

Impact of Rain on 3D Reconstruction with Multi-View Stereo, Neural Radiance Fields and Gaussian Splatting

Ivana Petrovska* and Boris Jutzi

Institute of Photogrammetry and Remote Sensing, Karlsruhe Institute of Technology (KIT), Karlsruhe, Germany -
ivana.petrovska@partner.kit.edu, boris.jutzi@kit.edu

KEY WORDS: Multi-View Stereo, Neural Radiance Fields, Gaussian Splatting, 3D Reconstruction, Rain, Dynamic Occlusion.

ABSTRACT:

Image-based 3D reconstruction uncovers many applications in documenting the geometry of the environment. Nonetheless, the assumption that images are captured in clear air rarely holds in real-world settings where adverse weather conditions are inevitable. We are particularly interested in rain as dynamic occlusion which degrades image quality and can hinder complete and accurate 3D scene reconstruction of the underlying features. In this contribution we analyze the geometry behind rain reconstructed by traditional Multi-View Stereo (MVS) and radiance field methods, namely: Neural Radiance Fields (NeRFs), 3D Gaussian Splatting (3DGS) and 2D Gaussian Splatting (2DGS). To assess the impact of rain to the 3D reconstruction we consider occlusion masks with different mask coverage. The results demonstrate that although MVS shows lowest accuracy errors, the completeness declines with rain. NeRFs manifest robustness in the reconstruction with high completeness, while 2DGS achieves second best accuracy results outperforming NeRFs and 3DGS. We demonstrate that radiance field methods can compete against MVS, indicating robustness in the geometric reconstruction under rainy conditions, allowing applicability to large-scale scenes, city modeling, digital twins and urban planning which is important for a multidisciplinary approach in problem-solving environmental challenges.

1. INTRODUCTION

Traditional Multi-View Stereo (MVS) recovers depth and geometry from multiple images with known camera poses by searching correspondences along epipolar lines, which is inefficient. Moreover, MVS struggles with homogeneous areas (Remondino et al., 2023) and occlusions (Petrovska and Jutzi, 2024, Petrovska and Jutzi, 2025c) assuming Lambertian surfaces. Overcoming these constraints, Neural Radiance Fields (NeRFs) (Mildenhall et al., 2021) involve training a neural network to create a volumetric radiance field where the scene is represented implicitly by predicting density and color at any given 3D point from images associated with viewing direction. However, NeRFs assume controlled conditions and are computationally demanding due to the costly ray marching. Furthermore, the continuous density distribution stores points in empty space and may contain multiple peaks, making it difficult for accurate surface reconstruction (Wan et al., 2023b, Petrovska et al., 2023, Petrovska and Jutzi, 2025a, Petrovska and Jutzi, 2025d). Reconstructing scenes as radiance fields without the slow neural rendering part, 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) emerged as an explicit point-based 3D representation. Starting from sparse points usually produced during camera calibration, the scene is represented by 3D Gaussians with position, rotation, scale, opacity and color. During optimization the Gaussians are densified undergoing adjustments in color and shape, until the photometric error between rendered and training images is minimized (Petrovska and Jutzi, 2025b).

Nonetheless, the assumption that images are captured in clear air rarely holds in real-world setting where adverse weather conditions are inevitable. Being a part of our environment, we are particularly interested in rain as dynamic occlusion which degrades captured image quality and can hinder complete and accurate 3D scene reconstruction of the underlying geometry.

In this contribution we qualitatively and quantitatively analyze the geometry under rainy conditions reconstructed by traditional MVS and radiance field methods, namely: NeRFs, 3DGS and 2DGS due to the different geometric representation addressing accuracy and completeness. Our investigations are based on real-world scenarios captured without and with camera flash to investigate how different type and level of rain occlusions affect the geometric reconstruction. The rain is always in front of the object to investigate how the methods allow to reconstruct the underlying geometry behind occlusions, thus investigating if radiance field methods can challenge traditional MVS for scenarios where the latter falls short. To assess the impact of rain to the 3D reconstruction we exclude the rain by considering occlusion masks with different mask coverage. For comparability we include an indoor scenario without occlusions. With a view to aid the benchmark progress, the images and masks are available at <https://github.com/squirrel3/STELLA> because there is no existing dataset dedicated to this problem setting.

In summary, our main contributions are:

- We show that radiance field methods: NeRFs, 3DGS and 2DGS can compete against MVS showing robustness for a complete reconstruction under rainy conditions.
- We provide a comprehensive qualitative and quantitative 3D comparison among MVS, NeRFs, 3DGS and 2DGS reporting accuracy and completeness.
- We introduce two scenarios tackling rainfall for the 3D reconstruction task.

We give an overview of radiance field reconstruction under rainy conditions in Section 2. In Section 3 the applied 3D reconstruction methods are presented. The data capturing is explained in Section 4. In Section 5 information about the evaluation metrics along with implementation details is provided. The qualitative

* Corresponding author

and quantitative results addressing accuracy and completeness are presented in Section 6. The discussion is laid out in Section 7 and Section 8 concludes this contribution.

2. RELATED WORK

In the following, we summarize radiance field geometric reconstruction under rainy conditions and provide an overview of current rain benchmark datasets.

Radiance Field Reconstruction with Rain. Tackling reconstruction in adverse weather conditions, DehazeNeRF (Chen et al., 2024) and ScatterNeRF (Ramazzina et al., 2023) learn a separate representation of the clear scene and the participating media for real-world scenarios. DerainNeRF (Li et al., 2024) addresses the challenge of raindrop removal by leveraging a NeRF-based network and a pre-trained detector. The rain is masked leading to rain-free reconstruction. RainyScape (Lyu et al., 2024) extends the reconstruction to rain streaks. WeatherGS (Qian et al., 2024) and DeRainGS (Liu et al., 2024) address the complexities of 3D reconstruction under rainy conditions in synthetic and real-world setup. The 3DGS-based framework learns rain-free novel views from rainy input images with raindrops and streaks. However, the rain is weak and the droplets are small with regular circular shape. Additionally, the methods are more suitable for novel view synthesis since the scenes have a small number of images and the rain occupies small image parts.

Rain Benchmarks. Several benchmark datasets (Bijelic et al., 2020, Xiao et al., 2023, Kenk and Hassaballah, 2020, Zhang et al., 2023, Wan et al., 2023a, Vora et al., 2023, Pham et al., 2020) focus on occluded parts behind rain droplets and streaks in real-world setup. Nevertheless, the images are non-overlapping using single image annotations and masking for semantic extraction and object detection in adverse weather environments.

We can conclude that NeRFs and Gaussian Splatting (GS) methods are primarily focused on novel view synthesis hence, the evaluation refers to the radiometric image quality. Moreover, a critical aspect still remains unexplored; evaluating the underlying geometry of object’s occluded parts behind rain occlusions. We bring the evaluation from image to 3D metric space through a point cloud comparison, addressing accuracy and completeness to investigate if NeRFs and GS can challenge traditional MVS. Current benchmark datasets tackling rain occlusions consist of non-overlapping images and thus are only suitable for semantic segmentation and object detection. In contrast, our two rain scenarios consist of overlapping images with different rain coverage, enabling 3D reconstruction and evaluation.

3. METHODOLOGY

We describe the applied 3D reconstruction methods in Section 3.1. The MVS dense scene representation is briefly summarized in Section 3.1.1, in Section 3.1.2 the geometric reconstruction of NeRFs is described, then the principles of 3DGS (Section 3.1.3) and 2DGS (Section 3.1.4) are laid out, followed by occlusion mask generation in Section 3.2.

3.1 3D Reconstruction Methods

3.1.1 Multi-View Stereo (MVS). After feature extraction and matching, to establish correspondences between the images, the camera poses are estimated through Structure-from-Motion (SfM) (Schonberger and Frahm, 2016) which additionally produces a sparse point cloud. Then, bundle adjustment to

minimize re-projection errors and refine the camera parameters and 3D points is applied. This allows triangulation of corresponding features in multiple images through pixel-wise computation of depth information with MVS (Schönberger et al., 2016). In a final step, the geometric consistent depth maps are fused into a dense point cloud. Additionally, we consider the occlusion masks (Section 3.2) in MVS by setting a value of 0 for the masked parts in the input images. The poses are input for training NeRFs (Section 3.1.2), while the SfM sparse point cloud is used for 3DGS and 2DGS initialization (Section 3.1.3 and 3.1.4).

3.1.2 Neural Radiance Fields (NeRFs). Given a set of calibrated images with corresponding poses, NeRFs (Mildenhall et al., 2021) optimize an underlying continuous volumetric function. The multi-layer perceptron (MLP) neural network predicts volume density $\sigma(\mathbf{x}) \in \mathbb{R}$ and view-dependent radiance $c(\mathbf{x}, d) \in \mathbb{R}^3$ for a given 3D position $\mathbf{x} \in \mathbb{R}^3$ and viewing direction $d \in \mathbb{R}^3$. The color is calculated as a weighted average of accumulated radiance values along a ray $\mathbf{r}(t) = \mathbf{o} + t\mathbf{d}$, where the ray origin \mathbf{o} coincides with the camera’s projective center. The neural network is trained by minimizing the image reconstruction loss over training views through gradient descent, by leveraging the mean squared error between rendered $\hat{C}(\mathbf{r})$ and ground truth pixel color $C(\mathbf{r})$ for a batch of camera rays \mathcal{R} .

$$\mathcal{L}_{NeRF} = \sum_{\mathbf{r} \in \mathcal{R}} \left\| \hat{C}(\mathbf{r}) - C(\mathbf{r}) \right\|^2 \quad (1)$$

With the occlusion masks $M(u, v) \in [0, 1]$ (Section 3.2), we train NeRF by restricting the masked pixels not to contribute to the optimization in the photometric loss.

$$\mathcal{L}_{NeRF_MASK} = \sum_{\mathbf{r} \in \mathcal{R}} \left\| \left(\hat{C}(\mathbf{r}) - C(\mathbf{r}) \right) \odot (1 - M) \right\|^2 \quad (2)$$

where \odot represents element-wise multiplication, 1 is a vector with the same dimension as the mask M .

3.1.3 3D Gaussian Splatting (3DGS). Initialized on a sparse point cloud, the scene is represented with many differentiable 3D Gaussians (Kerbl et al., 2023) parameterized by position centered in the Gaussian mean $\mu \in \mathbb{R}^3$, covariance matrix $\Sigma \in \mathbb{R}^{3 \times 3}$ decomposed into a scaling vector $s \in \mathbb{R}^3$ and a rotation quaternion $q \in \mathbb{R}^4$, opacity $\alpha \in [0, 1]$ and color c represented via spherical harmonics (SH). During training the Gaussian parameters are optimized through gradient descent by many rendering iterations to best fit the training dataset. The algorithm progressively optimizes the scene by cloning, splitting and culling the Gaussians, resulting in a denser point cloud than the initial sparse point cloud. The 3D Gaussians are trained to minimize the photometric loss which is a combination of $DSSIM$ term (Wang et al., 2004) and \mathcal{L}_1 loss for per-pixel color differences computed between rendered \hat{C} and ground truth C images with $\lambda_{DSSIM} = 0.2$.

$$\mathcal{L}_{3DGS} = (1 - \lambda_{DSSIM}) \|\hat{C} - C\|_1 + \lambda_{DSSIM} DSSIM(\hat{C}, C) \quad (3)$$

We consider the occlusion masks (Section 3.2) in the loss function restricting the gradients to learn only from the unmasked

pixels. With an indexing operation the pixels corresponding to masked parts (with 0 value) are replaced in the ground truth image C_M :

$$\mathcal{L}_{3DGS_MASK} = (1 - \lambda_{DSSIM}) \|\hat{C} - C_M\|_1 + \lambda_{DSSIM} DSSIM(\hat{C}, C_M) \quad (4)$$

The Gaussian mean is considered object geometry as point cloud. For visualization purposes, we color code the point cloud by converting the SH back to RGB values.

3.1.4 2D Gaussian Splatting (2DGS). Employing ray-splat intersection, a 2D Gaussian splat (Huang et al., 2024) is characterized by its central point p_k , two principal tangential vectors t_u and t_v , rotation matrix $R = [t_u, t_v, t_w]$ and a scaling vector $S = (s_u, s_v)$ controlling the size and shape of the 2D Gaussian surfels, opacity $\alpha \in [0, 1]$ and view-dependent appearance c parameterized by SH. The parameters are learned through gradient descent between rendered and ground truth images. Since optimizing solely with photometric loss can lead to noisy reconstruction, two additional regularization terms are introduced: depth distortion \mathcal{L}_D to minimize the distance between ray-splat intersections and normal consistency \mathcal{L}_N which minimizes discrepancies between the rendered normal map and the rendered depth.

$$\mathcal{L}_{2DGS} = \mathcal{L}_{3DGS} + \alpha \mathcal{L}_D + \beta \mathcal{L}_N \quad (5)$$

Following the original implementation, $\alpha = 1000$ for bounded scenes, $\alpha = 100$ for unbounded scenes and $\beta = 0.05$ for all scenes. Subsequently, considering the occlusion masks (Section 3.2) the total loss is minimized guiding the algorithm not to propagate for the masked image pixels.

$$\mathcal{L}_{2DGS_MASK} = \mathcal{L}_{3DGS_MASK} + \alpha \mathcal{L}_D + \beta \mathcal{L}_N \quad (6)$$




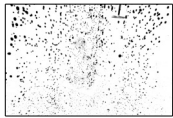
We consider the center of the Gaussian disks as point cloud representation and color code the point cloud by converting the SH back to RGB values.

3.2 Occlusion Masks

To investigate how different type and level of rain occlusions affect the geometric reconstruction, we consider binary masks $M(u, v) \in [0, 1]$ (Table 1). We first convert the images to gray scale and apply histogram equalization for contrast enhancement. We generate the binary masks by filtering through intensity values. For *General-Rain*: if gray value is between 150 and 200 set mask to 0, otherwise set to 1, while for *Illuminated-Rain* the threshold is between 250 and 255. However, due to the brightness the plate is also masked, so we derive annotations using Segment Anything Model (SAM) (Kirillov et al., 2023) in Roboflow¹, then with Boolean operations subtract that part from the masks. Captured without camera flash the rain appears as continuous streaks, while with flash as large elongated droplets. The zero-valued pixel coverage is 23.99 and 4.29 respectively. For simplicity and readability we round the percentages to 24 and 4. The masks are in the same 1840x1228px resolution and .png format as the images.

¹ <https://roboflow.com/>

Table 1. Occlusion rain masks for both scenarios without and with flash. The rain is always in front of the object with 24% and 4% masked pixels accordingly.

Scenario	General-Rain	Illuminated-Rain
RGB Image Preview		
Binary Mask Preview		
Description	Without Flash	With Flash
Percentage	24%	4%

4. DATA

The images for both rain scenarios are captured outdoors with Nikon D810 SLR digital camera with a 36MP sensor, 20mm focal length and f/8 aperture size. Due to the high resolution, the camera is mounted on a tripod to prevent shivering and vibrations during acquisition. The object behind rain whose geometry ought to be evaluated is a 0.7m tall Buddha statue (further on referred to as object) on a rectangular plate. The object is placed on a chair and the images are captured from two camera heights of approximately 0.6m and 1.4m (Figure 1) in a circular trajectory from a uniform distance around the object for full coverage (Figure 2). To generate rain occlusions, we pour water from a watering can in front of the object and capture images from same positions without and with flash, thus acquire two scenarios: *General-Rain* and *Illuminated-Rain*. Each scenario consists of 125 images with 1840x1228px resolution in .png format for lossless compression.




Figure 1. Image capturing setup from two camera heights of approximately (a) 0.6m and (b) 1.4m for full coverage. We simulate rain by pouring water from a watering can in front of the object and capture images from same positions without and with flash.

5. EXPERIMENTS

We first explain the 3D evaluation metrics, followed by the implementation details for each 3D reconstruction method.

3D Evaluation. As ground truth, we use a mesh generated using Structured Light Imaging (SLI) with 0.1mm accuracy. We remove redundant data and keep just the object. All point clouds are aligned in the same metric space as the ground truth for qualitative and quantitative evaluation addressing accuracy and completeness. The geometric distortions are estimated by: Mean Error (Mean), Standard Deviation (SD) and Root Mean

Squared Error (RMSE). Completeness (Cpl) is calculated as the ratio of covered points to the total number of points in the reference cloud within 5mm distance threshold. Thus, we converted the mesh into a point cloud by evenly subsampling 10M points on the mesh.

Implementation Details. We use COLMAP² as it implements an end-to-end MVS pipeline. As NeRF representative we train Nerfacto (Tancik et al., 2023) in Nerfstudio³ v1.1.0. with default parameters without pose refinement. In the original implementation, we use the default hyperparameters for 3DGS and 2DGS. NeRFs and GS are trained for 30.000 iterations with every 8th frame taken for test. Training and evaluation are performed on a Nvidia RTX3090. The accuracy is calculated in CloudCompare⁴, while the completeness in our python script.




Figure 2. Image capturing trajectory with visualized camera poses for both rain scenarios (a) without and (b) with flash. The red rectangle indicates (11) images not taken for pose estimation due to homogeneous and repetitive texture, challenged by the rain.

6. RESULTS

We report qualitative, where the geometric reconstruction as point clouds are visualized (Figure 3) as well as the cloud-to-mesh errors (Figure 4) and quantitative results addressing accuracy and completeness for each method separately (Table 2).

Overall, MVS demonstrates superior accuracy in all scenarios with lowest error displacements among the methods. However, it struggles to reconstruct object's geometry behind occlusions, especially in *Illuminated-Rain* and *Illuminated-Rain Masks* where only the edges are reconstructed resulting in low completeness scores of 31.57% and 18.89% respectively. Due to the costly dense matching, MVS reconstruction time is above 1h enough for third place. On the other hand, radiance field methods provide higher point coverage and perform better under rainy conditions, able to approximate the geometry of the occluded object parts. Although NeRF exhibits highest RMSE of above 10 and 20mm in *General-Rain* and *Illuminated-Rain* it shows robustness in the reconstruction as the surface representation remains stable with highest and second highest completeness scores. This is complimented by fastest training time within few minutes without and with masks. 3DGS struggles with accuracy and completeness, but maintains stable reconstruction in both rain scenarios and the masks don't have a strong impact on the evaluation metrics. Regarding training time, the performance is moderate with slightly more than 30min. 2DGS exhibits second best accuracy outperforming NeRF in all scenarios and highest completeness in *Illuminated-Rain* without and with masks. However, challenged by the complex light

transmission properties of the rain as it assumes surfaces with full opacity, it fails to reconstruct the object in *General-Rain Masks*. After around 14.000 iterations it runs Out Of Memory (OOM) caused by an explosion in Gaussian count as it adds Gaussians per-view optimizing in screen space. This explains the slow training time of above 6h in rainy conditions.

NeRF, 3DGS and 2DGS show lower accuracy than MVS since the optimization is based on minimizing the difference between predicted and actual pixel color from the input images through gradient descent without additional geometric constraints. On top of that, artifact points inside the object which are projections from input views that weren't moved into their correct place geometrically are present which distort the accuracy and don't contribute to the completeness. Thus, the error displacements tend to increase faster to the negative because those are the points behind the mesh surface. The errors with positive values lie above the nearest mesh triangle and thus most likely represent noise and outlier points. We can also observe color differences in the point clouds among the scenarios. *Original* is captured indoors under controlled lighting, which is different from *General-Rain* and *Illuminated-Rain* captured outdoors under different camera exposure, without and with camera flash.

7. DISCUSSION

The geometric reconstruction under rainy conditions depends on the 3D reconstruction method, image capturing setup and rain occlusion masks. In both scenarios, the rain acts as semi-transparent dynamic reflections that can blend into the background and may be misinterpreted as actual scene geometry, leading to incorrect depth estimates. Captured with flash, the rain appears as bright, high-contrast spots creating distinct occlusions imposing challenges for all methods leading to reduced accuracy and completeness metrics compared to without flash (Table 2).

During pose estimation, specular reflections cause mismatches between corresponding points in different images. Moreover, the presence of glare can distort or eliminate key features necessary for accurate feature extraction and matching. Complemented by homogeneous background and repetitive texture, in total 11 images captured from the right object side (Figure 2(b)) are not taken for pose estimation in *Illuminated-Rain*. Additionally, multi-path reflections where light bounces off multiple reflective surfaces before reaching the camera may confuse the matching algorithm to incorrectly associate features from different reflections, leading to erroneous depth information and distortions in the 3D reconstruction (Previtali et al., 2024). Thus, MVS struggles to reliably reconstruct the object geometry with low completeness scores (Figure 3).

Although NeRF achieves second best completeness score in all scenarios, the accuracy is lowest among the methods due to the ambiguity between light-reflecting solid surfaces and light-scattering atmospheric particles since both are interchangeably modeled as a light-emitting volume.

Overlapping rain streaks create uncertainty in Gaussian placement, especially for 2DGS. Since the 3D Gaussians are projected to screen space and evaluated in 2D, per-view Gaussians may be misaligned, leading to incorrect surface orientation. The raindrops create many small, high-intensity Gaussians which can collapse to very small points in screen space increasing computational complexity. This explains the long training

² <https://github.com/colmap/colmap>

³ <https://github.com/nerfstudio-project/nerfstudio>

⁴ <https://github.com/CloudCompare/CloudCompare>




Figure 3. Point cloud reconstruction for both rain scenarios without and with masks. MVS completeness is significantly affected, while NeRF exhibits highest point coverage indicating robustness in the reconstruction. 3DGS struggles with point coverage, but maintains stable reconstruction under masks. 2DGS fails in *General-Rain Masks* running Out Of Memory caused by an explosion in Gaussian count. Captured with flash the rain creates distinct occlusions, thus *Illuminated-Rain* imposes bigger challenges for all methods.



Figure 4. Geometric accuracy through cloud-to-mesh distances against the ground truth mesh. Positive values indicate points above the mesh and negative when behind. Except MVS, the reconstructed point clouds have artifact points inside the object with high errors degrading the accuracy. 2DGS fails in *General-Rain Masks* running Out Of Memory caused by an explosion in Gaussian count.

Table 2. Quantitative cloud-to-mesh results addressing accuracy and completeness along with training time. 2DGS exhibits second best accuracy just behind MVS, experiencing slowest reconstruction. With fastest training time, NeRF shows poor correspondence with the ground truth, however maintaining stable point coverage. The **first**, **second** and **third** best results are highlighted.

Scenario	Method	Accuracy (mm)			Completeness (%)		Time ↓
		Mean ↓	SD ↓	RMSE ↓	Npts	Cpl ↑	
<i>Original</i>	MVS	0.16	1.42	1.43	845.456	97.35	1h 22min
	NeRF	-2.10	3.78	4.32	842.308	96.13	16min
	3DGS	-2.02	4.29	4.74	352.298	82.75	33min
	2DGS	-1.22	3.59	3.79	156.465	90.73	54min
<i>General-Rain</i>	MVS	0.42	2.15	2.19	313.555	70.97	1h 13min
	NeRF	-5.98	10.37	11.97	330.299	86.55	14min
	3DGS	-2.93	7.89	8.42	53.995	41.79	34min
	2DGS	-1.07	7.19	7.27	41.470	59.19	6h 52min
<i>General-Rain Masks</i>	MVS	0.38	2.08	2.12	219.572	64.78	1h 22min
	NeRF	-5.21	9.62	10.94	232.297	85.87	13min
	3DGS	-2.86	8.01	8.50	42.493	40.53	32min
	2DGS	Failed					
<i>Illuminated-Rain</i>	MVS	0.38	5.45	5.46	46.918	31.57	1h 05min
	NeRF	-11.69	17.32	20.89	48.546	36.88	13min
	3DGS	-1.82	10.36	10.52	20.619	25.12	31min
	2DGS	-1.96	10.13	10.31	18.531	37.11	5h 41min
<i>Illuminated-Rain Masks</i>	MVS	-1.08	5.32	5.43	20.741	18.89	1h 13min
	NeRF	-11.39	16.72	20.23	20.672	26.51	12min
	3DGS	-1.79	10.19	10.35	20.352	25.13	32min
	2DGS	-2.00	10.12	10.32	18.450	34.63	7h 38min

time of above 6h in rainy conditions (Table 2). With high opacity, their movement can cause significant pixel changes, leading to pronounced positional gradients. Moreover, some points project smaller than one pixel, resulting in their covariance being replaced by a fixed value. Consequently, these points can't properly adjust their scale and rotation causing rapid gradient accumulation. This triggers exponential increase in Gaussian count as 2DGS keeps adding Gaussians per-view leading to OOM after approximately 14,000 iterations. Until 7,000 iterations the number of Gaussians is around 3.6M, almost the same as 3DGS after 30,000 iterations (Table 3).

Table 3. We use SfM for GS initialization (No. of Points). We report the number of Gaussians after 30,000 iterations, except for 2DGS in *General-Rain Masks* until 7,000 (*) because it runs OOM after around 14,000 iterations due to high Gaussian count.

Scenario	Method	No. of Points	Gaussian Count
<i>Original</i>	3DGS	42.580	1.893.679
	2DGS		822.797
<i>General-Rain</i>	3DGS	123.761	4.857.840
	2DGS		5.093.434
<i>General-Rain Masks</i>	3DGS	123.761	3.909.893
	2DGS		3.562.482*
<i>Illuminated-Rain</i>	3DGS	76.395	4.035.943
	2DGS		3.255.605
<i>Illuminated-Rain Masks</i>	3DGS	76.395	4.117.921
	2DGS		3.849.476

The occlusion masks also have limitations. Although we applied histogram equalization to adjust the pixel intensity distribution, the elongated rainfall effects are not masked in both scenarios, especially without flash (Figure 5) due to the low contrast and semi-transparency. This imposes challenges for the 3D reconstruction methods especially 2DGS which optimizes per-image causing it to fail in *General-Rain Masks*. Furthermore, object parts not being rain are masked, causing gaps in the geometry especially noticeable in MVS. Although the occlusion percentage for *General-Rain* is 24% (Table 1), some of the background and floor is masked and rain streaks are missed.

In contrast, in *Illuminated-Rain* captured with flash the mask percentage is 4 but the rain occupies bigger image part.




Figure 5. Mask overlay for both rain scenarios. The rain streaks are missed and non-occluded object parts are masked, challenging the 3D reconstruction methods.

8. CONCLUSION

We provide a qualitative and quantitative analysis of the accuracy and completeness of the 3D geometry behind rain reconstructed by traditional MVS and radiance field methods: NeRF, 3DGS and 2DGS. We acquire two real-world rain scenarios captured without and with flash from the same camera position for compatibility. To investigate the impact of rain occlusions on the 3D reconstruction we consider occlusion masks. Captured with flash, the rain appears as bright, high-contrast spots creating distinct occlusions imposing bigger challenges for all methods. MVS excels in pin-point accuracy, however the completeness declines under rain conditions, making it sensitive to occlusions. NeRF exhibits robustness in the reconstruction with high completeness. 3DGS struggles to reliably reconstruct the geometry behind rain, while 2DGS outperforms NeRF and

3DGS regarding accuracy in all scenarios. We demonstrate that radiance field methods can compete against traditional MVS, showing robustness in rainy conditions. Moreover, we indicate radiance fields ability to reliably reconstruct the geometry behind rain showing potential for larger geospatial scenes.

REFERENCES

- Bijelic, M., Gruber, T., Mannan, F., Kraus, F., Ritter, W., Dietmayer, K., Heide, F., 2020. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 11682–11692.
- Chen, W.-T., Yifan, W., Kuo, S.-Y., Wetzstein, G., 2024. Dehazenerf: Multi-image haze removal and 3d shape reconstruction using neural radiance fields. *2024 International Conference on 3D Vision (3DV)*, IEEE, 247–256.
- Huang, B., Yu, Z., Chen, A., Geiger, A., Gao, S., 2024. 2d gaussian splatting for geometrically accurate radiance fields. *ACM SIGGRAPH 2024 conference papers*, 1–11.
- Kenk, M. A., Hassaballah, M., 2020. DAWN: vehicle detection in adverse weather nature dataset. *arXiv:2008.05402*.
- Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G., 2023. 3D Gaussian Splatting for Real-Time Radiance Field Rendering. *ACM Trans. Graph.*, 42(4), 139–1.
- Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A. C., Lo, W.-Y. et al., 2023. Segment anything. *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4015–4026.
- Li, Y., Wu, J., Zhao, L., Liu, P., 2024. Derainnerf: 3d scene estimation with adhesive waterdrop removal. *2024 IEEE Intern. Conf. on Robotics and Automation (ICRA)*, IEEE, 2787–2793.
- Liu, S., Chen, X., Chen, H., Xu, Q., Li, M., 2024. DeRainGS: Gaussian Splatting for Enhanced Scene Reconstruction in Rainy Environments. *arXiv preprint arXiv:2408.11540*.
- Lyu, X., Liu, H., Hou, J., 2024. Rainyscape: Unsupervised rainy scene reconstruction using decoupled neural rendering. *Proceedings of the 32nd ACM International Conference on Multimedia*, 10920–10929.
- Mildenhall, B., Srinivasan, P. P., Tancik, M., Barron, J. T., Ramamoorthi, R., Ng, R., 2021. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1), 99–106.
- Petrovska, I., Jäger, M., Haitz, D., Jutzi, B., 2023. Geometric Accuracy Analysis between Neural Radiance Fields (NeRFs) and Terrestrial laser scanning (TLS). *ISPRS Archives of the Photogr., Remote Sens. and Spatial Inf. Sciences*, 48, 153–159.
- Petrovska, I., Jutzi, B., 2024. Vision through Obstacles-3D Geometric Reconstruction and Evaluation of Neural Radiance Fields (NeRFs). *Remote Sensing*, 16(7), 1188.
- Petrovska, I., Jutzi, B., 2025a. 3D Gaussian Splatting Methods for Real-World Scenarios. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, X-G-2025, 641–648.
- Petrovska, I., Jutzi, B., 2025b. 3D Gaussian Splatting Methods for Real-World Scenarios. *ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 641–648.
- Petrovska, I., Jutzi, B., 2025c. Seeing beyond vegetation: A comparative occlusion analysis between Multi-View Stereo, Neural Radiance Fields and Gaussian Splatting for 3D reconstruction. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, 100089.
- Petrovska, I., Jutzi, B., 2025d. Seeing beyond vegetation: A comparative occlusion analysis between Multi-View Stereo, Neural Radiance Fields and Gaussian Splatting for 3D reconstruction. *ISPRS Open Journal of Photogrammetry and Remote Sensing*, 16, 100089.
- Pham, Q.-H., Sevestre, P., Pahwa, R. S., Zhan, H., Pang, C. H., Chen, Y., Mustafa, A., Chandrasekhar, V., Lin, J., 2020. A* 3d dataset: Towards autonomous driving in challenging environments. *2020 IEEE International conference on Robotics and Automation (ICRA)*, IEEE, 2267–2273.
- Previtali, M., Barazzetti, L., Roncoroni, F., 2024. Orthophoto generation with gaussian splatting: mitigating reflective surface artifacts in imagery from low-cost sensors. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 48, 371–378.
- Qian, C., Guo, Y., Li, W., Markkula, G., 2024. WeatherGS: 3D Scene Reconstruction in Adverse Weather Conditions via Gaussian Splatting. *arXiv preprint arXiv:2412.18862*.
- Ramazzina, A., Bijelic, M., Walz, S., Sanvito, A., Scheuble, D., Heide, F., 2023. Scatternerf: Seeing through fog with physically-based inverse neural rendering. *Proceedings IEEE/CVF Intern. Conf. on Computer Vision*, 17957–17968.
- Remondino, F., Karami, A., Yan, Z., Mazzacca, G., Rigon, S., Qin, R., 2023. A critical analysis of NeRF-based 3D reconstruction. *Remote Sensing*, 15(14), 3585.
- Schonberger, J. L., Frahm, J.-M., 2016. Structure-from-motion revisited. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4104–4113.
- Schönberger, J. L., Zheng, E., Frahm, J.-M., Pollefeys, M., 2016. Pixelwise view selection for unstructured multi-view stereo. *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, Springer, 501–518.
- Tancik, M., Weber, E., Ng, E., Li, R., Yi, B., Wang, T., Kristoffersen, A., Austin, J., Salahi, K., Ahuja, A. et al., 2023. Nerfstudio: A modular framework for neural radiance field development. *ACM SIGGRAPH 2023 Conference Proceedings*, 1–12.
- Vora, J., Dutta, S., Jain, K., Karthik, S., Gandhi, V., 2023. Bringing generalization to deep multi-view pedestrian detection. *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 110–119.
- Wan, Y.-C., Shao, M.-W., Cheng, Y.-S., Liu, Y.-X., Bao, Z.-Y., 2023a. Restoring images captured in arbitrary hybrid adverse weather conditions in one go. *arXiv preprint arXiv:2305.09996*.
- Wan, Z., Richardt, C., Božić, A., Li, C., Rengarajan, V., Nam, S., Xiang, X., Li, T., Zhu, B., Ranjan, R. et al., 2023b. Learning neural duplex radiance fields for real-time view synthesis. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8307–8316.
- Wang, Z., Bovik, A. C., Sheikh, H. R., Simoncelli, E. P., 2004. Image quality assessment: from error visibility to structural similarity. *IEEE trans. on image processing*, 13(4), 600–612.
- Xiao, A., Huang, J., Xuan, W., Ren, R., Liu, K., Guan, D., El Saddik, A., Lu, S., Xing, E. P., 2023. 3d semantic segmentation in the wild: Learning generalized models for adverse-condition point clouds. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 9382–9392.
- Zhang, J., Liu, R., Shi, H., Yang, K., Reiß, S., Peng, K., Fu, H., Wang, K., Stiefelhagen, R., 2023. Delivering arbitrary-modal semantic segmentation. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1136–1147.