CoLight-GS: Coupled Radiometric Supervision for Low-Light Gaussian Splatting

Anonymous authors
Paper under double-blind review

Abstract

Low-light novel-view synthesis aims to reconstruct a well-exposed 3D scene from dark multi-view images. Exposure correction alters both the reconstruction targets and the scales of image residuals, linking brightness recovery to how scene details are fitted. We present CoLight-GS, a Gaussian reconstruction framework that couples bright-domain structural supervision with dark-domain observation fitting through a fixed invertible tone map. The map lifts dark observations into bright targets, while its analytic inverse maps the same scene rendering back to the observed domain, jointly supervising a single Gaussian scene. A shared gain defines a coherent exposure convention across views with matching capture conditions, while source-histogram anchoring accommodates heterogeneous exposures. Derivative-based weighting coordinates the residual scales of the bright and observed domains. CoLight-GS requires no pretrained depth or restoration models and renders novel views directly with standard Gaussian splatting. On RealX3D and LOM, CoLight-GS outperforms five compared methods in mean PSNR, SSIM, and LPIPS, reaching 18.10 dB and 23.45 dB PSNR, respectively.

1 Introduction

Low-light novel-view synthesis reconstructs a well-exposed scene from posed dark photographs. Beyond increasing brightness, it must preserve object boundaries and surface detail across viewpoints. Ordinary Gaussian splatting fits the recorded intensities and reproduces their low exposure. Recovering a bright, coherent scene therefore requires supervision that connects dark observations to the desired appearance.

Recent methods integrate appearance recovery into scene reconstruction. Aleth-NeRF and I2-NeRF model exposure through concealing or media interactions (Cui et al., 2024; Liu et al., 2025); LITA-GS introduces structural priors and denoising (Zhou et al., 2025); Luminance-GS adapts color curves across views (Cui et al., 2025). These strategies motivate a complementary question: how should an exposure correction supervise the reconstructed scene? The correction determines both the bright target and the scale assigned to an image error. Exposure assignment and structural fitting are therefore coupled design choices.

We present CoLight-GS, which couples exposure assignment and scene fitting through a fixed invertible tone map. The map lifts each dark observation into a canonical bright coordinate. The Gaussian scene renders there, and the analytic inverse maps that same rendering back to the observed coordinate, allowing both residuals to jointly supervise the scene. For captures with matching exposure conditions, a shared gain defines a coherent exposure convention across views; source-histogram anchoring accommodates heterogeneous captures. This formulation requires no pretrained depth or restoration models.

The inverse-map derivative relates rendering errors in the two domains and sets their relative loss weight. The bright branch fits local structure at the lifted intensity scale; the dark branch ties that structure to the recorded observation. An opacity–footprint penalty controls broad, opaque primitives. These operations turn exposure correction into a constraint on one reconstructed scene. Figure 1 shows recovered tool boundaries and backing detail on BlueHawaii.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

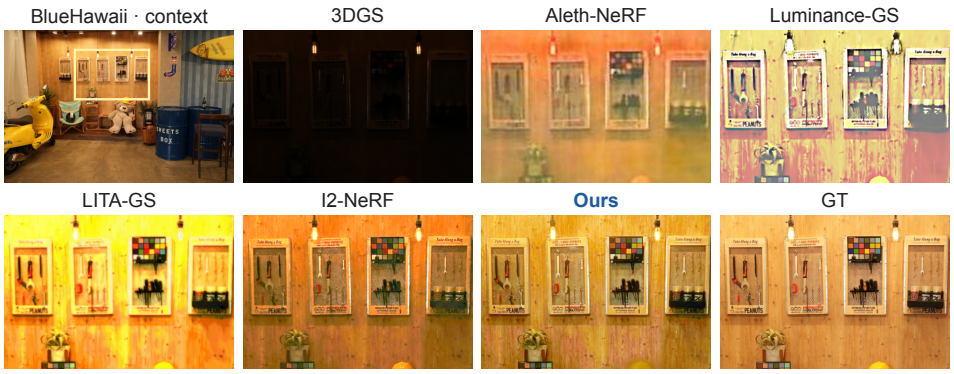


Figure 1: BlueHawaii tool detail across five official baselines, CoLight-GS and GT. The box in GT marks the crop shown for every method.

CoLight-GS obtains the best mean PSNR, SSIM and LPIPS on both benchmarks, with 18.10 dB on RealX3D and 23.45 dB on LOM. Controlled comparisons isolate exposure sharing under a fixed loss and coupled fitting under a fixed exposure map.

Our contributions are threefold:

- We formulate a radiometric coupling that supervises a single Gaussian scene in both the bright and observed domains through a fixed invertible tone map.
- We introduce derivative-based weighting to balance residual scales across domains, together with an opacity–footprint penalty to regularize primitive allocation.
- We demonstrate improved reconstruction on RealX3D and LOM, with controlled comparisons showing the benefits of exposure sharing and coupled fitting.

2 Related Work

2.1 Novel-view synthesis under adverse illumination

Radiance fields (Mildenhall et al., 2020; Barron et al., 2021; 2022), hash encodings (Müller et al., 2022) and dynamic representations (Pumarola et al., 2021; Luiten et al., 2024) support view synthesis on multiview benchmarks (Jensen et al., 2014; Mildenhall et al., 2019).

Illumination-aware fields recover appearance during reconstruction: RawNeRF works in RAW (Mildenhall et al., 2022), Aleth-NeRF uses a concealing field (Cui et al., 2024), and I2-NeRF models low light through virtual absorption (Liu et al., 2025). Related fields address decomposition, sharpening and RAW color (Wang et al., 2023a; Qu et al., 2024; Wang et al., 2025).

LITA-GS uses illumination-invariant structure, depth guidance and progressive denoising (Zhou et al., 2025); Luminance-GS jointly optimizes exposure and view-adaptive color curves (Cui et al., 2025). LL-Gaussian and SplatBright enhance Gaussian representations (Sun et al., 2025; Wen et al., 2025). CoLight-GS couples one rendering to bright targets and dark observations through a fixed analytic map and its inverse. The map sets the exposure, and its derivative sets the relative weight of the two residuals.

2.2 Restoration and Gaussian optimization

Enhancement uses illumination priors (Guo et al., 2017; Wei et al., 2018; Zhang et al., 2019) or learned restoration (Chen et al., 2018; Lore et al., 2017; Jiang et al., 2021; Zamir et al., 2023). Radiometric (Debevec & Malik, 1997; Grossberg & Nayar, 2004) and atmospheric models (Narasimhan & Nayar, 2002; He et al., 2011) describe acquisition.

Deblurring reconstruction addresses motion-corrupted inputs (Ma et al., 2022; Wang et al., 2023b; Cao et al., 2025); GC-GS controls gradients across restored views (Cao et al., 2026). HQGS guides

primitive placement and frequency supervision (Lin et al., 2025). Gaussian fitting admits stochastic relocation (Kheradmand et al., 2024), recurrent updates (Xu et al., 2026) and acceleration (Hahlbohm et al., 2026). CoLight-GS retains ordinary rendering (Kerbl et al., 2023), coupling analytic image targets with opacity–footprint control.

3 Method

3.1 Overview

Given posed dark images, CoLight-GS optimizes a Gaussian scene whose direct render is bright. Figure 2 distinguishes the transformed target from this rendering: one loss compares them, while a second compares the inverse-mapped render to the original observation. We define the radiometric coordinate, balanced image objective and primitive control in turn. A histogram rule provides exposure adaptation for heterogeneous captures.

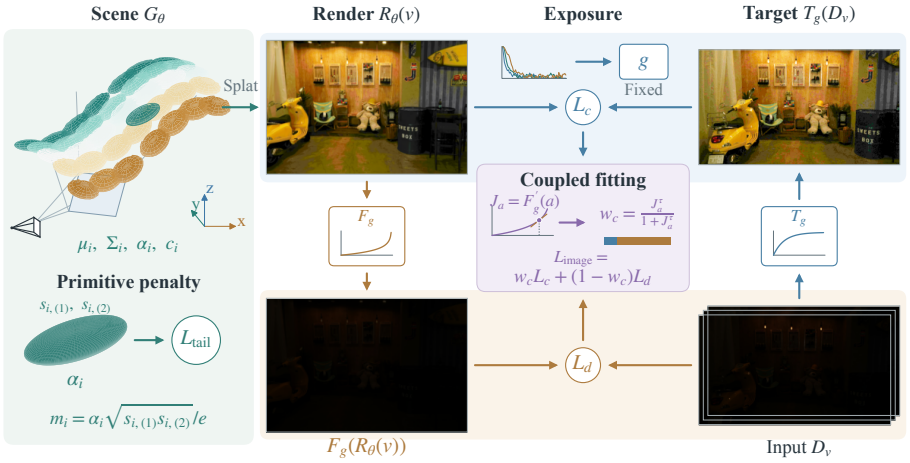


Figure 2: One Gaussian scene, two residual coordinates. The analytic lift supplies a bright target; its inverse returns the same rendering to the dark observation.

3.2 A common radiometric coordinate

Let $D_v \in [0,1]^{H\times W\times 3}$ be a dark training image at camera v, and let $R_\theta(v)$ denote a canonical RGB rendering of Gaussian scene parameters θ. We use the componentwise map

$$T_g(d) = 1 - \exp(-gd), \qquad F_g(r) = -\frac{\log(1-r)}{g}, \qquad g > 0. \tag{1}$$

T_g lifts the observed image, and F_g maps the canonical render back to its observed coordinate. Within their mutual range, $F_g(T_g(d)) = d$. The exponential map avoids the abrupt clipping of a linear gain. For numerical stability, the inverse input is clamped to $[0, 1-10^{-6}]$ and its output to $[0,1]$. The operators act on decoded, display-encoded RGB, with g specifying the brightness coordinate of the recovered render.

To choose a shared coordinate, let q_* be a robust upper-luminance statistic of the dark capture collection and let a be its desired brightness anchor. We set

$$g_* = -\log(1-a)/q_*. \tag{2}$$

We form q_* by aggregating image statistics within each scene and then across scenes. This gives scenes equal influence regardless of their number of views. The resulting scalar is fixed during optimization and shared by the scene renderings; the concrete statistic and anchor are specified in Experimental Setup.

A per-view percentile gain would map every view's own upper luminance quantile to the same anchor. That operation responds to scene contents as well as exposure. A camera turning toward a bright

object changes the percentile even when its capture settings remain fixed. Sharing the gain removes this degree of freedom from supervision: an RGB value has the same canonical target wherever it appears. The common map preserves this correspondence as geometry, occlusion and directional reflectance change across views.

3.3 Radiometrically coupled supervision

We compare the scene rendering to both image coordinates using

$$\ell(A, B) = 0.8\|A - B\|_1 + 0.2\big(1 - \mathrm{SSIM}(A, B)\big), \tag{3}$$

where the L_1 term averages over pixels and channels. Canonical and observed-coordinate losses are

$$L_c = \ell(R_\theta(v), T_g(D_v)), \qquad L_d = \ell(F_g(R_\theta(v)), D_v). \tag{4}$$

An exact canonical fit is also an exact observed fit. With finite capacity, noise and a finite optimization budget, the losses rank approximate fits differently. This matters most where small dark residuals correspond to large visible differences after lifting.

For a scalar dark pixel d and a small canonical perturbation ϵ, the observed residual satisfies

$$F_g(T_g(d) + \epsilon) - d = \frac{\epsilon}{g(1 - T_g(d))} + O(\epsilon^2). \tag{5}$$

Thus the same canonical error receives an intensity-dependent multiplier in the dark coordinate. A fixed mixture of the two losses can itself become unbalanced as the gain changes. We evaluate the inverse derivative at anchor a and set

$$J_a = \frac{1}{g(1 - a)}, \qquad w_c = \frac{J_a^\tau}{1 + J_a^\tau}, \qquad L_{\text{image}} = w_c L_c + (1 - w_c) L_d. \tag{6}$$

At $\tau = 1$, $w_c = (1 - w_c)J_a$, balancing the local scalar derivative factors at the anchor. Because SSIM and Gaussian projection couple neighboring pixels, this balance fixes only a local residual scale (Appendix B). A temperature below one moderates the weighting toward an equal mixture. The weight is fixed for a given gain and anchor.

The bright branch presents local contrast at a different scale, while the observed branch retains the original intensity comparison. Both depend on the same Gaussian parameters: changing a visible boundary also changes its predicted dark appearance. Equal weighting, derivative weighting and canonical-SSIM removal in the component study test how these choices affect the reconstructed scene.

The coupling acts through the rendered image before gradients reach the scene. Each branch contributes an image-space residual, and both contributions pass through the same rasterization Jacobian to update position, scale, opacity and color. The inverse branch additionally applies the inverse-map derivative F_g'. Thus the lifted neighborhoods and recorded intensities influence a common allocation of scene primitives. A Gaussian spanning an object boundary receives structural supervision at the bright scale together with the intensity constraint on the same prediction mapped back to the observation.

3.4 Opacity–footprint regularization

The renderer reduces residuals through colors, opacities and the allocation and extent of Gaussian primitives. To control the influence of opaque, broad primitives in this allocation, we form an opacity–footprint quantity. For primitive i, let α_i be opacity, $s_{i,(1)} \geq s_{i,(2)}$ the two largest positive scale axes, and e the scene extent. Define

$$m_i = \alpha_i \frac{\sqrt{s_{i,(1)} s_{i,(2)}}}{e}, \qquad L_{\text{tail}} = \frac{\frac{1}{N}\sum_i m_i^2}{\mathrm{sg}\big(\max\big\{\frac{1}{N}\sum_i m_i, 10^{-12}\big\}\big)}. \tag{7}$$

Here sg stops gradients through the normalization. The proxy combines primitive extent and opacity in a scene-normalized coordinate. We call primitives with large m_i carriers; squaring the numerator gives them more influence; the detached mean normalizes its magnitude without a compensating denominator gradient. The full objective is

$$L = L_{\text{image}} + \lambda_{\text{tail}} L_{\text{tail}} + L_{\text{reg}}, \tag{8}$$

where L_{reg} contains the standard opacity and scale regularizers. The penalty differentiates through opacity and the two largest scale axes of each primitive. It changes the cost of broad carriers; the densification schedule itself is unchanged. Setting $L_{\text{image}} = L_d$ and $\lambda_{\text{tail}} = 0$ gives the inverse-only variant.

3.5 Exposure anchors for heterogeneous scenes

For heterogeneous captures, we select the anchor before fitting the scene. First compute the scene's percentile gain with Equation 2, using that scene's statistic in place of q_*. When this gain exceeds a fixed threshold, use a mean anchor instead: solve for the gain whose mapped training RGB values reach a target mean. This scalar equation is monotone and can be solved by bisection on the source-image histogram. The selected anchor then determines both the map and its derivative weight. This preserves one coupled objective while adapting the visible appearance requested for each scene.

Percentile anchoring fixes an upper luminance level; mean anchoring fixes average mapped RGB. The choice is resolved from the source histogram before scene optimization. At inference, only the canonical Gaussian rendering is needed.

4 Experiments

4.1 Experimental setup

We evaluate normal-light novel views from dark posed training images. RealX3D (Liu et al., 2026) uses the official low-light split: nine scenes, 267 dark training views and 52 test views. Each scene has six test views except GearWorks and Ujikintoki, which have five. LOM (Cui et al., 2024) has five scenes with 160 dark training views and 21 test views. The target is the corresponding normal-light image at each test camera. Reconstruction uses the supplied poses and sparse geometry.

All main runs use 15,000 iterations. RealX3D has a one-million-primitive cap; LOM has a 200,000-primitive cap and starts from official point-cloud positions with neutral RGB. Spherical harmonics reach degree three. Densification follows a five-percent schedule, every 100 steps from step 600, with the same optimizer settings across variants. Appendix A gives remaining optimizer settings and the scene counts.

We report PSNR, SSIM (Wang et al., 2004) and VGG-based LPIPS (Zhang et al., 2018) using the original-3DGS evaluation protocol, averaging views within each scene and then scenes equally. On both datasets we compare 3DGS, Aleth-NeRF, Luminance-GS, LITA-GS and I2-NeRF. RealX3D uses the official nine-scene NVS averages. LOM uses published five-scene results (Zhou et al., 2025; Cui et al., 2025); the I2-NeRF average is computed from its five reported scene scores (Liu et al., 2025). In benchmark and RealX3D ablation tables, bold and underline denote best and second-best values; displayed ties share a rank.

For RealX3D, q_* is the median across scenes of their median per-image 90th-percentile luminance, computed from dark training views at width 500. Anchor $a = 0.8$ gives $q_* = 0.05186$ and $g_* = 31.03$. We set $\tau = 0.75$ and $\lambda_{\text{tail}} = 0.1$.

For LOM, the 90th luminance percentile is anchored at 0.8, the mapped RGB mean at 0.49, and the switching threshold at gain 32. We use $\tau = 1$ without a tail penalty.

4.2 Benchmark comparisons

CoLight-GS leads the five compared methods on all three metrics in both datasets (Table 1). On RealX3D, it reaches 18.10 dB, 0.5902 SSIM and 0.3709 LPIPS. LITA-GS is next in PSNR and LPIPS at 15.57 dB and 0.4880; I2-NeRF is next in SSIM at 0.5680. On LOM, CoLight-GS reaches 23.45 dB, 0.8350 SSIM and 0.2290 LPIPS. I2-NeRF is next in PSNR at 22.87 dB, while LITA-GS is next in SSIM and LPIPS at 0.8290 and 0.2330.

Figures 1 and 3 retain all five official baselines at matched cameras. Each case provides a full GT view to locate the common region, followed by identical crops from every method. BlueHawaii shows tool silhouettes and wire-mesh backing. Cupcake tests chair slats, tableware and overlapping

Table 1: Low-light novel-view synthesis.

Method	**RealX3D** 9 scenes			**LOM** 5 scenes		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
3DGS	6.660	0.05800	0.6590	7.070	0.1500	0.7130
Aleth-NeRF	12.99	0.4450	0.7040	19.87	0.7540	0.4170
Luminance-GS	10.05	0.4330	0.7080	18.34	0.7990	0.2940
LITA-GS	15.57	0.5420	0.4880	21.14	0.8290	0.2330
I2-NeRF	15.51	0.5680	0.5320	22.87	0.8056	0.2488
Ours	**18.10**	**0.5902**	**0.3709**	**23.45**	**0.8350**	**0.2290**

Table 2: RealX3D ablation of gain sharing, image supervision and primitive control. C-SSIM: canonical structural loss; Tail: opacity–footprint penalty.

Gain	Fitting	τ	C-SSIM	Tail	PSNR ↑	SSIM ↑	LPIPS ↓
None	Dark	–	–	–	7.391	0.06563	0.6771
Per-view	Inverse	–	–	–	17.47	0.5258	0.3914
	Equal	–	✓	–	17.56	0.5407	0.3834
	Derivative	1	✓	–	17.47	0.5788	0.3959
	Derivative	1	–	–	15.99	0.5543	0.5093
	Derivative	0.75	✓	–	17.64	0.5655	0.3811
	Derivative	0.75	✓	✓	17.60	0.5630	0.3781
Shared	Inverse	–	–	–	**18.10**	0.5532	0.3887
	Derivative	0.75	✓	✓	**18.10**	**0.5902**	**0.3709**

boundaries; Popcorn tests the elephant's trunk against shelving and a striped wall. CoLight-GS retains distinct contours and local tonal variation, while the comparisons reveal blur, flattened color and oversaturation in different baselines.

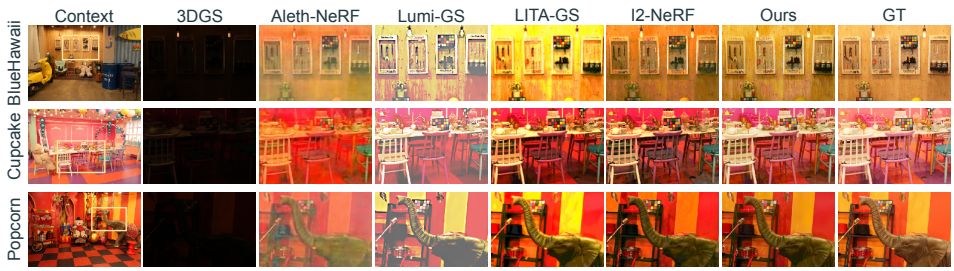


Figure 3: Matched details from BlueHawaii, Cupcake and Popcorn: tool silhouettes, chair slats and tableware, and the elephant against shelving. GT contexts locate each shared crop.

4.3 Exposure assignment and structural fitting

Table 2 shows two distinct improvements. Changing the per-view gain to the shared gain raises the full model's PSNR by 0.50 dB, SSIM by 0.027 and improves LPIPS by 0.007. The same direction holds for inverse fitting: sharing raises PSNR by 0.63 dB and SSIM by 0.027, with a smaller LPIPS improvement of 0.003.

At the same shared gain, full and inverse-only fitting reach 18.10 dB macro PSNR. Adding the derivative-weighted canonical branch and opacity–footprint control improves SSIM by 0.037 and reduces LPIPS by 0.018.

The two paired comparisons address different stages of reconstruction. Sharing changes the targets assigned to the cameras while retaining the fitting objective. Coupled fitting then changes how the

shared target is represented by the scene. Its structural gains accompany the same macro PSNR, locating the improvement in the arrangement of image content: boundaries, texture and local contrast. This is the intended role of the bright-coordinate structural loss, whose neighborhood statistics act on the lifted appearance throughout fitting.

Figure 4 separates these effects across nine scenes. Coupled fitting improves SSIM in every scene, with gains from 0.012 to 0.055, and reduces LPIPS in seven. MilkCookie gains 0.041 SSIM and reduces LPIPS by 0.042. Figure 4 links this structural improvement to matched renderings: fine boundaries and surface detail change even when the macro PSNR is nearly unchanged.

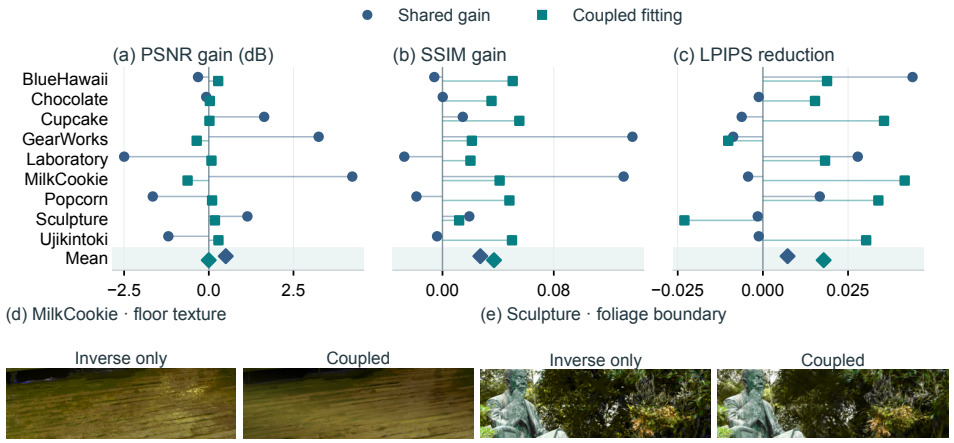


Figure 4: Exposure sharing and coupled fitting. (a–c) Nine-scene changes; diamonds are equal-scene means. (d,e) Matched inverse-only and coupled-fitting details show floor texture and foliage boundaries.

4.4 Structural supervision and primitive allocation

The per-view rows of Table 2 test the image losses at a fixed gain convention. Equal branch weights improve all three metrics over inverse-only fitting. Derivative scaling at temperature one reaches 0.5788 SSIM; temperature 0.75 reaches 17.64 dB PSNR and 0.3811 LPIPS. Figure 5 shows each component's effect by scene.

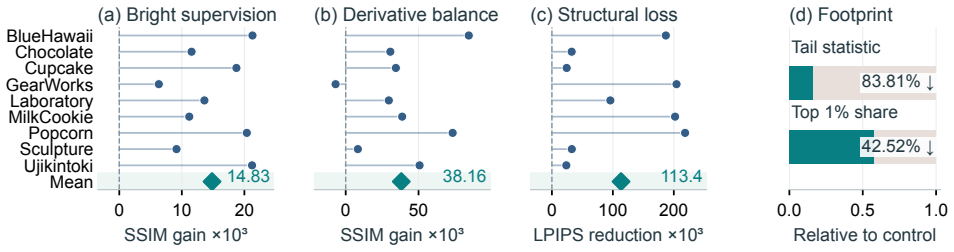


Figure 5: Per-scene effects of each loss component and of the tail penalty. (a–c) All nine scene differences and equal-scene means. (d) Tail statistic and top-one-percent carrier share relative to the no-penalty control.

Canonical SSIM raises PSNR from 15.99 to 17.47 dB and lowers LPIPS from 0.5093 to 0.3959 in the temperature-one objective. Both configurations retain the canonical pixel term and observed-domain objective. This comparison identifies the structural contribution of the bright branch: neighborhoods that are compressed in the dark image become more influential in the lifted coordinate.

Adding the opacity–footprint penalty to the temperature-0.75 objective lowers LPIPS from 0.3811 to 0.3781, with PSNR remaining near 17.60 dB. The tail statistic L_{tail} falls from 0.01581 to 0.002560 and the top-one-percent carrier share from 0.3410 to 0.1960. The penalty acts on the distribution

of primitive extent and opacity, complementing the image-domain changes tested above. Figure 6 connects the radiometric coordinate, its local sensitivity and these carrier measurements.

The two carrier statistics capture different aspects of this response. The tail statistic measures the normalized quadratic mass of the opacity–footprint proxy, so large carriers contribute disproportionately. The top-one-percent share measures concentration within the largest carriers. Their joint reduction shows that the regularized fit changes both the overall tail and its concentration. This matches the gradient of Equation 7: with its denominator detached, the derivative with respect to each carrier is proportional to that carrier's magnitude. Broad, opaque primitives consequently receive the strongest direct penalty.

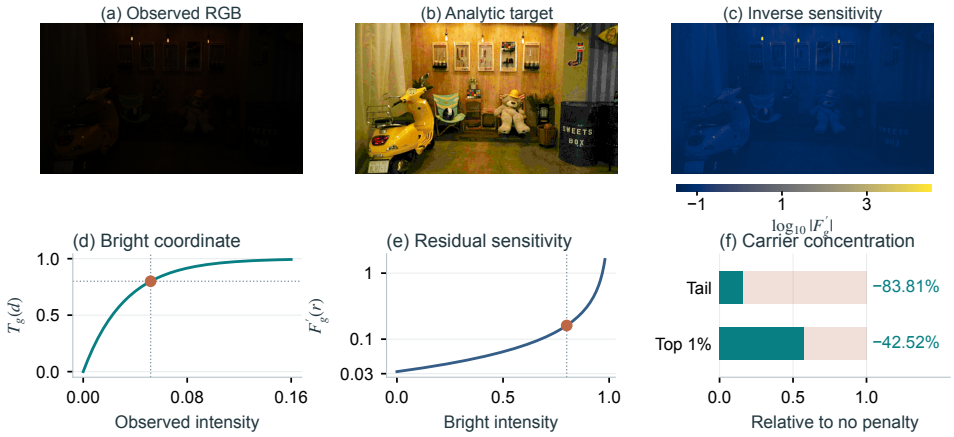


Figure 6: Radiometric conditioning and primitive allocation. Top: observed RGB, analytic target and log inverse sensitivity. Bottom: tone map, inverse derivative and carrier reductions from the footprint penalty.

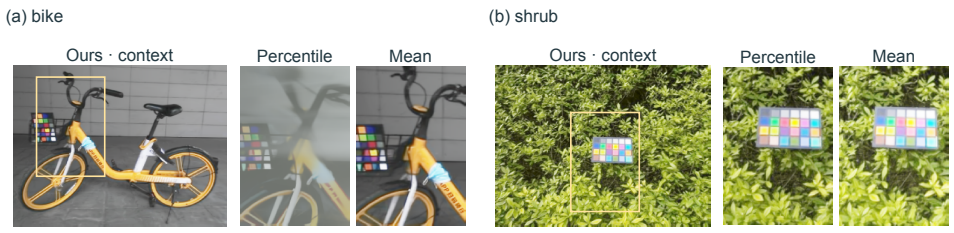


Figure 7: LOM anchor comparisons. Contexts locate identical percentile/mean crops: bike's frame and wall structure, and shrub's foliage and color chart.

4.5 Adapting the exposure anchor

Table 3: LOM fixed-anchor and loss comparisons.

Variant	PSNR ↑	SSIM ↑	LPIPS ↓
Percentile inverse	21.19	0.8059	0.2368
Percentile equal	21.20	0.8127	0.2355
Fixed percentile	21.24	0.8007	0.2825
Fixed mean	23.94	0.8322	0.2446
Inverse only	23.27	0.8164	0.2320
Ours	23.45	0.8350	0.2290

LOM tests exposure adaptation across heterogeneous captures. Figure 7 compares percentile and mean anchors under the same reconstruction objective. Mean anchoring recovers bike's frame and

wall structure; percentile anchoring preserves shrub's foliage and chart contrast. Table 3 compares the fixed-anchor and inverse-fitting controls with Ours; Percentile inverse/equal use the percentile map with inverse-only/equal-weight fitting. CoLight-GS attains 23.45 dB PSNR, 0.8350 SSIM and 0.2290 LPIPS.

Figure 8 compares the two anchors on all five scenes. With each scene's selected map fixed (Table 3, Inverse only vs. Ours), canonical supervision improves PSNR/SSIM/LPIPS from 23.27/0.8164/0.2320 to 23.45/0.8350/0.2290. SSIM improves in all five scenes.

The anchor comparison shows which part of the intensity distribution sets the desired appearance. Percentile anchoring holds an upper luminance level fixed, while mean anchoring accounts for the aggregate mapped RGB values. Their visible effects follow this distinction in bike and shrub. Once that coordinate is assigned, the inverse links its rendered appearance to the same dark inputs. The five-scene SSIM improvement with fixed maps shows the contribution of coupled fitting across these different exposure assignments.

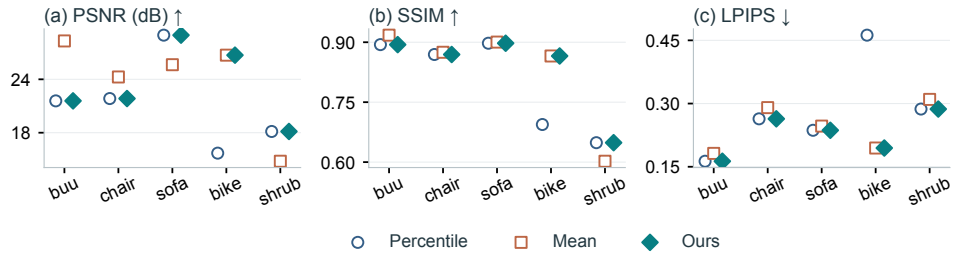


Figure 8: Fixed percentile, fixed mean and CoLight-GS anchoring on every LOM scene under the same reconstruction objective.

5 Discussion

Exposure correction defines both a target appearance and a reconstruction objective. Sharing the map improves the exposure convention requested by different cameras. With that map fixed, coupled fitting improves structure and perceptual similarity. The matched renderings and scene-resolved component effects connect these two roles to visible boundaries and surface detail.

Derivative weighting coordinates image residuals, while opacity–footprint control raises the cost of large opaque carriers. The canonical-SSIM removal and carrier measurements show these complementary roles. With gain sharing, they tie the requested appearance to the structure allocated by the Gaussian representation.

This separation provides a useful way to organize exposure-aware reconstruction. The map specifies the target appearance, the two image losses determine how local content is fitted, and the footprint term controls the primitives used to represent it. Each acts on a different object in the optimization and has a corresponding measurement: exposure-sharing comparisons, neighborhood and perceptual metrics, and carrier statistics.

A fixed analytic map supplies a consistent display-RGB coordinate across a common capture collection. The source-histogram rule adapts this coordinate across the heterogeneous LOM scenes. In both cases, the reconstruction stores the recovered appearance in the Gaussian scene and renders it directly from new cameras.

6 Conclusion

CoLight-GS reconstructs a bright Gaussian scene by coupling its rendered appearance to dark observations through an invertible map. Gain sharing sets a coherent exposure convention across views, while derivative-based weighting and opacity–footprint control guide its structural fit. Evaluations on RealX3D and LOM show that CoLight-GS outperforms five compared methods across all three metrics, while source-histogram anchoring adapts the exposure coordinate to heterogeneous captures. The recovered scene can be directly rendered with standard Gaussian splatting.

AI USAGE DISCLOSURE

AI assistance was used for literature search, evidence organization, manuscript drafting, derivation checking, figure scripting and document layout. The reported model runs and rendered scene images come from the archived research experiments; no experimental image was generated by an image synthesis model during manuscript preparation.

REFERENCES

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. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pp. 5835–5844. IEEE, 2021. doi: 10.1109/ICCV48922.2021.00580. URL https://doi.org/10.1109/ICCV48922.2021.00580.

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-nerf 360: Unbounded anti-aliased neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 5460–5469. IEEE, 2022. doi: 10.1109/CVPR52688.2022.00539. URL https://doi.org/10.1109/CVPR52688.2022.00539.

Qida Cao, Jiajun Ding, Zhenyang Liu, Zhenzhong Kuang, Yijie Shao, and Yilan Shen. VC-GS: view-consistent deblurring gaussian splatting via alternating branch optimization. *Vis. Comput.*, 41(10):7303–7317, 2025. doi: 10.1007/S00371-025-03806-Y. URL https://doi.org/10.1007/s00371-025-03806-y.

Qida Cao, Jiajun Ding, Qingyuan Tang, Tianning Zhao, Xiaoling Gu, Jianping Fan, and Zhou Yu. GC-GS: gradient control gaussian splatting with various image degradation. *Pattern Recognit.*, 171:112304, 2026. doi: 10.1016/J.PATCOG.2025.112304. URL https://doi.org/10.1016/j.patcog.2025.112304.

Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 3291–3300. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00347. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Chen_Learning_to_See_CVPR_2018_paper.html.

Ziteng Cui, Lin Gu, Xiao Sun, Xianzheng Ma, Yu Qiao, and Tatsuya Harada. Aleth-nerf: Illumination adaptive nerf with concealing field assumption. In Michael J. Wooldridge, Jennifer G. Dy, and Sriraam Natarajan (eds.), *Thirty-Eighth AAAI Conference on Artificial Intelligence, AAAI 2024, Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence, IAAI 2024, Fourteenth Symposium on Educational Advances in Artificial Intelligence, EAAI 2024, February 20-27, 2024, Vancouver, Canada*, pp. 1435–1444. AAAI Press, 2024. doi: 10.1609/AAAI.V38I2.27908. URL https://doi.org/10.1609/aaai.v38i2.27908.

Ziteng Cui, Xuangeng Chu, and Tatsuya Harada. Luminance-gs: Adapting 3d gaussian splatting to challenging lighting conditions with view-adaptive curve adjustment. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pp. 26472–26482. Computer Vision Foundation / IEEE, 2025. doi: 10.1109/CVPR52734.2025.02465. URL https://openaccess.thecvf.com/content/CVPR2025/html/Cui_Luminance-GS_Adapting_3D_Gaussian_Splatting_to_Challenging_Lighting_Conditions_with_CVPR_2025_paper.html.

Paul E. Debevec and Jitendra Malik. Recovering high dynamic range radiance maps from photographs. In G. Scott Owen, Turner Whitted, and Barbara Mones-Hattal (eds.), *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH 1997, Los Angeles, CA, USA, August 3-8, 1997*, pp. 369–378. ACM, 1997. doi: 10.1145/258734.258884. URL https://doi.org/10.1145/258734.258884.

Michael D. Grossberg and Shree K. Nayar. Modeling the space of camera response functions. *IEEE Trans. Pattern Anal. Mach. Intell.*, 26(10):1272–1282, 2004. doi: 10.1109/TPAMI.2004.88. URL https://doi.org/10.1109/TPAMI.2004.88.

Xiaojie Guo, Yu Li, and Haibin Ling. LIME: low-light image enhancement via illumination map estimation. *IEEE Trans. Image Process.*, 26(2):982–993, 2017. doi: 10.1109/TIP.2016.2639450. URL https://doi.org/10.1109/TIP.2016.2639450.

Florian Hahlbohm, Linus Franke, Martin Eisemann, and Marcus A. Magnor. Faster-gs: Analyzing and improving gaussian splatting optimization. *CoRR*, abs/2602.09999, 2026. doi: 10.48550/ARXIV.2602.09999. URL https://doi.org/10.48550/arXiv.2602.09999.

Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(12):2341–2353, 2011. doi: 10.1109/TPAMI.2010.168. URL https://doi.org/10.1109/TPAMI.2010.168.

Rasmus Ramsbøl Jensen, Anders Lindbjerg Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *2014 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2014, Columbus, OH, USA, June 23-28, 2014*, pp. 406–413. IEEE Computer Society, 2014. doi: 10.1109/CVPR.2014.59. URL https://doi.org/10.1109/CVPR.2014.59.

Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlightengan: Deep light enhancement without paired supervision. *IEEE Trans. Image Process.*, 30:2340–2349, 2021. doi: 10.1109/TIP.2021.3051462. URL https://doi.org/10.1109/TIP.2021.3051462.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139:1–139:14, 2023. doi: 10.1145/3592433. URL https://doi.org/10.1145/3592433.

Shakiba Kheradmand, Daniel Rebain, Gopal Sharma, Weiwei Sun, Yang-Che Tseng, Hossam Isack, Abhishek Kar, Andrea Tagliasacchi, and Kwang Moo Yi. 3d gaussian splatting as markov chain monte carlo. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/93be245fce00a9bb2333c17ceae4b732-Abstract-Conference.html.

Xin Lin, Shi Luo, Xiaojun Shan, Xiaoyu Zhou, Chao Ren, Lu Qi, Ming-Hsuan Yang, and Nuno Vasconcelos. HQGS: high-quality novel view synthesis with gaussian splatting in degraded scenes. In *The Thirteenth International Conference on Learning Representations, ICLR 2025, Singapore, April 24-28, 2025*. OpenReview.net, 2025. URL https://openreview.net/forum?id=25Zlvl7JxW.

Shuhong Liu, Lin Gu, Ziteng Cui, Xuangeng Chu, and Tatsuya Harada. I2-nerf: Learning neural radiance fields under physically-grounded media interactions. In Danielle Belgrave, Cheng Zhang, Laura N. Montoya, Hsuan-Tien Lin, Razvan Pascanu, Piotr Koniusz, Marzyeh Ghassemi, Nancy Chen, Iván Vladimir Meza Ruíz, and Arturo Loaiza-Bonilla (eds.), *Advances in Neural Information Processing Systems 38: Annual Conference on Neural Information Processing Systems 2025, NeurIPS 2025, San Diego, CA, USA, December 2-7, 2025 / Mexico City, Mexico, November 30 - December 5, 2025*, 2025. URL http://papers.nips.cc/paper_files/paper/2025/hash/77fa0e7d45c6687f1958de0b31e9fc05-Abstract-Conference.html.

Shuhong Liu, Chenyu Bao, Ziteng Cui, Yun Liu, Xuangeng Chu, Lin Gu, Marcos V. Conde, Ryo Umagami, Tomohiro Hashimoto, Zijian Hu, Tianhan Xu, Yuan Gan, Yusuke Kurose, and Tatsuya Harada. Realx3d: A physically-degraded 3d benchmark for multi-view visual restoration and reconstruction. *Int. J. Comput. Vis.*, 134(8):382, 2026. doi: 10.1007/S11263-026-02963-4. URL https://doi.org/10.1007/s11263-026-02963-4.

Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. Llnet: A deep autoencoder approach to natural low-light image enhancement. *Pattern Recognit.*, 61:650–662, 2017. doi: 10.1016/J.PATCOG.2016.06.008. URL https://doi.org/10.1016/j.patcog.2016.06.008.

Jonathon Luiten, Georgios Kopanas, Bastian Leibe, and Deva Ramanan. Dynamic 3d gaussians: Tracking by persistent dynamic view synthesis. In *International Conference on 3D Vision, 3DV 2024, Davos, Switzerland, March 18-21, 2024*, pp. 800–809. IEEE, 2024. doi: 10.1109/3DV62453.2024.00044. URL https://doi.org/10.1109/3DV62453.2024.00044.

Li Ma, Xiaoyu Li, Jing Liao, Qi Zhang, Xuan Wang, Jue Wang, and Pedro V. Sander. Deblur-nerf: Neural radiance fields from blurry images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 12851–12860. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01252. URL https://doi.org/10.1109/CVPR52688.2022.01252.

Ben Mildenhall, Pratul P. Srinivasan, Rodrigo Ortiz Cayon, Nima Khademi Kalantari, Ravi Ramamoorthi, Ren Ng, and Abhishek Kar. Local light field fusion: practical view synthesis with prescriptive sampling guidelines. *ACM Trans. Graph.*, 38(4):29:1–29:14, 2019. doi: 10.1145/3306346.3322980. URL https://doi.org/10.1145/3306346.3322980.

Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (eds.), *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part I*, volume 12346 of *Lecture Notes in Computer Science*, pp. 405–421. Springer, 2020. doi: 10.1007/978-3-030-58452-8_24. URL https://doi.org/10.1007/978-3-030-58452-8_24.

Ben Mildenhall, Peter Hedman, Ricardo Martin-Brualla, Pratul P. Srinivasan, and Jonathan T. Barron. Nerf in the dark: High dynamic range view synthesis from noisy raw images. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2022, New Orleans, LA, USA, June 18-24, 2022*, pp. 16169–16178. IEEE, 2022. doi: 10.1109/CVPR52688.2022.01571. URL https://doi.org/10.1109/CVPR52688.2022.01571.

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph.*, 41(4):102:1–102:15, 2022. doi: 10.1145/3528223.3530127. URL https://doi.org/10.1145/3528223.3530127.

Srinivasa G. Narasimhan and Shree K. Nayar. Vision and the atmosphere. *Int. J. Comput. Vis.*, 48 (3):233–254, 2002. doi: 10.1023/A:1016328200723. URL https://doi.org/10.1023/A:1016328200723.

Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-nerf: Neural radiance fields for dynamic scenes. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2021, virtual, June 19-25, 2021*, pp. 10318–10327. Computer Vision Foundation / IEEE, 2021. doi: 10.1109/CVPR46437.2021.01018. URL https://openaccess.thecvf.com/content/CVPR2021/html/Pumarola_D-NeRF_Neural_Radiance_Fields_for_Dynamic_Scenes_CVPR_2021_paper.html.

Zefan Qu, Ke Xu, Gerhard P. Hancke, and Rynson W. H. Lau. Lush-nerf: Lighting up and sharpening nerfs for low-light scenes. In Amir Globersons, Lester Mackey, Danielle Belgrave, Angela Fan, Ulrich Paquet, Jakub M. Tomczak, and Cheng Zhang (eds.), *Advances in Neural Information Processing Systems 37: Annual Conference on Neural Information Processing Systems 2024, NeurIPS 2024, Vancouver, BC, Canada, December 10 - 15, 2024*, 2024. URL http://papers.nips.cc/paper_files/paper/2024/hash/c69465280855cfe25d566e359da140c1-Abstract-Conference.html.

Hao Sun, Fenggen Yu, Huiyao Xu, Tao Zhang, and Changqing Zou. Ll-gaussian: Low-light scene reconstruction and enhancement via gaussian splatting for novel view synthesis. In Cathal Gurrin, Klaus Schoeffmann, Min Zhang, Luca Rossetto, Stevan Rudinac, Duc-Tien Dang-Nguyen, Wen-Huang Cheng, Phoebe Chen, and Jenny Benois-Pineau (eds.), *Proceedings of the 33rd ACM International Conference on Multimedia, MM 2025, Dublin, Ireland, October 27-31, 2025*, pp. 4261–4270. ACM, 2025. doi: 10.1145/3746027.3755375. URL https://doi.org/10.1145/3746027.3755375.

Haoyuan Wang, Xiaogang Xu, Ke Xu, and Rynson W. H. Lau. Lighting up nerf via unsupervised decomposition and enhancement. In *IEEE/CVF International Conference on Computer Vision,*

ICCV 2023, Paris, France, October 1-6, 2023, pp. 12598–12607. IEEE, 2023a. doi: 10.1109/ICCV51070.2023.01161. URL https://doi.org/10.1109/ICCV51070.2023.01161.

Min Wang, Xin Huang, Guoqing Zhou, Qifeng Guo, and Qing Wang. Bright-nerf: Brightening neural radiance field with color restoration from low-light RAW images. In Toby Walsh, Julie Shah, and Zico Kolter (eds.), *Thirty-Ninth AAAI Conference on Artificial Intelligence, Thirty-Seventh Conference on Innovative Applications of Artificial Intelligence, Fifteenth Symposium on Educational Advances in Artificial Intelligence, AAAI 2025, Philadelphia, PA, USA, February 25 - March 4, 2025*, pp. 7817–7825. AAAI Press, 2025. doi: 10.1609/AAAI.V39I8.32842. URL https://doi.org/10.1609/aaai.v39i8.32842.

Peng Wang, Lingzhe Zhao, Ruijie Ma, and Peidong Liu. Bad-nerf: Bundle adjusted deblur neural radiance fields. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2023, Vancouver, BC, Canada, June 17-24, 2023*, pp. 4170–4179. IEEE, 2023b. doi: 10.1109/CVPR52729.2023.00406. URL https://doi.org/10.1109/CVPR52729.2023.00406.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Trans. Image Process.*, 13(4):600–612, 2004. doi: 10.1109/TIP.2003.819861. URL https://doi.org/10.1109/TIP.2003.819861.

Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *British Machine Vision Conference 2018, BMVC 2018, Newcastle, UK, September 3-6, 2018*, pp. 155. BMVA Press, 2018. URL http://bmvc2018.org/contents/papers/0451.pdf.

Yue Wen, Liang Song, and Hesheng Wang. Splatbright: Generalizable low-light scene reconstruction from sparse views via physically-guided gaussian enhancement. *CoRR*, abs/2512.18655, 2025. doi: 10.48550/ARXIV.2512.18655. URL https://doi.org/10.48550/arXiv.2512.18655.

Haofei Xu, Daniel Barath, Andreas Geiger, and Marc Pollefeys. Resplat: Learning recurrent gaussian splatting. In Paolo Favaro, Zuzana Kukelova, Atsuto Maki, Anna Rohrbach, Konrad Schindler, and Federico Tombari (eds.), *Computer Vision - ECCV 2026 - 19th European Conference, Malmö, Sweden, September 8-12, 2026, Proceedings, Part XXXIX*, volume 17039 of *Lecture Notes in Computer Science*, pp. 364–382. Springer, 2026. doi: 10.1007/978-3-032-37435-6_20. URL https://doi.org/10.1007/978-3-032-37435-6_20.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for fast image restoration and enhancement. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(2):1934–1948, 2023. doi: 10.1109/TPAMI.2022.3167175. URL https://doi.org/10.1109/TPAMI.2022.3167175.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18-22, 2018*, pp. 586–595. Computer Vision Foundation / IEEE Computer Society, 2018. doi: 10.1109/CVPR.2018.00068. URL http://openaccess.thecvf.com/content_cvpr_2018/html/Zhang_The_Unreasonable_Effectiveness_CVPR_2018_paper.html.

Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In Laurent Amsaleg, Benoit Huet, Martha A. Larson, Guillaume Gravier, Hayley Hung, Chong-Wah Ngo, and Wei Tsang Ooi (eds.), *Proceedings of the 27th ACM International Conference on Multimedia, MM 2019, Nice, France, October 21-25, 2019*, pp. 1632–1640. ACM, 2019. doi: 10.1145/3343031.3350926. URL https://doi.org/10.1145/3343031.3350926.

Han Zhou, Wei Dong, and Jun Chen. LITA-GS: illumination-agnostic novel view synthesis via reference-free 3d gaussian splatting and physical priors. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2025, Nashville, TN, USA, June 11-15, 2025*, pp. 21580–21589. Computer Vision Foundation / IEEE, 2025. doi: 10.1109/CVPR52734.2025.02010. URL https://openaccess.thecvf.com/content/CVPR2025/html/Zhou_LITA-GS_Illumination-Agnostic_Novel_View_Synthesis_via_Reference-Free_3D_Gaussian_Splatting_CVPR_2025_paper.html.

A Implementation details

The radiometric operators are applied during camera-image preparation, and their inverse is applied to the differentiable rendered RGB during training. Novel-view evaluation uses the canonical rendering directly. Both branches use the same $0.8/0.2$ pixel/structure mixture. The no-canonical-SSIM variant removes the structural contribution from the bright branch while keeping its pixel term and the observed-domain objective.

Opacity and scale regularization coefficients are 0.01 each. Densification starts at step 600 and is scheduled every 100 steps up to 15,000. The position learning rate decays from 1.6×10^{-4} to 1.6×10^{-6}; the scale learning rate is 0.005. All runs initialize Python, NumPy and PyTorch with seed zero. Variants retain the same camera inputs, initialization and optimizer schedule. LOM uses official geometry with all initial RGB channels set to 127 in the stored byte representation. The scene extent in Equation 7 normalizes scale across reconstructions.

Table 4: Dark training and normal-light test views.

RealX3D scene	Train	Test	LOM scene	Train	Test
BlueHawaii	30	6	buu	22	3
Chocolate	31	6	chair	43	5
Cupcake	30	6	sofa	29	4
GearWorks	27	5	bike	36	4
Laboratory	30	6	shrub	30	5
MilkCookie	27	6			
Popcorn	30	6			
Sculpture	32	6			
Ujikintoki	30	5			
Total	267	52	Total	160	21

B Conditioning of the coupled objective

The radiometric map fixes the relationship between the two image targets. Its derivatives explain why the corresponding residuals have different scales:

$$T'_g(d) = g\exp(-gd), \qquad F'_g(r) = \frac{1}{g(1-r)}, \qquad F''_g(r) = \frac{1}{g(1-r)^2}. \tag{9}$$

At a canonical target $r_0 = T_g(d)$, expanding the observed prediction gives

$$F_g(r_0+\epsilon) - F_g(r_0) = \frac{\epsilon}{g(1-r_0)} + \frac{\epsilon^2}{2g(1-r_0)^2} + O(\epsilon^3). \tag{10}$$

For an L_1 discrepancy away from zero, the leading residual scale is $F'_g(r_0)$. Evaluating it at the exposure anchor gives J_a in Equation 6. The condition $w_c = (1-w_c)J_a$ then yields the temperature-one weight. With the RealX3D gain $g = 31.03$ and anchor $a = 0.8$, $J_a \approx 0.1611$. A temperature below one increases the canonical weight toward an equal mixture.

The implemented image discrepancy also contains SSIM, whose patch statistics couple neighboring pixels. The derivative calculation sets the local scalar scale used to mix the branches.

The map and objective can be varied independently. Changing the gain changes the requested canonical appearance. Holding the map fixed and changing the objective changes how the finite Gaussian representation fits it.