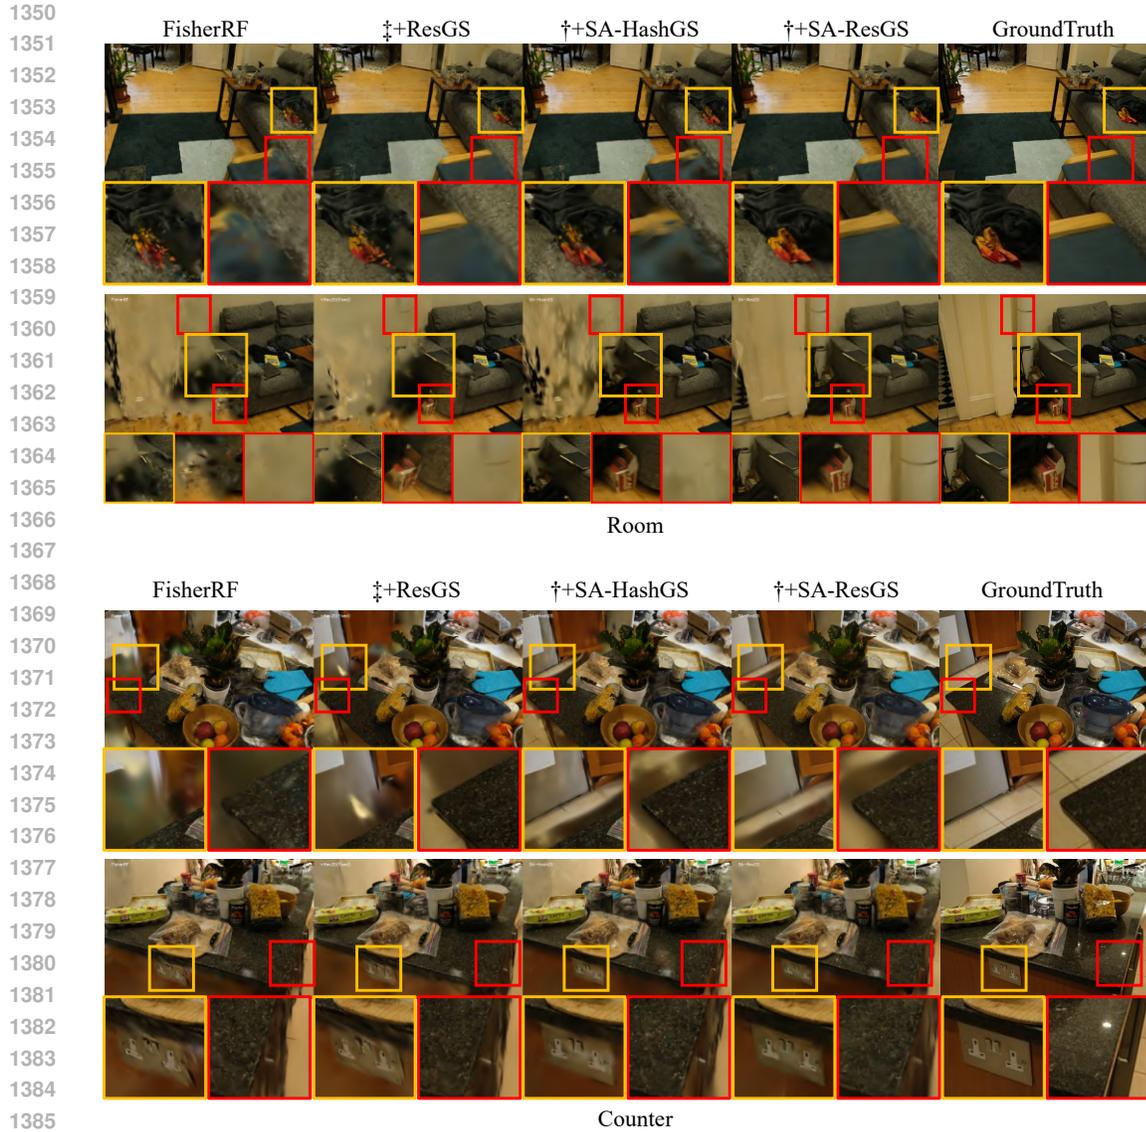
Results in Table S3 that under sparse conditions (7–10 views), our method reaches higher PSNR with fewer inputs compared to FisherRF, demonstrating the benefit of physically grounded filtering in stabilizing early-stage. As more views are added, the initial advantage narrows but new gains appear from 13 views onward, where residual learning further refines Gaussian parameters and improves geometry.

Overall, the method excels in low-view regimes while continuing to scale effectively with more observations, validating its robustness across varying view counts.

Selected Views	4	7	10	13	16	19
random	15.929	16.485	17.573	17.694	18.233	18.400
FisherRF	15.985	17.269	18.773	20.188	21.722	22.650
Ours	15.970	17.606	18.895	20.159	23.229	24.064

Table S3: **Performance scaling with the number of selected views.** Our method yields higher PSNR in sparse-view regimes (7–10 views) and continues to improve as more views are added, highlighting both early stability and long-term scalability.

Ablation on Hash encoding size. To assess the robustness of the hash-encoded voxel grid used in the coverage prefilter, we conducted an ablation study by varying the hash-table size from $2^{11} - 2^{19}$, as



1387 **Figure S9: Qualitative Self-Comparison for Ablation Study.** Results on the **Room** and **Counter**
 1388 scenes comparing FisherRF, ‡+ResGS, SA-HashGS, and SA-ResGS. **Orange boxes** highlight im-
 1389 improvements from coverage-guided view selection, while **red boxes** emphasize the effects of residual
 1390 supervision. Residual learning enhances geometric stability and reduces artifacts (e.g., jittering
 1391 surfaces), whereas SA-HashGS recovers unobserved (e.g., occluded or shadowed areas). The full
 1392 SA-ResGS combines both benefits, yielding the most complete and and structurally faithful recon-
 1393 structions.

1394

1395

1396

1397 well as a no-collision variant implemented with direct indexing. Despite more than a $200\times$ difference
 1398 in hash capacity, the reconstruction accuracy and coverage estimation remained remarkably stable
 1399 across all scenes, exhibiting no consistent trend of degradation, as shown in Table S4.

1400 This insensitivity arises from two structural properties of our pipeline. (1) Only occupied voxels are
 1401 hashed, while empty regions are skipped entirely; because occupancy varies widely across scenes, the
 1402 effective load factor of the hash table remains low even for relatively small hash sizes. (2) Residual
 1403 inconsistencies introduced by collisions are further mitigated by the subsequent Fisher-based fine
 selection stage, which provides an additional layer of error correction.

Together, these factors make the coverage prefilter highly robust to hash collisions in practice. While scenarios with extremely dense occupancy may increase collision likelihood, such cases did not arise in our benchmarks, and our empirical sweep demonstrates that the method maintains stable performance across a wide range of hash sizes.

Hash size	PSNR	SSIM	LPIPS
2^{11}	21.555	0.617	0.446
2^{13}	21.593	0.618	0.446
2^{15}	21.507	0.616	0.447
2^{17}	21.261	0.612	0.450
2^{19}	21.282	0.610	0.451
No Collision	21.325	0.610	0.450

Table S4: **Performance comparison across different hash feature sizes.**

π^3 **implementation.** To evaluate how the choice of dense correspondence model affects our pipeline, we replaced MAST3R with π^3 (Wang et al., 2025a), a recent VGGT-based matcher that predicts poses, point maps, and tracks directly from image pairs. Unlike triangulation-based MAST3R, π^3 provides dense 3D predictions without requiring multi-view geometry, but its outputs lie in an internal coordinate frame and thus require alignment to the global COLMAP frame via ICP.

A full re-evaluation using π^3 shows that reconstruction quality remains highly similar to the MAST3R-based variant, with differences typically within ± 0.1 across PSNR, SSIM, and LPIPS (see Table below). Interestingly, the two backbones introduce complementary error characteristics: MAST3R benefits from accurate scale and pairwise consistency inherited from COLMAP poses, whereas π^3 produces stable background predictions but is sensitive to alignment drift. These error modes counterbalance each other, resulting in comparable end-to-end performance.

Because coverage is measured over sparse, occupied voxels, the advantage of π^3 's denser predictions in background regions is not fully reflected in the final metrics. Overall, this ablation indicates that while the correspondence backbone affects intermediate behavior, it does not strongly influence the final reconstruction quality within our current setup. This observation suggests several promising extensions, including improved global alignment (e.g., SLAM-style refinement), joint optimization of camera poses and SA-Points, or lighter correspondence models for improved efficiency.

Backbone	PSNR \uparrow	SSIM \uparrow	LPIPS \downarrow
MASt3R (Ours)	21.325	0.610	0.450
π^3	21.179	0.609	0.451

Table S5: **Ablation on correspondence backbone.** Comparison between the MAST3R-based pipeline and the π^3 -based variant. Results show differences within ± 0.1 across PSNR, SSIM, and LPIPS.

Statistics of Active view selection for each dataset. We provide per-scene statistics for Active view selection on Mip-NeRF 360 dataset, NeRF Synthetic, and Extended dataset in Table S6 S7& S8. Each experiment is repeated four times with different random seeds.

F SUPPLEMENTARY VIDEO OVERVIEW

To complement the static visualizations provided in this document, we include a supplementary video that offers dynamic and comprehensive renderings of our results. This video is intended to provide a deeper visual understanding of the improvements achieved by our proposed SA-ResGS framework across various evaluation scenarios.

The video includes:

- **Novel View Rendering.** We present extended 360-degree novel view trajectories, captured along spiral and circular camera paths, which go beyond the discrete test views shown in the main paper

Methods	PSNR								
	bicycle	bonsai	counter	flowers	garden	kitchen	room	stump	treehill
random	18.642 ± 0.295	21.086 ± 1.547	20.544 ± 0.450	15.916 ± 0.583	21.131 ± 0.242	21.908 ± 0.671	22.256 ± 0.575	20.027 ± 0.841	18.211 ± 0.401
ACP	19.290 ± 0.213	22.158 ± 0.101	21.390 ± 0.382	16.543 ± 0.192	21.542 ± 0.105	21.279 ± 0.464	21.868 ± 0.604	20.575 ± 0.339	18.278 ± 0.455
MUSIQ	18.332 ± 0.146	20.088 ± 0.375	20.452 ± 0.190	16.946 ± 0.069	20.736 ± 0.245	22.804 ± 0.194	22.057 ± 0.097	19.217 ± 0.297	18.016 ± 0.269
CrossScore	18.485 ± 0.933	24.021 ± 0.330	22.480 ± 0.124	17.222 ± 0.221	21.876 ± 0.099	22.994 ± 0.117	22.924 ± 0.254	21.003 ± 0.544	18.681 ± 0.596
FisherRF	18.715 ± 0.153	23.125 ± 0.733	21.613 ± 0.110	16.616 ± 0.158	21.459 ± 0.101	23.123 ± 0.361	22.500 ± 0.642	20.230 ± 0.316	18.396 ± 0.209
Ours	18.182 ± 0.116	24.564 ± 0.025	22.742 ± 0.029	16.930 ± 0.154	22.182 ± 0.035	24.182 ± 0.111	24.513 ± 0.110	20.605 ± 0.224	18.789 ± 0.374

(a) Average PSNR on Mip-NeRF 360 (4 trials per scene)

Methods	SSIM								
	bicycle	bonsai	counter	flowers	garden	kitchen	room	stump	treehill
random	0.412 ± 0.008	0.758 ± 0.037	0.719 ± 0.011	0.318 ± 0.012	0.578 ± 0.008	0.774 ± 0.022	0.782 ± 0.017	0.457 ± 0.028	0.457 ± 0.009
ACP	0.429 ± 0.006	0.791 ± 0.011	0.746 ± 0.007	0.334 ± 0.005	0.596 ± 0.003	0.757 ± 0.010	0.779 ± 0.015	0.476 ± 0.012	0.458 ± 0.008
MUSIQ	0.400 ± 0.003	0.726 ± 0.016	0.723 ± 0.006	0.333 ± 0.002	0.549 ± 0.009	0.782 ± 0.009	0.780 ± 0.004	0.423 ± 0.011	0.459 ± 0.008
CrossScore	0.397 ± 0.034	0.836 ± 0.002	0.772 ± 0.003	0.340 ± 0.006	0.593 ± 0.006	0.803 ± 0.004	0.811 ± 0.004	0.494 ± 0.021	0.465 ± 0.010
FisherRF	0.411 ± 0.004	0.810 ± 0.016	0.751 ± 0.003	0.331 ± 0.005	0.573 ± 0.004	0.790 ± 0.007	0.773 ± 0.017	0.461 ± 0.015	0.457 ± 0.006
Ours	0.396 ± 0.003	0.841 ± 0.003	0.783 ± 0.001	0.334 ± 0.003	0.584 ± 0.002	0.822 ± 0.002	0.825 ± 0.003	0.473 ± 0.008	0.457 ± 0.004

(b) Average SSIM on Mip-NeRF 360

Methods	LPIPS								
	bicycle	bonsai	counter	flowers	garden	kitchen	room	stump	treehill
random	0.566 ± 0.001	0.364 ± 0.024	0.382 ± 0.010	0.612 ± 0.011	0.417 ± 0.002	0.288 ± 0.018	0.357 ± 0.009	0.556 ± 0.016	0.561 ± 0.006
ACP	0.557 ± 0.003	0.343 ± 0.008	0.362 ± 0.008	0.598 ± 0.003	0.410 ± 0.003	0.303 ± 0.008	0.358 ± 0.012	0.548 ± 0.006	0.563 ± 0.008
MUSIQ	0.579 ± 0.002	0.394 ± 0.014	0.384 ± 0.004	0.601 ± 0.002	0.429 ± 0.004	0.288 ± 0.007	0.381 ± 0.007	0.573 ± 0.005	0.567 ± 0.005
CrossScore	0.605 ± 0.045	0.318 ± 0.003	0.355 ± 0.003	0.609 ± 0.002	0.418 ± 0.006	0.269 ± 0.002	0.354 ± 0.003	0.541 ± 0.010	0.564 ± 0.007
FisherRF	0.571 ± 0.002	0.332 ± 0.013	0.356 ± 0.002	0.603 ± 0.003	0.420 ± 0.002	0.278 ± 0.005	0.370 ± 0.008	0.556 ± 0.009	0.562 ± 0.003
Ours	0.594 ± 0.002	0.313 ± 0.002	0.342 ± 0.001	0.618 ± 0.002	0.431 ± 0.002	0.277 ± 0.043	0.338 ± 0.002	0.570 ± 0.004	0.574 ± 0.003

(c) Average LPIPS on Mip-NeRF 360

Table S6: **Scene-wise quantitative results on Mip-NeRF 360.** Each subtable reports PSNR, SSIM, and LPIPS respectively. Values are averaged over four trials per scene.

and supplementary figures. These renderings highlight the effectiveness of our physically grounded view selection and residual supervision in preserving structural consistency and photometric quality across challenging viewpoints.

- **Ablation Study Comparisons.** To illustrate the impact of each component, we show side-by-side comparisons of different model variants under continuous camera movement. These scenes demonstrate the robustness and fidelity improvements from residual supervision and surface-aware physically grounded view selection, especially under sparse-view or occluded regions.

USE OF LARGE LANGUAGE MODELS

All substantive content was authored by the researchers, with the Large Language Models employed solely for auxiliary polishing of sentence-level expressions for better readability.

1512

1513

Methods	PSNR						
	chair	drums	hotdog	lego	materials	mic	ship
random	24.626 ± 3.489	20.918 ± 1.347	30.397 ± 1.161	28.812 ± 0.285	20.013 ± 0.646	23.813 ± 1.236	25.348 ± 0.830
ACP	25.989 ± 0.219	19.176 ± 0.604	24.409 ± 0.755	23.474 ± 0.104	19.125 ± 0.295	22.687 ± 0.476	24.168 ± 0.726
MUSIQ	27.899 ± 0.164	20.948 ± 1.108	30.308 ± 0.140	28.796 ± 0.159	20.268 ± 0.433	23.219 ± 0.437	25.219 ± 0.335
CrossScore	24.521 ± 0.531	18.817 ± 0.094	28.799 ± 0.661	27.629 ± 0.985	18.326 ± 0.187	23.011 ± 0.378	25.120 ± 0.325
FisherRF	27.066 ± 0.981	21.844 ± 0.608	30.998 ± 0.371	26.108 ± 2.142	20.516 ± 0.488	24.153 ± 0.668	25.645 ± 0.408
Ours	28.303 ± 0.280	22.948 ± 0.304	31.195 ± 0.166	29.703 ± 0.450	21.206 ± 0.187	26.267 ± 0.271	26.437 ± 0.042

1517

1518

(a) Average PSNR on NeRF Synthetic dataset

1519

1520

Methods	PSNR						
	chair	drums	hotdog	lego	materials	mic	ship
random	0.932 ± 0.019	0.878 ± 0.018	0.962 ± 0.005	0.937 ± 0.001	0.814 ± 0.014	0.896 ± 0.015	0.829 ± 0.014
ACP	0.920 ± 0.001	0.849 ± 0.012	0.920 ± 0.007	0.850 ± 0.002	0.804 ± 0.010	0.860 ± 0.005	0.785 ± 0.011
MUSIQ	0.940 ± 0.003	0.880 ± 0.010	0.961 ± 0.001	0.938 ± 0.002	0.817 ± 0.006	0.867 ± 0.003	0.820 ± 0.013
CrossScore	0.930 ± 0.003	0.826 ± 0.008	0.951 ± 0.003	0.925 ± 0.004	0.776 ± 0.007	0.848 ± 0.011	0.823 ± 0.007
FisherRF	0.945 ± 0.003	0.885 ± 0.012	0.965 ± 0.002	0.908 ± 0.021	0.810 ± 0.005	0.898 ± 0.007	0.834 ± 0.007
Ours	0.949 ± 0.002	0.903 ± 0.004	0.965 ± 0.001	0.941 ± 0.004	0.826 ± 0.010	0.918 ± 0.002	0.846 ± 0.002

1525

1526

(b) Average SSIM on NeRF Synthetic dataset

1527

1528

Methods	PSNR						
	chair	drums	hotdog	lego	materials	mic	ship
random	0.070 ± 0.016	0.111 ± 0.009	0.068 ± 0.004	0.076 ± 0.001	0.181 ± 0.008	0.108 ± 0.008	0.202 ± 0.007
ACP	0.075 ± 0.002	0.125 ± 0.006	0.109 ± 0.003	0.121 ± 0.002	0.198 ± 0.007	0.126 ± 0.002	0.213 ± 0.006
MUSIQ	0.062 ± 0.002	0.109 ± 0.007	0.071 ± 0.002	0.077 ± 0.002	0.190 ± 0.005	0.119 ± 0.003	0.204 ± 0.003
CrossScore	0.069 ± 0.002	0.125 ± 0.003	0.088 ± 0.004	0.086 ± 0.003	0.213 ± 0.005	0.123 ± 0.003	0.203 ± 0.004
FisherRF	0.059 ± 0.002	0.108 ± 0.006	0.066 ± 0.003	0.093 ± 0.011	0.177 ± 0.008	0.111 ± 0.003	0.199 ± 0.003
Ours	0.058 ± 0.001	0.099 ± 0.002	0.066 ± 0.001	0.078 ± 0.002	0.175 ± 0.002	0.095 ± 0.002	0.199 ± 0.001

1535

1536

(c) Average LPIPS on NeRF Synthetic dataset

1537

Table S7: **Scene-wise quantitative results on NeRF Synthetic dataset.** Each subtable reports PSNR, SSIM, and LPIPS respectively. Values are averaged over four trials per scene.

1538

1539

1540

Methods	PSNR						
	ballroom	barn	francis	horse	playroom	ponche	truck
random	16.970 ± 0.320	17.639 ± 0.592	18.474 ± 0.361	19.119 ± 0.513	18.506 ± 0.311	19.981 ± 0.899	21.737 ± 0.451
ACP	17.485 ± 0.107	18.804 ± 0.389	18.977 ± 0.551	20.251 ± 0.151	20.807 ± 0.574	21.578 ± 0.113	19.325 ± 0.137
MUSIQ	16.443 ± 0.235	18.021 ± 0.694	18.271 ± 0.708	18.693 ± 0.206	19.131 ± 0.162	18.438 ± 0.394	20.793 ± 0.169
CrossScore	18.099 ± 0.085	19.708 ± 0.253	18.546 ± 0.151	20.354 ± 0.119	19.842 ± 0.150	19.899 ± 0.809	21.517 ± 0.525
FisherRF	17.250 ± 0.326	19.208 ± 0.748	18.708 ± 0.139	19.834 ± 0.396	19.334 ± 0.179	19.833 ± 0.828	22.017 ± 0.096
Ours	18.281 ± 0.179	18.611 ± 0.963	19.803 ± 0.489	20.015 ± 0.232	20.689 ± 1.105	20.906 ± 1.067	22.115 ± 0.256

1545

1546

(a) Average PSNR on Extended datasets

1547

1548

Methods	PSNR						
	ballroom	barn	francis	horse	playroom	ponche	truck
random	0.551 ± 0.012	0.613 ± 0.013	0.747 ± 0.007	0.756 ± 0.011	0.679 ± 0.012	0.753 ± 0.018	0.757 ± 0.008
ACP	0.563 ± 0.003	0.632 ± 0.009	0.762 ± 0.014	0.792 ± 0.006	0.770 ± 0.008	0.752 ± 0.004	0.705 ± 0.008
MUSIQ	0.524 ± 0.008	0.624 ± 0.013	0.744 ± 0.013	0.749 ± 0.008	0.700 ± 0.005	0.723 ± 0.004	0.741 ± 0.002
CrossScore	0.605 ± 0.004	0.669 ± 0.006	0.774 ± 0.002	0.792 ± 0.003	0.724 ± 0.004	0.753 ± 0.012	0.759 ± 0.005
FisherRF	0.555 ± 0.013	0.653 ± 0.017	0.760 ± 0.005	0.777 ± 0.007	0.709 ± 0.006	0.777 ± 0.007	0.709 ± 0.006
Ours	0.614 ± 0.007	0.629 ± 0.018	0.768 ± 0.007	0.779 ± 0.007	0.741 ± 0.016	0.765 ± 0.013	0.761 ± 0.403

1553

1554

(b) Average SSIM on Extended datasets

1555

1556

Methods	PSNR						
	ballroom	barn	francis	horse	playroom	ponche	truck
random	0.385 ± 0.007	0.448 ± 0.013	0.410 ± 0.006	0.286 ± 0.009	0.348 ± 0.011	0.466 ± 0.019	0.388 ± 0.006
ACP	0.376 ± 0.002	0.430 ± 0.006	0.402 ± 0.010	0.256 ± 0.005	0.448 ± 0.006	0.397 ± 0.003	0.329 ± 0.008
MUSIQ	0.411 ± 0.006	0.450 ± 0.012	0.417 ± 0.011	0.298 ± 0.008	0.338 ± 0.003	0.506 ± 0.004	0.401 ± 0.001
CrossScore	0.352 ± 0.003	0.400 ± 0.006	0.383 ± 0.005	0.257 ± 0.002	0.318 ± 0.003	0.466 ± 0.014	0.386 ± 0.003
FisherRF	0.389 ± 0.009	0.416 ± 0.018	0.399 ± 0.005	0.273 ± 0.007	0.324 ± 0.004	0.479 ± 0.016	0.387 ± 0.001
Ours	0.350 ± 0.005	0.436 ± 0.018	0.394 ± 0.008	0.277 ± 0.015	0.338 ± 0.041	0.437 ± 0.031	0.002 ± 0.002

1562

1563

(c) Average LPIPS on Extended datasets

1564

1565

Table S8: **Scene-wise quantitative results on Extended datasets.** Each subtable reports PSNR, SSIM, and LPIPS respectively. Values are averaged over four trials per scene.