

# Supplementary Materials of “Seeing through Light and Darkness: Sensor-Physics Grounded Deblurring HDR NeRF from Single-Exposure Images and Events”

Yunshan Qi<sup>1</sup>    Lin Zhu<sup>2\*</sup>    Nan Bao<sup>1</sup>    Yifan Zhao<sup>1</sup>    Jia Li<sup>1\*</sup>

<sup>1</sup>State Key Laboratory of Virtual Reality Technology and Systems, SCSE & QRI, Beihang University

<sup>2</sup>School of Artificial Intelligence, Beijing Normal University

{qi.yunshan, nbao, zhaoyf, jiali}@buaa.edu.cn, linzhu@bnu.edu.cn

## A. Details of event calibration and division

**Photometric quantity calibration function  $h(\cdot)$ :** For the input event bin  $B(t_i, t_{i+1})$  between  $t_i$  and  $t_{i+1}$ , we use  $(x, y, t_{\text{first}}^i, p_{\text{first}}^i)$  and  $(x, y, t_{\text{last}}^i, p_{\text{last}}^i)$  to represent the first and last triggered event at pixel  $(x, y)$ , respectively. Then we can calculate the offset of each pixel for the input event bin  $B(t_i, t_{i+1})$  with:

$$h(B(t_i, t_{i+1})) = \phi\left(\frac{t_{i+1} - t_{\text{last}}^i}{t_{\text{first}}^{i+1} - t_{\text{last}}^i}, p_{\text{first}}^{i+1}\right) - \phi\left(\frac{t_i - t_{\text{last}}^{i-1}}{t_{\text{first}}^i - t_{\text{last}}^{i-1}}, p_{\text{first}}^i\right), \quad (16)$$

where  $\phi(\cdot)$  is a conditional function to determine whether there are adjacent events triggered before and after the current event bin, and whether the polarities are consistent:

$$\phi(z, p_{\text{first}}^i) = \begin{cases} zp_{\text{first}}^i & \text{if } p_{\text{first}}^i = p_{\text{last}}^{i-1} \\ 0 & \text{if } p_{\text{first}}^i \neq p_{\text{last}}^{i-1} \\ 0.5 & \text{if } \nexists p_{\text{first}}^i \vee \nexists p_{\text{last}}^{i-1} \end{cases} \quad (17)$$

Since the above operations are the same for all pixels, we omit the  $(x, y)$  in Eq. (16) and Eq. (17).

**Events division:** For each training view, we split the input event stream  $B(t_{\text{start}}, t_{\text{end}})$  that corresponding to the input LDR blurry image  $\mathcal{I}_{\text{LDR}}$  into  $b$  bins  $B(t_i, t_{i+1})\}_{i=0}^{b-1}$ , where  $t_0 = t_{\text{start}}$  and  $t_b = t_{\text{end}}$  are the start and end exposure time of the image.  $t_1, \dots, t_{b-1}$  are the time points that divide the event stream into  $b$  event bins with an equal number of events as in the E<sup>3</sup>NeRF [18], which leverages the temporal blur prior in event distribution for event loss optimization, achieving better performance.

\*Correspondence should be addressed to Jia Li and Lin Zhu. Website: <https://cvteam.buaa.edu.cn>

## B. Details of our dataset generation

### B.1. Synthetic dataset:

Our synthetic dataset consists of 8 HDR Blender [1] scenes (“bathroom”, “catroom”, “dogroom”, “diningroom”, “sofa”, “sponza”, “toyroom”, “warmroom”) in the synthetic data of HDR-NeRF [11], which contain both highlight and darkness areas in each scene, resulting in the rendered single-exposure images containing both overexposure and underexposure areas as in the first column of Fig. 12. All synthetic data are at a resolution of  $400 \times 400$ . Below, we introduce the training and test data generation process, respectively.

**Training images and events generation:** For each scene, there are 18 views of LDR blurry images with a single-exposure time  $\Delta t_2$  and corresponding events within the exposure time for training. For each view, we add camera shaking and generate 17 sharp HDR raw images  $\{\mathcal{I}_{\text{HDR}}^i\}_{i=0}^{16}$  at time  $t_i$  with camera poses  $p(t_i)$ . The camera pose  $p(t_i)$  changes over time, indicating the camera shake process. Note that  $t_{\text{start}} \leq t_i \leq t_{\text{end}}$ , where  $t_{\text{start}}$  and  $t_{\text{end}}$  are the start time and end time of the exposure as defined in the main paper. Then the 17 sharp HDR raw images are averaged to obtain the blurred HDR raw image  $\mathcal{I}_{\text{HDR}}^{\text{blur}}$ . We used the classic Reinhard tone mapping [20] algorithm in Eq. (18) to convert the blurred HDR raw image into an LDR blurred image with exposure time  $\Delta t_2$  for training as in HDR-NeRF [11]:

$$\mathcal{I}_{\text{LDR}}^{\Delta t_2} = \left(\frac{\Phi \Delta t_2 \mathcal{I}_{\text{HDR}}^{\text{blur}}}{\Phi \Delta t_2 \mathcal{I}_{\text{HDR}}^{\text{blur}} + 1}\right)^{\frac{1}{2.2}}. \quad (18)$$

$\Phi$  is a scale factor, and we take it as 62.5. The 17 sharp HDR raw images  $\{\mathcal{I}_{\text{HDR}}^i\}_{i=0}^{16}$  are also input into the event simulator v2e [10] with “noisy” option, which involves the event latency and noise simulation into the data generation. Then the events corresponding to the synthetic LDR blurry image  $\mathcal{I}_{\text{LDR}}^{\Delta t_2}$  is obtained:

$$B(t_{\text{start}}, t_{\text{end}}) = \mathbf{v2e}(\{\mathcal{I}_{\text{HDR}}^i\}_{i=0}^{16}). \quad (19)$$

Note that we add a Bayer filter in the v2e to generate the color events as in the DAVIS 346 Color [21] event camera.

**Additional training images:** For each training view, we also use Eq. (18) to tone map the 8-th raw sharp image  $\mathcal{I}_{\text{HDR}}^7$  in  $\{\mathcal{I}_{\text{HDR}}^i\}_{i=0}^{16}$  into LDR domain with exposure times  $\Delta t_0$ ,  $\Delta t_2$  and  $\Delta t_4$ , which are used to train the RGB-based HDR NVS works [3, 11, 15, 24] without deblurring effects and to train HDR-NeRF [11] with multi-exposure sharp LDR images as input for reference.

**Test images generation:** For each scene, there are 17 novel views for testing. For each test view, we render 1 ground-truth (GT) sharp HDR raw image  $\mathcal{I}_{\text{HDR}}^{\text{sharp}}$  with Blender for the HDR novel view synthesis (NVS) test. We also use Eq. (18) to transform the sharp HDR image  $\mathcal{I}_{\text{HDR}}^{\text{sharp}}$  into two GT sharp LDR images with novel exposure times  $\Delta t_1$  and  $\Delta t_3$  for the novel exposure NVS test.

## B.2. Real dataset:

Our real dataset consists of 5 real-world extreme lighting scenes (“bear”, “computer”, “computerhard”, “desktop”, “table”) with both highlight and darkness areas, resulting in the captured single-exposure LDR images containing both overexposure and underexposure areas. The upper part of Tab. 6 shows the minimum and maximum luminance of the scenes in our real data, which contains a very wide dynamic range. All data is obtained by a DAVIS 346 Color camera [21], which can capture spatial-temporal aligned images and corresponding events at resolution  $346 \times 260$ .

**Training images and events generation:** For each scene, there are 16 training views. For each view, we capture 1 LDR blurry image at a single-exposure time  $\Delta t_2$  and corresponding events within the exposure time with a handheld DAVIS 346 Color camera.

**Additional training images:** For each training view, we also capture three sharp LDR images with exposure times  $\Delta t_0$ ,  $\Delta t_2$ , and  $\Delta t_4$  with a tripod-fixed DAVIS 346 Color camera to train the RGB-based HDR NVS works [3, 11, 15, 24] without deblurring effects and to train HDR-NeRF [11] with multi-exposure sharp LDR images as input for reference.

**Test images generation:** For each scene, there are 12 novel views for testing. For each test view, we capture 5 LDR sharp images with exposure times  $\{\Delta t_i\}_{i=0}^4$  with a tripod-fixed DAVIS 346 Color camera. The LDR sharp images with novel exposure times  $\Delta t_1$  and  $\Delta t_3$  are used for the novel exposure NVS test, which is the same as the synthetic dataset. The GT HDR sharp image for the HDR NVS test is obtained by merging the five LDR sharp images with exposure times  $\{\Delta t_i\}_{i=0}^4$  with the classic multi-exposure HDR reconstruction algorithm Debevec [7] as in RawHDR [26].

## C. Implementation details

Our framework is built upon E<sup>2</sup>NeRF [17], and all experiments are conducted on a single RTX 3090 GPU. We set

Table 6. The illumination conditions of our real-world scenes and the exposure time settings of our synthetic and real-world datasets.

	Synthetic Scenes	Real-World Scenes				
		Bear	Computer	Computerhard	Desktop	Table
Min Scene luminance	-	70 lux	50 lux	30 lux	20 lux	10 lux
Max Scene luminance	-	2000 lux	4000 lux	5000 lux	1000 lux	1500 lux
Exposure Time $\Delta t_0$	0.001 s	0.05 s	0.02 s	0.002 s	0.05 s	0.05 s
Exposure Time $\Delta t_1$	0.004 s	0.10 s	0.04 s	0.004 s	0.10 s	0.10 s
Exposure Time $\Delta t_2$	0.016 s	0.15 s	0.06 s	0.006 s	0.15 s	0.15 s
Exposure Time $\Delta t_3$	0.064 s	0.20 s	0.08 s	0.008 s	0.20 s	0.20 s
Exposure Time $\Delta t_4$	0.256 s	0.25 s	0.10 s	0.010 s	0.25 s	0.25 s

the number of sampled points on each ray to 64 and 128 for the coarse and fine networks, respectively. The batch size of sampled rays is set to 512 and 1024 for the real data and synthetic data, respectively, and we take 50k iterations for each scene. The above NeRF network settings are the same for all the compared NeRF-based methods and our See-NeRF in the experiments. For the 3DGS-based methods, the number of iterations is set to 30000, and other parameters are set to their default parameter. For the detailed evaluation of hyperparameters  $b$ ,  $\lambda$ , and training time of our See-NeRF, please refer to Sec. D.2.

### C.1. Exposure time settings

The lower part of Tab. 6 shows the value of exposure times  $\{\Delta t_i\}_{i=0}^4$  of our synthetic and real datasets. For the Real-World-Challenge dataset [18] of the deblurring experiment, the exposure times of the input blurry images are 0.1 s, 0.12 s, 0.1 s, 0.1 s, 0.15 s, for the “corridor”, “lab”, “lobby”, “shelf”, “table” scenes, respectively. We use these exposure times to train and test our See-NeRF to obtain the corresponding deblurred novel view LDR images.

### C.2. Pose estimation

We use the pose generation framework of E<sup>2</sup>NeRF [17] to obtain the poses  $\{p(t_i)\}_{i=0}^b$  corresponding to the blurry images for all the compared methods and our See-NeRF, which ensures the input poses are the same, so the reconstruction result is only affected by the network architecture. Recent RGB-based deblurring NVS works provide different camera pose estimation strategies [13, 22, 25] and dynamic scene pose estimation strategies [2, 14]. Integrating the event supervision into these methods would be a promising future research direction for the ERGB-based HDR deblurring NVS task, as demonstrated in EBAD-NeRF [19].

### C.3. Evaluation of HDR results

For the HDR NVS task, we first use the “Enhancer” and “Compressor” options of Photomatrix Pro [9] to tone map the reconstructed and ground truth HDR images into the LDR domain for the synthetic and real data, respectively. Then, we use the tone-mapped images to quantitatively and qualitatively evaluate our results in our main paper and supplementary materials, as in [3, 11].

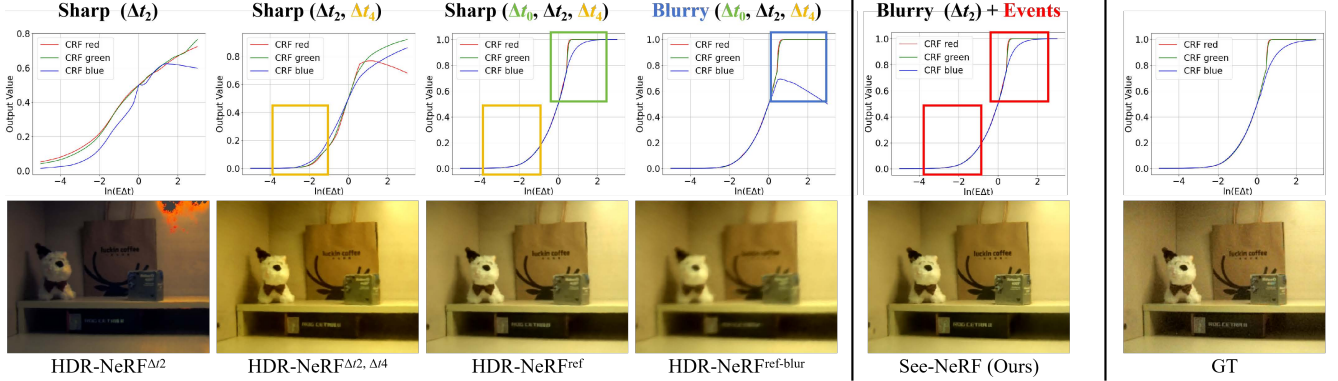


Figure 9. Evaluation of the HDR NVS performance boundaries between multi-exposure RGB-based HDR-NeRF and our single-exposure ERGB-based See-NeRF on the real “Table” scene. The upper and lower parts of the figure are the estimated CRF curves and HDR NVS results, respectively. See-NeRF can effectively leverage events to replace the CRF curve correction effect of the LDR images with exposure times  $\Delta t_0$  and  $\Delta t_4$  in HDR-NeRF<sup>ref</sup>. Besides, See-NeRF is also robust to the image motion blur compared to HDR-NeRF<sup>ref-blur</sup>.

Table 7. Evaluation of the HDR NVS performance boundaries between multi-exposure RGB-based HDR-NeRF and our single-exposure ERGB-based See-NeRF on the real “Table” scene.

	Real “Table” Scene of Our Dataset			HDR Results		
	Exposure Time	Image	Event	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
HDR-NeRF <sup>ref</sup>	$\Delta t_0, \Delta t_2, \Delta t_4$	Sharp	-	28.23	.9207	.1550
HDR-NeRF $\Delta t_2, \Delta t_4$	$\Delta t_2, \Delta t_4$	Sharp	-	21.78	.8404	.1867
HDR-NeRF $\Delta t_2$	$\Delta t_2$	Sharp	-	12.47	.6804	.3551
HDR-NeRF <sup>ref-blur</sup>	$\Delta t_0, \Delta t_2, \Delta t_4$	Blurry	-	25.54	.8640	.3017
See-NeRF	$\Delta t_2$	Blurry	✓	<b>27.12</b>	<b>.9093</b>	<b>.1548</b>

## D. Additional evaluations

### D.1. Single-exposure vs. multi-exposure

In this experiment, we train HDR-NeRF [11] with different input exposure ranges to evaluate the HDR NVS performance boundaries of multi-exposure RGB-based HDR-NeRF and our single-exposure ERGB-based See-NeRF.

**Quantitative results:** HDR-NeRF<sup>ref</sup> in Tab. 7 demonstrates that multi-exposure RGB-based HDR-NeRF only surpasses our single-exposure ERGB-Based See-NeRF in overall performance when the input LDR images are sharp and cover exposure times at all three scales: short ( $\Delta t_0$ ), medium ( $\Delta t_2$ ), and long ( $\Delta t_4$ ). When the input images are motion-blurred, the results of HDR-NeRF<sup>ref-blur</sup> will be significantly affected by the motion blur and degraded. In contrast, our See-NeRF achieves the best HDR NVS results with blurry single-exposure LDR images and corresponding events.

**Qualitative results:** HDR-NeRF $\Delta t_2$  in Fig. 9 illustrates that with images with single-exposure time  $\Delta t_2$  as input, the estimated CRF curve deviates completely from GT, and the HDR NVS results are also completely distorted. When adding the images with longer exposure time  $\Delta t_4$  for training, HDR-NeRF $\Delta t_2, \Delta t_4$  can accurately recover the CRF

curve of the low-light area. At the same time, there is still a deviation in the CRF curve of the highlight area, and the HDR NVS results will show a color cast in the highlight area. When further adding the images with shorter exposure time  $\Delta t_0$ , the CRF curve of the highlight area is further corrected and closer to the GT CRF curve, as shown in the results of HDR-NeRF<sup>ref</sup>. However, in the case of blurry input, the estimated CRF curves of HDR-NeRF<sup>ref-blur</sup> are distorted, and the reconstruction results are blurred. In comparison, See-NeRF achieves accurate CRF curve estimation and produces sharp, high-quality HDR NVS results. This demonstrates that See-NeRF effectively leverages spatial and temporal differential information of events in low-light, highlight, and blurry regions to infer the actual scene brightness to enhance CRF estimation and eliminate the effects of blurred input, which is consistent with the theoretical proof in Sec. 3.

Table 8. Evaluation of parameter  $b$  and training time.

$b$	HDR (“Desktop” scene)			Deblurring (“Lab” scene)			Time hours
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	
2	27.10	.9127	.1930	33.98	.9657	.1458	1.75
4	28.13	.9152	.1796	34.56	.9681	.1374	2.50
6	26.89	.8856	.1842	35.21	.9712	.1277	3.23
8	26.49	.8753	.1834	35.57	.9742	.1193	3.95

### D.2. Parameters and efficiency

**Parameter  $b$ :**  $b$  in Eq. (6) represents the discretization of Eq. (1) with  $b$  discrete time points. Therefore, a larger  $b$  leads to better simulation of the imaging process and generates better results, but at the same time requires more iterations and increases the training time as shown in Tab. 8. For the HDR NVS task, when  $b = 4$ , See-NeRF gets the best results. For the deblurring NVS task, when  $b > 6$ , the improvement is marginal. So we set  $b = 4$  for the HDR and novel exposure NVS experiments on our proposed real-world datasets and

Table 9. Results of the ablation study on previous related CRF models. The results are the averages of all scenes. LPF: Low-Pass Filter

	Event CRF Model	RGB CRF Model	Our Real Data (HDR)			Real-World-Challenge Data		
			PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
See-NeRF <sup>Pow</sup>	ours	Power Function $\mathcal{F}(x) = x^a$ CRF [6]	19.75	.7239	.2696	28.35	.9242	.2657
See-NeRF	ours	MLP-Based CRF (ours)	<b>26.49</b>	<b>.8953</b>	<b>.1638</b>	<b>32.70</b>	<b>.9564</b>	<b>.1574</b>
See-NeRF <sup>v2e</sup>	Explicitly Parameterized 2nd-Order LPF [10]	ours	22.10	.8644	.1938	30.51	.9441	.1748
See-NeRF <sup>Deblur e</sup>	Explicitly Parameterized 4th-Order LPF [16]	ours	24.28	.8859	.1847	31.84	.9521	.1664
See-NeRF	Implicitly MLP-Optimized 2nd-Order LPF (ours)	ours	<b>26.49</b>	<b>.8953</b>	<b>.1638</b>	<b>32.70</b>	<b>.9564</b>	<b>.1574</b>

$b = 6$  for the deblurring NVS experiments for the Real-World-Challenge dataset.

**Training and rendering time:** Compared to HDR-NeRF (2.1h), See-NeRF (2.5h) increases computation and training time by 19.05% under the same conditions (same NeRF-related parameters, number of iterations, and batch size; on a single RTX 3090 GPU; on our real dataset). The increased computation is mainly due to the event-related calculations. However, the rendering speed of See-NeRF (0.37 fps) is slightly faster than HDR-NeRF (0.35 fps) under a resolution of 346\*260. This is mainly because no events are involved during test time, and our 2D pixel-wise tone mapping is more efficient than the 3D points tone mapping of HDR-NeRF.




Figure 10. Evaluation of parameter  $\lambda$ .

**Parameter  $\lambda$ :** As in Fig. 10, we evaluate the impact of  $\lambda$  on the real-world “desktop” and “lab” scenes for the HDR and deblurring NVS tasks, respectively. When  $\lambda = 0.005$ , See-NeRF can achieve the best overall results on both tasks.

### D.3. Sensor-Physics grounded deblurring effect

Our See-NeRF significantly surpasses previous ERGB-based deblurring NVS methods E<sup>2</sup>NeRF [17], E<sup>3</sup>NeRF [18], E2GS [8], and Ev-DeblurNeRF [4] as shown in Tab. 16. This is because these methods all perform the motion blur and events synthesis directly on the nonlinear LDR images rendered by their 3D representation network, as shown in Fig. 11. In contrast, our See-NeRF performs the motion blur and events synthesis on the linear raw HDR values of the scene radiance rendered by the NeRF network, and then tone maps the blurry raw HDR image to the LDR domain, aligning with the physical imaging process of sensors in the real world. On the other hand, our event CRF model eliminates the effects of event delay and further improves the performance. Therefore, See-NeRF achieves SOTA results in the ERGB-based deblurring NeRF task.

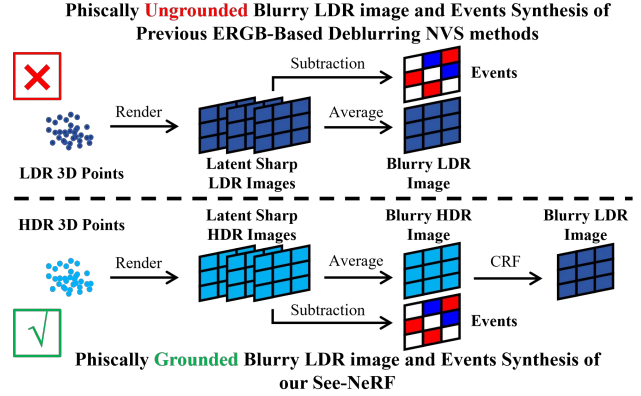


Figure 11. Analysis of sensor-physics grounded deblurring effect.

### D.4. Ablation on previous related CRF models

We conduct an ablation study on previous related CRF models on our real-world dataset and the Real-World-Challenge dataset, respectively, as shown in Tab. 9

**RGB CRF model of EvHDR-GS [6]:** Recently, EvHDR-GS uses 3DGS [12] as a backbone for the ERGB-based HDR video reconstruction task. Though it uses a similar design that RGB CRF is applied after volumetric rendering, it does not consider the image integration process and simply models the RGB CRF with one power function  $\mathcal{F}(x) = x^a$  for the three color channels during the 2D tone mapping, while See-NeRF address these two aspects with Eq. (6) in the main paper and a three-independent MLP-Based RGB CRF model, as shown in the Fig. 3 of the main paper. Therefore, applying its imaging model to See-NeRF degrades the performance significantly on both HDR and deblurring NVS tasks, as shown in the results of See-NeRF<sup>Pow</sup> in Tab. 9.

**Event CRF model of v2e [10] and Deblur e-NeRF [16]:** v2e and Deblur e-NeRF models intensity-dependent pixel bandwidth and latency with the explicitly parameterized 2nd-order and 4th-order low-pass filter, respectively. However, our See-NeRF incorporates scene/camera-related event delay factors into an MLP to calculate the delay coefficient for the 2nd-order low-pass filter. Replacing our event CRF model with the models of v2e and Deblur e-NeRF weakens the performance of See-NeRF on both deblurring and HDR tasks, as in the results of See-NeRF<sup>v2e</sup> and See-NeRF<sup>Deblur e</sup> in Tab. 9, proving the effectiveness of our event CRF model.

## E. Supplementary detailed results

### E.1. Detailed quantitative results

**Our synthetic dataset:** Tab. 10 and Tab. 11 illustrate the detailed quantitative novel exposure NVS results on each scene of our synthetic dataset. See-NeRF achieves the best overall performance. In certain scenes, the novel exposure results of the compared RGB-based HDR NVS methods are better than those of See-NeRF because it uses sharp images as input. Besides, the detailed quantitative results in Tab. 12 and Tab. 13 illustrate that their performance drops significantly on the HDR NVS task, indicating that they are overfitted to the LDR image supervision, thus learning a wrong scene HDR representation. In contrast, See-NeRF achieves the best HDR NVS results for all three metrics.

**Our real dataset:** Tab. 14 and Tab. 15 illustrate the detailed quantitative novel exposure and HDR NVS results on each scene of our real dataset. Some results of the compared RGB-based HDR NVS methods are better than those of See-NeRF on LPIPS due to the sharp images input, while our See-NeRF still achieves the best overall performance. For the HDR experiment, See-NeRF is even better than HDR-NeRF<sup>ref</sup> in “Bear”, “Computerhard”, and “Desktop” scenes, resulting in its average results in Tab. 1 of the main paper outperforming HDR-NeRF<sup>ref</sup>. This is likely attributable to the low imaging quality of the DAVIS 346 RGB sensor, which is not suitable for RGB-based methods.

**Real-World-Challenge dataset:** Tab. 16 illustrates the detailed quantitative deblurring NVS results on each scene of the Real-World-Challenge dataset [18]. Our See-NeRF realizes the SOTA performance compared to the previous RGB-based and ERGB-based deblurring NVS methods.

### E.2. Detailed qualitative results

**Our synthetic dataset:** Fig. 12 shows the detailed qualitative HDR NVS results on our synthetic dataset. The results of See-NeRF are closest to GT across all scenes. HDR-GS suffers from severe color shifts. HDR-NeRF fails to reconstruct the overexposed and underexposed areas. HDR-NeRF+ is limited by the HDREv-Net [23], which only achieves a brightening effect relative to the input image without fully restoring the scene color details. EvHDR-NeRF still suffers from slight color casts and motion blur.

**Our real dataset** Fig. 13 shows the detailed qualitative HDR NVS results on our real dataset, and the results of See-NeRF are closest to GT across all scenes. The compared methods show the same phenomenon as in the synthetic data.

**Real-World-Challenge dataset** Fig. 14 shows the detailed qualitative deblurring NVS results on the Real-World-Challenge dataset. See-NeRF shows impressive deblurring performance compared to other methods. Note that EvHDR-NeRF is significantly affected by the motion blur.

### E.3. Supplementary video results

We also provide supplementary video rendering results to show the impressive HDR and deblurring novel view synthesis results of See-NeRF and the robustness to noise events of our framework. We highly recommend that the reviewers refer to the attached video for details on our project web page: <https://icvteam.github.io/See-NeRF.html>.

## References

- [1] Blender Foundation. Blender 3.4.1, 2023. Software available from <https://www.blender.org>. 1
- [2] Minh-Quan Viet Bui, Jongmin Park, Jihyong Oh, and Munchul Kim. Moblurf: Motion deblurring neural radiance fields for blurry monocular video. *TPAMI*, 2025. 2
- [3] 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. *Nips*, 2024. 2, 6, 7
- [4] Marco Cannici and Davide Scaramuzza. Mitigating motion blur in neural radiance fields with events and frames. In *CVPR*, pages 9286–9296, 2024. 4
- [5] Zehao Chen, Zhanfeng Liao, De Ma, Huajin Tang, Qian Zheng, and Gang Pan. Evhdr-nerf: Building high dynamic range radiance fields with single exposure images and events. In *AAAI*, pages 2376–2384, 2025. 6, 7
- [6] Zehao Chen, Zhan Lu, De Ma, Huajin Tang, Xudong Jiang, Qian Zheng, and Gang Pan. Evhdr-gs: Event-guided hdr video reconstruction with 3d gaussian splatting. In *AAAI*, pages 2367–2375, 2025. 4
- [7] Paul E Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 643–652. 2023. 2
- [8] Hiroyuki Deguchi, Mana Masuda, Takuya Nakabayashi, and Hideo Saito. E2gs: Event enhanced gaussian splatting. In *ICIP*, pages 1676–1682. IEEE, 2024. 4
- [9] HDRsoft. Photomatix pro 7.1.2, 2024. Software available from <https://www.softpedia.com/>. 2
- [10] Yuhuang Hu, Shih-Chii Liu, and Tobi Delbruck. v2e: From video frames to realistic dvs events. In *CVPR*, 2021. 1, 4
- [11] Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. Hdr-nerf: High dynamic range neural radiance fields. In *CVPR*, 2022. 1, 2, 3, 6, 7
- [12] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *TOG*, 42(4):139–1, 2023. 4
- [13] Dogyoon Lee, Minhyeok Lee, Chajin Shin, and Sangyoun Lee. Dp-nerf: Deblurred neural radiance field with physical scene priors. In *CVPR*, pages 12386–12396, 2023. 2
- [14] Jungho Lee, Donghyeong Kim, Dogyoon Lee, Suhwan Cho, Minhyeok Lee, and Sangyoun Lee. Crim-gs: Continuous rigid motion-aware gaussian splatting from motion-blurred images. *arXiv preprint arXiv:2407.03923*, 2024. 2
- [15] 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 *CVPR*, 2025. 2, 6, 7

Table 10. Detailed quantitative novel exposure NVS results on each scene of our synthetic dataset.

		Input		Bathroom	Catroom	Diningroom	Dogroom
	Methods	Image	Event	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$
Reference	HDR-NeRF <sup>ref</sup>	Sharp ( $\Delta t_0, \Delta t_2, \Delta t_4$ )	-	29.69 .9148 .1798	26.56 .9089 .1589	29.02 .9493 .1053	25.20 .8537 .1028
RGB-Based HDR NVS	HDR-NeRF [11]	Sharp ( $\Delta t_2$ )	-	27.13 .9161 .1920	24.56 .8954 .1769	23.68 .9454 .1002	21.11 .8382 .1088
	HDR-GS [3]	Sharp ( $\Delta t_2$ )	-	17.17 .7536 .3440	15.66 .7041 .3088	12.73 .6970 .4790	13.24 .7189 .3377
	Gaussian-DK [24]	Sharp ( $\Delta t_2$ )	-	29.99 .9599 .1580	28.46 .9478 .1636	21.21 .9142 .1409	24.24 .9714 .1593
	GaussHDR [15]	Sharp ( $\Delta t_2$ )	-	15.30 .7288 .1925	14.37 .7367 .1858	12.73 .7483 .2230	12.03 .7115 .2473
ERGB-Based Deblurring HDR NVS	HDR-NeRF+	Blurry ( $\Delta t_2$ )	✓	19.50 .7311 .3197	18.58 .7288 .2970	14.59 .7336 .2699	14.74 .5893 .3586
	EvHDR-NeRF [5]	Blurry ( $\Delta t_2$ )	✓	26.43 .8729 .3124	25.11 .8855 .3249	20.45 .8412 .2847	22.81 .8229 .2903
	See-NeRF	Blurry ( $\Delta t_2$ )	✓	<b>28.13 .9141 .1596</b>	<b>27.45 .9214 .1774</b>	<b>23.28 .9315 .1314</b>	<b>24.81 .8678 .1430</b>

Table 11. Detailed quantitative novel exposure NVS results on each scene of our synthetic dataset.

		Input		Sofa	Sponza	Toyroom	Warmroom
	Methods	Image	Event	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$
Reference	HDR-NeRF <sup>ref</sup>	Sharp ( $\Delta t_0, \Delta t_2, \Delta t_4$ )	-	30.66 .9292 .1139	31.49 .9555 .1161	31.69 .9641 .0694	29.47 .9224 .1657
RGB-Based HDR NVS	HDR-NeRF [11]	Sharp ( $\Delta t_2$ )	-	20.48 .9199 .1248	29.48 .9504 .1252	25.48 .9569 .0934	24.97 .9200 .1583
	HDR-GS [3]	Sharp ( $\Delta t_2$ )	-	12.97 .7777 .2391	18.78 .7352 .4568	15.04 .7682 .3198	15.42 .7873 .3411
	Gaussian-DK [24]	Sharp ( $\Delta t_2$ )	-	24.46 .9787 .1351	20.27 .6563 .3864	28.63 .9681 .1526	29.74 .9609 .1378
	GaussHDR [15]	Sharp ( $\Delta t_2$ )	-	13.79 .7391 .2252	22.34 .8678 .2037	12.49 .7607 .2293	14.50 .7589 .2164
ERGB-Based Deblurring HDR NVS	HDR-NeRF+	Blurry ( $\Delta t_2$ )	✓	18.22 .7045 .2436	18.06 .6032 .2696	17.18 .8291 .1948	20.15 .8232 .2753
	EvHDR-NeRF [5]	Blurry ( $\Delta t_2$ )	✓	24.04 .9048 .2283	24.40 .8551 .2655	25.70 .9142 .2871	26.05 .8728 .2850
	See-NeRF	Blurry ( $\Delta t_2$ )	✓	<b>28.03 .9399 .1075</b>	<b>33.33 .9693 .0929</b>	<b>28.76 .9533 .0937</b>	<b>26.75 .9185 .1572</b>

- [16] Weng Fei Low and Gim Hee Lee. Deblur e-nerf: Nerf from motion-blurred events under high-speed or low-light conditions. In *ECCV*, pages 192–209. Springer, 2024. 4
- [17] Yunshan Qi, Lin Zhu, Yu Zhang, and Jia Li. E<sup>2</sup>nerf: Event enhanced neural radiance fields from blurry images. In *ICCV*, pages 13254–13264, 2023. 2, 4
- [18] Yunshan Qi, Jia Li, Yifan Zhao, Yu Zhang, and Lin Zhu. E<sup>3</sup>nerf: Efficient event-enhanced neural radiance fields from blurry images, 2024. 1, 2, 4, 5
- [19] Yunshan Qi, Lin Zhu, Yifan Zhao, Nan Bao, and Jia Li. Deblurring neural radiance fields with event-driven bundle adjustment. In *ACM MM*, pages 9262–9270, 2024. 2
- [20] Erik Reinhard, Michael Stark, Peter Shirley, and James Ferwerda. Photographic tone reproduction for digital images. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, pages 661–670. 2023. 1
- [21] Gemma Taverni, Diederik Paul Moeys, Chenghan Li, Celso Cavaco, Vasyl Motsnyi, David San Segundo Bello, and Tobi Delbruck. Front and back illuminated dynamic and active pixel vision sensors comparison. *IEEE Transactions on Circuits and Systems II: Express Briefs*, 65(5):677–681, 2018. 2
- [22] Peng Wang, Lingzhe Zhao, Ruijie Ma, and Peidong Liu. Bad-nerf: Bundle adjusted deblur neural radiance fields. In *CVPR*, pages 4170–4179, 2023. 2
- [23] Yixin Yang, Jin Han, Jinxiu Liang, Imari Sato, and Boxin Shi. Learning event guided high dynamic range video reconstruction. In *CVPR*, pages 13924–13934, 2023. 5
- [24] 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. Wiley Online Library, 2024. 2, 6, 7
- [25] Lingzhe Zhao, Peng Wang, and Peidong Liu. Bad-gaussians: Bundle adjusted deblur gaussian splatting. In *ECCV*, pages 233–250. Springer, 2024. 2
- [26] Yunhao Zou, Chenggang Yan, and Ying Fu. Rawhdr: High dynamic range image reconstruction from a single raw image. In *ICCV*, pages 12334–12344, 2023. 2

Table 12. Detailed quantitative HDR NVS results on each scene of our synthetic dataset.

		Input		Bathroom	Catroom	Diningroom	Dogroom
Methods		Image	Event	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$
Reference	HDR-NeRF <sup>ref</sup>	Sharp ( $\Delta t_0, \Delta t_2, \Delta t_4$ )	-	25.96 .8412 .2413	26.09 .8914 .2015	25.45 .9280 .1432	24.65 .8332 .1261
RGB-Based HDR NVS	HDR-NeRF [11]	Sharp ( $\Delta t_2$ )	-	19.06 .7886 .3384	19.84 .8531 .2722	11.15 .4079 .6993	14.65 .6379 .3291
	HDR-GS [3]	Sharp ( $\Delta t_2$ )	-	12.16 .6562 .5850	13.58 .7421 .5033	12.90 .6761 .4857	14.56 .6995 .4755
	Gaussian-DK [24]	Sharp ( $\Delta t_2$ )	-	13.57 .7088 .2689	12.04 .7207 .2380	10.71 .5619 .4420	12.87 .7302 .2583
	GaussHDR [15]	Sharp ( $\Delta t_2$ )	-	13.39 .7976 .1398	10.81 .7321 .1784	12.27 .7389 .1540	10.86 .7789 .1763
ERGB-Based Deblurring HDR NVS	HDR-NeRF+	Blurry ( $\Delta t_2$ )	✓	16.18 .6439 .3847	14.29 .6153 .3434	12.95 .5215 .6971	16.15 .5081 .3807
	EvHDR-NeRF [5]	Blurry ( $\Delta t_2$ )	✓	<b>22.85</b> .7915 .3624	23.56 .8594 .3784	23.11 .8841 .2452	17.62 .7442 .3509
	See-NeRF	Blurry ( $\Delta t_2$ )	✓	20.83 <b>.8282</b> <b>.2801</b>	<b>25.08</b> <b>.8956</b> <b>.2324</b>	<b>24.71</b> <b>.9219</b> <b>.1414</b>	<b>21.99</b> <b>.8638</b> <b>.1757</b>

Table 13. Detailed quantitative HDR NVS results on each scene of our synthetic dataset.

		Input		Sofa	Sponza	Toyroom	Warmroom
Methods		Image	Event	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$	PSNR $\uparrow$ SSIM $\uparrow$ LPIPS $\downarrow$
Reference	HDR-NeRF <sup>ref</sup>	Sharp ( $\Delta t_0, \Delta t_2, \Delta t_4$ )	-	28.34 .9000 .1504	26.24 .8841 .2189	30.66 .9525 .0994	27.94 .8948 .1828
RGB-Based HDR NVS	HDR-NeRF [11]	Sharp ( $\Delta t_2$ )	-	13.45 .7020 .3381	24.93 .8736 .2469	14.60 .8315 .2605	18.50 .8126 .2752
	HDR-GS [3]	Sharp ( $\Delta t_2$ )	-	12.29 .6713 .5124	17.00 .7045 .5966	13.30 .7606 .4901	12.53 .6837 .4284
	Gaussian-DK [24]	Sharp ( $\Delta t_2$ )	-	15.00 .7883 .2829	11.20 .5116 .5422	12.85 .7376 .2544	13.66 .7014 .2782
	GaussHDR [15]	Sharp ( $\Delta t_2$ )	-	13.89 .8521 .1467	16.48 .7860 .2726	12.92 .8112 .1398	13.03 .7470 .1647
ERGB-Based Deblurring HDR NVS	HDR-NeRF+	Blurry ( $\Delta t_2$ )	✓	17.93 .5986 .3015	17.90 .8002 .2967	17.77 .7607 .2310	15.45 .7002 .3215
	EvHDR-NeRF [5]	Blurry ( $\Delta t_2$ )	✓	20.63 .8312 .2983	20.91 .7667 .4215	21.24 .8450 .3948	23.93 .8315 .3056
	See-NeRF	Blurry ( $\Delta t_2$ )	✓	<b>23.55</b> <b>.9117</b> <b>.1447</b>	<b>26.74</b> <b>.9102</b> <b>.2127</b>	<b>26.01</b> <b>.9357</b> <b>.1604</b>	<b>24.13</b> <b>.8746</b> <b>.1850</b>

Table 14. Detailed quantitative novel exposure NVS results on each scene of our real dataset.

		Input		Bear	Computer	Computerhard	Desktop	Table
Methods		Image	Event	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 $\downarrow$
HDR-NeRF <sup>ref</sup>		Sharp ( $\Delta t_0, \Delta t_2, \Delta t_4$ )	-	32.89 .9611 .1208	28.95 .9451 .0675	32.69 .9685 .0440	35.71 .9612 .1353	36.97 .9732 .1072
HDR-NeRF		Sharp ( $\Delta t_2$ )	-	20.05 .9138 .1498	20.95 .8721 .1098	20.44 .8787 .0977	24.65 .9102 .1381	24.82 .9315 .1196
HDR-GS		Sharp ( $\Delta t_2$ )	-	18.35 .8736 .1901	18.76 .7913 .1863	19.19 .8257 .1537	21.03 .8965 .2218	18.83 .8163 .2382
Gaussian-DK		Sharp ( $\Delta t_2$ )	-	25.32 .9541 .1201	23.50 .8821 .1085	21.83 .7411 .1257	30.79 .9440 .1348	30.61 .9619 .1030
GaussHDR		Sharp ( $\Delta t_2$ )	-	25.07 .9454 .1862	24.63 .9236 .1777	23.28 .9101 .1763	30.13 .9454 .1829	30.74 .9561 .1626
HDR-NeRF+		Blurry ( $\Delta t_2$ )	✓	17.13 .7716 .3123	13.30 .5075 .2917	12.60 .3793 .4650	16.93 .6431 .3710	16.55 .6665 .2819
EvHDR-NeRF		Blurry ( $\Delta t_2$ )	✓	22.14 .9059 .1937	23.91 .8760 .1228	21.12 .8187 .1598	25.94 .8984 .1777	25.68 .9160 .1709
See-NeRF		Blurry ( $\Delta t_2$ )	✓	<b>30.24</b> <b>.9591</b> <b>.1202</b>	<b>27.04</b> <b>.9481</b> <b>.0713</b>	<b>24.15</b> <b>.8576</b> <b>.1154</b>	<b>31.84</b> <b>.9536</b> <b>.1155</b>	<b>32.27</b> <b>.9621</b> <b>.1076</b>

Table 15. Detailed quantitative HDR NVS results on each scene of our real dataset.

		Input		Bear	Computer	Computerhard	Desktop	Table
Methods		Image	Event	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 $\downarrow$
HDR-NeRF <sup>ref</sup>		Sharp ( $\Delta t_0, \Delta t_2, \Delta t_4$ )	-	30.79 .9479 .1045	28.08 .8976 .1226	17.03 .6914 .2024	27.15 .9006 .2036	28.23 .9207 .1550
HDR-NeRF		Sharp ( $\Delta t_2$ )	-	10.15 .4789 .4333	10.28 .4684 .4552	9.76 .4989 .4850	11.71 .6101 .3860	12.47 .6804 .3551
HDR-GS		Sharp ( $\Delta t_2$ )	-	17.11 .8773 .3282	14.29 .5777 .5406	9.25 .4036 .6376	15.00 .7993 .4103	15.24 .8140 .4334
Gaussian-DK		Sharp ( $\Delta t_2$ )	-	14.27 .8181 .3253	12.75 .6314 .4308	11.09 .4588 .5935	16.47 .8113 .3343	17.08 .8367 .3436
GaussHDR		Sharp ( $\Delta t_2$ )	-	19.55 .9288 .1939	16.35 .8305 .2451	17.40 .7771 .1689	24.21 .8934 .1432	24.43 .9047 .1102
HDR-NeRF+		Blurry ( $\Delta t_2$ )	✓	21.10 .8996 .2196	18.09 .7855 .2772	12.18 .6128 .4355	18.68 .8443 .3110	23.01 .9039 .1925
EvHDR-NeRF		Blurry ( $\Delta t_2$ )	✓	18.17 .7972 .2138	23.18 .8251 .2216	10.39 .1939 .4011	21.85 .8000 .2465	21.40 .8649 .2229
See-NeRF		Blurry ( $\Delta t_2$ )	✓	<b>31.91</b> <b>.9558</b> <b>.0999</b>	<b>25.85</b> <b>.8982</b> <b>.1379</b>	<b>19.43</b> <b>.7979</b> <b>.2469</b>	<b>28.13</b> <b>.9152</b> <b>.1796</b>	<b>27.13</b> <b>.9093</b> <b>.1548</b>

Table 16. Detailed quantitative deblurring NVS results on each scene of the Real-World-Challenge dataset.

	Methods	Corridor			Lab			Lobby			Shelf			Table		
		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 $\downarrow$
RGB-Based Deblurring NVS	BAD-NeRF	26.24	.8950	.4761	28.86	.9125	.2386	23.06	.8175	.6265	29.61	.8820	.2487	20.39	.7804	.6102
	DP-NeRF	29.64	.9508	.3361	31.40	.9452	.2350	27.52	.9168	.3093	29.92	.9036	.2937	25.79	.8964	.3336
ERGB-Based Deblurring NVS	EBAD-NeRF	28.74	.9288	.4018	30.10	.9354	.2052	26.66	.8867	.4048	29.88	.8965	.2646	23.34	.8574	.3981
	Ev-DeblurNeRF	31.27	.9546	.2960	31.64	.9570	.1790	22.48	.8611	.4917	31.53	.9285	.2141	22.23	.8647	.3975
	E2GS	23.83	.8699	.5516	32.84	.9571	.2131	27.92	.9132	.4155	31.81	.9205	.2348	26.91	.9218	.2807
	E <sup>2</sup> NeRF	30.13	.9584	.2228	32.16	.9581	.1717	28.94	.9258	.2635	31.23	.9109	.2578	27.15	.9196	.2622
ERGB-Based Deblurring HDR NVS	EvHDR-NeRF	29.98	.9444	.3701	27.34	.9078	.3400	27.24	.9064	.3544	28.08	.8639	.3345	23.32	.8575	.4665
	See-NeRF	<b>34.12</b>	<b>.9749</b>	<b>.1313</b>	<b>35.21</b>	<b>.9712</b>	<b>.1277</b>	<b>31.12</b>	<b>.9449</b>	<b>.2248</b>	<b>34.40</b>	<b>.9545</b>	<b>.1065</b>	28.64	<b>.9367</b>	<b>.1965</b>

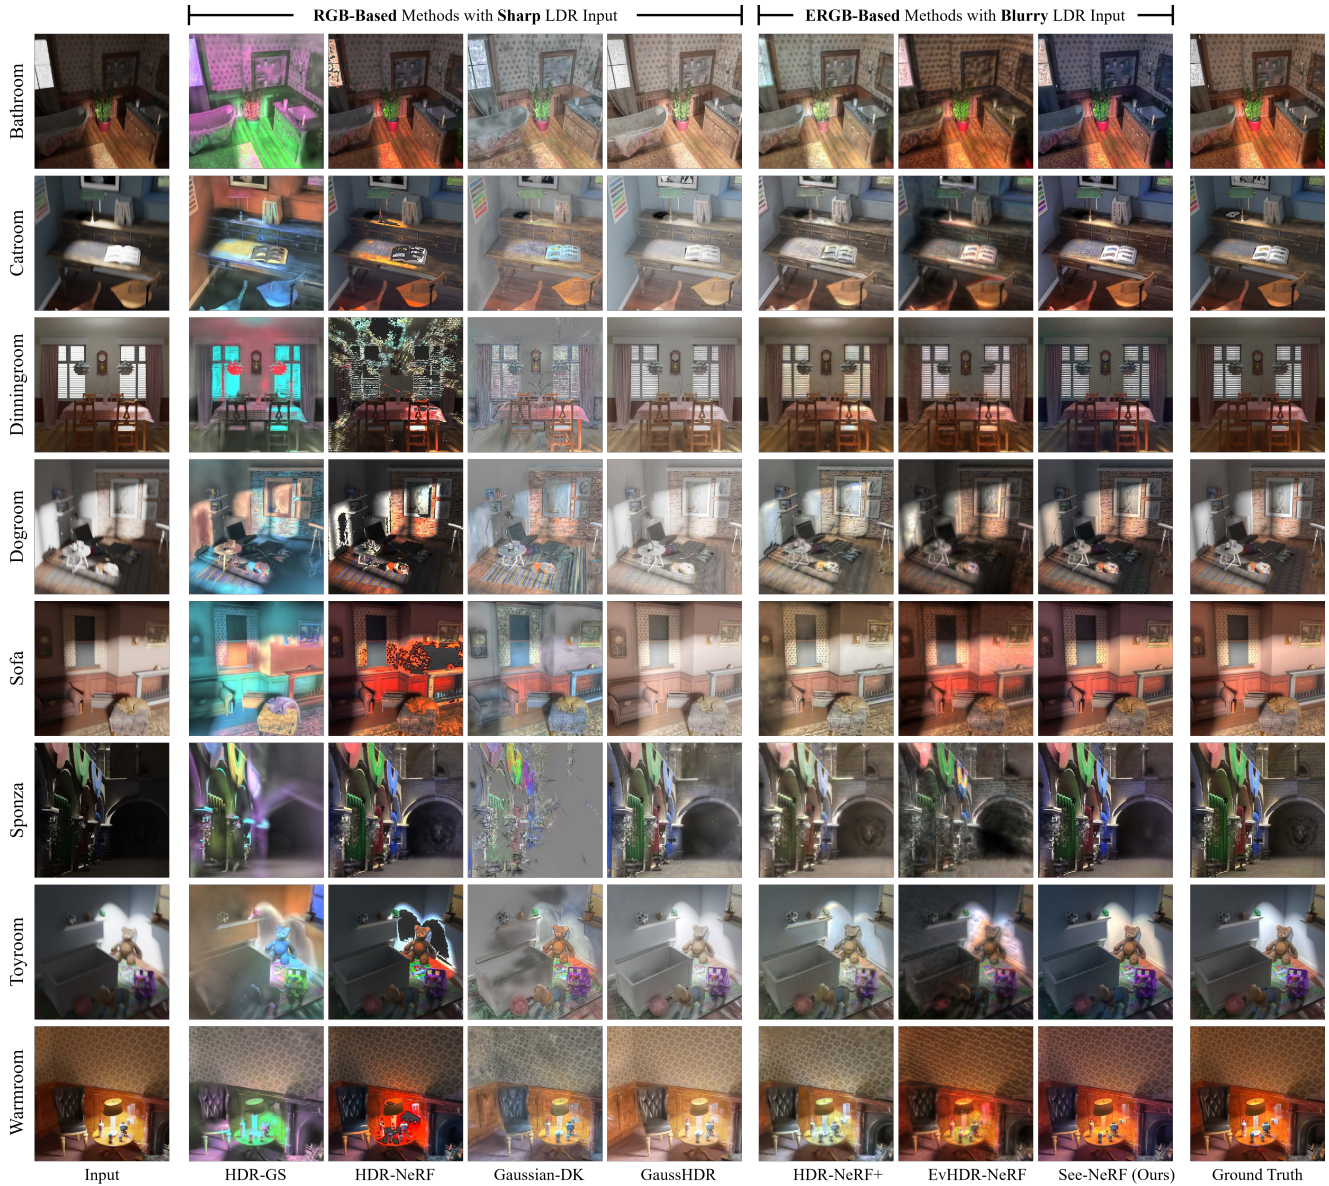


Figure 12. Detailed qualitative HDR NVS results on each scene of our synthetic dataset.

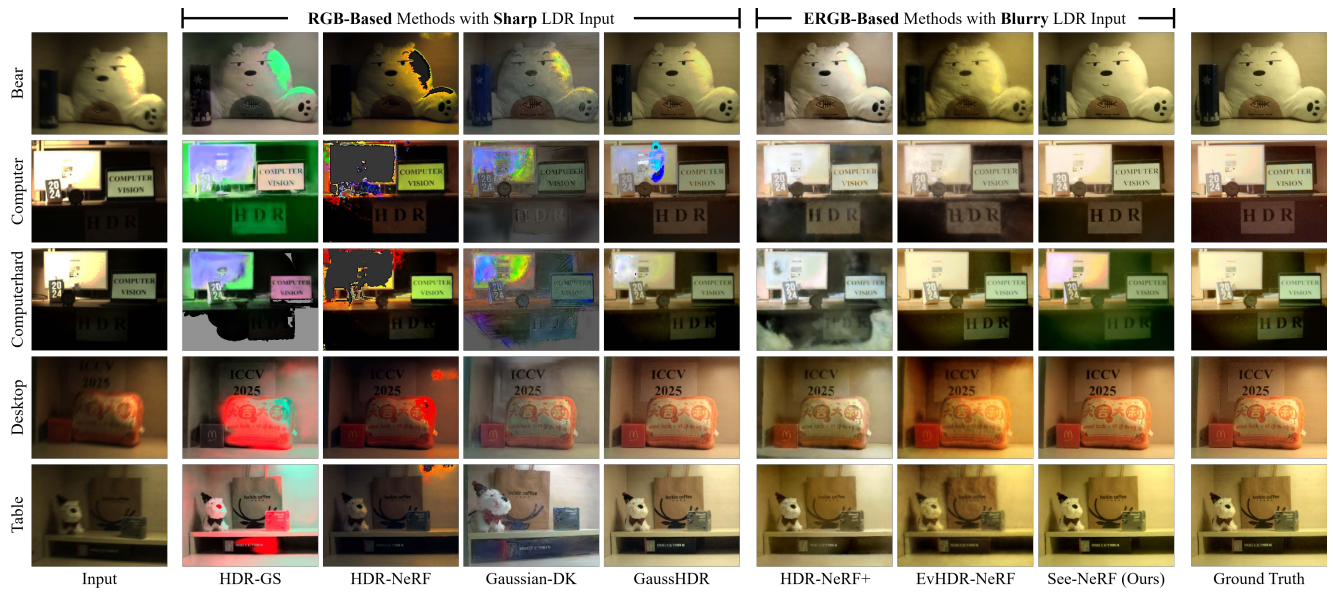


Figure 13. Detailed qualitative HDR NVS results on each scene of our real dataset.

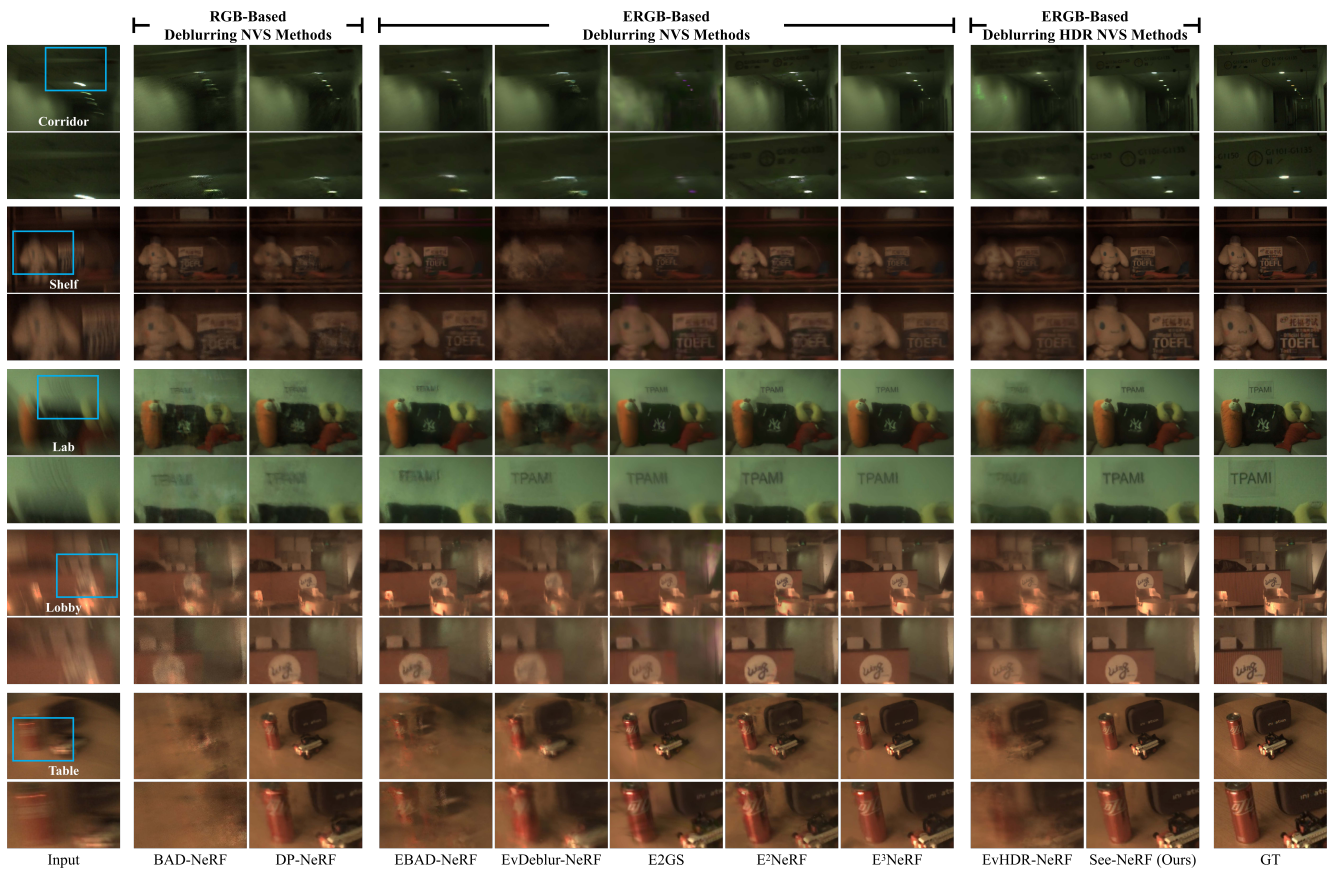


Figure 14. Detailed qualitative deblurring NVS results on each scene of the Real-World-Challenge dataset.