HIGH DYNAMIC RANGE 3D GAUSSIAN SPLATTING VIA LUMINANCE–CHROMATICITY DECOMPOSITION

Anonymous authors
Paper under double-blind review

ABSTRACT

High Dynamic Range (HDR) 3D reconstruction is pivotal for professional content creation in filmmaking and virtual production. Existing methods typically rely on multi-exposure Low Dynamic Range (LDR) supervision to constrain the learning process within vast brightness spaces, resulting in complex, dual-branch architectures. This work explores the feasibility of learning HDR 3D models exclusively in the HDR RGB data space to simplify model design. By analyzing 3D Gaussian Splatting (3DGS) under direct HDR supervision, we find that using a shared set of Spherical Harmonics (SHs) to jointly encode extreme radiance variations leads to difficult optimization and substantial underfitting. While increasing the maximum SH degree improves training fitting, it leads to severe overfitting and excessive parameter overhead. To address this, we propose *Luminance–Chromaticity Decomposition Gaussian Splatting* (LCD-GS). By decoupling luminance and chromaticity into independent parameters, LCD-GS significantly enhances HDR reconstruction with minimal parameter increase (*e.g.*, one extra scalar per primitive). Notably, LCD-GS maintains the original training and inference pipeline, requiring only a change in color representation. This explicit decomposition naturally enables primitive-level luminance editing during inference. Extensive experiments on synthetic and real datasets demonstrate that LCD-GS consistently outperforms state-of-the-art methods in reconstruction fidelity and dynamic-range preservation even with a simpler, more efficient architecture, providing an elegant paradigm for professional-grade HDR 3D modeling. Code and datasets will be released.

1 INTRODUCTION

3D scene modeling such as Novel View Synthesis (NVS) (Reddy et al., 2025; Cai & Lu, 2024) has achieved remarkable progress, enabling photorealistic rendering for gaming (Pan et al., 2024b), AR/VR (Van Bogaert et al., 2022), and autonomous driving (Ma et al., 2025). Existing NVS models (Kerbl et al., 2023; Dalal et al., 2024; Niu et al., 2024) are typically confined to LDR imagery, suffering from severe under- and over-exposure artifacts (see Fig. 2a-b) that prevent the recovery of accurate physical lighting. To alleviate this and reconstruct full dynamic range, High Dynamic Range Novel View Synthesis (HDR NVS) methods (Huang et al., 2022; Cai et al., 2024; Liu et al., 2025; Ye et al., 2026) have been built often leveraging multi-exposure LDR supervision. This paradigm is designed to constrain the learning process within vast brightness spaces, but consequently necessitates complex, dual-branch architectures with specialized inverse tone-mapping or fusion pipelines. Such designs significantly increase capture complexity, require aligned exposure stacks, and are highly sensitive to motion misalignment (Kalantari & Ramamoorthi, 2017).

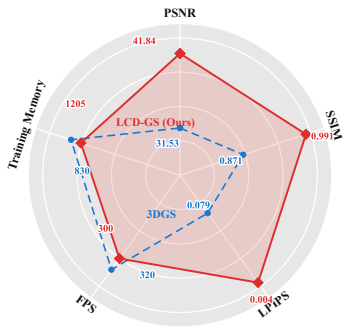
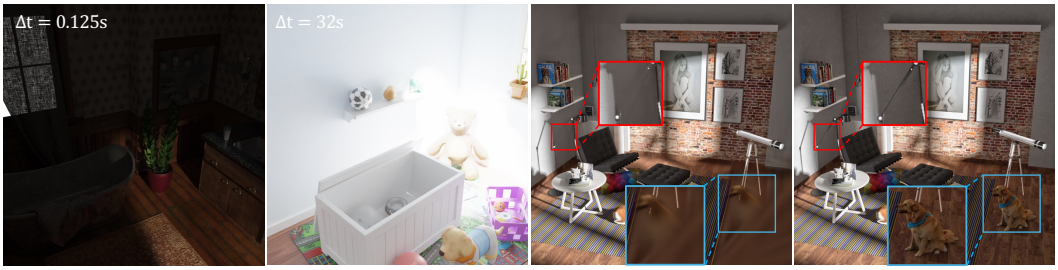


Figure 1: LCD-GS unlocks 3DGS in HDR 3D modeling while keeping the efficiency.

To overcome these structural limitations and simplify model design, here we propose to explore *the feasibility of learning HDR 3D models exclusively in the HDR data space*. The growing integration of HDR-oriented imaging pipelines in modern devices (*e.g.*, iPhones, Sony IMX490) makes acquiring such richer, direct radiometric observations increasingly accessible (Wang et al., 2016), bypassing the need for tedious multi-exposure setups. However, we observe that directly adapting 3DGS, a

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.



(a) Underexposure (b) Overexposure (c) NVS by 3DGS (d) Ground Truth

Figure 2: Example of HDR 3D modeling: (a) underexposure and (b) overexposure. (c) 3DGS (Kerbl et al., 2023) exhibits severe blurring artifacts. (d) ground truth. Δt: Exposure time.

state-of-the-art 3D model[1], to HDR supervision results in substantial underfitting: the resulting radiance fields become heavily biased towards bright regions and fail to learn darker, low-exposure areas (see Fig. 2c).

We conjecture this failure fundamentally stems from *the limited flexibility of 3DGS's color representation*. In conventional LDR imaging pipelines, 8-bit quantization, bounded display range, and clipping compress or discard extreme radiance values. Consequently, strongly view-dependent scene appearance cannot be faithfully recorded in the training images, and vanilla 3DGS is typically optimized from observations whose cross-view color variations are relatively moderate. Native HDR observations (*e.g.*, 12/32-bit per channel) preserve a substantially wider range of RGB radiance values than conventional LDR images. In vanilla 3DGS, a single set of low-order RGB SH coefficients must jointly determine both the luminance-like intensity of each primitive and its directional chromaticity. Under direct HDR supervision, this coupling assigns wide-range intensity and bounded angular color ratios to the same parameter block, making the representation difficult to optimize and leading to substantial underfitting in our experiments (see Fig. 8 and Tab. 2). Importantly, this behavior persists when 3DGS is trained with the same μ-law objective and positive output activation used by LCD-GS (see Appendix A.4), suggesting that it cannot be explained solely by the loss or output range. Increasing the maximum SH degree from 3 to 4/5 improves angular capacity but yields only limited reconstruction gains while increasing computational cost and introducing additional artifacts (see Fig. 3 and Tab. 2). These observations motivate an explicit luminance–chromaticity parameterization.

To address this inherent limitation without inducing prohibitive costs, we propose *Luminance-Chromaticity Decomposition Gaussian Splatting* (LCD-GS), inspired by the above insights, as well as human visual system's distinct functional pathways for processing luminance and chromaticity (Xing et al., 2015), LCD-GS explicitly *decouples luminance and chromaticity* in parameterization. We factorize the radiance of each 3D Gaussian primitive into a *view-independent positive luminance scalar*, which represents absolute radiance magnitude, and a *view-dependent chromaticity term*, which represents an objective specification of the color (regardless of its luminance). This luminance–chromaticity parameterization reduces the simultaneous roles assigned to the SHs: the luminance scalar accommodates wide per-primitive HDR intensity, while the normalized SH term models bounded directional color ratios. The chromaticity term can be represented with SH same as 3DGS's color. Beyond the resulting optimization benefits, this decoupling strategy further enables controllable luminance editing of scenes (*e.g.*, highlight suppression and shadow enhancement in Appendix A.5). Notably, LCD-GS preserves the original training and inference pipeline of 3DGS, requiring only a change in color representation and a minimal parameter increase (*e.g.*, one scalar for luminance per primitive).

Our **contributions** are: **(I)** We present the first systematic study of learning HDR 3D models *exclusively* in the HDR RGB data space, opening up a simplified, more practical paradigm for professional digital content creation. **(II)** We reveal the fundamental limitations of existing 3D models in processing HDR observations along with insightful analysis and discussions. **(III)** We propose LCD-GS, a novel 3DGS variant characterized by luminance–chromaticity decomposition for more scalable radiance representation and learning from HDR imagery. **(IV)** Extensive evaluations demonstrate that, even with simplified architecture and no LDR supervision, LCD-GS outperforms

[1]NeRF (Mildenhall et al., 2020) encounters similar representational limitations due to the entanglement of geometric density and radiance within a single neural network (see Tab. 2).

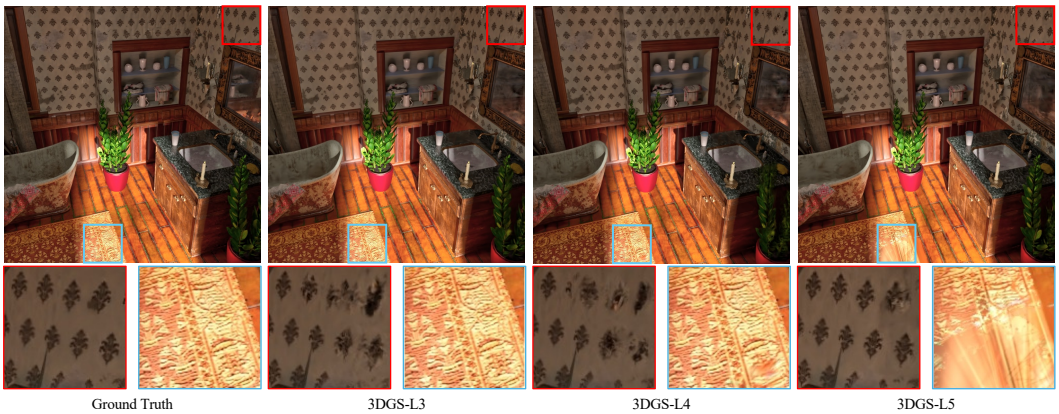


Figure 3: Elevating the maximum SH degree (D) of 3DGS mitigates blurring artifacts while simultaneously causing additional artifacts. Dx means the maximum SH degree is x, 3 by default.

all alternatives on both synthetic and real-world benchmarks, validating its radiometric accuracy and efficiency (see Fig. 1).

2 Related Work

Novel view synthesis (NVS) is a fundamental task in 3D vision, with applications in AR/VR (Van Bogaert et al., 2022), gaming (Pan et al., 2024b), and autonomous driving (Ma et al., 2025). Classical methods like Structure-from-Motion (SfM) (Pan et al., 2024a) and Multi-View Stereo (MVS) (Huang et al., 2025) rely on geometric cues from multi-view images but struggle with occlusions, textureless regions, and high computational costs (Feng et al., 2025). Recent learning-based approaches model scenes as continuous differentiable representations. NeRF (Mildenhall et al., 2020) pioneered this paradigm by using neural networks to map 3D coordinates and viewing directions to volume density and color via differentiable volume rendering. Subsequent works (Gao et al., 2024; Zhu et al., 2024; Yao et al., 2025; Müller et al., 2022; Yu et al., 2021; Fu et al., 2024) aim to improve efficiency and quality. Alternatively, 3DGS (Kerbl et al., 2023) and its variations (Ye et al., 2024; Yu et al., 2024; Cheng et al., 2024; Kulhanek et al., 2024) represent scenes as learnable 3D Gaussian primitives, bypassing volumetric rendering and enabling real-time performance. Despite progress, such existing state-of-the-art NVS methods are predominantly focused on LDR sRGB inputs, limiting their ability to reconstruct or render true HDR radiance from real-world scenes.

HDR novel view synthesis (HDR NVS) aims to reconstruct scenes with high dynamic range from multi-view observations. Huang *et al.* (Huang et al., 2022) introduced HDR-NeRF, the first HDR-NVS framework, which extends standard NeRF (Mildenhall et al., 2020) to learn mappings from physical radiance to HDR color using multi-exposure LDR inputs. However, its reliance on the NeRF architecture results in prohibitively slow inference. HDR-GS (Cai et al., 2024) addressed this limitation using 3DGS (Kerbl et al., 2023), achieving significantly faster rendering and improved visual quality. GaussHDR (Liu et al., 2025) further introduced a unified tone-mapping strategy that enables HDR novel view synthesis without requiring HDR supervision. Despite these advances, all aforementioned methods fundamentally rely on multi-exposure LDR image stacks, which entail complex capture protocols and substantial storage overhead.

The increasing availability of HDR-oriented imaging pipelines and HDR-capable sensors (*e.g.*, Sony IMX490) motivates our study (Wang et al., 2016). Compared with standard LDR imagery, such data provide richer radiometric information, making them attractive for HDR-aware 3D reconstruction. While prior HDR-aware methods such as RawNeRF (Mildenhall et al., 2022), HDRSplat (Singh et al., 2024), and LE3D (Jin et al., 2024) reconstruct HDR scenes from linear RAW observations, typically under noisy or low-light conditions, we instead target direct HDR reconstruction from RGB-domain HDR images. This setting removes the dependence on RAW sensor data and introduces a distinct representation challenge: how to faithfully model extreme radiance variation directly in HDR RGB space. To the best of our knowledge, we present the first study to directly operate on HDR RGB imagery without relying on LDR-based supervision. Furthermore, by introducing a principled

luminance–chromaticity decomposition, LCD-GS directly optimizes wide-range scene-referred HDR values and normalized directional chromaticity, enabling highly faithful HDR scene reconstruction and luminance editing. More discussion of illumination editing is provided in the Appendix A.3.

3 METHOD

Problem setting. We aim to learn an HDR 3D model $\mathcal{F}$ for a target scene from posed HDR images, $\mathcal{F} : (V) \rightarrow \mathbf{I}_V$, that can render an HDR image $\mathbf{I}_V$ for any viewpoint V. To that end, we capture a set of HDR training images $\mathbf{I} = \{\mathbf{I}_1, \cdots, \mathbf{I}_q, \cdots, \mathbf{I}_Q\}$, with $\mathbf{I}_q$ the q-th image taken under view V_q.

3.1 PRELIMINARY: COLOR REPRESENTATION IN 3DGS

In vanilla 3DGS (Kerbl et al., 2023), given a pixel located at $\boldsymbol{x}$ and its distances to all overlapping Gaussian primitives along the viewing ray, alpha compositing is employed to compute the final pixel color:

$$\boldsymbol{p}(\boldsymbol{x}) = \sum_{n=1}^{N} \boldsymbol{c}_n \alpha'_n \prod_{j=1}^{n-1} \left(1 - \alpha'_j\right), \tag{1}$$

where N denotes the number of Gaussian primitives intersected by the ray in a front-to-back order. Here, $\boldsymbol{c}_n \in \mathbb{R}^3$ represents the color of the n-th primitive, $\alpha'_n \in [0, 1]$ is its effective opacity at pixel $\boldsymbol{x}$ (obtained from the projected 2D Gaussian and its density), and $\prod_{j=1}^{n-1}(1 - \alpha'_j)$ denotes the accumulated transmittance accounting for occlusion by all preceding Gaussian primitives along the ray. The resulting $\boldsymbol{p}(\boldsymbol{x}) \in \mathbb{R}^3$ is the rendered RGB color of the pixel.

The color of each Gaussian primitive is modeled as a view-dependent function using spherical harmonics (SHs), as shown in Fig. 10 where the shape of each function tells the spatial distributions of these harmonic modes, with different colors representing the variation in the function's values (red / blue indicating positive / negative values). Specifically, each Gaussian primitive is associated with a set of SH coefficients $\mathbf{k} = \{\boldsymbol{k}_d^m | 0 \leq d \leq D, -d \leq m \leq d\} \in \mathbb{R}^{(D+1)^2 \times 3}$ that model view-dependent appearance, and each $\boldsymbol{k}_d^m \in \mathbb{R}^3$ is a set of three coefficients corresponding to the RGB components where the index m denotes the azimuthal order of the SH basis function Y_d^m, controlling its angular frequency and phase around the polar axis for a given degree d. D denotes the maximum SH degree, by default $D = 3$. Formally, the color of the n-th Gaussian primitive under view $\boldsymbol{v} = (\theta, \phi)$ is expressed as

$$\boldsymbol{c}_n(\boldsymbol{v}|\mathbf{k}) = \sum_{d=0}^{D} \sum_{m=-d}^{d} \boldsymbol{k}_d^m Y_d^m(\theta, \phi), \tag{2}$$

where $Y_d^m : \mathbb{S}^2 \rightarrow \mathbb{R}$ is the SH function that maps 3D points on the sphere to real numbers (Kerbl et al., 2023; Cai et al., 2024).

Remarks. The color formulation of 3DGS (*i.e.*, Eq. (2)) implicitly assumes that radiance variations of input images across views are moderate (Wang et al., 2006). Under such cases, the SH expansion provides a compact low-frequency approximation of the bidirectional color variation. However, when applied to HDR imagery, where radiance can vary by *orders of magnitude* across different views, this color representation becomes particularly susceptible, struggles to capture sharp, drastic color variations, leading to underfitting (see Fig. 8 and Tab. 2).

3.2 BASELINE: ELEVATING THE MAXIMUM SH DEGREE IN 3DGS

The above critique raises a natural question: *Can increasing the maximum SH degree alleviate these limitations*? An intuitive approach is to elevate the maximum SH degree in 3DGS (Kerbl et al., 2023), *e.g.*, increasing from the default maximum degree of 3 (16 coefficients) to 5 (36 coefficients, see Fig. 10 in the Appendix). We find that a higher maximum SH degree yields only marginal quantitative gains (see rows 2–4 of Tab. 2) without meaningfully enhancing perceptual rendering quality (see Fig. 3), yet it leads to increased computational overhead. These results show that increasing the maximum SH degree alone is insufficient to resolve the observed HDR reconstruction difficulty, motivating us to separate luminance-like intensity from directional chromaticity.

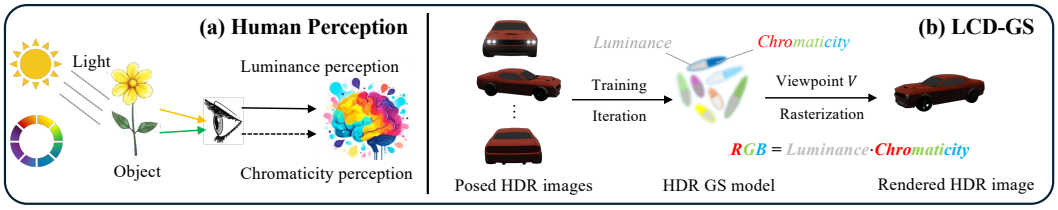


Figure 4: *Overview of LCD-GS*. With posed HDR images of a scene, we aim to learn an HDR Gaussian Splatting model. We draw inspiration from **(a)** human visual perception on tackling drastic radiance variation: luminance and chromaticity are perceived *separately*. In analogy, **(b)** LCD-GS decomposes the parameterization of luminance and chromaticity for color representation, in contrast to the conventional SH color representation of 3DGS which tightly entangles the two aspects using a single set of parameters. Concretely, LCD-GS assigns the luminance and directional chromaticity of each primitive to separate parameters, reducing the simultaneous roles placed on a single set of SHs.

3.3 PROPOSED: LCD-GS

To address the above limitations, we propose LCD-GS, illustrated in Fig. 4, which introduces a luminance–chromaticity decomposition for the color modeling of Gaussian primitives. Instead of directly representing the color of Gaussian primitives using RGB spherical harmonics, we decompose the radiance of each Gaussian primitive into a view-independent positive luminance scalar and a view-dependent bounded chromaticity term. Formally, the color emitted by the n-th Gaussian primitive along the viewing direction $\boldsymbol{v} = (\theta, \phi)$ is defined as:

$$\boldsymbol{c}_n(\boldsymbol{v}|\mathbf{k}, L) = \underbrace{L}_{\text{Luminance}} \underbrace{\cdot \mathcal{N}\left(\left[\sum_{d=0}^{D}\sum_{m=-d}^{d} \boldsymbol{k}_d^m Y_d^m(\theta,\phi)\right]_+ + \epsilon\right)}_{\text{Chromaticity}}, \tag{3}$$

where the luminance term $L \in \mathbb{R}^+$ is a view-independent positive scalar controlling the overall radiance magnitude and ϵ denotes a small positive constant introduced for numerical stability (1×10^{-6} by default). To enforce non-negativity, we parameterize it as $L = g(l)$, where l is a learnable scalar parameter and $g : \mathbb{R} \to \mathbb{R}^+$ is a positive activation function, such as $\text{softplus}(l) = \log(1 + e^l)$ (see Appendix A.4 for more options and ablation study). The chromaticity term is a normalized SH-based component, which captures view-dependent color variation in a bounded color space. Here, $\boldsymbol{k}_d^m$ denotes the learnable SH coefficient of degree d and order m, and $Y_d^m(\theta, \phi)$ is the corresponding SH basis evaluated along the viewing direction $\boldsymbol{v} = (\theta, \phi)$ (see Sec. 3.1). The operator $[\cdot]_+ = \max(\cdot, 0)$ is applied channel-wise to ensure non-negative chromatic responses, while $\mathcal{N}(\cdot)$ denotes L_1 normalization, producing a scale-invariant color representation per primitive.

Discussion. The proposed luminance–chromaticity decomposition offers several advantages. *First*, it explicitly decouples view-independent radiance scale from view-dependent chromatic variation, which are tightly entangled in the original SH color representation of vanilla 3DGS. This makes the formulation particularly suitable for HDR representation: the luminance component can accommodate a wide dynamic range through an unbounded positive scalar parameterization, while the chromaticity component remains in a normalized, stable, and physically plausible color space. *Second*, the formulation can be easily integrated into existing 3DGS methods (Kerbl et al., 2023; Fu et al., 2024; Ye et al., 2024; Yu et al., 2024) by replacing only the color parameterization in Eq. (2) with Eq. (3), while keeping the rest of the pipeline unchanged. This facilitates the reuse of previous open source codes and software. *Third*, our luminance component provides a simple and interpretable handle for luminance editing, while keeping chromaticity, opacity, and geometry unchanged, enabling exposure adjustment and HDR radiance control without collecting extra training data, model retraining, nor introducing specialized networks and computational costs (Liang et al., 2024; Shi et al., 2025) (see Sec. 3.5).

3.4 MODEL OPTIMIZATION

To mitigate optimization instabilities caused by the high dynamic range of radiance values, we apply a μ-law compression function $\mathcal{T}_\mu(\boldsymbol{p}) = \frac{\log(1+\mu\cdot\boldsymbol{p})}{\log(1+\mu)}$ during training, where $\boldsymbol{p}$ denotes a pixel RGB



(a) Reference (b) (c) (d) (e)

Figure 5: Luminance editing enabled by LCD-GS. (a) Reference image. (b)-(e) Shadow enhancement by scaling the luminance of a subset of low-luminance Gaussian primitives.

value in HDR and $\mu \in R^+$ controls the compression strength, following established works (Cai et al., 2024; Zhang et al., 2025). Following this transformation, the overall training objective is defined as:

$$\mathcal{L} = \lambda \cdot \mathcal{L}_1(\hat{\mathbf{I}}_{\text{pred}}, \hat{\mathbf{I}}_{\text{gt}}) + (1-\lambda) \cdot \mathcal{L}_{SSIM}(\hat{\mathbf{I}}_{\text{pred}}, \hat{\mathbf{I}}_{\text{gt}}), \tag{4}$$

where $\hat{\mathbf{I}}_{\text{pred}}$ and $\hat{\mathbf{I}}_{\text{gt}} \in \mathbb{R}^{H\times W\times C}$ represent rendered and HDR ground truth, respectively, both transformed by the μ-law compression, and λ controls the trade-off between the two loss terms.

3.5 Luminance Editing

A unique advantage of our LCD-GS is its inherent capability for luminance editing. As detailed in Sec. 3.3, LCD-GS explicitly decouples the radiance of each 3D Gaussian primitive into a view-independent luminance scalar L and a normalized chromaticity term. Since the chromaticity information is completely isolated from the radiance magnitude, we can effortlessly manipulate the scene's luminance without destroying the underlying colors, scene geometry, or introducing hue shifts. Specifically, luminance editing can be performed by applying a user-defined mapping function $f : \mathbb{R}^+ \to \mathbb{R}^+$ to the learned luminance values of part or all of Gaussian primitives. During rendering, the original luminance L is replaced with $L' = f(L)$, while the chromaticity term remains unchanged. As shown in Fig. 5, this enables global enhancement of shadowed regions by adjusting Gaussian primitives whose luminance values fall below a predefined threshold. Additional details and discussions are provided in Appendix A.5.

4 Experiments

Datasets. To validate the performance of LCD-GS, we conduct extensive experiments on both synthetic and real-world data. We adopt the synthetic HDR dataset from HDR-NeRF (Huang et al., 2022), comprising 8 scenes rendered at 800×800 resolution in Blender (Blender Foundation, 2025), named as ***HDR-Syn-8S***. Each scene contains 35 multi-view HDR images, with every HDR image explicitly paired with 5 corresponding LDR captures taken at distinct shutter speeds. Critically, this multi-exposure LDR acquisition protocol which demands precise exposure bracketing, vibration-free camera rigs, and per-scene recalibration represents a costly and tedious process that severely limits real-world deployment scalability. For real-world evaluation, we introduce ***iHDR-4S***, a new multi-view HDR dataset comprising 4 indoor scenes with geometrically and photometrically complex characteristics, which is captured by a single iPhone 14 Pro mounted on a tripod to ensure accuracy and stability, with the ProCam application[2] in ProRAW format. Unlike existing datasets that rely on synthetic HDR data or simple exposure bracketing, iHDR-4S provides HDR captures with precise camera poses, making it a valuable resource for HDR-aware 3D reconstruction research. Tab. 1 summarizes the dataset statistics and Fig. 9 shows some example HDR images of different scenes. Please refer to Appendix Sec. A.1 for more details.

Table 1: Statistics of iHDR-4S.

Scenes	Training	Testing	Total	Resolution
Box	39	5	44	4032×3024
Computer	48	6	54	
Desk	49	7	56	
Plushtoy	59	8	67	

Evaluation metrics. We use PSNR and SSIM for pixel-level primary, and LPIPS to assess perceptual similarity. For efficiency analysis, we report training time, peak training memory consumption, and inference speed in frames per second (fps). *Results are averaged over all scenes*. Following

[2]https://www.procamapp.com/

Table 2: Results on the HDR-Syn-8S dataset. Computational costs are measured on a single RTX 4090 GPU. The **best** and second-best results are highlighted.

Row	Method	Need LDR Supervision	Training		HDR			Inference
			Time (min)	Memory (MB)	PSNR↑	SSIM↑	LPIPS↓	speed (fps)
1	NeRF (Mildenhall et al., 2020)	×	264	11575	15.20	0.388	0.753	0.42
2	3DGS (Kerbl et al., 2023)	×	**7**	**830**	31.53	0.871	0.079	**320**
3	3DGS-D4	×	18	1110	31.67	0.873	0.072	262
4	3DGS-D5	×	19	1508	32.56	0.880	0.068	226
5	HDR-NeRF (Huang et al., 2022)	✓	543	12373	36.40	0.936	0.018	0.12
6	HDR-GS (Cai et al., 2024)	✓	21	3761	38.29	0.968	0.014	126
7	Mono-HDR-GS (Zhang et al., 2025)	✓	44	6329	38.57	0.970	0.013	137
8	Mono-HDR-NeRF (Zhang et al., 2025)	✓	439	9281	32.86	0.940	0.068	0.26
9	GaussHDR (Liu et al., 2025)	✓	31	5277	39.02	0.976	0.010	23
10	**LCD-GS (Ours)**	×	12	1205	**41.84**	**0.991**	**0.004**	300

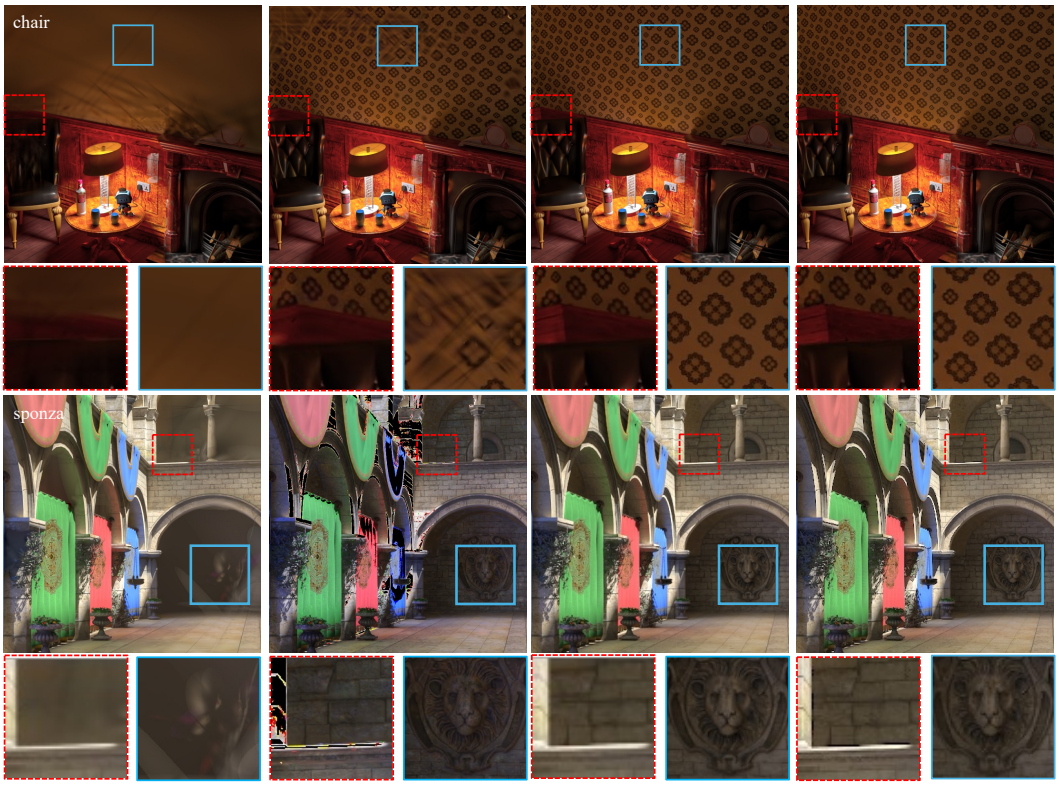


Figure 6: Visual comparisons on the HDR-Syn-8S dataset.

established practice (Huang et al., 2022; Cai et al., 2024; Liu et al., 2025; Zhang et al., 2025), we quantitatively evaluate rendered HDR images in the tone-mapped domain and qualitatively show HDR results tone-mapped by Photomatix Pro (HDRsoft Team, 2025).

4.1 Quantitative Evaluation

Competitors. On the synthetic HDR-Syn-8S dataset, we compare LCD-GS with (1) NeRF (Mildenhall et al., 2020) and (2) 3DGS (Kerbl et al., 2023) variants using SHs with different maximum degrees that serve to examine whether simply extending conventional LDR-oriented NVS methods to HDR inputs is sufficient for HDR scene reconstruction; (3) HDR-NeRF (Huang et al., 2022) and (4) HDR-GS (Cai et al., 2024) that aim to reconstruct HDR scenes from multi-exposure LDR images; (5) Mono-HDR-GS and (6) Mono-HDR-NeRF (Zhang et al., 2025), which aim to reconstruct HDR scenes from single-exposure LDR inputs; and (7) GaussHDR (Liu et al., 2025), a recent HDR NVS method designed to improve HDR reconstruction quality from multi-exposure LDR images.

On the real-world iHDR-4S dataset, we compare LCD-GS with (1) NeRF and (2) 3DGS variants, as well as (3) Mono-HDR-GS and (4) Mono-HDR-NeRF. For Mono-HDR-GS and Mono-HDR-NeRF,

Table 3: Results on the iHDR-4S dataset. Computational costs are measured on a single RTX 4090 GPU. The **best** and second-best results are highlighted.

Row	Method	LDR Supervision	Training		HDR			Inference
			Time (min)	Memory (MB)	PSNR↑	SSIM↑	LPIPS↓	speed (fps)
1	NeRF (Mildenhall et al., 2020)	✗	2147	14160	12.58	0.372	0.793	0.02
2	3DGS (Kerbl et al., 2023)	✗	**14**	**2949**	25.96	0.737	0.296	**258**
3	3DGS-D4	✗	16	2975	26.69	0.741	0.317	248
4	3DGS-D5	✗	18	3062	26.82	0.742	0.317	241
5	Mono-HDR-GS (Zhang et al., 2025)	✓	156	16127	16.63	0.668	0.337	61
6	Mono-HDR-NeRF (Zhang et al., 2025)	✓	570	19165	14.49	0.493	0.842	0.02
7	**LCD-GS (Ours)**	✗	22	3311	**27.74**	**0.749**	**0.295**	197

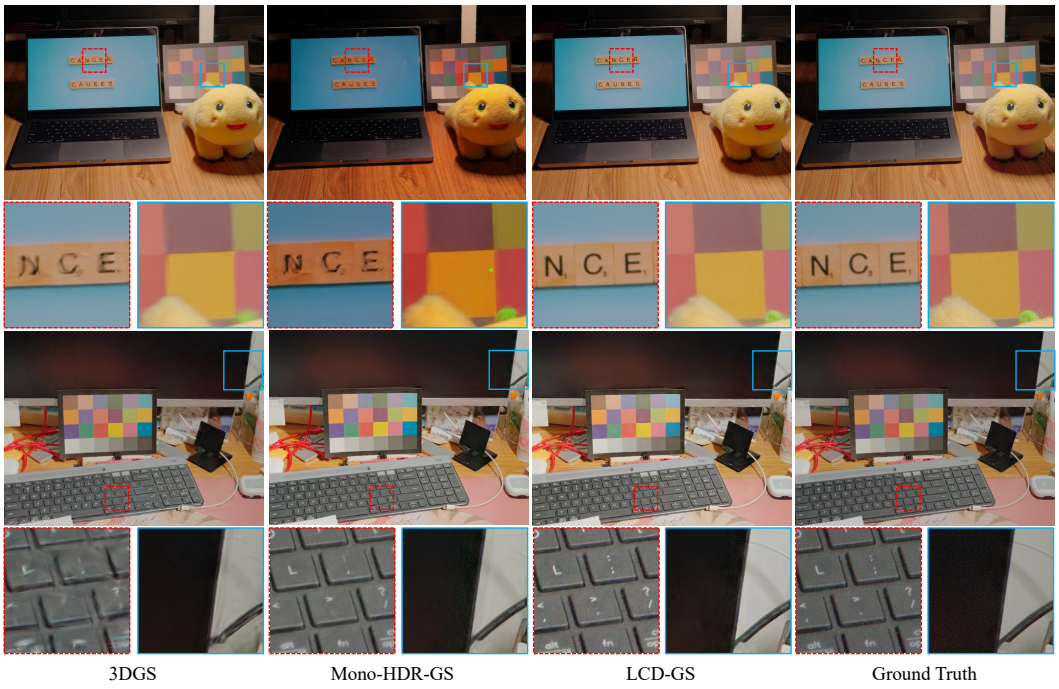


Figure 7: Visual comparisons on the iHDR-4S dataset.

we convert the HDR images in iHDR-4S into single-exposure LDR counterparts using a standard ISP pipeline, enabling their training under the required single-exposure LDR input setting.

Tab. 2 and Tab. 3 present the results on the HDR-Syn-8S and our iHDR-4S datasets, respectively (see Appendix Sec. A.4 for ablations). We highlight the following key points:

(I) Superior HDR reconstruction quality. Our model significantly outperforms all LDR-focused competitors (Kerbl et al., 2023; Mildenhall et al., 2020). Although we replace the bounded sigmoid output with softplus to allow unbounded HDR prediction, NeRF (Mildenhall et al., 2020) still struggles under direct HDR supervision, likely due to the difficulty of optimizing HDR radiance and geometry jointly through an implicit MLP representation. Relative to vanilla 3DGS (Kerbl et al., 2023), LCD-GS achieves a substantial +10.31 dB PSNR gain on the synthetic dataset and +1.78 dB PSNR gain on the real-world dataset, attributable to our luminance-chromaticity decomposition that resolves the fundamental limitation of entangled SH representations in HDR regimes. Increasing the maximum SH degree (*e.g.*, to 5th-degree) yields only marginal quantitative gains (*e.g.*, +1.03 dB PSNR on HDR-Syn-8S) while introducing significant computational overhead and may exhibit pronounced overfitting artifacts, as shown in Fig. 3 and discussed in Sec. 3.2. In contrast, our approach preserves the explicit, optimization-friendly structure of 3DGS while decoupling luminance from chromaticity, enabling seamless integration into any SH-based 3DGS framework without architectural overhaul.

(II) Advantages of native HDR representation. LCD-GS consistently surpasses HDR-NeRF (Huang et al., 2022) and other HDR NVS methods that depend on multi-exposure LDR inputs (Cai et al., 2024; Liu et al., 2025) across all evaluated metrics. This comparison highlights the benefit of using native HDR observations: when such data are available, they provide direct radiometric

cues and avoid the exposure calibration, view alignment, and tone-mapping ambiguities introduced by LDR-to-HDR pipelines (Huang et al., 2022; Cai et al., 2024; Liu et al., 2025). Rather than reconstructing HDR radiance through an intermediate LDR domain, LCD-GS optimizes the 3D representation directly in the HDR observation space, reducing error propagation and eliminating the need for additional tone-mapping or fusion modules. These results indicate that native HDR supervision offers a simpler and more reliable route to HDR scene reconstruction, enabling more accurate radiance estimation and more stable color reconstruction.

(III) Computational efficiency. Relative to vanilla 3DGS, LCD-GS requires slightly more training time and somewhat higher memory usage, while maintaining comparable inference speed across both synthetic and real-world datasets, demonstrating that the proposed luminance–chromaticity decomposition incurs only modest computational overhead. Compared to prior HDR methods, LCD-GS is 2,500× faster than HDR-NeRF (Huang et al., 2022) and approximately 2.4× faster than HDR-GS (Cai et al., 2024). This efficiency stems from our explicit parametric representation: all HDR attributes are optimized as direct Gaussian primitive parameters rather than latent network outputs, eliminating the need for auxiliary neural networks while preserving 3DGS's real-time rendering capability and extending its dynamic range coverage.

4.2 Qualitative Results

As shown in Fig. 6, both 3DGS and HDR-GS struggle to recover accurate structures and colors in HDR scenes, even though HDR-GS is trained with both HDR and multi-exposure LDR supervision. In contrast, LCD-GS reconstructs these areas with clearer and more complete details. Similarly, Fig. 7 shows that both 3DGS (Kerbl et al., 2023) and Mono-HDR-GS (Zhang et al., 2025) tend to produce floaters or blurred structures, whereas LCD-GS yields more faithful scene representations while preserving finer details. Fig. 8 further analyzes the relative ℓ_1 reconstruction error as a function of luminance on the desk scene which is defined as $|\mathbf{I}^{\text{lum}}_{\text{pred}} - \mathbf{I}^{\text{lum}}_{\text{gt}}|/\mathbf{I}^{\text{lum}}_{\text{gt}}$ where $\mathbf{I}^{\text{lum}}_{\text{pred}}$ and $\mathbf{I}^{\text{lum}}_{\text{gt}}$ denote rendered and ground-truth luminance images computed as a weighted sum of RGB channels, respectively. Vanilla 3DGS exhibits rapidly increasing relative ℓ_1 error in high-luminance HDR regions, suggesting that its color representation struggles to model large radiance variations. HDR-GS (Cai et al., 2024) and GaussHDR (Liu et al., 2025) suffer from high relative error across luminance, indicating inaccurate HDR luminance reconstruction. In contrast, LCD-GS maintains much lower error across the luminance range. This advantage comes from the proposed luminance–chromaticity decomposition, where the luminance term explicitly captures radiance magnitude and the bounded chromaticity term models view-dependent color variation, enabling more faithful HDR reconstruction for lighting-sensitive applications such as AR/VR (Van Bogaert et al., 2022).

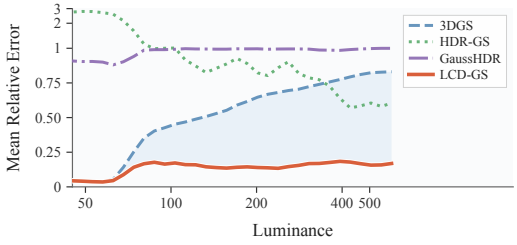


Figure 8: Mean ℓ_1 relative error across the luminance. LCD-GS maintains substantially lower error than 3DGS, HDR-GS and GaussHDR across the luminance range, particularly in bright regions.

5 Conclusion

We introduced LCD-GS, a lightweight HDR 3DGS framework that learns directly from native HDR observations without relying on multi-exposure LDR supervision. Our analysis shows that directly adapting vanilla 3DGS to HDR supervision requires a shared set of SHs to jointly encode wide-range luminance-like intensity and directional chromaticity, leading to substantial underfitting. LCD-GS explicitly separates these roles using a luminance–chromaticity decomposition strategy. This simple modification preserves the original 3DGS pipeline while substantially improving HDR reconstruction quality. Extensive experiments demonstrate that LCD-GS achieves superior reconstruction fidelity and dynamic-range preservation compared with LDR-oriented baselines and existing HDR NVS methods, while maintaining efficient training and real-time rendering. Its explicit luminance component further provides a simple handle for luminance editing naturally.

ETHICS STATEMENT

This work advances HDR 3D scene reconstruction and novel-view rendering from posed native HDR images. LCD-GS is intended for research and legitimate creative applications, including virtual reality, cinematic content creation, and immersive telepresence.

The captured real-world iHDR-4S dataset contains four indoor scenes captured using a tripod-mounted iPhone 14 Pro in ProRAW format. These scenes were captured with the informed consent of relevant participants and property owners, and the released data contain no personally identifiable information. Although LCD-GS is not designed for harmful use, neural scene reconstruction may be misused to create misleading or privacy-invasive visual content. We therefore encourage responsible data collection and release, transparency when synthetic media are used, and appropriate safeguards in downstream deployments.

REPRODUCIBILITY STATEMENT

We are committed to the reproducibility of LCD-GS. The code and datasets will be publicly released upon final acceptance. The method formulation, including the luminance–chromaticity decomposition and training objective, is specified in Sec. 3. Sec. 4 describes the datasets, baselines, metrics, and evaluation protocol; Appendix A.1 documents iHDR-4S, Appendix A.2 reports the implementation and optimization settings, and Appendix A.4 provides the luminance-activation ablation. Together, these details enable independent reproduction of the reported reconstruction and efficiency results.

AI USE STATEMENT

Generative AI tools were used to assist with manuscript formatting and migration to the ICLR template. The authors reviewed the resulting manuscript and take responsibility for its final content.

REFERENCES

Blender Foundation. Blender. `https://www.blender.org`, 2025.

Mark Boss, Raphael Braun, Varun Jampani, Jonathan T Barron, Ce Liu, and Hendrik Lensch. Nerd: Neural reflectance decomposition from image collections. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12684–12694, 2021.

Jintong Cai and Huimin Lu. NeRF-based multi-view synthesis techniques: A survey. In *International Wireless Communications and Mobile Computing*, pp. 208–213. IEEE, 2024.

Yuanhao Cai, Zihao Xiao, Yixun Liang, Minghan Qin, Yulun Zhang, Xiaokang Yang, Yaoyao Liu, and Alan L. Yuille. HDR-GS: Efficient high dynamic range novel view synthesis at 1000x speed via Gaussian splatting. *Advances in Neural Information Processing Systems*, 37:68453–68471, 2024.

Kai Cheng, Xiaoxiao Long, Kaizhi Yang, Yao Yao, Wei Yin, Yuexin Ma, Wenping Wang, and Xuejin Chen. GaussianPro: 3D Gaussian splatting with progressive propagation. In *Proceedings of Machine Learning Research*, volume 235, pp. 8123–8140, 2024.

Ziteng Cui, Xuangeng Chu, and Tatsuya Harada. Luminance-gs: Adapting 3d gaussian splatting to challenging lighting conditions with view-adaptive curve adjustment. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 26472–26482, 2025.

Anurag Dalal, Daniel Hagen, Kjell G. Robbersmyr, and Kristian Muri Knausgård. Gaussian splatting: 3D reconstruction and novel view synthesis: A review. *IEEE Access*, 12:96797–96820, 2024.

Kang Du, Zhihao Liang, Yulin Shen, and Zeyu Wang. Gs-id: Illumination decomposition on gaussian splatting via adaptive light aggregation and diffusion-guided material priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 26220–26229, 2025.

Shengjie Feng, Xiaoqun Wu, and Jian Cao. A survey of multi-view stereo 3D reconstruction algorithms based on deep learning. *Digital Signal Processing*, pp. 105291, 2025.

Yang Fu, Sifei Liu, Amey Kulkarni, Jan Kautz, Alexei A. Efros, and Xiaolong Wang. COLMAP-free 3D Gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 20796–20805, June 2024.

Xuan Gao, Wei Li, and Baojie Fan. Mip-NeRF+: Multi-scale 3D scene synthesis. In *International Conference on Intelligent Robotics and Applications*, pp. 174–188. Springer, 2024.

HDRsoft Team. Photomatix Pro. `https://www.hdrsoft.com/`, 2025.

Dahang Huang, Chengliang Wang, Ji Liu, Tao Cheng, and Jianbin Chen. Emissive-gs: Relightable reconstruction and emission editing via gaussian splatting. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 10317–10321. IEEE, 2026.

Hongbo Huang, Xiaoxu Yan, Yaolin Zheng, Jiayu He, Longfei Xu, and Dechun Qin. Multi-view stereo algorithms based on deep learning: A survey. *Multimedia Tools and Applications*, 84(6): 2877–2908, 2025.

Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, Xuan Wang, and Qing Wang. HDR-NeRF: High dynamic range neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18398–18408, 2022.

Arthur Hubert, Gamal Elghazaly, and Raphaël Frank. Editing implicit and explicit representations of radiance fields: a survey. *Machine Vision and Applications*, 36(6):119, 2025.

Xin Jin, Pengyi Jiao, Zheng-Peng Duan, Xingchao Yang, Chongyi Li, Chun-Le Guo, and Bo Ren. Lighting every darkness with 3DGS: Fast training and real-time rendering for HDR view synthesis. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Nima Khademi Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. *ACM Transactions on Graphics*, 36(4):144:1–144:12, 2017.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D Gaussian Splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023.

Jonas Kulhanek, Songyou Peng, Zuzana Kukelova, Marc Pollefeys, and Torsten Sattler. WildGaussians: 3D Gaussian splatting in the wild. In *Advances in Neural Information Processing Systems*, volume 37, pp. 21271–21288, 2024.

Zhihao Liang, Qi Zhang, Ying Feng, Ying Shan, and Kui Jia. Gs-ir: 3d gaussian splatting for inverse rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21644–21653, 2024.

Zhihao Liang, Hongdong Li, Kui Jia, Kailing Guo, and Qi Zhang. Gus-ir: Gaussian splatting with unified shading for inverse rendering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Jinfeng Liu, Lingtong Kong, Bo Li, and Dan Xu. GaussHDR: High dynamic range Gaussian splatting via learning unified 3D and 2D local tone mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5991–6000, 2025.

Xin Ma, Jiguang Zhang, Peng Lu, Shibiao Xu, and Chengwei Pan. Novel view synthesis under large-deviation viewpoint for autonomous driving. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 6000–6008, 2025.

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 *Proceedings of the European Conference on Computer Vision*, pp. 405–421. Springer, 2020.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16190–16199, 2022.

Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics*, 41(4):102:1–102:15, 2022.

Yuandong Niu, Limin Liu, Fuyu Huang, Siyuan Huang, and Shuangyou Chen. Overview of image-based 3D reconstruction technology. *Journal of the European Optical Society – Rapid Publications*, 20(1):18, 2024.

Linfei Pan, Dániel Baráth, Marc Pollefeys, and Johannes L. Schönberger. Global structure-from-motion revisited. In *Proceedings of the European Conference on Computer Vision*, pp. 58–77. Springer, 2024a.

Xiaonan Pan, Qilei Sun, Jia Wang, and Eng Gee Lim. Game engine based multi-view video dataset synthesis for pedestrian detection and tracking. In *IEEE International Conference on Metaverse Computing, Networking, and Applications*, pp. 259–264. IEEE, 2024b.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. PyTorch: An imperative style, high-performance deep learning library. *Advances in Neural Information Processing Systems*, 32: 8024–8035, 2019.

Lukas Radl, Michael Steiner, Andreas Kurz, and Markus Steinberger. Laenerf: local appearance editing for neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4969–4978, 2024.

Santosh Reddy, H. Abhiram, and K. S. Archish. A survey of 3D Gaussian splatting: Optimization techniques, applications, and AI-driven advancements. In *International Conference on Intelligent and Innovative Technologies in Computing, Electrical and Electronics*, pp. 1–6. IEEE, 2025.

Viktor Rudnev, Mohamed Elgharib, William Smith, Lingjie Liu, Vladislav Golyanik, and Christian Theobalt. Nerf for outdoor scene relighting. In *European Conference on Computer Vision*, pp. 615–631. Springer, 2022.

Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4104–4113, 2016.

Yahao Shi, Yanmin Wu, Chenming Wu, Xing Liu, Chen Zhao, Haocheng Feng, Jian Zhang, Bin Zhou, Errui Ding, and Jingdong Wang. Gir: 3d gaussian inverse rendering for relightable scene factorization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Shreyas Singh, Aryan Garg, and Kaushik Mitra. HDRSplat: Gaussian Splatting for high dynamic range 3D scene reconstruction from raw images. In *Proceedings of the 35th British Machine Vision Conference*, 2024.

Pratul P Srinivasan, Boyang Deng, Xiuming Zhang, Matthew Tancik, Ben Mildenhall, and Jonathan T Barron. Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 7495–7504, 2021.

Hanxiao Sun, Yupeng Gao, Jin Xie, Jian Yang, and Beibei Wang. Svg-ir: Spatially-varying gaussian splatting for inverse rendering. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 16143–16152, 2025.

Jia-Mu Sun, Tong Wu, Yong-Liang Yang, Yu-Kun Lai, and Lin Gao. Sol-nerf: Sunlight modeling for outdoor scene decomposition and relighting. In *SIGGRAPH Asia 2023 Conference Papers*, pp. 1–11, 2023.

Laurie Van Bogaert, Daniele Bonatto, Sarah Fernades Pinto Fachada, and Gauthier Lafruit. Novel view synthesis in embedded virtual reality devices. *Electronic Imaging*, 34:1–6, 2022.

Jian-Gang Wang, Lubing Zhou, Zhiwei Song, and Miaolong Yuan. Real-time vehicle signal lights recognition with HDR camera. In *IEEE International Conference on Internet of Things and IEEE Green Computing and Communications and IEEE Cyber, Physical and Social Computing and IEEE Smart Data*, pp. 355–358. IEEE, 2016.

Rui Wang, John Tran, and David Luebke. All-frequency relighting of glossy objects. *ACM Transactions on Graphics*, 25(2):293–318, 2006.

Zirui Wu, Jianteng Chen, Laijian Li, Shaoteng Wu, Zhikai Zhu, Kang Xu, Martin R Oswald, and Jie Song. 3d gaussian inverse rendering with approximated global illumination. *arXiv preprint arXiv:2504.01358*, 2025.

Dajun Xing, Ahmed Ouni, Stephanie Chen, Hinde Sahmoun, James Gordon, and Robert Shapley. Brightness–color interactions in human early visual cortex. *Journal of Neuroscience*, 35(5): 2226–2232, 2015.

Xingting Yao, Qinghao Hu, Fei Zhou, Tielong Liu, Zitao Mo, Zeyu Zhu, Zhengyang Zhuge, and Jian Cheng. SpiNeRF: Direct-trained spiking neural networks for efficient neural radiance field rendering. *Frontiers in Neuroscience*, 19:1593580, 2025.

Dingqiang Ye, Jiacong Xu, Jianglu Ping, Yuxiang Guo, Chao Fan, and Vishal M Patel. Instanthdr: Single-forward gaussian splatting for high dynamic range 3d reconstruction. *arXiv preprint arXiv:2603.11298*, 2026.

Weicai Ye, Shuo Chen, Chong Bao, Hujun Bao, Marc Pollefeys, Zhaopeng Cui, and Guofeng Zhang. Intrinsicnerf: Learning intrinsic neural radiance fields for editable novel view synthesis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 339–351, 2023.

Zongxin Ye, Wenyu Li, Sidun Liu, Peng Qiao, and Yong Dou. AbsGS: Recovering fine details in 3D Gaussian splatting. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 1053–1061, 2024.

Alex Yu, Ruilong Li, Matthew Tancik, Hao Li, Ren Ng, and Angjoo Kanazawa. PlenOctrees for real-time rendering of neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5752–5761, 2021.

Zehao Yu, Anpei Chen, Binbin Huang, Torsten Sattler, and Andreas Geiger. Mip-splatting: Alias-free 3D Gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 19447–19456, 2024.

Kaixuan Zhang, Hu Wang, Minxian Li, Mingwu Ren, Mao Ye, and Xiatian Zhu. High dynamic range novel view synthesis with single exposure. In *Proceedings of the 42nd International Conference on Machine Learning*, 2025.

Lin Zhu, Kangmin Jia, Yifan Zhao, Yunshan Qi, Lizhi Wang, and Hua Huang. SpikeNeRF: Learning neural radiance fields from continuous spike stream. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6285–6295, 2024.

A APPENDIX

A.1 DATASETS

This section provides a detailed description of the iHDR-4S dataset introduced in this work. Our iHDR-4S observations are *three-channel RGB HDR images*, rather than Bayer-mosaicked sensor measurements. Specifically, they are captured in Apple ProRAW and stored as 12-bit Linear DNG files. Unlike conventional Bayer RAW, Apple ProRAW images are already demosaiced and partially processed by the camera pipeline; their scene-referred, linearizable RGB values retain high-dynamic-range radiometric information. We therefore operate directly in the RGB-domain HDR space, without reconstructing RGB values from a Bayer mosaic. For COLMAP (Schönberger & Frahm, 2016), which does not support these HDR RGB DNG images as input, we generate LDR RGB counterparts only for sparse-point reconstruction and camera-pose estimation.



Figure 9: Dataset composition for four scenes of iHDR-4S. Each row corresponds to a distinct scene with HDR images tone-mapped by Photomatix Pro for display.

A.2 IMPLEMENTATION DETAILS

LCD-GS is built upon PyTorch (Paszke et al., 2019) and trained with Adam optimizer. The learning rate of luminance attribute is set to 0.05, and the activation function g for luminance is the softplus function. All the experiments are conducted with a single NVIDIA RTX 4090 GPU. λ and μ are set

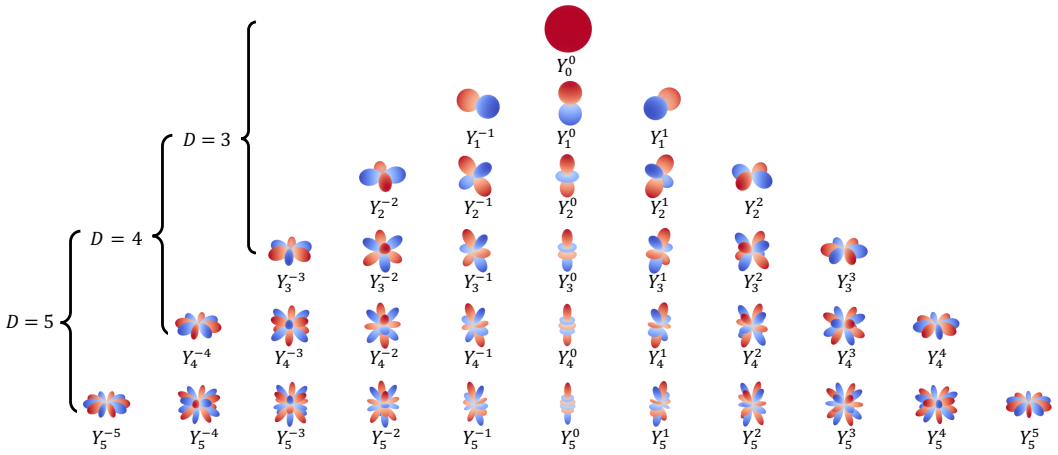


Figure 10: Spherical harmonic basis functions: maximum degrees $D = 3, 4, 5$ correspond to $(D+1)^2 = 16, 25, 36$ coefficients, respectively.

to 0.8 and 5000 during all the experiments. During initialization, the luminance term L in Eq. 3 is initialized as the sum of RGB values from the SfM point cloud.

When adapting NeRF (Mildenhall et al., 2020) and 3DGS (Kerbl et al., 2023) for HDR image training on the synthetic dataset, we replace their output activation function with softplus (see Tab. 4 for more results of 3DGS with different activation functions), enabling the model to predict unbounded radiance values. For a fair comparison, the loss function employed for 3DGS (as well as it variants, *e.g.*, 3DGS-D4/5) is identical to that defined in Eq. (4). If necessary, compared methods (Cai et al., 2024; Liu et al., 2025; Zhang et al., 2025) are granted access to LDR inputs per their original implementations. For all experiments, we uniformly downsample the input images by a factor of 2 to reduce computational cost.

A.3 Additional Related Work

Illumination editing in neural 3D representations is commonly studied through relighting and inverse rendering (Hubert et al., 2025). NeRF-based methods (Srinivasan et al., 2021; Boss et al., 2021; Rudnev et al., 2022; Sun et al., 2023) typically decompose scene appearance into geometry, reflectance, visibility, and illumination, enabling novel lighting synthesis by modifying the estimated lighting or material components. Recent works (Ye et al., 2023; Radl et al., 2024) further extend this direction to intrinsic decomposition and lighting-aware neural fields, aiming to separate reflectance and shading for controllable appearance editing. More recently, illumination-aware editing has also been explored in 3D Gaussian Splatting (Liang et al., 2024; Shi et al., 2025; Liang et al., 2025; Sun et al., 2025; Du et al., 2025; Cui et al., 2025). These approaches are powerful for general relighting, but they often introduce extra networks (Liang et al., 2024; Shi et al., 2025), physically based rendering assumptions (Wu et al., 2025), or dedicated illumination estimation modules (Huang et al., 2026).

Unlike these relighting-oriented methods, LCD-GS does not aim to recover a full physical decomposition of lighting, material, and visibility, nor does it target arbitrary light-source editing. Instead, it provides a lightweight and direct form of *luminance-based HDR scene editing*. By explicitly decoupling each Gaussian primitive into a view-independent luminance scalar and a normalized chromaticity term, LCD-GS enables manipulation of radiance magnitude while keeping chromaticity unchanged. This makes our editing operation particularly suitable for HDR-specific adjustments, such as enhancing shadowed regions (see Fig. 5) or compressing overly bright areas (see Fig. 12), without introducing undesired hue shifts.

A.4 Ablation Study

Activation functions. We evaluate three luminance activation functions (*i.e.*, softplus, exponential and ReLU) on both the HDR-Syn-8S and iHDR-4S datasets. Softplus achieves the best reconstruction quality and inference speed on both datasets while remaining competitive in training time and memory consumption. Its smooth, non-saturating gradients facilitate stable convergence and accurate

Table 4: Ablation studies of different activation functions for LCD-GS and 3DGS on both HDR-Syn-8S and iHDR-4S datasets. The **best** and second-best results on each dataset are highlighted.

Method	Activation	Dataset	Training Time (min) ↓	Training Memory (MB) ↓	HDR PSNR ↑	HDR SSIM ↑	HDR LPIPS ↓	Inference speed (fps) ↑
LCD-GS	Exponential	HDR-Syn-8S	**12**	1278	39.98	0.975	0.009	288
	ReLU		13	1655	37.03	0.978	0.019	255
	Softplus		**12**	**1205**	**41.84**	**0.991**	**0.004**	**300**
LCD-GS	Exponential	iHDR-4S	23	3378	27.56	0.746	0.303	180
	ReLU		**22**	**3257**	26.23	0.739	0.323	181
	Softplus		**22**	3311	**27.74**	**0.749**	**0.295**	**197**
3DGS	Exponential	HDR-Syn-8S	**7**	**808**	30.66	0.867	0.108	**354**
	ReLU		8	979	**31.82**	**0.897**	**0.064**	350
	Softplus		**7**	830	31.53	0.871	0.079	320
3DGS	Exponential	iHDR-4S	**11**	**2580**	25.23	0.733	0.339	**277**
	ReLU		12	2883	**26.18**	**0.749**	0.300	240
	Softplus		14	2949	25.96	0.737	**0.296**	258

Table 5: Ablation studies of μ-law function for LCD-GS and 3DGS on the HDR-Syn-8S dataset.

Method	μ-law	Training Time (min) ↓	Training Memory (MB) ↓	HDR PSNR ↑	HDR SSIM ↑	HDR LPIPS ↓	Inference speed (fps) ↑
3DGS	✓	**7**	**830**	**31.53**	**0.871**	**0.079**	**320**
	✗	19	1558	30.47	0.851	0.124	202
LCD-GS	✓	**12**	**1205**	**41.84**	**0.991**	**0.004**	**300**
	✗	24	2061	39.96	0.982	0.005	194

luminance estimation. In contrast, ReLU's thresholding behavior for negative inputs results in gradient truncation, which impedes optimization. While the exponential function preserves gradients, its unbounded nature causes slight numerical instabilities. As a result, softplus provides the most favorable overall trade-off for HDR luminance modeling, as demonstrated in Tab. 4.

μ-law function. To verify the effectiveness of the proposed LCD-GS while ruling out the influence of the μ-law function, we conduct additional experiments on the HDR-Syn-8S dataset. As shown in Tab. 5, the decomposition is the primary source of improvement, while μ-law provides a complementary optimization benefit that does not account for the large gap between representations.

A.5 Luminance Editing

This section provides additional analysis and examples of the luminance-editing capability enabled by LCD-GS. Unlike physically based relighting or inverse rendering methods, our goal is not to recover explicit light sources, materials, visibility, or illumination transport. Instead, LCD-GS provides a simple interface for HDR radiance-magnitude editing by manipulating the learned luminance scalar of each Gaussian primitive while keeping its normalized chromaticity unchanged. Thanks to the explicit separation of luminance and chromaticity, global exposure control can be implemented as a direct transformation of the luminance component. In this work, we use a linear mapping as an illustrative example:

$$L'_{global} = \gamma \cdot L, \tag{5}$$

where $\gamma \in \mathbb{R}^+$ is the luminance scaling factor. $\gamma > 1$ increases the overall radiance magnitude, simulating higher exposure or stronger illumination, while $\gamma < 1$ produces a dimming effect, as shown in Fig. 11.

LCD-GS also supports selective luminance editing at the Gaussian-primitive level. As shown in Fig. 5 and Fig. 12, we modify the luminance values of selected primitives during rendering while keeping chromaticity, opacity, and geometry unchanged. This enables targeted HDR radiance manipulation, such as suppressing highlights on the book in Fig. 12b–e and enhancing low-luminance regions ($L \leq \tau$, where τ is a predefined threshold) on the floor in Fig. 5b–e. Since these edits are performed before alpha compositing through the explicit luminance component, they provide a controllable post-reconstruction editing interface that cannot be achieved by simply scaling the rendered image.



(a) Reference (b) $\gamma = 0.1$ (c) $\gamma = 0.2$ (d) $\gamma = 0.3$ (e) $\gamma = 0.4$



(f) $\gamma = 0.5$ (g) $\gamma = 0.6$ (h) $\gamma = 0.7$ (i) $\gamma = 0.8$ (j) $\gamma = 0.9$



(k) $\gamma = 1.1$ (l) $\gamma = 1.2$ (m) $\gamma = 1.3$ (n) $\gamma = 1.4$ (o) $\gamma = 1.5$



(p) $\gamma = 1.6$ (q) $\gamma = 1.7$ (r) $\gamma = 1.8$ (s) $\gamma = 1.9$ (t) $\gamma = 2.0$

Figure 11: Luminance editing with LCD-GS. (a) Reference image. (b)-(t) are obtained by scaling the learned luminance values of all Gaussian primitives with different γ.



(a) Reference (b) (c) (d) (e)

Figure 12: Luminance editing enabled by LCD-GS. (a) Reference image. (b)-(e) Highlight suppression by scaling the luminance of a subset of Gaussian primitives.