Blur Issue Matters for Thermal Novel View Synthesis: A Floating Gaussian Suppression Approach

Mingyuan Xie[1] Sheng Wan[2] Jin Xie[3] Hang Yang[1] Liangyun Sun[1] Chen Gong[4]*

[1] Nanjing University of Science and Technology [2] Nanjing Agricultural University
[3] Nanjing University [4] Shanghai Jiao Tong University

* Correspondence: chen.gong@sjtu.edu.cn

NeurIPS 2026

Introduction

Existing thermal 3D Gaussian Splatting (3DGS) methods often produce local blur. Near-camera floating Gaussians become over-opaque and occlude distant textures.

FGS controls this opacity growth during early training while preserving thermal appearance through separate luminance updates.

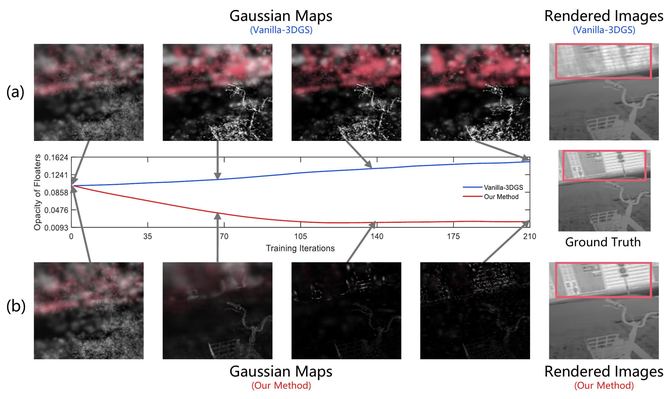
Why Does Local Blur Occur?

Single-channel intensity errors can be reduced by increasing opacity, even for floaters in front of sparsely represented backgrounds.

$$\frac{\partial \mathcal{L}}{\partial \sigma_i} = \sum_{u \in \mathcal{U}} 2(\widehat{\mathbf{C}}^u - \mathbf{C}^u)^\top (\mathbf{c}_i - \widehat{\mathbf{C}}_i^u) \cdot \prod_{j=1}^{i-1} (1 - \alpha_j^u) \cdot G_i^u.$$

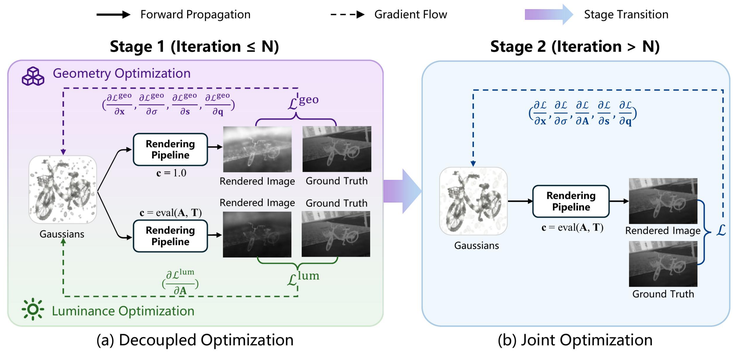
When the background contribution is small and the rendered intensity is too low, gradient descent tends to increase floater opacity. RGB observations impose stronger cross-channel constraints.

Evolution of Floating Gaussians



Red points mark near-camera floaters. FGS lowers their opacity and restores distant details (Figure 1).

Floating Gaussian Suppression



Two-stage optimization suppresses floater opacity while preserving view-dependent thermal appearance.

Stage 1: Decoupled Optimization

Geometry optimization (GO) fixes luminance at 1.0 and updates position, opacity, scale and rotation. This over-bright rendering encourages floaters to become transparent.

Luminance optimization (LO) updates only the spherical harmonic (SH) coefficients, using normal view-dependent luminance.

Stage 2: Joint Optimization

Jointly update all parameters. Reduced floater opacity increases transmittance to valid background Gaussians, helping recover scene detail.

Built on Veta-GS. Stage 1 lasts 500 iterations for most scenes.

Ablation Study

PSNR (dB) ↑	TI-NSD	ThermoScenes
Veta-GS (no GO / LO)	35.703	34.511
FGS without GO	35.751	33.002
FGS without LO	35.976	35.521
Full FGS	**36.330**	**35.908**

GO and LO provide complementary gains (Table 2).

Experiments

TI-NSD: 20 scenes. ThermoScenes: 10 scenes.
Metrics: PSNR, SSIM and LPIPS.

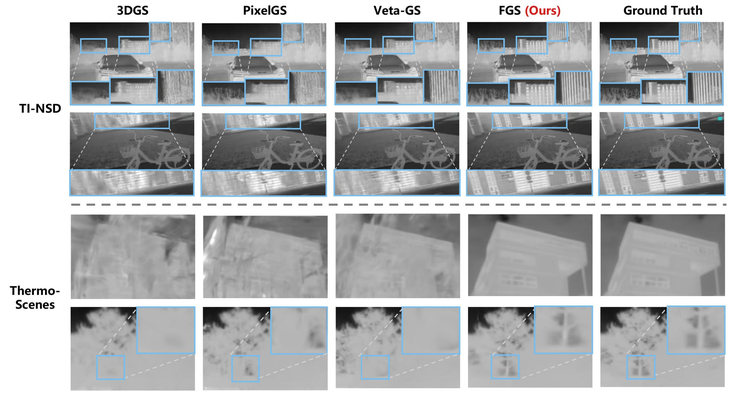
Average over both benchmarks

Method	PSNR ↑	SSIM ↑	LPIPS ↓
InstantNGP-Base	25.618	0.847	0.271
InstantNGP-Big	25.973	0.848	0.264
Plenoxel	23.877	0.839	0.315
3DGS	32.144	0.950	0.156
PixelGS	33.613	0.961	0.142
AD-GS	33.618	0.959	0.153
NTR-Gaussian	27.138	0.882	0.222
Thermal3D-GS	34.114	0.961	0.147
Veta-GS	35.305	0.964	0.137
FGS (Ours)	**36.189**	**0.967**	**0.132**

+0.884 dB over Veta-GS

FGS leads all compared methods on the three reported average metrics (Table 1).

Qualitative Comparison



FGS reduces local blur and preserves sharper textures and edges across both datasets (Figure 3).