

# Physically Inspired Gaussian Splatting for HDR Novel View Synthesis

## –Supplementary Materials–

This supplementary document is organized as follows:

- Section 1 provides a detailed explanation and pseudo-code to clarify the training procedure of the proposed PhysHDR-GS.
- Section 2 reports per-scene quantitative results on all adopted datasets.
- Section 3 includes more qualitative comparisons on both LDR and HDR views.
- Section 4 introduces further experimental details, including dataset preparation and implementation settings.
- Section 5 includes our observation for HDR-NVS under low-light conditions and discusses potential solutions to address such scenarios.

### 1. Training Procedure

The pseudo-code of the training procedure is summarized in Algorithm 1. PhysHDR-GS is built from three components: physical radiance composition, self-consistent HDR fusion and illumination-guided gradient scaling. Physical radiance modeling incorporates an image–exposure (IE) branch and a Gaussian–illumination (GI) branch. The IE branch composes intrinsic reflectance and ambient illumination into per-Gaussian color and synthesizes the HDR Gaussian primitives, while the GI branch adjusts the illumination and produces relit HDR Gaussian primitives. At the self-consistent HDR fusion stage, the projected HDR image from the IE branch is re-exposed to match the brightness of the relit HDR image of the GI branch. These HDR intermediates are tone-mapped and fused as the final LDR output. During training, the cross-branch HDR consistency loss and illumination-guided gradient scaling are adopted to enable self-supervision for HDR learning and alleviate under-densified representations in over-/under-exposed regions, respectively.

### 2. Per-scene Quantitative Results

We report per-scene quantitative results on HDR-NeRF-Real [2], HDR-Plenoxels-Real [3], and HDR-NeRF-Syn [2] in Table 1, Table 2, Table 3, and Table 4. On HDR-NeRF-Real, Ours<sup>†</sup> achieves leading performance, obtaining the best scores on most scenes, while the 3DGS variant (Ours) also shows competitive results and delivers the

best on several scenes. Notably, compared with the strong GaussHDR variants [5], our method delivers higher rendering speed (e.g., Ours<sup>†</sup> is  $1.37\times$  faster than GaussHDR<sup>†</sup>) and lower memory usage during training. On HDR-Plenoxels-Real, Ours<sup>†</sup> again shows overall superior performance over all compared methods. On the synthetic HDR-NeRF-Syn dataset, Ours<sup>†</sup> consistently outperforms baselines across both exp3 and exp1 exposure settings, achieving the best results on most scenes. This highlights our superiority in reconstructing HDR details, as well as the tone mapping ability to faithfully preserve standard sensor observations.

### 3. More Qualitative Results

We provide additional qualitative LDR results in Figure 1 and HDR results in Figure 2. For each method, we also visualize the residual map with respect to the ground truth to intuitively show the difference (from blue to red, higher values indicate larger errors). For LDR results, existing methods struggle to balance reconstructing dynamic range content and preserving fine structures (e.g., the texture details of HDR-GS [1] at the highlight region are missing). In contrast, our method preserves both dynamic-range detail and local structures, and even reflects environment-conditioned appearance changes (e.g., the reflection of the fireplace), which we attribute to the complementary IE and GI branches that decouple exposure handling and illumination modeling. For HDR results, HDR-NeRF [2] and HDR-GS often fail to reach the correct illumination level due to the lack of HDR ground truth during training, while our approach produces more accurate dynamic-range estimations and sharper reconstructions (e.g., leaf textures in the 4th row), demonstrating the ability to restore dynamic ranges and keep structural details.

### 4. Experimental Settings

**Datasets and settings.** For the HDR-Plenoxels-Real [3] dataset, LDR-NE ( $\{t_2, t_4\}$ ) views are not provided for the *Character* and *Coffee* scenes. Therefore, for LDR-NE results on HDR-Plenoxels-Real, we report the average performance over the remaining *Desk* and *Plant* scenes. For HDR-NeRF-Real [2], HDR-Plenoxels-Real [3], and the synthetic HDR-NeRF-Syn [2], we follow the previous

---

**Algorithm 1** Training Procedure of PhysHDR-GS

---

**Input:** Reflectance  $H_r$ , illumination  $L_a$ , opacity  $\{\alpha_i\}_{i=1}^N$ , centers  $\{\mu_i\}_{i=1}^N$ , covariances  $\{\Sigma_i\}_{i=1}^N$ , input LDR views  $\{I_{gt}^k\}_{k=1}^K$ , lighting level  $\{l^k\}_{k=1}^K$ , exposure  $\{t^k\}_{k=1}^K$

**Output:** Optimized HDR 3D Gaussian primitives  $G_{HDR}^{3D}$

```
1: Initialize  $H_r, L_a$ 
2: while iteration  $\leq$  MaxIteration do
3:    $I_{gt}, l, t \leftarrow$  SampleTrainingView( $\{I_{gt}^k\}_{k=1}^K, \{l^k\}_{k=1}^K, \{t^k\}_{k=1}^K$ );
4:
5:    $c = g(L_a, H_r)$ ; /* Sec. 4.1: Physical Radiance Composition */
6:    $G_{HDR}^{3D} = \{(\mu_i, \Sigma_i, \alpha_i, c_i)\}_{i=1}^N$ ; /* Image-exposure (IE) branch */
7:    $I_{HDR} \leftarrow$  Rasterize( $G_{HDR}^{3D}$ );
8:    $I_{HDR} \times t \leftarrow$  RescaleExposure( $I_{HDR}, t$ );
9:
10:   $\hat{L}_a = \varphi(L_a, l)$ ; /* Gaussian-illumination (GI) branch */
11:   $\hat{c} = g(\hat{L}_a, H_r)$ ; /* Virtual illumination */
12:   $\hat{G}_{HDR}^{3D} = \{(\mu_i, \Sigma_i, \alpha_i, \hat{c}_i)\}_{i=1}^N$ ; /* Relit Gaussians */
13:   $\hat{I}_{HDR} \leftarrow$  Rasterize( $\hat{G}_{HDR}^{3D}$ );
14:
15:   $\{I_{LDR}^{IG}, I_{LDR}^{GI}, I_{LDR}\} \leftarrow$  ToneMapping( $f, I_{HDR} \times t, \hat{I}_{HDR}$ ); /* Sec. 4.2: Self-Consistent HDR Fusion */
16:   $\mathcal{L}_{rec} \leftarrow$  ReconstructionLoss( $\{I_{LDR}, I_{LDR}^{IG}, I_{LDR}^{GI}\}, I_{gt}$ ); /* Tone-mapped LDR learning */
17:   $\mathcal{L}_{cons} \leftarrow$  HDRConsistencyLoss( $I_{HDR}^{IE}, \hat{I}_{HDR}$ ); /* Self-consistent HDR learning */
18:   $\mathcal{L}_{total} \leftarrow$  WeightedLossSum( $\mathcal{L}_{rec}, \mathcal{L}_{cons}, \mathcal{L}_{unit}$ );
19:   $L_a, H_r, \alpha, c, \mu, \Sigma \leftarrow$  AdamOptimize( $\nabla \mathcal{L}_{total}$ )
20:
21:   $s_a \leftarrow$  GradientScaling( $s, L_a, \hat{L}_a$ ); /* Sec. 4.3: Illumination-Guided Gradient Scaling */
22:  if  $s_a \nabla \mathcal{L}_{total} > \tau_p$  then
23:    Densification;
24:  end if
25: end while
```

---

method [5] to downsample the original resolutions by factors of 1/4, 1/6, and 1/2, yielding training resolutions of  $804 \times 534$ ,  $992 \times 746$ , and  $400 \times 400$ , respectively.

**Implementation details.** During training, we adopt the Adam optimizer [4] to optimize the Gaussian parameters, the radiance composer  $g$ , the illumination modulator  $\varphi$ , and the tone mapper  $f$ . The initial learning rates for  $g$  and  $\varphi$  are set to  $6 \times 10^{-5}$ , and the learning rate for the tone mapper is set to  $2 \times 10^{-4}$ . All learning rates are scheduled with a cosine annealing scheme [6]. The densification threshold  $\tau_p$  is set to  $2 \times 10^{-4}$ .

## 5. Discussions

During training, we observe that LDR views captured in low-light conditions may contain significant sensor noise. Fitting view-specific noise can introduce floaters/thin Gaussians. This can be mitigated by applying a noise-prior guided regularization on per-view reconstruction to discourage fitting high-frequency noise. Meanwhile, stronger geometry backbones (e.g., the adopted Scaffold-GS [7]) can

further suppress floaters by constraining Gaussians around anchor points.

## References

- [1] Yuanhao Cai, Zihao Xiao, Yixun Liang, Minghan Qin, Yulun Zhang, Xiaokang Yang, Yaoyao Liu, and Alan L. Yuille. Hdr-gs: Efficient high dynamic range novel view synthesis at 1000x speed via gaussian splatting. *Advances in Neural Information Processing Systems*, 37:68453–68471, 2024. 1, 3, 4, 5, 6
- [2] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18398–18408, 2022. 1, 3, 4, 5, 6
- [3] Kim Jun-Seong, Kim Yu-Ji, Moon Ye-Bin, and Tae-Hyun Oh. Hdr-plenoxels: Self-calibrating high dynamic range radiance fields. In *European Conference on Computer Vision*, pages 384–401. Springer, 2022. 1, 4
- [4] Diederik P Kingma. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 2

Table 1. Per-scene quantitative comparisons of each scene on HDR-NeRF-Real [2] dataset. For each scene, the best and second-best results are highlighted in red and yellow. LDR-OE and LDR-NE denote the LDR results with exposure  $\{t_1, t_3, t_5\}$  and  $\{t_2, t_4\}$ , respectively.

Method		<i>Box</i>			<i>Computer</i>			<i>Flower</i>			<i>Luckycat</i>			
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
exp3	LDR-OE	HDR-NeRF [2]	35.13	0.9612	0.055	34.28	0.9486	0.076	33.21	0.9525	0.058	34.47	0.9504	0.063
		HDR-GS [1]	36.21	0.9797	0.011	36.02	0.9720	0.017	32.44	0.9603	0.039	34.82	0.9668	0.018
		GaussHDR [5]	37.17	0.9821	0.009	36.24	0.9735	0.015	34.61	0.9681	0.020	36.19	0.9718	0.014
		GaussHDR $\dagger$ [5]	37.18	0.9834	0.008	36.37	0.9737	0.013	35.06	0.9757	0.012	36.67	0.9750	0.012
		Ours	37.94	0.9845	0.007	35.67	0.9761	0.012	33.91	0.9732	0.015	37.23	0.9779	0.010
		Ours $\dagger$	37.88	0.9836	0.007	37.12	0.9746	0.011	35.32	0.9764	0.011	37.32	0.9760	0.009
	LDR-NE	HDR-NeRF [2]	31.20	0.9518	0.081	33.56	0.9487	0.080	30.05	0.9439	0.072	33.81	0.9458	0.063
		HDR-GS [1]	29.04	0.9697	0.024	33.23	0.9696	0.021	29.08	0.9539	0.049	32.75	0.9612	0.021
		GaussHDR	33.18	0.9774	0.013	35.12	0.9716	0.018	30.94	0.9648	0.024	34.72	0.9690	0.015
		GaussHDR $\dagger$ [5]	33.19	0.9784	0.011	35.37	0.9715	0.015	31.71	0.9730	0.014	35.10	0.9722	0.013
		Ours	33.75	0.9794	0.011	34.53	0.9738	0.014	30.51	0.9671	0.020	35.92	0.9747	0.011
		Ours $\dagger$	33.42	0.9779	0.011	35.87	0.9721	0.013	31.68	0.9730	0.014	35.62	0.9719	0.011
exp1	LDR-OE	HDR-NeRF [2]	35.13	0.9612	0.055	34.26	0.9486	0.076	33.18	0.9525	0.058	34.46	0.9504	0.063
		HDR-GS [1]	33.63	0.9692	0.019	34.50	0.9672	0.020	30.14	0.9425	0.051	33.56	0.9597	0.022
		GaussHDR [5]	36.21	0.9780	0.012	34.85	0.9679	0.019	32.78	0.9617	0.024	34.52	0.9637	0.018
		GaussHDR $\dagger$ [5]	35.44	0.9780	0.010	34.38	0.9679	0.016	33.23	0.9695	0.015	35.46	0.9687	0.015
		Ours	36.55	0.9800	0.009	34.48	0.9692	0.016	32.43	0.9638	0.022	35.10	0.9699	0.014
		Ours $\dagger$	35.69	0.9766	0.010	35.07	0.9693	0.013	33.42	0.9685	0.014	35.17	0.9677	0.013
	LDR-NE	HDR-NeRF [2]	31.17	0.9518	0.081	31.17	0.9518	0.081	30.05	0.9439	0.072	33.80	0.9457	0.063
		HDR-GS [1]	27.35	0.9578	0.033	31.20	0.9643	0.025	28.61	0.9403	0.055	31.47	0.9516	0.027
		GaussHDR [5]	33.18	0.9742	0.014	34.31	0.9671	0.021	30.86	0.9605	0.026	33.56	0.9593	0.019
		GaussHDR $\dagger$ [5]	32.90	0.9741	0.013	34.04	0.9670	0.018	31.90	0.9677	0.016	34.32	0.9646	0.017
		Ours	32.33	0.9746	0.013	33.60	0.9677	0.018	30.55	0.9605	0.026	34.06	0.9660	0.015
		Ours $\dagger$	33.02	0.9730	0.013	34.55	0.9683	0.015	32.13	0.9665	0.015	33.92	0.9630	0.015




Figure 1. Qualitative results on LDR views, where residual maps between the rendered results and ground truth are visualized to highlight the difference. HDR-GS struggles to balance between dynamic-range content reconstruction and fine-structure preservation (e.g., characters in the highlight region of the 2nd row are missing), whereas our approach delivers accurate dynamic-range estimations and sharper details.

Table 2. Per-scene quantitative comparisons on the HDR-Plenoxels-Real [3] dataset. HDR-NeRF [2] is not included, as it is evaluated on HDR-Plenoxels-Real. For each scene, the best and second-best results are highlighted in red and yellow. LDR-OE and LDR-NE denote LDR results with exposure sets  $\{t_1, t_3, t_5\}$  and  $\{t_2, t_4\}$ , respectively.

Method		Character			Coffee			Desk			Plant			
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
exp3	LDR-OE	HDR-GS [1]	36.08	0.9799	0.024	28.13	0.9468	0.046	30.07	0.9366	0.039	30.39	0.9403	0.049
		GaussHDR [5]	36.47	0.9813	0.024	29.10	0.9512	0.041	29.79	0.9355	0.038	30.63	0.9439	0.043
		GaussHDR $\dagger$ [5]	39.09	0.9847	0.017	29.82	0.9597	0.029	30.52	0.9417	0.031	32.05	0.9499	0.036
		Ours	39.45	0.9854	0.015	28.94	0.9538	0.035	29.79	0.9351	0.037	30.56	0.9402	0.046
		Ours $\dagger$	40.34	0.9870	0.013	29.51	0.9585	0.026	30.66	0.9428	0.028	31.72	0.9486	0.034
	LDR-NE	HDR-GS [1]	-	-	-	-	-	-	28.12	0.9260	0.044	29.09	0.9309	0.057
		GaussHDR [5]	-	-	-	-	-	-	28.26	0.9272	0.042	29.59	0.9380	0.047
		GaussHDR $\dagger$ [5]	-	-	-	-	-	-	28.61	0.9330	0.034	30.95	0.9451	0.040
		Ours	-	-	-	-	-	-	28.12	0.9259	0.041	29.47	0.9339	0.051
		Ours $\dagger$	-	-	-	-	-	-	28.91	0.9339	0.030	30.61	0.9439	0.037
exp1	LDR-OE	HDR-GS [1]	34.85	0.9773	0.029	25.21	0.9194	0.070	29.24	0.9250	0.044	29.16	0.9275	0.059
		GaussHDR [5]	36.36	0.9801	0.024	28.02	0.9432	0.048	29.08	0.9242	0.044	30.03	0.9355	0.048
		GaussHDR $\dagger$ [5]	38.24	0.9830	0.018	29.31	0.9549	0.032	29.91	0.9331	0.035	31.48	0.9448	0.040
		Ours	38.13	0.9836	0.017	28.69	0.9479	0.040	28.78	0.9240	0.040	29.19	0.9301	0.050
		Ours $\dagger$	38.68	0.9839	0.015	29.54	0.9551	0.029	29.87	0.9326	0.035	31.26	0.9444	0.037
	LDR-NE	HDR-GS [1]	-	-	-	-	-	-	27.12	0.9134	0.052	27.52	0.9167	0.069
		GaussHDR [5]	-	-	-	-	-	-	27.56	0.9163	0.046	28.96	0.9297	0.053
		GaussHDR $\dagger$ [5]	-	-	-	-	-	-	27.21	0.9225	0.038	30.55	0.9405	0.044
		Ours	-	-	-	-	-	-	27.37	0.9152	0.042	27.82	0.9234	0.053
		Ours $\dagger$	-	-	-	-	-	-	27.96	0.9227	0.037	30.17	0.9399	0.040

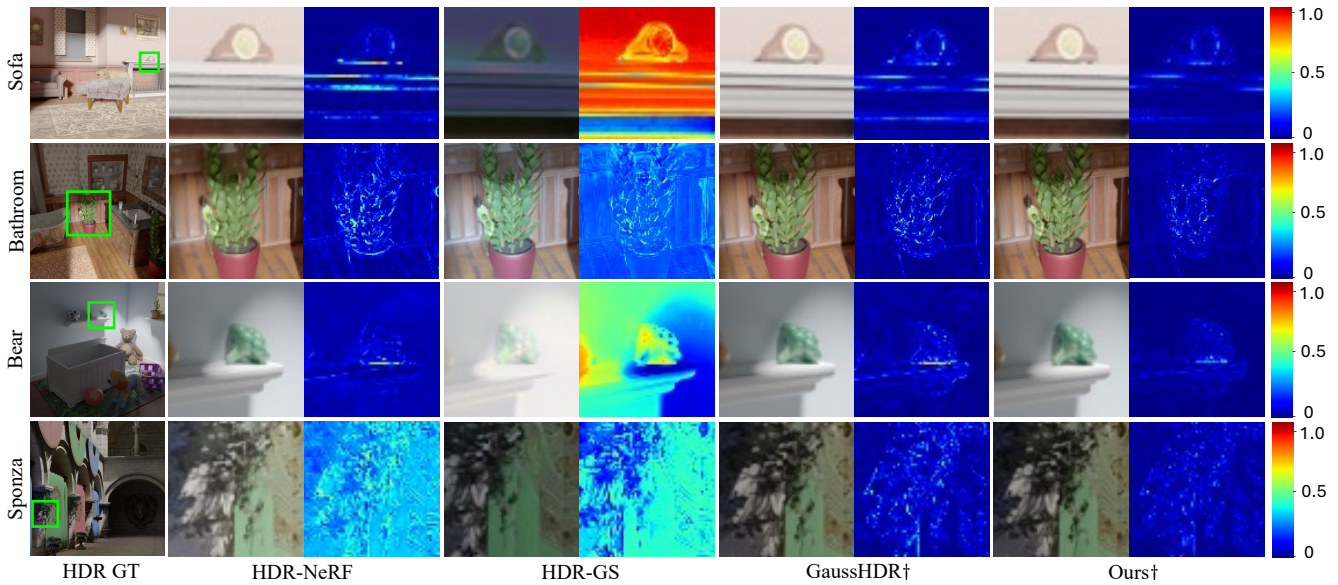


Figure 2. Qualitative results on HDR views, where residual maps between the rendered results and ground truth are visualized to highlight the differences. HDR-NeRF and HDR-GS fail to reach the target illumination level due to the absence of HDR ground truth during training, whereas our results yield accurate dynamic-range estimations and sharper reconstructions (e.g., leaf textures in the 4th row).

[5] Jinfeng Liu, Lingtong Kong, Bo Li, and Dan Xu. Gausshdr: High dynamic range gaussian splatting via learning unified 3d and 2d local tone mapping. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 5991–

6000, 2025. 1, 2, 3, 4, 5, 6  
 [6] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 2

Table 3. Per-scene quantitative comparisons on the HDR-NeRF-Syn [2] dataset under the exp3 exposure setting. For each scene, the best and second-best results are highlighted in red and yellow. LDR-OE and LDR-NE denote LDR results with exposure sets  $\{t_1, t_3, t_5\}$  and  $\{t_2, t_4\}$ , respectively.

Method		<i>Bathroom</i>			<i>Bear</i>			<i>Chair</i>			<i>Desk</i>		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
LDR-OE	HDR-NeRF [2]	33.98	0.9067	0.108	43.96	0.9917	0.007	34.36	0.9286	0.064	38.78	0.9735	0.021
	HDR-GS [1]	42.29	0.9796	0.008	38.61	0.9835	0.015	37.75	0.9734	0.013	43.62	0.9885	0.004
	GaussHDR [5]	42.38	0.9803	0.006	45.60	0.9934	0.002	37.44	0.9671	0.015	42.84	0.9888	0.004
	GaussHDR $\dagger$ [5]	42.97	0.9844	0.006	46.71	0.9945	0.001	38.48	0.9746	0.011	43.77	0.9905	0.003
	Ours	42.34	0.9790	0.006	46.76	0.9942	0.002	37.89	0.9708	0.012	43.78	0.9881	0.003
	Ours $\dagger$	43.25	0.9839	0.005	47.14	0.9945	0.001	38.85	0.9745	0.010	44.50	0.9910	0.002
LDR-NE	HDR-NeRF [2]	33.30	0.9041	0.115	42.46	0.9906	0.009	33.55	0.9214	0.072	39.47	0.9748	0.019
	HDR-GS [1]	37.81	0.9745	0.024	12.70	0.7217	0.196	32.73	0.9561	0.040	37.24	0.9815	0.014
	GaussHDR [5]	42.04	0.9813	0.006	44.92	0.9928	0.002	36.81	0.9651	0.016	42.45	0.9887	0.004
	GaussHDR $\dagger$ [5]	42.49	0.9849	0.006	45.51	0.9937	0.002	37.73	0.9729	0.013	43.12	0.9902	0.003
	Ours	41.96	0.9801	0.006	45.80	0.9935	0.002	37.21	0.9693	0.013	43.31	0.9883	0.003
	Ours $\dagger$	42.39	0.9844	0.005	45.40	0.9937	0.002	38.14	0.9733	0.011	43.12	0.9906	0.002
HDR	HDR-NeRF [2]	23.06	0.9556	0.092	26.35	0.9631	0.024	26.38	0.9555	0.057	44.11	0.9936	0.007
	HDR-GS [1]	20.77	0.8493	0.095	9.10	0.6961	0.241	14.30	0.2937	0.404	29.40	0.8144	0.102
	GaussHDR [5]	35.35	0.9491	0.021	42.09	0.9879	0.006	37.15	0.9604	0.020	44.50	0.9932	0.006
	GaussHDR $\dagger$ [5]	36.45	0.9565	0.016	42.83	0.9891	0.004	38.42	0.9685	0.016	44.86	0.9933	0.005
	Ours	36.05	0.9495	0.016	42.61	0.9887	0.004	38.00	0.9655	0.016	44.83	0.9935	0.004
	Ours $\dagger$	36.63	0.9558	0.015	42.86	0.9891	0.004	38.55	0.9686	0.014	45.08	0.9938	0.004
		<i>Diningroom</i>			<i>Dog</i>			<i>Sofa</i>			<i>Sponza</i>		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
LDR-OE	HDR-NeRF [2]	43.33	0.9894	0.008	39.50	0.9840	0.013	39.83	0.9818	0.013	36.80	0.9701	0.020
	HDR-GS [1]	32.96	0.9254	0.090	41.65	0.9920	0.005	43.45	0.9921	0.003	41.96	0.9901	0.004
	GaussHDR [5]	40.44	0.9813	0.015	43.58	0.9927	0.003	43.82	0.9917	0.003	42.09	0.9868	0.008
	GaussHDR $\dagger$ [5]	45.65	0.9948	0.002	44.59	0.9948	0.002	44.96	0.9934	0.002	43.80	0.9922	0.003
	Ours	44.62	0.9930	0.003	44.34	0.9939	0.002	44.66	0.9931	0.002	40.50	0.9660	0.034
	Ours $\dagger$	45.94	0.9949	0.001	45.18	0.9953	0.002	45.38	0.9942	0.002	43.85	0.9910	0.004
LDR-NE	HDR-NeRF [2]	42.38	0.9892	0.008	38.47	0.9823	0.014	38.70	0.9806	0.013	36.23	0.9702	0.020
	HDR-GS [1]	27.38	0.8968	0.135	15.10	0.7027	0.336	17.30	0.7788	0.258	36.28	0.9832	0.010
	GaussHDR [5]	39.37	0.9795	0.017	42.71	0.9921	0.004	43.11	0.9918	0.003	41.76	0.9884	0.006
	GaussHDR $\dagger$ [5]	44.26	0.9942	0.002	42.13	0.9937	0.003	43.34	0.9930	0.003	43.34	0.9925	0.003
	Ours	43.90	0.9927	0.003	43.31	0.9935	0.002	44.00	0.9931	0.002	40.36	0.9713	0.036
	Ours $\dagger$	44.72	0.9944	0.002	43.78	0.9949	0.002	44.62	0.9942	0.002	43.37	0.9914	0.004
HDR	HDR-NeRF [2]	22.84	0.9360	0.047	18.08	0.8729	0.066	22.01	0.9541	0.058	26.63	0.9523	0.046
	HDR-GS [1]	22.98	0.8760	0.149	20.10	0.8444	0.202	7.94	0.4471	0.347	15.46	0.7648	0.099
	GaussHDR [5]	35.26	0.9664	0.027	36.50	0.9745	0.015	36.97	0.9675	0.015	34.42	0.9642	0.025
	GaussHDR $\dagger$ [5]	39.26	0.9844	0.007	37.52	0.9781	0.011	37.44	0.9704	0.012	35.82	0.9733	0.014
	Ours	38.89	0.9826	0.008	37.73	0.9784	0.011	37.48	0.9705	0.011	34.53	0.9389	0.056
	Ours $\dagger$	39.27	0.9842	0.006	37.85	0.9794	0.010	37.68	0.9718	0.011	35.74	0.9720	0.014

[7] Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-gs: Structured 3d gaussians for view-adaptive rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20654–20664, 2024. 2

Table 4. Per-scene quantitative comparisons on the HDR-NeRF-Syn [2] dataset under the exp1 exposure setting. For each scene, the best and second-best results are highlighted in red and yellow. LDR-OE and LDR-NE denote LDR results with exposure sets  $\{t_1, t_3, t_5\}$  and  $\{t_2, t_4\}$ , respectively.

Method		<i>Bathroom</i>			<i>Bear</i>			<i>Chair</i>			<i>Desk</i>		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
LDR-OE	HDR-NeRF [2]	33.97	0.9066	0.108	43.81	0.9915	0.007	34.31	0.9284	0.064	38.26	0.9701	0.023
	HDR-GS [1]	40.62	0.9734	0.011	36.31	0.9745	0.020	35.21	0.9610	0.022	41.34	0.9847	0.006
	GaussHDR [5]	41.02	0.9731	0.009	44.02	0.9912	0.003	37.01	0.9678	0.014	42.46	0.9872	0.004
	GaussHDR $\dagger$ [5]	42.08	0.9807	0.006	45.40	0.9934	0.002	37.65	0.9712	0.012	43.03	0.9896	0.003
	Ours	40.75	0.9740	0.009	45.42	0.9933	0.002	37.36	0.9695	0.013	42.31	0.9884	0.004
	Ours $\dagger$	42.17	0.9801	0.006	45.49	0.9933	0.002	37.64	0.9712	0.012	42.30	0.9895	0.003
LDR-NE	HDR-NeRF [2]	33.29	0.9040	0.115	42.35	0.9903	0.009	33.51	0.9212	0.072	36.47	0.9611	0.027
	HDR-GS [1]	34.99	0.9653	0.035	12.58	0.7105	0.215	27.47	0.9108	0.087	33.47	0.9697	0.029
	GaussHDR [5]	40.81	0.9749	0.009	43.48	0.9906	0.003	36.54	0.9673	0.015	42.36	0.9874	0.004
	GaussHDR $\dagger$ [5]	41.70	0.9818	0.007	44.43	0.9927	0.002	37.14	0.9698	0.013	42.76	0.9895	0.003
	Ours	40.49	0.9754	0.009	44.36	0.9927	0.002	36.90	0.9686	0.014	42.07	0.9888	0.004
	Ours $\dagger$	41.85	0.9814	0.006	44.50	0.9926	0.002	37.28	0.9703	0.012	42.10	0.9899	0.003
HDR	HDR-NeRF [2]	23.09	0.9558	0.090	26.56	0.9649	0.023	26.36	0.9554	0.057	44.01	0.9932	0.007
	HDR-GS [1]	20.01	0.8370	0.113	14.01	0.7845	0.191	12.64	0.0876	0.591	22.79	0.4358	0.305
	GaussHDR [5]	34.60	0.9401	0.025	41.22	0.9855	0.007	36.76	0.9621	0.020	44.01	0.9931	0.006
	GaussHDR $\dagger$ [5]	36.21	0.9522	0.017	42.22	0.9881	0.005	37.79	0.9659	0.017	43.83	0.9932	0.005
	Ours	35.53	0.9436	0.019	42.12	0.9881	0.005	37.64	0.9653	0.017	44.05	0.9932	0.005
	Ours $\dagger$	36.31	0.9517	0.015	42.31	0.9880	0.004	37.92	0.9663	0.015	43.96	0.9934	0.005
		<i>Diningroom</i>			<i>Dog</i>			<i>Sofa</i>			<i>Sponza</i>		
		PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
LDR-OE	HDR-NeRF [2]	43.29	0.9894	0.008	39.49	0.9840	0.013	39.71	0.9815	0.013	36.61	0.9683	0.020
	HDR-GS [1]	30.17	0.8914	0.130	40.20	0.9888	0.007	40.98	0.9888	0.006	40.91	0.9880	0.005
	GaussHDR [5]	40.47	0.9810	0.017	43.17	0.9921	0.003	43.50	0.9911	0.003	41.73	0.9874	0.005
	GaussHDR $\dagger$ [5]	45.21	0.9942	0.002	43.76	0.9941	0.002	43.52	0.9919	0.003	42.91	0.9911	0.004
	Ours	44.02	0.9926	0.003	43.12	0.9929	0.003	43.56	0.9922	0.003	39.39	0.9649	0.016
	Ours $\dagger$	44.96	0.9941	0.002	43.25	0.9939	0.002	44.08	0.9932	0.002	42.48	0.9896	0.004
LDR-NE	HDR-NeRF [2]	42.36	0.9892	0.008	38.46	0.9822	0.014	38.56	0.9802	0.013	36.04	0.9683	0.021
	HDR-GS [1]	25.62	0.8598	0.175	15.44	0.7160	0.305	18.00	0.7981	0.229	35.75	0.9810	0.013
	GaussHDR [5]	39.62	0.9798	0.020	42.40	0.9916	0.004	42.90	0.9913	0.003	41.56	0.9890	0.005
	GaussHDR $\dagger$ [5]	43.33	0.9934	0.003	42.12	0.9933	0.003	42.04	0.9914	0.003	42.60	0.9916	0.004
	Ours	43.36	0.9925	0.003	42.15	0.9925	0.003	42.96	0.9923	0.002	39.46	0.9707	0.016
	Ours $\dagger$	43.58	0.9938	0.002	42.40	0.9938	0.002	43.49	0.9935	0.002	41.83	0.9906	0.004
HDR	HDR-NeRF [2]	30.17	0.9873	0.017	22.84	0.9359	0.047	18.16	0.8749	0.063	21.68	0.9514	0.062
	HDR-GS [1]	20.88	0.8194	0.205	20.16	0.8476	0.185	9.85	0.6003	0.296	11.10	0.4262	0.234
	GaussHDR [5]	35.49	0.9666	0.029	36.30	0.9734	0.016	36.81	0.9666	0.016	34.36	0.9655	0.021
	GaussHDR $\dagger$ [5]	38.61	0.9832	0.008	37.51	0.9780	0.012	37.13	0.9691	0.013	35.64	0.9720	0.014
	Ours	38.57	0.9817	0.008	37.37	0.9774	0.011	37.22	0.9693	0.011	34.26	0.9418	0.030
	Ours $\dagger$	39.04	0.9837	0.006	37.68	0.9790	0.010	37.46	0.9710	0.011	35.60	0.9718	0.014