FSGS: Real-Time Few-shot View Synthesis using Gaussian Splatting

Zehao Zhu^{®*}, Zhiwen Fan^{*†}[®], Yifan Jiang[®], Zhangyang Wang[®]

[†] Project Lead The University of Texas at Austin

Abstract. Novel view synthesis from limited observations remains a crucial and ongoing challenge. In the realm of NeRF-based few-shot view synthesis, there is often a trade-off between the accuracy of the synthesized view and the efficiency of the 3D representation. To tackle this dilemma, we introduce a $\underline{\mathbf{F}}$ ew- $\underline{\mathbf{S}}$ hot view synthesis framework based on 3D Gaussian Splatting, which facilitates real-time, photo-realistic synthesis from a minimal number of training views. FSGS employs an innovative Proximity-guided Gaussian Unpooling, specifically designed for sparse-view settings, to bridge the gap presented by the sparse initial point sets. This method involves the strategic placement of new Gaussians between existing ones, guided by a Gaussian proximity score, enhancing the adaptive density control. We have identified that Gaussian optimization can sometimes result in overly smooth textures and a propensity for overfitting when training views are limited. To mitigate these issues, FSGS introduces the synthesis of virtual views to replicate the parallax effect experienced during training, coupled with geometric regularization applied across both actual training and synthesized viewpoints. This strategy ensures that new Gaussians are placed in the most representative locations, fostering more accurate and detailed scene reconstruction. Our comprehensive evaluation across various datasets-including NeRF-Synthetic, LLFF, Shiny, and Mip-NeRF360 datasets---illustrates that FSGS not only delivers exceptional rendering quality but also achieves an inference speed more than 2000 times faster than existing state-of-the-art methods for sparse-view synthesis. Project webpage: https://zehaozhu.github.io/FSGS/.

Keywords: Neural Rendering \cdot Gaussian Splatting \cdot Sparse View

1 Introduction

Novel view synthesis (NVS) from a set of view collections, as demonstrated by recent works [9,18,42], has played a critical role in the domain of 3D vision and is pivotal in many applications, e.g., VR/AR and autonomous driving. Despite its effectiveness in photo-realistic rendering, the requirement of dense support views has hindered its practical usages [39]. Previous studies have focused on

^{*} Equal Contribution.

2 Zehao Zhu[®], Zhiwen Fan[®], Yifan Jiang[®], Zhangyang Wang[®]



Fig. 1: Real-Time Few-shot Novel View Synthesis. We present a optimizationbased framework that is initialized from COLMAP point clouds, achieving a significantly faster rendering speed $(2900 \times)$ while enhancing the visual quality (from 0.684 to 0.745, in SSIM) compared to the previous SparseNeRF [56].

reducing the view requirements by leveraging Neural Radiance Field (NeRF) [36], a powerful implicit 3D representation that captures scene details, combined with volume rendering techniques [14]. Depth regularization [12, 39, 53, 63] within the density field, additional supervision from 2D pre-trained models [26, 56], largescale pre-training [8, 67], and frequency annealings [64] have been proposed and adopted to address the challenge of few-shot view synthesis. While these NeRFbased approaches are promising, they often lead to substantial computational demands, which can affect real-time performance adversely. Subsequent research has managed to reduce training time in real-world scenarios from days to mere hours [47, 50, 58, 66], and even minutes in some cases [37]. However, a noticeable gap persists between attaining real-time rendering speeds and the desired photorealistic, high-resolution output quality.

In our research, we delve into the advancements in efficient 3D Gaussian Splatting (3D-GS) [29] and examine the challenges associated with deploying 3D-GS from sparse viewpoints. A crucial aspect for the effectiveness of 3D-GS is the densification process, which transforms the sparse initial point cloud into a more detailed representation of the 3D environment. However, the placement of new Gaussians, dictated by the spatial gradient, tends to be noisy and unrepresentative, especially in sparse-view scenarios. Additionally, the reliance on photometric loss with limited view counts often results in overly smooth textures when adhering to the conventional densification approach.

FSGS

3

To address these issues, we introduce Proximity-guided Gaussian Unpooling, a novel strategy designed specifically for sparse inputs. This method enhances the Gaussian representation by inserting new Gaussians between existing ones, based on the proximity to their neighbors. This strategic placement, combined with the initialization using observations from existing Gaussians, significantly improves scene representation by increasing the Gaussian density. Furthermore, we advocate for view augmentation through the generation of virtual camera. not present during training, to apply additional constraints in sparse setups. Incorporating monocular depth priors helps regularize the new Gaussians, steering them towards a plausible solution while enhancing texture detail, essential for accurate relative positioning in both actual training and synthetic camera views. We have conducted thorough evaluations of our Few-Shot Gaussian Splatting (FSGS) framework across a variety of few-shot Novel View Synthesis benchmarks. These include the object-centric NeRF-Synthetic datasets, the forward-facing LLFF datasets, the Shiny datasets with intricate lighting conditions, and the unbounded Mip-NeRF360 datasets. Our experiments demonstrate that FSGS sets a new benchmark in rendering quality and operates at a realtime speed (203 FPS), making it suitable for real-world applications. The efficacy of our method allows FSGS to outperform 3D-GS even with fewer Gaussians, enhancing both the efficiency and quality on the rendered scenes.

Our key contributions are as follows:

- We propose a novel point-based framework, FSGS, for few-shot view synthesis that densifies new Gaussians via Proximity-guided Gaussian Unpooling. This method effectively increases the density of Gaussians, ensuring detailed and comprehensive scene representation.
- FSGS addresses the overfitting challenge inherent in sparse-view Gaussian splatting. It achieves this by generating unseen viewpoints during training and incorporating distance correspondences on both training and synthesized pseudo views. This strategy directs the Gaussian optimization process toward solutions that are both highly accurate and visually compelling.
- FSGS significantly enhances the visual quality, and also facilitates real-time rendering speeds (over 200 FPS) leading to a viable option for practical implementation in various real-world applications.

2 Related Works

2.1 Neural Representations for 3D Reconstruction

The recent advancement of neural rendering techniques, such as Neural Radiance Fields (NeRFs) [36], has shown encouraging progress for novel view synthesis. NeRF learns an implicit neural scene representation that utilizes a MLP to map 3D coordinates (x, y, z) and view dependency (θ, ϕ) to color and density through a volume rendering function. Tremendous works focus on improving its efficiency [7,17,19,29,37,45,50], quality [1,2,4,10,22,49,54,59], generalizing to unseen scenes [8,11,24,27,51,60,67], feature distillation [16,30,31,41,69], applying artistic effects [15,25,55,68] and 3D generation [5,6,20,23,28,33,34,40,48,52]. In particular, Reiser *et al.* [45] accelerate NeRF's training by splitting a big MLP into thousands of tiny MLPs. MVSNeRF [8] constructs a 3D cost volume [21,65] and renders high-quality images from novel viewpoints. Moreover, Mip-NeRF [1] adopts conical frustum rather than a single ray in order to mitigate aliasing. Mip-NeRF 360 [3] further extends it to the unbounded scenes. While these NeRFlike models present strong performance on various benchmarks, they generally require several hours of training time. Muller *et al.* [37] adopt a multiresolution hash encoding technique that reduces the training time significantly. Kerbl *et al.* [29] propose to use a 3D Gaussian Splatting pipeline that achieves real-time rendering for either objects or unbounded scenes. The proposed FSGS approach is based on the 3D Gaussian Splatting framework but largely reduces the required training views.

2.2 Novel-View Synthesis Using Sparse Views

The original neural radiance field takes more than one hundred images as input, largely prohibiting its practical usage. To tackle this issue, several works have attempted to reduce the number of training views. Specifically, Depth-NeRF [13] applies additional depth supervision to improve the rendering quality. RegNeRF [38] proposes a depth smoothness loss as geometry regularization to stabilize training. DietNeRF [26] adds supervision on the CLIP embedding space [43], to constraint the rendered unseen views. PixelNeRF [67] trains a convolution encoder to capture context information and learns to predict 3D representation from sparse inputs. More recently, FreeNeRF [64] proposes a dynamic frequency controlling module for few-shot NeRF. SparseNeRF [56] proposes a new spatial continuity loss to distill spatial coherence from monocular depth estimators. Other works either handle single perspective images or utilize panoramic formats for 3D reconstruction, as explored in previous literature [32, 57, 63, 70]. Concurrent work ReconFusion [62] employs diffusion models to synthesize novel views, which may not always adhere to view consistency. ReconFusion jointly train a Zip-NeRF with synthesized views under a sparse-view setting. In contrast, our method improves the optimization process of Gaussian Splatting, and facilitates both the rendering speed and rendering quality.

3 Method

4

Overview. This section provides an overview of the FSGS framework, as illustrated in Fig. 2. FSGS processes a limited set of images captured from a static scene. The camera poses and point clouds are derived using the Multi-View Stereo software, COLMAP [46]. The initialization of 3D Gaussians is based on a sparse point cloud, incorporating attributes such as color, position, and a predefined conversion rule for shape and opacity. The issue of sparse initial points is tackled through the implementation of Proximity-guided Gaussian Unpooling. Our method densifies Gaussians and populates the empty spaces by assessing



Fig. 2: FSGS Pipeline. 3D Gaussians are initialized from COLMAP, with a few images (**black cameras**). For the sparsely placed Gaussians, we propose densifying new Gaussians to enhance scene coverage by unpooling existing Gaussians into new ones, with properly initialized Gaussian attributes. Monocular depth priors, enhanced by sampling unobserved views (**red cameras**), guide the optimization of grown Gaussians towards a reasonable geometry. The final loss consists of a photometric loss term, and a geometric regularization term calculated as depth relative correspondence.

the proximity between existing Gaussians and positioning new ones in the most representative areas, thereby enhancing scene details. To mitigate overfitting in standard 3D-GS with sparse-view data, we introduce the generation of pseudo camera viewpoints around the training cameras. This approach, coupled with the geometry regularization, steers the model towards accurately reconstructing the scene's geometry.

3.1 Preliminary and Problem Formulation

3D Gaussian Splatting (3D-GS), as delineated in Kerbl et al. [29], represents an 3D scene explicitly through a collection of 3D Gaussians, with attributes: a position vector $\boldsymbol{\mu} \in \mathbb{R}^3$ and a covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{3\times 3}$. Each Gaussian influences a point \boldsymbol{x} in 3D space following the 3D Gaussian distribution:

$$G(\boldsymbol{x}) = \frac{1}{(2\pi)^{3/2} |\Sigma|^{1/2}} e^{-\frac{1}{2} (\boldsymbol{x} - \boldsymbol{\mu})^T \Sigma^{-1} (\boldsymbol{x} - \boldsymbol{\mu})}$$
(1)

To ensure that Σ is positive semi-definite and holds practical physical significance, Σ is decomposed into two learnable components by $\Sigma = RSS^TR^T$, where R is a quaternion matrix representing rotation and S is a scaling matrix.

Each Gaussians also store an opacity logit $o \in \mathbb{R}$ and the appearance feature represented by n spherical harmonic (SH) coefficients $\{c_i \in \mathbb{R}^3 | i = 1, 2, ..., n\}$ where $n = D^2$ is the number of coefficients of SH with degree D. To render the 2D image, 3D-GS orders all the Gaussians that contributes to a pixel and blends the ordered Gaussians overlapping the pixels using the following function:

$$c = \sum_{i=1}^{n} c_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$$
(2)

6



Fig. 3: Points Sparsity vs. Synthesized Quality. The initial points from COLMAP using 3-views (Bottom Left) is significantly sparse than full-view(Top Left). 3D-GS with sparse initial points will decrease its quality when the training view number decreases.

where c_i is the color computed from the SH coefficients of the i^{th} Gaussian. α_i is given by evaluating a 2D Gaussian with covariance $\Sigma' \in \mathbb{R}^{2 \times 2}$ multiplied by the opacity. The 2D covariance matrix Σ' is calculated by $\Sigma' = JW\Sigma W^T J^T$, projecting the 3D covariance Σ to the camera coordinates. Here, J denotes the Jacobian of the affine approximation of the projective transformation, Wis the view transformation matrix. A heuristic Gaussian densification scheme is introduced in 3D-GS [29], where Gaussians are densified based on an average magnitude of view-space position gradients which exceed a threshold. Although this method is effective when initialized with comprehensive point cloud, it is insufficient for fully covering the entire scene with an sparse point cloud, from sparse-view input images. Additionally, some Gaussians tend to grow towards extremely large volumes, leading to results that overfit the training views and generalize badly to novel viewpoints (See Fig. 3).

The performance of 3D-GS strongly relies on both the quantity and accuracy of the initialized points. Although the subsequent Gaussian densification [29] can increase the number of Gaussians in both under-reconstructed and overreconstructed regions, this straightforward strategy falls short in few-shot settings: it suffers from inadequate initialization, leading to oversmoothed outcomes and a tendency to overfit on training views.

3.2 Proximity-guided Gaussian Unpooling

The granularity of the modeled scene depends heavily on the quality of the 3D Gaussians representing the scene; therefore, addressing the limited 3D scene coverage is crucial for effective sparse-view modeling.

Proximity Score and Graph Construction. During Gaussian optimization, we construct a directed graph, referred to as the proximity graph, to connect



Fig. 4: Gaussian Unpooling Illustration. We show a 2D toy case for Gaussian Unpooling with depth guidance, where the 1D depth provides priors on the relative distance of the Gaussians from the viewing direction, guide the Gaussian deformation toward a better solution.

each existing Gaussian with its nearest K neighbors by computing the proximity (a.k.a. Euclidean distance). Specifically, we denote the originating Gaussian at the head as the "source" Gaussian, while the one at the tail as the "destination" Gaussian, which is one of the source's K neighbors. These "destination" Gaussians are determined via the rule:

$$D_i^K = K \operatorname{-min}(d_{ij}), \quad \forall j \neq i \tag{3}$$

Here, d_{ij} is calculated via $d_{ij} = \|\mu_i - \mu_j\|$, representing the Euclidean distance among the centers of Gaussian G_i and Gaussian G_j . The assigned proximity score P_i to Gaussian G_i is calculated as the average distance to its K nearest neighbors. The proximity graph is continuously updated after each densification or pruning iteration.

$$P_i = \frac{1}{K} \sum_{j=1}^K D_i^K \tag{4}$$

Gaussian Unpooling. Inspired by the vertex-adding strategy of the mesh subdivision algorithm [71] which is widely used in computer graphics, we propose unpooling Gaussians based on the proximity graph and the proximity score of each Gaussian. Specifically, if the proximity score of a Gaussian exceeds the threshold t_{prox} , our method will grow a new Gaussian at the center of each edge, connecting the "source" and "destination" Gaussians, as shown in Fig. 4. The attributes of scale and opacity are set the same with those of the "destination" Gaussians. Meanwhile, other attributes such as rotation and SH coefficients are initialized to zero. The newly added Gaussians may distant from the "source" and "destination" Gaussians, due to the scarcity of initial points. We empirically found that duplicating the rotation and SH coefficients led to suboptimal scene parameterization. We initialize them to zero in practice. The Gaussian unpooling strategy encourages the newly densified Gaussians to be distributed around the representative locations and progressively fill observation gaps during optimization.

3.3 Geometry Guidance for Gaussian Optimization

Having achieved dense coverage by unpooling Gaussians, a photometric loss with sparse-view clues is applied for optimizing Gaussians. However, the insufficient parallax in the sparse-view setting limit the 3D Gaussians to be optimized toward a globally consistent direction where it tend to overfit on training views, and poor generalization to novel views. To inject more regularization to the optimization, we propose to create some virtual cameras that unseen in training, and apply the pixel wise geometric correspondences as additional regularization.

Synthesize Pseudo Views. To address the inherent issue of overfitting to sparse training views, we employ unobserved (pseudo) view augmentation to incorporate more prior knowledge within the scene derived from a 2D prior model. The synthesized view is sampled from the two closest training views in Euclidean space, calculating the averaged camera orientation and interpolating a virtual one between them. A random noise is applied to the 3 degrees-of-freedom (3DoF) camera location as shown in Eq. 5, and then images are rendered.

$$\mathbf{P}' = (\mathbf{t} + \varepsilon, \mathbf{q}), \quad \varepsilon \sim \mathcal{N}(0, \delta) \tag{5}$$

Here, $t \in P$ denotes camera location, while q is a quaternion representing the rotation averaged from the two cameras. This approach of synthesizing online pseudo-views enables dynamic geometry updates, as the 3D Gaussians will update progressively, reducing the risk of overfitting.

Inject Geometry Coherence from Monocular Depth. We generate the monocular D_{est} depth maps at both training and pseudo views by using the pre-trained Dense Prediction Transformer (DPT) [44], trained with 1.4 million image-depth pairs as a handy yet effective choice. To mitigate the scale ambiguity between the true scene scale and the estimated depth, we introduce a relaxed relative loss, *Pearson* correlation, on the estimated and rendered depth maps. It measures the distribution difference between 2D depth maps and follows the below function:

$$\operatorname{Corr}(\hat{\boldsymbol{D}}_{\operatorname{ras}}, \hat{\boldsymbol{D}}_{\operatorname{est}}) = \frac{\operatorname{Cov}(\hat{\boldsymbol{D}}_{\operatorname{ras}}, \hat{\boldsymbol{D}}_{\operatorname{est}})}{\sqrt{\operatorname{Var}(\hat{\boldsymbol{D}}_{\operatorname{ras}})\operatorname{Var}(\hat{\boldsymbol{D}}_{\operatorname{est}})}}$$
(6)

This soften constraint allows for the alignment of depth structure without being hindered by the inconsistencies in absolute depth values.

8

Differentiable Depth Rasterization. To enable the backpropogation from depth prior to guide Gaussian training, we implement a differentiable depth rasterizor, allowing for receiving the error signal between the rendered depth D_{ras} and the estimated depth D_{est} . Specifically, we utilize the alpha-blending rendering in 3D-GS for depth rasterization, where the z-buffer from the ordered Gaussians contributing to a pixel is accumulated for producing the depth value:

$$d = \sum_{i=1}^{n} d_i \alpha_i \prod_{j=1}^{i-1} (1 - \alpha_j)$$
(7)

Here d_i represents the z-buffer of the i^{th} Gaussians and α is identical to that in Eq. 2. This differentiable implementation enables the depth correlation loss.

3.4 Optimization

Combining all together, we can summarize the training loss:

$$\mathcal{L}(\boldsymbol{G},\boldsymbol{C}) = \lambda_1 \underbrace{\left\|\boldsymbol{C} - \hat{\boldsymbol{C}}\right\|_1}_{\mathcal{L}_1} + \lambda_2 \underbrace{\text{D-SSIM}(\boldsymbol{C}, \hat{\boldsymbol{C}})}_{\mathcal{L}_{\text{ssim}}} + \lambda_3 \underbrace{\left\|\text{Corr}(\boldsymbol{D}_{\text{ras}}, \boldsymbol{D}_{\text{est}})\right\|_1}_{\mathcal{L}_{\text{regularization}}}$$
(8)

where \mathcal{L}_1 , and \mathcal{L}_{ssim} stands for the photometric loss term between predicted image \hat{C} and ground-truth image C. $\mathcal{L}_{regularization}$ represents the geometric regularization term on both the training views and synthesized pseudo views. We set λ_1 , λ_2 , λ_3 as 0.8, 0.2, 0.05 respectively by grid search. The pseudo views sampling is enabled after 2,000 iterations to ensure the Gaussians can roughly represent the scene.

4 Experiments

4.1 Experimental Settings

LLFF Datasets [35] consist of eight forward-facing real-world scenes. The main testbed of our work follows previous settings on few-shot NeRF [38, 56, 64]. We select every eighth image as the test set, and evenly sample three views from the remaining images for training. We report experimental results on "1/8 Resolution" (504×378) and "1/4 Resolution" (1008×756).

Mip-NeRF360 Datasets [3] consist of nine scenes, each featuring a complex central object or area against a detailed background. We utilize 24 training views for comparison, with images downsampled to $4 \times$ and $8 \times$. Test images are selected the same with LLFF Datasets. We aim to establish this challenge benchmark for testing few-shot view synthesis in complex outdoor scenarios.

NeRF-Synthetic Datasets (Blender) [36] have eight objects with realistic images synthesized by Blender. We align with DietNeRF [26], where we use 8 images for training and 25 for testing, at resolution of 400×400 .

Shiny Datasets [61] contain more challenging view-dependent effects, like the rainbow reflections on a CD and refraction through a liquid bottle. We evenly select 3 views from the Shiny datasets for training at resolutions of 504×378 .

Table 1: Quantitative Comparison in LLFF Datasets, with 3 Training Views. FSGS achieves the best performance in terms of rendering accuracy and inference speed across all resolutions. Significantly, FSGS runs $2,180 \times$ faster than the previous best, SparseNeRF, while improving the SSIM from 0.624 to 0.652, at the resolution of 504×378 . We color each cell as **best**, **second best**, and **third best**.

Methods	1/8 Resolution				1/4 Resolution			
	$\text{FPS}\uparrow$	$PSNR\uparrow$	$SSIM\uparrow$	LPIPS↓	$\text{FPS}\uparrow$	$PSNR\uparrow$	$SSIM\uparrow$	$LPIPS\downarrow$
Mip-NeRF	0.21	16.11	0.401	0.460	0.14	15.22	0.351	0.540
3D-GS	385	17.43	0.522	0.321	312	16.94	0.488	0.402
DietNeRF	0.14	14.94	0.370	0.496	0.08	13.86	0.305	0.578
$\operatorname{RegNeRF}$	0.21	19.08	0.587	0.336	0.14	18.06	0.535	0.411
FreeNeRF	0.21	19.63	0.612	0.308	0.14	18.73	0.562	0.384
SparseNeRF	0.21	19.86	0.624	0.328	0.14	19.07	0.564	0.401
Ours	458	20.31	0.652	0.288	351	19.88	0.612	0.340



Fig. 5: Qualitative Results on LLFF Datasets. We demonstrate novel view results produced by 3D-GS [29], Mip-NeRF360 [3], SparseNeRF [56] and our approach for comparison. We can observe that NeRF-based methods generate floaters (Scene: *Flower*) and show aliasing results (Scene: *Leaves*) due to limited observation. 3D-GS produces oversmoothed results, caused by overfitting on training views. Our method produces pleasing appearances while demonstrating detailed thin structures.

Baselines. We compare FSGS with several few-shot NVS methods on these three dataset, including DietNeRF [26], RegNeRF [39], FreeNeRF [64], and SparseNeRF [56]. Additionally, we include comparisons with the high-performing Mip-NeRF [2], primarily designed for dense-view training, and point-based 3D-GS, following its original dense-view training recipe. Following [29, 39, 56], we report the average PSNR, SSIM, LPIPS scores and FPS for all the methods.

Table 2: Quantitative Comparison in Mip-NeRF360 Datasets, with 24 Training Views. Our FSGS shows obvious advantages over NeRF-based methods, with an improvement of more than 0.05 in SSIM and running $4,142 \times$ faster. Additionally, our method not only performs better than 3D-GS in rendering metrics but also shows improvement in FPS (from 223 to 290), thanks to the Gaussian unpooling which motivates Gaussians to expand to unseen regions more accurately.

Methods	1/8 Resolution				1/4 Resolution			
	$FPS\uparrow$	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}{\downarrow}$	$\text{FPS}\uparrow$	$PSNR\uparrow$	$SSIM\uparrow$	$\mathrm{LPIPS}{\downarrow}$
Mip-NeRF360	0.12	21.23	0.613	0.351	0.07	19.78	0.530	0.431
3D-GS	223	20.89	0.633	0.317	145	19.93	0.588	0.401
DietNeRF	0.05	20.21	0.557	0.387	0.03	19.11	0.482	0.452
$\operatorname{RegNeRF}$	0.07	22.19	0.643	0.335	0.04	20.55	0.546	0.398
FreeNeRF	0.07	22.78	0.689	0.323	0.04	21.04	0.587	0.377
SparseNeRF	0.07	22.85	0.693	0.315	0.04	21.13	0.600	0.389
Ours	290	23.70	0.745	0.220	203	22.82	0.693	0.293

Implementation Details. We implemented FSGS using the PyTorch framework, with initial point cloud from patch-match stereo in COLMAP, using only the training views. During optimization, we densify the Gaussians every 100 iterations and start densification after 500 iterations. The total optimization steps are set to 10,000, requiring approximately 9.5 minutes on LLFF datasets, and ~24 minutes on Mip-NeRF360 datasets. We set proximity threshold $t_{\rm prox}$ to 10, and the pseudo views are sampled after 2,000 iterations. We utilize the pre-trained DPT model [44] for depth estimation. All results are obtained using a NVIDIA A6000 GPU.

4.2 Comparisons to other Few-shot Methods

Comparisons on LLFF Datasets. As shown in Tab. 1, our method FSGS, despite trained from messy initial points, provides the best quantitative results and effectively addresses the insufficient scene coverage in the initialization. Our method surpasses SparseNeRF by 0.45dB and 0.81dB in PSNR at both test resolutions, while inferencing 2180 times faster, which makes FSGS a viable choice for practical usages. FSGS also outperforms 3D-GS by 2.88dB in PSNR and boost the FPS from 385 to 458, demonstrating that our refined Gaussians are more compact for scene representation from sparse views.

The qualitative analysis, as presented in Fig. 5, demonstrates that Mip-NeRF and 3D-GS struggle with the extreme sparse view problem; Mip-NeRF [1] leads to degraded geometric modeling, and 3D-GS produces blurred results in areas with complex geometry. The geometry regularization in RegNeRF [38], SparseN-eRF [56] and frequency annealing in FreeNeRF [64] do improve the quality to some extent, but still exhibit insufficient visual quality. In contrast, our proposed *Proximity-Guided Gaussian Unpooling*, and the relative geometric regularizations on both training and synthesized virtual views, pulls more Gaussians to the unobserved regions, and thus recovers more textural and structural details.



Fig. 6: Qualitative Results on Mip-NeRF360 Datasets. Comparisons were conducted with 3D-GS [29], Mip-NeRF360 [3], and SparseNeRF [56]. Our method continues to produce visually pleasing results with sharper details than other methods in large-scale scenes.

It is worth to note that, the rendering speed of the 3D Gaussian representation is contingent upon the number of optimized Gaussians, where a reduced Gaussian count leads to faster rendering speed. FSGS outperforms 3D-GS while using less Gaussian counts, by effectively places and optimizes the Gaussians to the most representative positions than 3D-GS. With the same initialization method as 3D-GS, the average optimized Gaussian count from FSGS is **57,513** on LLFF dataset, considerably lower than **63,219** from 3D-GS, which results in a faster rendering speed of FSGS than 3D-GS.

Comparisons on Mip-NeRF360 Datasets. As shown in Tab. 2, methods requiring dense view coverage (Mip-NeRF360, 3D-GS) are outperformed by ours in terms of rendering speed and metrics, across the two resolutions. Methods employing regularizations from their respective geometry and appearance fields (DietNeRF, RegNeRF, FreeNeRF, SparseNeRF) still fall short in rendering quality, while remains far from achieving real-time speed. Our FSGS significantly outperforms NeRF-based approaches, boosting PNSR by 0.85dB and improving FPS from 0.07 to 290 at 1/8 resolution. We provide a qualitative comparison in Fig. 6, where we observe that Mip-NeRF360 and SparseNeRF fail to capture the intricate details of scenes and tend to overfit on sparse training views, most notably in areas far away from cameras. In comparison, FSGS recovers the finegrained details such as the leaves on the ground (Scene: *Stump*) and the piano keys (Scene: *Bonsai*), aligning well with the ground truth. *Comparisons on Blender and Shiny Datasets.* We report both the quantitative and qualitative results on the supplementary materials.



Baseline (3D-GS)

+ Gaussian Unpool.

+ Pseudo-views + Depth Reg.

Fig. 7: Ablation Study by Visualization. 3D-GS [29] (1st column) shows that the baseline method is significantly degraded when the view coverage is insufficient. *Gaussian Unpooling* provides extra capacity to 3D Gaussians to model the scene, but the learned geometry may not be accurate (2nd column). Adding *Relative Depth Correspondence* regularization (3rd column) can further improve the modeled details.

4.3 Ablation Studies

We ablate our design choices on the LLFF dataset under the 3-view setting.

Effectiveness of Promity-guided Gaussian Unpooling. As shown in the second row of Tab. 3, our Promity-guided Gaussian Unpooling expands the scene geometry caused by limited training views, resulting in a PSNR improvement of 1.21dB compared to 3D-GS. We also visualize its visual effects in Fig. 7. The heuristic Gaussian densification leads to blurring results, particularly noticeable in areas like bush and grass, our approach enriches structural and visual details.

Impact of Relative Depth Regularization. The third row of Tab. 3 demonstrates the improvement by introducing depth priors, guiding the Gaussian unpooling towards more plausible geometry. In Fig. 7, we observe that the depth regularization effectively eliminates the artifacts in grassy regions, and enforces more consistent and solid surfaces with geometric coherence. We also display the rendered depth map, where depth regularization leads to depths aligning better with the actual geometric structures.

14 Zehao Zhu[®], Zhiwen Fan[®], Yifan Jiang[®], Zhangyang Wang[®]

Table 3: Ablation Study on proposed components. Starting from 3D-GS [29] (1st row), we find that our proposed *Gaussian Unpooling* (2nd row) is more effective than the densification scheme in 3D-GS for few-shot view synthesis. Applying additional supervision from a monocular depth estimator further regularizes the Gaussian optimization towards a better solution (3rd row). Introducing pseudo-view augmentation to apply additional regularization when optimizing Gaussians further enhances the results in a few-shot scenario.

Gaussian Unpooling	Geometry Guidance	Pseudo Views	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	LPIPS↓
×	X	×	17.43	0.522	0.321
1	×	×	18.64	0.580	0.311
1	1	×	19.83	0.634	0.297
✓	1	1	20.31	0.652	0.288

Pseudo-view Matters in Few-shot Modeling. Tab. 3 validates the impact of synthesizing more unseen views during training, which anchors the Gaussians to a plausible geometry and further enhances the modeling quality when the geometry in densification is not accurate.

5 Conclusion and Limitation

In this work, we present a real-time few-shot framework, FSGS, for novel views synthesis within an insufficiently view overlapping. Starting from COLMAP point clouds, FSGS adopts the 3DGS representation and proposes an effective *Proximity-guided Gaussian Unpooling* by measuring the proximity of each Gaussian to its neighbors. The overfitting issue in few-view 3D-GS can be alleviated by the adoption of pseudo-view generation and monocular relative depth correspondences to guide the expanded scene geometry toward a better solution. FSGS is capable of generating photo-realistic images with as few as three images, and perform inference at more than 200FPS, offering new avenues for real-time rendering and more cost-effective capture methods.

Although FSGS notably enhances the quality and efficiency of real-time fewshot neural rendering, it cannot generalize to occluded views that are unobserved during training. We hope that our proposed approach drives new research towards few-shot novel view synthesis in arbitrary 3D scenes.

Acknowledgement

The work is supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/Interior Business Center (DOI/IBC) contract number 140D0423C0074. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DOI/IBC, or the U.S. Government.

References

- Barron, J.T., Mildenhall, B., Tancik, M., Hedman, P., Martin-Brualla, R., Srinivasan, P.P.: Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5855–5864 (2021)
- 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)
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-NeRF 360: Unbounded Anti-Aliased Neural Radiance Fields. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 5460–5469 (2022)
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Zip-nerf: Anti-aliased grid-based neural radiance fields. ICCV (2023)
- Cao, Y., Cao, Y.P., Han, K., Shan, Y., Wong, K.Y.K.: Dreamavatar: Text-andshape guided 3d human avatar generation via diffusion models. arXiv preprint arXiv:2304.00916 (2023)
- Chan, E., Monteiro, M., Kellnhofer, P., Wu, J., Wetzstein, G.: pi-gan: Periodic implicit generative adversarial networks for 3d-aware image synthesis. In: arXiv (2020)
- Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: European Conference on Computer Vision. pp. 333–350. Springer (2022)
- Chen, A., Xu, Z., Zhao, F., Zhang, X., Xiang, F., Yu, J., Su, H.: Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14124–14133 (2021)
- 9. Chen, G., Wang, W.: A survey on 3d gaussian splatting. arXiv preprint arXiv:2401.03890 (2024)
- Chen, T., Wang, P., Fan, Z., Wang, Z.: Aug-nerf: Training stronger neural radiance fields with triple-level physically-grounded augmentations. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15191– 15202 (2022)
- 11. Chibane, J., Bansal, A., Lazova, V., Pons-Moll, G.: Stereo radiance fields (srf): Learning view synthesis from sparse views of novel scenes. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR). IEEE (jun 2021)

- 16 Zehao Zhu[©], Zhiwen Fan[©], Yifan Jiang[©], Zhangyang Wang[©]
- Deng, C., Jiang, C., Qi, C.R., Yan, X., Zhou, Y., Guibas, L., Anguelov, D., et al.: Nerdi: Single-view nerf synthesis with language-guided diffusion as general image priors. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20637–20647 (2023)
- 13. Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised nerf: Fewer views and faster training for free. arXiv preprint arXiv:2107.02791 (2021)
- Drebin, R.A., Carpenter, L., Hanrahan, P.: Volume rendering. ACM Siggraph Computer Graphics 22(4), 65–74 (1988)
- Fan, Z., Jiang, Y., Wang, P., Gong, X., Xu, D., Wang, Z.: Unified implicit neural stylization. In: European Conference on Computer Vision. pp. 636–654. Springer (2022)
- Fan, Z., Wang, P., Jiang, Y., Gong, X., Xu, D., Wang, Z.: Nerf-sos: Anyview self-supervised object segmentation on complex scenes. arXiv preprint arXiv:2209.08776 (2022)
- Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5501–5510 (2022)
- Gao, K., Gao, Y., He, H., Lu, D., Xu, L., Li, J.: Nerf: Neural radiance field in 3d vision, a comprehensive review (2023)
- Garbin, S.J., Kowalski, M., Johnson, M., Shotton, J., Valentin, J.: Fastnerf: Highfidelity neural rendering at 200fps. arXiv preprint arXiv:2103.10380 (2021)
- Gu, J., Liu, L., Wang, P., Theobalt, C.: Stylenerf: A style-based 3d aware generator for high-resolution image synthesis. In: International Conference on Learning Representations (2022)
- Gu, X., Fan, Z., Zhu, S., Dai, Z., Tan, F., Tan, P.: Cascade cost volume for high-resolution multi-view stereo and stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2495– 2504 (2020)
- Guo, Y.C., Kang, D., Bao, L., He, Y., Zhang, S.H.: Nerfren: Neural radiance fields with reflections. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18409–18418 (2022)
- Höllein, L., Cao, A., Owens, A., Johnson, J., Nießner, M.: Text2room: Extracting textured 3d meshes from 2d text-to-image models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 7909–7920 (October 2023)
- Irshad, M.Z., Zakharov, S., Liu, K., Guizilini, V., Kollar, T., Gaidon, A., Kira, Z., Ambrus, R.: Neo 360: Neural fields for sparse view synthesis of outdoor scenes (2023)
- Jain, A., Mildenhall, B., Barron, J.T., Abbeel, P., Poole, B.: Zero-shot text-guided object generation with dream fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 867–876 (2022)
- Jain, A., Tancik, M., Abbeel, P.: Putting nerf on a diet: Semantically consistent few-shot view synthesis. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 5885–5894 (2021)
- 27. Johari, M.M., Lepoittevin, Y., Fleuret, F.: Geonerf: Generalizing nerf with geometry priors. Proceedings of the IEEE international conference on Computer Vision and Pattern Recognition (CVPR) (2022)
- Karnewar, A., Vedaldi, A., Novotny, D., Mitra, N.: Holodiffusion: Training a 3D diffusion model using 2D images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2023)

- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. ACM Transactions on Graphics (ToG) 42(4), 1-14 (2023)
- Kerr, J., Kim, C.M., Goldberg, K., Kanazawa, A., Tancik, M.: Lerf: Language embedded radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 19729–19739 (2023)
- Kobayashi, S., Matsumoto, E., Sitzmann, V.: Decomposing nerf for editing via feature field distillation. Advances in Neural Information Processing Systems 35, 23311–23330 (2022)
- Li, R., Pan, P., Yang, B., Xu, D., Zhou, S., Zhang, X., Li, Z., Kadambi, A., Wang, Z., Fan, Z.: 4k4dgen: Panoramic 4d generation at 4k resolution. arXiv preprint arXiv:2406.13527 (2024)
- Lin, C.H., Gao, J., Tang, L., Takikawa, T., Zeng, X., Huang, X., Kreis, K., Fidler, S., Liu, M.Y., Lin, T.Y.: Magic3d: High-resolution text-to-3d content creation. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2023)
- Liu, R., Wu, R., Hoorick, B.V., Tokmakov, P., Zakharov, S., Vondrick, C.: Zero-1to-3: Zero-shot one image to 3d object (2023)
- 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)
- 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). https://doi.org/10.1145/ 3503250
- Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. ACM Transactions on Graphics (ToG) 41(4), 1– 15 (2022)
- Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. arXiv preprint arXiv:2112.00724 (2021)
- Niemeyer, M., Barron, J.T., Mildenhall, B., Sajjadi, M.S., Geiger, A., Radwan, N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5480–5490 (2022)
- Poole, B., Jain, A., Barron, J.T., Mildenhall, B.: Dreamfusion: Text-to-3d using 2d diffusion. arXiv preprint arXiv:2209.14988 (2022)
- Qin, M., Li, W., Zhou, J., Wang, H., Pfister, H.: Langsplat: 3d language gaussian splatting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 20051–20060 (2024)
- 42. Rabby, A.S.A., Zhang, C.: Beyondpixels: A comprehensive review of the evolution of neural radiance fields (2023)
- Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. pp. 8748–8763. PMLR (2021)
- Ranftl, R., Bochkovskiy, A., Koltun, V.: Vision transformers for dense prediction. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12179–12188 (2021)

- 18 Zehao Zhu[©], Zhiwen Fan[©], Yifan Jiang[©], Zhangyang Wang[©]
- Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14335–14345 (2021)
- 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)
- 47. Schwarz, K., Sauer, A., Niemeyer, M., Liao, Y., Geiger, A.: Voxgraf: Fast 3d-Aware Image Synthesis With Sparse Voxel Grids. ArXiv Preprint ArXiv:2206.07695 (2022)
- Seo, J., Jang, W., Kwak, M.S., Ko, J., Kim, H., Kim, J., Kim, J.H., Lee, J., Kim, S.: Let 2d diffusion model know 3d-consistency for robust text-to-3d generation. arXiv preprint arXiv:2303.07937 (2023)
- Suhail, M., Esteves, C., Sigal, L., Makadia, A.: Light field neural rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8269–8279 (2022)
- Sun, C., Sun, M., Chen, H.T.: Direct Voxel Grid Optimization: Super-Fast Convergence for Radiance Fields Reconstruction. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5459–5469 (2022)
- 51. T, M.V., Wang, P., Chen, X., Chen, T., Venugopalan, S., Wang, Z.: Is attention all that neRF needs? In: The Eleventh International Conference on Learning Representations (2023), https://openreview.net/forum?id=xE-LtsE-xx
- 52. Tang, J., Wang, T., Zhang, B., Zhang, T., Yi, R., Ma, L., Chen, D.: Make-it-3d: High-fidelity 3d creation from a single image with diffusion prior (2023)
- 53. Truong, P., Rakotosaona, M.J., Manhardt, F., Tombari, F.: Sparf: Neural radiance fields from sparse and noisy poses. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4190–4200 (2023)
- Verbin, D., Hedman, P., Mildenhall, B., Zickler, T., Barron, J.T., Srinivasan, P.P.: Ref-nerf: Structured view-dependent appearance for neural radiance fields. arXiv preprint arXiv:2112.03907 (2021)
- 55. Wang, C., Chai, M., He, M., Chen, D., Liao, J.: Clip-nerf: Text-and-image driven manipulation of neural radiance fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3835–3844 (2022)
- Wang, G., Chen, Z., Loy, C.C., Liu, Z.: Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. arXiv preprint arXiv:2303.16196 (2023)
- Wang, G., Wang, P., Chen, Z., Wang, W., Loy, C.C., Liu, Z.: Perf: Panoramic neural radiance field from a single panorama. arXiv preprint arXiv:2310.16831 (2023)
- Wang, L., Zhang, J., Liu, X., Zhao, F., Zhang, Y., Zhang, Y., Wu, M., Yu, J., Xu, L.: Fourier PlenOctrees for Dynamic Radiance Field Rendering in Real-Time. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 13524–13534 (2022)
- Wang, P., Liu, Y., Chen, Z., Liu, L., Liu, Z., Komura, T., Theobalt, C., Wang, W.: F2-nerf: Fast neural radiance field training with free camera trajectories. CVPR (2023)
- 60. Wang, Q., Wang, Z., Genova, K., Srinivasan, P.P., Zhou, H., Barron, J.T., Martin-Brualla, R., Snavely, N., Funkhouser, T.: Ibrnet: Learning multi-view image-based rendering. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4690–4699 (2021)
- Wizadwongsa, S., Phongthawee, P., Yenphraphai, J., Suwajanakorn, S.: Nex: Realtime view synthesis with neural basis expansion. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021)

- Wu, R., Mildenhall, B., Henzler, P., Park, K., Gao, R., Watson, D., Srinivasan, P.P., Verbin, D., Barron, J.T., Poole, B., et al.: Reconfusion: 3d reconstruction with diffusion priors. arXiv preprint arXiv:2312.02981 (2023)
- Xu, D., Jiang, Y., Wang, P., Fan, Z., Shi, H., Wang, Z.: Sinnerf: Training neural radiance fields on complex scenes from a single image. In: European Conference on Computer Vision. pp. 736–753. Springer (2022)
- 64. Yang, J., Pavone, M., Wang, Y.: Freenerf: Improving few-shot neural rendering with free frequency regularization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8254–8263 (2023)
- Yao, Y., Luo, Z., Li, S., Fang, T., Quan, L.: Mvsnet: Depth inference for unstructured multi-view stereo. In: Proceedings of the European conference on computer vision (ECCV). pp. 767–783 (2018)
- Yu, A., Fridovich-Keil, S., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance Fields Without Neural Networks. 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) pp. 5491–5500 (2022)
- 67. Yu, A., Ye, V., Tancik, M., Kanazawa, A.: pixelnerf: Neural radiance fields from one or few images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4578–4587 (2021)
- Zhang, K., Kolkin, N., Bi, S., Luan, F., Xu, Z., Shechtman, E., Snavely, N.: Arf: Artistic radiance fields (2022)
- Zhou, S., Chang, H., Jiang, S., Fan, Z., Zhu, Z., Xu, D., Chari, P., You, S., Wang, Z., Kadambi, A.: Feature 3dgs: Supercharging 3d gaussian splatting to enable distilled feature fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21676–21685 (2024)
- Zhou, S., Fan, Z., Xu, D., Chang, H., Chari, P., Bharadwaj, T., You, S., Wang, Z., Kadambi, A.: Dreamscene360: Unconstrained text-to-3d scene generation with panoramic gaussian splatting. arXiv preprint arXiv:2404.06903 (2024)
- Zorin, D., Schröder, P., Sweldens, W.: Interpolating subdivision for meshes with arbitrary topology. In: Proceedings of the 23rd annual conference on Computer graphics and interactive techniques. pp. 189–192 (1996)