

# 000 RESPLAT: DEGRADATION-AGNOSTIC FEED- 001 FORWARD GAUSSIAN SPLATTING VIA SELF-GUIDED 002 RESIDUAL DIFFUSION 003

006 **Anonymous authors**

007 Paper under double-blind review

## 011 ABSTRACT

013 Recent advances in novel view synthesis (NVS) have predominantly focused on  
014 ideal, clear input settings, limiting their applicability in real-world environments  
015 with common degradations such as blur, low-light, haze, rain, and snow. While  
016 some approaches address NVS under specific degradation types, they are often tai-  
017 lored to narrow cases, lacking the generalizability needed for broader scenarios.  
018 To address this issue, we propose Restoration-based feed-forward Gaussian Splat-  
019 tting, named *ReSplat*, a novel framework capable of handling degraded multi-view  
020 inputs. Our model jointly estimates restored images and gaussians to represent the  
021 clear scene for NVS. We enable multi-view consistent universal image restoration  
022 by utilizing the 3d gaussians generated during the diffusion sampling process as  
023 self-guidance. This results in sharper and more reliable novel views. Notably, our  
024 framework adapts to various degradations without prior knowledge of their spe-  
025 cific types. Extensive experiments demonstrate that ReSplat significantly outper-  
026 forms existing methods across challenging conditions, including blur, low-light,  
027 haze, rain, and snow, delivering superior quality and robust NVS performance.

## 029 1 INTRODUCTION

031 Novel View Synthesis (NVS) is a task aimed at generating novel views of a scene from a known  
032 set of views. NVS strives to accurately estimate the geometry and appearance of a scene, enabling  
033 the rendering of realistic images from unseen perspectives. In recent years, Neural Radiance Fields  
034 (NeRF) Mildenhall et al. (2021) have revolutionized NVS by utilizing neural networks to repre-  
035 sent scenes in a continuous volumetric format, producing highly realistic results. However, NeRF’s  
036 slow rendering speed has limited its practicality, especially in real-time applications. Solutions like  
037 InstantNGP Müller et al. (2022) and TensoRF Chen et al. (2022) have addressed these speed limita-  
038 tions, and Gaussian Splatting Kerbl et al. (2023), introduced later, further accelerated the rendering  
039 process. By representing scenes with Gaussian ellipsoids instead of dense point samples like NeRF,  
040 Gaussian Splatting maintains competitive visual quality while enabling faster rendering.

041 Despite the impressive results of NeRF and Gaussian Splatting, generalizable approaches have be-  
042 come a major focus area. Generalizable NeRF aims to synthesize new views without retraining on  
043 each new scene, enhancing model flexibility across diverse datasets Wang et al. (2021); Yu et al.  
044 (2021b); Wang et al. (2022b); Suhail et al. (2022). Similarly, generalizable Gaussian Splatting ex-  
045 tends this concept, offering a faster and adaptable solution for unseen scenes Charatan et al. (2023);  
046 Chen et al. (2025); Liu et al. (2025); Ziwen et al. (2024). However, these methods have primarily  
047 been developed to work on clean multi-view images captured from controlled environments.

048 Against this backdrop, scene reconstruction using corrupted images has gained attention. Some  
049 studies Ma et al. (2022); Wang et al. (2022a); Yoon & Yoon (2023); Wang et al. (2023); Chen et al.  
050 (2023b) are designed to handle specific types of corruption. GAURA Gupta et al. (2024), on the  
051 other hand, leverages the capacity of feed-forward NVS models to be pre-trained on large multi-  
052 view datasets, proposing a generalizable NeRF model that operates under a variety of degradations.  
053 However, GAURA excludes the image restoration capabilities developed in the 2D domain, which  
limits its performance potential.

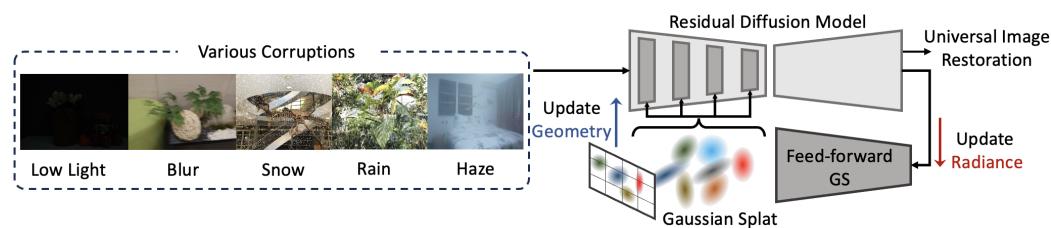


Figure 1: Proposed degradation-agnostic feed-forward Gaussian Splatting (GS) framework. Our framework achieves high-performance universal image restoration and novel view synthesis results through mutual information exchange between the universal image restoration model and the generalizable GS model.

Simply adopting an image restoration model does not fully address this limitation. Universal image restoration is a severely ill-posed problem, with countless possible solutions. This has led to various methods using denoising diffusion models—prominent examples of stochastic models—to address image restoration Fei et al. (2023); Özdenizci & Legenstein (2023). Research has shown that training to predict residual images enables effective image restoration Zhang et al. (2017; 2018); Zamir et al. (2021); Anwar & Barnes (2020), enhancing performance through diffusion-based residual learning.

In this paper, we propose a new generalizable gaussian splatting framework, **ReSplat**, aimed at degradation-agnostic novel view synthesis. At the core of our framework is a method that leverages the model priors of a diffusion-based unified image restoration network through Gaussian splatting. Unlike NeRF’s representation, Gaussian splatting uses a point-based representation that enables explicit scene geometry extraction during training. As shown in Fig. 1, generalizable Gaussian splatting models Charatan et al. (2023); Chen et al. (2025); Liu et al. (2025) inherently estimate Gaussian centroids (geometry) using multi-view stereo (MVS) and radiance (color) through multi-view image aggregation. In our framework, a diffusion model iteratively estimates Gaussian centroids, or 3D geometry, leveraging this information to achieve 3D-consistent image restoration.

Our framework specifically adapts a 3D cross-attention module to the residual diffusion model, enabling it to utilize the location information of Gaussian centroids. Here, Gaussian centroids are derived from the point clouds of restored images estimated in the previous time-step. Second, our model performs multi-view aligned pre-filtering when generating Gaussian ellipsoids. This process involves calculating a weight map that is applied to the image features used to generate the Gaussian ellipsoids, helping to achieve artifact-free novel view synthesis. Through these techniques, our model retains the advantages of a generalizable method that operates without a scene optimization process, working effectively even in sparse-view settings while remaining degradation-agnostic. This makes it a more practical NVS model, demonstrating superior NVS and image restoration performance in multiple degradation settings compared to other approaches.

In summary, our contributions are summarized as follows:

1. We propose *ReSplat*, a novel framework for multi-view image restoration using 3DGS.
2. We introduce a multi-view aligned denoising diffusion model for universal image restoration.
3. Our method outperforms other methods in novel view synthesis and image restoration tasks.

## 2 RELATED WORKS

### 2.1 GENERALIZABLE RADIANCE FIELDS

Generating realistic images has been a central research topic for many years. Neural scene representations, such as Neural Radiance Fields (NeRF) Mildenhall et al. (2021), have emerged as effective solutions for view synthesis, achieving remarkable results. Subsequent NeRF-based approaches have further improved rendering quality Roessle et al. (2022); Wei et al. (2021); Deng et al. (2022), as well as optimization and rendering speed Sun et al. (2022); Chen et al. (2022); Fridovich-Keil et al. (2022); Yu et al. (2021a); Müller et al. (2022). However, NeRF still requires optimization for each new scene to synthesize novel views. To address this, various studies have proposed generalizable NeRF models Yu et al. (2021b); Wang et al. (2021); Liu et al. (2022); Wang et al. (2022b); Suhail et al. (2022); Cao et al. (2022), enabling cross-scene generalization by learning a view in-

108 terpolation function from source images. In these generalizable NeRFs, a common technique is  
 109 to apply volume rendering for aggregating information from images, such as deep features, depth  
 110 maps, or cost volumes Wang et al. (2021); Liu et al. (2022); Johari et al. (2022); Chen et al. (2021);  
 111 Xu et al. (2023). GPNR Suhail et al. (2022) and GNT Wang et al. (2022b) utilize transformers to  
 112 aggregate features, enhancing information interaction along a ray to directly predict RGB values for  
 113 each pixel. PixelSplat Charatan et al. (2023) and MVSplat Chen et al. (2025) propose generalizable  
 114 volume rendering techniques that utilize scene parameterization with 3D Gaussian primitives Kerbl  
 115 et al. (2023). We note that existing generalizable radiance fields have predominantly been studied  
 116 on clean images, and we aim to address this limitation by developing a universal model utilizing a  
 117 residual diffusion model.

## 118 2.2 NOVEL VIEW SYNTHESIS WITH DEGRADATIONS

120 Some research has advanced novel view synthesis (NVS) by leveraging radiance fields with physics-  
 121 based multi-view geometry techniques, targeting cases where train-view images require enhancement.  
 122 NeRF-W Martin-Brualla et al. (2021) addresses variations in illumination and transient  
 123 occlusions by relaxing strict assumptions on consistency across inputs. Deblur-NeRF Ma et al.  
 124 (2022) introduces a spatially-varying blur kernel model to handle blurry inputs effectively. RawN-  
 125 eRF Mildenhall et al. (2022) facilitates high-dynamic range (HDR) view synthesis by training NeRF  
 126 on raw input data and generating raw-format outputs. Similarly, HDR-NeRF Huang et al. (2022)  
 127 supports exposure control and HDR image synthesis by learning two distinct implicit functions:  
 128 one for the radiance field and another for tone mapping. LLNeRF Wang et al. (2023) and Aleth-  
 129 NeRF Cui et al. (2024) conducted research on novel view synthesis under low-light conditions.  
 130 More recently, DiET-GS Lee & Lee (2025) and DiSR-NeRF Lee et al. (2024) leverage diffusion  
 131 priors to improve 3D representations from degraded inputs, but they are designed for specific cor-  
 132 ruption types such as motion blur or low resolution. HQGS Lin et al. and RobustGS Wu et al. (2025)  
 133 further study Gaussian Splatting under various degraded conditions and propose task-specific 3DGS  
 134 pipelines to boost robustness in these scenarios. Overall, these studies do not explicitly leverage  
 135 a pretrained universal image restoration model and remain tailored to specific degradation regimes  
 136 or 3D configurations. In contrast, we target a degradation-agnostic framework that actively uses a  
 137 pretrained universal restoration prior within a feed-forward 3DGS pipeline, so that a single model  
 138 can handle diverse and mixed degradations.

## 139 2.3 UNIVERSAL IMAGE RESTORATION

140 Developing a unified model capable of handling multiple degradations has become a growing area of  
 141 interest. AiRnet Li et al. (2022) introduces a module to align various distributions into a shared dis-  
 142 tribution using contrastive learning, though this approach can be challenging to train and may limit  
 143 performance. IDR Zhang et al. (2023) identifies that distinct degradation types can be separated  
 144 using singular value decomposition (SVD), allowing for clean image reconstruction through refor-  
 145 mulation of singular values and vectors. PromptIR Potlapalli et al. (2024) enhances performance by  
 146 employing a prompt block to capture degradation-specific features. multi-task DINO-based restora-  
 147 tion Lin et al. (2023) and mask-based blind restoration Qin et al. (2024) exploit strong visual priors  
 148 from foundation models. Adair Cui et al. (2025) and Perceive-IR Zhang et al. (2025) further improve  
 149 all-in-one restoration by adaptively modeling degradation-specific frequency cues and enhancing  
 150 degradation perception, respectively. Methods such as ProRes Ma et al. (2023) and DA-CLIP Luo  
 151 et al. (2023) leverage prompt learning to fully utilize the power of large-scale models. Daclip-IR Luo  
 152 et al. (2024) incorporates a CLIP-based encoder to identify degradation types, extracting semantic  
 153 information from distorted images to guide a diffusion model in generating high-quality outputs.  
 154 DiffUIR Zheng et al. (2024) introduces selective hourglass mapping to adapt residual denoising  
 155 diffusion models Liu et al. (2024) as a comprehensive image restoration approach.

## 156 3 METHODS

157 In this section, we provide an overview of the residual denoising diffusion model (RDDM), univer-  
 158 sal RDDM, and 3D Gaussian splatting as preliminaries. Next, we propose an overall framework for  
 159 degradation-agnostic feed-forward Gaussian splatting. Additionally, we detail two modules specifi-  
 160 cally designed to enhance NVS performance.

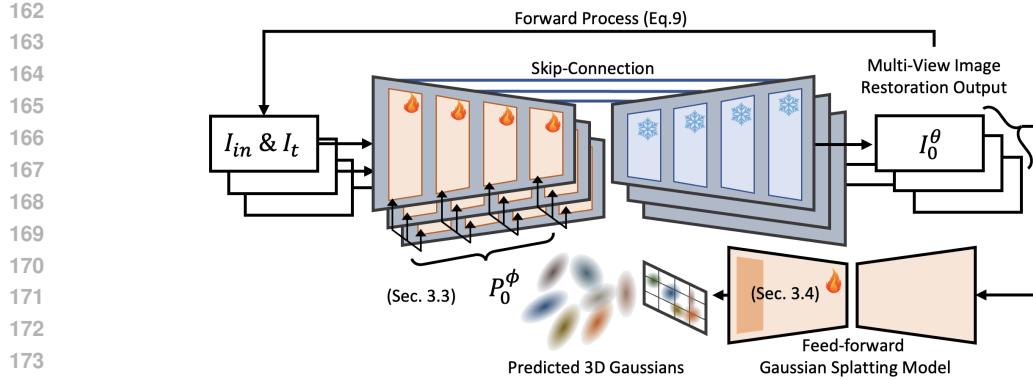


Figure 2: The overall framework for degradation-agnostic feed-forward gaussian splatting (GS). A diffusion-based image restoration model restores the original image by iteratively estimating the residual image. During this process, feed-forward GS is performed using the original image generated in the intermediate stages of diffusion sampling. By utilizing the Gaussian points information obtained in this process, the diffusion model receives multi-view information in the next diffusion step, enabling more accurate image restoration.

### 3.1 PRELIMINARIES

**3D Gaussian Splatting** 3D-GS Kerbl et al. (2023) models a scene using a collection of anisotropic 3D Gaussians, which retain the differential characteristics of volumetric representations while enabling efficient rendering through a tile-based rasterization approach. Beginning with points derived from Structure-from-Motion (SfM), each point serves as the position (mean)  $\mu$  of a 3D gaussian ellipsoids.

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

where  $x$  represents a specific point in the 3D scene, and  $\Sigma$  is the covariance matrix of the 3D Gaussian.  $\Sigma$  is constructed from a scaling matrix  $S$  and a rotation matrix  $R$  with the equation  $\Sigma = RSS^T R^T$ . For performing tile-based rasterization, the 3D Gaussians  $G(x)$  are projected onto the image plane as 2D Gaussians  $G'(x)$ . The rasterizer then sorts these 2D Gaussians and applies alpha blending:

$$C(x') = \sum_{i \in N} c_i \sigma_i \prod_{j=1}^{i-1} (1 - \sigma_j), \quad \sigma_i = \alpha_i G'_i(x') \quad (2)$$

$x'$  represents the queried pixel position, and  $N$  denotes the number of sorted 2D Gaussians associated with that pixel.

**Feed-forward 3D Gaussian Splatting** While vanilla 3DGS optimizes Gaussian parameters per scene, recent feed-forward 3DGS models Liu et al. (2025) predict them in a single forward pass from a few posed views. Given  $N$  input images  $\{I_{in}^i\}_{i=1}^N$  and their camera poses  $\{\Pi^i\}_{i=1}^N$ , the network  $\phi$  maps multi-view features to per-pixel Gaussian primitives:

$$\phi : \{(I_{in}^i, \Pi^i)\}_{i=1}^N \mapsto \{(\mu_j, \Sigma_j, \alpha_j, c_j)\}_{j=1}^{H \times W \times N}, \quad (3)$$

where  $(\mu_j, \Sigma_j, \alpha_j, c_j)$  denote the center, covariance, opacity, and color of candidate Gaussians. These predictions are then pruned and merged into the explicit Gaussian set  $P_\phi^0$ , which is rendered using the standard splatting formulation above. In addition, the feed-forward 3DGS predicts per-view aggregation weights  $W^i$  that are used to combine warped multi-view features at each novel-view pixel; in Sec. 3.4, we modulate these weights with our degradation-aware pre-filtering module.

**Residual Denoising Diffusion Model** RDDM Liu et al. (2024) uses a standard  $T$ -step diffusion model that includes both a forward and a reverse process. In the forward process, one-step noising is formulated as a Markov chain:

$$q(I_t | I_{t-1}, I_{res}) = \mathcal{N}(I_t; I_{t-1} + \alpha_t I_{res}, \beta_t^2 \mathbf{I}) \quad (4)$$

where  $\alpha_t$  and  $\beta_t$  are the noise coefficients for  $I_{res}$  and gaussian noise.  $I_t$  is the result at timestep  $t$ , and  $I_{res}$  represents the residual between the degraded image  $I_{in}$  and the clean image  $I_0$ , with  $I_{res} =$

216

**Algorithm 1** Training

217

**Input:** Clean image, Degraded image:  $I_0, I_{in}$ ;  
 GT novel view image :  $I_{nv}$ ;  
 GT residual map:  $I_{res} = I_{in} - I_0$ ;  
 Image Restoration Model:  $\theta(*)$ ;  
 Feed-forward GS Model:  $\phi(*)$ ;

1: **repeat**  
 2:    $I_0 \sim q(I_0)$ ;  
 3:    $P_0^\phi = \phi(I_0, I_{in})$ ;  
 4:    $t \sim \text{Uniform}(1, \dots, T)$ ;  
 5:    $\epsilon \sim \mathcal{N}(0, I)$ ;  
 6:    $I_t = I_0 + \bar{\alpha}_t I_{res} + \bar{\beta}_t \epsilon - \bar{\delta}_t I_{in}$ ;  
 7:   Take the gradient descent step on  

$$\nabla_\theta \|I_{res} - I_{res}^\theta(P_0^\phi, I_t, I_{in}, t)\|_1 +$$

$$\nabla_\phi \|I_{nv} - I_{ren}^\phi(I_{in} - I_{res}^\theta, I_{in})\|_1$$

8: **until** converged

221

222

223

224

225

226

227

228

229

230

231

232

233

234

235

236

$I_{in} - I_0$ . In the reverse process, RDDM approximates the true generative distribution  $p_\theta(I_{t-1}|I_t)$  by using  $q(I_{t-1}|I_t, I_0, I_{res})$ , which is also formulated as a Markov chain when deterministic implicit sampling using DDIM Song et al. (2020):

237

238

$$p_\theta(I_{t-1}|I_t) = \mathcal{N}(I_{t-1}; I_0^\theta + \bar{\alpha}_{t-1} I_{res}^\theta + \bar{\beta}_{t-1} \epsilon^\theta, 0 \cdot \mathbf{I}) \quad (5)$$

239

In summary, the relations between  $I_t$  and  $I_{t-1}$  in both the forward and reverse processes is as follows:

240

241

242

243

$$I_t = I_{t-1} + \alpha_t I_{res} + \beta_t \epsilon_{t-1} \quad (6)$$

$$I_{t-1} = I_t - (\bar{\alpha}_t - \bar{\alpha}_{t-1}) I_{res}^\theta - (\bar{\beta}_t - \bar{\beta}_{t-1}) \epsilon^\theta \quad (7)$$

244

245

246

**Universal Residual Denoising Diffusion Model DiffUIR** Zheng et al. (2024) utilizes the conditioning mechanism from RDDM and incorporate a shared distribution term (SDT) within the diffusion algorithm for universal image restoration. They adjust the forward process as follows:

247

$$I_t = I_{t-1} + \alpha_t I_{res} + \beta_t \epsilon_{t-1} - \delta_t I_{in} \quad (8)$$

248

249

where  $\delta_t I_{in}$  represents the SDT, and  $\delta$  is the shared distribution coefficient. The reverse process is also as follows:

250

251

$$I_{t-1} = I_t - \alpha_t I_{res}^\theta + \delta_t I_{in} - (\beta_t^2 / \bar{\beta}_t) \epsilon^\theta \quad (9)$$

252

253

Finally, in the deterministic implicit sampling process,  $I_{t-1}$  and the pseudo clean image  $I_0^\theta$  can be derived using the following equation:

254

255

256

257

$$I_{t-1} = I_0^\theta + \bar{\alpha}_{t-1} I_{res}^\theta - \bar{\delta}_{t-1} I_{in} \quad \text{s.t. } I_0^\theta = I_{in} - I_{res}^\theta \quad (10)$$

### 3.2 OVERALL FRAMEWORK

258

259

260

261

262

263

264

265

266

We aim to develop a novel view synthesis model that can be performed under arbitrary degradation. Most of the existing novel view synthesis studies have been conducted on clean images without corruption, and even in the case of studies on situations with corruption, models specialized for specific degradation types are being developed Ma et al. (2022); Wang et al. (2022a); Yoon & Yoon (2023); Wang et al. (2023); Chen et al. (2023b). These studies solve the problem by simultaneously optimizing scene optimization and physical characteristics by implementing physical characteristics that cause specific degradation as a rendering process. Therefore, there is a need for new degradation-agnostic novel view synthesis (NVS) studies. To address this, we propose a new framework, *ReSplat*, the NVS model that leverages the diffusion prior studied in the field of 2D image restoration.

267

268

269

**Training stage** As shown in Fig. 2, we combine the feed-forward gaussian splatting (GS) model and the universal image restoration (UIR) model. The GS and UIR models support complementary roles. Unlike NeRF, feed-forward GS inevitably performs multi-view stereo (MVS) within the model because it needs to explicitly extract point clouds. This enables acquisition of 3D scene geometry

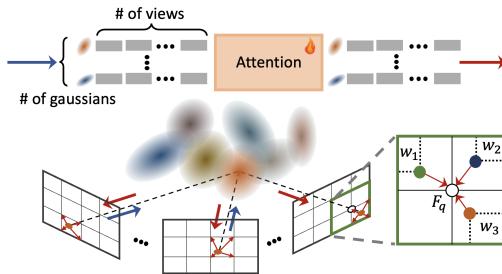


Figure 3: **GS-guided multi-view alignment.** Module embedded in the residual diffusion model that shares info between adjacent views using Gaussian centers.

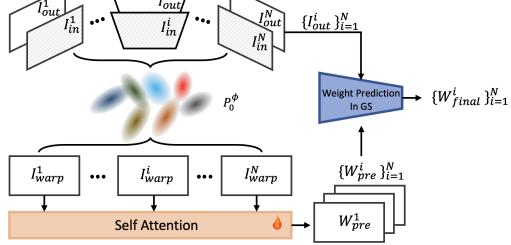


Figure 4: **Pre-filtering with warped features.** Warped inputs are self-attended to form pre-filtering weights for feature aggregation.

information and helps the UIR model find corresponding points for adjacent multi-view images. Meanwhile, the UIR model performs degradation-agnostic image restoration to help feed-forward GS perform NVS using images with corruption removed. In addition, we adopt DiffUIR Zheng et al. (2024), a diffusion model-based UIR model, to perform iterative scene geometry extraction and iterative image refinement so that the UIR results can be gradually refined. The training process can be found in Algorithm 1. The first term  $\|I_{\text{res}} - I_{\text{res}}^\theta(P_0^\phi, I_t, I_{\text{in}}, t)\|_1$  corresponds to the universal restoration loss  $\mathcal{L}_{\text{UIR}}$ , which supervises the residual prediction of the UIR model  $\theta$ . The second term  $\|I_{\text{nv}} - I_{\text{ren}}^\theta(I_{\text{in}} - I_{\text{res}}, I_{\text{in}})\|_1$  defines the novel-view rendering loss  $\mathcal{L}_{\text{NV}}$ , which supervises the feed-forward GS model  $\phi$  using the ground-truth clean novel view  $I_{\text{nv}}$ .

**Sampling stage** The specific sampling process of ReSplat is shown in Algorithm 2. Given  $N$  multi-view input images  $\{I_{\text{in}}^i\}_{i=1}^N$ , the  $t^{\text{th}}$  noise images  $\{I_t^i\}_{i=1}^N$  are generated according to the forward process of the DiffUIR. We generate predicted clean images  $\{I_0^i\}_{i=1}^N$  from predicted  $\{I_{\text{res}}^i\}_{i=1}^N$  by the UIR model. The predicted clean images are used to generate explicit point cloud  $P_0^\phi$  by the MVS module of the feed-forward GS model. Meanwhile, we generate  $\{I_{t-1}^i\}_{i=1}^N$  to perform the next diffusion step. After that, we perform a 3d aligned diffusion reverse process using  $P_0^\phi$  (sec. 3.3). Through this, we regenerate the refined  $\{I_0^i\}_{i=1}^N$  and  $P_0^\phi$ . We repeat the process and perform the feed-forward GS overall process using the finally generated  $\{I_0^i\}_{i=1}^N$ . In this process, we perform a feature pre-filtering process conditioned on the original corrupted images  $\{I_{\text{in}}^i\}_{i=1}^N$  to remove points where artifacts exist before the multi-view feature aggregation process, thereby generating a more robust GS output (sec. 3.4). Through this, we can obtain a rendered output for the novel view point.

### 3.3 GS GUIDED MULTI-VIEW ALIGNMENT

Since the original UIR model is designed for a single image, it is necessary to design a module for enabling multi-view image interaction. As shown in Fig. 3, we adapt a module that performs feature attention in space to UIR by utilizing  $P_0^\phi$ , a pseudo geometry generated during the sampling process. Multi-view features are projected toward each gaussian center in  $P_0^\phi$ . When there are  $N$  multi-view feature vectors  $\{f_i^j\}_{j=1}^N$  for the  $i^{\text{th}}$  center point  $p_i$ , we perform self-attention between the corresponding vectors. This process is repeated in the encoder of the diffusion model and helps ensure the 3D consistency of multi-view images. The processed feature vector  $f_{i,\text{rep}}^j$  is reprojected to the original pixel coordinates. However, since the reprojected point is located in continuous coordinates, not discrete coordinates, it is necessary to propagate to the surrounding discrete coordinates. Therefore, we perform a weighted sum by applying 2D interpolation weights  $\{w_i\}$  to all reprojected points existing between adjacent pixels. Each weight is determined by the area of the opposite region, ensuring that features closer to the query point have a higher influence. Therefore, when there is a discrete point  $q$ , the multi-view feature  $F_q$  that  $q$  obtains is as follows.

$$F_q = \sum_i w_i f_{i,\text{rep}}^j \quad \text{where } i \in Q \quad (11)$$

and  $Q$  is the set of the index of all points that exist within the smallest rectangle surrounding the point  $q$ .

324 Table 1: Novel View Synthesis (NV) and Image Restoration (IR) results of five corruption types on  
325 LLFF degradation dataset with **three multi-view** inputs. The **best** scores and **second best** scores  
326 are highlighted with their respective colors.

327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343	Method	Operation	Year	Corruption Type	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343			327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343		
					327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343	327 328 329 330 331 332 333 334 335 336 337 338 339 340 341 342 343			
327 328 329 330 331 332	AiRnet	IR $\rightarrow$ NV	2022	Motion Blur	20.11	0.6896	0.4250	21.99	0.7543	0.3769
					20.04	0.6872	0.4208	22.14	0.7526	0.3668
					21.28	0.7198	0.4343	-	-	-
					22.75	0.7824	0.3209	26.34	0.8640	0.2951
					23.15	0.8049	0.3151	27.14	0.8850	0.2713
					20.22	0.6852	0.3026	21.57	0.8184	0.2159
333 334 335 336 337 338	PromptIR	IR $\rightarrow$ NV	2023	Snow	20.54	0.7067	0.2737	23.21	0.8578	0.1912
					20.48	0.7044	0.3195	-	-	-
					24.24	0.8549	0.1826	31.20	0.9627	0.1019
					24.46	0.8614	0.1677	32.07	0.9685	0.0886
					9.159	0.3841	0.3892	8.871	0.4155	0.2949
					9.784	0.4651	0.3508	9.585	0.5251	0.2280
339 340 341 342 343	GAURA	Only NV	2024	Haze	17.22	0.7205	0.4516	-	-	-
					21.56	0.8392	0.1857	25.57	0.9612	0.0701
					21.99	0.8471	0.1750	26.45	0.9680	0.0619
					9.526	0.1364	0.6041	6.388	0.0859	0.7654
					6.367	0.0805	0.6240	6.298	0.0805	0.5789
					15.28	0.6627	0.5177	-	-	-
341 342 343	DiffUIR	IR $\rightarrow$ NV	2024	Low-light	18.87	0.8241	0.2429	21.88	0.9374	0.1647
					19.76	0.8276	0.2433	22.82	0.9452	0.1605
					20.49	0.6988	0.3416	23.09	0.8008	0.3018
					20.71	0.7175	0.2992	24.78	0.8555	0.2499
					21.78	0.7578	0.4110	-	-	-
					23.51	0.8313	0.2538	29.69	0.9357	0.1919
343	ReSplat	IR w/ NV	2025	Rain	24.11	0.8505	0.2140	31.28	0.9538	0.1533

344 Table 2: Novel View Synthesis results and multi-view image restoration results of three types  
345 (rain+motion blur, snow+motion blur, and haze+snow) on LLFF **mixed degradation** dataset with  
346 three multi-view inputs. The **best** scores are highlighted.

347 348 349 350 351 352	Method	Corruption Type	347 348 349 350 351 352			347 348 349 350 351 352			347 348 349 350 351 352		
			347 348 349 350 351 352	347 348 349 350 351 352	347 348 349 350 351 352						
353 354 355 356 357 358 359	DiffUIR	Rain+Motion Blur	20.07	0.6910	0.4885	20.41	0.7083	0.4653	353 354 355 356 357 358 359	353 354 355 356 357 358 359	353 354 355 356 357 358 359
			20.44	0.7090	0.4555	20.74	0.7220	0.4507			
359 360 361 362	ReSplat	Snow+Motion Blur	21.63	0.7407	0.4076	22.51	0.7757	0.3848	359 360 361 362	359 360 361 362	359 360 361 362
			22.00	0.7594	0.3782	22.90	0.7908	0.3661			
362 363 364 365 366 367 368 369 370 371	DiffUIR	Haze+Snow	15.38	0.6978	0.3488	15.52	0.7702	0.2843	362 363 364 365 366 367 368 369 370 371	362 363 364 365 366 367 368 369 370 371	362 363 364 365 366 367 368 369 370 371
			20.17	0.7730	0.3148	19.92	0.8067	0.2808			

### 3.4 PRE-FILTERING WITH WARPED FEATURES

354 The final outputs of the UIR,  $\{I_{\text{out}}^i\}_{i=1}^N$ , are first depth-warped toward the novel pose using  $P_0^\phi$ . The  
355 feed-forward GS backbone then produces per-view aggregation weights  $\{W_{\text{final}}^i\}_{i=1}^N$  for combining  
356 the  $N$  warped multi-view features at each novel-view pixel. Since these weights have a critical  
357 impact on determining the radiance of the final Gaussian ellipsoids, we introduce a pre-filtering  
358 module that is additionally conditioned on the corrupted inputs  $\{I_{\text{in}}^i\}_{i=1}^N$ .

359 As illustrated in Fig. 4, the pre-filtering module takes the warped restored and degraded images  
360 as input and predicts a per-view reliability map  $\{W_{\text{pre}}^i\}_{i=1}^N$ , independently of the occlusion-based  
361 weights from the GS model. We then modulate the original GS weights by this reliability map:

$$363 \quad W_{\text{final}}^i(x) = W_{\text{pre}}^i(x) \cdot W^i(x), \quad (12)$$

364 and use the updated  $W_{\text{final}}^i$  in the splatting renderer. In other words, the pre-filtering module acts  
365 as a soft, degradation-aware gate on top of the standard visibility weights: regions where residual  
366 artifacts (e.g., remaining rain streaks, snow blobs, or haze fragments) are strong or inconsistent  
367 across views receive lower  $W_{\text{pre}}^i$  and are down-weighted, while geometry-consistent, clean structures  
368 are preserved, leading to a more robust radiance field and improved NVS quality. In practice, we  
369 simply replace the original per-view aggregation weights  $W^i$  with the updated  $W_{\text{final}}^i$ . These weights  
370 are then used in its standard multi-view feature aggregation and Gaussian rendering pipeline to  
371 determine the contribution of each input view at every novel-view location.

## 372 4 EXPERIMENTS

### 373 4.1 EXPERIMENT SETTINGS

374 **Datasets** For training our model, we utilize the synthetic multi-degradation generation pipeline pro-  
375 posed by GAURA Gupta et al. (2024) to construct a multi-view degradation dataset. We use a train-

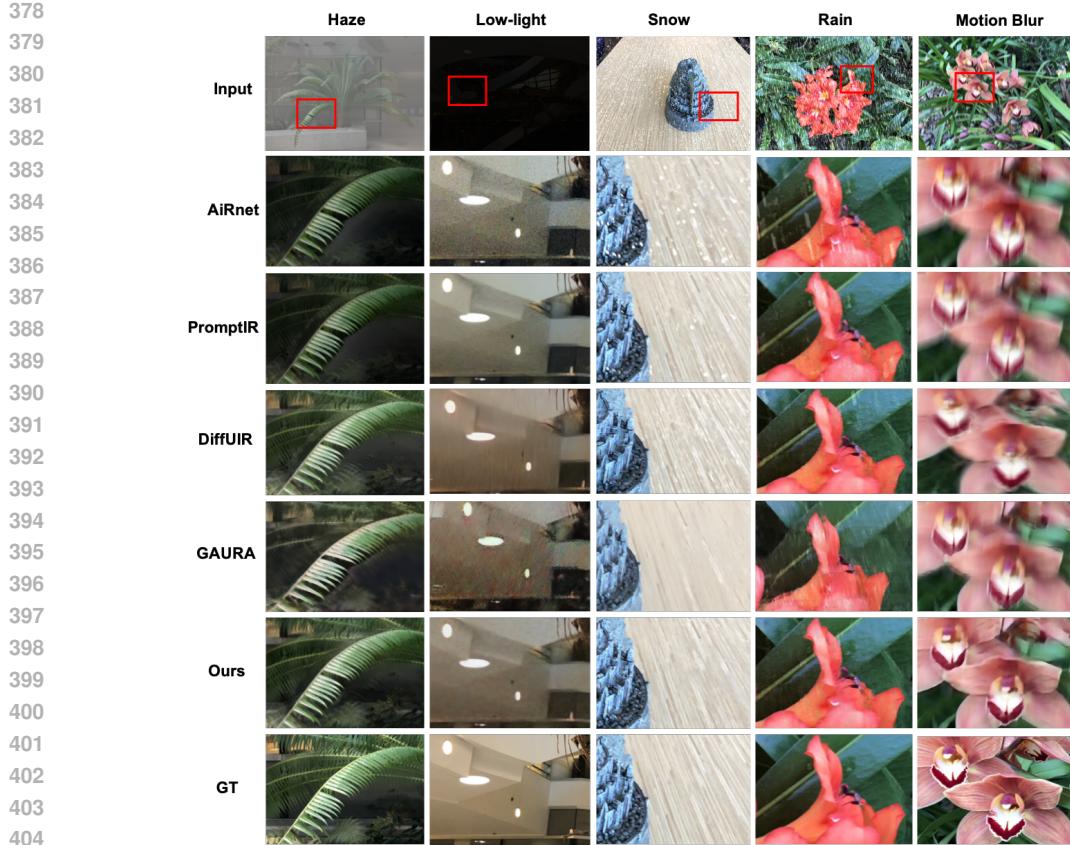


Figure 5: Comparisons of novel view synthesis results of five types (motion blur, snow, haze, low-light, rain) on LLFF degradation dataset.

dataset provided by IBRNet Wang et al. (2021), commonly used in novel view synthesis task. The test sets are divided into synthetic and real-world datasets. The synthetic dataset is generated using synthetic degradations applied to the LLFF Mildenhall et al. (2019) dataset. For real-world scenarios, we evaluate our model using the DeblurNeRF Ma et al. (2022) dataset for motion blur, the REVIDE Zhang et al. (2021) dataset for haze, and the LLNeRF Wang et al. (2023) dataset for low-light conditions.

**Network** We use DiffUIR Zheng et al. (2024), a residual diffusion model, as our baseline for image restoration. We also use MVSGaussian Liu et al. (2025), one of the state-of-the-art models, as the feed-forward GS. To accelerate the training process, MVSGaussian is first trained on our training dataset without image restoration process.

We conduct a comparison of universal image restoration with AiRnet, PromptIR, and DiffUIR. The goal of our approach is to develop an adapter that transforms a UiR model to handle multi-view inputs. Therefore, we utilize a network pretrained with a single-view UiR. For a fair comparison, all models are fine-tuned on our training dataset.

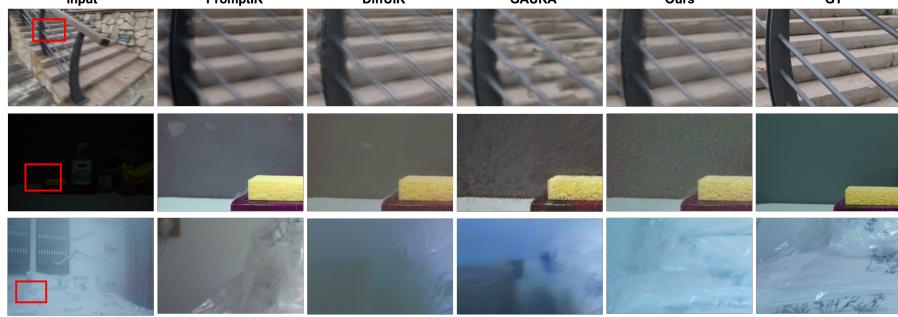
During the inference time, ReSplat uses DDIM sampling strategy with a total of three sampling steps fixed. The inference process for the three multi-view inputs can be completed within one second. For more details, please refer to the supplementary material.

## 4.2 QUANTITATIVE ANALYSIS

**Synthetic Degradation** As shown in Table 1, we evaluate ReSplat against baselines (AiRnet, PromptIR, DiffUIR, GAURA) across five corruption types: motion blur, snow, haze, low-light, and rain, for both novel view synthesis and multi-view image restoration. Performance is measured using PSNR, SSIM, and LPIPS. For novel view synthesis, ReSplat consistently outperforms other methods, especially in motion blur, snow, and rain scenarios, producing sharper, more perceptually

432 Table 3: Novel view synthesis results of  
 433 three types (motion blur, haze, low-light)  
 434 on real-world degradation datasets.

435 Method	436 Type	437 Novel View Synthesis		
		438 PSNR( $\uparrow$ )	439 SSIM( $\uparrow$ )	440 LPIPS( $\downarrow$ )
441 AiRnet	442 Motion Blur	18.59	0.6429	0.4009
		18.42	0.6289	0.3959
		21.54	0.7711	0.3909
		22.76	0.8090	0.2988
		22.91	0.8145	0.2922
443 AiRnet	444 Haze	15.91	0.7189	0.3290
		15.32	0.7106	0.3224
		16.90	0.8397	0.3920
		17.26	0.8451	0.1900
		17.75	0.8511	0.1968
445 AiRnet	446 Low-light	9.526	0.1364	0.6040
		17.10	0.8322	0.5091
		19.07	0.8503	0.6301
		22.00	0.8165	0.4958
		22.92	0.8578	0.4759



447 Figure 6: Visual Comparisons of novel view synthesis results of 3 types (motion blur, haze, low-  
 448 light) on real-world degradation dataset (DeblurNeRF, REVIDE, and LLNeRF dataset).

451 accurate views. In multi-view image restoration, ReSplat excels in high-corruption cases, achieving  
 452 the best overall visual fidelity and structural similarity. Notably, it handles complex degradations  
 453 like heavy rain and motion blur more effectively than competing models, preserving both fine details  
 454 and global consistency.

457 **Mixed Degradation** Table 2 presents a comparison between our method, ReSplat, and the  
 458 strongest baseline, DiffUIR, under various mixed degradation scenarios, including Rain+Motion  
 459 Blur, Snow+Motion Blur, and Haze+Snow. ReSplat consistently achieves the best performance  
 460 across all conditions, significantly outperforming DiffUIR in both novel view synthesis and multi-  
 461 view image restoration. We conduct a direct comparison between the top-performing method (Re-  
 462 Splat) and the next best (DiffUIR) to highlight the effectiveness of our approach. Notably, in more  
 463 challenging scenarios such as Snow+Motion Blur and Haze+Snow, ReSplat delivers considerably  
 464 higher fidelity, as reflected by higher PSNR and SSIM values and lower LPIPS.

467 **Real-World Degradation** As shown in Table 3, we evaluate ReSplat against baselines for novel  
 468 view synthesis under real-world corruptions: motion blur (DeblurNeRF Ma et al. (2022)), haze (RE-  
 469 VIDE Zhang et al. (2021)), and low-light (LLNeRF Wang et al. (2023)) dataset. ReSplat achieves  
 470 the best overall results for motion blur, preserving structural and perceptual quality. In haze, it  
 471 yields the lowest LPIPS, indicating superior perceptual quality despite similar PSNR/SSIM scores  
 472 with DiffUIR. Under low-light conditions, ReSplat balances structural integrity and perceptual fi-  
 473 delity, performing competitively across all metrics. In addition, Table 5 and Table 6 report results on  
 474 in-the-wild rain (NTURain Chen et al. (2018)) and snow (RSVD Chen et al. (2023a)) datasets, where  
 475 ReSplat consistently outperforms UIR and GS baselines, demonstrating robust generalization.

### 478 4.3 QUALITATIVE COMPARISON AND ANALYSIS

481 **Synthetic Degradation** Figure 5 shows the results on synthetic degradations including haze, low-  
 482 light, snow, rain, and motion blur. Our method consistently outperforms competing methods like  
 483 AiRnet, PromptIR, DiffUIR, and GAURA by producing clearer images with better detail preserva-  
 484 tion and color balance. Notably, it excels in haze removal, low-light enhancement, and rain streak  
 485 removal, closely matching the ground truth.

486  
487  
488  
489  
490  
491  
492  
493  
494  
495  
496  
497  
498  
499  
500  
501  
502  
503  
504  
505  
506  
507  
508  
509  
510  
511  
512  
513  
514  
515  
516  
517  
518  
519  
520  
521  
522  
523  
524  
525  
526  
527  
528  
529  
530  
531  
532  
533  
534  
535  
536  
537  
538  
539

Table 5: Novel View Synthesis (NV) results of **Rain** corruption on the real-world deraining dataset with three multi-view inputs. The best scores and second best scores are highlighted.

Method	Type	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
Rain		24.05	0.8183	0.1955
		24.23	0.8230	0.1801
		19.39	0.6602	0.3987
		23.99	0.8145	0.2094
		24.35	0.8232	0.1772

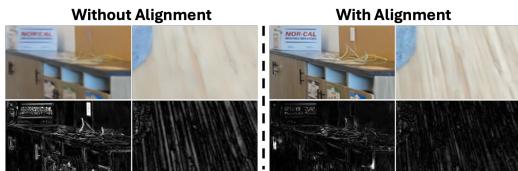


Figure 7: Qualitative comparison of the alignment module. The top row shows the restored RGB outputs, while the bottom row visualizes the corresponding error maps.

Table 6: Novel View Synthesis (NV) results of **Snow** corruption on the real-world desnowing dataset with three multi-view inputs. The best scores and second best scores are highlighted.

Method	Type	PSNR( $\uparrow$ )	SSIM( $\uparrow$ )	LPIPS( $\downarrow$ )
Snow		20.23	0.7035	0.3103
		21.27	0.7192	0.3007
		20.22	0.7578	0.3647
		22.12	0.8215	0.2277
		22.45	0.8263	0.2175

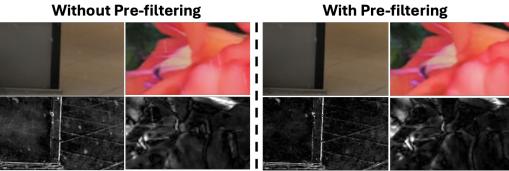


Figure 8: Qualitative comparison of pre-filtering module. The top row shows the restored RGB outputs, while the bottom row visualizes the corresponding error maps.

**Real-World Degradation** Figure 6 presents results on real-world degradations such as blur, low-light scenes, and haze. Our approach effectively reconstructs details and reduces artifacts, outperforming other methods which often introduce noise or fail to restore fine structures. The results demonstrate that our method achieves superior restoration, preserving natural colors and sharpness in challenging real-world scenarios.

#### 4.4 ABLATION STUDIES

We conduct an ablation study using four model variants: Model 1 as baseline, Model 2 with pre-filtering, Model 3 with alignment, and Model 4 with both. Metrics are averaged over five LLFF degradation datasets. As shown in Table 4, each component contributes to improved performance in novel view synthesis. Adding pre-filtering (Model 2) increases PSNR and reduces LPIPS, indicating a modest improvement in reconstruction quality. Alignment alone (Model 3) further enhances PSNR and other quality metrics compared to the baseline. When both alignment and pre-filtering are applied (Model 4), the model achieves the best overall results, with a PSNR of 22.69, demonstrating a clear effect across various degradations. Qualitative comparisons in Fig. 7 and Fig. 8 illustrate that alignment reduces multi-view geometric inconsistencies, while pre-filtering suppresses residual artifacts without destroying fine structures in the rendered novel views.

## 5 LIMITATIONS

Although ReSplat performs well across various degradation scenarios, several limitations remain. The diffusion-based refinement increases computational cost and memory usage compared to purely feed-forward 3DGS, making efficient high-resolution scaling an important direction for future work. The method also inherits the representational biases of Gaussian Splatting, showing reduced fidelity in scenes with strong transparency, specularities, or pronounced view-dependent effects. Finally, ReSplat depends on a pretrained universal image restoration model for guidance; while this introduces a dependency, it also enables the framework to naturally benefit from future advances in universal restoration without architectural changes.

## 6 CONCLUSION

We present a feed-forward gaussian splatting framework for degradation-agnostic novel view synthesis. By integrating a residual diffusion model with 3D cross-attention and multi-view pre-filtering, our method robustly restores images and improves geometry estimation, outperforming existing approaches in both novel view synthesis and universal multi-view image restoration.

## 540 REFERENCES

541  
 542 Saeed Anwar and Nick Barnes. Densely residual laplacian super-resolution. *IEEE Transactions on*  
 543 *Pattern Analysis and Machine Intelligence*, 44(3):1192–1204, 2020.

544 Ang Cao, Chris Rockwell, and Justin Johnson. Fwd: Real-time novel view synthesis with forward  
 545 warping and depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 546 *Recognition*, pp. 15713–15724, 2022.

547  
 548 David Charatan, Sizhe Li, Andrea Tagliasacchi, and Vincent Sitzmann. pixelsplat: 3d gaussian splats  
 549 from image pairs for scalable generalizable 3d reconstruction. *arXiv preprint arXiv:2312.12337*,  
 550 2023.

551 Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su.  
 552 Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings*  
 553 *of the IEEE/CVF International Conference on Computer Vision*, pp. 14124–14133, 2021.

554 Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. Tensorrf: Tensorial radiance  
 555 fields. In *European Conference on Computer Vision*, pp. 333–350. Springer, 2022.

556  
 557 Haoyu Chen, Jingjing Ren, Jinjin Gu, Hongtao Wu, Xuequan Lu, Haoming Cai, and Lei Zhu. Snow  
 558 removal in video: A new dataset and a novel method. In *2023 IEEE/CVF International Conference*  
 559 *on Computer Vision (ICCV)*, pp. 13165–13176. IEEE, 2023a.

560 Jie Chen, Cheen-Hau Tan, Junhui Hou, Lap-Pui Chau, and He Li. Robust video content alignment  
 561 and compensation for rain removal in a cnn framework. In *Proceedings of the IEEE conference*  
 562 *on computer vision and pattern recognition*, pp. 6286–6295, 2018.

563  
 564 Wei-Ting Chen, Wang Yifan, Sy-Yen Kuo, and Gordon Wetzstein. Dehazenerf: Multiple im-  
 565 age haze removal and 3d shape reconstruction using neural radiance fields. *arXiv preprint*  
 566 *arXiv:2303.11364*, 2023b.

567 Yuedong Chen, Haofei Xu, Chuanxia Zheng, Bohan Zhuang, Marc Pollefeys, Andreas Geiger, Tat-  
 568 Jen Cham, and Jianfei Cai. Mvsplat: Efficient 3d gaussian splatting from sparse multi-view  
 569 images. In *European Conference on Computer Vision*, pp. 370–386. Springer, 2025.

570  
 571 Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz  
 572 Khan. Adair: Adaptive all-in-one image restoration via frequency mining and modulation. In  
 573 *13th International Conference on Learning Representations, ICLR 2025*, pp. 57335–57356. In-  
 574 ternational Conference on Learning Representations, ICLR, 2025.

575 Ziteng Cui, Lin Gu, Xiao Sun, Xianzheng Ma, Yu Qiao, and Tatsuya Harada. Aleth-nerf: Illumina-  
 576 tion adaptive nerf with concealing field assumption. In *Proceedings of the AAAI Conference on*  
 577 *Artificial Intelligence*, volume 38, pp. 1435–1444, 2024.

578 Kangle Deng, Andrew Liu, Jun-Yan Zhu, and Deva Ramanan. Depth-supervised nerf: Fewer views  
 579 and faster training for free. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*  
 580 *Pattern Recognition*, pp. 12882–12891, 2022.

581  
 582 Ben Fei, Zhaoyang Lyu, Liang Pan, Junzhe Zhang, Weidong Yang, Tianyue Luo, Bo Zhang, and  
 583 Bo Dai. Generative diffusion prior for unified image restoration and enhancement. In *Proceedings*  
 584 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9935–9946, 2023.

585 Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo  
 586 Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF*  
 587 *Conference on Computer Vision and Pattern Recognition*, pp. 5501–5510, 2022.

588  
 589 Vinayak Gupta, Rongali Simhachala Venkata Girish, Ayush Tewari, Kaushik Mitra, et al. Gaura:  
 590 Generalizable approach for unified restoration and rendering of arbitrary views. *arXiv preprint*  
 591 *arXiv:2407.08221*, 2024.

592 Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High  
 593 dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer*  
 594 *Vision and Pattern Recognition*, pp. 18398–18408, 2022.

594 Mohammad Mahdi Johari, Yann Lepoittevin, and François Fleuret. Geonerf: Generalizing nerf with  
 595 geometry priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 596 *Recognition*, pp. 18365–18375, 2022.

597 Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splat-  
 598 ting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023.

600 Jie Long Lee, Chen Li, and Gim Hee Lee. Disr-nerf: Diffusion-guided view-consistent super-  
 601 resolution nerf. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 602 *Recognition*, pp. 20561–20570, 2024.

603 Seungjun Lee and Gim Hee Lee. Diet-gs: Diffusion prior and event stream-assisted motion de-  
 604 blurring 3d gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition*  
 605 *Conference*, pp. 21739–21749, 2025.

606 Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restora-  
 607 tion for unknown corruption. In *Proceedings of the IEEE/CVF conference on computer vision*  
 608 *and pattern recognition*, pp. 17452–17462, 2022.

609 Xin Lin, Shi Luo, Xiaojun Shan, Xiaoyu Zhou, Chao Ren, Lu Qi, Ming-Hsuan Yang, and Nuno  
 610 Vasconcelos. Hqgs: High-quality novel view synthesis with gaussian splatting in degraded scenes.  
 611 In *The Thirteenth International Conference on Learning Representations*.

612 Xin Lin, Jingtong Yue, Kelvin CK Chan, Lu Qi, Chao Ren, Jinshan Pan, and Ming-Hsuan Yang.  
 613 Multi-task image restoration guided by robust dino features. *arXiv preprint arXiv:2312.01677*,  
 614 2023.

615 Jiawei Liu, Qiang Wang, Huijie Fan, Yinong Wang, Yandong Tang, and Liangqiong Qu. Residual  
 616 denoising diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision*  
 617 *and Pattern Recognition*, pp. 2773–2783, 2024.

618 Tianqi Liu, Guangcong Wang, Shoukang Hu, Liao Shen, Xinyi Ye, Yuhang Zang, Zhiguo Cao,  
 619 Wei Li, and Ziwei Liu. Mvsgaussian: Fast generalizable gaussian splatting reconstruction from  
 620 multi-view stereo. In *European Conference on Computer Vision*, pp. 37–53. Springer, 2025.

621 Yuan Liu, Sida Peng, Lingjie Liu, Qianqian Wang, Peng Wang, Christian Theobalt, Xiaowei Zhou,  
 622 and Wenping Wang. Neural rays for occlusion-aware image-based rendering. In *Proceedings of*  
 623 *the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7824–7833, 2022.

624 Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling  
 625 vision-language models for universal image restoration. *arXiv preprint arXiv:2310.01018*, 3(8),  
 626 2023.

627 Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Photo-realistic  
 628 image restoration in the wild with controlled vision-language models. In *Proceedings of the*  
 629 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6641–6651, 2024.

630 Jiaqi Ma, Tianheng Cheng, Guoli Wang, Qian Zhang, Xinggang Wang, and Lefei Zhang. Prores:  
 631 Exploring degradation-aware visual prompt for universal image restoration. *arXiv preprint*  
 632 *arXiv:2306.13653*, 2023.

633 Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-  
 634 nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF Conference on*  
 635 *Computer Vision and Pattern Recognition*, pp. 12861–12870, 2022.

636 Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy,  
 637 and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collec-  
 638 tions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,  
 639 pp. 7210–7219, 2021.

640 Ben Mildenhall, Pratul P Srinivasan, Rodrigo Ortiz-Cayon, Nima Khademi Kalantari, Ravi Ra-  
 641 mamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: Practical view synthesis with  
 642 prescriptive sampling guidelines. *ACM Transactions on Graphics (TOG)*, 38(4):1–14, 2019.

648 Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and  
 649 Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications*  
 650 *of the ACM*, 65(1):99–106, 2021.

651

652 Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P Srinivasan, and Jonathan T Barron.  
 653 Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *Proceedings of*  
 654 *the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16190–16199, 2022.

655 Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics prim-  
 656 itives with a multiresolution hash encoding. *arXiv preprint arXiv:2201.05989*, 2022.

657

658 Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-  
 659 based denoising diffusion models. *IEEE Transactions on Pattern Analysis and Machine Intelli-  
 660 gence*, 45(8):10346–10357, 2023.

661 Vaishnav Potlapalli, Syed Waqas Zamir, Salman H Khan, and Fahad Shahbaz Khan. Promptir:  
 662 Prompting for all-in-one image restoration. *Advances in Neural Information Processing Systems*,  
 663 36, 2024.

664 Chu-Jie Qin, Rui-Qi Wu, Zikun Liu, Xin Lin, Chun-Le Guo, Hyun Hee Park, and Chongyi Li. Re-  
 665 store anything with masks: Leveraging mask image modeling for blind all-in-one image restora-  
 666 tion. In *European Conference on Computer Vision*, pp. 364–380. Springer, 2024.

667

668 Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner.  
 669 Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the*  
 670 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12892–12901, 2022.

671 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv*  
 672 *preprint arXiv:2010.02502*, 2020.

673

674 Mohammed Suhail, Carlos Esteves, Leonid Sigal, and Ameesh Makadia. Generalizable patch-based  
 675 neural rendering. In *European Conference on Computer Vision*, pp. 156–174. Springer, 2022.

676 Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast conver-  
 677 gence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF conference on computer*  
 678 *vision and pattern recognition*, pp. 5459–5469, 2022.

679

680 Chen Wang, Xian Wu, Yuan-Chen Guo, Song-Hai Zhang, Yu-Wing Tai, and Shi-Min Hu. Nerf-  
 681 sr: High quality neural radiance fields using supersampling. In *Proceedings of the 30th ACM*  
 682 *International Conference on Multimedia*, pp. 6445–6454, 2022a.

683

684 Haoyuan Wang, Xiaogang Xu, Ke Xu, and Rynson WH Lau. Lighting up nerf via unsupervised  
 685 decomposition and enhancement. In *Proceedings of the IEEE/CVF International Conference on*  
 686 *Computer Vision*, pp. 12632–12641, 2023.

687

688 Peihao Wang, Xuxi Chen, Tianlong Chen, Subhashini Venugopalan, Zhangyang Wang, et al. Is  
 689 attention all nerf needs? *arXiv preprint arXiv:2207.13298*, 2022b.

690

691 Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T  
 692 Barron, Ricardo Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-  
 693 view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision*  
 694 *and Pattern Recognition*, pp. 4690–4699, 2021.

695

696 Yi Wei, Shaohui Liu, Yongming Rao, Wang Zhao, Jiwen Lu, and Jie Zhou. Nerfingmvs: Guided op-  
 697 timization of neural radiance fields for indoor multi-view stereo. In *Proceedings of the IEEE/CVF*  
 698 *International Conference on Computer Vision*, pp. 5610–5619, 2021.

699

700 Anran Wu, Long Peng, Xin Di, Xueyuan Dai, Chen Wu, Yang Wang, Xueyang Fu, Yang Cao, and  
 701 Zheng-Jun Zha. Robustgs: Unified boosting of feedforward 3d gaussian splatting under low-  
 702 quality conditions. *arXiv preprint arXiv:2508.03077*, 2025.

703

704 Haofei Xu, Anpei Chen, Yuedong Chen, Christos Sakaridis, Yulun Zhang, Marc Pollefeys, Andreas  
 705 Geiger, and Fisher Yu. Murf: Multi-baseline radiance fields. *arXiv preprint arXiv:2312.04565*,  
 706 2023.

702 Youngho Yoon and Kuk-Jin Yoon. Cross-guided optimization of radiance fields with multi-view  
 703 image super-resolution for high-resolution novel view synthesis. In *Proceedings of the IEEE/CVF*  
 704 *conference on computer vision and pattern recognition*, pp. 12428–12438, 2023.

705 Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. Plenocubes for  
 706 real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International*  
 707 *Conference on Computer Vision*, pp. 5752–5761, 2021a.

708 Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from  
 709 one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern*  
 710 *Recognition*, pp. 4578–4587, 2021b.

711 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-  
 712 Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the*  
 713 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 14821–14831, 2021.

714 Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao.  
 715 Ingredient-oriented multi-degradation learning for image restoration. In *Proceedings of the*  
 716 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5825–5835, 2023.

717 Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser:  
 718 Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26  
 719 (7):3142–3155, 2017.

720 Xinyi Zhang, Hang Dong, Jinshan Pan, Chao Zhu, Ying Tai, Chengjie Wang, Jilin Li, Feiyue Huang,  
 721 and Fei Wang. Learning to restore hazy video: A new real-world dataset and a new method.  
 722 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.  
 723 9239–9248, 2021.

724 Xu Zhang, Jiaqi Ma, Guoli Wang, Qian Zhang, Huan Zhang, and Lefei Zhang. Perceive-ir: Learning  
 725 to perceive degradation better for all-in-one image restoration. *IEEE Transactions on Image*  
 726 *Processing*, 2025.

727 Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for  
 728 image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern*  
 729 *recognition*, pp. 2472–2481, 2018.

730 Dian Zheng, Xiao-Ming Wu, Shuzhou Yang, Jian Zhang, Jian-Fang Hu, and Wei-Shi Zheng. Selec-  
 731 tive hourglass mapping for universal image restoration based on diffusion model. In *Proceedings*  
 732 *of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25445–25455,  
 733 2024.

734 Chen Ziwen, Hao Tan, Kai Zhang, Sai Bi, Fujun Luan, Yicong Hong, Li Fuxin, and Zexiang Xu.  
 735 Long-lrm: Long-sequence large reconstruction model for wide-coverage gaussian splats. *arXiv*  
 736 *preprint arXiv:2410.12781*, 2024.

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755