

# DeblurGS: Gaussian Splatting for Camera Motion Blur

Jeongtaek Oh<sup>1</sup>, Jaeyoung Chung<sup>2</sup>, Dongwoo Lee<sup>2</sup>, and Kyoung Mu Lee<sup>1,2</sup>

<sup>1</sup> IPAI, Seoul National University, Korea

<sup>2</sup> Dept. of ECE & ASRI, Seoul National University, Korea  
 {ohjtgood, robot0321, dongwoo.lee, kyoungmu}@snu.ac.kr



**Fig. 1: Novel View Synthesis with Blurry Views.** Our DeblurGS achieves state-of-the-art performance in novel view synthesis and deblurring compared to previous approaches.

**Abstract.** Although significant progress has been made in reconstructing sharp 3D scenes from motion-blurred images, a transition to real-world applications remains challenging. The primary obstacle stems from the severe blur which leads to inaccuracies in the acquisition of initial camera poses through Structure-from-Motion, a critical aspect often overlooked by previous approaches. To address this challenge, we propose DeblurGS, a method to optimize sharp 3D Gaussian Splatting from motion-blurred images, even with the noisy camera pose initialization. We restore a fine-grained sharp scene by leveraging the remarkable reconstruction capability of 3D Gaussian Splatting. Our approach estimates the 6-Degree-of-Freedom camera motion for each blurry observation and synthesizes corresponding blurry renderings for the optimization process. Furthermore, we propose Gaussian Densification Annealing strategy to prevent the generation of inaccurate Gaussians at erroneous locations during the early training stages when camera motion is still imprecise. Comprehensive experiments demonstrate that our DeblurGS achieves state-of-the-art performance in deblurring and novel view synthesis for real-world and synthetic benchmark datasets, as well as field-captured blurry smartphone videos.

**Keywords:** 3D Gaussian Splatting · Camera Motion Deblurring

## 1 Introduction

3D reconstruction and photorealistic novel view synthesis are longstanding interests for industrial applications, such as AR/VR, autonomous driving, large scene reconstruction, and robot navigation. Neural radiance fields (NeRF) [29] and its numerous following works [1, 3, 14, 45] achieve great progress in such 3D vision tasks. The practical deployment of NeRFs often involves cumbersome processes of capturing a number of individual images from different viewpoints. Therefore, recording scenes through video is a more applicable and user-friendly setting. However, multi-view input data acquired from video capturing is usually accompanied by camera motion blur, and the 3D reconstruction quality can be compromised by the blurry inputs.

We investigate the problem of novel view synthesis from motion blurred images. Several works [23, 24, 27, 42] attempt to reconstruct 3D scenes from blurry multi-view images by introducing deblur modules to NeRF’s framework. Although the aforementioned works demonstrate significant improvements in sharp reconstruction of scenes, they rely on the assumption that the exact poses of each image are known through the use of Blender [11] or a beam-splitter camera [36]. However, this assumption is unrealistic, as Structure-from-Motion (SfM) with blurry images estimates camera poses with large errors.

More recently, 3D Gaussian Splatting (3DGS) [20] has attracted attention because of its real-time and high-quality rendering with fast training time. While original 3DGS is designed to reconstruct 3D scenes from a set of clean images, we also find it an attractive option for restoring sharp scenes from blurry images because of its exceptional capability for fine-detailed reconstruction. Despite the potential of 3DGS to capture fine-grained patterns, to the best of our knowledge, the problem of optimizing clean 3DGS from blurry images has not been explored yet.

In this paper, we present DeblurGS, the first Gaussian Splatting framework designed to reconstruct 3D scenes from images blurred by camera motion. The proposed DeblurGS pipeline jointly optimizes the latent camera motion of blurry images and sharp 3D scenes represented by Gaussian Splatting. Specifically, we reconstruct blurry images by accumulating 3DGS rendered images following the estimated camera motion of training views and optimize the reconstructed blurry images to converge with input blurry images.

DeblurGS improves the deblurring performance beyond existing NeRF-based deblurring approaches by leveraging the fine-detailed rendering performance of 3DGS. Furthermore, we observe that while previous methods perform well under the ideal condition of accurately known camera poses for blurry images, they falter when blur-induced error disrupts the initial camera poses from SfM. We analyze the behavior of 3DGS when initial camera poses are inaccurate due to blur, and propose a Gaussian Densification Annealing strategy that enables stable optimization of the sharp 3D scene. In addition, we have found that different blurry images can be reconstructed depending on how sub-frame camera poses are sampled along the estimated camera motion trajectory. To achieve more ac-

curate blur optimization, we propose the sub-frame alignment parameters, which control discrete sampling intervals on the continuous motion trajectory.

The proposed method, DeblurGS, outperforms previous methods in deblurring across all existing multi-view deblurring datasets. In particular, DeblurGS is the only method capable of restoring high-quality sharp 3D scenes from the noisy camera poses obtained through SfM on existing datasets. Furthermore, our experiments for real-world captured videos demonstrate the effectiveness of DeblurGS for practical applications.

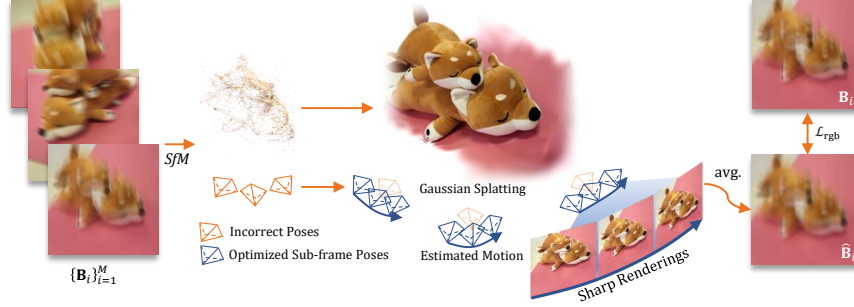
In summary, our key contributions are:

- We propose DeblurGS, the first 3D Gaussian Splatting pipeline that optimizes a sharp 3D scene from motion blur images.
- We adopt the Gaussian Densification Annealing strategy to optimize noisy initial camera poses of input blurry images.
- We demonstrate the practicality of DeblurGS through experiments on field-captured, fast-moving, real-world videos.

## 2 Related Work

**Image Deblurring** Image deblurring is one of the fundamental tasks in the image restoration field. The conventional deep learning approach to restore sharp images from blurry input is CNN or transformer-based supervised learning [5, 8, 10, 22, 39, 40, 43, 48, 50]. However, these methods require a large amount of training data paired with sharp ground truth images [31, 32, 35, 36], and the quality of deblurring is dependent on the scale of the data. Moreover, these methods often struggle with generalization across different conditions due to the domain gap problem, making their performance inconsistent in diverse real-world scenarios. In contrast, our DeblurGS framework deviates from the traditional data-driven paradigm, eliminating the need for a pre-trained network trained by large-scale datasets and inherently unaffected by a domain-gap problem.

**NeRF and 3D Gaussian Splatting** NeRF [29] draws great attention in 3D vision fields on account of its photo-realistic view synthesis results. The core strategy of NeRF is to optimize neural implicit representation with a differentiable volume rendering technique. Several follow-up approaches aim to improve the rendering quality [2, 3, 19, 44, 49], while other branches of the study are committed to mitigate the time-consuming training and rendering speed [6, 7, 9, 13, 15, 26, 30, 34, 38, 41, 47] of NeRF’s framework, improving the rendering speed by several orders of magnitude. Recently, 3DGS [20] enhances variants of the radiance fields model and achieves both detailed reconstruction performance and real-time rendering speed. By replacing NeRF’s ray marching [28] with rasterization that is efficient and deterministic, 3DGS achieves real-time rendering without loss of visual quality. Our DeblurGS is also built on a 3DGS pipeline to restore fine-grained patterns of latent sharp scenes from blurry observations and perform real-time rendering of reconstructed scenes.



**Fig. 2: Training pipeline of DeblurGS.** We simulate the physical blur operation while the camera is moving. In our optimization, the blurry images  $\{\mathbf{B}_i\}_{i=1}^M$  are reconstructed by accumulating rendered images along with the estimated camera trajectories. We minimize L1 loss  $\mathcal{L}_{\text{rgb}}$  between the input blurry images  $\{\mathbf{B}_i\}_{i=1}^M$  and reconstructed blurry images  $\{\hat{\mathbf{B}}_i\}_{i=1}^M$  to jointly optimize the camera motion trajectories and the sharp 3D scene.

**3D Reconstruction from Blur** Recently, NeRF-based deblurring approaches [23, 24, 27, 42] attempt to reconstruct a sharp 3D scene from blurry multi-view images. The NeRF based approaches jointly optimize the blur operation of each image with the sharp 3D scene that explains all blurry inputs. DeblurNeRF [27] and DP-NeRF [23] adopts a 2D pixel-wise blur kernel estimator, and BAD-NeRF [42] and ExBluRF [24] directly estimate camera trajectories of each input image. Despite impressive novel view synthesis and deblurring performance, NeRF based methods confront challenges optimizing scenes from inaccurate initial poses, which is a natural assumption considering SfM pipeline estimates erroneous camera poses if blurry views are given. Our DeblurGS establishes restoration of sharp 3D scenes from erroneous poses by adopting our Gaussian Densification Annealing strategy.

### 3 Method

We introduce DeblurGS, a method to optimize sharp Gaussian Splatting based 3D scenes from camera motion blurred images. Given multi-view observations which are blurry due to camera motion, our goal is to restore a sharp 3D scene. To this end, we adopt 3DGS [20] as a scene representation to accomplish photo-realistic recovery of 3D scenes in fine-grained detail, and jointly optimize 3DGS with the latent camera motion. We present a brief review of 3DGS in Sec. 3.1. Next, we introduce a method for blurry view synthesis by estimating camera motion and accumulating sub-frame rendering from the approximated motion in Sec. 3.2. We present the optimization process in Sec. 3.3, highlighting Gaussian Densification Annealing strategy for optimization from the erroneous poses. Lastly, we derive the loss terms for the optimization in Sec. 3.4. We provide the overview of the training procedure in Fig. 2

### 3.1 Preliminary: 3D Gaussian Splatting

3D Gaussian Splatting [20] characterizes a 3D scene as a set of Gaussian primitives  $G = \{(\mathbf{x}_i, \Sigma_i, c_i, \sigma_i)\}_{i=1}^{|G|}$ , where  $\mathbf{x}_i$ ,  $\Sigma_i$ ,  $c_i$ , and  $\sigma_i$  depict 3D coordinates, covariance matrix, RGB color and opacity of  $i$ -th primitive, respectively. Omitting the intrinsic notation of the camera for simplicity, we obtain a 2D projection of 3D Gaussian represented by mean  $\mathbf{x}_i$  and covariance  $\Sigma_i$  given camera pose  $\mathbf{P} \in SE(3)$  by:

$$\mathbf{x}_{2D,i} = \mathbf{P}\mathbf{x}_i, \quad \text{and} \quad \Sigma_{2D,i} = J\mathbf{P}\Sigma_i\mathbf{P}^\top J^\top, \quad (1)$$

where  $\mathbf{x}_{2D,i}$  and  $\Sigma_{2D,i}$  are mean and covariance of projected 2D Gaussian, respectively, and  $J$  is a Jacobian of the affine approximation of the projective transformation [52]. After projecting all Gaussians to image space, we calculate the opacity  $\alpha_i = \sigma_i \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{x}_{2D,i})^\top \Sigma_{2D,i}^{-1}(\mathbf{x} - \mathbf{x}_{2D,i})\right)$  and obtain the pixel RGB value  $\hat{c}$  at pixel  $\mathbf{x}$  by  $\alpha$ -blending:

$$\hat{c}(\mathbf{x}) = \sum_{i=1}^{|G|} \left( \prod_{j=1}^{i-1} (1 - \alpha_j(\mathbf{x})) \right) \alpha_i(\mathbf{x}) c_i. \quad (2)$$

The order of Gaussians for  $\alpha$ -blending is decided by sorting by distance from the camera to Gaussians. Consequently, we represent the rendering  $\mathcal{R}_G$  of Gaussian Splatting  $G$  from the pose  $\mathbf{P}$  as:

$$\mathcal{R}_G(\mathbf{P}) = \{\hat{c}(\mathbf{x}) \mid \mathbf{x} \in \mathcal{P}\}, \quad (3)$$

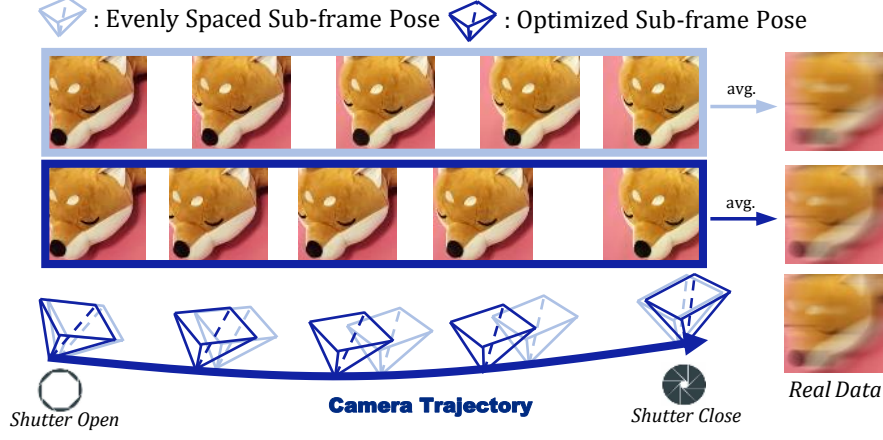
where  $\mathcal{P}$  stands for a set of all pixels on the image space.

Given multi-view images with known camera poses, the optimization of 3DGS is done by minimizing the L1 distance between the rasterized view  $\mathcal{R}_G(\mathbf{P})$  and the ground truth image from the pose  $\mathbf{P}$ . If given multi-view images are blurry caused by camera motion, naïvely employing Gaussian Splatting training procedure results in poor 3D reconstruction as shown in Fig. 1 because minimization of the L1 norm hardly converges with blurry observation. Our approach integrates auxiliary parameters that simulate blur operation into the original 3DGS pipeline to guide the Gaussian splattings converge to represent a sharp scene.

### 3.2 Blurry View Synthesis

We aim to optimize sharp Gaussian Splatting using given motion blurred inputs. Physically, the camera motion blur is generated from the integration of irradiance during camera movement [33] such as hand-shake or trembling. Therefore, the acquisition of blurry image  $\mathbf{B}$  is represented by an integration of irradiance for the time-varying 6-DOF camera pose  $\mathbf{P}_\tau \in SE(3)$  within the exposure time  $\tau \in [\tau_o, \tau_c]$ :

$$\mathbf{B} = \int_{\tau_o}^{\tau_c} \mathbf{I}(\mathbf{P}_\tau) d\tau \approx \frac{1}{N} \sum_{i=1}^N \mathbf{I}(\mathbf{P}_{\tau_i}), \quad (4)$$



**Fig. 3: Illustration of the Sub-frame Alignment Parameters.** With the estimated camera trajectory, the resulting blurry image changes based on the sampling intervals of sub-frame images. Even if the latent camera trajectory is well-optimized, the evenly sampled blurry image (1<sup>st</sup> row) does not correspond to the blurry observation of real data. We align the sub-frames intervals by optimizing the sub-frame alignment parameters to accumulate precise blurry image (2<sup>nd</sup> row).

where  $\mathbf{I}(\mathbf{P}_\tau)$  denotes sharp image acquired from the pose  $\mathbf{P}_\tau$ . Here, we approximate the integral calculation as a finite sum of  $N$  irradiance  $\mathbf{I}(\mathbf{P}_{\tau_i})$ , where the sub-interval timestamp  $\tau_i = \tau_o + \frac{i-1}{N-1}(\tau_c - \tau_o)$ .

We estimate the 6-DOF camera trajectory that explains motion blur to accurately simulate the blur operation described in Eq. (4). Following ExBluRF [24], we parameterize the rigid motion of the camera with the Bézier curve in the Lie algebra space  $\mathfrak{se}(3)$ . However, we have found that even if we accurately estimate the camera trajectory using a Bézier curve, the blurry image is not synthesized uniquely due to differences in the sampling positions of camera poses along the trajectory, as shown in Fig. 3. Therefore, we define the sub-frame alignment parameter  $\nu = \{\nu_i\}_{i=1}^N$  which calibrates each camera pose  $\hat{\mathbf{P}}(\nu_i)$  on the estimated trajectory to be aligned with the latent camera poses at time  $\tau_i$ :

$$\hat{\mathbf{P}}(\nu_i) \triangleq \mathbf{P}_{\tau_i} \quad \forall i \in \{1, 2, \dots, N\}. \quad (5)$$

We apply the definition of the alignment parameter  $\nu$  (Eq. 5) to formulate motion blur image  $\mathbf{B}$  using Eq. 4:

$$\mathbf{B} \approx \frac{1}{N} \sum_{i=1}^N \mathbf{I}(\mathbf{P}_{\tau_i}) = \frac{1}{N} \sum_{i=1}^N \mathbf{I}(\hat{\mathbf{P}}(\nu_i)), \quad (6)$$

where we abbreviate  $\hat{\mathbf{P}}(\nu_i)$  as  $\hat{\mathbf{P}}_i$  for brevity. By introducing the sub-frame alignment parameter  $\nu$ , we overcome the challenge of non-unique blur synthesis from a single trajectory.

We synthesize a motion blurred rendering by adopting Gaussian Splatting rasterization described in Sec. 3.1 to the motion blur formulation. Given approximate 6-DOF camera motion, we retrieve the sub-frame poses  $\{\hat{\mathbf{P}}_i\}_{i=1}^N$  corresponding to each sub-interval  $\{\tau_i\}_{i=1}^N$ . We generate a blurry rendering by accumulating sharp rendering of Gaussian Splatting from the sub-frame poses. We also apply a gamma correction function to the synthesized blurry view to correctly replicate the image acquisition process of the camera. When a camera converts the irradiance to a digital image, a nonlinear function is applied to brighten the raw image [12]. We choose to leverage the widely used gamma correction function  $\gamma(x) = x^{1/2.2}$  to the synthesized blurry view, following previous works on blur synthesis [24, 27, 36]. Finally, applying the rasterization process depicted in Eq. 2 to the blur operation (Eq. 6), we formulate the generation of blurry rendering  $\hat{\mathbf{B}}$  as integration of sharp rendering  $\mathcal{R}_G(\hat{\mathbf{P}}_i)$  of the 3DGS scene  $G$  from the retrieved sub-frame poses  $\hat{\mathbf{P}}_i$ , followed by the gamma correction function  $\gamma(\cdot)$ :

$$\hat{\mathbf{B}}(G) = \gamma \left( \frac{1}{N} \sum_{i=1}^N \mathcal{R}_G(\hat{\mathbf{P}}_i) \right). \quad (7)$$

We minimize the distance between synthesized blurry rendering  $\hat{\mathbf{B}}(G)$  and the blurry observation to optimize the latent sharp scene.

### 3.3 Optimization from Inaccurate Poses

Given a set of  $M$  blurry observations  $\{\mathbf{B}_i\}_{i=1}^M$ , our goal is to find 3DGS parameters  $G$  which represent a sharp scene. To achieve this, we define a set of camera motion parameters (Sec. 3.2) to estimate camera trajectory and alignment of sub-frames. For  $i$ -th camera, we assign the camera motion parameter to retrieve sub-frame poses  $\{\hat{\mathbf{P}}_j^{(i)}\}_{j=1}^N$  from the estimated motion, where  $\hat{\mathbf{P}}_j^{(i)}$  denotes the  $j$ -th sub-frame pose of  $i$ -th camera. We initialize the camera trajectory from the camera poses obtained by executing the SfM pipeline COLMAP [37] for all blurry observations and the  $N$  alignment parameter  $\{\nu_j^{(i)}\}_{j=1}^N$  to be evenly spaced within the range of  $[0, 1]$ . Subsequently, the optimization process adjusts the camera trajectory and the sub-frame parameter  $\nu^{(i)}$  toward latent camera motion. We accumulate sharp renderings from each sub-frame pose and reconstruct blurry rendering  $\hat{\mathbf{B}}_i$  as outlined in Eq. 7.

**Gaussian Densification Annealing** The initial poses obtained by COLMAP are erroneous since the conventional feature matching algorithm is conducted on noisy features from blurry images. Given the inaccurate initial poses, joint estimation of camera motion leads to the generation of the Gaussians in incorrect positions during the early stages of optimization. Specifically, if the gradient with respect to the position of the Gaussian exceeds the densification threshold  $\theta$ , the

Gaussian is split into two. Consequently, the Gaussians at incorrect positions disrupt the optimization process by trying to fit the training images from the incorrect locations. To prevent the generation of premature Gaussian splitting in an incorrect location, we employ an annealing strategy of the densification threshold  $\theta$ . We gradually anneal  $\theta$  down from a higher initial  $\theta$ , allowing for a more refined densification when the camera motion is sufficiently optimized. The annealing strategy for  $\theta$  prioritizes the camera motion’s accurate optimization before the scene attempts to represent fine details through densification, thereby mitigating the disruptions caused by Gaussians at an erroneous location.

### 3.4 Loss Functions

**Reconstruction Loss** We employ a reconstruction loss  $\mathcal{L}_{\text{rgb}}$  to converge the blurry rendering  $\hat{\mathbf{B}}_i$  to the blurry observation  $\mathbf{B}_i$ :

$$\mathcal{L}_{\text{rgb}} = \sum_i \|\mathbf{B}_i - \hat{\mathbf{B}}_i\|_1. \quad (8)$$

The reconstruction loss  $\mathcal{L}_{\text{rgb}}$  leads joint optimization of the motion trajectories, the sub-frame alignment parameters and the sharp Gaussian Splatting.

**Temporal Smoothness Loss** The reconstruction loss  $\mathcal{L}_{\text{rgb}}$  minimizes the error between the integration of sharp renderings and the corresponding blurry observation. However, the optimization relying solely on  $\mathcal{L}_{\text{rgb}}$  can converge to unrealistic latent Gaussians and motion trajectories due to the ill-posedness of the deblurring problem. To regularize the ill-posedness of the blurry view synthesis in Eq. (7), we apply a temporal smoothness loss:

$$\mathcal{L}_{\text{smooth}} = \frac{1}{N} \sum_{i,j} \|\mathcal{R}_G(\hat{\mathbf{P}}_{j+1}^{(i)}) - \mathcal{R}_G(\hat{\mathbf{P}}_j^{(i)})\|_2, \quad (9)$$

where  $\mathcal{L}_{\text{smooth}}$  penalizes the RGB difference of neighboring sub-frame renderings to prevent abrupt changes in the rendering of adjacent viewpoints. We finalize our optimization objective  $\mathcal{L}$  by combining  $\mathcal{L}_{\text{rgb}}$  to  $\mathcal{L}_{\text{smooth}}$ :

$$\mathcal{L} = \mathcal{L}_{\text{rgb}} + \lambda \mathcal{L}_{\text{smooth}}. \quad (10)$$

We assign  $\lambda = 0.05$ , gradually annealing down to  $\lambda = 0.01$  during the optimization. After the training procedure, we only keep sharp Gaussian Splatting  $G$  and all camera motion parameters are not used in sharp view rendering at inference time.

## 4 Experiment

### 4.1 Experimental Setup

**Datasets** We perform comprehensive experiments on two real-world benchmark datasets, Real-motion-blur [27] and ExBlur [24], and also synthetic dataset Synthetic Extreme Blur [24]. All datasets contain multiple scenes, and each scene

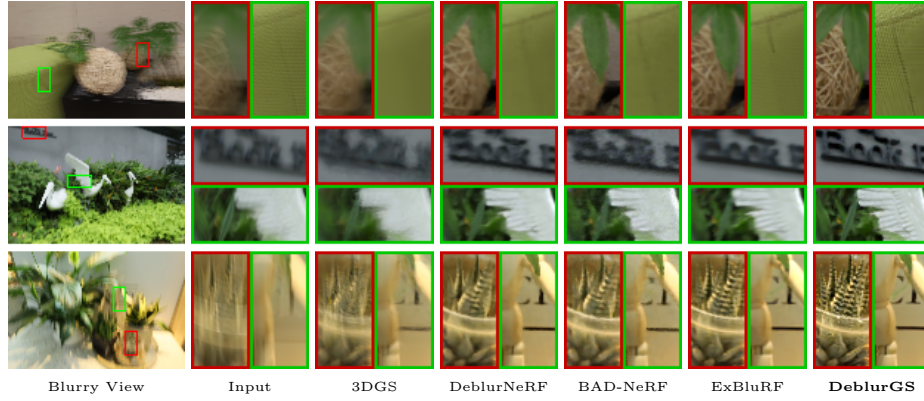
**Table 1:** Quantitative comparison of novel view synthesis.

Method	Real Motion Blur [27]			Synthetic Extreme Blur [24]		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3DGS [20]	21.69	0.684	0.281	20.39	0.591	0.454
DeblurNeRF [27]	25.49	0.763	0.182	23.98	0.687	0.301
BAD-NeRF [42]	22.59	0.633	0.252	21.94	0.590	0.349
ExBluRF [24]	25.93	0.775	0.198	27.81	0.823	0.227
<b>Ours</b>	<b>26.28</b>	<b>0.862</b>	<b>0.086</b>	<b>30.23</b>	<b>0.890</b>	<b>0.078</b>
Method	ExBlur-CP [24]			ExBlur-NP [24]		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
3DGS [20]	23.79	0.643	0.538	22.34	0.611	0.532
DeblurNeRF [27]	28.87	0.709	0.402	24.59	0.656	0.425
BAD-NeRF [42]	27.15	0.647	0.453	25.31	0.645	0.377
ExBluRF [24]	30.17	0.757	0.284	24.37	0.659	0.472
<b>Ours</b>	<b>31.75</b>	<b>0.806</b>	<b>0.172</b>	<b>30.11</b>	<b>0.768</b>	<b>0.192</b>

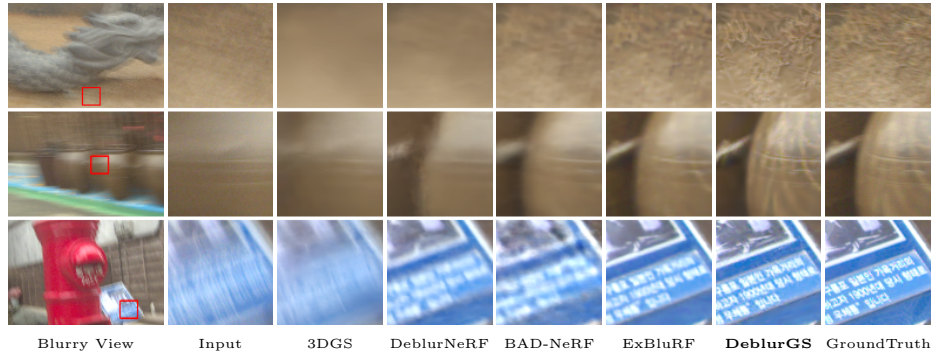
contains 20 to 40 blurry training images with hand-shaking camera motion blur, as well as 4 to 6 sharp testing images. For the Real-motion-blur dataset, the camera poses are obtained by COLMAP [37] with a set of training and testing images. For the ExBlur dataset, the amount of blur is more severe than the Real-motion-Blur dataset. Each blurry image in the ExBlur dataset has its sharp pair obtained from beam-splitter [36]. The camera poses of the blurry observation are calculated from the sharp pair in the ExBlur dataset. For the Synthetic Extreme Blur dataset also contains more severe motion blur than the Real-motion-blur dataset, and the camera pose is known by generating from Blender [11].

To better evaluate sharp 3D reconstruction capabilities in realistic scenarios, we divide the ExBlur dataset into two experimental settings based on the accuracy of initial poses and points cloud. The first setup, referred to as ExBlur-CleanPose or *ExBlur-CP*, is unchanged from the dataset’s public release, which is comprised of accurate camera pose and point clouds from COLMAP with sharp image pairs. On the other hand, ExBlur-NoisyPose or *ExBlur-NP* setting exclusively utilizes only blurry images for COLMAP processing, notably excluding sharp test images as well. To obtain the camera poses and point clouds from the inaccuracies induced by the blur, we loosen pre-defined COLMAP hyperparameters such as the re-projection error threshold and the number of inlier feature points. We detail the modifications made for COLMAP with blurry images in the supplementary materials. Given the significant blur that affects all images, not every image is guaranteed to be registered successfully in the process of Structure-from-Motion. However, we proceed our experiments with a minimum of 16 images for registration, which is at least 60% of the total images.

**Implementation Details** Our DeblurGS pipeline is built upon the official implementation of 3DGS [20]. We set the number of sub-frames for blurry rendering synthesis to  $N = 21$ , the same as ExBluRF [24] for a fair comparison. Addition-

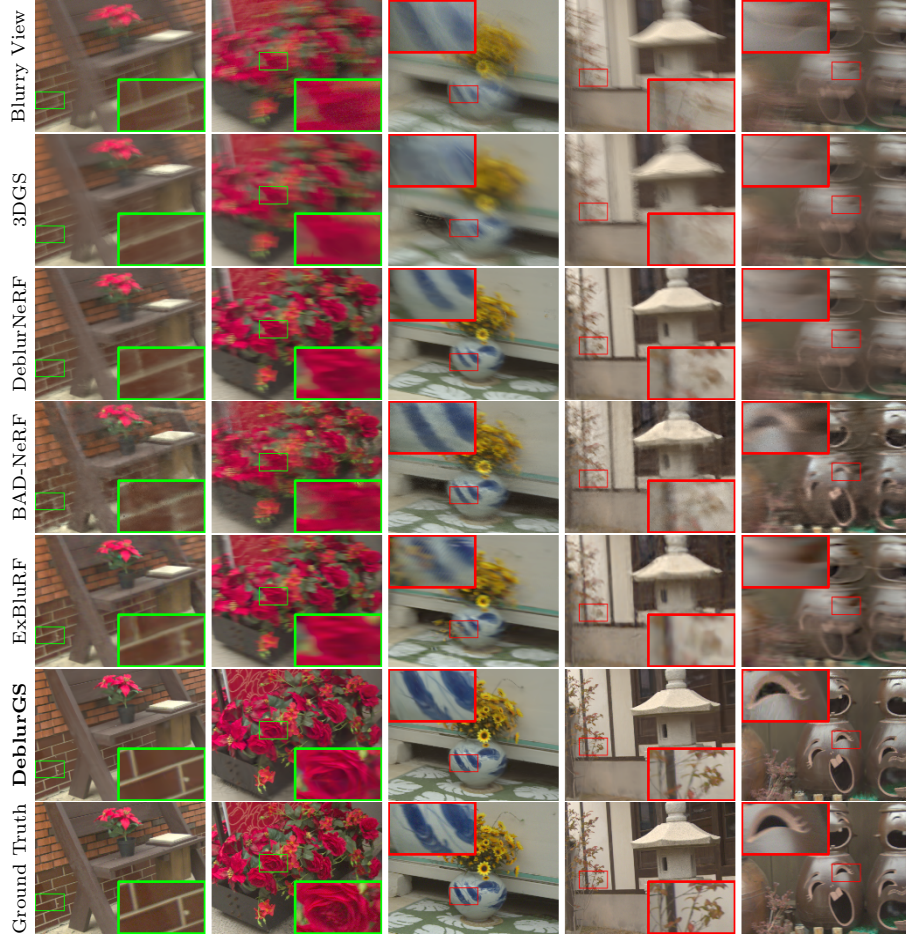


**Fig. 4: Qualitative Comparison of Deblurring on Real Motion Blur [27].** Note that blurry views in Real Motion Blur do not have their ground-truth pairs.



**Fig. 5: Qualitative Comparison of Deblurring on ExBlur-CP [24].** The camera poses are initialized from COLMAP [37] with sharp pairs corresponding to each blurry observation.

ally, we use the 9-th order Bézier curve in our camera trajectory estimation. We use Adam optimizer [21] with default configuration, and all the learning rate schema regarding the 3DGS model is unchanged from its built-in setting. For the camera motion parameters, we schedule the learning rate of  $1 \times 10^{-2}$  for the translation part of the camera trajectory,  $1 \times 10^{-3}$  for the rotation, and  $3 \times 10^{-3}$  for the alignment parameter  $\nu$ , all of which decay exponentially to be halved for every 15k iterations. Each scene is optimized for 150k iterations using a single NVIDIA GeForce RTX 4090 GPU.



**Fig. 6: Qualitative comparison of sharp 3D rendering on ExBlur-NP [24].** The camera poses are initialized from COLMAP [37] with blurry observations only.

## 4.2 Results

We compare the novel view synthesis performance to vanilla 3D Gaussian Splatting [20], DeblurNeRF [27], BAD-NeRF [42], and ExBluRF [24]. DeblurNeRF approximates the blur kernel on the 2D plane and optimizes with Neural Radiance Fields [29]. BAD-NeRF and ExBluRF jointly estimate the 6-DOF camera trajectory as a form of a linear [42] and Bèzier curve [24] in the Lie algebra ( $\mathfrak{se}(3)$ ) space, respectively.

**Evaluation Setup** The performance of novel view synthesis is evaluated by measuring PSNR, SSIM [17], and LPIPS [51] between rendering and sharp test

**Table 2: Ablation study.** We describe the ablation studies on each element of proposed method on "Camellia" and "Stone Lantern" scenes in ExBlur dataset with noisy camera pose setup. We ablate the effectiveness of the temporal smoothness loss  $\mathcal{L}_{\text{smooth}}$ , Gaussian densification annealing strategy, sub-frame alignment parameters  $\nu$ , and the gamma correction  $\gamma(\cdot)$ .

Method				Camellia			Stone Lantern		
$\mathcal{L}_{\text{smooth}}$	Anneal	$\theta$ learn	$\nu$ $\gamma(\cdot)$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
$\times$	$\times$	$\times$	$\times$	26.78	0.686	0.215	24.92	0.688	0.277
$\times$	$\checkmark$	$\checkmark$	$\checkmark$	27.22	0.653	0.252	26.20	0.716	0.275
$\checkmark$	$\times$	$\checkmark$	$\checkmark$	28.12	0.680	0.243	24.35	0.699	0.269
$\checkmark$	$\checkmark$	$\times$	$\checkmark$	28.12	0.675	0.253	28.20	0.772	0.234
$\checkmark$	$\checkmark$	$\checkmark$	$\times$	27.35	0.699	0.225	25.48	0.729	0.256
$\checkmark$	$\checkmark$	$\checkmark$	$\checkmark$	<b>28.97</b>	<b>0.711</b>	<b>0.189</b>	<b>29.07</b>	<b>0.779</b>	<b>0.194</b>

images. For all settings, since test poses are misaligned due to the joint optimization of 6-DOF camera motion, we perform the post-optimization process to align the test image poses, following previous 3D reconstruction works on pose-free settings [4, 16, 18, 25, 46] and blurry-image settings [24, 42]. In the ExBlur-NP setup, we fix the training camera poses and generate sharp renderings from these viewpoints to conduct COLMAP’s image registration and acquire poses for test images. After the test image registration, we align the test image poses by the post-optimization process similar to the evaluation process for the other setups. Note that the scene parameters such as NeRF, voxels and 3DGS remain frozen during post-optimization, and the purpose of the test-time optimization is strictly limited to factoring out the pose misalignment of test view images.

**Evaluation on Deblurring** As shown in Tab. 1, quantitative results show that our DeblurGS outperforms previous methods across all metrics. Comparing with the results from vanilla 3DGS, we can evident that previous NeRF-based deblurring methods are capable of restoring sharp images under conditions of moderate blur or when camera poses computed from sharp images. While deblurring performance differences among methods on the Real-motion-blur datasets are minimal, this gap significantly increases with the synthetic and ExBlur-CP datasets, which feature more severe motion blur. Notably, the proposed DeblurGS outperforms ExBluRF by a large margin, even though ExBluRF is explicitly designed to handle extreme motion blur. The qualitative results in Fig. 4 and Fig. 5 visually support the performance gain presented in Tab. 1. Interestingly, we have found that DeblurGS can render images even sharper than their ground-truth images in certain scenes of ExBlur dataset.

However, the results from accurate initial poses represent an ideal scenario, because the accurate camera poses cannot be obtained by SfM from blurry images. We demonstrate that the proposed DeblurGS is the only deblurring method capable of being applied in real-world scenarios, as evidenced by the difference in

deblurring performance between ExBlur-NP and ExBlur-CP setups. BAD-NeRF incorporates a bundle-adjustment mechanism in their optimization, so the performance drops are relatively smaller compared to other NeRF-based methods. However, Fig. 6 shows that only our method can restore scenes that are close to the ground truth sharp images. DeblurGS maintains its plausible deblurring performance even with the inaccurate pose initialization and performs almost equally to ExBluRF with the clean pose initialization in terms of PSNR, while exhibiting better performance in SSIM and LPIPS.

### 4.3 Ablation

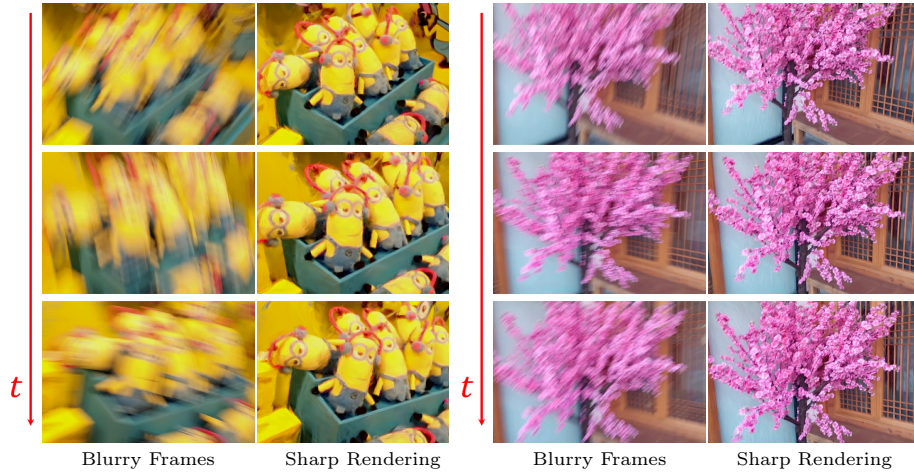
We present an ablation study on two scenes, Camellia and Stone Lantern, in the Exblur-NP dataset to verify and analyze the effectiveness of each component of the proposed approach. Tab. 2 indicates that each component plays a crucial role in sharp scene reconstruction and the highest performance is achieved when all elements are combined.

**Temporal Smoothness Loss** We find that ignoring the smoothness between adjacent sub-frame renderings leads to rendering with jittering artifacts, and degrades the metric as shown in Tab. 2. Without the smooth transitions between neighboring sub-frame renderings, we observe that the accumulation of unrealistic renderings with heavy artifacts reconstructs the blurry observation perfectly, due to the ill-posedness of the deblurring problem.

**Gaussian Densification Annealing** We observe that maintaining a constant threshold for Gaussian densification during the observation process generates Gaussians at inaccurate locations before camera motion converges. Subsequently, the Gaussians in erroneous positions attempt to fit the observation and cause floating artifacts, which deteriorate the quality of reconstruction as detailed in Tab. 2.

**Alignment Parameter** We detect a slight drop in effectiveness as recorded in Tab. 2 when employing uniformly sampled poses from an estimated trajectory to synthesize blur instead of sampling sub-frame poses from learnable alignment parameters. Due to the lack of a complete formulation of physical motion, convergence to latent camera motion is challenging.

**Gamma Correction** Omitting a gamma correction function on synthesized blur overlooks a process occurring in real cameras, resulting in areas naturally dark appearing unexpectedly brighter. Consequently, as illustrated in Tab. 2, elevating the scene’s brightness leads to a substantial decline in performance.



**Fig. 7: Sharp scene reconstruction from field-captured video.** We record the real-world scene as a video with smartphone, swiftly to capture the whole environment within 6 seconds. Our method successfully reconstructs sharp scene with a real-world blurry video.

#### 4.4 Real-World Video Analysis

To evaluate the practical applicability of our DeblurGS method, we conduct experiments using videos captured with a fast-moving smartphone. Using a Galaxy Ultra 22 device, we capture the surrounding environment at 30 fps for 4 to 6 seconds. We deploy our framework by taking each frame of the video as input and rendering a sharp view after the optimization process. As exhibited in Fig. 7, our method reconstructs sharp scenes even starting with field-captured video with blurry frames. DeblurGS effectively reconstructs sharp details in real-world environments since our method robustly operates on erroneous COLMAP [37] initialization, showcasing the effectiveness for practical applications.

## 5 Conclusion

In this paper, we present DeblurGS, a method for reconstructing a sharp 3D scene from a collection of motion blurred images. We simulate the camera motion to synthesize a blurry view, and optimize 3D Gaussian Splatting by minimizing distance between given blurry observation and generated blur. With our Gaussian Densification Annealing strategy, the camera motion converges to the latent camera movement, even initialized from noisy camera poses, the natural outcome of SfM from blurry observations. The capability to optimize from imprecise poses underscores the practicality of our framework, showcasing the successful deblurring of videos captured with smartphones. DeblurGS outperforms all existing methods in the sharp 3D scene reconstruction task, achieving state-of-the-art in both experimental and practical settings.

## 6 Acknowledgement

This work was supported in part by the SNU-NAVER Hyperscale AI Center, IITP grants [No. 2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University), No.2021-0-02068, and No.2023-0-00156], the NRF grant [No.2021M3A9E4080782] funded by the Korean government (MSIT).

## References

- Adamkiewicz, M., Chen, T., Caccavale, A., Gardner, R., Culbertson, P., Bohg, J., Schwager, M.: Vision-only robot navigation in a neural radiance world. *IEEE Robotics and Automation Letters* **7**(2), 4606–4613 (2022)
- 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: *CVPR* (2021)
- Barron, J.T., Mildenhall, B., Verbin, D., Srinivasan, P.P., Hedman, P.: Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In: *CVPR* (2022)
- Bian, W., Wang, Z., Li, K., Bian, J., Prisacariu, V.A.: Nope-nerf: Optimising neural radiance field with no pose prior. In: *CVPR* (2023)
- Chakrabarti, A.: A neural approach to blind motion deblurring. In: *ECCV* (2016)
- Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., et al.: Efficient geometry-aware 3d generative adversarial networks. In: *CVPR* (2022)
- Chen, A., Xu, Z., Geiger, A., Yu, J., Su, H.: Tensorf: Tensorial radiance fields. In: *ECCV* (2022)
- Chen, L., Chu, X., Zhang, X., Sun, J.: Simple baselines for image restoration. In: *ECCV* (2022)
- Chen, Z., Funkhouser, T., Hedman, P., Tagliasacchi, A.: Mobilenerf: Exploiting the polygon rasterization pipeline for efficient neural field rendering on mobile architectures. In: *CVPR* (2023)
- Cho, S.J., Ji, S.W., Hong, J.P., Jung, S.W., Ko, S.J.: Rethinking coarse-to-fine approach in single image deblurring. In: *ICCV* (2021)
- Community, B.O.: Blender - a 3D modelling and rendering package. Blender Foundation, Stichting Blender Foundation, Amsterdam (2018), <http://www.blender.org>
- Conzales, R.C., Woods, R.E.: Multiple view geometry in computer vision. Prentice Hall (2008)
- Deng, K., Liu, A., Zhu, J.Y., Ramanan, D.: Depth-supervised NeRF: Fewer views and faster training for free. In: *CVPR* (2022)
- Deng, N., He, Z., Ye, J., Duinkharjav, B., Chakravarthula, P., Yang, X., Sun, Q.: Fov-nerf: Foveated neural radiance fields for virtual reality. *IEEE Transactions on Visualization and Computer Graphics* (2022)
- Fridovich-Keil, S., Yu, A., Tancik, M., Chen, Q., Recht, B., Kanazawa, A.: Plenoxels: Radiance fields without neural networks. In: *CVPR* (2022)
- Fu, Y., Liu, S., Kulkarni, A., Kautz, J., Efros, A.A., Wang, X.: Colmap-free 3d gaussian splatting (2023)
- Hore, A., Ziou, D.: Image quality metrics: Psnr vs. ssim. In: *ICPR* (2010)
- Jeong, Y., Ahn, S., Choy, C., Anandkumar, A., Cho, M., Park, J.: Self-calibrating neural radiance fields. In: *ICCV* (2021)

19. Jiang, Y., Hedman, P., Mildenhall, B., Xu, D., Barron, J.T., Wang, Z., Xue, T.: Alignerf: High-fidelity neural radiance fields via alignment-aware training. CVPR (2023)
20. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3d gaussian splatting for real-time radiance field rendering. TOG (2023)
21. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)
22. Kupyn, O., Martyniuk, T., Wu, J., Wang, Z.: Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In: ICCV (2019)
23. Lee, D., Lee, M., Shin, C., Lee, S.: Dp-nerf: Deblurred neural radiance field with physical scene priors. In: CVPR (2023)
24. Lee, D., Oh, J., Rim, J., Cho, S., Lee, K.M.: Exblurf: Efficient radiance fields for extreme motion blurred images. In: ICCV (2023)
25. Lin, C.H., Ma, W.C., Torralba, A., Lucey, S.: Barf: Bundle-adjusting neural radiance fields. In: ICCV (2021)
26. Lindell, D.B., Martel, J.N., Wetzstein, G.: Autoint: Automatic integration for fast neural volume rendering. In: CVPR (2021)
27. Ma, L., Li, X., Liao, J., Zhang, Q., Wang, X., Wang, J., Sander, P.V.: Deblur-nerf: Neural radiance fields from blurry images. In: CVPR (2022)
28. Max, N.: Optical models for direct volume rendering. IEEE Transactions on Visualization and Computer Graphics pp. 99–108 (1995). <https://doi.org/10.1109/2945.468400>
29. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. In: ECCV (2020)
30. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multiresolution hash encoding. TOG (2022)
31. Nah, S., Baik, S., Hong, S., Moon, G., Son, S., Timofte, R., Lee, K.M.: Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In: CVPR Workshops (2019)
32. Nah, S., Kim, T.H., Lee, K.M.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In: CVPR (2017)
33. Park, H., Lee, K.M.: Joint estimation of camera pose, depth, deblurring, and super-resolution from a blurred image sequence. In: ICCV (2017)
34. Reiser, C., Peng, S., Liao, Y., Geiger, A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In: ICCV (2021)
35. Rim, J., Kim, G., Kim, J., Lee, J., Lee, S., Cho, S.: Realistic blur synthesis for learning image deblurring. In: ECCV (2022)
36. Rim, J., Lee, H., Won, J., Cho, S.: Real-world blur dataset for learning and benchmarking deblurring algorithms. In: Proceedings of the European Conference on Computer Vision (ECCV) (2020)
37. Schonberger, J.L., Frahm, J.M.: Structure-from-motion revisited. In: CVPR (2016)
38. Sun, C., Sun, M., Chen, H.T.: Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. CVPR (2022)
39. Sun, J., Cao, W., Xu, Z., Ponce, J.: Learning a convolutional neural network for non-uniform motion blur removal. In: CVPR (2015)
40. Tao, X., Gao, H., Shen, X., Wang, J., Jia, J.: Scale-recurrent network for deep image deblurring. In: CVPR (2018)
41. Wang, H., Ren, J., Huang, Z., Olszewski, K., Chai, M., Fu, Y., Tulyakov, S.: R2l: Distilling neural radiance field to neural light field for efficient novel view synthesis. In: ECCV (2022)

42. Wang, P., Zhao, L., Ma, R., Liu, P.: BAD-NeRF: Bundle Adjusted Deblur Neural Radiance Fields. In: CVPR (2023)
43. Wieschollek, P., Hirsch, M., Scholkopf, B., Lensch, H.: Learning blind motion deblurring. In: ICCV (2017)
44. Wu, L., Lee, J.Y., Bhattad, A., Wang, Y., Forsyth, D.: Diver: Real-time and accurate neural radiance fields with deterministic integration for volume rendering (2022)
45. Xiangli, Y., Xu, L., Pan, X., Zhao, N., Rao, A., Theobalt, C., Dai, B., Lin, D.: Bungeenerf: Progressive neural radiance field for extreme multi-scale scene rendering. In: ECCV (2022)
46. Yen-Chen, L., Florence, P., Barron, J.T., Rodriguez, A., Isola, P., Lin, T.Y.: iNeRF: Inverting neural radiance fields for pose estimation. In: IROS (2021)
47. Yu, A., Li, R., Tancik, M., Li, H., Ng, R., Kanazawa, A.: Plenotrees for real-time rendering of neural radiance fields. In: ICCV (2021)
48. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: CVPR (2022)
49. Zhang, K., Riegler, G., Snavely, N., Koltun, V.: NeRF++: Analyzing and improving neural radiance fields. arXiv:2010.07492 (2020)
50. Zhang, K., Luo, W., Zhong, Y., Ma, L., Stenger, B., Liu, W., Li, H.: Deblurring by realistic blurring. In: CVPR (2020)
51. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: CVPR (2018)
52. Zwicker, M., Pfister, H., Baar, J.V., Gross, M.: Surface splatting. Proceedings of the 28th annual conference on Computer graphics and interactive techniques (2001)