

MeteGS: Meteorology-Guided Gaussian Splatting for Scene Rendering and Recovery in Adverse Weather Conditions

Sha Fan¹, Xinhua Shan¹, Mingyu Liang¹, Ningjie Bao¹, Wei Liu² and Ying Fu¹

¹Chengdu University of Information Technology

²Shanghai Jiao Tong University

3240608055@stu.cuit.edu.cn, 3240608029@stu.cuit.edu.cn, 1215062016@qq.com,
1742422074@qq.com, weiluicv@sjtu.edu.cn, fuying@cuit.edu.cn*

Abstract

3D Gaussian Splatting enables efficient, high-fidelity novel view synthesis with explicit Gaussians and differentiable rendering. However, adverse weather introduces rain streaks and droplets as well as volumetric scattering, producing view-dependent, spatially varying degradations that break the clean multi-view consistency assumption and lead to geometric drift and unstable appearance. Existing approaches are largely confined to the 2D image domain and seldom model the 3D degradation formation process, limiting controllable weather synthesis and consistent restoration in 3DGS. To our knowledge, we are the first to incorporate real meteorological observations as an external prior and propose a 3D-level unified rendering–restoration framework for joint deraining and dehazing within a single 3DGS pipeline. Precipitation is mapped to interpretable parameters controlling the density and morphology of rain Gaussians, while an atmospheric-scattering extinction strength drives 3D fog generation for reproducible weather modeling. We adopt a closed-loop dual-branch optimization where the rendering branch fits degraded observations to capture weather degradation patterns, and the restoration branch regularizes scene Gaussians toward a clean domain with multi-scale perceptual consistency, suppressing artifacts and improving cross-view detail fidelity. Experiments across diverse scenes and weather conditions consistently outperform strong baselines on all three metrics, with a particularly notable PSNR improvement of 1.78 dB on average.

1 Introduction

In practical applications such as autonomous driving, mobile robotics, virtual reality, and digital twins, rainy and foggy weather introduces rain streaks, lens-adhered droplets that cause localized blur and refraction, and volumetric scattering that reduces contrast and compresses the effective visibility range. Consequently, observed images suffer substantial degradation in structural detail, texture discriminability, and color consistency. More importantly, these degrada-

tions are often spatially non-uniform and vary across view-points, which undermines the cross-view consistency constraints fundamental to 3D representation learning and novel view synthesis. This results in suppressed geometric details and unstable appearance, ultimately limiting the reliability of 3D systems under adverse weather.

Photorealistic reconstruction and controllable rendering in complex weather are essential for real-world deployment. However, existing rain and fog modeling largely relies on heuristic 2D effects or simplified volumetric approximations, with limited systematic coupling to real meteorological observations. As a result, it remains difficult to simultaneously achieve physical plausibility, multi-view consistency, and interpretable, controllable parameterization. With the success of 3DGS [Kerbl *et al.*, 2023] in efficient and high-fidelity real-time rendering, a key question arises, namely how to introduce weather effects driven by meteorological priors into 3DGS [Kerbl *et al.*, 2023] representations while effectively enforcing perceptual quality during learning. Meanwhile, prior deraining and dehazing research has primarily focused on the 2D image domain [Chen *et al.*, 2025; Li *et al.*, 2019; Zamir *et al.*, 2022]. Although strong performance has been reported on single-frame metrics, these methods typically do not explicitly model multi-view geometric consistency or cross-view radiometric consistency, and thus tend to produce view-dependent artifacts and inconsistent results in multi-view reconstruction and neural rendering. In addition, many 3D representations assume stable imaging conditions [Mildenhall *et al.*, 2021; Martin-Brualla *et al.*, 2021], and applying per-view 2D restoration as a preprocessing step often propagates restoration errors into the 3D model and amplifies cross-view inconsistency. Therefore, robust 3D modeling in rainy and foggy conditions requires unified correction directly in the 3D representation space rather than frame-wise 2D postprocessing.

To address this challenge, we propose a restoration framework that performs joint deraining and dehazing directly in the 3D representation domain. Specifically, we propose a dedicated weather module within 3DGS [Kerbl *et al.*, 2023] that jointly models rain streaks, lens-droplet-induced local blur, and atmospheric scattering, enabling geometrically consistent 3D weather synthesis and rendering. On the meteorological side, we establish a continuous mapping from real precipitation measurements to rendering parameters, provid-

ing interpretable and controllable weather strength. On the restoration side, we leverage a clean 3D scene prior by predicting and correcting residual updates of Gaussian parameters, and impose multi-scale perceptual consistency on rendered outputs. This yields a restored 3D scene with strong cross-view consistency, detail preservation, and stable appearance, while avoiding the view inconsistencies introduced by per-view 2D restoration.

Our main contributions are summarized as follows:

- We propose a real meteorological data-driven parameter mapping module that maps meteorological measurements to interpretable weather parameters, providing precise and controllable inputs for 3D weather rendering.
- We develop a weather rendering model based on 3D Gaussian representations that unifies rain streaks, droplet-induced blur, and atmospheric scattering within a single 3D Gaussian field and supports adjustable weather strength.
- We present an end-to-end 3D reconstruction framework for adverse rain-fog conditions that incorporates multi-scale perceptual constraints into 3D Gaussian optimization and refines geometry and appearance via residual correction and cross-scale consistency.
- We conduct a comprehensive evaluation across diverse scenes and weather settings, demonstrating consistent gains over prior baselines in both quantitative metrics and visual quality.

2 Related Work

2.1 3D Reconstruction

Neural Radiance Fields (NeRF) [Mildenhall *et al.*, 2021] achieve multi-view novel view synthesis via differentiable volumetric rendering, yet their volume integration and dense sampling typically incur high training cost and limited real-time performance [Müller *et al.*, 2022; Fridovich-Keil *et al.*, 2022]. In contrast, 3D Gaussian Splatting (3DGS) [Kerbl *et al.*, 2023] represents a scene as a set of 3D Gaussians with anisotropic covariance and appearance parameters, and leverages differentiable splatting to strike a favorable balance between rendering quality and efficiency, making it a representative approach for real-time novel view synthesis. Existing 3DGS [Kerbl *et al.*, 2023] studies mainly focus on geometric and appearance modeling, extensions to dynamic scenes [Wu *et al.*, 2024], and editability [Chen *et al.*, 2024], and are commonly built upon the assumption of stable, clean observations without significant atmospheric degradation. Under adverse weather with time-varying degradations, this assumption is violated, and optimization can be driven by corrupted signals, resulting in geometric drift and unstable appearance with degraded cross-view detail fidelity. Overall, these methods largely adopt a forward synthesis paradigm, focusing on controllable weather generation and photorealistic 3D rendering.

2.2 3D Synthesis of Rain and Fog

Conventional rain and fog synthesis is mostly conducted in the 2D image domain by overlaying rain streaks or noise, or by applying simplified scattering approximations to generate degraded samples [Garg and Nayar, 2006; Garg and Nayar, 2007; Narasimhan and Nayar, 2003; Narasimhan and Nayar, 2002]. While simple and efficient, such strategies are decoupled from true 3D geometry and light transport, making it difficult to guarantee multi-view consistency. To improve consistency, recent work has incorporated particle distributions, volumetric fog fields, or depth-guided scattering into 3D rendering pipelines [Narasimhan and Nayar, 2003; Nishita *et al.*, 1993]. In addition, visibility-oriented scattering formulations are commonly used to relate attenuation strength to a physically meaningful notion of visibility range [Lee and Shang, 2016]. Some studies further integrate physically motivated weather effects into 3D Gaussian representations for controllable and scalable weather rendering [Kerbl *et al.*, 2023; Dai *et al.*, 2025]. Overall, these methods emphasize interpretable degradation generation and efficient photorealistic rendering.

2.3 3D Scene Restoration under Adverse Weather

3D scene restoration under adverse weather has mainly progressed along two lines. One representative direction, largely built on NeRF, explicitly incorporates atmospheric effects or occlusion processes into differentiable rendering and jointly optimizes degradation formation together with a clean radiance field, enabling integrated restoration and novel view synthesis from degraded multi-view observations while maintaining geometric consistency [Mildenhall *et al.*, 2021; Li *et al.*, 2024]. With the popularity of 3D Gaussian Splatting (3DGS) [Kerbl *et al.*, 2023], another direction enhances robustness by suppressing transient artifacts during 3DGS optimization, often via visibility reasoning or mask-guided fitting, to improve stability and reconstruction quality in unconstrained captures [Martin-Brualla *et al.*, 2021; Chen *et al.*, 2022; Yang *et al.*, 2023; Zhang *et al.*, 2024; Sabour *et al.*, 2025; Kulhanek *et al.*, 2024]. Recently, weather-oriented 3DGS frameworks further explore reconstructing clean scenes from adverse-weather observations by explicitly handling precipitation-induced artifacts and lens effects [Qian *et al.*, 2025; Liu *et al.*, 2025]. In parallel, extensive progress has been made in 2D deraining and adverse weather removal [Ren *et al.*, 2019; Xiao *et al.*, 2022; Li *et al.*, 2020; Valanarasu *et al.*, 2022], but these per-view approaches typically do not explicitly enforce multi-view geometric and radiometric consistency, which can lead to view-dependent artifacts when integrated into multi-view reconstruction pipelines. Overall, existing 3D-level methods either co-estimate degradations and clean 3D representations or mitigate transient corruptions within 3DGS [Kerbl *et al.*, 2023], yet they still leave open the need for a unified 3D-level restoration framework that directly recovers a clean, cross-view-consistent scene representation from degraded observations.




Figure 1: **Overview of the MeteGS framework.** We adopt 3DGS as a differentiable rendering backbone and augment it with a meteorology-driven weather Gaussian module. Meteorological information (Rain) is aggregated and mapped to the precipitation rate P (mm/h), which is then used to generate rain-streak Gaussians and fog Gaussians. The weather Gaussians and scene Gaussians are jointly rendered in a differentiable manner to synthesize adverse-weather observations. Meanwhile, a frozen feature extractor together with training objectives imposes constraints on the rendered results, and gradients are backpropagated to optimize trainable components, enabling unified optimization of weather-induced degradation modeling and scene reconstruction.

3 Method

3.1 Preliminaries

We adopt 3D Gaussian Splatting (3DGS) [Kerbl *et al.*, 2023] as the foundational framework for scene representation and differentiable rendering. The scene is modeled as a set of 3D Gaussian primitives endowed with geometric attributes (e.g., position and shape) as well as opacity and appearance parameters. Given a camera pose, these 3D Gaussians are projected onto the image plane and rendered via differentiable splatting and alpha compositing to produce the final image. Since the entire rendering process is differentiable with respect to the Gaussian parameters, the model can be optimized end-to-end by backpropagating gradients from the discrepancy between the rendered images and the corresponding real observations, thereby updating the Gaussian parameters. As shown in Fig. 1, we further augment this differentiable rendering pipeline with a meteorology-driven weather Gaussian module, which generates weather components such as rain streaks and fog conditioned on meteorological cues and jointly renders them with the scene Gaussians to match adverse-weather observations. Meanwhile, additional constraints encourage the scene representation to converge toward the clean domain, enabling unified modeling of weather-induced degradation and scene reconstruction within a single framework.

3.2 Meteorology-Driven Weather Gaussian

To enable fine-grained, meteorology-driven control of rendering parameters, we design a mapping module grounded in physically motivated empirical formulations. The input to this module is the precipitation rate P (mm/h), derived from accumulated rainfall amounts obtained from meteorological services, either from observations or forecasts. When

multiple measurements are available within a temporal window, we aggregate them using $\text{Agg}(\cdot) \in \{\text{mean}, \text{max}, \text{min}\}$ to characterize the precipitation intensity at the current time, and normalize the result into a bounded scalar:

$$s_t = \left[\frac{\text{Agg}(P_t) - P_{\min}}{P_{\max} - P_{\min}} \right]_0^1, \quad \text{Agg} \in \{\text{mean}, \text{max}, \text{min}\} \quad (1)$$

Here, P_t denotes the precipitation rate (mm/h) at time t from meteorological observations or forecasts, and $\text{Agg}(\cdot)$ optionally aggregates values (mean/max/min) to represent the current intensity. Here, $P_{\min}$ and $P_{\max}$ are the normalization bounds, and $[\cdot]_0^1$ clips to $[0, 1]$. The normalized factor s_t is then mapped to weather variables that control the generation and rendering of rain and fog Gaussians. By making weather parameters functions of meteorological inputs, we maintain a physically meaningful link between measured intensity and synthesized degradation, avoiding arbitrary aesthetic tuning.

We denote the meteorology-driven parameter mapping as:

$$\theta(P) = \{\theta_{\text{rain}}, \theta_{\text{fog}}\} \quad (2)$$

Here, $\theta(P)$ denotes the weather-parameter mapping function (or the resulting mapped parameters) driven by the precipitation rate P . θ_{rain} represents the parameter set of the Rain Gaussians, and θ_{fog} represents the parameter set of the Fog Gaussians. We adopt the precipitation-intensity categories provided by the World Meteorological Organization (WMO) [Jarraud, 2008] as an external prior, and construct a monotonic nonlinear mapping with respect to P such that stronger precipitation yields progressively stronger synthesized weather effects. In the following, we elaborate the specific parameterization and formal definitions of θ_{rain} and θ_{fog} , respectively.

Concretely, the number of Rain Gaussians is written as:

$$N_{\text{rain}} = \text{clamp}(\text{base} + k \cdot P^\gamma, 0, N_{\text{max}}), \quad \gamma \in (0, 1] \quad (3)$$

Here, N_{rain} is the number of Rain Gaussians controlling rain density, and P is the meteorological precipitation rate. base provides a baseline rain level, k scales the response to precipitation, and $\gamma \in (0, 1]$ introduces a sub-linear nonlinearity to avoid excessive growth under heavy rain. The $\text{clamp}(\cdot, 0, N_{\text{max}})$ operator enforces the valid range $[0, N_{\text{max}}]$, where N_{max} limits the maximum number of primitives for stable and efficient rendering.

We parameterize fog by θ_{fog} , taking inspiration from the classical atmospheric scattering model [Narasimhan and Nayar, 2002]:

$$I(d) = J(d)t(d) + A(1 - t(d)), \quad t(d) = \exp(-\beta d) \quad (4)$$

where $I(d)$ denotes the observed foggy radiance at depth d , $J(d)$ is the fog-free radiance, $t(d)$ is the transmission, A is the atmospheric light, and β is the extinction coefficient governing attenuation [Bohren and Huffman, 2008]. Rather than explicitly enforcing this image-plane formation model, we use β as a physically grounded control signal within θ_{fog} to condition the density and opacity of 3D fog Gaussians, producing progressively stronger fog effects as β increases. In summary, our meteorology-driven design provides interpretable and reproducible control signals for weather Gaussian generation and rendering.

3.3 Joint Optimization Framework

We adopt a closed-loop joint optimization scheme composed of a weather rendering branch and a clean restoration branch, which enables us to explicitly model weather degradations while enforcing clean-scene consistency in the 3D domain. Let the scene be represented by a set of 3D Gaussians, and let the weather effects be represented by a meteorology-driven weather Gaussian, where θ denotes meteorology-conditioned weather parameters controlling both rain and fog effects. Given a camera viewpoint, we construct two parallel forward branches.

Weather rendering branch. This branch feeds the shared scene 3D Gaussians and the meteorology-driven weather Gaussians into the renderer, producing a weather-degraded synthetic image through standard 3DGS [Kerbl *et al.*, 2023] rasterization and alpha compositing:

$$\hat{I}_w = R(G_s \cup G_w(\theta), c) \quad (5)$$

Here, $R(\cdot, c)$ denotes the differentiable 3DGS [Kerbl *et al.*, 2023] renderer under camera c , G_s is the set of shared scene Gaussians, and $G_w(\theta)$ is the set of meteorology-driven weather Gaussians parameterized by θ . The union $G_s \cup G_w(\theta)$ indicates that weather primitives are composited with the scene representation during rendering to synthesize the adverse-weather image $\hat{I}_w$.

If the training data provide the corresponding real weather observation I_{obs} for the same viewpoint, we supervise the realism of the synthesized weather appearance via a photometric consistency constraint:

$$\mathcal{L}_{\text{photo}} = (1 - \lambda) \left\| \hat{I}_w - I_{\text{obs}} \right\|_1 + \lambda \left(1 - \text{SSIM}(\hat{I}_w, I_{\text{obs}}) \right) \quad (6)$$

where I_{obs} is the real adverse-weather observation from the same viewpoint, $\|\cdot\|_1$ measures pixel-wise differences, and $\lambda \in [0, 1]$ balances the two terms. Minimizing $\mathcal{L}_{\text{photo}}$ enforces $\hat{I}_w$ to match I_{obs} and provides gradients to the shared scene Gaussians for explicitly modeling weather effects.

Clean restoration branch. In the same iteration, we render a clean image by disabling the weather effects and using only the shared scene Gaussians. This branch enforces clean-domain regularization on the shared 3D representation, encouraging it to capture stable geometry and intrinsic appearance while separating out transient weather-induced degradations. During optimization, it counteracts potential drifts introduced by the weather rendering branch, thereby improving cross-view consistency and stabilizing novel-view synthesis.

We further incorporate a multi-scale perceptual consistency constraint by aligning structural and textural statistics in feature space to suppress over-smoothing and improve detail fidelity. We therefore employ a pretrained feature extractor with frozen parameters [Simonyan and Zisserman, 2014] to avoid training drift. The clean-branch output $\hat{I}_{\text{clean}}$ and the corresponding clean reference image I_{gt} are fed into the extractor, and features are collected from multiple intermediate layers. We compute feature discrepancies and aggregate them across layers to obtain the perceptual loss [Johnson *et al.*, 2016]:

$$\mathcal{L}_{\text{perc}} = \sum_l \left\| \Phi_l(\hat{I}_{\text{clean}}) - \Phi_l(I_{\text{gt}}) \right\|_1 \quad (7)$$

Here, $\hat{I}_{\text{clean}}$ denotes the rendered clean image, I_{gt} is the corresponding clean reference, $\Phi_l(\cdot)$ is the feature map from the l -th frozen layer, and $\|\cdot\|_1$ measures element-wise differences. This loss aligns local contrast, edge shapes, and texture distributions across scales, enabling the clean branch not only to approach the reference at the pixel level but also to preserve textural details, which improves detail fidelity and cross-view appearance stability. Moreover, since deep features are sensitive to texture and noise statistics relevant to restoration, this constraint helps suppress the leakage of residual rain-fog artifacts into the clean rendering. By combining the closed-loop dual-branch mechanism with perceptual consistency, we unify explicit degradation accounting and stable clean-scene recovery under a single optimization objective.

The overall training loss is formulated as a weighted combination of a photometric consistency term from the weather rendering branch, a perceptual consistency term from the clean restoration branch, and an optional clean-domain photometric supervision term when clean references are available:

$$\mathcal{L} = \lambda_{\text{photo}} \mathcal{L}_{\text{photo}} + \lambda_{\text{perc}} \mathcal{L}_{\text{perc}} + \lambda_{\text{clean}} \mathcal{L}_{\text{clean}} \quad (8)$$

Here, The scalars λ_{photo} , λ_{perc} , and λ_{clean} weight the relative contributions of these terms. The weather-branch photometric term encourages the synthesized weather appearance to match real degradations and provides optimization signals to the shared 3D scene representation under explicit weather-conditioned rendering. In contrast, the perceptual and clean supervision terms regularize the shared scene representation toward the clean domain, promoting stable geometry and




Figure 2: Qualitative results on datasets under rainy (top) and foggy (bottom) conditions.

intrinsic appearance. Through end-to-end backpropagation, these losses jointly optimize the shared scene representation, and weather effects are injected as explicit 3D components conditioned on meteorological cues during rendering.

Overall, the weather branch strengthens the ability to account for degraded observations, while the clean branch strengthens the recovery of intrinsic scene content. The two branches mutually balance each other, enabling cross-view consistent 3D-level deraining/dehazing and detail restoration, and ultimately producing a high-quality clean 3D scene representation for novel view synthesis and downstream tasks.

3.4 Depth-aware Pruning and Point Compensation

We adopt the depth-aware pruning and point compensation strategy from Deblur-3DGS [Lee *et al.*, 2024] and adapt it to our dual-branch closed-loop framework. For stable structural optimization, densification and pruning statistics are computed only from the shared scene Gaussians (e.g., visibility and projected radius), while weather Gaussians are excluded to prevent transient effects from biasing decisions. In addition, we perform point compensation in sparse regions at early stages by inheriting appearance and local features from

neighboring Gaussians, improving the reconstruction of small objects and thin structures and accelerating convergence.

4 Experiments

4.1 Experimental Setup

Datasets. To evaluate the effectiveness of 3D reconstruction under adverse weather, we construct a novel, diverse, and challenging benchmark. Following the base scenes provided by Deblur-NeRF [Ma *et al.*, 2022], we build rainy and foggy inputs for four static multi-view scenes (*factory*, *wine*, *pool*, and *tanabata*). We further incorporate real precipitation observations from rainfall records of Chinese meteorological stations as the meteorological input for driving the generation of weather parameters. More details about the meteorological data are provided in the supplementary material. Given the clean images and the collected meteorological parameters, we synthesize corresponding rain, fog, and combined rain-fog degraded observations using the weather modeling in Section 3.2, which serve as the degraded inputs for training and testing on the synthetic benchmark. In addition, we evaluate our method on the real-world HydroViews [Liu *et al.*, 2025]

Model	factory			wine			pool			tanabata		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
NeRF	15.305	0.276	0.543	16.415	0.351	0.265	23.324	0.671	0.283	21.842	0.586	0.328
DerainNeRF	17.473	0.301	0.592	23.177	0.712	0.300	20.483	0.586	0.394	23.760	<u>0.731</u>	0.462
3DGS	22.285	0.673	0.304	24.468	0.806	0.241	<u>26.874</u>	0.764	0.252	23.173	0.675	0.327
Restormer+3DGS	22.712	0.635	0.379	<u>24.851</u>	0.605	0.262	25.132	0.713	0.263	24.327	0.526	0.284
GS-W	<u>24.548</u>	0.682	0.309	22.104	0.750	0.242	24.526	0.762	0.201	22.628	0.637	0.392
WeatherGS	24.218	<u>0.740</u>	<u>0.235</u>	23.646	0.745	<u>0.208</u>	25.405	<u>0.817</u>	<u>0.198</u>	<u>24.351</u>	0.632	0.215
Ours	26.037	0.786	0.213	25.738	<u>0.790</u>	0.179	27.042	0.865	0.182	25.829	0.782	<u>0.232</u>

Table 1: Comparison of PSNR, SSIM, and LPIPS for Different Models on Rainy Scenes

Model	factory			wine			pool			tanabata		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
NeRF	17.060	0.659	0.345	18.751	0.584	<u>0.251</u>	20.263	0.721	0.283	21.307	0.643	0.262
3DGS	20.226	0.725	0.274	21.478	0.682	0.391	22.057	<u>0.741</u>	0.266	<u>24.472</u>	0.801	<u>0.205</u>
DehazeFormer+3DGS	<u>24.795</u>	0.667	0.303	22.574	<u>0.705</u>	0.263	<u>23.568</u>	0.730	<u>0.198</u>	22.196	0.761	0.291
GS-W	23.684	<u>0.783</u>	<u>0.217</u>	<u>23.362</u>	0.643	0.328	21.848	0.692	0.259	22.341	0.735	0.248
Ours	25.524	0.812	0.206	24.571	0.776	0.224	24.915	0.746	0.192	25.228	<u>0.794</u>	0.203

Table 2: Comparison of PSNR, SSIM, and LPIPS for Different Models on Fog Scenes

dataset for qualitative comparison. For HydroViews [Liu *et al.*, 2025], we do not apply any weather rendering. Instead, we directly restore the captured rainy or foggy scenes and compare the results under real conditions.

Implementation details. All methods are compared under the same train and test split and the same input view configuration. Our reconstruction backbone is built upon the 3DGS [Kerbl *et al.*, 2023] framework and implemented in PyTorch. We further employ a lightweight learnable MLP (Gaussian Transformation Network), whose main body consists of three fully connected layers with width 64 and ReLU activations, followed by several linear prediction heads. Its input is formed by concatenating the positional encodings of 3D position and viewing direction (with frequencies 3 and 10, respectively) with scale and rotation features. All experiments are conducted on a single RTX 3090 GPU for 30,000 training iterations, and all input images are resized to a resolution of 600×400 . Unless otherwise specified, we set the loss weights to $\lambda_{\text{photo}} = 1.0$, $\lambda_{\text{perc}} = 0.05$, and $\lambda_{\text{clean}} = 0.1$.

Comparison and metrics. To evaluate visual quality, we compare renderings produced by our method with those from state-of-the-art 3D reconstruction approaches, including NeRF [Mildenhall *et al.*, 2021], DerainNeRF [Li *et al.*, 2024], 3DGS [Kerbl *et al.*, 2023], GS-W [Zhang *et al.*, 2024], and WeatherGS [Qian *et al.*, 2025]. We also consider two-stage pipelines that combine 2D restoration with 3D reconstruction, using Restormer+3DGS for rainy scenes [Zamir *et al.*, 2022; Kerbl *et al.*, 2023] and DehazeFormer+3DGS for foggy scenes [Song *et al.*, 2023; Kerbl *et al.*, 2023]. For quantitative evaluation, we benchmark all methods using three standard metrics, namely peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), and learned perceptual image patch similarity (LPIPS).



Figure 3: Qualitative Comparison on Real-World Rainy Dataset.

4.2 Qualitative Analysis

To analyze differences in fine-structure recovery and appearance stability, we present representative qualitative comparisons on rainy and foggy scenes in Fig. 2. NeRF [Mildenhall *et al.*, 2021] and 3DGS [Kerbl *et al.*, 2023] often exhibit geometric drift, unstable textures, or residual artifacts under degraded observations. While two-stage pipelines (e.g., Restormer [Zamir *et al.*, 2022] and DehazeFormer [Song *et al.*, 2023] combined with 3DGS [Kerbl *et al.*, 2023]) can alleviate some degradations, per-frame restoration errors tend to accumulate across views, leading to cross-view-inconsistent “floating” textures. In contrast, our method explicitly models weather degradations in the 3D representation space and imposes clean-domain regularization, enabling more faithful recovery of structural edges and material textures, while maintaining consistent color and contrast across viewpoints.

To further assess generalization in real-world conditions, Fig. 3 reports qualitative results on captured rainy scenes. These examples exhibit challenging degradations such as dense, non-uniform rain streaks and local blur, which often lead to texture instability and residual artifacts for the 3DGS [Kerbl *et al.*, 2023] baseline. In comparison, our method produces cleaner reconstructions with sharper structures and more stable textures, and the zoomed-in regions highlight its improved recovery of fine details.

4.3 Quantitative Analysis

Table 1 reports quantitative results on rainy scenes. Overall, our method achieves the most consistent performance across all four scenes and attains the best or second-best results under all metrics, demonstrating robust recovery of geometry and appearance under rain streaks and droplet-induced blur. Compared with the strongest 3D-level baseline, WeatherGS [Qian *et al.*, 2025], our approach provides more stable improvements across scenes. We also observe scene-dependent trade-offs in prior methods, 3DGS [Kerbl *et al.*, 2023] attains relatively higher SSIM on *wine* but falls behind in overall quality, whereas WeatherGS yields lower LPIPS on *tanabata* at the cost of reduced SSIM. These results suggest that suppressing transient artifacts alone is insufficient; in contrast, our method better balances structural fidelity and texture recovery, leading to more reliable restoration under rainy conditions. Table 2 reports quantitative comparisons

	3DGS	Weather Branch	Clean Sup	Perceptual	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
(a)	✓	×	×	×	22.287	0.675	0.312
(b)	✓	✓	×	×	23.094	0.702	0.281
(c)	✓	✓	✓	×	23.716	0.719	0.241
(d)	✓	✓	×	✓	24.625	0.774	0.219
(e)	✓	✓	✓	✓	25.513	0.823	0.125

Table 3: Ablation on core modules of MeteGS .

	Depth-based pruning	Add extra points	Metrics		
			PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
(a)	×	×	20.971	0.636	0.269
(b)	×	✓	22.609	0.679	0.211
(c)	✓	×	22.335	0.651	0.213
(d)	✓	✓	25.479	0.791	0.138

Table 4: Ablation on geometric refinement strategies

on foggy scenes. Overall, our method consistently ranks best across the four scenes under standard criteria, demonstrating stronger perceptual quality and more faithful reconstruction than prior approaches. On average, our method improves PSNR by 1.78 dB over the strongest baseline and reduces LPIPS by 0.057, with particularly notable gains in perceptual quality. These results suggest that explicitly introducing extinction-driven 3D fog modeling better explains foggy observations and helps disentangle intrinsic scene appearance from scattering-induced degradations. Moreover, the perceptual consistency constraint effectively alleviates typical post-dehazing artifacts, such as over-smoothing and residual haze veils, leading to more stable textures and cleaner novel-view renderings.

4.4 Ablation Studies

We further conduct ablation studies to quantify the contribution of each component (Table 3). Overall, enabling the Weather Branch already improves reconstruction under degraded observations. Adding the Perceptual constraint further enhances texture and structural fidelity, while introducing

Clean Sup brings more pronounced gains in overall restoration quality. With all components enabled, the full model achieves the best performance. Fig. 4 also shows that it most effectively suppresses residual rain-fog artifacts and recovers finer local details with better appearance consistency.

Table 4 further analyzes the impact of depth-guided pruning and extra point sampling on optimization stability. Both strategies yield stable benefits, extra point sampling alleviates under-representation caused by sparse initialization, whereas depth-guided pruning reduces noisy gradients from degraded observations via more reliable visibility filtering. Combining them performs best, suggesting that jointly improving visibility filtering and sample supplementation effectively mitigates erroneous gradients under adverse weather and leads to higher-quality reconstructions. Due to space limitations, more experimental results are provided in the supplementary material.




Figure 4: **Qualitative ablation of MeteGS core modules.** From left to right, the figure shows (a) the weather branch only ($\mathcal{L}_{\text{photo}}$), (b) with the perceptual constraint ($\mathcal{L}_{\text{perc}}$), (c) with clean supervision ($\mathcal{L}_{\text{clean}}$), (d) the full model, and (e) the ground-truth clean image. The perceptual constraint improves texture and edge fidelity, clean supervision enhances detail recovery, and the full model most effectively suppresses residual weather artifacts.

5 Conclusion

We present a meteorology-driven unified 3D Gaussian rendering–restoration framework for robust 3D modeling and novel view synthesis under rain and fog. By mapping precipitation observations to interpretable rain parameters and estimating an extinction-strength cue for fog generation, our method enables controllable and reproducible weather modeling. We further introduce a dual-branch closed-loop optimization in which a weather rendering branch fits degraded observations to account for adverse-weather effects, while a restoration branch regularizes the shared scene Gaussians toward intrinsic clean appearance via clean-domain supervision and multi-scale perceptual consistency, yielding cross-view-consistent deraining and dehazing with improved detail fidelity. Extensive experiments demonstrate consistent gains on both rainy and foggy scenes, producing fewer residual artifacts and more stable geometry and textures, and further validate meteorological priors as effective and interpretable conditional control signals. Future work will incorporate richer meteorological factors, extend the framework to temporally consistent and dynamically updated 3D representations, and integrate it with downstream applications such as autonomous driving, robotics, and digital twins.

Acknowledgements

This work was supported in part by NSFC under Grant 62402318 and Grant 24Z990200676 and in part by Artificial Intelligence (AI) for Science Seed Program of Shanghai Jiao Tong University under Grant WH410261901/002.

References

- [Bohren and Huffman, 2008] Craig F Bohren and Donald R Huffman. *Absorption and scattering of light by small particles*. John Wiley & Sons, 2008.
- [Chen *et al.*, 2022] Xingyu Chen, Qi Zhang, Xiaoyu Li, Yue Chen, Ying Feng, Xuan Wang, and Jue Wang. Hallucinated neural radiance fields in the wild. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12943–12952, 2022.
- [Chen *et al.*, 2024] Yiwen Chen, Zilong Chen, Chi Zhang, Feng Wang, Xiaofeng Yang, Yikai Wang, Zhongang Cai, Lei Yang, Huaping Liu, and Guosheng Lin. Gaussianeditor: Swift and controllable 3d editing with gaussian splatting. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 21476–21485, 2024.
- [Chen *et al.*, 2025] Xiang Chen, Jinshan Pan, Jiangxin Dong, and Jinhui Tang. Towards unified deep image deraining: A survey and a new benchmark. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.
- [Dai *et al.*, 2025] Qiyu Dai, Xingyu Ni, Qianfan Shen, Wenzheng Chen, Baoquan Chen, and Mengyu Chu. Rainygs: Efficient rain synthesis with physically-based gaussian splatting. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 16153–16162, 2025.
- [Fridovich-Keil *et al.*, 2022] Sara Fridovich-Keil, Alex Yu, Matthew Tancik, Qinhong Chen, Benjamin Recht, and Angjoo Kanazawa. Plenoxels: Radiance fields without neural networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5501–5510, 2022.
- [Garg and Nayar, 2006] Kshitiz Garg and Shree K Nayar. Photorealistic rendering of rain streaks. *ACM Transactions on Graphics (TOG)*, 25(3):996–1002, 2006.
- [Garg and Nayar, 2007] Kshitiz Garg and Shree K Nayar. Vision and rain. *International Journal of Computer Vision*, 75(1):3–27, 2007.
- [Jarraud, 2008] M Jarraud. Guide to meteorological instruments and methods of observation (wmo-no. 8). *World Meteorological Organisation: Geneva, Switzerland*, 29, 2008.
- [Johnson *et al.*, 2016] Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European conference on computer vision*, pages 694–711. Springer, 2016.
- [Kerbl *et al.*, 2023] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023.
- [Kulhanek *et al.*, 2024] Jonas Kulhanek, Songyou Peng, Zuzana Kukelova, Marc Pollefeys, and Torsten Sattler. Wildgaussians: 3d gaussian splatting in the wild. *arXiv preprint arXiv:2407.08447*, 2024.
- [Lee and Shang, 2016] Zhongping Lee and Shaoling Shang. Visibility: how applicable is the century-old koschmieder model? *Journal of the Atmospheric Sciences*, 73(11):4573–4581, 2016.
- [Lee *et al.*, 2024] Byeonghyeon Lee, Howoong Lee, Xiangyu Sun, Usman Ali, and Eunbyung Park. Deblurring 3d gaussian splatting. In *European Conference on Computer Vision*, pages 127–143. Springer, 2024.
- [Li *et al.*, 2019] Siyuan Li, Iago Breno Araujo, Wenqi Ren, Zhangyang Wang, Eric K Tokuda, Roberto Hirata Junior, Roberto Cesar-Junior, Jiawan Zhang, Xiaojie Guo, and Xiaochun Cao. Single image deraining: A comprehensive benchmark analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3838–3847, 2019.
- [Li *et al.*, 2020] Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. All in one bad weather removal using architectural search. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3175–3185, 2020.
- [Li *et al.*, 2024] Yunhao Li, Jing Wu, Lingzhe Zhao, and Peidong Liu. Derainnerf: 3d scene estimation with adhesive waterdrop removal. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 2787–2793. IEEE, 2024.
- [Liu *et al.*, 2025] Shuhong Liu, Xiang Chen, Hongming Chen, Quanfeng Xu, and Mingrui Li. Deraings: Gaussian splatting for enhanced scene reconstruction in rainy environments. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pages 5558–5566, 2025.
- [Ma *et al.*, 2022] Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V Sander. Deblur-nerf: Neural radiance fields from blurry images. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12861–12870, 2022.
- [Martin-Brualla *et al.*, 2021] Ricardo Martin-Brualla, Noha Radwan, Mehdi SM Sajjadi, Jonathan T Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 7210–7219, 2021.
- [Mildenhall *et al.*, 2021] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.
- [Müller *et al.*, 2022] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM transactions on graphics (TOG)*, 41(4):1–15, 2022.

- [Narasimhan and Nayar, 2002] Srinivasa G Narasimhan and Shree K Nayar. Vision and the atmosphere. *International journal of computer vision*, 48(3):233–254, 2002.
- [Narasimhan and Nayar, 2003] Srinivasa G. Narasimhan and Shree K. Nayar. Contrast restoration of weather degraded images. *IEEE transactions on pattern analysis and machine intelligence*, 25(6):713–724, 2003.
- [Nishita et al., 1993] Tomoyuki Nishita, Takao Sirai, Katsumi Tadamura, and Eihachiro Nakamae. Display of the earth taking into account atmospheric scattering. In *Proceedings of the 20th annual conference on Computer graphics and interactive techniques*, pages 175–182, 1993.
- [Qian et al., 2025] Chenghao Qian, Yuhu Guo, Wenjing Li, and Gustav Markkula. Weathers: 3d scene reconstruction in adverse weather conditions via gaussian splatting. In *2025 IEEE International Conference on Robotics and Automation (ICRA)*, pages 185–191. IEEE, 2025.
- [Ren et al., 2019] Dongwei Ren, Wangmeng Zuo, Qinghua Hu, Pengfei Zhu, and Deyu Meng. Progressive image de-raining networks: A better and simpler baseline. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3937–3946, 2019.
- [Sabour et al., 2025] Sara Sabour, Lily Goli, George Kopanas, Mark Matthews, Dmitry Lagun, Leonidas Guibas, Alec Jacobson, David Fleet, and Andrea Tagliasacchi. Spotlessplats: Ignoring distractors in 3d gaussian splatting. *ACM Transactions on Graphics*, 44(2):1–11, 2025.
- [Simonyan and Zisserman, 2014] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [Song et al., 2023] Yuda Song, Zhuqing He, Hui Qian, and Xin Du. Vision transformers for single image dehazing. *IEEE Transactions on Image Processing*, 32:1927–1941, 2023.
- [Valanarasu et al., 2022] Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M Patel. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2353–2363, 2022.
- [Wu et al., 2024] Guanjun Wu, Taoran Yi, Jiemin Fang, Lingxi Xie, Xiaopeng Zhang, Wei Wei, Wenyu Liu, Qi Tian, and Xinggang Wang. 4d gaussian splatting for real-time dynamic scene rendering. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 20310–20320, 2024.
- [Xiao et al., 2022] Jie Xiao, Xueyang Fu, Aiping Liu, Feng Wu, and Zheng-Jun Zha. Image de-raining transformer. *IEEE transactions on pattern analysis and machine intelligence*, 45(11):12978–12995, 2022.
- [Yang et al., 2023] Yifan Yang, Shuhai Zhang, Zixiong Huang, Yubing Zhang, and Mingkui Tan. Cross-ray neural radiance fields for novel-view synthesis from unconstrained image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15901–15911, 2023.
- [Zamir et al., 2022] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022.
- [Zhang et al., 2024] Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d gaussian splatting for unconstrained image collections. In *European Conference on Computer Vision*, pages 341–359. Springer, 2024.