

Supplementary of Luminance-GS: Adapting 3D Gaussian Splatting to Challenging Lighting Conditions with View-Adaptive Curve Adjustment

Ziteng Cui¹, Xuangeng Chu¹, Tatsuya Harada^{1,2}

¹ The University of Tokyo ² RIKEN AIP

(cui, xuangeng.chu, harada)@mi.t.u-tokyo.ac.jp

A. Detailed Experimental Results

Due to page limitations, we could not fully present the per-scene experimental results in the main text. Instead, we provide the complete experimental results in the supplementary part: the per-scene results for the LOM dataset low-light scene are shown in Table. II, it can be observed that our method achieves excellent performance in most scenarios. Additionally, we found that Aleth-NeRF [4] is sensitive to hyper-parameter, for instance, if the enhance degree of Aleth-NeRF is set to a slightly lower value (e.g., reduced from 0.45 to 0.4), its PSNR value significantly decreases as well (see Table. II). The results for overexposure scenes in the LOM dataset are presented in Table. III, and the other 4 scenes results for our synthesized varying exposure dataset are shown in Table. I. We can see that our Luminance-GS both achieves SOTA performance in PSNR, SSIM and LPIPS.

More visualization results are shown in Fig. III, Fig. IV and Fig. V, we can found that sometimes image restoration modules easily lead to multi-view inconsistency, which ultimately causes floaters during rendering (see Fig. IV “MSEC [1] + 3DGS” for example). Meanwhile, our method achieves better detail reconstruction results compared to other approaches. We have zoomed in on random areas to enlarge the details and demonstrate the superior performance of our detail recovery.

B. Ablation Analysis of Loss Functions

We further assess the effectiveness of various loss functions in our Luminance-GS model. Figure II illustrates this through the “sofa” scene, showcasing examples under both low-light and overexposure conditions.

From Fig. II, we can find that spatial loss \mathcal{L}_{spa} (Eq.8 in main text) plays a crucial role in maintaining multi-view consistency, after removing loss \mathcal{L}_{spa} , the rendered scenes tend to exhibit large areas of floaters, which become particularly pronounced under low-light conditions. Meanwhile, in curve loss \mathcal{L}_{curve} (Eq.10 in main text), the cumulative distribution function (CDF) \mathbb{L}_{cdf} of the histogram-equalized



Figure I. The limitation of our method, in some cases, Luminance-GS fails to correctly render colors, resulting in color discrepancies in certain areas and the occurrence of pixelation.

(HE) $C^{in}(x)$ is essential for controlling the degree of illumination. Without \mathbb{L}_{cdf} , efforts to enhance or attenuate illumination are often unsuccessful. Additionally, the predefined curve shape $\mathbb{L}_{po} \cdot \mathbb{L}_s$ (Eq.9 in main text) maintain the smoothness of the generated images, reducing the likelihood of large areas of pixels collapsing into a single value, ensuring the generated images more aligned with human visual perception. Ultimately, with the assistance of all the aforementioned losses, we can achieve satisfactory rendered novel views, as shown in the last column of Fig. II.

C. Experimental Setup

For training settings, we trained Luminance-GS on a single Nvidia Tesla V100 GPU using the Adam optimizer. The learning rates for the various parameters were set as follows:

- For the basic 3DGS parameters $G_i = \{\mu_i, c_i, o_i, \Sigma_i\}$, we adopted the default settings from GS-Splat [10].
- The learning rates for the color adjustment parameters \mathbf{a}_i and \mathbf{b}_i were set to 2.5×10^{-3} .
- The learning rate for the color space mapping matrix \mathcal{M}_k was set to 2.5×10^{-4} , with a weight decay of 1×10^{-5} .
- The global curve \mathbb{L}^g was optimized with a learning rate of 1×10^{-3} and a weight decay of 1×10^{-4} .

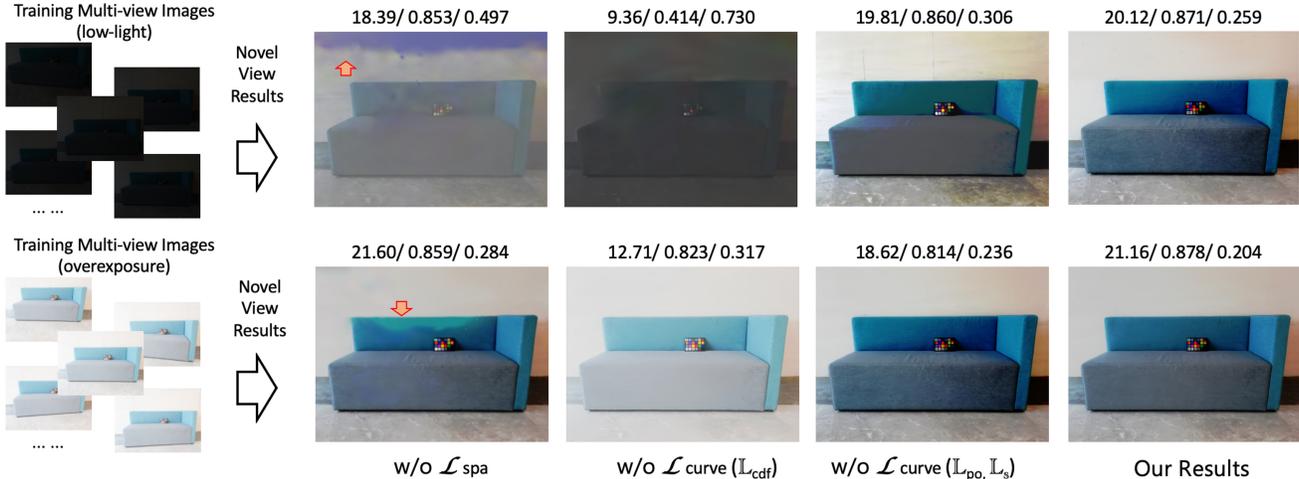


Figure II. Ablation analysis of different loss functions in Luminance-GS, we denote PSNR \uparrow / SSIM \uparrow / LPIPS \downarrow value upon the figures.

Table I. Other 4 scenes results on varying exposure unbounded dataset, we show the PSNR \uparrow , SSIM \uparrow and LPIPS \downarrow . Red indicates the best result, while blue indicates the second-best result.

Methods	“bonsai”	“kitchen”	“room”	“stump”
3DGS	19.05/ 0.699/ 0.403	18.54/ 0.703/ 0.398	19.66/ 0.812/ 0.356	18.12/ 0.698/ 0.366
NeRF-W	14.11/ 0.529/ 0.633	12.87/ 0.451/ 0.502	13.09/ 0.355/ 0.498	14.11/ 0.495/ 0.688
Aleth-NeRF	10.09/ 0.377/ 0.709	9.58/ 0.410/ 0.698	7.22/ 0.308/ 0.742	10.51/ 0.365/ 0.703
GS-W	19.78/ 0.698/ 0.363	18.55/ 0.722/ 0.369	19.32/ 0.794/ 0.386	18.35/ 0.721/ 0.350
Ours	19.77/ 0.709/ 0.351	18.47/ 0.751/ 0.348	20.44/ 0.811/ 0.331	18.64/ 0.732/ 0.344

- The learning rates for the two attention blocks (view-adaptive curve generator and view-adaptive parameters generator, as shown in Fig. 4 of the main text) were set to 1×10^{-5} , with a weight decay of 1×10^{-5} .

Training was conducted for a total of 10,000 iterations, with Gaussian refinement stopping at 8,000 iterations. For other settings, such as Gaussian reset steps [6], we adhered to the default configuration provided by GS-Splat [10].

D. Limitation and Future Discussion

Some failure cases are shown in Fig. I. In certain scenes, Luminance-GS may lose fine details, such as the leaves of plants disappearing (see Fig. I above). Additionally, Luminance-GS sometimes renders incorrect colors and can exhibit pixelated artifacts, as seen in the chair’s color in Fig. I below. This could be due to errors in the pseudo-labels generated by curve \mathbb{L} , and we hope that future research can optimize both the training strategy and the pseudo-label generation solution.

For future research directions, we think that it would be valuable to consider more scenarios of internal camera degradation, such as inconsistent white balance settings. Exploring how to enable scene generalization with

Luminance-GS is also a promising direction. Additionally, we believe that extending Luminance-GS to real-world applications, such as autonomous driving and underground coal mining scenarios, would be highly meaningful.

References

- [1] Mahmoud Afifi, Konstantinos G. Derpanis, Bjorn Ommer, and Michael S. Brown. Learning multi-scale photo exposure correction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9157–9167, 2021. 1, 3, 5
- [2] Jonathan T. Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P. Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields, 2021. 6
- [3] Ziteng Cui, Kunchang Li, Lin Gu, Shenghan Su, Peng Gao, ZhengKai Jiang, Yu Qiao, and Tatsuya Harada. You only need 90k parameters to adapt light: a light weight transformer for image enhancement and exposure correction. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. 3, 5
- [4] Ziteng Cui, Lin Gu, Xiao Sun, Xianzheng Ma, Yu Qiao, and Tatsuya Harada. Aleth-nerf: Illumination adaptive nerf with concealing field assumption. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2024. 1, 3, 4, 5, 6
- [5] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 1780–1789, 2020. 3
- [6] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023. 2, 3, 4, 5, 6
- [7] Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image

Table II. Per-scene experimental results (PSNR \uparrow , SSIM \uparrow , LPIPS \downarrow) on LOM dataset [4] low-light subset, we compare with various enhancement methods [5, 7, 9, 12, 13] and NeRF-based methods [4, 15]. (*: The results of work [15] are directly taken from their paper). Red indicates the best result, while blue indicates the second-best result.

Method	“ <i>buu</i> ”	“ <i>chair</i> ”	“ <i>sofa</i> ”	“ <i>bike</i> ”	“ <i>shrub</i> ”	<i>mean</i>
	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS
3DGS [6]	7.53/ 0.299/ 0.442	6.06/ 0.151/ 0.742	6.31/ 0.216/ 0.723	6.37/ 0.077/ 0.781	8.15/ 0.044/ 0.620	6.88/ 0.157/ 0.662
Image Enhancement Methods + 3DGS						
3DGS + Z-DCE [5]	18.02/ 0.834/ 0.303	12.55/ 0.725/ 0.478	14.66/ 0.822/ 0.460	10.26/ 0.509/ 0.491	12.93/ 0.468/ 0.309	13.64/ 0.672/ 0.408
Z-DCE [5] + 3DGS	17.83/ 0.874 / 0.350	12.47/ 0.762/ 0.399	13.86/ 0.841/ 0.308	10.37/ 0.544/ 0.441	12.74/ 0.487/ 0.248	13.45/ 0.702/ 0.349
3DGS + SCI [7]	13.80/ 0.845/ 0.339	19.70/ 0.812/ 0.455	19.63 / 0.851/ 0.455	12.86/ 0.621/ 0.463	16.14/ 0.600/ 0.442	15.22/ 0.748/ 0.430
SCI [7] + 3DGS	7.68/ 0.690/ 0.523	11.69/ 0.794/ 0.419	10.02/ 0.770/ 0.365	13.55/ 0.667/ 0.390	15.72/ 0.538/ 0.339	11.73/ 0.692/ 0.407
3DGS + NeRCo [9]	16.64/ 0.765/ 0.401	19.24/ 0.759/ 0.466	16.77/ 0.834/ 0.399	16.33/ 0.700/ 0.427	17.07/ 0.503/ 0.411	17.21/ 0.712/ 0.421
NeRCo [9] + 3DGS	16.69/ 0.802/ 0.330	19.11/ 0.773/ 0.376	18.04/ 0.868 / 0.381	16.16/ 0.703/ 0.397	17.97 / 0.502/ 0.399	17.59/ 0.727/ 0.345
Video Enhancement Methods + 3DGS						
LLVE [12] + 3DGS	19.67/ 0.868/ 0.253	15.29/ 0.805/ 0.424	17.18/ 0.858/ 0.379	14.01/ 0.677/ 0.452	15.98/ 0.430/ 0.488	16.43/ 0.728/ 0.399
SGZ [13] + 3DGS	19.21/ 0.832/ 0.270	12.30/ 0.755/ 0.377	14.54/ 0.815/ 0.329	10.61/ 0.563/ 0.375	14.04/ 0.565 / 0.416	14.14/ 0.706/ 0.353
NeRF-based Enhancement Methods						
AME-NeRF* [15]	19.89 / 0.854/ 0.312	17.05/ 0.751/ 0.381	17.93/ 0.847/ 0.378	18.14/ 0.732 / 0.437	15.23/ 0.462/ 0.518	17.65/ 0.729/ 0.405
Aleth-NeRF [4](0.45)	20.22 / 0.859/ 0.315	20.93 / 0.818 / 0.468	19.52/ 0.857/ 0.354	20.46 / 0.727/ 0.499	18.24 / 0.511/ 0.448	19.87 / 0.754 / 0.417
Aleth-NeRF [4](0.4)	19.14/ 0.839/ 0.306	16.96/ 0.793/ 0.483	16.97/ 0.847/ 0.367	17.56/ 0.719/ 0.468	17.55/ 0.484/ 0.451	17.64/ 0.736/ 0.415
Our Proposed Method						
Luminance-GS	18.09/ 0.877 / 0.193	19.82 / 0.835 / 0.367	20.12 / 0.871 / 0.259	18.27 / 0.749 / 0.411	15.40/ 0.666 / 0.241	18.34 / 0.799 / 0.294

Table III. Per-scene experimental results (PSNR \uparrow , SSIM \uparrow , LPIPS \downarrow) on LOM dataset [4] overexposure scene, we compare with exposure correction methods [1, 3, 14] and NeRF-based methods [4]. Red indicates the best result, while blue indicates the second-best result.

Method	“ <i>buu</i> ”	“ <i>chair</i> ”	“ <i>sofa</i> ”	“ <i>bike</i> ”	“ <i>shrub</i> ”	<i>mean</i>
	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS	PSNR/ SSIM/ LPIPS
3DGS [6]	6.96/ 0.674/ 0.609	11.14/ 0.790/ 0.362	10.17/ 0.790/ 0.369	9.58/ 0.730/ 0.323	10.34/ 0.646/ 0.299	9.64/ 0.726/ 0.392
Exposure Correction Methods + 3DGS						
3DGS + MSEC [1]	16.03/ 0.806 / 0.517	20.81/ 0.851 / 0.408	20.65/ 0.862/ 0.397	22.10/ 0.826/ 0.305	18.21/ 0.678/ 0.289	19.56/ 0.805/ 0.382
MSEC [1] + 3DGS	15.08/ 0.804/ 0.440	16.63/ 0.797/ 0.416	20.09/ 0.828/ 0.335	17.57/ 0.739/ 0.368	16.61/ 0.666/ 0.255	17.20/ 0.767/ 0.363
3DGS + IAT [3]	15.34/ 0.804/ 0.522	21.96 / 0.833/ 0.292	20.23/ 0.872 / 0.402	22.36/ 0.832/ 0.291	21.24 / 0.765/ 0.226	20.23/ 0.821/ 0.347
IAT [3] + 3DGS	15.86/ 0.803/ 0.387	18.61/ 0.830/ 0.367	17.42/ 0.833/ 0.348	19.17/ 0.801/ 0.235	16.74/ 0.731/ 0.219	17.56/ 0.800/ 0.311
3DGS + MSLT [14]	15.34/ 0.798/ 0.473	21.69/ 0.823/ 0.304	23.05 / 0.830/ 0.317	23.37/ 0.830/ 0.317	18.89 / 0.779/ 0.214	20.39 / 0.815/ 0.345
MSLT [14] + 3DGS	16.35/ 0.805/ 0.333	20.93/ 0.828/ 0.275	21.65 / 0.847/ 0.259	24.03 / 0.841 / 0.244	18.29/ 0.797 / 0.199	20.25/ 0.824 / 0.262
NeRF-based Exposure Correction Method						
Aleth-NeRF [4]	16.78 / 0.805/ 0.611	20.08/ 0.820/ 0.499	17.85/ 0.852/ 0.458	19.85/ 0.773/ 0.392	15.91/ 0.477/ 0.483	18.09/ 0.745/ 0.488
Our Proposed Method						
Luminance-GS	19.67 / 0.811 / 0.311	22.63 / 0.856 / 0.207	21.16/ 0.878 / 0.204	24.05 / 0.851 / 0.216	16.04/ 0.780 / 0.173	20.71 / 0.835 / 0.222

enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5637–5646, 2022. 3, 4

Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections. In *CVPR*, 2021. 6

[8] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi,

[9] Shuzhou Yang, Moxuan Ding, Yanmin Wu, Zihan Li, and

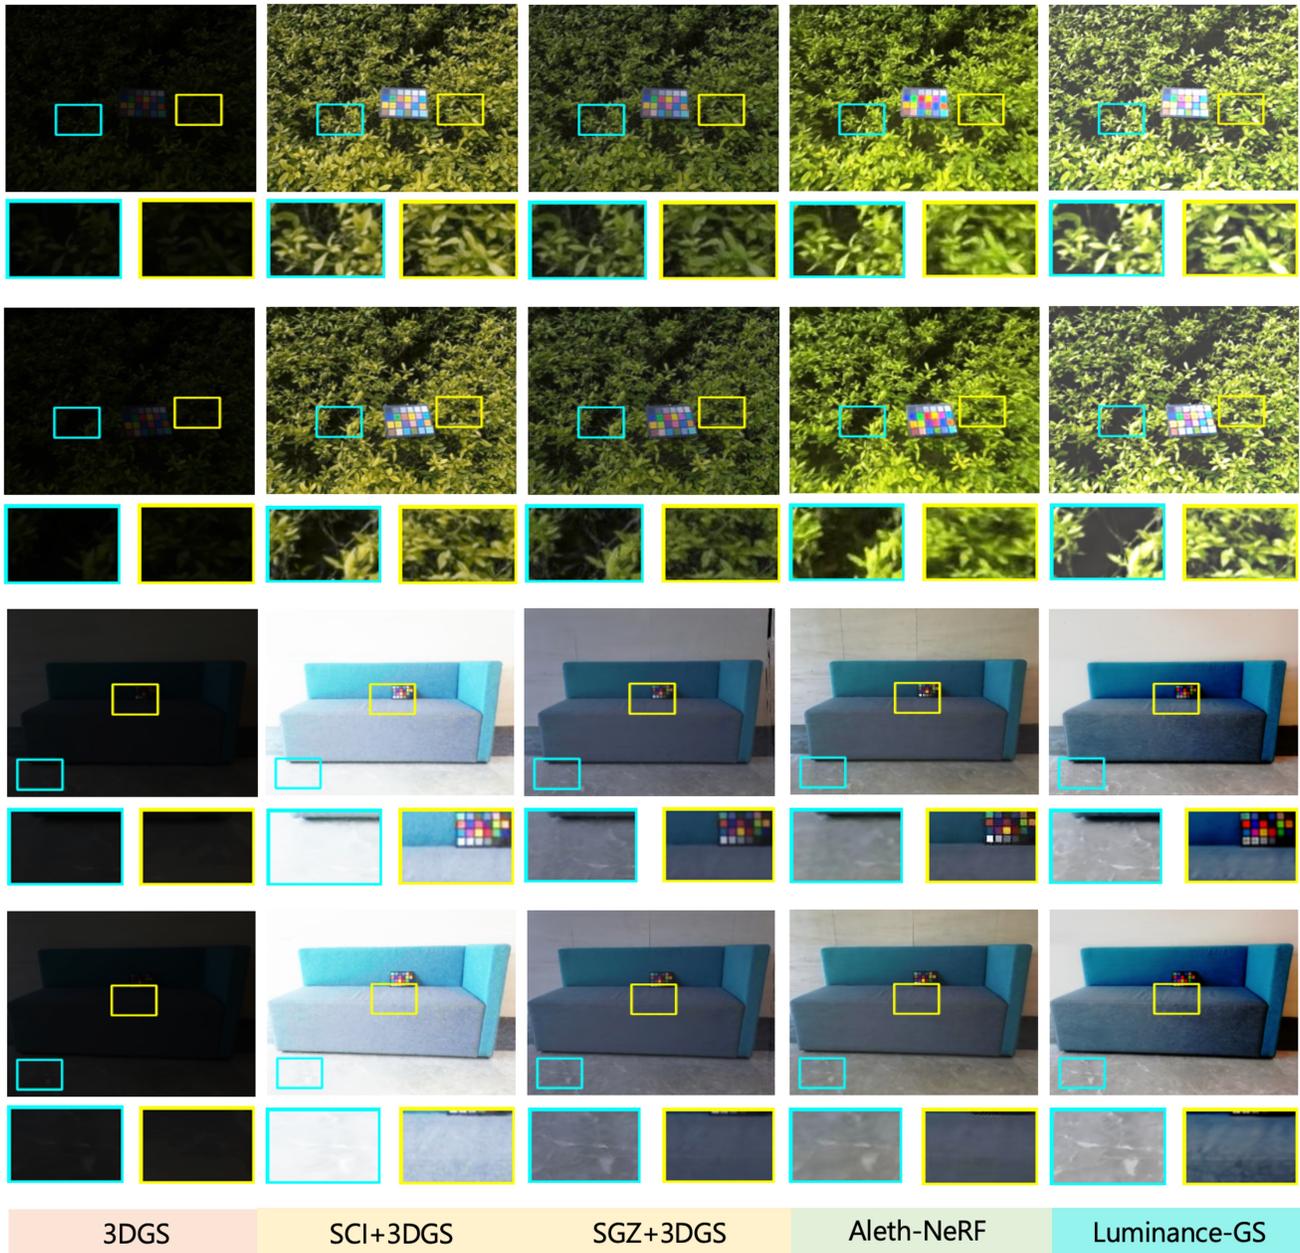


Figure III. Novel view synthesis results in LOM dataset low-light “*shrub*” and “*sofa*” scenes, we show the comparison results with 3DGS [6], combination of low-light enhancement methods (SCI [7], SGZ [13]) with 3DGS and Aleth-NeRF [4].

Jian Zhang. Implicit neural representation for cooperative low-light image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12918–12927, 2023. 3

[10] Vickie Ye, Ruilong Li, Justin Kerr, Matias Turkulainen, Brent Yi, Zhuoyang Pan, Otto Seiskari, Jianbo Ye, Jeffrey Hu, Matthew Tancik, and Angjoo Kanazawa. gsplat: An open-source library for Gaussian splatting. *arXiv preprint*

arXiv:2409.06765, 2024. 1, 2

[11] Dongbin Zhang, Chuming Wang, Weitao Wang, Peihao Li, Minghan Qin, and Haoqian Wang. Gaussian in the wild: 3d gaussian splatting for unconstrained image collections. *ECCV*, 2024. 6

[12] Fan Zhang, Yu Li, Shaodi You, and Ying Fu. Learning temporal consistency for low light video enhancement from single images. In *Proceedings of the IEEE/CVF Conference*



Figure IV. Novel view synthesis results in LOM dataset over-exposure “*chair*” and “*bike*” scenes, we show the comparison results with 3DGS [6], combination of exposure correction methods (MSEC [1], IAT [3]) with 3DGS and Aleth-NeRF [4].

on *Computer Vision and Pattern Recognition (CVPR)*, pages 4967–4976, 2021. 3

[13] Shen Zheng and Gaurav Gupta. Semantic-guided zero-shot learning for low-light image/video enhancement. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 581–590, 2022. 3, 4

[14] Yijie Zhou, Chao Li, Jin Liang, Tianyi Xu, Xin Liu, and Jun Xu. 4k-resolution photo exposure correction at 125 fps with 8k parameters. In *Winter Conference on Applications of Computer Vision (WACV)*, 2024. 3

[15] Yang Zou, Xingyuan Li, Zhiying Jiang, and Jinyuan Liu. Enhancing neural radiance fields with adaptive multi-exposure fusion: A bilevel optimization approach for novel view synthesis. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 7882–7890, 2024. 3

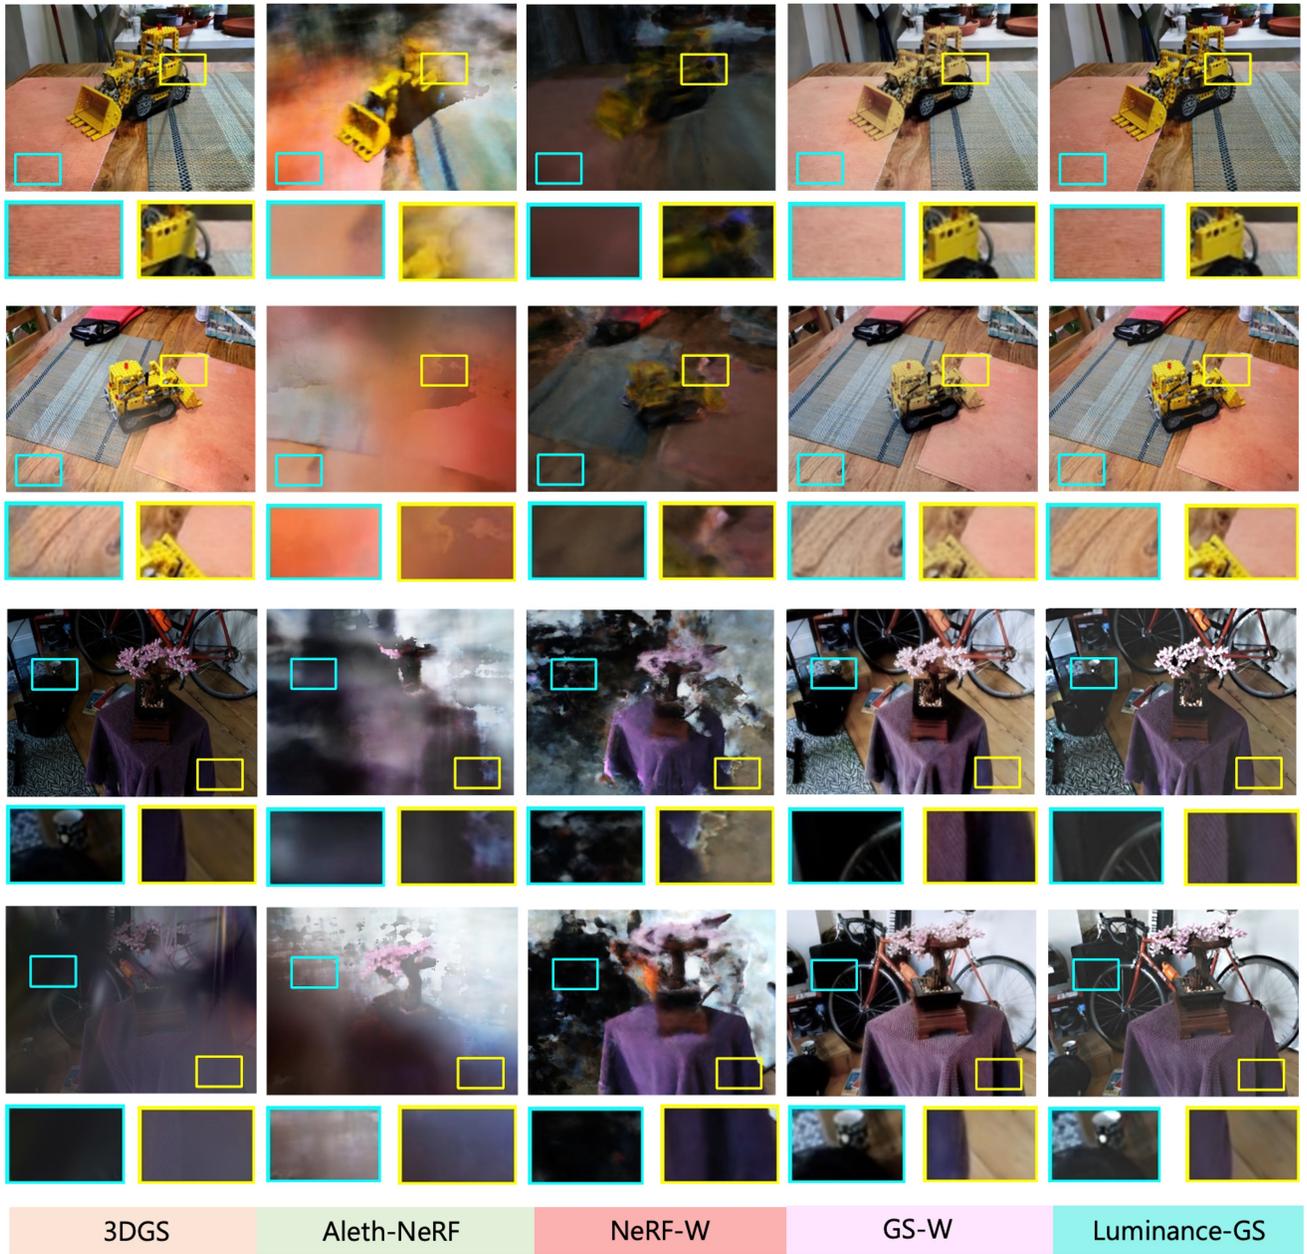


Figure V. Novel view synthesis results on our synthesized varying exposure unbounded dataset (from Mip-NeRF 360 dataset [2]) “*kitchen*” and “*bonsai*” scenes, with comparison of 3DGS [6], Aleth-NeRF [4], NeRF-W [8] and GS-W [11].