

Improving Sparse-View 3DGS Generalization via Flat Minima Optimization

Kangmin Seo[✉], Sangeek Hyun[✉], MinKyu Lee[✉], and Jae-Pil Heo^{★✉}

Sungkyunkwan University, South Korea

Abstract. Recent advances in neural rendering have established 3D Gaussian Splatting (3DGS) as a highly efficient representation for novel view synthesis, enabling fast training and real-time rendering with strong fidelity. However, when supervision is limited to sparse input views, 3DGS tends to overfit to the observed images and generalize poorly to unseen viewpoints. We address this challenge from the perspective of flat minima (FM) optimization, which seeks solutions that remain stable under small parameter perturbations. Viewing Gaussian parameters as trainable weights, we adapt FM principles to the geometric and dynamic nature of 3DGS with a lightweight training framework. Our method regularizes optimization with controlled Gaussian perturbations that account for each Gaussian’s anisotropy and the training progress, preserving fine details while improving robustness to sparse-view overfitting. To further stabilize this flat minima optimization process, we introduce periodic reinitialization, which temporarily returns non-positional parameters to their initial states for a short window. Together, these techniques integrate seamlessly into existing 3DGS pipelines without architectural changes. Experiments on LLFF and Mip-NeRF360 datasets demonstrate improved quantitative metrics and perceptual quality under sparse-view supervision, producing reconstructions that are sharper, more stable, and better generalized to novel viewpoints.

Keywords: 3D Gaussian Splatting · Sparse-View 3D Reconstruction · Novel View Synthesis

1 Introduction

Recent advances in neural rendering have been driven by representative methods such as Neural Radiance Fields (NeRF) [20] and, more recently, 3D Gaussian Splatting (3DGS) [12]. While these approaches have achieved remarkable progress in novel view synthesis (NVS), sparse-view scenarios, where only a few input images are available, remain highly challenging. In such settings, models are prone to overfitting to the input views, leading to poor generalization to novel viewpoints.

In this paper, we aim to tackle this through the lens of flat minima (FM) optimization. Originally developed in the context of neural networks, FM optimization improves generalization by encouraging solutions that reside in flatter

[★] Corresponding author

regions of the loss landscape, where small parameter perturbations do not significantly increase the loss [5, 9, 13]. Inspired by this principle, we reinterpret 3DGS as a supervised learning system where camera poses serve as inputs, rendered images as outputs, and Gaussian parameters play the role of learnable weights. Within this formulation, overfitting in 3DGS corresponds to sharp-minima solutions that are highly sensitive to parameter shifts; resembling the failure modes addressed by FM theory.

Motivated by these insights, we propose to introduce flat minima optimization into the sparse-view 3DGS setting. Specifically, we apply stochastic perturbations to Gaussian parameters during training, following the spirit of prior FM techniques that encourage robustness through parameter noise injection [17, 18]. At each iteration, random noise is added to Gaussian parameters, and the model is trained to minimize reconstruction loss under these perturbations, guiding it toward flatter minima and more generalizable solutions.

Beyond this baseline application of flat minima optimization, we further tailor the approach to the specific characteristics of 3DGS. While directly adding perturbations can improve generalization, naïve strategies often suppress fine-grained geometry, leading to underfitting of local structures. To address this, we adapt FM optimization to better align with the geometric properties and training dynamics of 3DGS.

First, we introduce a Scale-Adaptive Perturbation (SAP) strategy that takes into account the scale and anisotropic shape of each Gaussian. Perturbation noise is sampled in proportion to the Gaussian’s spatial extent, applying larger perturbations along longer axes and smaller ones along shorter axes. This ensures that perturbations remain meaningful across primitives of varying sizes, preventing destabilization for small Gaussians while still enforcing robustness for larger ones. In particular, this design helps preserve fine-grained scene details while maintaining the regularization benefits of flat minima optimization.

We adopt a stochastic perturbation scheme inspired by prior FM studies [17], which combine losses from perturbed and unperturbed models to encourage convergence to flatter solutions. Building on this idea, we instead apply SAP in a stochastic manner. That is, for each iteration, each Gaussian is probabilistically perturbed or left unchanged. This avoids the need to render both perturbed and unperturbed models separately, while keeping the training objective simple and lightweight.

In addition, we complement positional perturbations with periodic parameter reinitialization. Scale, rotation, and higher-order SH coefficients are periodically returned to their initial state, following the same initialization used when constructing Gaussians from SfM points, and kept unchanged for a short window. Opacity is reset at each interval following the standard mechanism in 3DGS [12], while positions remain unaffected. This acts as an additional regularizer that mitigates overfitting and stabilizes training under sparse-view supervision.

Overall, our framework is lightweight, requires no architectural modifications, and integrates seamlessly into existing pipelines. Experiments demonstrate that it consistently improves both quantitative metrics and perceptual quality, pro-

ducing reconstructions that are sharper, more stable, and better generalized to novel viewpoints.

2 Related Work

3D Gaussian Splatting has attracted substantial attention for novel view synthesis, but sparse-view settings remain challenging because limited input views make the optimization underconstrained, often leading to degraded novel view quality. A broad set of recent methods improve robustness by introducing additional geometric cues, imposing constraints on geometric structure, or regularizing primitive interactions.

A common direction is to leverage geometric priors, especially depth-related supervision or normalization, to stabilize geometry under sparse observations. Representative examples include depth-aided training and normalization strategies that reduce geometric ambiguity and improve reconstruction and rendering in sparse-view settings [4, 8, 16, 25, 34, 36]. Related methods also refine geometry through additional constraints or alternative training schedules, such as co-regularization and alternating densification behaviors that affect how primitives are created, split, and pruned during training [22, 31]. Other works focus on sparse-view reconstruction with geometry-prioritized formulations, including approaches that emphasize surface consistency or surface reconstruction objectives [28, 29].

Another line of work improves sparse-view generalization through stochastic regularization of the primitive set. DropGaussian and DropoutGS randomly remove Gaussians during training to discourage brittle reliance on a small subset of primitives and to promote more distributed explanations [21, 30]. The co-adaptation analysis further formalizes this phenomenon by quantifying inter-primitive entanglement and alleviating it with Gaussian dropout and opacity noise injection [3]. These methods primarily regularize the representation structurally by changing which primitives contribute and how primitives interact.

Self-ensembling and consistency-based approaches also improve robustness. SE-GS introduces uncertainty-aware perturbations and self-ensembling objectives to stabilize sparse-view optimization [33], while SCGaussian uses matching priors to enforce structure consistency [23].

A recent method leverages a video diffusion prior with scene-grounding guidance to synthesize additional observations for sparse inputs, which are then used to optimize 3DGS [35]. Difx3D+ enhances reconstruction and novel-view synthesis via single-step diffusion models as a learned prior that helps correct artifacts under sparse-view conditions [27]. MAtCha Gaussians proposes an atlas-of-charts surface representation rendered with 2D Gaussian surfels, initialized with monocular depth and refined for photorealistic rendering from sparse views [7]. These pipelines can be highly effective, but they typically rely on external priors and additional components beyond standard 3DGS optimization.

Finally, structural regularization has been studied to reduce degenerate Gaussian configurations. Prior methods regularize geometry, primitive structure, or

inter-primitive interactions by introducing geometric cues or constraints, penalizing undesirable covariance shapes, or changing how Gaussian primitives contribute and interact [3, 4, 8, 10, 16, 21, 25, 30, 33, 34, 36].

3 Method

3.1 Preliminaries

3D Gaussian Splatting (3DGS) represents a scene using a set of anisotropic Gaussians, each defined by learnable parameters including position $x_i \in \mathbb{R}^3$, scale s_i , rotation r_i , color c_i , and opacity o_i , where i indexes the i -th Gaussian. These Gaussians are projected and composited through differentiable rasterization to render photorealistic images from arbitrary viewpoints.

Given a set of training images and corresponding camera parameters, the Gaussian parameters are optimized by minimizing a photometric reconstruction loss:

$$\theta \leftarrow \theta - \eta \cdot \nabla \mathcal{L}(\theta), \quad (1)$$

where θ encompasses all learnable parameters and $\mathcal{L}(\theta)$ denotes the rendering loss, typically computed as a combination of L1 and SSIM losses between the rendered and corresponding ground-truth images, and η is the learning rate.

Flat Minima (FM) Optimization aims to improve generalization by encouraging the model to converge to regions in the loss landscape where small perturbations do not significantly increase the loss. Here, the representative frameworks are adversarial and stochastic perturbation-based optimization. Sharpness-Aware Minimization (SAM) [6] is a representative of the former method, raising the adversarial perturbation scheme. It minimizes the worst-case loss within a small neighborhood around the parameters as following:

$$\mathcal{L}_{\text{SAM}}(\theta) = \min_{\theta} \max_{\|\epsilon\|_2 \leq \rho} \mathcal{L}(\theta + \epsilon). \quad (2)$$

This min-max formulation aggressively biases optimization toward flatter regions of the loss landscape, but requires an additional gradient computation, which significantly increases the training cost. Meanwhile, Random Weight Perturbation (RWP) [2, 18] provides a more efficient alternative by replacing the worst-case loss with an expected loss over random perturbations:

$$\mathcal{L}_{\text{RWP}}(\theta) = \mathbb{E}_{\epsilon \sim \mathcal{N}(0, \sigma^2 I)} [\mathcal{L}(\theta + \epsilon)]. \quad (3)$$

By averaging losses across randomly perturbed weight configurations, RWP effectively smooths the loss landscape without additional backpropagation. Later works [17] further refined this approach by mixing perturbed and unperturbed losses and adapting noise magnitudes, achieving stronger performance with minimal cost.

In this work, we build upon these insights and adapt flat minima optimization to the context of 3D Gaussian Splatting, as summarized in Fig. 1.




Fig. 1: Overview of our flat minima optimization framework for sparse-view 3D Gaussian Splatting (3DGS). Our approach enhances generalization by introducing perturbations to Gaussian positions during training. To account for the geometric structure, we scale perturbations according to the anisotropic extent of each Gaussian and apply them stochastically at the level of individual Gaussians. In addition, we periodically reinitialize non-positional parameters by temporarily returning scale, rotation, and higher-order SH coefficients to their initial state, following the same initialization method used when constructing Gaussians from SfM points. These parameters are kept unchanged for a short window, while opacity is reset following the standard mechanism in 3DGS.

3.2 Flat Minima Optimization in 3DGS

We begin by formulating flat minima (FM) optimization in the context of 3D Gaussian Splatting (3DGS). While conventional FM approaches in feed-forward neural networks inject random perturbations into all weights to encourage flatter solutions, directly adopting this strategy in 3DGS may overlook its unique geometric structure. Instead, we reinterpret FM optimization by perturbing the 3D position coordinates $\mathbf{x}_i$ of Gaussians during training. From a geometric perspective, overfitting in sparse-view 3DGS manifests as sharp minima that are highly sensitive to positional shifts, which in turn causes vulnerability in novel view rendering. By introducing position-level perturbations, we guide the optimization toward flatter minima where Gaussian positions are more robust under spatial variation. Accordingly, we extend this idea with perturbation and reinitialization techniques that account for the geometric properties and dynamic behavior of 3DGS.

Scale-Adaptive Perturbation (SAP). Prior works on random perturbation strategies [2, 17] have demonstrated the effectiveness of stochastic perturbations in im-




Fig. 2: Optimization procedure of our method. At iteration t , we temporarily perturb the Gaussian model to obtain $\hat{\theta}_t$, render it, and compute the standard 3DGS reconstruction loss $\mathcal{L}(\hat{\theta}_t)$ without introducing an additional loss term or auxiliary objective. The perturbation is then discarded, while $\nabla \mathcal{L}(\hat{\theta}_t)$ updates the original model from θ_t to θ_{t+1} .

proving generalization. To adapt this principle to the geometric nature of 3DGS, we introduce a Scale-Adaptive Perturbation (SAP) scheme that tailors the perturbation magnitude to the anisotropic scale of each Gaussian. By aligning perturbations with the spatial extent of each primitive, SAP effectively balances robustness and detail preservation in 3D reconstruction.

Formally, the perturbed parameters $\hat{\theta}$ are defined as

$$\begin{aligned} \mathbf{x}'_i &= \mathbf{x}_i + \delta_i, \quad \delta_i \sim \mathcal{N}(0, \gamma^2 \Sigma_i), \\ \Sigma_i &= \mathbf{R}_i \mathbf{S}_i \mathbf{S}_i^\top \mathbf{R}_i^\top, \quad \hat{\theta} = \{\mathbf{x}'_i, s_i, r_i, c_i, o_i\}, \end{aligned} \quad (4)$$

where $\mathbf{x}_i$ denotes the 3D position of the i -th Gaussian, s_i is its anisotropic scale, r_i its rotation, c_i its color, and o_i its opacity. Here, $\mathbf{R}_i$ is the rotation matrix constructed from r_i , and $\mathbf{S}_i = \text{diag}(s_i)$ is the scaling matrix whose diagonal elements correspond to the per-axis scales of the Gaussian. As illustrated in Fig. 2, gradients are then computed at the perturbed parameters $\hat{\theta}$, and the update at iteration t is performed accordingly:

$$\theta_{t+1} \leftarrow \theta_t - \eta \cdot \nabla \mathcal{L}(\hat{\theta}_t). \quad (5)$$

Unlike 3DGS-MCMC [14], which injects noise into the parameter update ($\theta_{t+1} = \theta_t - \eta \nabla \mathcal{L}(\theta_t) + \epsilon_t$) so that stochasticity accumulates along the trajectory, our method perturbs only the loss evaluation while keeping the parameters clean, locally smoothing the loss landscape toward flatter minima.

By modulating perturbations according to the anisotropic scale of each Gaussian, SAP enforces stronger regularization on larger or elongated primitives while applying finer noise to smaller ones. This design preserves high-frequency details in compact Gaussians while still promoting robustness across the broader scene geometry.

Stochastic Application of SAP. To further adapt flat minima optimization to 3DGS, we introduce stochastic perturbations, where each Gaussian is probabilistically perturbed or left unchanged. This design is motivated by recent Random

Weight Perturbation strategies [17], where perturbed and unperturbed objectives are mixed to encourage flatter solutions. Instead of explicitly rendering two models and combining their losses, we achieve a similar effect more efficiently by perturbing each Gaussian with probability p at every iteration. In practice, this avoids the need to perform two forward passes, keeping the training objective simple and lightweight.

Formally, for each Gaussian i , the perturbed position $\mathbf{x}'_i$ is defined as

$$\mathbf{x}'_i = \begin{cases} \mathbf{x}_i + \delta_i, & \text{with probability } p, \\ \mathbf{x}_i, & \text{with probability } 1 - p, \end{cases} \quad (6)$$

Applying perturbations in this localized, probabilistic manner offers an additional advantage over global perturbation strategies. In prior FM-inspired methods, where perturbed losses are always incorporated into optimization, we empirically observe a tendency to suppress fine-grained geometry, leading to underfitting of local structures. Similar observations can also be found in our ablation results (see Tab. 3), where the loss-mixing strategy shows degradation in rendering fidelity. By contrast, our per-Gaussian stochastic formulation limits the extent to which perturbations dominate the training signal, allowing the model to retain delicate scene details while still benefiting from the regularization effect. This not only reduces overfitting but also enhances the stability of optimization without incurring additional computational overhead.

Perturbation Magnitude Scheduling. We linearly scale the perturbation magnitude from 0 at the beginning to a target value at the final iteration. This gradual increase avoids excessive noise in the early stage, while ensuring that the desired perturbation strength is applied once the model reaches a more stable regime. Empirically, we find that applying strong perturbations too early, before Gaussians have sufficiently captured the scene structure, tends to degrade performance, motivating this progressive schedule.

Formally, the stochastic perturbation is reformulated to incorporate this scheduling scheme, following the linear annealing strategy used in prior work such as [21]:

$$\mathbf{x}'_i = \begin{cases} \mathbf{x}_i + \alpha(t) \delta_i, & \text{with probability } p, \\ \mathbf{x}_i, & \text{with probability } 1 - p, \end{cases} \quad \text{where } \alpha(t) = \frac{t}{T}, \quad (7)$$

where t is the current iteration, T is the total number of iterations, and $\alpha(t)$ linearly increases the perturbation magnitude from 0 to 1 over training.

Periodic Gaussian Reinitialization. In addition to position-level perturbations, which serve as our main mechanism for flat minima optimization, we periodically reinitialize a subset of Gaussian parameters during training. Specifically, at fixed intervals we temporarily return scale, rotation, and higher-order SH coefficients to their initial state, following the same initialization procedure used when constructing Gaussians from a Structure-from-Motion (SfM) [24] point cloud, and

Table 1: Quantitative comparison on the LLFF dataset under 3-view, 6-view, and 9-view settings. Our method achieves consistently strong performance across all metrics, outperforming or matching prior baselines in PSNR, SSIM, and LPIPS. The results demonstrate the effectiveness of our flat-minima optimization framework in improving both photometric accuracy and perceptual quality.

Method	3-view			6-view			9-view		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3DGS	19.22	0.649	0.229	23.80	0.814	0.125	25.44	0.860	0.096
DNGaussian	19.12	0.591	0.294	22.18	0.755	0.198	23.17	0.788	0.180
FSGS	20.43	0.682	0.248	24.09	0.823	0.145	25.31	0.860	0.122
CoR-GS	20.45	0.712	0.196	24.49	0.837	0.115	26.06	0.874	0.089
DropGaussian	20.76	0.713	0.200	24.74	0.837	0.117	26.21	0.874	0.088
Ours	20.88	0.731	0.184	24.76	0.840	0.114	26.23	0.875	0.088

Table 2: Quantitative comparison on MipNeRF360 dataset under 12-view and 24-view settings.

Method	12-view			24-view		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3DGS	18.52	0.523	0.415	22.80	0.708	0.276
FSGS	18.80	0.531	0.418	23.70	0.745	0.230
CoR-GS	19.52	0.558	0.418	23.39	0.727	0.271
DropGaussian	19.74	0.577	0.364	24.13	0.762	0.225
Ours	19.50	0.584	0.348	24.19	0.771	0.219

keep them unchanged for W iterations. Opacity is reset at each interval following the standard reset mechanism in 3DGS [12]. Positions and the total number of Gaussians remain unaffected. This design complements position perturbations by providing an additional regularization effect that mitigates overfitting under sparse-view supervision. In particular, this return-to-initialization step reduces the effective degrees of freedom of non-positional parameters during the window, which stabilizes optimization. By periodically applying this process during training, we enhance stability and further guide optimization toward flatter minima.

4 Experiments

Experimental Setup. We base our implementation on the original 3DGS framework [12], while following most hyperparameters from DropGaussian [21]. We set the perturbation probability to $p_{\max} = 0.3$, perturbation coefficient $\gamma = 2$, and reinitialize scale, rotation, and higher-order SH coefficients every 1,000 iterations, returning them to their SfM-based initial state for $W = 100$ iterations; opacity follows the standard reset in [12] at each interval. To ensure stability,




Fig. 3: Qualitative comparison in LLFF and Mip-NeRF360 datasets. Compared to prior methods, our approach reconstructs sharper details and more geometrically consistent structures, especially in under-constrained regions such as planar surfaces and occlusion boundaries. While baselines often suffer from spatial misalignment or blurring, our method produces visually coherent and high-fidelity results, demonstrating improved generalization to novel viewpoints.

perturbed positions are clamped so that the displacement does not exceed the corresponding Gaussian’s scale. Experiments were conducted on NVIDIA RTX TITAN and A6000 GPU.

Datasets and Evaluation Metrics. We evaluate our method on two widely used benchmarks for novel view synthesis: LLFF [19] and Mip-NeRF360 [1], both downsampled by a factor of 8. For LLFF, we use 3, 6, and 9 input views, and for Mip-NeRF360, we use 12 and 24 views, following prior works. The reconstruction performance is evaluated by three commonly adopted image quality metrics: peak signal-to-noise ratio (PSNR), structural similarity index (SSIM) [26], and learned perceptual image patch similarity (LPIPS) [32]. Together, these metrics provide a comprehensive evaluation of both photometric accuracy and perceptual quality of novel view synthesis.

Baselines. We compare our method with recent approaches for sparse-view novel view synthesis, including 3DGS [12], CoR-GS [31], and DropGaussian [21]. We also include DNGaussian [16] and FSGS [36], which incorporate additional geometric supervision using depth priors.




Fig. 4: Perturbation robustness analysis on LLFF dataset under 3-view setting. Our method shows smaller performance degradation across perturbation magnitudes compared to 3DGS [12] and DropGaussian [21].

4.1 Quantitative Evaluation

Tab. 1 and Tab. 2 summarize the quantitative results across all sparse-view benchmarks. Our method outperforms the compared prior approaches across datasets and view settings. Notably, improvements are observed in PSNR, SSIM and perceptual quality as measured by LPIPS. These results demonstrate the effectiveness of our flat minima framework in enhancing both reconstruction fidelity and generalization under sparse-view supervision.

4.2 Qualitative Evaluation

Fig. 3 presents qualitative comparisons across various sparse-view scenarios. Compared to existing baselines, our method consistently generates sharper and more structurally coherent renderings, particularly under challenging geometric conditions.

Baseline methods often suffer from noticeable blurring in background regions where supervision is scarce. These methods also tend to exhibit structural inconsistencies, such as spatial misalignment or distorted geometry. For instance, CoR-GS and DropGaussian occasionally fail to maintain consistent geometry, leading to spatial shifts in novel viewpoints, or inaccurately rendering planar surfaces and straight lines (*e.g.*, second row). Such artifacts indicate insufficient regularization or overfitting to the limited input views.

In contrast, our method produces clean and sharper reconstructions across all scenes. The recovered geometry remains consistent under significant viewpoint changes, preserving fine-grained details such as object boundaries and linear structures. These improvements highlight the effectiveness of our flat minima framework in promoting generalizable geometry, and its robustness under sparse-view supervision.

4.3 Perturbation Robustness Comparison

In line with prior works that evaluate the sharpness of minima through a model’s sensitivity to parameter perturbations [5, 6, 13], we examine how the reconstruction

Table 3: Ablation studies on the LLFF dataset under the 3-view setting. † indicates our default setting.

Method	PSNR	SSIM	LPIPS
<i>Noise distribution</i>			
Anisotropic†	20.88	0.731	0.184
Isotropic (max)	20.73	0.725	0.188
Isotropic (mean)	20.67	0.724	0.185
Isotropic (fixed)	20.54	0.714	0.188
<i>Perturbed parameter</i>			
Position†	20.88	0.731	0.184
Rotation	20.23	0.710	0.195
Scale	20.22	0.707	0.195
Opacity	20.21	0.709	0.194
Position + Scale	20.33	0.719	0.199
<i>Training strategy</i>			
Full Method†	20.88	0.731	0.184
Stochastic → Mixed Loss	20.84	0.714	0.212
w/o Scheduling	20.69	0.721	0.191
w/o Reinitialization	20.58	0.713	0.197

tion quality of a trained 3DGS model changes when its Gaussian positions are perturbed after optimization. We apply position perturbations using our Scale-Adaptive Perturbation (SAP) with magnitudes $\gamma \in \{0.0, 0.2, 0.4, 0.6, 0.8, 1.0\}$ and evaluate the resulting models on both training and test views.

As shown in Fig. 4, our method shows substantially smaller degradation than both 3DGS [12] and DropGaussian [21] across all perturbation magnitudes. On test views, DropGaussian nearly matches the degradation trend of vanilla 3DGS, while our method shows substantially smaller drops at every scale. The difference is more evident in training views as vanilla 3DGS and DropGaussian exhibit larger degradation, indicating a stronger tendency to overfit to training images compared to our method. This suggests that our optimization strategy leads to flatter and more robust minima than existing 3DGS baselines.

4.4 Ablation Studies

Perturbation Design. In Tab. 3, we evaluate different strategies for injecting perturbations into Gaussian positions. Our default configuration leverages anisotropic noise that is proportional to the scale of Gaussian axes, promoting both detail preservation and robustness. For comparison, we test isotropic variants where perturbations are scaled by the longest axis of each Gaussian, by the mean axis length, or fixed uniformly across all Gaussians at half the size threshold for determining clone/split targets. Among these, anisotropic perturbation achieves the best balance between fidelity and regularization, while isotropic variants either

oversmooth fine structures or under-regularize large primitives. These results confirm the importance of accounting for the geometric shape of each Gaussian when designing perturbations. Nevertheless, we find that any perturbation strategy yields better generalization than the vanilla 3DGS baseline, highlighting that introducing perturbation itself acts as a strong form of regularization, even when the design is not optimal.

Effect of Perturbing Different Parameters. To assess the sensitivity of different Gaussian attributes, we applied perturbations individually to position, scale, rotation, and opacity, as reported in Tab. 3. For scale, noise was added proportional to each axis length, following our SAP strategy. Rotation was perturbed by altering each Gaussian’s orientation with small random angular offsets around the three principal axes. Opacity was perturbed by adding Gaussian noise directly to its values.

The results indicate that perturbing position is the most effective for improving generalization, consistently outperforming other parameter choices. Interestingly, applying perturbations to both position and scale does not further enhance performance; instead, its performance degraded compared to perturbing position alone. These observations suggest that inducing flatness in the position space is particularly impactful for promoting stable and generalizable geometry in 3DGS, and that geometry-based perturbation plays an important role in our framework, while other parameter groups may still be beneficial under different schedules or formulations and thus remain worth exploring in future work.

Stochastic Perturbation Strategy. We further investigate the role of stochasticity in perturbation design. Prior flat minima studies [17] improve generalization by mixing losses from perturbed and unperturbed models, but this requires rendering twice per iteration. Instead, we realize a similar effect by applying perturbations stochastically during training. Concretely, each Gaussian is randomly perturbed with some probability at every iteration, which introduces variability into the optimization while keeping the training objective simple and lightweight. We found that directly relying on perturbed losses at every step tends to oversmooth high-frequency details, whereas our stochastic scheme alleviates this issue and preserves fine structures. This can also be observed in our ablation results (Tab. 3).

Perturbation Magnitude Scheduling. We further validate the role of perturbation scheduling by comparing models trained with and without the progressive scaling of perturbation magnitude (Tab. 3). Without scheduling, strong perturbations are applied from the very beginning of training, when Gaussians have not yet captured the coarse scene structure. As a result, optimization is hindered, leading to unstable convergence and degraded reconstruction quality. In contrast, our scheduling strategy gradually increases perturbation strength as training progresses, allowing Gaussians to first establish a stable representation of the scene before stronger perturbations are introduced. This results in consistently better quantitative performance, supporting our intuition that early




Fig. 5: Training and test PSNR trajectories with and without Gaussian reinitialization on the LLFF fern scene under the 3-view setting. Without reinitialization, training PSNR keeps rising steadily while test PSNR plateaus, indicating overfitting. In contrast, periodic reinitialization moderates the growth of training PSNR but enables continuous improvements in test PSNR, suggesting better generalization. The sharper drops observed at each reset interval are due to reinitializing more diverse parameters, but performance quickly recovers and ultimately surpasses the non-reinitialized baseline. This aligns with the role of periodic reinitialization as an additional regularizer that mitigates overfitting and improves generalization.

excessive noise is detrimental, while a progressive schedule provides a more effective balance between regularization and fidelity.

Periodic Gaussian Reinitialization. To validate the effect of our reinitialization strategy, we conducted an ablation study on the LLFF dataset under the 3-view setting (Tab. 3, Fig. 5). When reinitialization is disabled, the additional periodic return-to-initialization of scale, rotation, and higher-order SH coefficients is deactivated, while opacity continues to follow the original reset mechanism introduced in [12]. Under this setting, performance drops across all metrics, with PSNR decreasing by 0.3. This indicates that augmenting the standard opacity reset with periodic return-to-initialization provides a mild regularization effect during training. It serves as a complementary mechanism alongside Scale-Adaptive Perturbation, jointly mitigating overfitting and stabilizing optimization.

4.5 Plug-in Compatibility

Since our method acts on the optimization process rather than introducing architectural changes, it can be naturally integrated into existing sparse-view 3DGS pipelines as a plug-in module. To verify this, we apply our method to optimization-based methods (FSGS [36] and DropGaussian [21]), diffusion-based artifact correction framework (Difix3D+ [27]), and feedforward initialization followed by short post-optimization (AnySplat [11]), without modifying their original pipeline designs. As shown in Tab. 4, incorporating our optimization strategy

Table 4: Plug-in compatibility across optimization-based, diffusion-enhanced, and feed-forward sparse-view 3DGS pipelines.

Dataset	Method	PSNR	SSIM	LPIPS
<i>Optimization-based</i>				
LLFF [19] 3-view	FSGS [36]	20.43	0.682	0.248
	FSGS + Ours	21.03	0.725	0.196
	DropGaussian [21]	20.76	0.713	0.200
	DropGaussian + Ours	21.05	0.734	0.192
<i>Diffusion-enhanced correction</i>				
Mip-NeRF360 [1] 12-view	Difix3D+ [27]	19.25	0.554	0.375
	Difix3D+ + Ours	19.63	0.575	0.386
Tanks & Temples [15] 12-view	Difix3D+ [27]	14.07	0.499	0.494
	Difix3D+ + Ours	14.35	0.517	0.516
<i>Feed-forward</i>				
Mip-NeRF360 [1] 12-view	AnySplat [11]	15.76	0.307	0.481
	AnySplat + 3DGS 3k	16.97	0.338	0.432
	AnySplat + Ours 3k	17.36	0.384	0.427
	AnySplat [11]	12.67	0.308	0.506
Tanks & Temples [15] 12-view	AnySplat + 3DGS 3k	13.58	0.325	0.469
	AnySplat + Ours 3k	14.07	0.394	0.445

consistently improves FSGS, DropGaussian, and AnySplat, while also improving PSNR and SSIM when combined with Difix3D+. This suggests that the flat minima regularization provided by our approach is complementary to existing techniques, and can be readily adopted to further improve their performance under sparse-view supervision.

5 Conclusion

We addressed the challenge of sparse-view generalization in 3D Gaussian Splatting (3DGS) from the perspective of flat minima optimization. Viewing overfitting as convergence to sharp solutions that are sensitive to positional perturbations, we proposed a lightweight training framework that adapts flat minima principles to 3DGS through scale-adaptive perturbation with stochastic application and magnitude scheduling, together with periodic reinitialization of non-positional parameters. These components regularize optimization while preserving fine details, requiring no architectural changes or external priors. Experiments on LLFF and Mip-NeRF360 datasets demonstrate improved stability and generalization under sparse-view supervision, producing sharper and more robust reconstructions. Our results also suggest that further exploring perturbation strategies for other Gaussian parameters may provide additional opportunities to improve 3DGS optimization.

Acknowledgements

This work was supported in part by MSIT/IITP (No. RS-2022-II220680, RS-2020-II201821, RS-2019-II190421, RS-2024-00459618, RS-2024-00360227, RS-2024-00437633, RS-2024-00437102, RS-2025-25442569), MSIT/NRF (No. RS-2024-00357729), and KNPA/KIPoT (No. RS-2025-25393280).

References

1. Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mipnerf 360: Unbounded anti-aliased neural radiance fields. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5470–5479 (2022)
2. Bisla, D., Wang, J., Choromanska, A.: Low-pass filtering SGD for recovering flat optima in the deep learning optimization landscape. In: International Conference on Artificial Intelligence and Statistics. pp. 8299–8339. PMLR (2022)
3. Chen, K., Zhong, Y., Li, Z., Lin, J., Chen, Y., Qin, M., Wang, H.: Quantifying and alleviating co-adaptation in sparse-view 3d gaussian splatting. *Advances in Neural Information Processing Systems* **38**, 115939–115968 (2026)
4. Chung, J., Oh, J., Lee, K.M.: Depth-regularized optimization for 3d gaussian splatting in few-shot images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 811–820 (2024)
5. Dinh, L., Pascanu, R., Bengio, S., Bengio, Y.: Sharp minima can generalize for deep nets. In: International Conference on Machine Learning. pp. 1019–1028. PMLR (2017)
6. Foret, P., Kleiner, A., Mobahi, H., Neyshabur, B.: Sharpness-aware minimization for efficiently improving generalization. *arXiv preprint arXiv:2010.01412* (2020)
7. Guédon, A., Ichikawa, T., Yamashita, K., Nishino, K.: Matcha gaussians: Atlas of charts for high-quality geometry and photorealism from sparse views. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 6001–6011 (2025)
8. Guo, Z., Wang, P., Chen, Z., Kong, X., Lyu, Y., Gao, G., Han, L.: Ugod: Uncertainty-guided differentiable opacity and soft dropout for enhanced sparse-view 3dgs. *arXiv preprint arXiv:2508.04968* (2025)
9. Hochreiter, S., Schmidhuber, J.: Flat minima. *Neural computation* **9**(1), 1–42 (1997)
10. Hyung, J., Hong, S., Hwang, S., Lee, J., Choo, J., Kim, J.H.: Effective rank analysis and regularization for enhanced 3d gaussian splatting. *Advances in Neural Information Processing Systems* **37**, 110412–110435 (2024)
11. Jiang, L., Mao, Y., Xu, L., Lu, T., Ren, K., Jin, Y., Xu, X., Yu, M., Pang, J., Zhao, F., et al.: Anysplat: Feed-forward 3d gaussian splatting from unconstrained views. *ACM Transactions on Graphics (TOG)* **44**(6), 1–16 (2025)
12. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., et al.: 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.* **42**(4), 139–1 (2023)
13. Keskar, N.S., Mudigere, D., Nocedal, J., Smelyanskiy, M., Tang, P.T.P.: On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836* (2016)
14. Kheradmand, S., Rebain, D., Sharma, G., Sun, W., Tseng, Y.C., Isack, H., Kar, A., Tagliasacchi, A., Yi, K.M.: 3d gaussian splatting as markov chain monte carlo. *Advances in Neural Information Processing Systems* **37**, 80965–80986 (2024)

15. Knapitsch, A., Park, J., Zhou, Q.Y., Koltun, V.: Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)* **36**(4), 1–13 (2017)
16. Li, J., Zhang, J., Bai, X., Zheng, J., Ning, X., Zhou, J., Gu, L.: Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 20775–20785 (2024)
17. Li, T., Tao, Q., Yan, W., Lei, Z., Wu, Y., Fang, K., He, M., Huang, X.: Revisiting random weight perturbation for efficiently improving generalization. *arXiv preprint arXiv:2404.00357* (2024)
18. Li, T., Yan, W., Lei, Z., Wu, Y., Fang, K., Yang, M., Huang, X.: Efficient generalization improvement guided by random weight perturbation. *arXiv preprint arXiv:2211.11489* (2022)
19. Mildenhall, B., Srinivasan, P.P., Ortiz-Cayon, R., Kalantari, N.K., Ramamoorthi, R., Ng, R., Kar, A.: Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM Transactions on Graphics (ToG)* **38**(4), 1–14 (2019)
20. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
21. Park, H., Ryu, G., Kim, W.: Dropgaussian: Structural regularization for sparse-view gaussian splatting. In: *Proceedings of the computer vision and pattern recognition conference*. pp. 21600–21609 (2025)
22. Patle, G., Girgaonkar, N., Somraj, N., Soundararajan, R.: Ad-gs: Alternating densification for sparse-input 3d gaussian splatting. In: *Proceedings of the SIGGRAPH Asia 2025 Conference Papers*. pp. 1–11 (2025)
23. Peng, R., Xu, W., Tang, L., Liao, L., Jiao, J., Wang, R.: Structure consistent gaussian splatting with matching prior for few-shot novel view synthesis. *Advances in Neural Information Processing Systems* **37**, 97328–97352 (2024)
24. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4104–4113 (2016)
25. Song, M., Lin, X., Zhang, D., Li, H., Li, X., Du, B., Qi, L.: D²gs: Depth-and-density guided gaussian splatting for stable and accurate sparse-view reconstruction. *arXiv preprint arXiv:2510.08566* (2025)
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
27. Wu, J.Z., Zhang, Y., Turki, H., Ren, X., Gao, J., Shou, M.Z., Fidler, S., Gojcic, Z., Ling, H.: Difx3d+: Improving 3d reconstructions with single-step diffusion models. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 26024–26035 (2025)
28. Wu, J., Li, R., Zhu, Y., Guo, R., Sun, J., Zhang, Y.: Sparse2dgs: Geometry-prioritized gaussian splatting for surface reconstruction from sparse views. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 11307–11316 (2025)
29. Xiong, H., Muttukuru, S., Xiao, H., Upadhyay, R., Chari, P., Zhao, Y., Kadambi, A.: Sparsegs: sparse view synthesis using 3d gaussian splatting. In: *2025 International Conference on 3D Vision (3DV)*. pp. 1032–1041. IEEE (2025)

30. Xu, Y., Wang, L., Chen, M., Ao, S., Li, L., Guo, Y.: Dropouts: Dropping out gaussians for better sparse-view rendering. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 701–710 (2025)
31. Zhang, J., Li, J., Yu, X., Huang, L., Gu, L., Zheng, J., Bai, X.: Cor-gs: sparse-view 3d gaussian splatting via co-regularization. In: European conference on computer vision. pp. 335–352. Springer (2024)
32. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
33. Zhao, C., Wang, X., Zhang, T., Javed, S., Salzmänn, M.: Self-ensembling gaussian splatting for few-shot novel view synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4940–4950 (2025)
34. Zheng, Y., Jiang, Z., He, S., Sun, Y., Dong, J., Zhang, H., Du, Y.: Nexusgs: Sparse view synthesis with epipolar depth priors in 3d gaussian splatting. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 26800–26809 (2025)
35. Zhong, Y., Li, Z., Chen, D.Z., Hong, L., Xu, D.: Taming video diffusion prior with scene-grounding guidance for 3d gaussian splatting from sparse inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6133–6143 (2025)
36. Zhu, Z., Fan, Z., Jiang, Y., Wang, Z.: Fsgs: Real-time few-shot view synthesis using gaussian splatting. In: European conference on computer vision. pp. 145–163. Springer (2024)