

DeRainGS: Gaussian Splatting for Enhanced Scene Reconstruction in Rainy Environments

Shuhong Liu¹, Xiang Chen², Hongming Chen³, Quanfeng Xu^{4,5}, Mingrui Li^{6*}

¹The University of Tokyo, 7-3-1 Hongo, Tokyo, Japan

²Nanjing University of Science and Technology, 200 Xiaolingwei Rd, Nanjing, China

³Dalian Maritime University, 1 Linghai Rd, Dalian, China

⁴Shanghai Astronomical Observatory, 80 Nandan Rd, Shanghai, China

⁵University of Chinese Academy of Sciences, 1 Yanqihu East Rd, Beijing, China

⁶Dalian University of Technology, 2 Linggong Rd, Dalian, China

Abstract

Reconstruction under adverse rainy conditions poses significant challenges due to reduced visibility and the distortion of visual perception. These conditions can severely impair the quality of geometric maps, which is essential for applications ranging from autonomous planning to environmental monitoring. In response to these challenges, this study introduces the novel task of 3D Reconstruction in Rainy Environments (3DRRE), specifically designed to address the complexities of reconstructing 3D scenes under rainy conditions. To benchmark this task, we construct the HydroViews dataset that comprises a diverse collection of both synthesized and real-world scene images characterized by various intensities of rain streaks and raindrops. Furthermore, we propose DeRainGS, the first 3DGS method tailored for reconstruction in adverse rainy environments. Extensive experiments across a wide range of rain scenarios demonstrate that our method delivers state-of-the-art performance, remarkably outperforming existing occlusion-free methods by a large margin.

Home Page and Datasets — <https://deraings.github.io>

Introduction

Images and videos captured in rainy conditions often exhibit significant quality degradation due to reduced visibility and distortions such as rain streaks and waterdrops. While such conditions pose challenges across various fields, they are particularly problematic for 3D reconstruction tasks, which rely on clear and accurate visual inputs to effectively model environments. In practical applications like autonomous vehicles (Hnewa and Radha 2020; Deng et al. 2023), aerial surveying (Chang et al. 2024; Deng et al. 2024), and outdoor robotics (Fumagalli et al. 2021), rain can severely distort the geometry, depth, and textures that are critical for accurate mapping, leading to blurred details and spatial distortions.

Remarkable strides in 3D scene reconstruction, notably through Neural Radiance Field (NeRF) (Mildenhall et al. 2021) and 3D Gaussian Splatting (3DGS) (Kerbl et al.

2023), have substantially enhanced the quality and efficiency of 3D scene representation and rendering. NeRF employs a neural implicit representation using Multi-layer Perceptron (MLP), providing a continuous volumetric scene representation that enables high-quality novel view synthesis (NVS). In contrast, 3DGS utilizes an explicit representation with spherical harmonics (SH) to model scene geometry and appearance, delivering high-fidelity reconstruction with fine-grained details while achieving exceptional rendering speed through efficient Gaussian rasterization.

Building on these advancements, recent studies (Martin-Brualla et al. 2021; Chen et al. 2022b; Cui et al. 2024; Ren et al. 2024; Kulhanek et al. 2024; Xu, Mei, and Patel 2024; Wang, Wang, and Qi 2024; Lee et al. 2025) have expanded to facilitate reconstruction in unstructured, “in-the-wild” environments, tackling challenges posed by transient or dynamic objects and adapting to versatile styles and appearances. However, these approaches have only been demonstrated to be effective in mitigating typical transient and dynamic occluders. When it comes to addressing adverse weather conditions such as rain, these methods encounter significant limitations. Rain introduces unique visual artifacts that are not merely transient distractors but continuous disturbances that affect visibility throughout the entire scene. Unlike typical occluders, which can be effectively segmented and removed (Yang et al. 2023; Zhang et al. 2024), the pervasive impact of rain on visibility and the complex interplay of light and water remarkably challenge existing reconstruction techniques.

On the other hand, current image deraining approaches, which predominantly focus on single image deraining (Wang et al. 2019; Wei et al. 2021; Chen et al. 2022a; Chen, Pan, and Dong 2024) or video deraining for static scenes (Wei et al. 2017; Jiang et al. 2018; Cho et al. 2020; Yue et al. 2021), face certain limitations in the context of 3D reconstruction. Despite showing efficacy in mitigating rain effects, these methods struggle to fully eliminate rain-induced distortions due to the irregular, complex nature of real-world rain streaks and the ill-posed nature of the problem. Crucially, these methods focus primarily on pixel-level manipulations in 2D images, neglecting the critical aspect of geometry accuracy that is heavily affected by rain.

To overcome these challenges, we introduce a new task:

*Corresponding Author: mmclmr@mail.dlut.edu.cn
Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

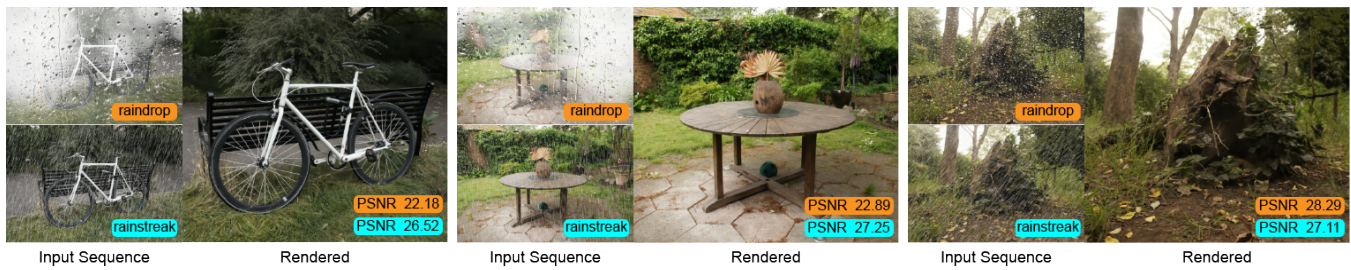


Figure 1: DeRainGS reconstructs clean scenes from input rainy images. The visualization displays three different scenes along with their reconstruction outcomes of DeRainGS on our HydroViews dataset.

3D Reconstruction in Rainy Environments (3DRRE).

This novel formulation aims at reconstructing 3D scenes under adverse rainy conditions. To benchmark this task, we have developed the HydroViews dataset, which includes both synthesized and real-world scene images featuring a variety of rain streaks and raindrops. Building upon this dataset, we propose DeRainGS, a novel reconstruction pipeline tailored for the 3DRRE problems. DeRainGS integrates unique rainy image enhancement and occlusion masking modules to explicitly tackle the impact of rain on visual features and geometric consistency.

To summarize, our key contributions include:

- We establish a new task in 3D reconstruction, specifically aiming at constructing clean scenes from rainy image input. This task is supported by our newly constructed HydroViews dataset, featuring both synthesized and real-world scenes to cover common practical scenarios.
- We propose DeRainGS, the first 3D reconstruction method solely utilizing 3DGS for high-quality reconstruction under adverse rainy conditions.
- DeRainGS effectively incorporates novel rainy image enhancement and occlusion masking modules into a unified pipeline, collectively contributing to reconstructing clean scenes from rainy environments.

Related Work

Image Deraining

In recent years, image deraining has become a crucial pre-processing step in high-level vision tasks. It can be further categorized into single image deraining and video deraining (Chen et al. 2023c). Video deraining utilizes temporal redundancy and rain dynamics, whereas single image deraining primarily relies on spatial information from neighboring pixels and the visual characteristics of rain and background scenes. Deep learning-based methods have become predominant in image deraining due to their powerful representational capacity compared to conventional prior-based methods. Researchers can opt for deep models based on convolutional neural networks (CNNs) (Ren et al. 2019; Yang et al. 2019) or Transformer architectures (Chen et al. 2023b,a) to directly estimate a clear image from a rainy one.

Compared to CNNs, Transformers have an advantage in image deraining, especially when dealing with long spatial rain effects. Transformers can effectively capture and model

the dependencies over longer distances in the image, allowing them to better understand and remove complex rain patterns that extend across larger spatial areas. Xiao et al. (Xiao et al. 2022) introduce an image deraining Transformer that incorporates spatial-based and window-based self-attention modules. Recently, Chen et al. (Chen, Pan, and Dong 2024) are the first to explore integrating implicit neural representations into Transformers to enhance the robustness of image deraining. However, these methods can only handle specific forms of rain degradation. In real-world scenarios, rain is diverse and uncertain in terms of density, shape, location, and size. Therefore, improving the generalization capability of deraining models across different types of rain is crucial.

3D Reconstruction in the Wild

Rain, particularly waterdrops that adhere to camera lenses, can be treated as semi-transparent occluders to the views. Other occlusions, such as moving pedestrians and vehicles, are frequently presented in real-world image collections. Regardless of varying illumination across frames, a primary concern arises from these dynamic obstacles which can introduce substantial artifacts in reconstructed scenes. NeRF-W (Martin-Brualla et al. 2021) tackled transient occlusions with a dual NeRF setup and aleatoric uncertainty. Ha-NeRF (Chen et al. 2022b) and CR-NeRF (Yang et al. 2023) use learned priors for improved occlusion removal. In 3DGS, Wild-GS (Xu, Mei, and Patel 2024) and GS-W (Zhang et al. 2024) use U-Net (Ronneberger, Fischer, and Brox 2015) segmentation masks for occlusion supervision. SpotlessPlats (Sabour et al. 2024) employs an MLP classifier for dynamic distractors, and WildGaussians (Kulhanek et al. 2024) uses DINOv2 (Oquab et al. 2023) features to mask transient objects. While these in-the-wild methods effectively address floaters from transient obstacles, they falter with the widespread distortion caused by rain which uniquely affects visibility and scene integrity. To address this challenge, we incorporate a novel image enhancement network that is essential for mitigating rain-induced disturbances.

Three closely related works to ours, Occlusion-free NeRF (Zhu et al. 2023), DerainNeRF (Li et al. 2024), and RainyScape (Lyu, Liu, and Hou 2024) attempt to mitigate rain effects in 3D reconstruction. Occlusion-free NeRF (Zhu et al. 2023) employed bidirectional volume rendering to distinguish foreground obstacles. DerainNeRF (Li et al. 2024) trivially combines a deraining AttGAN network (Qian et al.

2018) with NeRF, removing the waterdrops using the hard-coded masks. RainyScape (Lyu, Liu, and Hou 2024) incorporates a NeRF model and directional embedding for unsupervised rainstreak removal. However, these methods are validated only in scenarios with limited camera angles that enable clear foreground-background separation, making them impractical for real-world applications. Moreover, while these approaches all incorporate NeRF models, our method is purely based on 3DGS, offering superior training and real-time rendering capabilities.

Dataset Construction

We observe that real-world rain effects exhibit a wide range of appearance characteristics, such as variations in rain shape, length, and direction. Existing rain datasets (Chen et al. 2023c) typically create rainy images using Photoshop software to randomly add rain streaks onto clear backgrounds. In fact, synthesizing rain in 2D images and 3D scenes involves distinct processes and challenges, with a significant difference being the continuity across multiple camera views in 3D scenes. In this study, we propose a new pipeline for synthesizing rainy images tailored for 3D scenes that include both rain streaks and raindrops. As depicted in Figure 2, we synthesize data derived from real-world scenes captured by MipNeRF-360 (Barron et al. 2022) and Tanks-and-Temples (Knapitsch et al. 2017), along with real-world rainy environments that we have independently collected.




Figure 2: The visualization of our HydroViews dataset. From top to bottom, it includes synthesized raindrops, rain streaks, and real-world scenes.

Rain Streak Generation

We synthesize rain streaks using motion blur, which involves creating streaks that mimic the natural motion of falling rain. This technique leverages the motion blur effect to simulate the linear, directional movement of rain, capturing the essence of rain streaks by considering their repeatability and directionality. Mathematically, this can be represented as:

$$\mathbf{S} = \mathbf{K}(l, \theta, w) \otimes \mathbf{N}(n), \quad (1)$$

where \mathbf{N} represents the rain streak layers generated by random noise. We utilize uniform random numbers and thresholds to control the noise level n . The length and angle of the motion blur kernel $\mathbf{K} \in \mathbb{R}^{p \times p}$ are denoted by l and θ . Additionally, we incorporate a rotated diagonal kernel with Gaussian blur to achieve the desired rain thickness w . The values for noise quantity n , rain length l , rain angle θ , and rain

thickness w are sampled from the ranges $[100, 300]$, $[20, 40]$, $[40^\circ, 120^\circ]$, and $[3, 7]$. The symbol \otimes denotes the spatial convolution operator. To ensure the continuity of rain synthesis in 3D environments, we keep these parameters within fixed ranges in the same scene.

Raindrop Generation

As an atmospheric process, raindrops are individual spherical water droplets, while rain streaks are elongated trails formed by the motion blur of falling raindrops. Here, we use the open-source 3D graphics engine Blender¹ to simulate continuous raindrops in realistic scenes, utilizing its fluid motion model and rain generator plugin to render raindrops with accurate transparency. We model the blurred or occluded effects D of raindrops in dispersed, small-sized, locally coherent regions and combine these with background images B to create the degraded raindrop images R_d . Mathematically, the synthesis process can be expressed as:

$$R_d = (1 - M_d) \odot B + D, \quad (2)$$

where \odot means element-wise multiplication, D denotes the occlusion or blurring effects of raindrops, and B is a binary mask. When $M_d(x) = 1$, the pixel x in the mask is part of the raindrop region; otherwise, it belongs to the background.

Real-world Data Collection

We collect videos of real-world rainy scenes using a SONY-A7R3 camera in Full HD format. The scene images are extracted from the video at a rate of one frame per second. The footage is captured across six different scenes in Tokyo, with most scenes specifically filmed in the evening to enhance the contrast of rain streaks through ambient light. To ensure privacy protection, all pedestrians and any identifiable personal elements are manually erased using Photoshop software.

Proposed Method

Preliminaries: 3D Gaussian Splatting (3DGS)

3DGS (Kerbl et al. 2023) is a breakthrough method for reconstructing and rendering 3D scenes from an image collection. It employs a set of 3D Gaussian SHs G_i to explicitly represent the scene. Each Gaussian G_i is characterized by its mean position $\mu_i \in \mathbb{R}^3$, a 3D covariance matrix $\sigma_i \in \mathbb{R}^{3 \times 3}$, an opacity α_i , and view-dependent color coefficients c_i . The mean positions μ_i of the Gaussian points are initialized using point clouds derived from Structure from Motion. The covariance matrix Σ_i is decomposed into a scaling matrix S and a rotation matrix R to ensure positive semi-definiteness:

$$\Sigma_i = RSS^T R^T. \quad (3)$$

During rendering, these 3D Gaussians are projected onto the image plane. Given a world-to-camera transformation W and the Jacobian J of an affine approximation of perspective projection, the 2D covariance matrix Σ'_i is computed as:

$$\Sigma'_i = (JW\Sigma_i W^T J^T). \quad (4)$$

¹<https://blendermarket.com/products/rain-generator>

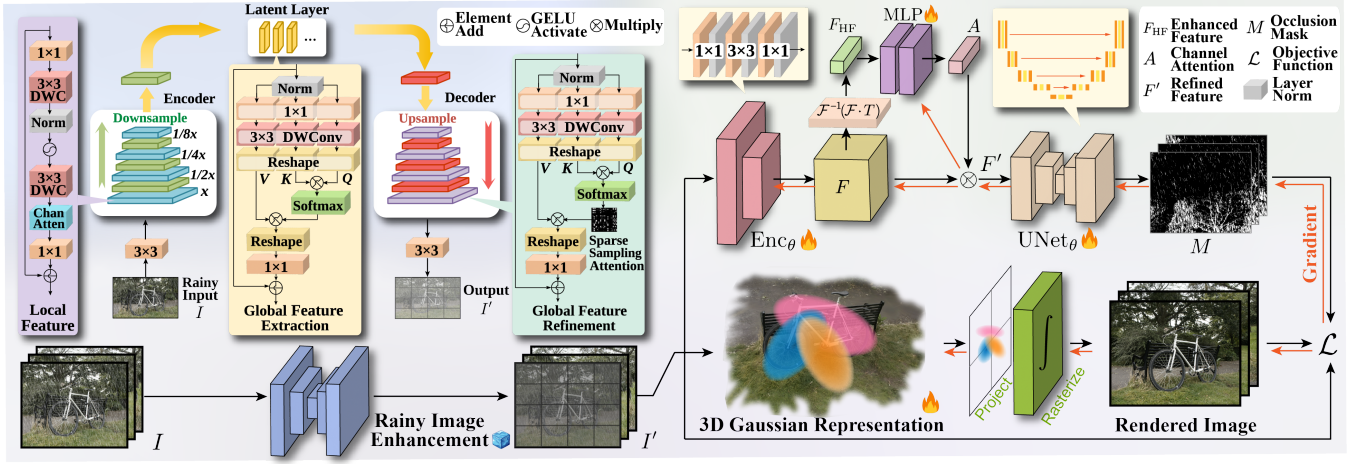


Figure 3: The illustration of the DeRainGS pipeline. The left part displays the rainy image enhancement procedure and its network structure, which is pretrained before being applied to a scene. The right part demonstrates the reconstruction by 3DGS, utilizing learned occlusion masks to handle rain-induced artifacts.

The color of each pixel is then determined through alpha compositing, traversing the Gaussians from front to back

$$C = \sum_i c_i(\mathbf{r}) \alpha'_i \prod_{j=1}^{i-1} (1 - \alpha'_j). \quad (5)$$

In this context, $c_i(\mathbf{r})$ is the view-dependent color computed using SHs, and α'_i is the blending weight equal to the product of the Gaussian’s opacity and its 2D projection.

Rainy Image Enhancement

We first perform rainy image enhancement as a preprocessing step to tackle rain effects, providing an enhanced image to guide subsequent scene reconstruction. To enhance the robustness of the deraining model across different scenes, we model complex rain distribution by combining local and non-local information. In the enhancement network, we utilize a 5-level encoder-decoder architecture. The overall framework is presented on the left of Figure 3. Assuming I_k is the original image, it undergoes local feature extraction, global feature extraction, and global feature refinement stages, resulting in I'_k as the enhanced image.

For local feature extraction, it contains $\{4, 6, 7, 8\}$ convolutional blocks in the network encoder. Specifically, we begin by doubling the channel dimension, and then use the DWConvLN-GELU-DWConv design to capture rain information. In the bottleneck layer of the encoder-decoder, we stack 8 vanilla Transformer modules (Xiao et al. 2022) for global deep feature extraction. For the network decoder, we adopt $\{3, 6, 7, 8\}$ sparse Transformer modules (Chen et al. 2023a) to reconstruct the final outputs. Here, sparse Transformer modules are employed to eliminate irrelevant feature interactions among tokens to generate high-quality outputs, while vanilla Transformer modules are employed to ensure necessary information flows through the whole network. In summary, our goal is to effectively integrate their comple-

mentary features by CNN and Transformer to better achieve comprehensive rain distribution prediction.

The end-to-end rainy image enhancement network is trained on the rain streak dataset 4K-Rain13k (Chen et al. 2024) and raindrop dataset UAV-Rain1k (Chang et al. 2024), and the model is frozen during the reconstruction process.

Scene Reconstruction

The varied shapes and distortions caused by rain pose substantial challenges for a derainer, making the complete removal of rain effects extremely difficult. Moreover, beyond the residual effects of rain, the end-to-end enhancement process may also introduce additional artifacts as a result of over-correction for distortion. As shown in Figure 4, these artifacts, which often appear as high-frequency patterns, can adversely impact the quality of the reconstructed scene.

To address the intricate problem of high-frequency artifacts, we propose an unsupervised learning approach to predict masks of artifacts. As depicted in the right of Figure 3, our method strategically leverages spectral pooling within a channel attention module, which is designed to enhance sensitivity to high-frequency details that could potentially manifest as artifacts. Thereafter, refined features are processed by a U-Net model to generate masks. Unlike typical methods that segment by categories (Xu, Mei, and Patel 2024; Yang et al. 2023), our network is designed to identify rain artifacts directly. This design leverages the fact that raindrops typically do not conform to consistent multiview geometry, enabling clearer reconstructions from alternate views.

Frequency-based Feature Channel Attention The initial step in the pipeline involves a CNN encoder Enc_θ that processes the enhanced image I'_k and produces a feature map $F_k \in \mathbb{R}^{C \times H' \times W'}$. The encoder captures a broad array of features across different channels. Following feature extraction, we utilize spectral pooling to manipulate these features to accentuate high-frequency components that are crucial for




Figure 4: Different types of artifacts, produced either by the image enhancement process or the persisting rain effect.

detecting artifacts. This process involves transforming the F_k into the frequency domain using the Fourier Transform:

$$\tilde{F}_k = \mathcal{F}(F_k) = \sum_{x=0}^{H'-1} \sum_{y=0}^{W'-1} F_k[x, y] \cdot e^{-2\pi i(\frac{ux}{H'} + \frac{vy}{W'})} \quad (6)$$

where \mathcal{F} denotes the Fourier Transform applied to each channel of F_k and produces frequency components $\tilde{F}_k \in \mathbb{R}^{C \times H' \times W'}$. In the frequency domain, a high-pass filtering operation is employed to isolate and enhance high-frequency features that are potentially indicative of artifacts:

$$\tilde{F}_{\text{HF}} = \tilde{F}_k \cdot T \quad (7)$$

T represents the high-pass filter mask. Subsequently, we utilize the Inverse Fourier Transform to convert these emphasized frequency-domain features back to the spatial domain:

$$\begin{aligned} F_{\text{HF}} &= \mathcal{F}^{-1}(\tilde{F}_{\text{HF}}) \\ &= \frac{1}{H'W'} \sum_{u=0}^{H'-1} \sum_{v=0}^{W'-1} \tilde{F}_{\text{HF}}[u, v] \cdot e^{2\pi i(\frac{ux}{H'} + \frac{vy}{W'})} \end{aligned} \quad (8)$$

The resulting feature map $F_{\text{HF}} \in \mathbb{R}^{C \times H' \times W'}$ predominantly contains high-frequency information that corresponds to potential artifacts in the image.

A channel attention mechanism (Woo et al. 2018) is then employed to refine the focus on these high-frequency details. An MLP processes the high-frequency emphasized features F_{HF} to compute channel-wise attention scores:

$$A_k = \text{Sigmoid}(\text{MLP}_\theta(F_{\text{HF}})) \quad (9)$$

The attention scores $A_k \in \mathbb{R}^{C \times 1 \times 1}$ serve to modulate the significance of each channel in the original feature map F_k based on their relevance to high-frequency features.

Mask Prediction We apply the attention scores A_k back to the original feature map to enhance channels identified as significant, resulting in a refined feature map:

$$F'_k = A_k \odot F_k \quad (10)$$

The refined feature map $F'_k \in \mathbb{R}^{C \times H' \times W'}$, is then fed into a U-Net architecture to predict a mask M'_k :

$$M'_k = \text{UNet}_\theta(F'_k) \quad (11)$$

Subsequently, the mask M'_k is resized to $M_k \in \mathbb{R}^{1 \times H \times W}$, aligning with the dimensions of the original image I_k . The mask M_k aims to highlight the locations and extents of high-frequency artifacts, which is crucial for excluding erroneous pixels that could degrade the quality of reconstruction through optimizing photometric residuals.

Optimization DeRainGS is trained in an end-to-end manner. The rainy image enhancement network is frozen during the training. We optimize all parameters of Gaussians \mathcal{G} , the feature extractor Enc_θ , the channel attention layer MLP_θ , and the mask predictor UNet_θ . The training objective for the 3DGS scene representation is defined as:

$$\mathcal{L}_C = (1 - \lambda_{\text{SSIM}})(1 - M) \odot \|\hat{C} - C\| \quad (12)$$

$$+ \lambda_{\text{SSIM}}(1 - M) \odot \mathcal{L}_{\text{SSIM}}(\hat{C}, C) \quad (13)$$

where the mask M is integrated within the photometric loss to specifically exclude pixels affected by rain. The SSIM loss is crucial for improving the structural accuracy of the rendered scenes, which are broadly affected by the rain effect. The final objective function can be formalized as:

$$\mathcal{L} = \mathcal{L}_C + \lambda_{\text{reg}} \mathcal{L}_{\text{reg}} \quad (14)$$

where $\mathcal{L}_{\text{reg}} = \sum_i^H \sum_j^W M_{ij}^2$ is the mask regularization loss to prevent the predicted masks from excessively masking everything. λ_{SSIM} and λ_{reg} are weighting hyperparameters.

Experiments

Datasets We evaluate our method on the waterdrop scenes from DerainNeRF (Li et al. 2024) and our HydroViews dataset, where each synthesized scene features three distinct patterns of raindrops and rain streaks. We evaluate all methods across three patterns and report the averaged results.

Metrics To assess the quality of the reconstruction, we adhere to common metrics including PSNR, SSIM, and LPIPS by comparing the rendered scenes with the clean images.

Baselines We compare our method with deraining and in-the-wild reconstruction methods, including NeRF-W (Martin-Brualla et al. 2021), DerainNeRF (Li et al. 2024), GS-W (Zhang et al. 2024) and WildGaussians (Kulhanek et al. 2024). In the case of DerainNeRF, which employs an AttGAN (Qian et al. 2018) derainer, the weights of the original network are not publicly available. To ensure a fair comparison, we retrained the AttGAN network on the same datasets that were used to train our rainy image enhancement network. Best results are shaded as **first**, **second**, and **third** for each metric.

Implementation Details We implement DeRainGS in PyTorch using Adam optimizer with learning rates of 1.6e-4, 5e-4, and 2.5e-3 for \mathcal{G} 's means, scaling, and SH features, and 1e-3 for tuning occlusion masking modules. We minimize \mathcal{L} over 35,000 iterations. All experiments were conducted on a single NVIDIA A100-80GB GPU. Additional implementation details are provided in the supplementary material.

Raindrops and Rain Streaks

Table 1 presents the quantitative comparison of DeRainGS with baseline methods on our HydroViews dataset. In challenging raindrop scenarios, such as scenes containing heavy occlusions like bicycle and garden, DeRainGS benefits from the effective rainy image enhancement strategy and consistently delivers high-quality reconstructions, surpassing baseline methods that struggle with mist on the lens. For scenes




Figure 5: The visualization compares the rendering outcomes of DeRainGS with baseline methods on selected scenes featuring raindrops (top two rows) or rain streaks (bottom row) of our HydroViews dataset.

Scene	NeRF-W			DerainNeRF			WildGaussians			GS-W			Ours			
	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	
Raindrop	Bicycle	12.88	0.317	0.622	12.72	0.367	0.595	12.25	0.269	0.664	13.27	0.416	0.507	22.18	0.687	0.307
	Bonsai	17.16	0.554	0.458	17.66	0.606	0.422	18.88	0.713	0.270	18.78	0.749	0.238	27.41	0.877	0.173
	Counter	19.23	0.617	0.414	20.24	0.701	0.358	22.22	0.741	0.281	22.51	0.769	0.239	28.83	0.894	0.137
	Garden	8.43	0.282	0.667	8.52	0.354	0.629	10.73	0.396	0.510	10.11	0.417	0.468	22.89	0.692	0.278
	Kitchen	16.54	0.522	0.421	17.88	0.598	0.409	17.02	0.729	0.208	<u>16.73</u>	0.751	0.204	26.27	0.854	0.141
	Room	18.97	0.640	0.395	15.53	0.645	0.355	23.21	0.754	0.324	22.37	0.760	0.278	31.28	0.899	0.180
	Stump	18.54	0.368	0.617	18.68	0.408	0.596	<u>18.31</u>	0.468	0.501	20.46	0.520	0.487	28.69	0.841	0.200
Rain streak	Bicycle	17.80	0.386	0.575	13.24	0.286	0.665	16.48	0.417	0.504	18.20	0.459	0.489	26.52	0.834	0.220
	Garden	19.78	0.474	0.474	15.25	0.393	0.566	18.55	0.604	0.329	19.89	0.621	0.338	25.42	0.705	0.216
	Kitchen	19.46	0.622	0.373	17.35	0.602	0.364	20.12	0.750	0.230	20.39	0.783	0.228	27.40	0.800	0.164
	Stump	17.60	0.372	0.586	14.19	0.318	0.554	16.60	0.411	0.564	17.87	0.423	0.530	27.11	0.773	0.234
	Horse	17.76	0.756	0.375	14.97	0.578	0.516	17.85	0.788	0.337	19.32	0.789	0.328	27.20	0.889	0.194
	Train	17.42	0.607	0.456	14.39	0.512	0.542	17.20	0.618	0.449	19.18	0.643	0.414	24.74	0.818	0.264
	Truck	19.34	0.700	0.442	16.36	0.609	0.503	19.89	0.747	0.360	20.05	0.727	0.378	26.40	0.895	0.216

Table 1: The quantitative evaluation of our method and baseline approaches on the HydroViews dataset. Results are averaged over three different types of raindrops or rain streaks within each scene. An underline indicates one or more failures occurred in the reconstruction process among three trials of different rain effects for a scene.

affected by rain streaks, DeRainGS delivers even higher performance in scene recovery. In contrast, DerainNeRF, which lacks occlusion handling, suffers from inconsistent geometry and performs even worse than NeRF-W. It is important to note that the exceptional performance of DeRainGS is not primarily derived from the high-fidelity rendering capabilities of 3DGS, as the performance of the 3DGS baselines is on par with that of the NeRF approaches. This indi-

cates that the enhancements and masking modules specific to our method are crucial in achieving superior results. Figure 5 visualizes the reconstruction outcomes, demonstrating that our method achieves rendering quality close to clean scenes. In contrast, the baseline methods struggle with distortion and occlusion induced by rainstreak and mist. Furthermore, Table 2 presents the rendering evaluation on DerainNeRF’s dataset. DeRainGS continues to exhibit state-of-

	Scene	NeRF-W			DerainNeRF			WildGaussians			GS-W			Ours		
		PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS	PSNR	SSIM	LPIPS
Waterdrop	Church	29.76	0.934	0.071	30.10	0.968	0.043	32.88	0.975	0.070	31.68	0.970	0.068	33.23	0.975	0.063
	Factory	29.03	0.896	0.078	30.85	0.923	0.056	31.70	0.965	0.081	31.25	0.958	0.076	32.67	0.971	0.074
	Tanabata	25.58	0.873	0.121	26.04	0.887	0.108	33.58	0.975	0.077	33.16	0.973	0.071	34.62	0.980	0.068
	Doraemon	27.83	0.894	0.117	28.74	0.915	0.102	33.11	0.957	0.085	32.46	0.951	0.092	33.94	0.967	0.076

Table 2: The quantitative evaluation of our method and baseline approaches on the DerainNeRF dataset (Li et al. 2024).

the-art performance compared to the NeRF and 3DGS-based methods. However, since this dataset predominantly features lightweight waterdrops, the differences in performance between each method are marginal. This underscores the need to construct more realistic datasets for benchmarking.

Real-world Rainy Environments

To demonstrate the generalization of DeRainGS, Figure 6 compares our method with baseline approaches on selected real-world scenes from our HydroViews dataset. These real-world scenes present significant challenges posed by irregularly shaped rain streaks. DeRainGS effectively mitigates the rain effect, surpassing both NeRF and 3DGS-based methods, which struggle with accurate scene reconstruction and tackling complex rain effects.



Figure 6: The visualization compares the rendering quality of DeRainGS and baselines on real-world scenes.

Ablation Study

We conduct ablation studies on both the raindrop and rain streak scenes for each component of DeRainGS and present the average results for quantitative evaluation. As indicated in Table 3, we find that utilizing either the enhancement process as in (d) or the occlusion masking strategy as in (c) can individually improve the reconstruction quality to a certain degree. Moving to a full implementation as in (f) can thereby significantly enhance performance when compared with vanilla 3DGS. When comparing (c) and (f), we notice that the frequency-based channel attention module also plays a key role in effectively predicting rain occlusions. To straightforwardly show each component’s contribution, Figure 7 illustrates the rendering outcomes of configurations (b)

	DeRainGS’s Components				Metrics
	3DGS	Image Enh.	Mask Mod.	Chan. Attn.	PSNR / SSIM
(b)	✓				17.56 / 0.659
(c)	✓		✓	✓	19.89 / 0.726
(d)	✓	✓			20.42 / 0.721
(e)	✓	✓	✓		23.25 / 0.770
(f)	✓	✓	✓	✓	24.67 / 0.782

Table 3: Ablation study of each component of DeRainGS.

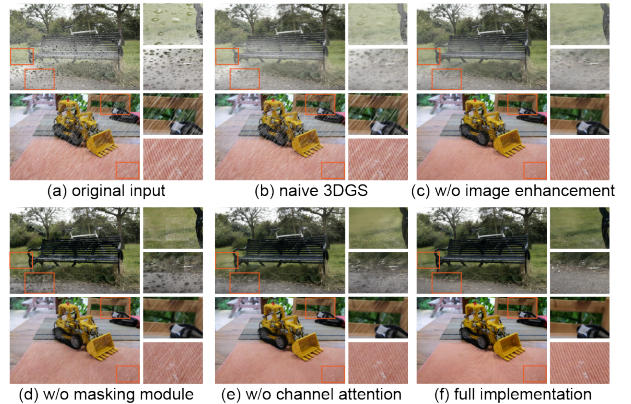


Figure 7: Visualization of rendering quality for ablations.

to (f). We observe that the enhancement process effectively mitigates mist caused by raindrops in (d), while the masking strategy with the channel attention module can adeptly handle high-frequency rain streaks, as shown in (c).

Conclusion

In this study, we introduce 3DRRE, a new task aimed at 3D scene reconstruction under adverse rainy conditions. To support this initiative, we developed the HydroViews dataset, featuring both synthesized and real-world scenes impacted by rain. Building upon this foundation, we propose DeRainGS, the first 3DGS-based method tailored specifically for reconstruction in rainy environments. DeRainGS significantly outperforms existing occlusion-free baselines under various rainy conditions. Together, these contributions enhance the capability to reconstruct high-quality 3D scenes in rain-impacted environments and pave the way for further research in weather-adaptive 3D reconstruction techniques.

References

- Barron, J. T.; Mildenhall, B.; Verbin, D.; Srinivasan, P. P.; and Hedman, P. 2022. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5470–5479.
- Chang, W.; Chen, H.; He, X.; Chen, X.; and Shen, L. 2024. UAV-Rain1k: A Benchmark for Raindrop Removal from UAV Aerial Imagery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15–22.
- Chen, H.; Chen, X.; Wu, C.; Zheng, Z.; Pan, J.; and Fu, X. 2024. Towards Ultra-High-Definition Image Deraining: A Benchmark and An Efficient Method. *arXiv preprint arXiv:2405.17074*.
- Chen, S.; Ye, T.; Bai, J.; Chen, E.; Shi, J.; and Zhu, L. 2023a. Sparse sampling transformer with uncertainty-driven ranking for unified removal of raindrops and rain streaks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 13106–13117.
- Chen, X.; Li, H.; Li, M.; and Pan, J. 2023b. Learning a sparse transformer network for effective image deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5896–5905.
- Chen, X.; Pan, J.; and Dong, J. 2024. Bidirectional multi-scale implicit neural representations for image deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 25627–25636.
- Chen, X.; Pan, J.; Dong, J.; and Tang, J. 2023c. Towards unified deep image deraining: A survey and a new benchmark. *arXiv preprint arXiv:2310.03535*.
- Chen, X.; Pan, J.; Jiang, K.; Li, Y.; Huang, Y.; Kong, C.; Dai, L.; and Fan, Z. 2022a. Unpaired deep image deraining using dual contrastive learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2017–2026.
- Chen, X.; Zhang, Q.; Li, X.; Chen, Y.; Feng, Y.; Wang, X.; and Wang, J. 2022b. Hallucinated neural radiance fields in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12943–12952.
- Cho, J.; Kim, S.; Min, D.; and Sohn, K. 2020. Single image deraining using time-lapse data. *IEEE Transactions on Image Processing*, 29: 7274–7289.
- Cui, Z.; Gu, L.; Sun, X.; Ma, X.; Qiao, Y.; and Harada, T. 2024. Aleth-nerf: Illumination adaptive nerf with concealing field assumption. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 1435–1444.
- Deng, T.; Liu, S.; Wang, X.; Liu, Y.; Wang, D.; and Chen, W. 2023. Prosgnerf: Progressive dynamic neural scene graph with frequency modulated auto-encoder in urban scenes. *arXiv preprint arXiv:2312.09076*.
- Deng, T.; Zhou, Y.; Wu, W.; Li, M.; Huang, J.; Liu, S.; Song, Y.; Zuo, H.; Wang, Y.; Yue, Y.; et al. 2024. Multi-Modal UAV Detection, Classification and Tracking Algorithm—Technical Report for CVPR 2024 UG2 Challenge. *arXiv preprint arXiv:2405.16464*.
- Fumagalli, G.; Huber, Y.; Dymczyk, M.; Siegwart, R.; and Dubé, R. 2021. Fast Image-Anomaly Mitigation for Autonomous Mobile Robots. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2436–2442. IEEE.
- Hnewa, M.; and Radha, H. 2020. Object detection under rainy conditions for autonomous vehicles: A review of state-of-the-art and emerging techniques. *IEEE Signal Processing Magazine*, 38(1): 53–67.
- Jiang, T.-X.; Huang, T.-Z.; Zhao, X.-L.; Deng, L.-J.; and Wang, Y. 2018. Fastderain: A novel video rain streak removal method using directional gradient priors. *IEEE Transactions on Image Processing*, 28(4): 2089–2102.
- Kerbl, B.; Kopanas, G.; Leimkühler, T.; and Drettakis, G. 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 42(4): 139–1.
- Knapitsch, A.; Park, J.; Zhou, Q.-Y.; and Koltun, V. 2017. Tanks and Temples: Benchmarking Large-Scale Scene Reconstruction. *ACM Transactions on Graphics*, 36(4).
- Kulhanek, J.; Peng, S.; Kukulova, Z.; Pollefeys, M.; and Sattler, T. 2024. WildGaussians: 3D Gaussian Splatting in the Wild. *arXiv preprint arXiv:2407.08447*.
- Lee, B.; Lee, H.; Sun, X.; Ali, U.; and Park, E. 2025. Deblurring 3d gaussian splatting. In *European Conference on Computer Vision*, 127–143. Springer.
- Li, Y.; Wu, J.; Zhao, L.; and Liu, P. 2024. DerainNeRF: 3D Scene Estimation with Adhesive Waterdrop Removal. *arXiv preprint arXiv:2403.20013*.
- Lyu, X.; Liu, H.; and Hou, J. 2024. Rainyscape: Unsuper-vised rainy scene reconstruction using decoupled neural rendering. In *Proceedings of the 32nd ACM International Conference on Multimedia*, 10920–10929.
- Martin-Brualla, R.; Radwan, N.; Sajjadi, M. S.; Barron, J. T.; Dosovitskiy, A.; and Duckworth, D. 2021. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7210–7219.
- Mildenhall, B.; Srinivasan, P. P.; Tancik, M.; Barron, J. T.; Ramamoorthi, R.; and Ng, R. 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1): 99–106.
- Oquab, M.; Darcet, T.; Moutakanni, T.; Vo, H.; Szafraniec, M.; Khalidov, V.; Fernandez, P.; Haziza, D.; Massa, F.; El-Nouby, A.; et al. 2023. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*.
- Qian, R.; Tan, R. T.; Yang, W.; Su, J.; and Liu, J. 2018. Attentive generative adversarial network for raindrop removal from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2482–2491.
- Ren, D.; Zuo, W.; Hu, Q.; Zhu, P.; and Meng, D. 2019. Progressive image deraining networks: A better and simpler baseline. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3937–3946.
- Ren, W.; Zhu, Z.; Sun, B.; Chen, J.; Pollefeys, M.; and Peng, S. 2024. NeRF On-the-go: Exploiting Uncertainty for Distractor-free NeRFs in the Wild. In *Proceedings of*

- the *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8931–8940.
- Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5–9, 2015, proceedings, part III 18*, 234–241. Springer.
- Sabour, S.; Goli, L.; Kopanas, G.; Matthews, M.; Lagun, D.; Guibas, L.; Jacobson, A.; Fleet, D. J.; and Tagliasacchi, A. 2024. SpotlessSplats: Ignoring Distractors in 3D Gaussian Splatting. *arXiv preprint arXiv:2406.20055*.
- Wang, T.; Yang, X.; Xu, K.; Chen, S.; Zhang, Q.; and Lau, R. W. 2019. Spatial attentive single-image deraining with a high quality real rain dataset. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 12270–12279.
- Wang, Y.; Wang, J.; and Qi, Y. 2024. WE-GS: An In-the-wild Efficient 3D Gaussian Representation for Unconstrained Photo Collections. *arXiv preprint arXiv:2406.02407*.
- Wei, W.; Yi, L.; Xie, Q.; Zhao, Q.; Meng, D.; and Xu, Z. 2017. Should we encode rain streaks in video as deterministic or stochastic? In *Proceedings of the IEEE International Conference on Computer Vision*, 2516–2525.
- Wei, Y.; Zhang, Z.; Wang, Y.; Xu, M.; Yang, Y.; Yan, S.; and Wang, M. 2021. Deraincyclegan: Rain attentive cyclegan for single image deraining and rainmaking. *IEEE Transactions on Image Processing*, 30: 4788–4801.
- Woo, S.; Park, J.; Lee, J.-Y.; and Kweon, I. S. 2018. Cbam: Convolutional block attention module. In *Proceedings of the European conference on computer vision (ECCV)*, 3–19.
- Xiao, J.; Fu, X.; Liu, A.; Wu, F.; and Zha, Z.-J. 2022. Image de-raining transformer. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(11): 12978–12995.
- Xu, J.; Mei, Y.; and Patel, V. M. 2024. Wild-GS: Real-Time Novel View Synthesis from Unconstrained Photo Collections. *arXiv preprint arXiv:2406.10373*.
- Yang, W.; Tan, R. T.; Feng, J.; Guo, Z.; Yan, S.; and Liu, J. 2019. Joint rain detection and removal from a single image with contextualized deep networks. *IEEE transactions on pattern analysis and machine intelligence*, 42(6): 1377–1393.
- Yang, Y.; Zhang, S.; Huang, Z.; Zhang, Y.; and Tan, M. 2023. Cross-ray neural radiance fields for novel-view synthesis from unconstrained image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 15901–15911.
- Yue, Z.; Xie, J.; Zhao, Q.; and Meng, D. 2021. Semi-supervised video deraining with dynamical rain generator. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 642–652.
- Zhang, D.; Wang, C.; Wang, W.; Li, P.; Qin, M.; and Wang, H. 2024. Gaussian in the Wild: 3D Gaussian Splatting for Unconstrained Image Collections. *arXiv preprint arXiv:2403.15704*.
- Zhu, C.; Wan, R.; Tang, Y.; and Shi, B. 2023. Occlusion-free scene recovery via neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20722–20731.