

---

# Structure Consistent Gaussian Splatting with Matching Prior for Few-shot Novel View Synthesis

---

Rui Peng<sup>1,2</sup> Wangze Xu<sup>1</sup> Luyang Tang<sup>1,2</sup> Liwei Liao<sup>1</sup> Jianbo Jiao<sup>3</sup> Ronggang Wang<sup>✉,1,2</sup>

<sup>1</sup>Guangdong Provincial Key Laboratory of Ultra High Definition Immersive Media Technology, Peking University Shenzhen Graduate School

<sup>2</sup>Peng Cheng Laboratory <sup>3</sup>University of Birmingham  
ruipeng@stu.pku.edu.cn rgwang@pkusz.edu.cn

## Abstract

Despite the substantial progress of novel view synthesis, existing methods, either based on the Neural Radiance Fields (NeRF) or more recently 3D Gaussian Splatting (3DGS), suffer significant degradation when the input becomes sparse. Numerous efforts have been introduced to alleviate this problem, but they still struggle to synthesize satisfactory results efficiently, especially in the large scene. In this paper, we propose *SCGaussian*, a Structure Consistent Gaussian Splatting method using matching priors to learn 3D consistent scene structure. Considering the high interdependence of Gaussian attributes, we optimize the scene structure in two folds: rendering geometry and, more importantly, the position of Gaussian primitives, which is hard to be directly constrained in the vanilla 3DGS due to the non-structure property. To achieve this, we present a hybrid Gaussian representation. Besides the ordinary non-structure Gaussian primitives, our model also consists of ray-based Gaussian primitives that are bound to matching rays and whose optimization of their positions is restricted along the ray. Thus, we can utilize the matching correspondence to directly enforce the position of these Gaussian primitives to converge to the surface points where rays intersect. Extensive experiments on forward-facing, surrounding, and complex large scenes show the effectiveness of our approach with state-of-the-art performance and high efficiency. Code is available at <https://github.com/prstrive/SCGaussian>.

## 1 Introduction

Few-shot novel view synthesis (NVS) aims to reconstruct the scene given only a sparse collection of views, which has always been a cornerstone and challenging task in computer vision. Neural radiance field (NeRF) [29], emerged as an excellent 3D representation, has shown great success in rendering photo-realistic novel views. However, such impressive results require an expensive and time-consuming collection of dense images which impedes many practical applications, *e.g.*, the input is typically much sparser in autonomous driving, robotics, and virtual reality. Although many attempts have been proposed to solve this challenging few-shot rendering problem from the aspect of pre-training [6, 52, 60, 17], regularization terms [20, 32, 58, 43], external priors [8, 39, 51, 46, 15], *etc.*, these NeRF-based methods still suffer from low rendering speed and high computational cost, *i.e.*, each scene requires hours or even days of training time.

Recently, an efficient representation 3D Gaussian Splatting (3DGS) [19] is proposed to leverage a set of Gaussian primitives (initialized from the Structure-from-motion (SfM) [41, 42] points) along with some attributes to explicitly model the 3D scene. Through replacing the cumbersome volume rendering in NeRF methods with the efficient differentiable splatting, which directly projects the Gaussian primitives onto the 2D image plane, 3DGS has expressed remarkable improvement in both

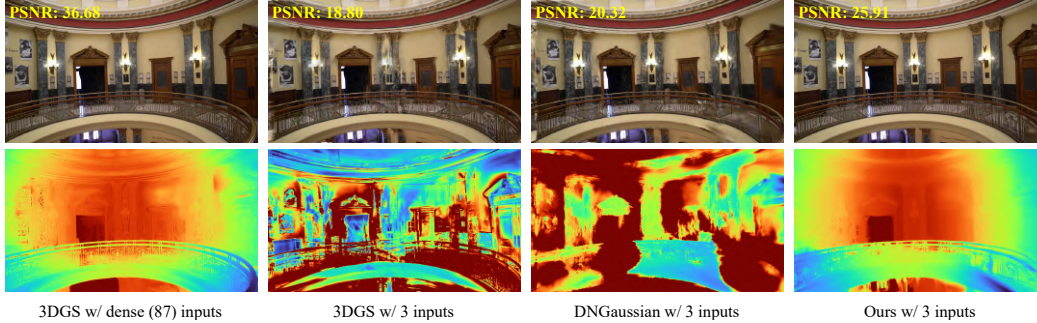


Figure 1: **Comparisons in view synthesis and geometry rendering.** 3DGS [19] can synthesize high-quality novel views and plausible geometry with excessive inputs, but suffers from significant degradation in the sparse scenario. Even using the monocular depth prior, DNGaussian [23] still struggles to generate accurate geometry and novel views. In contrast, our method can learn the more consistent scene structure and render the more realistic images.

rendering quality and speed, *i.e.*, high-resolution images can be rendered in real-time. Even with this unprecedented performance, 3DGS still relies on dense image captures and faces the same problem of novel view degeneration with NeRF methods, when only a few inputs are available.

In this paper, we aim to address this issue by establishing a few-shot 3DGS model with a consistent structure to pursue high-quality and efficient novel view synthesis. Compared with the dense counterpart, this few-shot system introduces more challenging problems, *e.g.*, the trivial multi-view constraints that the model can only be supervised from sparse viewpoints, and the high interdependence between Gaussian attributes make their optimization ambiguity, *i.e.*, optimizing the position *vs* optimizing the shape. Although some recent efforts [23, 65] have attempted to use monocular depth priors [36, 37] to stabilize the optimization, *e.g.*, the monocular depth consistency of sampled virtual viewpoints [65] and the hard-soft monocular depth regularization [23], as shown in Fig. 1, the inherent scale and multi-view inconsistency of monocular depth make it hard to guarantee a consistent scene structure and lead to unsatisfactory rendering results, especially in complex scenes.

To this end, we are motivated to exploit the matching prior, which exhibits worthwhile characteristics indicating the ray/pixel correspondence between views and the multi-view visible region. Based on this, we propose *SCGaussian*, a framework that leverages matching priors to explicitly enforce the optimization of scene structure to be 3D consistent. A straightforward idea for this purpose is to use the ray correspondence to supervise the reprojection error of the rendering depth. However, we observe that the rendering geometry is not always consistent with the scene structure due to the interdependence of Gaussian attributes. In this paper, we argue that in addition to the rendering geometry, the more important aspect to ensure the consistency of the scene structure is to optimize the *position of Gaussian primitives*. To achieve this, we design a hybrid representation, which consists of ray-based Gaussian primitives besides the ordinary non-structure Gaussian primitives. For these rays-based ones, we bind them to matching rays and restrict the optimization of their position along the ray, thus we can utilize the matching correspondence to optimize the position of Gaussian primitives to converge to the consistent surface position where rays intersect. In this dual optimization solution, both the position and shape of the Gaussian primitives can be constrained properly.

Extensive experiments on LLFF [30], IBRNet [52], DTU [16], Tanks and Temples [21] and Blender [29] datasets show the effectiveness of our *SCGaussian*, which is capable of synthesizing detail and accurate novel views in these forward-facing, surrounding, and complex large scenes, achieving new state-of-the-art performance in both rendering quality (**3 – 5 dB** improvement on challenging complex scenes [21]) and efficiency (**~200 FPS** rendering and 1-minute convergence speed).

## 2 Related Works

**Novel view synthesis.** Novel view synthesis is a task to render realistic images of unseen views given a set of training images. Many methods are proposed to address this problem in both traditional [7, 4, 45, 22, 55] and deep-learning based [64, 63, 11, 10] manners. In particular, NeRF [29] achieves photo-realistic rendering and has become one of the most popular methods in recent years, which

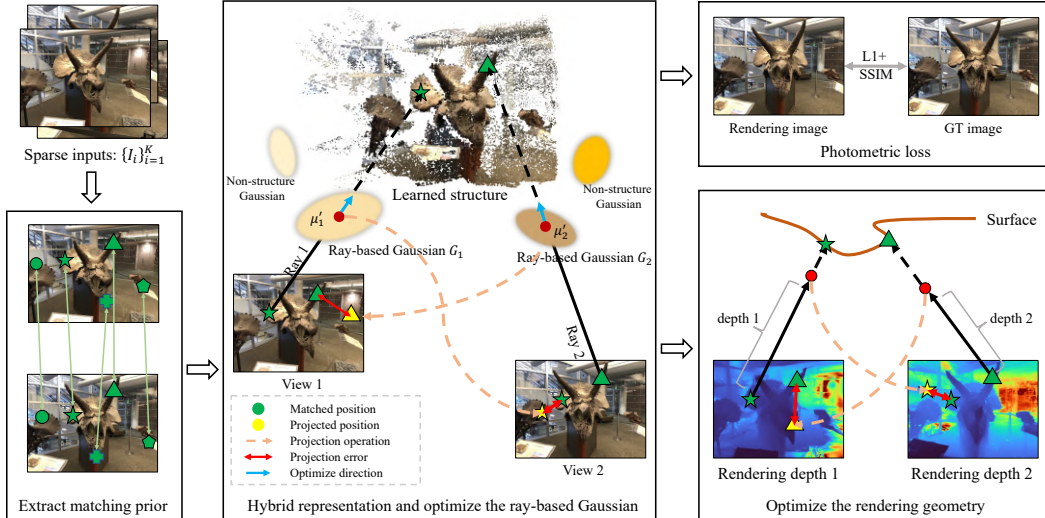


Figure 2: **Framework of SCGaussian.** We first extract the matching prior from the sparse input, and randomly initialize the hybrid Gaussian representation. The ray-based Gaussian primitives are bound to matching rays, and are explicitly optimized using the matching correspondence. The rendering geometry optimization is further conducted to optimize the shape of all types of Gaussian primitives. Combined with the ordinary photometric loss, SCGaussian can learn the consistent scene structure.

successfully combines multi-layer perceptrons (MLP) and volume rendering. The following works try to improve NeRF in many aspects, *e.g.*, quality [1, 2], pose-free [54, 28, 25, 50], dynamic view synthesis [35, 12, 24, 33], training and rendering efficiency [14, 31, 13, 47, 5, 26]. And more recently, a point-based method 3D Gaussian Splatting [19] represents the scene as 3D Gaussians and significantly improve the rendering speed to a real-time level. And it has shown an advantage in many aspects [56, 59, 61, 18, 27, 49, 3] compared with NeRF-based methods. However, these methods need dense input views, which makes them unsuitable for many practical applications.

**Few-shot novel view synthesis.** Compared with the ordinary NVS, the few-shot NVS is a more practical task but also more challenging. The original rendering methods always suffer from dramatic degradation when applied directly to the sparse scenario. Many works have attempted to solve this problem. Specifically, one thread of works [60, 52, 17, 6] attempt to pre-train a generalizable model on the large-scale datasets first and apply it to the target scene with sparse inputs. Another alternate approach is to optimize the model from scratch for each scene. [8, 39] try to add the depth supervision from the SFM points or depth completion model, and [62, 51] adopt the more practical monocular depth prior. To exploit smoothness and semantic priors, works [32, 15] choose to render some patches first and introduce the geometry and appearance regularization. These methods are all based on the NeRF and rely on volume rendering to synthesize novel views, which is always time-consuming. Some recent methods [65, 23, 57] combine the efficient 3DGS representation with monocular depth [36] and multi-view stereo [34] prior to improve the efficiency of the few-shot NVS task. However, since 3DGS relies on the initialization of sparse SFM points and adequate multi-view constraints, which are hard to observe in the sparse scenario, how to learn the globally consistent structure is the crucial bottleneck.

### 3 Methodology

In this section, we introduce the proposed new few-shot approach, SCGaussian, which can learn consistent 3D scene structure using matching priors. The overall framework of our model is illustrated in Fig. 2. In Sec. 3.1, we first review the 3DGS. Then we elaborate on the challenge of few-shot 3DGS and the motivation of using matching priors in Sec. 3.2, and the design of our Structure Consistent Gaussian Splatting will be introduced in Sec. 3.3. The full loss function and training detail will be described in Sec. 3.4.

### 3.1 Preliminary of 3D Gaussian Splatting

3DGS [19] explicitly represents the 3D scene through a collection of anisotropic 3D Gaussians. Each Gaussian is defined by a center vector  $\mu \in \mathbb{R}^3$  and a covariance matrix  $\Sigma \in \mathbb{R}^{3 \times 3}$ , and the influence for a position  $x$  is defined as:

$$G(x) = e^{-\frac{1}{2}(x-\mu)^T \Sigma^{-1}(x-\mu)}. \quad (1)$$

To ensure positive semi-definiteness and effective optimization, the covariance matrix is decomposed into a scaling matrix  $S$  and a rotation matrix  $R$  as:

$$\Sigma = R S S^T R^T, \quad (2)$$

where these two matrices can be derived by a scaling factor  $s \in \mathbb{R}^3$  and a rotation factor  $r \in \mathbb{R}^4$ . Additionally, each Gaussian also contains the appearance feature  $sh \in \mathbb{R}^k$  represented by a set of spherical harmonics (SH) coefficients and an opacity value  $\alpha \in \mathbb{R}$ .

To render the image, the 3DGS is projected to the 2D image plane via a view transformation matrix  $W$  and the Jacobian of the affine approximation of the projective transformation  $J$ :

$$\Sigma' = J W \Sigma W^T J^T. \quad (3)$$

Using the point-based rendering, the color  $C$  of a pixel can be computed by blending  $N$  ordered Gaussians overlapping the pixel:

$$C = \sum_{i \in N} c_i \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (4)$$

where  $c_i$  denotes the view-dependent color of  $i$ -th Gaussian, and  $\alpha'_i$  is determined by the multiplication of  $\Sigma'$  and the opacity  $\alpha_i$ . Similarly, we can also render the depth image  $D$  through:

$$D = \sum_{i \in N} d_i \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j), \quad (5)$$

where  $d_i$  refers to the depth of  $i$ -th Gaussian.

In summary, each Gaussian point  $\theta$  can be characterized by the set of attributes:  $\theta = \{\mu, s, r, \alpha, sh\}$ . To optimize the model, 3DGS takes the photometric loss, which is measured with the combination of L1 and SSIM [53], between the rendering image  $\hat{I}$  and the ground-truth image  $I$ :

$$L_{photo} = (1 - \lambda)L_1(\hat{I}, I) + \lambda L_{ssim}(\hat{I}, I), \quad (6)$$

where  $\lambda$  is always set to 0.2. Facilitated by the highly-optimized rasterization pipeline, 3DGS can achieve remarkably fast rendering and training speed and enables real-time view synthesizing.

### 3.2 Motivation of Matching Priors

Through mimicking the image-formation process at training views, the model aims to find a set of optimal Gaussians  $\mathcal{G} = \{\theta_i\}_{i=1}^N$  to build a photo-realistic mapping function  $f_{\mathcal{G}} : P \rightarrow \hat{I}$ , i.e., mapping the image for arbitrary camera pose  $P$ :

$$\mathcal{G} = \arg \min_{\mathcal{G}} L_{photo}(\mathcal{G}). \quad (7)$$

With adequate training views, the optimized model is capable of generating great novel view rendering results. However, as shown in Fig. 3, when the input becomes sparse, the 3DGS model always overfits training views and suffers from a significant degradation in test poses.

We observe that the challenge of the few-shot 3DGS mainly comes from the failure of learning the 3D consistent scene structure, e.g., the learned Gaussian primitives cannot distribute over the accurate surface region and the rendering geometry is multi-view inconsistent. In the sparse scenario, the supervision signal only comes from a few training poses, and this trivial multi-view constraint makes it hard to bias the model towards learning a 3D consistent solution. Conversely, as shown in Fig. 3 (a), 3DGS model tends to learn the inconsistent Gaussians for each view separately, e.g.,

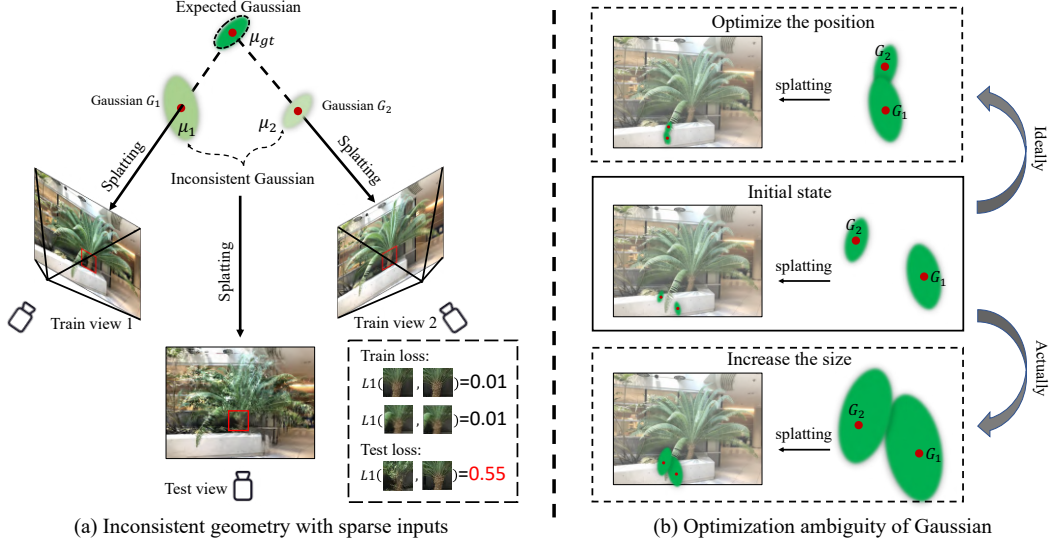


Figure 3: **Visualization of some challenges faced by few-shot 3DGS.** (a) The expected Gaussian in the surface region cannot be learned, and the model tends to learn the inconsistent Gaussian and overfit the training views. While the training loss is small enough, the testing error is pretty bad. (b) The attributes of Gaussian primitives are interdependent and the model tends to increase the size to cover the pixels rather than correct the position.

the model learns a “wall” extremely close to each camera, in which case the training loss is still small. Furthermore, we find that 3DGS has an obvious optimization ambiguity due to the high interdependence of attributes, *e.g.*, shape versus position. Theoretically, the model needs to learn more small-sized Gaussian primitives over the texture region to recover high-frequency details, but in practice, it prefers to increase the size of Gaussian primitives to cover these pixels as shown in Fig. 3 (b), resulting in an overly smooth view synthesis. To ensure that the learned structure is consistent, a heuristic strategy in the vanilla 3DGS is to use sparse SFM points as initialization and guide the model’s optimization, which is especially crucial for complex scenes. However, in the sparse scenario, it’s pretty hard to stably extract enough SFM points, and usually, only random initialization can be used like [23]. This amplifies the challenge of learning the consistent structure.

Although some methods [65, 23] attempt to use the monocular depth to regularize the geometry, the inherent scale and multi-view ambiguity of monocular depth make it difficult to solve the aforementioned problems. Thus, we are interested in the question: *how can we make 3DGS without SFM point initialization to learn 3D consistent scene structure under sparse input?* In this paper, we consider exploiting matching priors using the pre-trained matching model [44], which doesn’t face the ambiguity problem like monocular depth. Matching priors have two important characteristics: ray correspondence and ray position.

**Ray correspondence.** The pair of matching rays represent the corresponding 2D position of a consistent 3D point in different views, which can serve as the prominent multi-view constraint, *i.e.*, the matching rays should theoretically intersect at the same surface position. Given a pair of matching rays  $\{r_i, r_j\}$  at image  $I_i$  and  $I_j$ , and the corresponding pixel coordinates are  $\{p_i, p_j\}$ , supposing we have computed the position of the surface point intersect with each ray as  $X_i$  and  $X_j$ , we can get the following equation:

$$X_i = X_j. \quad (8)$$

Meanwhile, given the camera intrinsics  $\{K_i, K_j\}$  and extrinsics  $\{[R_i, t_i], [R_j, t_j]\}$ , we can further project the surface point to another 2D image plane and get the projected pixel coordinate, *e.g.*, the projection from  $i$  to  $j$  can be modeled as:

$$p_{i \rightarrow j}(X_i) = \pi(K_j R_j^T (X_i - t_j)), \quad (9)$$

where  $\pi$  is the projection operator  $\pi([x, y, z]^T) = [x/z, y/z]^T$ . Thus we have the equation in pixel coordinate:  $p_j = p_{i \rightarrow j}$ , and similarly, we have  $p_i = p_{j \rightarrow i}$ .



**Ray position.** In the matching prior, the position of matching rays exactly indicates the region that is commonly visible to at least two views. This multi-view visible region plays a crucial role in the reconstruction model, as it’s meaningless when there is no overlapping region between views. In the sparse scenario, the stereo correspondence is insignificant and the non-overlapping region can even harm the model training, while the importance of the multi-view visible region is magnified.

### 3.3 Structure Consistent Gaussian Splatting

To fully exploit the characteristics of matching prior, our SCGaussian explicitly optimizes the scene structure in two folds: the position of Gaussian primitive and the rendering geometry. Optimizing the position of Gaussian primitive is non-trivial due to the non-structural properties of Gaussian primitives. To address this, we present a hybrid Gaussian representation. Besides ordinary non-structure Gaussian primitives used to recover the background region visible in a single view, our model also consists of ray-based Gaussian primitives which are bound to matching rays, in which case their positions are restricted to be optimized along the ray.

**Initialization and densification.** Different from existing methods that initialize with either SFM points [65] or random points [23], we initialize with ray-based Gaussian primitives and bind them to matching rays. For convenience, here we discuss two input images  $I_i$  and  $I_j$ . Suppose we have  $N$  pairs of matching rays  $\{r_i^k, r_j^k\}_{k=1}^N$ , we can initialize  $N$  pairs of ray-based Gaussian primitives  $\{\mathcal{G}_i^k, \mathcal{G}_j^k\}_{k=1}^N$ . Similar to 3DGS, each primitive is equipped with a set of learnable attributes but with a different position representation. The position of the ray-based Gaussian primitive  $\mu^l$  is defined as:

$$\mu^l = o + zd, \quad (10)$$

where  $o$  and  $d$  refer to the camera center and ray direction respectively, and  $z$  is a learnable distance factor, which is randomly initialized.

For densification, we follow the same strategy in [19] to determine the “under-reconstruction” candidates using the average magnitude of view-space position gradients, and generate the non-structure Gaussian primitives, whose positions can be optimized in arbitrary directions.

**Optimize the position of Gaussian primitives.** As analyzed in Sec. 3.2, the accurate position of Gaussian primitives plays a fundamental role in the learned scene structure. Since the matching correspondence between ray-based Gaussian primitives can be constructed using the binding strategy, we can conveniently optimize their positions.

For a pair of matching rays  $\{r_i, r_j\}$  in image  $I_i$  and  $I_j$ , thanks to our binding strategy, we can get a pair of binding Gaussian primitives  $\{\mathcal{G}_i, \mathcal{G}_j\}$ , and their positions in 3D space are  $\mu_i^l = o_i + z_i d_i$  and  $\mu_j^l = o_j + z_j d_j$  respectively. According to Eq. (8) and Eq. (9), we can get the projected 2D coordinate from  $i$  to  $j$ :  $p_{i \rightarrow j}(\mu_i^l)$  and from  $j$  to  $i$ :  $p_{j \rightarrow i}(\mu_j^l)$ . Thus we can get the projection error of this pair of Gaussian primitives as:

$$\begin{cases} L_{gp}^{i \rightarrow j} = \|p_j - p_{i \rightarrow j}(\mu_i^l)\| \\ L_{gp}^{j \rightarrow i} = \|p_i - p_{j \rightarrow i}(\mu_j^l)\|. \end{cases} \quad (11)$$

The final Gaussian position loss  $L_{gp}$  is computed as the average error of all binding Gaussian pairs.

**Optimize the rendering geometry.** Due to the interdependence of Gaussian attributes, the rendering geometry is not consistent with Gaussian positions, *e.g.*, the incorrect scaling or rotation can lead to wrong rendering geometry and affect the rendering results even with the correct Gaussian position.

We first render the depth image  $D_i$  and  $D_j$  through Eq. (5) and get the estimated depth for the pair of matching rays  $\{D_i(p_i), D_j(p_j)\}$ . Then we lift the pixel coordinate to 3D space:

$$\nu_i = R_i(D_i(p_i)K_i^{-1}\tilde{p}_i) + t_i, \quad (12)$$

where  $\tilde{p}$  refers to the 2D homogeneous of  $p$ . Similarly, we can get the 3D position  $\nu_j$  of ray  $j$ . Then we get the projected 2D coordinate  $p_{i \rightarrow j}(\nu_i^l)$  and  $p_{j \rightarrow i}(\nu_j^l)$  according to Eq. (9) as mentioned above and compute the projection error based on the rendering depth as:

$$\begin{cases} L_{rg}^{i \rightarrow j} = \|p_j - p_{i \rightarrow j}(\nu_i^l)\| \\ L_{rg}^{j \rightarrow i} = \|p_i - p_{j \rightarrow i}(\nu_j^l)\|, \end{cases} \quad (13)$$

and we take the average error of all ray pairs as the final rendering geometry loss  $L_{rg}$ .

Table 1: **Quantitative comparisons on the LLFF and IBRNet datasets with 3 training views.** Best results are in **bold**. We run our method 5 times and report the error bar in the appendix.

Method	Approach	LLFF				IBRNet			
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	AVG $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	AVG $\downarrow$
Mip-NeRF [1]	NeRF-based	14.62	0.351	0.495	0.246	15.83	0.406	0.488	0.223
RegNeRF [32]		19.08	0.587	0.336	0.149	19.05	0.542	0.377	0.152
FreeNeRF [58]		19.63	0.612	0.308	0.134	19.76	0.588	0.333	0.135
SparseNeRF [51]		19.86	0.624	0.328	0.127	19.90	0.593	0.364	0.137
3DGS [19]	3DGS-based	16.46	0.440	0.401	0.192	17.79	0.538	0.377	0.166
FSGS [65]		20.43	0.682	0.248	-	19.84	0.648	0.306	0.130
DNGaussian [23]		19.12	0.591	0.294	0.132	19.01	0.616	0.374	0.151
SCGaussian (Ours)		<b>20.77</b>	<b>0.705</b>	<b>0.218</b>	<b>0.105</b>	<b>21.59</b>	<b>0.731</b>	<b>0.233</b>	<b>0.097</b>

### 3.4 Overall pipeline

**Loss function.** Our loss function consists of three parts: the ordinary photometric loss  $L_{photo}$ , the Gaussian position loss  $L_{gp}$ , the rendering geometry loss  $L_{rg}$ , and the full function is defined as:

$$L = L_{photo} + \beta L_{gp} + \delta L_{rg}. \quad (14)$$

**Training details.** During training, we set  $\beta = 1.0$ . To avoid the model falling into sub-optimization in the early stage of training, we set  $\delta = 0$  and then increase it to  $\delta = 0.3$  after 1k iterations. To ensure that the Gaussian primitive converges to the optimal position, we use a caching strategy in the first 1k iterations, *i.e.*, cache the position with the minimum Gaussian position loss  $L_{gp}$  at each iteration. Meanwhile, considering there are some mismatched ray pairs in the matching prior, we further filter out those primitives with large Gaussian position loss  $L_{gp} > \eta$ . During optimization, the ray-based primitive will not be pruned. We build our model based on the official 3DGS codebase, and train the model for 3k iterations with the same setting as 3DGS but set the learning rate of the learnable distance factor  $z$  to 0.1 at the beginning and decrease to  $1.6 \times 10^{-6}$ .

## 4 Experiments

In this section, we demonstrate the performance of our model in popular datasets and conduct ablation studies to verify the effectiveness of our designs. Next, we first describe the common datasets and the selected baselines for comparison, then analyze the results.

**Datasets & metrics.** We evaluate our model on forward-facing, complex large-scale and surrounding datasets under the sparse setting: LLFF [30], IBRNet [52], Tanks and Temples (T&T) [21], DTU [16] and NeRF Blender Synthetic dataset (Blender) [29]. LLFF dataset contains 8 real scenes, and following previous methods [32, 51], every 8-th images are held out for testing, and sparse views are evenly sampled from the remaining images for training. IBRNet dataset is also a real forward-facing dataset, and we select 9 scenes for evaluation and adopt the same split as in LLFF. T&T is a large-scale dataset collected from more complex realistic environments containing both indoor and outdoor scenes, and we use 8 scenes for evaluation and also apply the same split as in LLFF. DTU is an object-centric dataset, which contains more texture-poor scenes. We use the same evaluation strategy as [32] on DTU. For Blender, containing 8 object-centric synthetic scenes, we follow [58] to train with 8 images and test on 25 images. We report PSNR, SSIM, and LPIPS scores to measure our reconstruction quality and also report the geometric average (AVG) of  $MSE = 10^{-PSNR/10}$ ,  $\sqrt{1 - SSIM}$  and LPIPS as in [32].

**Baselines.** We compare our model against both NeRF-based and 3DGS-based few-shot NVS methods. For NeRF-based methods, we compare with methods with relatively high performance, including MipNeRF [1], DietNeRF [15], RegNeRF [32], FreeNeRF [58] and SparseNeRF [51]. For 3DGS-based methods, we compare with the vanilla 3DGS and its recent few-shot follow-ups like FSGS [65] and DNGaussian [23].

### 4.1 Results

**Results on LLFF and IBRNet.** We use the aforementioned split method to sample 3 images for training. The quantitative comparisons on two datasets with recent SOTA methods are summarized in Tab. 1. Although 3DGS-based methods natively have a weakness in invisible areas due to their




Figure 4: **Qualitative comparisons on LLFF (first two rows) and IBRNet (last two rows) datasets with 3 training views.** The reconstruction of our method is more accurate and exhibits finer details.

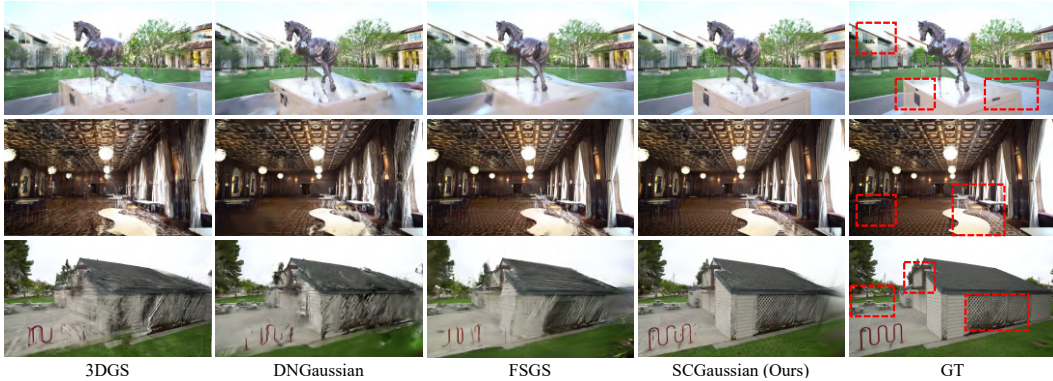


Figure 5: **Qualitative comparisons on Tanks and Temples dataset with 3 training views.**

discrete properties as discussed in [23], our SCGaussian still achieves the best performance in all metrics. Note that FSGS [65] uses the sparse SFM points for initialization, even though, our model holds remarkable superiority. Moreover, our advantage against previous methods is amplified in the IBRNet dataset, which has more low-texture scenes. Some qualitative comparisons are shown in Fig. 4, from which we can see that our method can recover more accurate high-frequency details.

**Results on T&T.** To evaluate the performance of our model on complex large scenes, we conduct further comparisons on the T&T dataset. Using the same split strategy as LLFF, we quantitatively compare with existing methods with 3 training views in Tab. 2. With the large difference in camera poses and the unbounded scene range, previous NeRF-based methods [51, 58, 32, 1], mostly designed for the bounded scenes, are hard to reconstruct plausible results. Among them, methods [32, 51] using geometric regularization perform better. Even combined with explicit point representation, the recent 3DGS still struggles on this large scene with sparse inputs. Although some recent efforts apply the monocular depth prior [23] or the initialization of sparse SFM point [65] to 3DGS, they have limited reconstruction quality. From the qualitative comparisons shown in Fig. 5, we can see that our method can synthesize the novel view with more accurate and complete details. Benefiting from our novel design in

Table 2: **Quantitative comparisons on the T&T dataset with 3 training views.**

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
MipNeRF [1]	12.57	0.241	0.623
RegNeRF [32]	13.12	0.268	0.618
FreeNeRF [58]	12.30	0.308	0.636
SparseNeRF [51]	13.66	0.331	0.615
3DGS [19]	17.14	0.493	0.397
FSGS [65]	20.01	0.652	0.323
DNGaussian [23]	18.59	0.573	0.437
<b>SCGaussian (Ours)</b>	<b>22.17</b>	<b>0.752</b>	<b>0.257</b>






Figure 6: **Qualitative comparisons on Blender dataset.**

Table 3: **Quantitative comparisons on the Blender dataset.**

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
NeRF [29]	14.934	0.687	0.318
NeRF Simple [15]	20.092	0.822	0.179
Mip-NeRF [1]	20.890	0.830	0.168
DietNeRF [15]	23.147	0.866	0.109
DietNeRF + ft [15]	23.591	0.874	0.097
FreeNeRF [58]	24.259	0.883	0.098
SparseNeRF [51]	22.410	0.861	0.119
3DGS [19]	22.226	0.858	0.114
FSGS [65]	24.640	<b>0.895</b>	0.095
DNGaussian [23]	24.305	0.886	0.088
<b>SCGaussian (Ours)</b>	<b>25.618</b>	<b>0.894</b>	<b>0.086</b>

Table 4: **Quantitative comparisons on the DTU dataset with 3 training views.**

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	AVG $\downarrow$
FreeNeRF	19.92	0.787	0.182	0.098
SparseNeRF	19.55	0.769	0.201	0.102
DNGaussian	18.91	0.790	0.176	0.102
<b>Ours</b>	<b>20.56</b>	<b>0.864</b>	<b>0.122</b>	<b>0.078</b>

Table 5: **Comparisons with methods of directly initializing with triangulation points, on LLFF.**

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	AVG $\downarrow$
3DGS (baseline)	16.46	0.440	0.401	0.192
Triang. init + 3DGS	19.11	0.643	0.335	0.140
Triang. init + ScaffoldGS	19.41	0.699	0.217	0.118
Triang. init + OctreeGS	19.61	0.710	0.210	0.114
<b>Ours</b>	<b>20.77</b>	<b>0.705</b>	<b>0.218</b>	<b>0.105</b>

hybrid representation and explicit optimization of rendering geometry and position of Gaussian primitives, our model can learn more consistent structure as shown in Fig. 1, and show great generalization ability on these large scenes.

**Results on DTU.** We further conduct more experiments on DTU dataset to prove the robustness on more texture-poor scenes. The quantitative results in Tab. 4 indicate that our method achieves the best performance on all metrics, which proves that our model is still robust on those texture-poor scenes. The qualitative results in Fig. 7 also demonstrate that our method can recover more accurate details.

**Results on Blender.** We test on the Blender dataset to verify our performance in the surrounding scenario. While [65] uses its uppooling strategy to clone more Gaussian primitives and gets the best SSIM score, our method achieves the best PSNR and LPIPS scores, as the quantitative results reported in Tab. 3. We visualize more qualitative comparisons in Fig. 6 to demonstrate our superiority, and we can see that our method has a clear advantage in recovering fine details and reconstructing complete structures. This further demonstrates our generalization ability in different scenes.

## 4.2 Analysis

**Ablation studies.** We conduct a few ablation studies on LLFF and T&T datasets to understand how our model performs with different settings. Our baseline is the vanilla 3DGS. From the results shown in Tab. 6, we can see that using only the hybrid representation (Hybrid Rep.), our model can already bring more than 2dB improvement to the baseline, which verifies that our model can indeed mitigate the risk of overfitting. Combined with the optimization of rendering geometry (Rend. Geo.), the performance can be further improved. When we optimize both the rendering geometry and the position of Gaussian primitives (Dual optim.), the model can learn the more consistent scene structure and render more convincing novel views. These results prove our motivation for learning the 3D consistent structure. And the adopted cache & filter strategy further mitigates the impact of wrong matching priors. To assess the performance of the model with different numbers of views, we conduct the comparison in T&T dataset, as shown in Fig. 8. Our model can

Table 6: **Ablation studies on LLFF dataset.**

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Baseline	16.46	0.440	0.401
w/ Hybrid Rep.	18.62	0.607	0.273
w/ Rend. Geo.	19.40	0.634	0.259
w/ Dual optim.	20.69	0.703	<b>0.205</b>
w/ Cache & filter	<b>20.77</b>	<b>0.705</b>	0.218

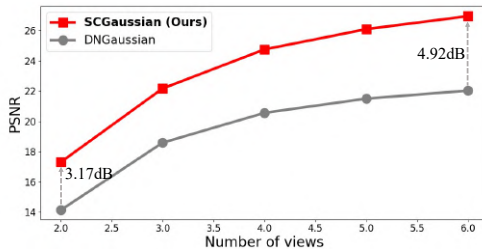


Figure 8: **PSNR vs view number on T&T.**




Figure 7: **Qualitative comparisons on DTU dataset with 3 training views.**

consistently outperform the SOTA method [23], and the advantage becomes more significant as the number of views increases.

**Triangulation initialization.** To prove the effectiveness of our optimization strategy, we perform some comparisons with methods that directly use the triangulation points of matched pixels for initialization. The results are shown in Tab. 5, which indicate that using the triangulation initialization can improve the performance of the baseline especially equipped with more structured ScaffoldGS [27] or OctreeGS [38] models. Even though, our model still achieves the best performance and demonstrates the effectiveness of our model.

**Robustness to matching models.** We perform more experiments in Tab. 7 to verify the robustness of our model to different pre-trained matching models. Concretely, we use the same optimization and testing configuration for all models, and additionally use the DKM [9], LoFTR [48] and SuperGlue [40] models to extract the matching prior. The results in Tab. 7 show that all these matching models can bring a satisfactory improvement to the baseline, and our method can even achieve better performance when using weaker matching models (*e.g.*, GIM vs. LoFTR). These results prove that our strategy is robust to different matching models.

Table 7: **Robustness to different matching models**, conducted on the LLFF dataset.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	AVG $\downarrow$
Ours + GIM	20.77	0.705	0.218	0.105
Ours + DKM	20.92	0.732	0.189	0.099
Ours + LoFTR	<b>20.94</b>	<b>0.737</b>	<b>0.182</b>	<b>0.097</b>
Ours + SuperGlue	20.25	0.689	0.221	0.110

**Efficiency.** With a single NVIDIA RTX 3090 GPU, the training of our method consumes about 3GB memories and converges within 1 minute on LLFF 3-view setting, which is much faster than existing methods, *e.g.*, [58, 51] need about 10 hours and [65] needs about 10 minutes. Our method also achieves a real-time inference speed of over 200FPS at  $504 \times 378$  resolution, superior to NeRF-based methods (*e.g.*, [58] at 0.04FPS) and comparable to 3DGS-based methods (*e.g.*, [23] at 181FPS).

**Limitation.** Following the common pipeline in the research field of few-shot NVS, our model requires an accurate camera pose, which may not always be available. Thus liberating this limitation could further improve our work to be more practical, and we will investigate this in future work.

## 5 Conclusion

In this paper, we observed the main challenge of few-shot 3DGS is learning the 3D consistent scene structure, and we exploited the matching prior to construct a Structure Consistent Gaussian Splatting method named *SCGaussian*. Due to the optimization ambiguity of Gaussian attributes between the position and shape, we presented two approaches to optimize the scene structure: explicitly optimize the rendering geometry and the position of Gaussian primitives. While directly constraining the position is non-trivial in the vanilla 3DGS, we introduced a hybrid Gaussian representation, consisting of ordinary non-structure Gaussian primitives and ray-based Gaussian primitives. In this way, both the position and shape of Gaussian primitives can be optimized to be 3D consistent. To evaluate our method as comprehensively as possible, we conducted experiments on forward-facing, complex large-scale, and surrounding datasets. The results consistently demonstrate that our method achieves new state-of-the-art performance while being highly efficient.

## Acknowledgments and Disclosure of Funding

This work is financially supported by the Outstanding Talents Training Fund in Shenzhen, this work is also supported by the National Natural Science Foundation of China U21B2012, Shenzhen Science and Technology Program-Shenzhen Cultivation of Excellent Scientific and Technological Innovation Talents project(Grant No. RCJC20200714114435057). J. Jiao is supported by the Royal Society Short Industry Fellowship (SIF\R1\231009) and the Amazon Research Award. In addition, we sincerely thank all assigned anonymous reviewers, whose comments were constructive and very helpful to our writing and experiments.

## References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *ICCV*, pages 5855–5864, 2021.
- [2] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, pages 5470–5479, 2022.
- [3] David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats from image pairs for scalable generalizable 3d reconstruction. In *CVPR*, 2023.
- [4] Gaurav Chaurasia, Sylvain Duchene, Olga Sorkine-Hornung, and George Drettakis. Depth synthesis and local warps for plausible image-based navigation. *ACM TOG*, 32(3):1–12, 2013.
- [5] Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorf: Tensorial radiance fields. In *ECCV*, pages 333–350. Springer, 2022.
- [6] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *ICCV*, pages 14124–14133, 2021.
- [7] Paul E Debevec, Camillo J Taylor, and Jitendra Malik. Modeling and rendering architecture from photographs: A hybrid geometry-and image-based approach. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 465–474. 2023.
- [8] Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views and faster training for free. In *CVPR*, pages 12882–12891, 2022.
- [9] Johan Edstedt, Ioannis Athanasiadis, Mårten Wadenbäck, and Michael Felsberg. Dkm: Dense kernelized feature matching for geometry estimation. In *CVPR*, pages 17765–17775, 2023.
- [10] John Flynn, Michael Broxton, Paul Debevec, Matthew DuVall, Graham Fyffe, Ryan Overbeck, Noah Snavely, and Richard Tucker. Deepview: View synthesis with learned gradient descent. In *CVPR*, pages 2367–2376, 2019.
- [11] John Flynn, Ivan Neulander, James Philbin, and Noah Snavely. Deepstereo: Learning to predict new views from the world’s imagery. In *CVPR*, pages 5515–5524, 2016.
- [12] Sara Fridovich-Keil, Giacomo Meanti, Frederik Rahbæk Warburg, Benjamin Recht, and Angjoo Kanazawa. K-planes: Explicit radiance fields in space, time, and appearance. In *CVPR*, pages 12479–12488, 2023.
- [13] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *CVPR*, pages 5501–5510, 2022.
- [14] Stephan J Garbin, Marek Kowalski, Matthew Johnson, Jamie Shotton, and Julien Valentin. Fastnerf: High-fidelity neural rendering at 200fps. In *ICCV*, pages 14346–14355, 2021.
- [15] Ajay Jain, Matthew Tancik, and Pieter Abbeel. Putting nerf on a diet: Semantically consistent few-shot view synthesis. In *ICCV*, pages 5885–5894, 2021.

- [16] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *CVPR*, pages 406–413, 2014.
- [17] Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. Geonerf: Generalizing nerf with geometry priors. In *CVPR*, pages 18365–18375, 2022.
- [18] Nikhil Keetha, Jay Karhade, Krishna Murthy Jatavallabhula, Gengshan Yang, Sebastian Scherer, Deva Ramanan, and Jonathon Luiten. Splatam: Splat, track & map 3d gaussians for dense rgb-d slam. In *CVPR*, 2024.
- [19] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 42(4):1–14, 2023.
- [20] Mijeong Kim, Seonguk Seo, and Bohyung Han. Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *CVPR*, pages 12912–12921, 2022.
- [21] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4):1–13, 2017.
- [22] Marc Levoy and Pat Hanrahan. Light field rendering. In *SIGGRAPH*, pages 441–452. 2023.
- [23] Jiahe Li, Jiawei Zhang, Xiao Bai, Jin Zheng, Xin Ning, Jun Zhou, and Lin Gu. Dngaussian: Optimizing sparse-view 3d gaussian radiance fields with global-local depth normalization. In *CVPR*, 2024.
- [24] Tianye Li, Mira Slavcheva, Michael Zollhoefer, Simon Green, Christoph Lassner, Changil Kim, Tanner Schmidt, Steven Lovegrove, Michael Goesele, Richard Newcombe, et al. Neural 3d video synthesis from multi-view video. In *CVPR*, pages 5521–5531, 2022.
- [25] Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. Barf: Bundle-adjusting neural radiance fields. In *ICCV*, pages 5741–5751, 2021.
- [26] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. volume 33, pages 15651–15663, 2020.
- [27] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffoldgs: Structured 3d gaussians for view-adaptive rendering. In *CVPR*, 2024.
- [28] Quan Meng, Anpei Chen, Haimin Luo, Minye Wu, Hao Su, Lan Xu, Xuming He, and Jingyi Yu. Gnerf: Gan-based neural radiance field without posed camera. In *ICCV*, pages 6351–6361, 2021.
- [29] B Mildenhall, PP Srinivasan, M Tancik, JT Barron, R Ramamoorthi, and R Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020.
- [30] Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with prescriptive sampling guidelines. *ACM TOG*, 38(4):1–14, 2019.
- [31] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM TOG*, 41(4):1–15, 2022.
- [32] Michael Niemeyer, Jonathan T Barron, Ben Mildenhall, Mehdi SM Sajjadi, Andreas Geiger, and Noha Radwan. Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *CVPR*, pages 5480–5490, 2022.
- [33] Keunhong Park, Utkarsh Sinha, Jonathan T Barron, Sofien Bouaziz, Dan B Goldman, Steven M Seitz, and Ricardo Martin-Brualla. Nerfies: Deformable neural radiance fields. In *ICCV*, pages 5865–5874, 2021.
- [34] Rui Peng, Rongjie Wang, Zhenyu Wang, Yawen Lai, and Ronggang Wang. Rethinking depth estimation for multi-view stereo: A unified representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8645–8654, 2022.



- [35] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *CVPR*, pages 10318–10327, 2021.
- [36] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *ICCV*, pages 12179–12188, 2021.
- [37] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE TPAMI*, 44(3):1623–1637, 2020.
- [38] Kerui Ren, Lihan Jiang, Tao Lu, Mulin Yu, Linning Xu, Zhangkai Ni, and Bo Dai. Octree-gs: Towards consistent real-time rendering with lod-structured 3d gaussians. *arXiv preprint arXiv:2403.17898*, 2024.
- [39] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *CVPR*, pages 12892–12901, 2022.
- [40] Paul-Edouard Sarlin, Daniel DeTone, Tomasz Malisiewicz, and Andrew Rabinovich. Superglue: Learning feature matching with graph neural networks. In *CVPR*, pages 4938–4947, 2020.
- [41] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *CVPR*, pages 4104–4113, 2016.
- [42] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, pages 501–518. Springer, 2016.
- [43] Seunghyeon Seo, Donghoon Han, Yeonjin Chang, and Nojun Kwak. Mixnerf: Modeling a ray with mixture density for novel view synthesis from sparse inputs. In *CVPR*, pages 20659–20668, 2023.
- [44] Xuelun Shen, Zhipeng Cai, Wei Yin, Matthias Müller, Zijun Li, Kaixuan Wang, Xiaozhi Chen, and Cheng Wang. Gim: Learning generalizable image matcher from internet videos. In *ICLR*, 2024.
- [45] Sudipta Sinha, Drew Steedly, and Rick Szeliski. Piecewise planar stereo for image-based rendering. In *ICCV*, pages 1881–1888, 2009.
- [46] Nagabhushan Somraj, Adithyan Karanayil, and Rajiv Soundararajan. Simplenerf: Regularizing sparse input neural radiance fields with simpler solutions. In *SIGGRAPH Asia*, pages 1–11, 2023.
- [47] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *CVPR*, pages 5459–5469, 2022.
- [48] Jiaming Sun, Zehong Shen, Yuang Wang, Hujun Bao, and Xiaowei Zhou. Loftr: Detector-free local feature matching with transformers. In *CVPR*, pages 8922–8931, 2021.
- [49] Stanislaw Szymanowicz, Christian Rupprecht, and Andrea Vedaldi. Splatter image: Ultra-fast single-view 3d reconstruction. In *CVPR*, 2023.
- [50] Prune Truong, Marie-Julie Rakotosaona, Fabian Manhardt, and Federico Tombari. Sparf: Neural radiance fields from sparse and noisy poses. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4190–4200, 2023.
- [51] Guangcong Wang, Zhaoxi Chen, Chen Change Loy, and Ziwei Liu. Sparsenerf: Distilling depth ranking for few-shot novel view synthesis. In *ICCV*, pages 9065–9076, 2023.
- [52] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *CVPR*, pages 4690–4699, 2021.
- [53] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004.

- [54] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. Nerf-: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021.
- [55] Daniel N Wood, Daniel I Azuma, Ken Aldinger, Brian Curless, Tom Duchamp, David H Salesin, and Werner Stuetzle. Surface light fields for 3d photography. In *SIGGRAPH*, pages 487–496. 2023.
- [56] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *CVPR*, 2024.
- [57] Wangze Xu, Huachen Gao, Shihe Shen, Rui Peng, Jianbo Jiao, and Ronggang Wang. Mvpgs: Excavating multi-view priors for gaussian splatting from sparse input views. In *ECCV*, 2024.
- [58] Jiawei Yang, Marco Pavone, and Yue Wang. Freenerf: Improving few-shot neural rendering with free frequency regularization. In *CVPR*, pages 8254–8263, 2023.
- [59] Zeyu Yang, Hongye Yang, Zijie Pan, Xiatian Zhu, and Li Zhang. Real-time photorealistic dynamic scene representation and rendering with 4d gaussian splatting. In *ICLR*, 2024.
- [60] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, pages 4578–4587, 2021.
- [61] Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3d gaussian splatting. In *CVPR*, 2024.
- [62] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *NeurIPS*, 35:25018–25032, 2022.
- [63] Tinghui Zhou, Richard Tucker, John Flynn, Graham Fyffe, and Noah Snavely. Stereo magnification: Learning view synthesis using multiplane images. *ACM TOG*, 2018.
- [64] Tinghui Zhou, Shubham Tulsiani, Weilun Sun, Jitendra Malik, and Alexei A Efros. View synthesis by appearance flow. In *ECCV*, pages 286–301. Springer, 2016.
- [65] Zehao Zhu, Zhiwen Fan, Yifan Jiang, and Zhangyang Wang. Fsgs: Real-time few-shot view synthesis using gaussian splatting. *arXiv preprint arXiv:2312.00451*, 2023.

## A Appendix

### A.1 More experimental details

Our experiments are performed following the common solution of existing methods. LLFF dataset contains eight different scenes, and we perform the training and inference at  $8\times$  downsampling scale with a resolution of  $504 \times 378$ . IBRNet is another forward-facing dataset collected by [52] that contains larger camera motion and more scenes. We select nine scenes for evaluation, which include ‘giraffe\_plush’, ‘yamaha\_piano’, ‘sony\_camera’, ‘Japanese\_camilia’, ‘scaled\_model’, ‘dumbbell\_jumprope’, ‘hat\_on\_fur’, ‘roses’ and ‘plush\_toys’. We use the same setting as LLFF. Tanks and Temples is a large-scale complex dataset, which has large camera motion. For comparison, we use eight scenes, namely ‘Ballroom’, ‘Barn’, ‘Church’, ‘Family’, ‘Francis’, ‘Horse’, ‘Ignatius’, and ‘Museum’, both indoors and outdoors. We train and infer at the resolution of  $960 \times 540$ . For the Blender dataset, we use the common solution in existing methods and run at the resolution of  $400 \times 400$  ( $2\times$  downsampling).

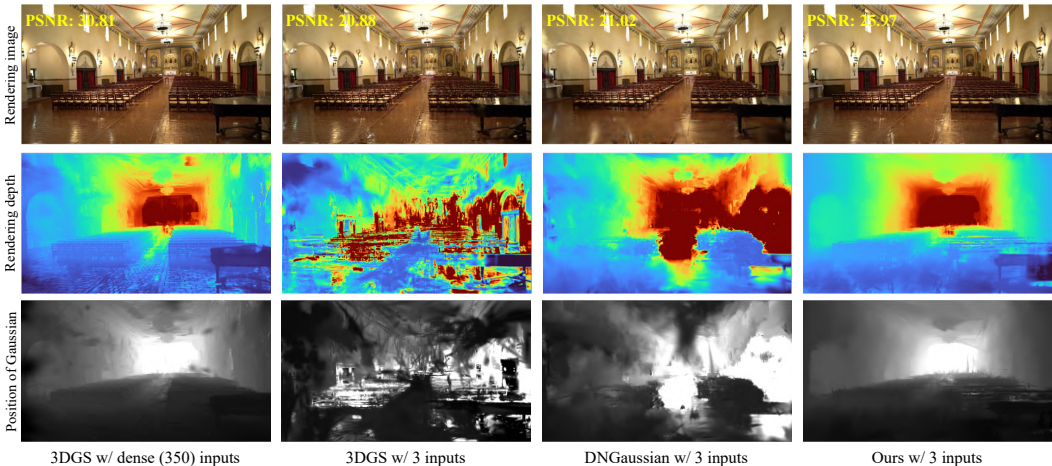


Figure 9: **Comparisons in view synthesis, rendering geometry and position of Gaussian primitive.** The first, second and third rows show the synthesized novel view of four methods, their rendering depth and the position of nearest Gaussian primitives, respectively.

### A.2 More results of consistent structures

The consistent scene structure is important for reconstruction models, including the NeRF-based and the 3DGS-based, and the inaccurate structure, *e.g.*, floaters or walls, can lead to extremely poor novel view synthesis results. Thus in this paper, we propose SCGaussian to solve the challenge of learning consistent structure in few-shot 3DGS models. From the quantitative and qualitative results shown in our main paper, we can see that our method can synthesize more complete novel views, especially in the high-frequency regions, and these results just demonstrate our learned scene structure is more 3D consistent.

We show some comparisons of rendering geometry in Fig. 1, and we can see that when the input becomes sparse, existing methods fail to learn the plausible geometry while our method can still render the accurate geometry. Here, we show more comparisons in Fig. 9 to understand the results of the position of Gaussian primitives. To visualize these positions, we first fix the opacity of all primitives to a large value (1.0 in our setting) and then we render the distance of Gaussian primitives using the Gaussian rasterization. In this way, the rendering results can indicate the position of the nearest Gaussian primitives (third row in Fig. 9). We can see that both the rendering geometry and the position of Gaussian primitives of our method are more accurate and 3D consistent.

Meanwhile, we find that our learned structure in texture-less regions (*e.g.*, the wall region shown in Fig. 9 3rd row) is even better (smoother) than the dense version (with way more views of inputs) of 3DGS. We suspect the main reason is that the proposed method has better control over the number of Gaussian primitives, *i.e.*, adaptively allocates more primitives in high-textured regions while fewer primitives in the texture-less regions.

### A.3 Effectiveness of dual optimization based on the hybrid representation

In this paper, we propose a dual optimization strategy to separately optimize the rendering geometry and position of Gaussian primitives based on our hybrid Gaussian representation. While we show some quantitative results in Tab. 6, here we show more visual results in Fig. 10. The model ‘w/ Matching prior’ in Tab. 6 refers to the straightforward combination of matching priors and the vanilla 3DGS, and the model ‘w/ Dual optim’ corresponds to the model optimizes both rendering geometry and position of Gaussian primitives based on the hybrid representation.

We can see that the straightforward solution is still hard to synthesize accurate novel views and still suffers from obvious inconsistencies in its rendering geometry. And our solution, optimizing both the rendering geometry and position of Gaussian primitives based on our hybrid representation, can synthesize more accurate novel views and render more consistent depth.

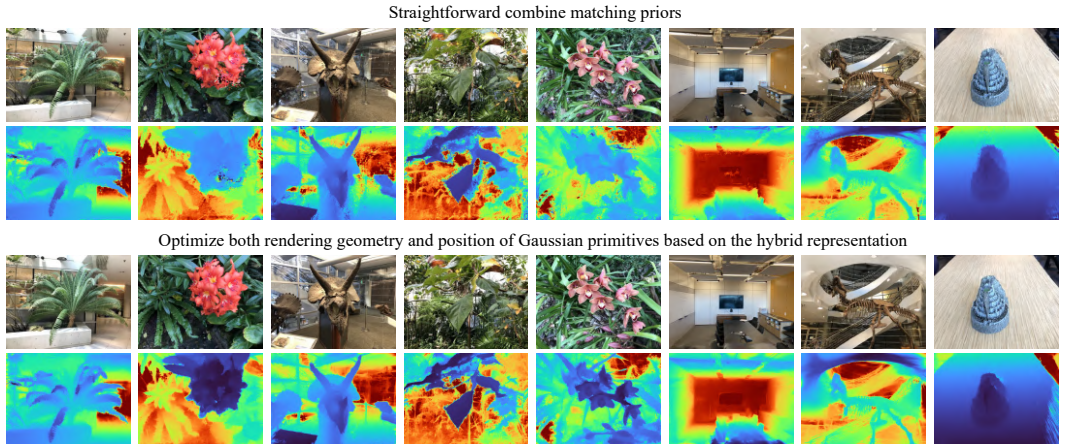


Figure 10: Results of using the straightforward combination and the dual optimization.

### A.4 Results at different resolutions

While 3DGS model can synthesize high-resolution images efficiently, we here conduct more comparisons with existing methods at a higher resolution ( $1008 \times 756$ ) on LLFF dataset. As the quantitative results shown in Tab. 8, our method still achieves the best in all metrics. Compared with the previous few-shot 3DGS method [23], our advantage is amplified at the higher resolution.

We further show some visual comparisons in Fig. 11. We can see that our method can recover more high-frequency details with the best accuracy, while previous 3DGS-based method [23] even loses some structures. Compared with the NeRF-based methods [58, 51], which have a slow rendering speed and smooth reconstruction, our advantage is more significant.

Table 8: Quantitative comparisons on LLFF under different resolutions. Experiments are conducted with 3 training views.

Method	res $8 \times (504 \times 378)$			res $4 \times (1008 \times 756)$		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
MipNeRF [1]	14.62	0.351	0.495	15.53	0.416	0.490
RegNeRF [32]	19.08	0.587	0.336	18.40	0.545	0.405
FreeNeRF [58]	19.63	0.612	0.308	19.12	0.568	0.393
SparseNeRF [51]	19.86	0.620	0.329	19.30	0.565	0.413
3DGS [19]	16.46	0.401	0.440	15.92	0.504	0.370
DNGaussian [23]	19.12	0.591	0.294	18.03	0.574	0.394
<b>SCGaussian (Ours)</b>	<b>20.77</b>	<b>0.705</b>	<b>0.218</b>	<b>20.09</b>	<b>0.679</b>	<b>0.252</b>





Figure 11: **Qualitative comparisons at the high resolution** ( $1008 \times 756$ ).

### A.5 Results for different view numbers

As shown in Fig. 8, our method can consistently outperform existing methods with different numbers of inputs. Here, we show some visual comparisons in Fig. 12 to qualitatively evaluate our advantage. We can see that our method can recover more details with both 3 and 6 training views. Furthermore, we report some quantitative comparisons in Tab. 9. The results show that the 3DGS-based methods perform better than the NeRF-based method on the complex scene. FreeNeRF [58] propose efficient frequency regularization terms to improve the few-shot performance in the bounded scene, but we can see that this simple strategy does not work well in the unbound scene, while SparseNeRF [51] uses the monocular depth prior to achieve better performance than FreeNeRF. This suggests that using external priors may be a better option in complex scenarios. Meanwhile, we find a situation that the vanilla 3DGS performs better than DNGaussian, which adopts the monocular depth to regularize the geometry. We analyze that the main reason is the inherent scale and multi-view inconsistency of monocular depth. This situation further demonstrates the superiority of the matching prior that we adopt.

Table 9: **Quantitative comparisons on T&T with different numbers of inputs.**

Method	3 views			6 views		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
FreeNeRF [58]	12.30	0.308	0.636	14.34	0.375	0.586
SparseneRF [51]	13.66	0.331	0.615	17.50	0.454	0.539
3DGS [19]	17.14	0.493	0.397	22.27	0.702	0.275
DNGaussian [23]	18.59	0.573	0.437	22.03	0.687	0.382
<b>SCGaussian (Ours)</b>	<b>22.17</b>	<b>0.752</b>	<b>0.257</b>	<b>26.95</b>	<b>0.869</b>	<b>0.149</b>



Figure 12: Qualitative comparisons with 3 and 6 training views.

### A.6 More discussion on the hybrid representation

Optimizing the position of Gaussian primitives to the 3D consistent surface position is fundamental for the novel view synthesis task. However the Gaussian primitive in the vanilla 3DGS is non-structure and is hard to be controlled, whose position can be moved to arbitrary directions. While the dense counterpart can leverage the initialization of SFM points to guide the optimization of the Gaussian primitive, our few-shot model with sparse inputs can only start from the random initialization, which obviously makes the optimization of the position of Gaussian primitives become more difficult. Thus, it would be an ideal solution if there was a method that could directly control the position of the Gaussian primitive.

With the ray correspondence in the matching prior, we can assume that there is a surface point in the matching ray. Therefore, we propose to bind Gaussian primitives to matching rays, restrict the optimization of their positions along the ray and enforce them to converge to the surface position. This approach makes the optimization of Gaussian primitives more controllable. Meanwhile, we notice that there are still regions that are not multi-view visible, only using these ray-based Gaussian primitives makes it hard to cover the complete scene, as shown in Fig. 13. Here, we treat these regions only visible to a single view as the ‘background’, and we use the ordinary non-structure Gaussian primitives to recover them and propose the hybrid representation. We report the ablation results of the hybrid representation in Tab. 8.

Table 10: Ablation results of the hybrid representation on LLFF with 3 training views.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
only non-structure	19.40	0.634	0.259
only ray-based	20.50	0.684	0.231
hybrid rep.	<b>20.77</b>	<b>0.705</b>	<b>0.218</b>




Figure 13: Visualization of matching points on two scenes.

## A.7 Error bars

Although most previous don't provide the error bar, here, to enhance the experimental significance, we run all methods 5 times and report error bars of SparseNeRF [51], 3DGS [19], DNGaussian [23] and our method in Fig. 14. We can see that the results of the baseline model 3DGS [19] have the largest fluctuations in all metrics. Our method gets the best score on all metrics and has relatively satisfactory stability.



Figure 14: Error bars of SparseNeRF [51], 3DGS [19], DNGaussian [23] and our method.

## A.8 Video comparisons

To better illustrate the effectiveness of our method, here we further provide visual comparisons in video format. Please refer to <https://drive.google.com/drive/folders/1sTpuRRV4YYJOPTpQYb-ZChrck37GtU6h?usp=sharing> for more details.



## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [Yes]

Justification: We summarize our main contributions and the main experimental performance in our abstract and introduction.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [Yes]

Justification: We discuss our limitations in Sec. 4.2.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when the image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [Yes]



Justification: We have clarified and checked our formulas.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

#### 4. **Experimental Result Reproducibility**

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: We detail our loss functions and training details in our paper and we promise to release our code when the paper is accepted.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
  - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

#### 5. **Open access to data and code**

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [No]

Justification: We don't introduce new dataset. And due to the rush of merging the final manuscripts, we don't have enough time to prepare our codebase, and we promise to release our code in the future.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We have detailed our training details and experimental settings in our paper.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: Although most previous few-shot NVS methods don't provide their error bars, we run our model for multiple times and report the error bar in Fig. 14 of our Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

## 8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: The experiments of our method are conducted on a single NVIDIA 3090 GPU.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

## 9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: We confirm that our research conducted in the paper conforms with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

## 10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: Our paper don't have obvious societal impacts currently.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

## 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper poses no such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

## 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: We cite and comply with the licenses of the public datasets we use.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

### 13. **New Assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: Our paper doesn’t release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and Research with Human Subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: Our paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.