RestorGS: Depth-aware Gaussian Splatting for Efficient 3D Scene Restoration

Supplementary Material

1. Additional Details

We provide more details of our RestorGS. Specifically, we modify the diffi-gauss-rasterization used in the original 3DGS to render depth maps. To maintain a proper visual perception, we set the exposure value to 0.4 to control the brightness of the restoration results. The corresponding ablation studies are shown in Fig. 2. The ablation results in the main paper are evaluated using the average metrics in the underwater scene. We also show the training details of RestorGS in Algorithm 1.

Algorithm 1 Training steps of RestorGS.

Input: $I(x)$: degraded image; θ_i: learnable degradation factors; G_I : Gaussian parameter $\{\mu_j, \Sigma_j, \alpha_j, SH_j\}$; MLP: color MLP; f_{CNN}: convolutional neural network; β: scattering coefficient

Output: $\hat{J}(x)$: restored image; $\hat{I}(x)$: rendered images; θ_G: trained model;

Set the initialization parameters: λ_1 = 10; λ_2 = 0.001; α_1 = 0.001; α_2 = 0.0005; $total\_i$ = 15000;
for $i = 1$ to $total\_i$ **do**;
    Appearance Decoupling:
    $\eta, \delta \leftarrow MLP(\gamma(\mu_j), \gamma(SH_j))$;
    $SH_j^c \leftarrow \eta \cdot SH_j + \delta$;
    $SH_j^d \leftarrow [\theta_0, \theta_1, \cdots, \theta_n]$;
    $G_C \leftarrow \{\mu_j, \Sigma_j, \alpha_j, SH_j^c\}$;
    $G_D \leftarrow \{\mu_j, \Sigma_j, \alpha_j, SH_j^d\}$;
    Depth-Guided Modeling:
    $\hat{J}(x), \hat{d}(x) \leftarrow rasterized(G_C)$;
    $F_d \leftarrow rasterized(G_D)$;
    $\hat{d}_{norm}(x) \leftarrow norm(\hat{d}(x))$;
    $\hat{A}(x) \leftarrow e^{-\beta \hat{d}_{norm}(x)}$;
    $\hat{L} \leftarrow f_{CNN}[F_d, A(x)]$;
    $\hat{I}(x) \leftarrow \hat{J}(x) \cdot \hat{A}(x) + \hat{L} \cdot (1 - \hat{A}(x))$;
end for
$\mathcal{L}_{total} \leftarrow \mathcal{L}_{re} + \lambda_1 \mathcal{L}_{exp} + \lambda_2 \mathcal{L}_{depth}$;
Optimize the network with $\mathcal{L}_{total}$;
return θ_G.

2. More Quantitative Comparisons

In Table 1 and Table 2, we further conduct quantitative comparisons with 3DGS [2], Gaussian-DK [5], and WaterSplatting [4] using the non-reference image quality assessment metrics NIQE, MUSIQ, and PI. The experimental results show that our RestorGS outperforms these schemes in most of the metrics, verifying its significant advantages in image recovery quality. It is worth noting that although RestorGS slightly underperforms 3DGS in some specific metrics, 3DGS is unable to effectively restore degraded images. The main reason is that 3DGS is only suitable for high-quality scenes, which greatly limits its practical application. In contrast, RestorGS can not only handle degraded scenes, but also ensure that the restored image is closer to the real scene in terms of detail retention, color reproduction, and overall naturalness. In addition, we compare the average training time and FPS of the different methods in Table 3. As mentioned in the main paper, our method achieves high-quality restoration while enabling real-time rendering speed. In Fig. 1, we also show the average quantitative scores of the different methods in underwater, nighttime and hazy scenes. It can be seen that our RestorGS significantly outperforms existing methods.

3. More Ablation Results

We similarly conduct ablation studies for exposure loss. As shown in Fig. 2, with the increase of the exposure value E, the intensity of the luminance map is enhanced accordingly, indicating the significant effect of exposure adjustment on the image brightness. Meanwhile, the global brightness of the restored image shows a decreasing trend with the increase of the exposure value E. This is due to the complementary relationship between the brightness of the restored image and the luminance map. It can be seen that when the exposure value E is set to 0.4, the restoration results reach an ideal balance between brightness and color correction, and avoids overexposure or underexposure phenomenon to mask the details. This result fully demonstrates that a reasonable exposure value can promote a more natural restoration of brightness and color.

4. More Visual Results

Visual Comparison. In Fig. 3, Fig. 4, and Fig. 5, we show more visual comparison results to further validate the superiority of our method. It can be clearly observed that existing methods usually face significant challenges in dealing with degraded scenes with inconsistent viewpoints, resulting in residual effects in the scattering medium and inaccuracies in light correction. Such residual effects can lead to noticeable haze or blurring in the restored images, while inappropriate light correction may trigger localized overexposure or underexposure of the images, significantly affecting the quality of visual perception. In contrast, our method is able to adaptively adjust the brightness and contrast, thus avoiding the distortion problem and presenting a good visual perception.

Table 1. Quantitative comparison on the SeaThru-NeRF dataset [3]. The best and second-best score are marked as red and yellow, respectively. ↓: lower is better, ↑: higher is better.

Method	Curacao			IUI3 Red Sea			Panama			J.G. Red Sea		
	NIQE↓	MUSIQ↑	PI↓	NIQE↓	MUSIQ↑	PI↓	NIQE↓	MUSIQ↑	PI↓	NIQE↓	MUSIQ↑	PI↓
3DGS [2]	2.931	59.378	2.132	3.101	60.002	2.225	3.601	53.086	2.380	2.726	55.865	1.957
Gaussian-DK [5]	3.925	48.923	4.187	3.667	60.388	2.604	4.124	52.807	3.616	4.460	36.989	4.227
WaterSplatting [4]	2.975	58.741	3.362	3.337	57.948	2.854	3.327	49.883	3.637	3.531	56.452	2.878
Ours	2.695	61.222	2.083	2.704	62.377	2.340	3.739	55.376	3.072	3.054	58.562	2.180

Table 2. Quantitative comparison on the Gaussian-DK dataset [5]. The best and second-best score are marked as red and yellow, respectively. ↓: lower is better, ↑: higher is better.

Method	Bicycle			Dormitory			Street			Kitchen		
	NIQE↓	MUSIQ↑	PI↓	NIQE↓	MUSIQ↑	PI↓	NIQE↓	MUSIQ↑	PI↓	NIQE↓	MUSIQ↑	PI↓
3DGS [2]	4.130	60.771	3.416	4.464	67.981	3.986	2.703	64.915	2.498	4.010	69.504	3.167
Gaussian-DK [5]	5.193	66.781	3.978	5.087	72.223	3.661	2.986	70.557	2.569	5.125	74.606	3.800
WaterSplatting [4]	4.953	62.841	4.275	5.248	70.661	3.736	2.895	68.834	2.894	4.827	68.036	3.465
Ours	4.081	68.372	4.172	4.914	73.974	3.633	2.529	71.064	2.975	4.538	74.931	3.307

Table 3. Comparison of average training time and FPS. ↓: lower is better, ↑: higher is better.

Ablation	Training Time↓	FPS↑
SeaThru-NeRF [3]	12h	0.08
3DGS [2]	22.6min	162
Gaussian-DK [5]	28.4min	69
WaterSplatting [4]	23.8min	48
Ours	34.4min	52

Visualization Results. We further visualize more examples of the learned depth map, attenuation map and illumination map in Fig. 6. Note that despite some noise in the depth map, 3D scene restoration relies more on the complementarity of multi-view depth than absolute accuracy in a single view. It can be seen that each physical component can accurately reflect the degradation change with scene depth. This reveals the physical mechanism of Eq. (9) that the transmitted light decreases with the increase of scene depth after medium attenuation.

5. Limitation and Future Work

Although our RestorGS has achieved superior capabilities in restoring 3D degraded scenes, it may suffer from limitations when dealing with extremely severe degradation situations. The primary reason is that these severely degraded scenes are usually accompanied by the superposition of multiple complex factors, such as extremely dense scattering media, extreme lighting conditions (e.g., overexposure or complete darkness), and extremely blurred or missing depth information. These factors can dramatically increase the difficulty of image restoration and scene reconstruction, posing significant challenges to existing models.

To address these issues, one potential solution is to seamlessly combine large-scale pre-trained 2D image restoration models with 3DGS. By utilizing the strong image restoration capabilities of 2D models to complement 3DGS, it is expected to address the shortcomings when dealing with extremely degraded scenes. However, this approach inevitably increases the computational complexity significantly, especially in application scenarios that require real-time processing. Therefore, achieving a balance between recovery performance and model complexity is an important direction 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5470–5479, 2022. 4

[2] 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. 1, 2

[3] Deborah Levy, Amit Peleg, Naama Pearl, Dan Rosenbaum, Derya Akkaynak, Simon Korman, and Tali Treibitz. Seathru-nerf: Neural radiance fields in scattering media. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 56–65, 2023. 2, 3

[4] Huapeng Li, Wenxuan Song, Tianao Xu, Alexandre Elsig, and Jonas Kulhanek. Watersplatting: Fast underwater 3d scene reconstruction using gaussian splatting. *arXiv preprint arXiv:2408.08206*, 2024. 1, 2

[5] Sheng Ye, Zhen-Hui Dong, Yubin Hu, Yu-Hui Wen, and Yong-Jin Liu. Gaussian in the dark: Real-time view synthesis from inconsistent dark images using gaussian splatting. In *Computer Graphics Forum*, page e15213, 2024. 1, 2, 4

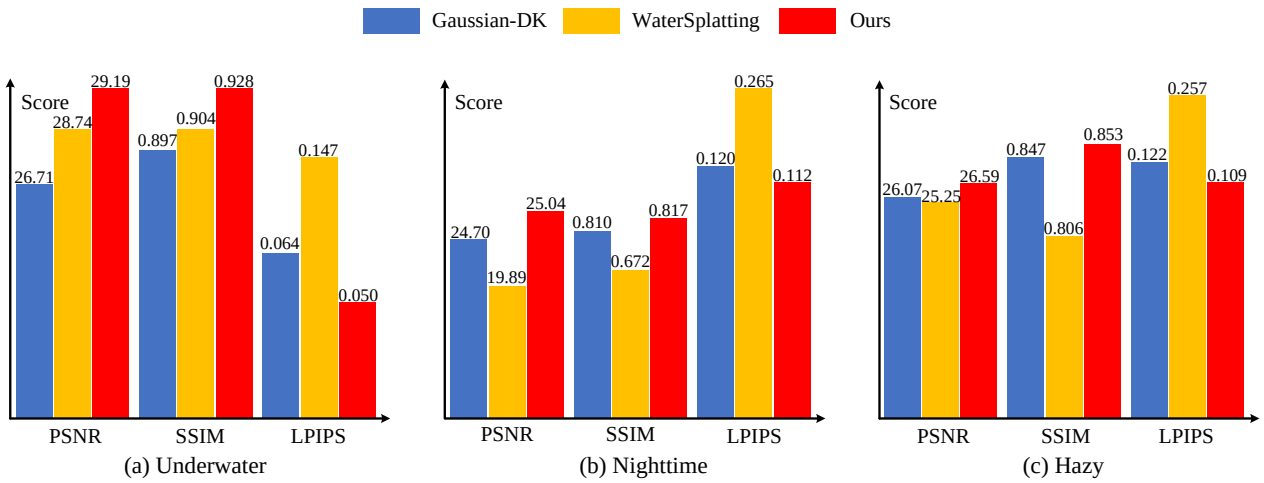

Figure 1. Average scores of different methods in underwater, nighttime and hazy scenes.

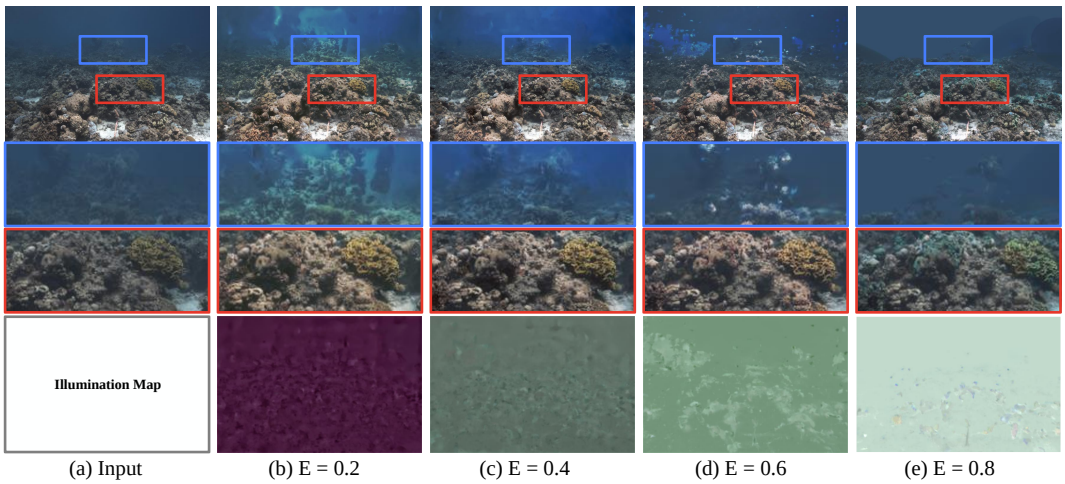

Figure 2. Ablation studies of exposure loss. As the exposure value increases, the intensity of the luminance map increases while the brightness of the restored image decreases.

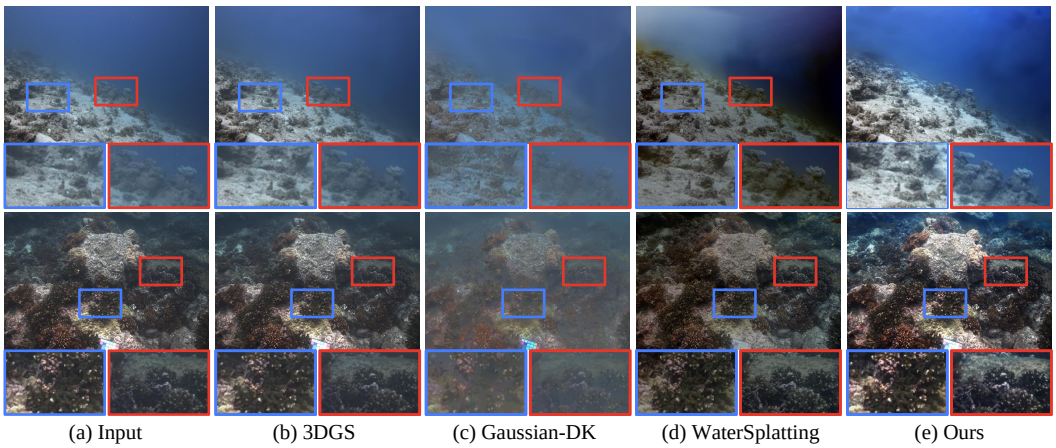

Figure 3. Qualitative comparison of different methods on the SeaThru-NeRF [3] dataset.

(a) Input (b) 3DGS (c) Gaussian-DK (d) WaterSplatting (e) Ours

Figure 4. Qualitative comparison of different methods on the Gaussian-DK dataset [5].

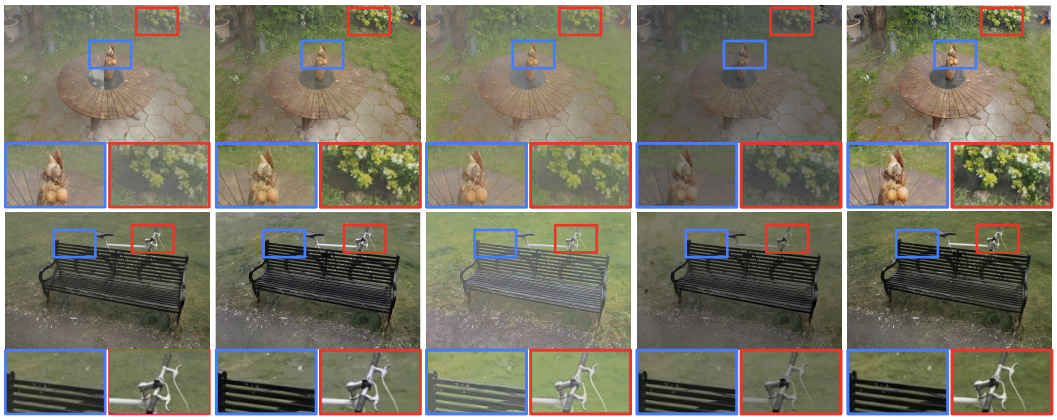

(a) Input (b) 3DGS (c) Gaussian-DK (d) WaterSplatting (e) Ours

Figure 5. Qualitative comparison of different methods on the Mip-NeRF360 [1] dataset.

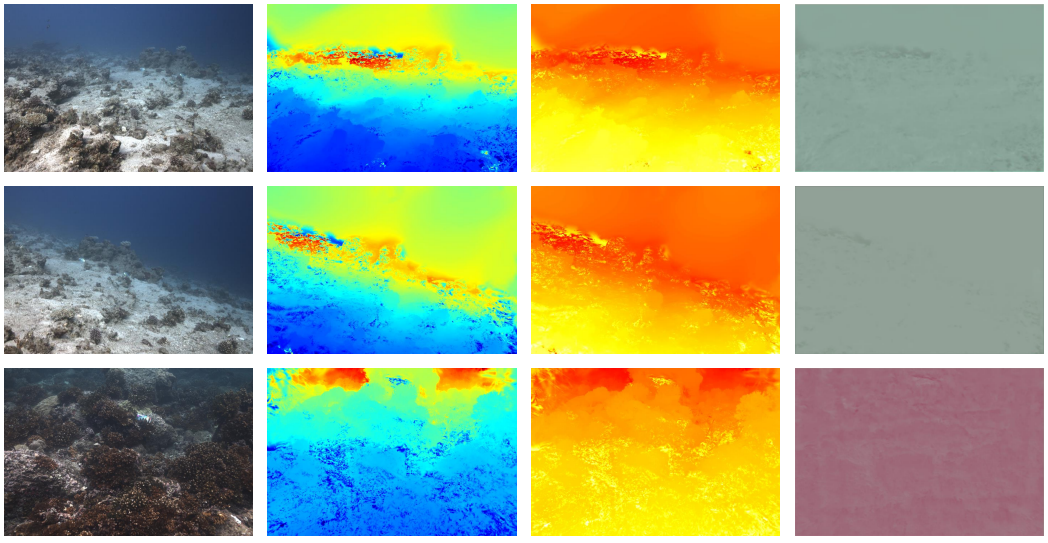

(a) Input (b) Depth Map (c) Attenuation Map (d) Illumination Map

Figure 6. More Visualization. The learned physical components can accurately reflect the degradation distribution.