

# NimbusGS: Unified 3D Scene Reconstruction under Hybrid Weather – Supplementary Materials –

Yanying Li<sup>1</sup>, Jinyang Li<sup>1</sup>, Shengfeng He<sup>2</sup>, Yangyang Xu<sup>3</sup>, Junyu Dong<sup>1</sup>, Yong Du<sup>1,4,2\*</sup>

<sup>1</sup>School of Computer Science and Technology, Ocean University of China,

<sup>2</sup>Singapore Management University, <sup>3</sup>Harbin Institute of Technology (Shenzhen),

<sup>4</sup>Sanya Oceanographic Institution, Ocean University of China

## 1. Additional Implementation Details

### 1.1. Training Configuration

During the geometry initialization stage, the learning rate of the Gaussian scale parameters is reduced to  $2 \times 10^{-3}$  for stable early optimization, while the other learning rate settings follow 3DGS. In the joint optimization stage, all Gaussian attributes revert to the default 3DGS settings.

For the extinction field  $\beta$ , we build an axis-aligned bounding box (AABB) from the Gaussian centers and enlarge it by  $2\times$  along all axes so that the entire scene and all valid ray segments fall within the volume. A uniform voxel grid of size  $128^3$  is constructed inside this AABB, with each voxel initialized as a single scalar drawn from a Gaussian distribution (mean = 0, std = 0.01). The extinction field and the scattering-color MLP are optimized with learning rates of  $5 \times 10^{-3}$  and  $5 \times 10^{-4}$ , respectively.

Details regarding ray sampling and volumetric rendering are provided in the main paper and omitted here for brevity.

### 1.2. Adverse-weather Degradation Synthesis

Since no public dataset provides multi-view scenes with controllable hybrid adverse-weather effects, we generate the required degradations using the composite model of OneRestore [4]. Clean images are taken from the outdoor scenes of Mip-NeRF360 [1] and Deblur-NeRF [10]. Haze (H) is synthesized using the standard atmospheric scattering model, where the extinction coefficient and airlight intensity are randomly sampled within the ranges defined by the OneRestore implementation. Depth maps are predicted by the pretrained Depth-Anything v2 [13] network and used to compute transmission. Rain and snow components follow the public OneRestore pipeline: rain streaks are sampled from RainStreakGen [3], and snow masks are taken from Snow100K [8] and blended with the clean images. This procedure yields six degradation types: haze (H), snow (S), rain+snow (R+S), rain+haze (R+H), snow+haze (S+H), and

---

### Algorithm 1: Training of NimbusGS

---

**Input:** scene images  $\{I_{in}\}$ , initialized 3D Gaussians  $\mathcal{G}$ , sampling count  $K$ , reference radius  $r_0$ , loss weights  $\lambda_r$ , iteration numbers  $M_{init}$  (geometry initialization) and  $M_{joint}$  (joint optimization), residual refresh interval  $Z_{ref}$ .  
**Output:** optimized Gaussians  $\mathcal{G}$ , extinction field  $\beta$ , and particulate layer  $R$ .

```
// Stage 1: Geometry Initialization
1 for  $i = 1$  to  $M_{init}$  do
2   Render continuous-medium output  $I_{con}$ ;
3   Update model parameters using  $\mathcal{L}_{ini}(I_{in}, I_{con})$ ;
4 end
5 Compute initial particulate layer:  $R = \text{ReLU}(I_{in} - I_{con})$ ;

// Stage 2: Joint Optimization
6 for  $j = 1$  to  $M_{joint}$  do
7   Render degraded output  $I_{deg}$  using CSM and PLM;
8   Update  $\mathcal{G}$  with  $\mathcal{L}$  and GGS;
9   Update extinction field  $\beta$  and the airlight module by  $\mathcal{L}$ ;
10  if  $j \bmod Z_{ref} = 0$  then
11    Recompute  $I_{con}$  from current  $\mathcal{G}$  and medium components;
12    Refresh particulate layer:  $R = \text{ReLU}(I_{in} - I_{con})$ ;
13  end
14 end
```

---

rain+snow+haze (R+S+H).

Note that real-world haze and snow datasets do not provide multi-view scenes or camera poses and are therefore unsuitable for our 3D reconstruction setting, so these weather effects are synthesized. In contrast, the rain (R) setting is directly taken from the multi-view synthetic rainy scenes in the Rainyscape dataset rather than generated with OneRestore.

\*Corresponding author (csyongdu@ouc.edu.cn).




Figure 1. Qualitative ablation of Particulate Layer Modeling. Best viewed zoomed in.




Figure 2. Qualitative ablation of Continuous Scattering Modeling.

### 1.3. Training Pipeline

We provide the complete training workflow in Algorithm 1. It offers a compact procedural description for implementing NimbusGS.

## 2. Additional Ablation Studies

### 2.1. Continuous Scattering Modeling

Fig. 2 visualizes the ablation of our Continuous Scattering Modeling (CSM) under the haze setting, which offers a clean environment for examining the extinction field. We omit the Uniform Extinction variant since its spatially constant  $\beta$  contains no meaningful structure. For the variants that yield spatial variation, we compare the  $\beta$  estimated by CSM (CSM- $\beta$ ) with the Gaussian-Driven Extinction baseline (GD- $\beta$ ), where  $\beta$  is directly optimized as Gaussian parameters, and we also show their corresponding reconstruction results (CSM-Recon and GD-Recon). Each  $\beta$  map is normalized independently for visualization because the two approaches follow fundamentally different numerical distributions.

The GD- $\beta$  map shows strong geometry dependence, with high values clustering near bright surfaces and object boundaries and unstable fluctuations in distant regions. This indicates entanglement between  $\beta$  and surface radiance, resulting in incomplete haze removal and geometry-coupled artifacts in the GD-Recon output.

In contrast, CSM- $\beta$  exhibits smooth and depth-coherent variation and effectively avoids texture leakage. Foreground regions maintain low extinction, while background areas

progressively accumulate higher values consistent with volumetric scattering. The resulting CSM-Recon preserves sharper structures and yields more stable far-field appearance, demonstrating that CSM produces a physically meaningful and depth-coherent extinction field.

### 2.2. Particulate Layer Modeling

Fig. 1 provides qualitative examples of PLM. We visualize on the H+R setting, where particle effects are more clearly distinguishable.

Without PLM (w/o PLM), the model is forced to explain both continuous-medium effects and particle streaks using the same volumetric formulation. This entangles the two degradation types and hampers the estimation of the extinction field, leading to incomplete haze removal and leaving noticeable rain streaks in the final reconstruction.

We further evaluate the Restored-Derived Residuals (RDR) alternative (w/ RDR), where an all-in-one restoration method (OneRestore) produces  $I_{\text{restore}}$  as an estimate of  $I_{\text{con}}$ , and the residual is computed as  $I_{\text{deg}} - I_{\text{restore}}$ . Such all-in-one models tend to suppress multiple degradation types simultaneously, introducing heterogeneous artifacts into the residual. As 2D restorers, they may further produce view-dependent color variations that propagate into the 3D reconstruction and appear as color shifts or overly smoothed regions.

NimbusGS instead renders a multi-view-consistent  $I_{\text{con}}$  from the Gaussian field and separates particle effects via PLM, yielding cleaner removal of rain streaks and more stable appearance consistency.

### 2.3. Geometry-Guided Gradient Scaling

We further study the behavior of GGS through qualitative and factor-wise quantitative analyses.

**Qualitative Ablation of GGS.** Fig. 3 shows a comparison on an H+R scene under two settings: w/o GGS and w/ GGS. Without GGS, distant regions suffer from pronounced detail loss and appear overly blurred, mainly due to weakened supervision and the inability of coarse far-field Gaussians to trigger timely densification. With GGS, the depth-, radius-, and error-aware scaling better prioritizes refinement in these challenging areas, yielding sharper geometry and noticeably finer background details.




Figure 3. Qualitative ablation of Geometry-Guided Gradient Scaling. Best viewed zoomed in.

| Variant                  | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
|--------------------------|-----------------|-----------------|--------------------|
| w/o Depth                | 21.98           | 0.731           | 0.214              |
| w/o Projected Radius     | 22.14           | 0.733           | 0.209              |
| w/o Reconstruction Error | 22.09           | 0.729           | 0.211              |
| <b>Ours</b>              | <b>22.25</b>    | <b>0.742</b>    | <b>0.202</b>       |

Table 1. Factor-wise ablation of GGS. Best results are marked in **bold**.

**Factor-wise Quantitative Analysis.** We further conduct a factor-level ablation on four hybrid-weather scenes (Bicycle, Garden, Factory, and Tanabata) under the H+R+S setting. The averaged results are reported in Tab. 1. Removing any single factor from GGS leads to a consistent drop in reconstruction quality, indicating that depth, projected radius, and reconstruction error jointly contribute to effective densification guidance.

Among the three factors, the depth-based component shows the largest impact, as it enhances the supervision of distant structures where weather-induced attenuation makes optimization particularly challenging. The projected-radius component and the reconstruction-error term offer comparable additional gains: the former helps prioritize large Gaussians that dominate far-field regions, while the latter highlights areas that remain under-refined, providing modest but stable improvements across metrics.

## 2.4. Loss Terms

To assess the influence of each loss, we individually remove  $\mathcal{L}_{DCP}$  and  $\mathcal{L}_{TV}$ , extend  $\mathcal{L}_{DCP}$  to the second stage, and evaluate these variants on four hybrid-weather scenes (Bicycle, Garden, Factory, and Tanabata) rendered under the H+R+S setting. The averaged results are reported in Tab. 2. Removing  $\mathcal{L}_{DCP}$  produces the most significant drop in performance. Although the model can still separate clean

| Variant                          | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
|----------------------------------|-----------------|-----------------|--------------------|
| w/o $\mathcal{L}_{DCP}$          | 17.32           | 0.610           | 0.233              |
| w/ $\mathcal{L}_{DCP}$ (stage 2) | 20.92           | 0.718           | 0.220              |
| w/o $\mathcal{L}_{TV}$           | 22.17           | 0.735           | 0.211              |
| <b>Ours</b>                      | <b>22.25</b>    | <b>0.742</b>    | <b>0.202</b>       |

Table 2. Ablation of loss terms. Best results are marked in **bold**.

appearance from weather scattering through its underlying structure, the absence of this prior weakens this decomposition and leads to noticeably lower reconstruction accuracy across all metrics. However, when  $\mathcal{L}_{DCP}$  is kept in the second stage, the reconstruction quality also degrades, typically producing darker outputs. This behavior is consistent with the bias of the dark channel prior, which tends to favor locally dark pixels and may therefore drive the optimization toward overly dark solutions in later training. Removing  $\mathcal{L}_{TV}$  results in a moderate decrease, but introduces small spatial irregularities in the extinction field, which degrade reconstruction stability.

## 2.5. Geometry Initialization Schedule

We ablate the geometry-initialization iteration count  $M_{init}$  on two scenes from Mip-NeRF360 (Bicycle and Garden) under the H+R+S setting, where particle-based degradations naturally reveal whether the initialization properly separates static geometry from transient weather artifacts. The averaged results are shown in Tab. 3. Very short schedules (*e.g.*, 1K–2K) do not allow the Gaussian field to establish reliable geometry, causing part of the static structure to leak into the particulate residuals and preventing a clean separation between geometry and particle artifacts. Conversely, excessively long schedules (*e.g.*, 10K–15K) cause the Gaussians to absorb particle patterns as if they were static content, leading to floating artifacts and reducing the

| $M_{\text{init}}$ | Bicycle         |                 |                    | Garden          |                 |                    |
|-------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                   | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 1K                | 17.69           | 0.521           | 0.408              | 20.64           | 0.651           | 0.280              |
| 2K                | 18.01           | 0.573           | 0.345              | 21.75           | 0.679           | 0.256              |
| 4K                | <b>19.74</b>    | <b>0.603</b>    | <b>0.327</b>       | <b>22.12</b>    | <b>0.708</b>    | <b>0.209</b>       |
| 10K               | 19.16           | 0.593           | 0.335              | 22.09           | 0.703           | 0.237              |
| 15K               | 19.24           | 0.591           | 0.334              | 22.07           | 0.702           | 0.241              |

Table 3. Ablation of geometry-initialization iteration count. Best results are marked in **bold**.

| Method             | Params. | Memory | Training | FPS |
|--------------------|---------|--------|----------|-----|
| 3DGS [5]           | 94.1M   | 5.3G   | 29min    | 171 |
| SeaSplat [12]      | 29.4M   | 5.1G   | 126min   | 169 |
| WaterSplatting [6] | 38.9M   | 5.4G   | 19min    | 32  |
| RainyScape [9]     | 86.8M   | 8.3G   | 169min   | 167 |
| RobustSplat [2]    | 81.0M   | 8.9G   | 29min    | 168 |
| WeatherGS [11]     | 108.0M  | 19.9G  | 28min    | 159 |
| Ours               | 78.4M   | 9.3G   | 77min    | 162 |

Table 4. Model complexity and runtime comparison under the H+R+S setting. Params., Memory, Training, and FPS denote parameter count, training memory usage, training time, and rendering speed, respectively.

completeness of the residual map produced by the particulate layer modeling. A setting of 4K iterations achieves the best balance, providing well-formed geometry while keeping particle effects out of the Gaussian representation and consistently delivering the highest reconstruction quality.

## 2.6. Model Complexity and Runtime Analysis

We compare the model complexity, training cost, and rendering efficiency of NimbusGS with several baselines in Tab. 4. Although NimbusGS introduces additional modules beyond the original 3DGS, its overall parameter count is lower, *i.e.*, 78.4M versus 94.1M. Specifically, the CSM module contributes 2.1M parameters and the airlight MLP adds 0.2M, while the PLM component contains no learnable parameters. The remaining parameters are mainly associated with Gaussian primitives. Compared with 3DGS, NimbusGS requires fewer Gaussians because degradations are more effectively removed, reducing the need to compensate for them with additional primitives. In contrast, 3DGS tends to trigger additional densification to fit residual artifacts, leading to more Gaussians and a higher parameter count.

In terms of efficiency, NimbusGS uses 9.3G training memory, requires 77 minutes for training, and achieves 162 FPS rendering. These results indicate that the proposed degradation modeling introduces only moderate computational overhead while maintaining competitive training and rendering efficiency among weather-aware Gaussian splat-

| Method             | H+R             |                 |                    |                 |                 |                    |
|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                    | Stump           |                 |                    | Wine            |                 |                    |
|                    | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 3DGS [5]           | 17.31           | 0.572           | 0.335              | 15.09           | 0.551           | 0.368              |
| SeaSplat [12]      | 16.36           | 0.533           | 0.368              | 12.77           | 0.531           | 0.367              |
| WaterSplatting [6] | 17.49           | 0.538           | 0.340              | 14.45           | 0.518           | 0.305              |
| DerainNeRF [7]     | 17.39           | 0.569           | 0.332              | 15.10           | 0.543           | 0.387              |
| RainyScape [9]     | 17.48           | 0.564           | 0.307              | 15.51           | 0.600           | 0.228              |
| RobustSplat [2]    | 17.29           | 0.586           | 0.344              | 15.53           | 0.629           | 0.315              |
| WeatherGS [11]     | 16.15           | 0.490           | 0.377              | 15.10           | 0.507           | 0.394              |
| OR [4]+3DGS [5]    | 15.50           | 0.510           | 0.351              | 19.64           | 0.624           | 0.279              |
| <b>Ours</b>        | 21.33           | 0.648           | 0.277              | 20.45           | 0.782           | 0.181              |

| Method             | H+S             |                 |                    |                 |                 |                    |
|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                    | Stump           |                 |                    | Wine            |                 |                    |
|                    | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 3DGS [5]           | 16.82           | 0.575           | 0.322              | 15.78           | 0.642           | 0.256              |
| SeaSplat [12]      | 16.08           | 0.498           | 0.417              | 16.15           | 0.617           | 0.264              |
| WaterSplatting [6] | 16.74           | 0.534           | 0.345              | 16.31           | 0.627           | 0.187              |
| DerainNeRF [7]     | 16.00           | 0.499           | 0.401              | 16.03           | 0.630           | 0.294              |
| RainyScape [9]     | 17.02           | 0.564           | 0.309              | 15.12           | 0.594           | 0.222              |
| RobustSplat [2]    | 17.07           | 0.597           | 0.342              | 16.10           | 0.668           | 0.214              |
| WeatherGS [11]     | 17.10           | 0.578           | 0.376              | 17.56           | 0.631           | 0.281              |
| OR [4]+3DGS [5]    | 15.00           | 0.529           | 0.318              | 20.29           | 0.747           | 0.169              |
| <b>Ours</b>        | 20.27           | 0.657           | 0.257              | 20.89           | 0.753           | 0.152              |

| Method          | R+S             |                 |                    |                 |                 |                    |
|-----------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                 | Stump           |                 |                    | Wine            |                 |                    |
|                 | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 3DGS [5]        | 19.87           | 0.579           | 0.348              | 21.59           | 0.698           | 0.292              |
| DerainNeRF [7]  | 19.95           | 0.574           | 0.335              | 21.62           | 0.701           | 0.284              |
| RainyScape [9]  | 20.07           | 0.571           | 0.316              | 22.35           | 0.738           | 0.123              |
| RobustSplat [2] | 19.90           | 0.588           | 0.329              | 21.58           | 0.702           | 0.289              |
| WeatherGS [11]  | 21.12           | 0.523           | 0.348              | 21.41           | 0.674           | 0.297              |
| OR [4]+3DGS [5] | 15.17           | 0.536           | 0.334              | 22.33           | 0.756           | 0.229              |
| <b>Ours</b>     | 21.63           | 0.600           | 0.303              | 23.27           | 0.811           | 0.142              |

Table 5. Quantitative comparisons on the additional Stump and Wine scenes. The three blocks correspond to the H+R, H+S, and R+S settings. The **best** and **second-best** scores are color-encoded for clarity.

ting methods.

## 3. Additional Weather-Effects Evaluation

This section extends the hybrid-weather evaluation beyond the configurations considered in the main paper. For clarity, the main paper adopts a task-aligned comparison protocol. Single-weather methods are evaluated only on the corresponding single-weather tasks, and methods designed for hybrid-weather scenarios are evaluated on hybrid-weather settings. This ensures a fair and meaningful comparison within the scope of each task. In the supplementary material, we relax this restriction and include all baselines that are able to generate images under the corresponding degradations. This broader comparison allows a more comprehensive examination of model behavior under mixed adverse conditions.

We include three additional hybrid compositions, H+R, H+S, and R+S, to more thoroughly stress-test NimbusGS

| Method             | Bicycle         |                 |                    | Garden          |                 |                    | Factory         |                 |                    | Tanabata        |                 |                    | Average         |                 |                    |
|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                    | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 3DGS [5]           | 15.24           | 0.491           | 0.370              | 15.95           | 0.530           | 0.376              | 13.90           | 0.534           | 0.406              | 13.58           | 0.445           | 0.428              | 14.66           | 0.500           | 0.395              |
| SeaSplat [12]      | 15.48           | 0.523           | 0.354              | 13.56           | 0.558           | 0.359              | 14.64           | 0.621           | 0.329              | 14.58           | 0.564           | 0.324              | 14.56           | 0.566           | 0.341              |
| WaterSplatting [6] | 16.69           | 0.537           | 0.360              | 17.42           | 0.592           | 0.312              | 14.71           | 0.602           | 0.307              | 14.45           | 0.518           | 0.309              | 15.81           | 0.562           | 0.322              |
| DerainNeRF [7]     | 16.10           | 0.499           | 0.401              | 16.27           | 0.556           | 0.350              | 14.29           | 0.583           | 0.372              | 14.31           | 0.503           | 0.407              | 15.24           | 0.535           | 0.382              |
| RainyScape [9]     | 16.66           | 0.504           | 0.387              | 18.22           | 0.653           | 0.267              | 14.75           | 0.642           | 0.268              | 14.65           | 0.527           | 0.306              | 16.07           | 0.581           | 0.307              |
| RobustSplat [2]    | 16.28           | 0.498           | 0.379              | 17.43           | 0.603           | 0.322              | 13.97           | 0.542           | 0.404              | 13.63           | 0.443           | 0.429              | 15.32           | 0.521           | 0.383              |
| WeatherGS [11]     | 15.94           | 0.479           | 0.414              | 17.56           | 0.509           | 0.436              | 14.62           | 0.569           | 0.395              | 14.60           | 0.533           | 0.405              | 15.68           | 0.522           | 0.412              |
| OR [4]+3DGS [5]    | 17.01           | 0.552           | 0.353              | 20.28           | 0.636           | 0.311              | 18.43           | 0.698           | 0.300              | 20.11           | 0.637           | 0.338              | 18.95           | 0.630           | 0.325              |
| <b>Ours</b>        | 19.71           | 0.651           | 0.321              | 22.96           | 0.727           | 0.200              | 21.87           | 0.836           | 0.151              | 22.17           | 0.722           | 0.229              | 21.67           | 0.734           | 0.225              |

Table 6. Quantitative comparisons on H+R scenes. The best and second-best scores are color-encoded for clarity.

| Method             | Bicycle         |                 |                    | Garden          |                 |                    | Factory         |                 |                    | Tanabata        |                 |                    | Average         |                 |                    |
|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                    | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 3DGS [5]           | 17.19           | 0.554           | 0.327              | 14.21           | 0.507           | 0.411              | 13.87           | 0.601           | 0.282              | 15.16           | 0.615           | 0.261              | 15.10           | 0.569           | 0.320              |
| SeaSplat [12]      | 15.73           | 0.525           | 0.348              | 14.66           | 0.513           | 0.400              | 13.98           | 0.626           | 0.247              | 15.18           | 0.625           | 0.339              | 14.88           | 0.572           | 0.333              |
| WaterSplatting [6] | 16.47           | 0.530           | 0.325              | 14.10           | 0.512           | 0.422              | 14.42           | 0.613           | 0.276              | 15.37           | 0.623           | 0.251              | 15.09           | 0.569           | 0.318              |
| DerainNeRF [7]     | 17.12           | 0.558           | 0.331              | 14.32           | 0.576           | 0.378              | 14.12           | 0.611           | 0.279              | 15.16           | 0.617           | 0.288              | 15.18           | 0.590           | 0.319              |
| RainyScape [9]     | 17.33           | 0.567           | 0.256              | 14.52           | 0.600           | 0.369              | 14.08           | 0.631           | 0.241              | 15.75           | 0.641           | 0.211              | 15.42           | 0.609           | 0.269              |
| RobustSplat [2]    | 17.41           | 0.556           | 0.323              | 14.17           | 0.552           | 0.416              | 13.86           | 0.600           | 0.275              | 15.34           | 0.637           | 0.224              | 15.19           | 0.586           | 0.309              |
| WeatherGS [11]     | 17.92           | 0.579           | 0.312              | 14.34           | 0.566           | 0.419              | 13.82           | 0.571           | 0.407              | 15.57           | 0.536           | 0.339              | 15.41           | 0.563           | 0.369              |
| OR [4]+3DGS [5]    | 16.24           | 0.475           | 0.396              | 23.09           | 0.693           | 0.295              | 21.34           | 0.818           | 0.178              | 23.25           | 0.801           | 0.159              | 20.98           | 0.696           | 0.257              |
| <b>Ours</b>        | 19.11           | 0.625           | 0.251              | 23.15           | 0.709           | 0.232              | 21.48           | 0.839           | 0.136              | 23.50           | 0.809           | 0.116              | 21.81           | 0.745           | 0.183              |

Table 7. Quantitative comparisons on H+S scenes. The best and second-best scores are color-encoded for clarity.

| Method          | Bicycle         |                 |                    | Garden          |                 |                    | Factory         |                 |                    | Tanabata        |                 |                    | Average         |                 |                    |
|-----------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|-----------------|-----------------|--------------------|
|                 | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ | PSNR $\uparrow$ | SSIM $\uparrow$ | LPIPS $\downarrow$ |
| 3DGS [5]        | 19.25           | 0.590           | 0.332              | 23.39           | 0.714           | 0.245              | 18.29           | 0.685           | 0.306              | 22.79           | 0.765           | 0.241              | 20.93           | 0.688           | 0.281              |
| DerainNeRF [7]  | 18.37           | 0.533           | 0.369              | 23.12           | 0.712           | 0.295              | 18.62           | 0.675           | 0.297              | 22.89           | 0.762           | 0.255              | 20.75           | 0.670           | 0.304              |
| RainyScape [9]  | 19.57           | 0.614           | 0.285              | 24.25           | 0.729           | 0.216              | 22.12           | 0.805           | 0.127              | 24.07           | 0.808           | 0.117              | 22.50           | 0.739           | 0.186              |
| RobustSplat [2] | 19.56           | 0.588           | 0.348              | 23.70           | 0.718           | 0.271              | 18.99           | 0.701           | 0.295              | 22.86           | 0.785           | 0.244              | 21.27           | 0.698           | 0.289              |
| WeatherGS [11]  | 17.47           | 0.539           | 0.413              | 20.71           | 0.618           | 0.442              | 19.11           | 0.562           | 0.364              | 22.01           | 0.733           | 0.243              | 19.82           | 0.613           | 0.365              |
| OR [4]+3DGS [5] | 16.91           | 0.533           | 0.368              | 24.02           | 0.713           | 0.252              | 22.13           | 0.757           | 0.219              | 24.35           | 0.814           | 0.177              | 21.85           | 0.704           | 0.254              |
| <b>Ours</b>     | 20.03           | 0.644           | 0.278              | 24.81           | 0.751           | 0.194              | 23.82           | 0.859           | 0.120              | 24.81           | 0.845           | 0.107              | 23.36           | 0.774           | 0.174              |

Table 8. Quantitative comparisons on R+S scenes. The best and second-best scores are color-encoded for clarity.

under mixed adverse conditions that combine volumetric degradations with particle-based artifacts. To further enrich scene diversity under these challenging settings, we also incorporate two additional scenes used exclusively for the supplementary experiments, broadening the spatial layouts and weather interactions while keeping the training and evaluation protocol identical to that in the main benchmark.

### 3.1. Quantitative Results

We first report quantitative results on the four main scenes (Bicycle, Garden, Factory, and Tanabata) under the H+R, H+S, and R+S configurations in Tab. 6, Tab. 7, and Tab. 8. Across all three hybrid-weather settings, NimbusGS consistently achieves the best performance in terms of PSNR, SSIM, and LPIPS. In contrast, directly applying 3DGS or recent weather-aware Gaussian and NeRF variants leads to a substantial quality drop once heterogeneous

haze–rain–snow interactions are present. The single-image baseline OR [4] combined with 3DGS (*i.e.*, OR+3DGS) forms a strong competitor and typically ranks second, but still falls short of NimbusGS by a notable margin, particularly on challenging scenes such as Bicycle and Tanabata where both structural consistency and perceptual quality are severely affected by mixed degradations.

To further assess generalization under different spatial layouts, we additionally evaluate two new scenes, Stump from Mip-NeRF360 [1] and Wine from Deblur-NeRF [10], as shown in Tab. 5. The performance trends match those observed on the main benchmark: NimbusGS maintains the top results across all hybrid-weather combinations, while OR+3DGS remains competitive but continues to underperform our method. These supplementary experiments confirm that NimbusGS scales reliably to diverse scene struc-

tures and more complex mixtures of weather degradations.

### 3.2. Additional Visual Comparisons

We provide additional qualitative comparisons for all seven weather settings, covering both the single-weather tasks and the hybrid-weather tasks. Figs. 4 and 5 show supplementary visual examples for these two groups of settings. Due to space limitations, each figure includes the top-performing baselines for the corresponding task.

Specifically, Fig. 4 presents the single-weather comparisons, where NimbusGS preserves geometry and appearance quality across haze-, rain-, and snow-induced degradations. Fig. 5 shows the results for the four hybrid-weather settings that combine volumetric scattering and particle-based artifacts, demonstrating the effectiveness of NimbusGS under more complex adverse weather conditions.

### 4. Limitation

Despite its effectiveness under adverse weather conditions, NimbusGS may drop in sparse-view settings. This is because removing particulate residuals reveals regions without appearance cues, while the model lacks priors to recover them. These regions then rely on cross-view information, often insufficient with few views, which remains an open challenge for future work.


### References

- [1] Jonathan T Barron, Ben Mildenhall, Dor Verbin, Pratul P Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *CVPR*, pages 5470–5479, 2022. 1, 5
- [2] Chuanyu Fu, Yuqi Zhang, Kunbin Yao, Guanying Chen, Yuan Xiong, Chuan Huang, Shuguang Cui, and Xiaochun Cao. Robustsplat: Decoupling densification and dynamics for transient-free 3dgs. In *ICCV*, 2025. 4, 5
- [3] Kshitiz Garg and Shree K Nayar. Photorealistic rendering of rain streaks. *ACM Transactions on Graphics*, 25(3):996–1002, 2006. 1
- [4] Yu Guo, Yuan Gao, Yuxu Lu, Huilin Zhu, Ryan Wen Liu, and Shengfeng He. Onerestore: A universal restoration framework for composite degradation. In *ECCV*, pages 255–272, 2024. 1, 4, 5
- [5] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM TOG*, 42(4):139–1, 2023. 4, 5
- [6] Huapeng Li, Wenxuan Song, Tianao Xu, Alexandre Elsig, and Jonas Kulhanek. Watersplating: Fast underwater 3d scene reconstruction using gaussian splatting. In *3DV*, pages 969–978, 2025. 4, 5
- [7] Yunhao Li, Jing Wu, Lingzhe Zhao, and Peidong Liu. Derainnerf: 3d scene estimation with adhesive waterdrop removal. In *ICRA*, pages 2787–2793, 2024. 4, 5
- [8] Yun-Fu Liu, Da-Wei Jaw, Shih-Chia Huang, and Jenq-Neng Hwang. Desnownet: Context-aware deep network for snow removal. *IEEE Transactions on Image Processing*, 27(6): 3064–3073, 2018. 1
- [9] Xianqiang Lyu, Hui Liu, and Junhui Hou. Rainyscape: Unsupervised rainy scene reconstruction using decoupled neural rendering. In *ACM MM*, pages 10920–10929, 2024. 4, 5
- [10] 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 *CVPR*, pages 12861–12870, 2022. 1, 5
- [11] Chenghao Qian, Yuhu Guo, Wenjing Li, and Gustav Markkula. Weathersg: 3d scene reconstruction in adverse weather conditions via gaussian splatting. In *ICRA*, pages 185–191, 2025. 4, 5
- [12] Daniel Yang, John J Leonard, and Yogesh Girdhar. Seasplat: Representing underwater scenes with 3d gaussian splatting and a physically grounded image formation model. In *ICRA*, pages 7632–7638, 2025. 4, 5
- [13] Lihe Yang, Bingyi Kang, Zilong Huang, Zhen Zhao, Xiaogang Xu, Jiashi Feng, and Hengshuang Zhao. Depth anything v2. *NeurIPS*, 37:21875–21911, 2024. 1

### Hazy Scenes



### Rainy Scenes



### Snowy Scenes

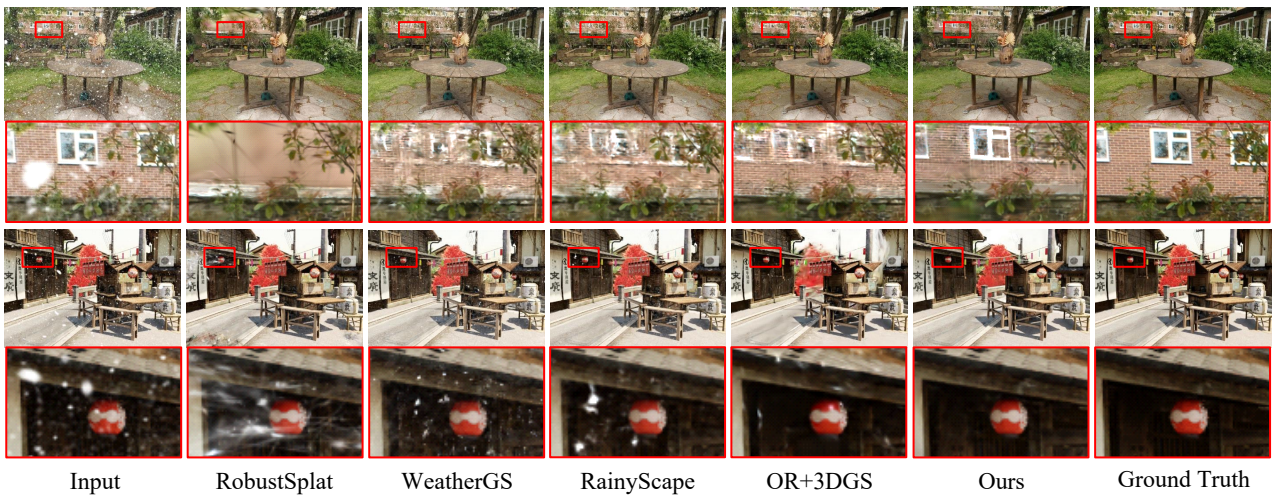
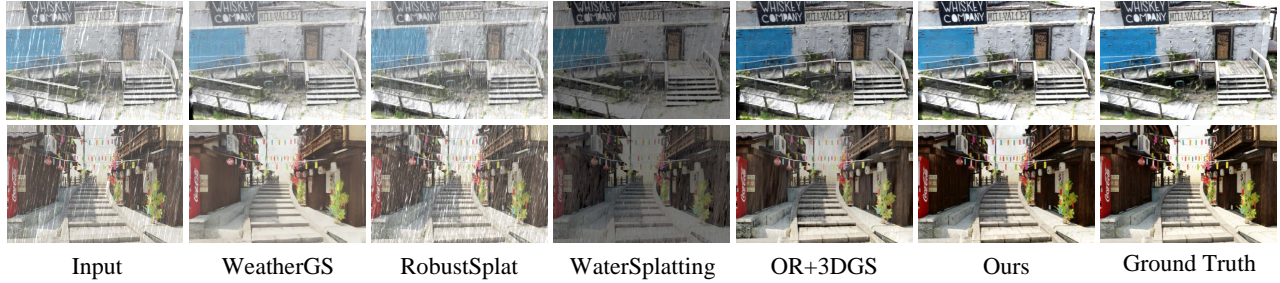
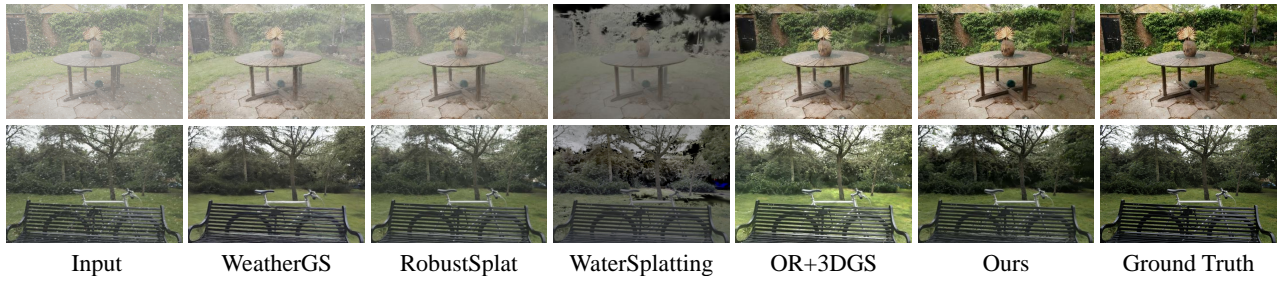


Figure 4. Qualitative comparisons on single-weather scenes (haze, rain, snow).

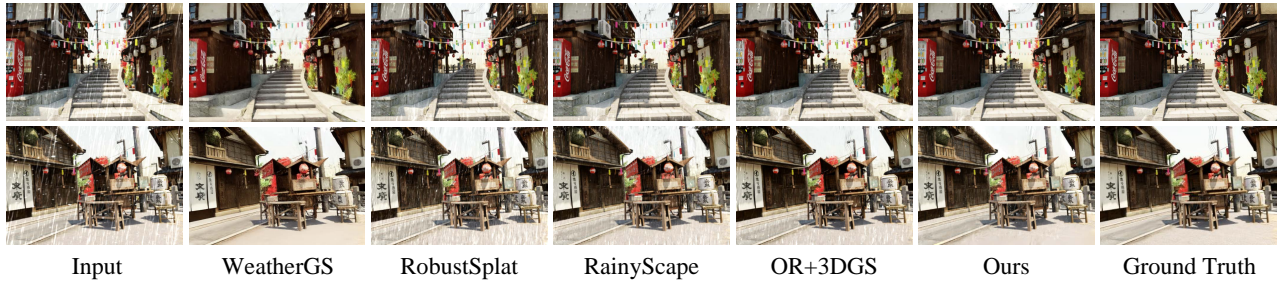
### H+R Setting



### H+S Setting



### R+S Setting



### H+R+S Setting




Figure 5. Qualitative comparisons on hybrid-weather scenes (H+R, H+S, R+S, H+R+S). Best viewed zoomed in.