PRIOR-GUIDED SINGLE-EXPOSURE HDR NOVEL VIEW SYNTHESIS VIA BASE-GAIN REPRESENTATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Single-exposure high dynamic range (HDR) novel view synthesis (NVS) recovers HDR scenes from low dynamic range (LDR) views captured at a single exposure and estimates camera response function (CRF) for LDR re-exposure. Without HDR ground-truth, using single-exposure LDR views as the sole source of direct supervision underconstrains HDR reconstruction and cannot ensure its faithful recovery, motivating the need for learned LDR-to-HDR reconstruction priors. Moreover, single-exposure LDR views supervise CRF only over a limited range, necessitating a learnable prior for reliable re-exposure beyond it. We propose PriorGS, a 3D Gaussian Splatting framework that addresses both challenges by (1) guiding HDR recovery with pseudo-HDR supervision from 2D single-exposure HDR reconstruction methods while anchoring it to input LDR views through a closed reconstruction loop; and (2) regularizing CRF modeling with a learnable prior that guarantees monotonicity and promotes plausible responses. Yet cross-view inconsistencies in independently generated pseudo-HDR targets impair scene reconstruction when HDR is encoded directly as Gaussian's color attribute. We therefore introduce a 3D base–gain HDR representation that factors HDR radiance into an LDR base and an overlaid 3D gain field, promoting robustness to such inconsistencies and enabling 3D HDR appearance editing. Extensive experiments across public datasets demonstrate PriorGS's substantial superiority over baselines. A short demo video is included in the supplementary materials.

1 INTRODUCTION

High dynamic range (HDR) novel view synthesis (NVS) aims to reconstruct a 3D scene's geometry and HDR radiance from multi-view low dynamic range (LDR) observations, with the dual objective of synthesizing novel HDR views and rendering their LDR counterparts at unseen exposures. Compared with conventional LDR-NVS (Mildenhall et al., 2021; Kerbl et al., 2023), HDR-NVS represents a wider range of scene illumination, preserving textures in both highlights and shadows. This richer representation improves the visual appeal and realism of reconstructed scenes and benefits many downstream applications (Debevec, 1998; Zhang et al., 2021; Gomez-Ojeda et al., 2018). Most HDR-NVS methods rely on multi-exposure LDR captures (Huang et al., 2022; Cai et al., 2024; Liu et al., 2025) that typically require static scenes and tedious acquisition to avoid motion artifacts. Motivated by this limitation, recent methods, including ours, target HDR-NVS with more accessible single-exposure LDR inputs (Li et al., 2025; Zhang et al., 2025). Also, as calibrated per-view HDR capture is rarely practical, we target HDR recovery without access to HDR ground-truth (GT).

HDR scene radiance represents linear scene-referred light over a wide intensity range. During LDR capture, this radiance is converted to sensor exposure, transformed by a nonlinear camera response function (CRF), quantized, and partially lost through highlight clipping and shadow noise Grossberg & Nayar (2004). Recovering HDR radiance from this degraded single-exposure observation is hence fundamentally underconstrained (Eilertsen et al., 2021), motivating 2D single-exposure HDR reconstruction methods that learn LDR-to-HDR reconstruction priors from large data (Dille et al., 2025; Liu et al., 2020). The same challenge arises in single-exposure HDR-NVS: using single-exposure views as the only source of direct supervision leaves HDR reconstruction underconstrained and cannot ensure its accurate recovery. This motivates the need for learned LDR-to-HDR reconstruction priors to guide HDR recovery in single-exposure HDR-NVS, a critical factor overlooked by prior works. Figure 1 illustrates this: when trained on the same single-exposure inputs, existing methods

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.

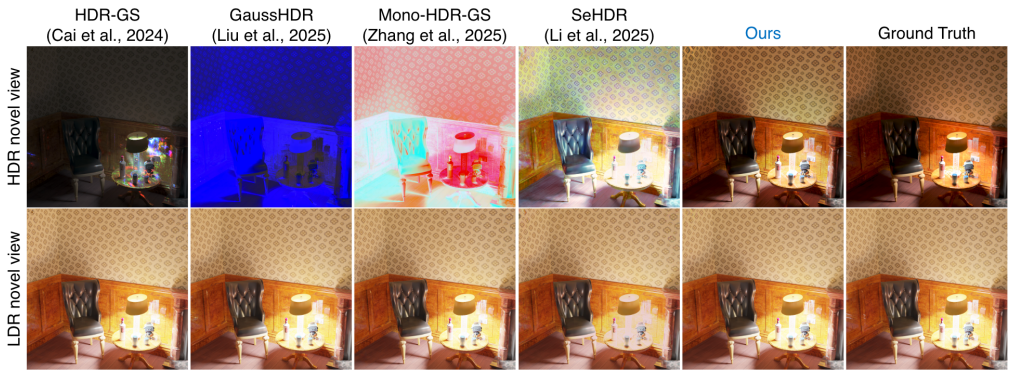

Figure 1: Given the same single-exposure inputs, methods lacking learned LDR-to-HDR priors reconstruct the LDR novel view faithfully but fail at HDR radiance recovery. We leverage this prior to recover radiance much closer to the GT. HDR results are tone-mapped identically for visualization.

lacking learned LDR-to-HDR priors reproduce the observed exposure yet recover an incorrect HDR radiance. Our method, however, achieves high-fidelity HDR recovery by leveraging this prior.

Beyond high-fidelity HDR recovery, a holistic HDR-NVS method must accurately re-expose HDR radiance to LDR views at unseen exposures through CRF modeling, an objective often bypassed by existing single-exposure methods. Unlike multi-exposure captures that enable supervision for a wide CRF range, single-exposure views supervise only the camera's limited unsaturated range at one exposure, leaving the CRF unconstrained elsewhere. Existing single-exposure HDR-NVS methods either (1) learn flexible MLP-based CRFs (Zhang et al., 2025) that fit the observed exposure yet neither generalize well to unseen ones nor ensure the expected monotonicity (i.e. LDR intensity grows with input irradiance) (Grossberg & Nayar, 2004); or (2) fix a heuristically selected CRF per scene (Li et al., 2025) that risks response mismatch and precludes training-time adaptation. Single-exposure HDR-NVS thus needs a learnable CRF prior that regularizes response beyond the supervised range, guarantees monotonicity, and promotes plausible responses while retaining scene adaptability.

In response, we propose PriorGS, a novel single-exposure HDR-NVS method based on 3D Gaussian Splatting (3DGS) (Kerbl et al., 2023) that incorporates *LDR-to-HDR reconstruction* and *CRF priors* within a *base–gain HDR representation*. To incorporate learned LDR-to-HDR reconstruction priors, we supervise HDR rendering at training views using pseudo-HDR targets generated by 2D single-exposure HDR reconstruction methods. However, as these targets are estimated separately per views, they may exhibit cross-view inconsistencies. As shown in Figures 2 and 16, when HDR radiance is encoded directly as Gaussian's color attribute (direct HDR encoding), such inconsistencies impair scene reconstruction, causing severe color and geometry artifacts. To mitigate this, we represent HDR radiance as an LDR base supervised by LDR views and extended to HDR through a 3D gain field. This base-gain representation (1) anchors the scene to the more reliable LDR views while modeling the less reliable HDR extension by the 3D gain field, thereby promoting robustness to cross-view inconsistencies (Figures 2 and 16); and (2) enables post-training HDR appearance editing in 3D without altering the LDR base. Next, we regularize CRF modeling with a learnable prior derived from a database of CRFs (Grossberg & Nayar, 2004). The CRF

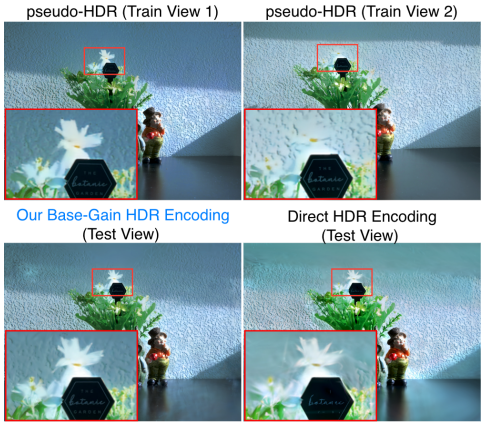

Figure 2: (*top row*) Cross-view inconsistencies in pseudo-HDR targets (*bottom row*) produce severe artifacts w/ direct HDR encoding, while our base–gain representation promotes robustness to such inconsistencies.

then re-exposes HDR renderings to match the same LDR views, closing an LDR-base → HDR → re-exposed LDR loop that anchors HDR recovery to accurate LDR reconstruction at both ends.

Our contributions can be summarized as: (1) We propose PriorGS, a novel 3DGS-based single-exposure HDR-NVS framework that guides HDR reconstruction with learned LDR-to-HDR reconstruction priors while enforcing consistent LDR reconstruction before and after HDR recovery; (2) We introduce a base-gain 3D HDR representation that factors scene radiance into an LDR base and an overlaid 3D gain field, promoting robustness to cross-view pseudo-HDR inconsistencies and enabling HDR appearance editing in 3D; (3) We regularize CRF modeling in single-exposure HDR-NVS with a learnable prior, guaranteeing monotonicity and promoting plausible responses while retaining scene adaptability; and (4) Extensive experiments demonstrate that PriorGS substantially outperforms existing methods both quantitatively and qualitatively.

2 Related Works

HDR Novel View Synthesis seeks to reconstruct HDR scenes from either (1) LDR captures (Huang et al., 2022; 2024; Cai et al., 2024; Niemeyer et al., 2025); (2) RAW images (Mildenhall et al., 2022; Jin et al., 2024; Li et al., 2024; Singh et al., 2024); or (3) event–image pairs (Chen et al., 2026; Qi et al., 2026). We target the first as LDR images are easier to acquire and more widely available. Most LDR-based methods require multi-exposure LDR captures to jointly learn HDR radiance and LDR re-exposure (Cai et al., 2024), with recent approaches exploring unified 3D and 2D local tone-mapping (Liu et al., 2025), physically-inspired appearance modeling (Zeng et al., 2026), and color-geometry-exposure disentanglement (Song et al., 2026). However, these methods rely on multi-exposure supervision and thus do not address the underconstrained nature of single-exposure HDR-NVS. This limitation, together with tedious acquisition of multi-exposure images has motivated recent HDR-NVS methods using more accessible single-exposure LDR inputs (Li et al., 2025; Zhang et al., 2025). Among single-exposure methods, SeHDR (Li et al., 2025) fuses bracketed 3D Gaussians into an HDR representation, whereas Mono-HDR-GS (Zhang et al., 2025) learns MLP-based LDR-to-HDR and HDR-to-LDR mappings. However, without HDR GT, their HDR radiance recovery remains underconstrained, lacking direct guidance for accurate HDR reconstruction and thus leading to unreliable recovery. PriorGS overcomes this by guiding HDR recovery with learned LDR-to-HDR reconstruction priors. Outside exposure-controllable HDR-NVS, PanoHDR-NeRF (Gera et al., 2022) and GaSLight (Bolduc et al., 2025) use LDR-to-HDR predictions to train 3D HDR representations for relighting. However, they neither enforce consistency with original LDR views nor support LDR re-exposure. In contrast, PriorGS anchors prior-guided HDR recovery to accurate LDR reconstruction and re-exposure through closed-loop reconstruction. Finally, HDR-NVS requires reliable CRF modeling for re-exposure at unseen exposures. Multi-exposure methods employ either neural tone mappers (Cai et al., 2024; Liu et al., 2025; Ye et al., 2026) or explicit response parameterizations (Jun-Seong et al., 2022). These CRFs are constrained by multi-exposure inputs across a wide response range, while single-exposure inputs supervise only a narrow unsaturated region, leaving the remaining range unconstrained. Among single-exposure methods, Mono-HDR-GS (Zhang et al., 2025) learns a flexible MLP-based CRF that fits the observed exposure but offers poor generalization beyond it and cannot ensure monotonicity. SeHDR (Li et al., 2025) instead heuristically fixes a CRF per scene but risks response mismatch and forfeits training-time adaptation. To address both limitations, PriorGS regularizes CRF modeling with a learnable prior that guarantees monotonicity expected of CRFs (Grossberg & Nayar, 2004) and promotes plausible responses.

Single-exposure 2D HDR Reconstruction aims to recover an HDR image from a single LDR observation by reversing camera pipeline and inferring radiance degraded or lost through noise, quantization, and saturation. Approaches range from CNN- or Transformer-based prediction (Eilertsen et al., 2017; Marnerides et al., 2018; Liu et al., 2020; Dille et al., 2025; Santos et al., 2020; Zhang et al., 2023) to GAN- and diffusion-based generation (Wang et al., 2025; 2023; Liu et al., 2026; Wu et al., 2026), with recent works adopting dual-layer LDR-gain formulation that generates a 2D gain map over an LDR base (Liao et al., 2025; Guan et al., 2025). A solution for single-exposure HDR-NVS is to use these methods to generate pseudo-HDR targets for views and train 3DGS with direct HDR encoding–an encoding prevalent in prior works (Cai et al., 2024; Liu et al., 2025). However, cross-view inconsistencies in per-view generated pseudo-HDR targets impair scene reconstruction. PriorGS instead extends the 2D LDR–gain formulation to 3D, yielding a base–gain HDR representation that promotes robustness to such inconsistencies and enables 3D HDR appearance editing.

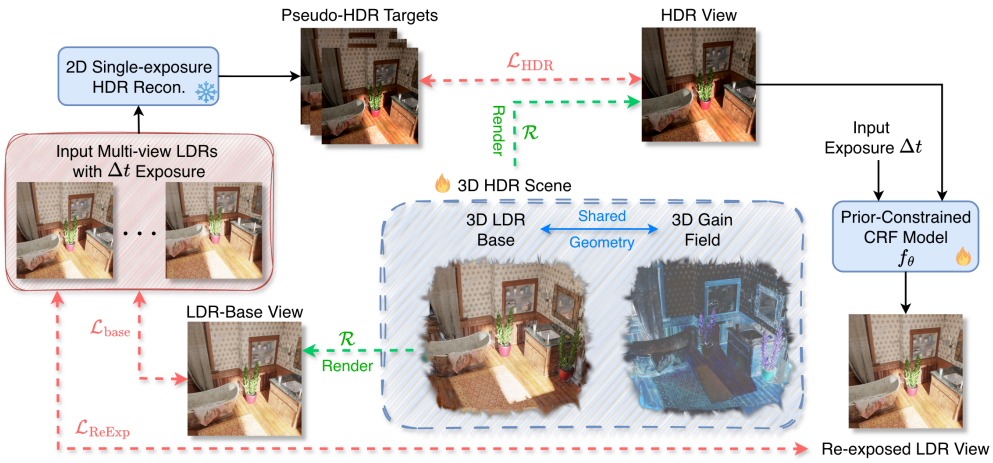

Figure 3: Overview of **PriorGS**. Given single-exposure multi-view LDR captures of a scene, we reconstruct the 3D HDR scene as an LDR base extended by a 3D gain field. Input views supervise the LDR base, while pseudo-HDR targets generated by a 2D single-exposure HDR reconstruction model impose LDR-to-HDR prior on HDR recovery. We then re-expose the HDR view by a prior-constrained learnable CRF to match the same input view, which closes the reconstruction loop.

3 Proposed Method

Problem Statement. Given V posed LDR views L^v of a 3D scene, $\{(L^v, \boldsymbol{\pi}^v)\}_{v=1}^V$, captured with camera parameters $\boldsymbol{\pi}^v$ at the same exposure Δt, we aim to jointly recover the HDR scene as a set of N Gaussian primitives $\mathbb{G}$ (Kerbl et al., 2023) together with a scene-specific CRF $f_{\boldsymbol{\theta}}$, where

$$\mathbb{G} = \{G_i(\boldsymbol{\mu}_i, \boldsymbol{\Sigma}_i, \alpha_i, \boldsymbol{\beta}_i^c, \boldsymbol{\beta}_i^g)\}_{i=1}^N. \tag{1}$$

Here, G_i represents the i-th Gaussian and $\boldsymbol{\theta}$ parameterizes learnable CRF $f_{\boldsymbol{\theta}}$ for LDR re-exposure. Each Gaussian G_i is composed of the following learnable parameters: (1) a position $\boldsymbol{\mu}_i \in \mathbb{R}^3$; (2) a 3D covariance $\boldsymbol{\Sigma}_i = \boldsymbol{R}_i \boldsymbol{S}_i \boldsymbol{S}_i^{\mathrm{T}} \boldsymbol{R}_i^{\mathrm{T}}$ represented by rotation and scale matrices $\boldsymbol{R}_i, \boldsymbol{S}_i \in \mathbb{R}^{3\times3}$; (3) an opacity $\alpha_i \in \mathbb{R}$ controlling transparency; (4) RGB spherical harmonics (SH) coefficients $\boldsymbol{\beta}_i^c$ encoding view-dependent LDR-base color $\boldsymbol{c}_i^v \in \mathbb{R}^3$; and (5) a second set of RGB SH coefficients $\boldsymbol{\beta}_i^g$ encoding view-dependent log-gain $\boldsymbol{g}_i^v \in \mathbb{R}^3$ that lifts the LDR base to HDR radiance. We follow the standard 3DGS rasterization of Kerbl et al. (2023) and refer readers there for details.

Overview. Figure 3 shows an overview of PriorGS. We represent the 3D HDR scene using shared Gaussian geometry with two color components: an LDR base and a 3D gain field. Each Gaussian's gain lifts its LDR-base color to HDR radiance before rasterization, producing LDR-base and HDR views from the same geometry. A pretrained single-exposure HDR reconstruction method generates a pseudo-HDR target for each training view. The input LDR view and its pseudo-HDR target then supervise the LDR-base and HDR renderings, respectively. We next re-expose the HDR rendering with a prior-constrained learnable CRF and require the output to match the same LDR view.

3.1 Prior-Guided HDR Reconstruction via Base–Gain HDR Representation

Pseudo-HDR Guidance. In Section 1, we discussed the necessity of using learned LDR-to-HDR reconstruction priors for accurate HDR reconstruction from single-exposure LDR views. We incorporate this prior through direct pseudo-HDR supervision: before training, we feed each training LDR view to a pretrained 2D single-exposure HDR reconstruction method and use the resulting pseudo-HDR target to supervise the rendered HDR view during training.

Base-Gain HDR Representation. As pseudo-HDR targets are generated independently for each view, they may have cross-view inconsistencies. When HDR radiance is encoded directly as a free-standing Gaussian color attribute (direct HDR encoding), such inconsistencies can impair scene reconstruction by driving unrestrained erroneous updates to Gaussians' attributes, resulting in severe color and geometry artifacts (see Figures 2 and 16). To mitigate this and inspired by the modern

LDR-gain 2D HDR formats (ISO/TC 42, 2025; Android Developers, 2024; Canham et al., 2025), we propose a base-gain 3D HDR representation that factors HDR radiance into an LDR base and an overlaid 3D gain field. Specifically, we implement each Gaussian G_i to carry two sets of RGB SH coefficients, $\boldsymbol{\beta}_i^c$ and $\boldsymbol{\beta}_i^g$, whose evaluations for view v produce its LDR-base color $\boldsymbol{c}_i^v$ and log-gain $\boldsymbol{g}_i^v$, respectively. We then derive its view-dependent HDR radiance $\boldsymbol{h}_i^v \in \mathbb{R}^3$ as

$$\boldsymbol{h}_i^v = (\boldsymbol{l}_i^v + \delta) \odot 2^{\,\boldsymbol{g}_i^v} - \delta, \tag{2}$$

where $\boldsymbol{l}_i^v$ is the linear-light value obtained from $\boldsymbol{c}_i^v$ via inverse gamma encoding, and $\delta = 1/64$ facilitates gain application near blacks ($\boldsymbol{l}_i^v = 0$). Finally, for a camera $\boldsymbol{\pi}^v$, we render the LDR base $\widehat{L}_{\mathrm{B}}^v$ and HDR radiance $\widehat{H}^v$ with shared geometry but different color attributes $\boldsymbol{c}_i^v$ and $\boldsymbol{h}_i^v$:

$$\widehat{L}_{\mathrm{B}}^v = \mathcal{R}\big(\mathbb{G}, \boldsymbol{\pi}^v; \{\boldsymbol{c}_i^v\}_{i=1}^N\big)\,, \qquad \widehat{H}^v = \mathcal{R}\big(\mathbb{G}, \boldsymbol{\pi}^v; \{\boldsymbol{h}_i^v\}_{i=1}^N\big)\,. \tag{3}$$

This enables joint supervision of LDR-base and HDR renderings by LDR views and pseudo-HDR targets, respectively. Our base-gain representation thus models HDR radiance as a gain-based extension of LDR base rather than as a free color field. The reliable LDR views therefore anchor the base appearance and shared geometry, while 3D gain field supplies dedicated capacity for modeling the less reliable HDR extension, thereby promoting robustness to cross-view inconsistencies (see Figures 2 and 16). Appendix A.1 provides further discussions.

3.2 Prior-Constrained CRF Modeling & Closed-loop Reconstruction

Prior-Constrained CRF Modeling. Section 1 motivated regularized CRF modeling in single-exposure HDR-NVS with a learnable prior. We derive this prior from DoRF (Grossberg & Nayar, 2004), a database of CRFs, by adopting the two-parameter monotonic approximation of its mean response proposed by Eilertsen et al. (2017). As DoRF responses and the approximation are defined over normalized irradiance $x \in [0, 1]$, we prepend a learnable soft normalization that preserves monotonicity and enhances representational capacity on DoRF responses (see Appendix A.2 for proof and capacity analysis). Let $\boldsymbol{\theta} = (\boldsymbol{\phi}, \boldsymbol{\psi})$, $\boldsymbol{\phi} = (\sigma, n)$ and $\boldsymbol{\psi} = (a, b)$. Our CRF is defined as

$$\begin{aligned} f_{\boldsymbol{\theta}}(E) &= t_{\boldsymbol{\phi}}(\eta_{\boldsymbol{\psi}}(E))\,, \\ t_{\boldsymbol{\phi}}(x) = (1+\sigma)\frac{x^n}{x^n+\sigma}, &\qquad \eta_{\boldsymbol{\psi}}(E) = \mathrm{sigmoid}(a\ln(E+\epsilon)+b)\,. \end{aligned} \tag{4}$$

Here $\sigma, n > 0$ parameterize the mean response approximation, while $a > 0$ and $b \in \mathbb{R}$ govern the normalization. We learn all four parameters per scene to achieve scene-adaptability and share the learned CRF across RGB channels. Our CRF is strictly monotonic by construction, ensuring LDR intensity increases with irradiance, while DoRF-derived form of $t_{\boldsymbol{\phi}}$ promotes plausible responses.

Closed-loop Reconstruction. As shown in Figure 3, we re-expose the HDR rendering $\widehat{H}^v$ at the training exposure Δt using the learned CRF to obtain the re-exposed LDR view as

$$\widehat{L}_{\mathrm{R}}^v(\Delta t) = f_{\boldsymbol{\theta}}\left(\Delta t\, \widehat{H}^v\right). \tag{5}$$

At test time, varying Δt enables re-exposure at arbitrary exposures. By supervising the re-exposed view $\widehat{L}_{\mathrm{R}}^v(\Delta t)$ with the input LDR view, we close the reconstruction loop $\widehat{L}_{\mathrm{B}}^v \to \widehat{H}^v \to \widehat{L}_{\mathrm{R}}^v(\Delta t)$, where HDR reconstruction is anchored upstream to LDR base through 3D gain field and constrained downstream to recover the same LDR view after re-exposure. Consequently, HDR recovery is not only guided by LDR-to-HDR prior but also constrained to consistently reproduce input LDR views.

Training Objective. We jointly optimize reconstruction of LDR-base, HDR, and re-exposed LDR views. We follow prior works (Cai et al., 2024; Li et al., 2025) to use a common reconstruction loss that combines pixel-wise $\mathcal{L}_1$ loss with D-SSIM (Wang et al., 2004) loss as

$$D_\lambda(A, B) = \mathcal{L}_1(A, B) + \lambda\, \mathcal{L}_{\text{D-SSIM}}(A, B), \tag{6}$$

where λ balances the two terms. The overall training objective is then formulated as

$$\mathcal{L} = \mathcal{L}_{\text{base}} + w_1\, \mathcal{L}_{\text{HDR}} + w_2\, \mathcal{L}_{\text{ReExp}}, \tag{7}$$

$$= D_\lambda(\widehat{L}_{\mathrm{B}}^v, L^v) + w_1\, D_\lambda\left(T_\mu(\widehat{H}^v), T_\mu(\widetilde{H}^v)\right) + w_2\, D_\lambda\left(\widehat{L}_{\mathrm{R}}^v(\Delta t), L^v\right). \tag{8}$$

Here, $\widetilde{H}^v$ denotes the pseudo-HDR target, $w_1 = w_2 = 0.5$ control the effect of loss terms, and T_μ represents μ-law range compressor (Kalantari & Ramamoorthi, 2017) widely used in prior works for HDR benchmarking and loss optimization (Cai et al., 2024; Zhang et al., 2025).

Table 1: Quantitative results of (*top half*) re-exposed LDR-NVS on real scenes of HDR-NeRF and HDR-Plenoxels datasets and (*bottom half*) HDR-NVS and re-exposed LDR-NVS on synthetic scenes of HDR-NeRF dataset. red and yellow denote the best and second-best results, respectively.

Method	HDR-NeRF-Real – LDR			HDR-Plenoxels-Real – LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
HDR-GS	11.10	0.5212	0.370	4.70	0.2799	0.456
GaussHDR	16.28	0.7849	0.176	17.73	0.7141	0.197
Mono-HDR-GS	11.43	0.6627	0.301	9.79	0.5570	0.389
SeHDR	13.75	0.6991	0.182	14.21	0.6644	0.232
PriorGS (Ours)	24.31	0.9248	0.064	23.00	0.9021	0.092
Method	HDR-NeRF-Syn – HDR			HDR-NeRF-Syn – LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
HDR-GS	26.11	0.8939	0.113	10.86	0.3429	0.397
GaussHDR	26.96	0.8828	0.114	11.99	0.5848	0.302
Mono-HDR-GS	16.71	0.5642	0.359	12.44	0.5640	0.422
SeHDR	15.14	0.6405	0.262	15.76	0.7450	0.169
PriorGS (Ours)	31.22	0.9569	0.038	22.96	0.8937	0.081

4 Experiments

4.1 Experimental Settings

Datasets & Implementation Details. We evaluate PriorGS on two public datasets: HDR-NeRF (Huang et al., 2022) and HDR-Plenoxels (Jun-Seong et al., 2022). Each dataset contains 4 real scenes without HDR GT and captured with digital cameras. HDR-NeRF additionally provides 8 synthetic scenes created in Blender (Blender Foundation, 2026) with HDR GT. Note that HDR GTs are not involved in training of any method and are exclusively used for evaluation. In each scene, each view is available at 5 exposures $\{\Delta t_i\}_{i=1}^{5}$. We use a single-exposure training strategy, where we train all methods on one exposure and evalute re-exposed LDR-NVS at the remaining 4 exposures along with HDR-NVS. We repeat this protocol for all 5 training exposures and report the average results. We use 18 training and 17 test views per HDR-NeRF scene and 27 training and 13 test views per HDR-Plenoxels scene. We initialize the scenes using COLMAP (Schönberger & Frahm, 2016) and train them for 30K iterations on an RTX 3090 GPU. We select IntrinsicHDR (Dille et al., 2025) as our 2D HDR reconstruction method for its robustness. Appendix A.3 provides further details.

Metrics & Baselines. We evaluate HDR-NVS and re-exposed LDR-NVS using PSNR, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018). In HDR quantitative comparisons, we follow Hanji et al. (2022) and apply luminance correction once per scene to adjust the scale ambiguity inherent in single-exposure HDR reconstruction (Cao et al., 2024). Next, following Cai et al. (2024), we apply μ-law range compression and then compute HDR metrics. We omit both steps for HDR qualitative comparisons and visualize HDR images after only a fixed tone-mapping. We compare PriorGS with all existing single-exposure methods Mono-HDR-GS (Zhang et al., 2025) and SeHDR (Li et al., 2025), and multi-exposure methods GaussHDR (Liu et al., 2025) and HDR-GS (Cai et al., 2024) adapted to single-exposure setting. For fairness to other methods, we use GaussHDR's 3DGS variant (Kerbl et al., 2023) and assess its Scaffold-GS (Lu et al., 2024) variant in Appendix A.4. As SeHDR is closed-source, we reproduce it to our best ability. Appendix A.3 offers further details.

4.2 Performance Comparison & Application

Quantitative. Table 1 summarizes the quantitative results. PriorGS consistently outperforms all baselines by large margins across both datasets, real and synthetic scenes, and both HDR-NVS and re-exposed LDR-NVS. This shows the effectiveness of our LDR-to-HDR prior integration with closed-loop reconstruction, base–gain representation, and regularized CRF modeling. Appendix A.4 offers more results, including PriorGS's robustness to the choice of 2D HDR reconstruction method.

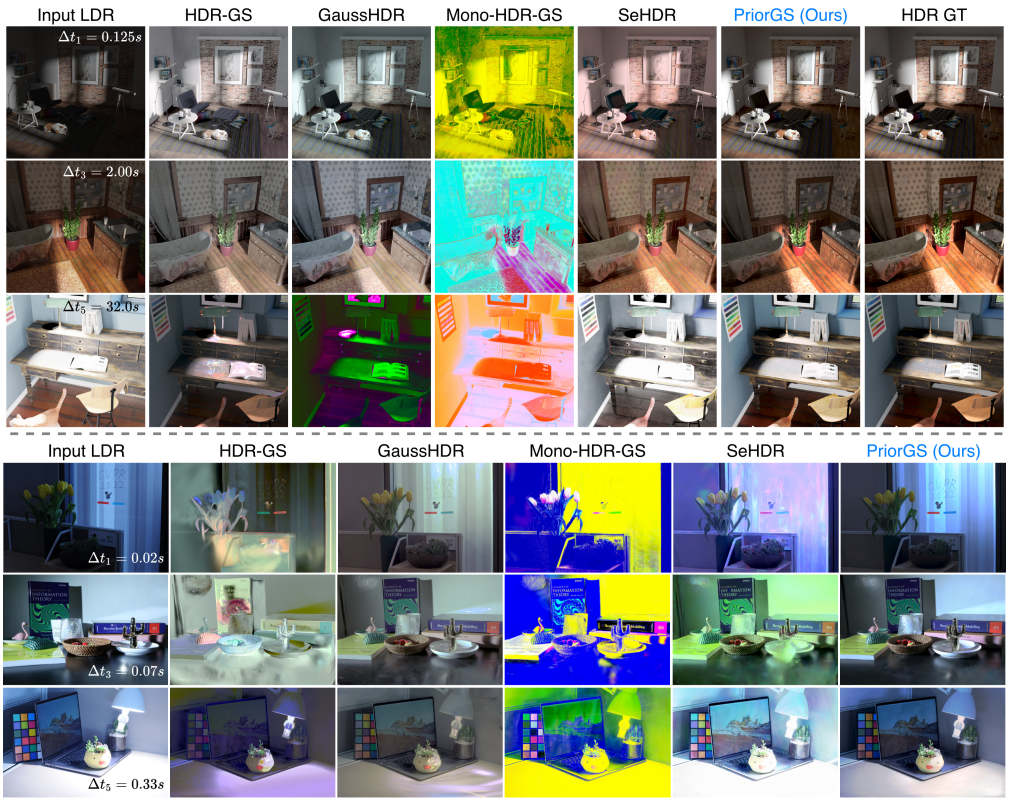

Figure 4: Qualitative HDR-NVS results given low- (Δt_1), medium- (Δt_3), and high-exposure (Δt_5) inputs across (*top half*) synthetic and (*bottom half*) real scenes. PriorGS substantially outperforms all baselines in high-fidelity HDR reconstruction across all scenes and exposure settings.

Qualitative. Figure 4 presents qualitative single-exposure HDR-NVS results from low-, medium-, and high-exposure inputs across real and synthetic scenes. Without learned LDR-to-HDR reconstruction priors, all baselines fail to recover scene radiance faithfully. However, PriorGS leverages this prior in a closed-loop reconstruction framework to produce substantially more accurate HDR reconstructions across all scenes and exposure settings. Moreover, in Figure 5, we present qualitative re-exposed LDR-NVS results across real and synthetic scenes under diverse input-output exposure setups. As observed, the flexible MLP-based CRFs of GaussHDR and Mono-HDR-GS generalize poorly to unseen exposures, producing overexposed or underexposed renderings and color artifacts. SeHDR mitigates this by heuristically fixing a CRF per scene but still suffers from inaccurate CRF estimation and limited scene adaptability. PriorGS instead regularizes CRF modeling with a learnable prior that guarantees monotonicity and promotes plausible responses, thereby achieving much more accurate re-exposure at unseen exposures and consistently outperforming all baselines. More qualitative results are offered in Appendix A.5 and in a short video in supplementary materials.

3D HDR Appearance Editing. Beyond performance improvement, our base-gain HDR representation enables continuous post-training control over HDR appearance in 3D without altering the LDR base scene. We introduce a coefficient $\rho \in \mathbb{R}_{\geq 0}$ that scales the 3D gain field: each Gaussian's log-gain SH coefficients $\boldsymbol{\beta}_i^g$ are scaled to $\rho\boldsymbol{\beta}_i^g$. This is equivalent to scaling evaluated log-gain $\boldsymbol{g}_i$ to $\rho\boldsymbol{g}_i$ in Equation 2 as SH evaluation is linear in SH coefficients. Varying ρ then controls HDR radiance enhancement relative to the LDR base: $\rho = 0$ removes the gain and recovers the LDR base, $0 < \rho < 1$ yields intermediate enhancement, $\rho = 1$ recovers the full HDR radiance, and $\rho > 1$ extrapolates the enhancement for creative HDR appearance editing. In Figure 6, we vary ρ from 0.0 to 1.5, report the corresponding dynamic range (DR), and show separate tone-mapped and low-exposure views of a novel HDR view. We observe that increasing ρ expands the DR and enhances the HDR appearance and chromatic contrast. Notably, $\rho = 0.5$ recovers highlight texture missing at $\rho = 0$. Furthermore, as shown in Figure 6, the edits are inherently view-consistent as ρ operates

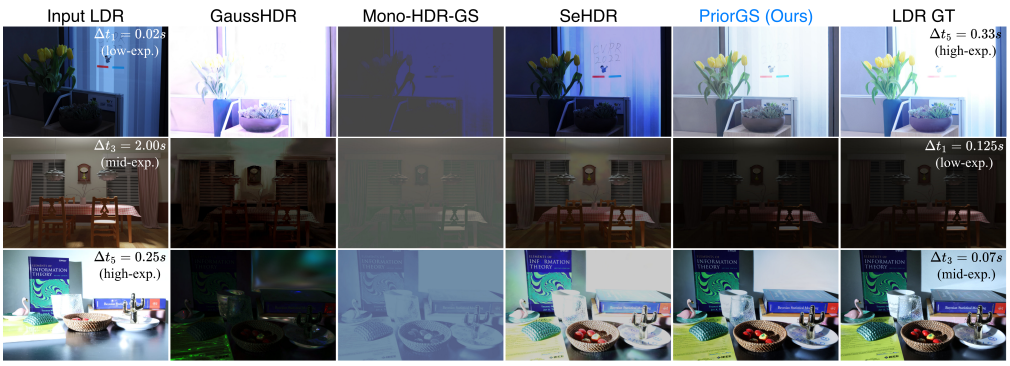

Figure 5: Qualitative re-exposed LDR-NVS results across scenes for three input → output exposure setups: (*top row*) $\Delta t_1 \rightarrow \Delta t_5$, (*middle row*) $\Delta t_3 \rightarrow \Delta t_1$, (*bottom row*) $\Delta t_5 \rightarrow \Delta t_3$. PriorGS consistently outperforms all baselines in novel-view LDR re-exposure at unseen exposures.

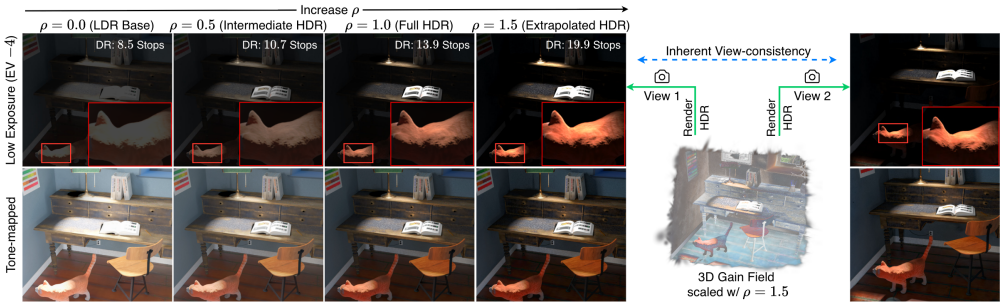

Figure 6: View-consistent 3D HDR appearance editing via gain field scaling with ρ. We sweep ρ and show an HDR view at a low-exposure (*top row*) and after a fixed tone-mapping (*bottom row*). Increasing ρ expands dynamic range (DR) and enhances HDR appearance and chromatic contrast.

directly on 3D Gaussians, offering a key advantage over independent per-view 2D editing. Note that, without a 3D gain field, the same editing capability cannot be achieved by scaling Gaussians' HDR radiance, as this merely amounts to an exposure shift and does not expand the DR (see Figure 12). Our base–gain representation thus offers a unique mechanism for 3D HDR appearance editing while preserving the LDR-base scene. Appendix A.6 offers extension to local edits and more scenes.

4.3 Ablation Studies

We ablate PriorGS and report qualitative results using all datasets and quantitative results on real HDR-NeRF and HDR-Plenoxels scenes, whose challenging scenes and diverse capture settings better reveal the contribution of components. Appendix A.7 offers further ablation results and details.

LDR-to-HDR Prior & Closed-Loop Reconstruction. In Table 2, we ablate (a) LDR-to-HDR reconstruction prior by removing pseudo-HDR supervision; and (b) closed-loop reconstruction by blocking gradients from the re-exposure loss to HDR representation, such that this loss optimizes only the CRF and does not impact HDR recovery. Both ablations degrade the performance, with removal of LDR-to-HDR prior causing a substantial drop. Together with qualitative results in Figures 7 and 15, these results prove the necessity of learned LDR-to-HDR priors for faithful HDR-NVS from single-exposure views and the contribution of closed-loop reconstruction to HDR-NVS quality.

Base-Gain HDR Representation. As shown in Table 2, we ablate our base-gain HDR representation using (c) direct HDR encoding with equal capacity (bigger SH degree); and MLP replacements for the gain field that predict Gaussians' HDR color from either (d) base color $\boldsymbol{c}$, position $\boldsymbol{\mu}$, and view direction $\boldsymbol{d}$ or (e) base color $\boldsymbol{c}$ alone. All ablations degrade the overall performance, with direct encoding causing the largest drop. Qualitatively, Figures 7 and 16 show that direct HDR encoding exhibits severe color and geometry artifacts, while Figures 7 and 17 reveal failure of MLPs in recon-

Table 2: Ablations on re-exposed LDR-NVS on real scenes of HDR-NeRF and HDR-Plenoxels, with best and second-best results denoted by red and yellow. $x \rightarrow y$ denotes replacing x with y.

(ID) Description	HDR-NeRF-Real – LDR			HDR-Plenoxels-Real – LDR		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
(a) Remove LDR-to-HDR prior	18.39	0.8503	0.099	15.00	0.7881	0.141
(b) Remove reconstruction loop	23.52	0.9180	0.073	22.97	0.9010	0.099
(c) Base-Gain $\rightarrow$ direct encoding	22.71	0.8981	0.105	15.05	0.6860	0.251
(d) Gain Field $\rightarrow$ MLP$(\boldsymbol{c}, \boldsymbol{\mu}, \boldsymbol{d})$	24.44	0.9238	0.066	22.64	0.8979	0.093
(e) Gain Field $\rightarrow$ MLP$(\boldsymbol{c})$	24.28	0.9268	0.066	22.07	0.8732	0.111
(f) CRF $f_{\boldsymbol{\theta}} \rightarrow$ MLP in 3D	23.26	0.8874	0.084	20.35	0.8296	0.147
(g) CRF $f_{\boldsymbol{\theta}} \rightarrow$ MLP in 2D	23.53	0.9071	0.075	22.34	0.8802	0.126
(h) CRF $f_{\boldsymbol{\theta}} \rightarrow$ fixed gamma	14.87	0.6784	0.222	12.57	0.5031	0.320
PriorGS (Ours)	24.31	0.9248	0.064	23.00	0.9021	0.092

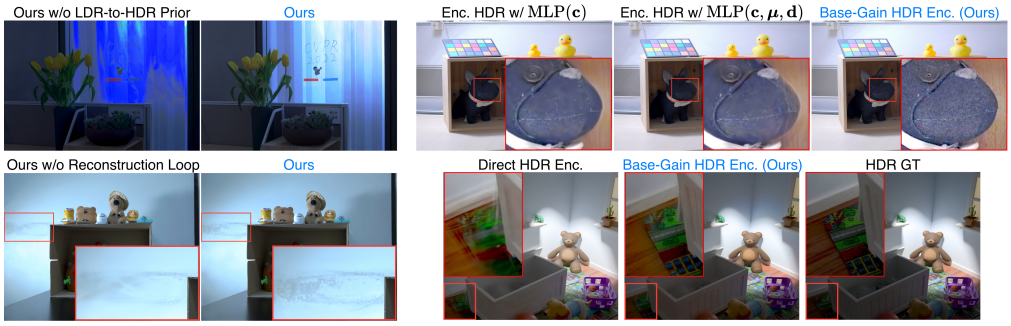

Figure 7: Ablations on HDR views for (*top left*) removing LDR-to-HDR prior; (*bottom left*) removing reconstruction loop; and replacing our base-gain HDR encoding by (*top right*) two MLP variants or (*bottom right*) direct HDR encoding. All ablations degrade HDR reconstruction quality.

structing high-frequency textures in challenging regions. Also, all configurations preclude 3D HDR editing enabled by the gain field. These results show the benefits of our base-gain representation.

Prior-Constrained CRF Modeling. In Table 2, we replace our CRF $f_{\boldsymbol{\theta}}$ with (f) an MLP applied in 3D on log-exposed per-Gaussian HDR colors; (g) an MLP applied in 2D on log-exposed HDR renderings; and (h) fixed gamma encoding $\gamma = 2.2$. Our prior-constrained CRF consistently outperforms flexible MLPs and the fixed CRF, validating the benefit of regularizing CRF with a learnable prior in single-exposure HDR-NVS. Appendix A.7 provides qualitative comparisons.

5 Discussion & Conclusion

We present a novel single-exposure HDR-NVS method based on 3DGS that leverages LDR-to-HDR reconstruction priors within a closed reconstruction loop, encodes HDR by a novel 3D base-gain representation, and regularizes CRF modeling with a learnable prior to achieve high-fidelity HDR-NVS and LDR re-exposure. We prove superiority of our method with extensive experiments and present its application in 3D HDR appearance editing. A few limitations present room for improvement. Our method is limited by the quality of pseudo-HDR targets generated by 2D HDR reconstruction methods: poor recovery of lost textures in the targets can degrade our HDR reconstruction. Also, while our base-gain representation improves robustness to cross-view inconsistencies in the targets, major discrepancies may still be encoded as fake view-dependent appearance, reducing scene realism. Future directions thus include (1) generative 3D HDR reconstruction that infers radiance directly in 3D using neighboring-view cues to improve reconstruction and cross-view consistency; and (2) mitigating inconsistency issue by uncertainty learning (Martin-Brualla et al., 2021).

AI USE STATEMENT

In this work, we used generative AI tools to implement methods, and assist in the writing of proofs. We have not used generative AI tools to generate synthetic data sets, help develop theoretical models or conceptual frameworks, formulate mathematical claims, provide critical ingredients for proving mathematical claims, propose or refine hypotheses, design or provide feedback on research methodology or experiments, assist with translation, clean and reformat dataset, support qualitative and thematic data analysis, interpret results, draft parts of the paper, and create or modify scientific figures or images. Additionally, we used generative AI tools for polishing the writing and enhancing readability, sourcing/searching for information, and brainstorming. We have reviewed all AI-assisted work in the following forms: (1) LLM-generated code was verified and tested for correctness by at least one author; (2) All LLM-generated text used for polishing the writing of the paper and the monotonicity proof was reviewed by at least two authors and manually refined repeatedly to ensure the intended message is delivered; and (3) All LLM-driven search results were manually reviewed to ensure factual accuracy. We take responsibility for the final content of this work, including text, claims or artifacts produced with the aid of generative AI.

REPRODUCIBILITY STATEMENT

Sections 3 and 4.1, together with Appendix A.3, provide the methodological and implementation details needed to reproduce our results. Our implementation builds on the publicly available HDR-GS implementation (Cai et al., 2024). We use the publicly available HDR-NeRF (Huang et al., 2022) and HDR-Plenoxels (Jun-Seong et al., 2022) datasets for training and evaluation. Pseudo-HDR targets are generated from their input LDR views using the publicly available 2D HDR reconstruction method IntrinsicHDR (Dille et al., 2025). Our implementation will be made publicly available on GitHub upon publication. All experiments can be reproduced using a single NVIDIA RTX 3090 GPU with 24GB of memory.

REFERENCES

Android Developers. Ultra HDR image format v1.1. https://developer.android.com/media/platform/hdr-image-format, 2024. Accessed: August 3, 2026.

Blender Foundation. Blender. https://www.blender.org/, 2026.

Christophe Bolduc, Yannick Hold-Geoffroy, and Jean-François Lalonde. GaSLight: Gaussian splats for spatially-varying lighting in HDR. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 29120–29130, 2025. doi: 10.1109/ICCV51701.2025.02704.

Yuanhao Cai, Zihao Xiao, Yixun Liang, Minghan Qin, Yulun Zhang, Xiaokang Yang, Yaoyao Liu, and Alan Yuille. HDR-GS: Efficient high dynamic range novel view synthesis at 1000x speed via gaussian splatting. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 68453–68471. Curran Associates, Inc., 2024. doi: 10.52202/079017-2186. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/7e83fd2ff7ae58485d418685521c9608-Paper-Conference.pdf.

Trevor D. Canham, SaiKiran Tedla, Michael J. Murdoch, and Michael S. Brown. Gain-MLP: Improving HDR gain map encoding via a lightweight MLP. In *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 18619–18628, 2025. doi: 10.1109/ICCV51701.2025.01730.

Peibei Cao, Rafal K. Mantiuk, and Kede Ma. Perceptual assessment and optimization of HDR image rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22433–22443, June 2024.

Zehao Chen, Binbin Zhou, and Zengwei Zheng. Hdr scene reconstruction from resolution-mismatched event streams and single-exposure images. *Sensors*, 26(14), 2026. ISSN 1424-8220. doi: 10.3390/s26144621. URL https://www.mdpi.com/1424-8220/26/14/4621.

Paul Debevec. Rendering synthetic objects into real scenes: bridging traditional and image-based graphics with global illumination and high dynamic range photography. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH '98, pp. 189–198, New York, NY, USA, 1998. Association for Computing Machinery. ISBN 0897919998. doi: 10.1145/280814.280864. URL https://doi.org/10.1145/280814.280864.

Sebastian Dille, Chris Careaga, and Yağız Aksoy. Intrinsic single-image HDR reconstruction. In Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol (eds.), *Computer Vision – ECCV 2024*, pp. 161–177, Cham, 2025. Springer Nature Switzerland. ISBN 978-3-031-73247-8.

Gabriel Eilertsen, Joel Kronander, Gyorgy Denes, Rafał K. Mantiuk, and Jonas Unger. HDR image reconstruction from a single exposure using deep cnns. *ACM Trans. Graph.*, 36(6), November 2017. ISSN 0730-0301. doi: 10.1145/3130800.3130816. URL https://doi.org/10.1145/3130800.3130816.

Gabriel Eilertsen, Saghi Hajisharif, Param Hanji, Apostolia Tsirikoglou, Rafał K. Mantiuk, and Jonas Unger. How to cheat with metrics in single-image HDR reconstruction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) Workshops*, pp. 3998–4007, October 2021.

Pulkit Gera, Mohammad Reza Karimi Dastjerdi, Charles Renaud, PJ Narayanan, and Jean-François Lalonde. Casual indoor HDR radiance capture from omnidirectional images. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA, 2022. URL https://pdf-files.bmvc2022.org/0305.pdf.

Ruben Gomez-Ojeda, Zichao Zhang, Javier Gonzalez-Jimenez, and Davide Scaramuzza. Learning-based image enhancement for visual odometry in challenging HDR environments. In *2018 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 805–811, 2018. doi: 10.1109/ICRA.2018.8462876.

M.D. Grossberg and S.K. Nayar. Modeling the space of camera response functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(10):1272–1282, 2004. doi: 10.1109/TPAMI.2004.88.

Yuanshen Guan, Ruikang Xu, Yinuo Liao, Mingde Yao, Lizhi Wang, and Zhiwei Xiong. HDR image generation via gain map decomposed diffusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 17536–17545, October 2025.

Param Hanji, Rafal Mantiuk, Gabriel Eilertsen, Saghi Hajisharif, and Jonas Unger. Comparison of single image HDR reconstruction methods — the caveats of quality assessment. In *ACM SIGGRAPH 2022 Conference Proceedings*, SIGGRAPH '22, New York, NY, USA, 2022. Association for Computing Machinery. ISBN 9781450393379. doi: 10.1145/3528233.3530729. URL https://doi.org/10.1145/3528233.3530729.

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 (CVPR)*, pp. 18398–18408, June 2022.

Xin Huang, Qi Zhang, Ying Feng, Hongdong Li, and Qing Wang. LTM-NeRF: Embedding 3d local tone mapping in HDR neural radiance field. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):10944–10959, 2024. doi: 10.1109/TPAMI.2024.3448620.

ISO/TC 42. Digital photography—gain map metadata for image conversion. ISO 21496-1:2025, International Organization for Standardization, July 2025. URL https://www.iso.org/standard/86775.html.

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 A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 80191–80219. Curran Associates, Inc., 2024. doi: 10.52202/079017-2549. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/92dd1adab39f362046f99dfe3c39d90f-Paper-Conference.pdf.

Kim Jun-Seong, Kim Yu-Ji, Moon Ye-Bin, and Tae-Hyun Oh. HDR-Plenoxels: Self-calibrating high dynamic range radiance fields. In Shai Avidan, Gabriel Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (eds.), *Computer Vision – ECCV 2022*, pp. 384–401, Cham, 2022. Springer Nature Switzerland. ISBN 978-3-031-19824-3.

Nima Khademi Kalantari and Ravi Ramamoorthi. Deep high dynamic range imaging of dynamic scenes. *ACM Trans. Graph.*, 36(4), July 2017. ISSN 0730-0301. doi: 10.1145/3072959.3073609. URL https://doi.org/10.1145/3072959.3073609.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkuehler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4), July 2023. ISSN 0730-0301. doi: 10.1145/3592433. URL https://doi.org/10.1145/3592433.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL http://arxiv.org/abs/1412.6980.

Yiyu Li, Haoyuan Wang, Ke Xu, Gerhard Petrus Hancke, and Rynson W.H. Lau. SEHDR: Single-exposure HDR novel view synthesis via 3d gaussian bracketing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 26045–26054, October 2025.

Zhihao Li, Yufei Wang, Alex Kot, and Bihan Wen. From chaos to clarity: 3DGS in the dark. In A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang (eds.), *Advances in Neural Information Processing Systems*, volume 37, pp. 94971–94992. Curran Associates, Inc., 2024. doi: 10.52202/079017-3010. URL https://proceedings.neurips.cc/paper_files/paper/2024/file/aca97732e30bcf1303bc22ac3924fd16-Paper-Conference.pdf.

Yinuo Liao, Yuanshen Guan, Ruikang Xu, Jiacheng Li, Shida Sun, and Zhiwei Xiong. Learning gain map for inverse tone mapping. In Y. Yue, A. Garg, N. Peng, F. Sha, and R. Yu (eds.), *International Conference on Learning Representations*, volume 2025, pp. 27246–27263, 2025. URL https://proceedings.iclr.cc/paper_files/paper/2025/file/43d2b7fbee8431f7cef0d0afed51c691-Paper-Conference.pdf.

Aoyu Liu, Zhen Liu, Ziyi Wang, Dian Chen, Bing Zeng, and Shuaicheng Liu. ExpoCM: Exposure-aware one-step generative single-image HDR reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 29909–29918, June 2026.

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 (CVPR)*, pp. 5991–6000, June 2025.

Yu-Lun Liu, Wei-Sheng Lai, Yu-Sheng Chen, Yi-Lung Kao, Ming-Hsuan Yang, Yung-Yu Chuang, and Jia-Bin Huang. Single-image HDR reconstruction by learning to reverse the camera pipeline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2020.

Tao Lu, Mulin Yu, Linning Xu, Yuanbo Xiangli, Limin Wang, Dahua Lin, and Bo Dai. Scaffold-GS: Structured 3d gaussians for view-adaptive rendering. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20654–20664, 2024. doi: 10.1109/CVPR52733.2024.01952.

D. Marnerides, T. Bashford-Rogers, J. Hatchett, and K. Debattista. ExpandNet: A deep convolutional neural network for high dynamic range expansion from low dynamic range content. *Computer Graphics Forum*, 37(2):37–49, 2018. doi: https://doi.org/10.1111/cgf.13340. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13340.

Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth. Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 7206–7215, 2021. doi: 10.1109/CVPR46437.2021.00713.

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. *Commun. ACM*, 65(1):99–106, December 2021. ISSN 0001-0782. doi: 10.1145/3503250. URL https://doi.org/10.1145/3503250.

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 *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16169–16178, 2022. doi: 10.1109/CVPR52688.2022.01571.

Michael Niemeyer, Fabian Manhardt, Marie-Julie Rakotosaona, Michael Oechsle, Christina Tsalicoglou, Keisuke Tateno, Jonathan T. Barron, and Federico Tombari. Learning neural exposure fields for view synthesis. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*, 2025. URL https://openreview.net/forum?id=N6ZzwaYSxQ.

Yunshan Qi, Lin Zhu, Nan Bao, Yifan Zhao, and Jia Li. Seeing through light and darkness: Sensor-physics grounded deblurring hdr nerf from single-exposure images and events. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 22723–22732, June 2026.

Marcel Santana Santos, Tsang Ing Ren, and Nima Khademi Kalantari. Single image HDR reconstruction using a cnn with masked features and perceptual loss. *ACM Trans. Graph.*, 39(4), August 2020. ISSN 0730-0301. doi: 10.1145/3386569.3392403. URL https://doi.org/10.1145/3386569.3392403.

Johannes L. Schönberger and Jan-Michael Frahm. Structure-from-motion revisited. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4104–4113, 2016. doi: 10.1109/CVPR.2016.445.

Shreyas Singh, Aryan Garg, and Kaushik Mitra. HDRSplat: Gaussian splatting for high dynamic range 3d scene reconstruction from raw images. In *35th British Machine Vision Conference 2024, BMVC 2024, Glasgow, UK, November 25-28, 2024*. BMVA, 2024. URL https://papers.bmvc2024.org/0022.pdf.

Chaoda Song, Yiren Lu, Xinpeng Li, Yunlai Zhou, Yanyan Zhang, Yu Yin, and Vipin Chaudhary. Expo-GS: Exposure-aware signed distance function in gaussian splatting for high dynamic range. In *Forty-third International Conference on Machine Learning*, 2026. URL https://openreview.net/forum?id=2itDye7kPO.

Chao Wang, Ana Serrano, Xingang Pan, Bin Chen, Karol Myszkowski, Hans-Peter Seidel, Christian Theobalt, and Thomas Leimkühler. GlowGAN: Unsupervised learning of HDR images from LDR images in the wild. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 10475–10485, 2023. doi: 10.1109/ICCV51070.2023.00964.

Chao Wang, Zhihao Xia, Thomas Leimkuhler, Karol Myszkowski, and Xuaner Zhang. LEDiff: Latent exposure diffusion for HDR generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 453–464, June 2025.

Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004. doi: 10.1109/TIP.2003.819861.

Ronghuan Wu, Wanchao Su, Kede Ma, Jing Liao, and Rafał K Mantiuk. X2HDR: HDR image generation in a perceptually uniform space. *arXiv preprint arXiv:2602.04814*, 2026.

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.

Huimin Zeng, Yue Bai, Hailing Wang, and Yun Fu. Physically inspired gaussian splatting for HDR novel view synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11808–11817, June 2026.

Kaixuan Zhang, Hu Wang, Minxian Li, Mingwu Ren, Mao Ye, and Xiatian Zhu. High dynamic range novel view synthesis with single exposure. In Aarti Singh, Maryam Fazel, Daniel Hsu, Simon Lacoste-Julien, Felix Berkenkamp, Tegan Maharaj, Kiri Wagstaff, and Jerry Zhu (eds.), *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pp. 75120–75131. PMLR, 13–19 Jul 2025. URL https://proceedings.mlr.press/v267/zhang25af.html.

Ning Zhang, Yuyao Ye, Yang Zhao, and Ronggang Wang. Revisiting the stack-based inverse tone mapping. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 9162–9171, June 2023.

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/CVF Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018. doi: 10.1109/CVPR.2018.00068.

Xiuming Zhang, Pratul P. Srinivasan, Boyang Deng, Paul Debevec, William T. Freeman, and Jonathan T. Barron. NeRFactor: neural factorization of shape and reflectance under an unknown illumination. *ACM Trans. Graph.*, 40(6), December 2021. ISSN 0730-0301. doi: 10.1145/3478513.3480496. URL https://doi.org/10.1145/3478513.3480496.

A Appendix

A.1 Discussion on Base-Gain HDR Representation

Our base-gain HDR representation anchors the base appearance and shared geometry to reliable LDR supervision and directs cross-view inconsistencies in pseudo-HDR targets toward the gain field. This provides a soft inductive bias that discourages, but does not strictly prevent, HDR supervision from updating the base appearance and geometry. We made this design choice after evaluating the following hard constraints:

- **Preventing HDR supervision from updating geometry.** In regions where LDR views are degraded by shadow noise or highlight clipping, pseudo-HDR targets may contain recovered textures. 3D reconstruction of these textures may require geometric refinement beyond what LDR supervision alone can provide. We therefore allow HDR supervision to update the shared geometry, a choice motivated by the above rationale and our empirical observations.
- **Preventing HDR supervision from updating base appearance.** We evaluated this constraint by implementing stop gradient on linearized LDR-base color $\boldsymbol{l}_i^v$ in Equation 2 and achieved no improvement in HDR reconstruction quality. We therefore favor the simpler design, avoiding an additional constraint that provides no observed benefit.

A.2 Representational Capacity Analysis & Monotonicity Proof of Our CRF

A.2.1 Representational Capacity Analysis

We evaluate the representational capacity of our CRF model (see Section 3.2) on all 201 camera response curves in DoRF (Grossberg & Nayar, 2004). Specifically, we follow DoRF's range-normalization convention to fit our CRF's normalized response to each DoRF curve by minimizing the mean squared error (MSE). Figure 8 compares the DoRF response curves with the resulting fitted responses. Qualitatively, the fitted responses closely reproduce the broad diversity of DoRF curves, yielding an overall response family whose distribution and envelope are nearly indistinguishable from those of the DoRF curves. Quantitatively, across all curves, our CRF model achieves a mean MSE of 2.081×10^{-5}, a median MSE of 4.328×10^{-6}, and a maximum MSE of 8.225×10^{-4}. For comparison, fitting the direct two-parameter response $t_{\boldsymbol{\phi}}(x)$ (see Equation 4) and a gamma response x^{γ} yields mean MSEs of 5.717×10^{-5} and 3.801×10^{-3}, respectively. Thus, our model reduces the mean MSE by 63.6% relative to the two-parameter response and by 99.45% relative to the gamma model. This demonstrates that our model represents empirical camera responses substantially more accurately than the baselines.

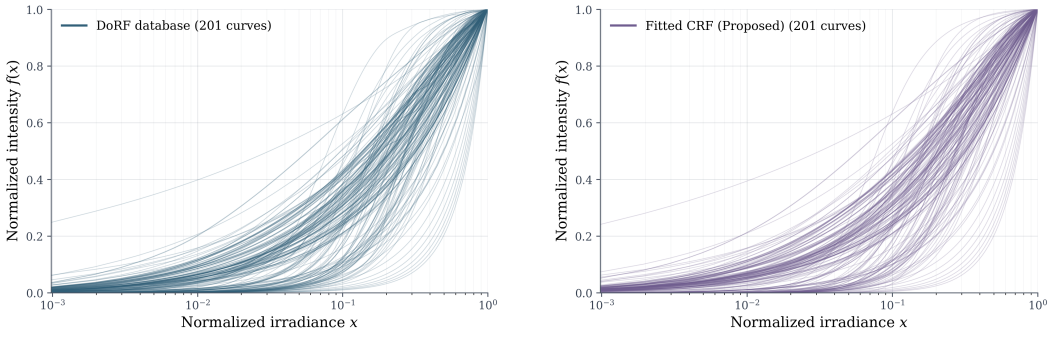

(a) DoRF responses

(b) Best-fitting responses of our CRF

Figure 8: Representational capacity of our proposed CRF model on the DoRF database. (a) the 201 DoRF response curves. (b) the independently fitted responses of our CRF. All curves are displayed using a logarithmic irradiance axis; the zero-irradiance samples are included during fitting but omitted from the visualization.

A.2.2 Monotonicity Proof

As characterized by Grossberg & Nayar (2004), a physically valid CRF should preserve the ordering of input irradiance: a greater irradiance must produce a bigger recorded pixel intensity. We show below that our proposed CRF satisfies this: It is strictly increasing over its domain by construction. Our proposed CRF is defined as

$$\begin{aligned} f(E) &= t\left(\eta(E)\right), \\ t(x) = (1+\sigma)\frac{x^n}{x^n+\sigma}, &\qquad \eta(E) = \mathrm{sigmoid}(a\ln(E+\epsilon)+b)\,. \end{aligned} \tag{9}$$

where $E \in \mathbb{R}_{\geq 0}$, $\sigma, n, a \in \mathbb{R}_{>0}$, $b \in \mathbb{R}$, and $\epsilon = 1 \times 10^{-8}$.

Lemma 1 (Composition of strictly increasing functions)**.** *Let $u\colon I \to J$ and $v\colon J \to \mathbb{R}$ be strictly increasing. Then $v \circ u$ is strictly increasing.*

Proof. Let $x_1, x_2 \in I$ with $x_1 < x_2$. Since u is strictly increasing,

$$u(x_1) < u(x_2).$$

Because $u(x_1), u(x_2) \in J$ and v is strictly increasing, it follows that

$$v(u(x_1)) < v(u(x_2)).$$

Thus,

$$(v \circ u)(x_1) < (v \circ u)(x_2),$$

so $v \circ u$ is strictly increasing on I. □

Proposition 1. *Under the parameter constraints above, f is strictly increasing on $\mathbb{R}_{\geq 0}$.*

Proof. Define $h(E) = a\ln(E+\epsilon) + b$. Since $E + \epsilon > 0$ and $a > 0$,

$$h'(E) = \frac{a}{E+\epsilon} > 0, \tag{10}$$

so h is strictly increasing on $\mathbb{R}_{\geq 0}$. Now, since sigmoid has range $(0, 1)$,

$$\mathrm{sigmoid}'(t) = \mathrm{sigmoid}(t)\big(1 - \mathrm{sigmoid}(t)\big) > 0. \tag{11}$$

Therefore, sigmoid is also strictly increasing, and thus $\eta = \mathrm{sigmoid} \circ h$ is strictly increasing due to Lemma 1 and has range $(0, 1)$. Finally, for every $x > 0$,

$$t'(x) = \frac{(1+\sigma)n\sigma x^{n-1}}{(x^n+\sigma)^2} > 0, \tag{12}$$

since $\sigma, n > 0$. Hence t is strictly increasing on $(0, \infty)$. Since $\eta(E) \in (0, 1)$, $f = t \circ \eta$ is a composition of strictly increasing functions. Lemma 1 thus shows that f is strictly increasing on $\mathbb{R}_{\geq 0}$. □

A.3 Experimental Setup Details

Details of Datasets. The real and synthetic HDR-NeRF (Huang et al., 2022) views have native resolutions of 3216×2136 and 800×800, respectively. We train and evaluate in the native resolution for synthetic scenes and downsample real scenes to 1608×1068 ($\frac{1}{2}$ scale) to reduce computational cost. Similarly, we downsample the HDR-Plenoxels (Jun-Seong et al., 2022) views from their native resolution of 5952×4480 to $\frac{1}{4}$ scale, resulting in a resolution of 1488×1120 for training and evaluation. In the HDR-Plenoxels dataset, the *plant* and *desk* scenes are available at all 5 exposures, $\{\Delta t_1, \Delta t_2, \Delta t_3, \Delta t_4, \Delta t_5\}$, whereas *coffee* and *character* scenes provide only 3 exposures $\{\Delta t_1, \Delta t_3, \Delta t_5\}$. Accordingly, we conduct five single-exposure training runs for *plant* and *desk* but only three for *coffee* and *character*. Reported results are averaged across runs for all training exposures and all scenes of a dataset.

Further Implementation Details. The Adam optimizer (Kingma & Ba, 2015) is used for training. We adopt the learning-rate schedules and densification/pruning hyperparameters of Kerbl et al. (2023) for training the Gaussian scene representations and use the same configuration across all baselines. We initialize our CRF model $f_{\boldsymbol{\theta}}$ to the mean DoRF response (Grossberg & Nayar, 2004; Eilertsen et al., 2017) and optimize its parameters with an initial learning rate of 5×10^{-4} that exponentially decays to 5×10^{-5}. To ensure that CRF parameters σ, n, and a remain strictly positive throughout training, we use exponential reparameterizations: $\sigma = \exp(\tilde{\sigma})$, $n = \exp(\tilde{n})$, $a = \exp(\tilde{a})$ and optimize the unconstrained log-parameters $\tilde{\sigma}, \tilde{n}, \tilde{a} \in \mathbb{R}$ jointly with b. We note that our CRF model $f_{\boldsymbol{\theta}}$ learns a single response curve shared across all three RGB channels. To improve optimization stability, we normalize pseudo-HDR targets to $[0, 1]$ by dividing them by the maximum radiance across all pseudo-HDR targets in the scene.

Further Details of Metrics & Baselines. As discussed in Section 4.1, in HDR quantitative comparisons, we follow the recommendation of Hanji et al. (2022) and apply luminance correction to adjust the scale ambiguity inherent in single-exposure HDR reconstruction (Cao et al., 2024) and enhance judgment of HDR metrics. To adapt their proposed correction to 3D settings, we compute luminance correction weights once per scene using its first indexed viewpoint and then reuse the weights across all evaluation viewpoints. We note that unlike per-view correction, this per-scene correction avoids masking of intensity flickering among views during evaluation. After luminance correction, we apply μ-law range compression ($\mu = 5000$) and compute HDR metrics in this tone-mapped domain, following prior works (Cai et al., 2024; Liu et al., 2025). For comparisons with multi-exposure baselines, we adapt the publicly available implementations of GaussHDR (Liu et al., 2025) and HDR-GS (Cai et al., 2024) to run on single-exposure inputs. The recent 3DGS-based multi-exposure methods PhysHDR-GS (Zeng et al., 2026) and Expo-GS (Song et al., 2026) are not included only because their implementations are not publicly available. All baselines, whether single-exposure or multi-exposure, were retrained given our HDR GT-free single-exposure training protocol (as explained in Section 4.1) to ensure fairness in comparisons.

A.4 Further Quantitative Results

PriorGS's Robustness across 2D Single-exposure HDR Reconstruction Methods. To evaluate PriorGS's robustness to the choice of 2D single-exposure HDR reconstruction method, we replace our primary 2D HDR reconstruction method, IntrinsicHDR (Dille et al., 2025), with three alternatives: DeepHDR (Santos et al., 2020), ExpandNet (Marnerides et al., 2018), and GMNet (Liao et al., 2025). Table 3 summarizes the results for these alternative configurations and compares them with our closest baselines, GaussHDR (Liu et al., 2025) and SeHDR (Li et al., 2025). We make the following observations: (1) Regardless of the 2D HDR reconstruction method integrated into PriorGS, every resulting configuration consistently outperforms both baselines across both datasets, including real and synthetic scenes, and on both HDR-NVS and re-exposed LDR-NVS, with a single exception: the GMNet configuration trails GaussHDR in PSNR on HDR-Plenoxels dataset while surpassing it by large margins in SSIM and LPIPS. These results further validate the effectiveness of PriorGS in single-exposure HDR-NVS and LDR re-exposure; (2) On synthetic scenes, the primary and all alternative 2D HDR reconstruction methods yield comparable PriorGS performance, demonstrating robustness to method choice. However, on the more challenging real scenes, which span diverse cameras and exposure settings, PriorGS exhibits more sensitivity to the selected reconstruction method; and (3) Our primary 2D HDR reconstruction method, IntrinsicHDR, yields the

most robust performance among all evaluated methods: it outperforms all alternatives on the more challenging real scenes while achieving comparable results on synthetic scenes.

Table 3: PriorGS's performance after integration with different 2D HDR reconstruction methods on (*top half*) re-exposed LDR-NVS on real scenes of HDR-NeRF and HDR-Plenoxels datasets and (*bottom half*) HDR-NVS and re-exposed LDR-NVS on synthetic scenes of HDR-NeRF dataset. red and yellow denote the best and second-best results, respectively.

Method	HDR-NeRF-Real – LDR			HDR-Plenoxels-Real – LDR		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
GaussHDR	16.28	0.7849	0.176	17.73	0.7141	0.197
SeHDR	13.75	0.6991	0.182	14.21	0.6644	0.232
PriorGS + DeepHDR	20.38	0.8604	0.085	18.43	0.8111	0.115
PriorGS + ExpandNet	21.18	0.8771	0.094	21.23	0.8624	0.100
PriorGS + GMNet	19.36	0.8482	0.095	16.94	0.7887	0.126
PriorGS + IntrinsicHDR (Primary)	24.31	0.9248	0.064	23.00	0.9021	0.092
Method	HDR-NeRF-Syn – HDR			HDR-NeRF-Syn – LDR		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
GaussHDR	26.96	0.8828	0.114	11.99	0.5848	0.302
SeHDR	15.14	0.6405	0.262	15.76	0.7450	0.169
PriorGS + DeepHDR	30.81	0.9512	0.052	23.76	0.8818	0.096
PriorGS + ExpandNet	30.35	0.9437	0.039	24.06	0.9125	0.067
PriorGS + GMNet	31.77	0.9540	0.043	23.64	0.8838	0.083
PriorGS + IntrinsicHDR (Primary)	31.22	0.9569	0.038	22.96	0.8937	0.081

LDR-NVS Performance at the Observed Exposure. In addition to our primary evaluation of re-exposed LDR-NVS at unseen exposures in Table 1, we evaluate LDR-NVS quality at the observed (training) exposure. Our base-gain HDR representation enables joint reconstruction of an LDR-base at the input exposure and an HDR scene. This, therefore, allows PriorGS to render a novel LDR view at the training exposure through two pathways: directly from the LDR-base or by re-exposing the HDR rendering through the learned CRF. We opt to use the direct LDR-base rendering as our primary output when LDR view at training exposure is requested. The results averaged across all five single-exposure training runs for real and synthetic scenes are summarized in Tables 4 and 5, respectively. PriorGS outperforms all baselines on the real HDR-NeRF and HDR-Plenoxels scenes and performs comparably to Mono-HDR-GS (Zhang et al., 2025) on the synthetic HDR-NeRF scenes. Together with Table 1, these results show that PriorGS not only generalizes substantially better to unseen exposures but also achieves excellent LDR-NVS performance at the observed exposure. This highlights the superiority of PriorGS and the benefit of our base-gain HDR representation in providing a clean LDR base scene while simultaneously maintaining a high-quality HDR scene representation.

Table 4: Quantitative results on LDR-NVS at the training exposure on real scenes of HDR-NeRF and HDR-Plenoxels datasets. red and yellow denote the best and second-best results, respectively.

Method	HDR-NeRF-Real – LDR			HDR-Plenoxels-Real – LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
HDR-GS	35.24	0.9667	0.0198	30.43	0.9428	0.0492
GaussHDR	34.42	0.9662	0.0366	30.54	0.9447	0.0500
Mono-HDR-GS	35.65	0.9656	0.0172	31.02	0.9465	0.0464
SeHDR	30.08	0.9632	0.0346	21.77	0.9037	0.1157
PriorGS (ours)	36.10	0.9700	0.0158	31.71	0.9482	0.0439

Table 5: Quantitative results on LDR-NVS at the training exposure on synthetic scenes of HDR-NeRF dataset. red and yellow denote the best and second-best results, respectively.

Method	HDR-NeRF-Syn – LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓
HDR-GS	41.56	0.9825	0.0088
GaussHDR	40.90	0.9795	0.0154
Mono-HDR-GS	41.80	0.9828	0.0083
SeHDR	32.53	0.9768	0.0199
PriorGS (ours)	41.77	0.9836	0.0087

Comparison Against Scaffold-GS variant of GaussHDR. Beyond the baselines introduced in Section 4.1, Table 6 compares PriorGS with Scaffold-GS variant of GaussHDR (Liu et al., 2025). We exclude this variant from the main comparison as all other methods, including ours, use standard 3DGS representation (Kerbl et al., 2023). Despite GaussHDR's superior Scaffold-GS representation, PriorGS consistently outperforms it on both datasets, across real and synthetic scenes, and in both HDR-NVS and re-exposed LDR-NVS, highlighting the effectiveness of our proposed method.

Table 6: Quantitative comparison of PriorGS vs. the 3DGS (Kerbl et al., 2023) and Scaffold-GS† (Lu et al., 2024) variants of GaussHDR (Liu et al., 2025) for (*top half*) re-exposed LDR-NVS on real scenes of HDR-NeRF and HDR-Plenoxels datasets and (*bottom half*) HDR-NVS and re-exposed LDR-NVS on synthetic scenes of HDR-NeRF dataset. red and yellow denote the best and second-best results, respectively.

Method	HDR-NeRF-Real – LDR			HDR-Plenoxels-Real – LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
GaussHDR†	22.76	0.8739	0.128	21.03	0.8635	0.115
GaussHDR	16.28	0.7849	0.176	17.73	0.7141	0.197
PriorGS (Ours)	24.31	0.9248	0.064	23.00	0.9021	0.092
Method	HDR-NeRF-Syn – HDR			HDR-NeRF-Syn – LDR		
	PSNR ↑	SSIM ↑	LPIPS ↓	PSNR ↑	SSIM ↑	LPIPS ↓
GaussHDR†	28.45	0.8970	0.087	12.56	0.6293	0.295
GaussHDR	26.96	0.8828	0.114	11.99	0.5848	0.302
PriorGS (Ours)	31.22	0.9569	0.038	22.96	0.8937	0.081

A.5 Further Qualitative Results

Additional HDR-NVS Results. Complementing Figure 4, Figure 9 presents additional qualitative HDR-NVS results across two scenes and three different input exposures. We observe that GaussHDR and SeHDR recover the HDR radiance inaccurately and inconsistently across different input exposures, displaying severe color and geometry artifacts particularly in under-exposed and over-exposed regions. In contrast, PriorGS behaves much more consistently across input exposure levels and substantially outperforms the baselines in faithful HDR radiance recovery, consistent with our observations in Figure 4. These results further demonstrate the effectiveness of LDR-to-HDR prior supervision within our closed-loop reconstruction framework.

Additional Re-exposed LDR-NVS Results. We complement Figure 5 by presenting additional qualitative re-exposed LDR-NVS comparisons across two scenes in Figure 10. Specifically, given single-exposure LDR inputs, we compare the methods' novel-view LDR re-exposures at all five exposure levels, including the training (input) exposure. We observe that HDR-GS, GaussHDR, and Mono-HDR-GS reproduce novel-view LDR images at the observed training exposure with high fidelity, showing that their flexible MLP-based CRFs can closely fit the supervised exposure. How-



Figure 9: Further qualitative HDR-NVS results given low- (Δt_1), medium- (Δt_3), and high-exposure (Δt_5) inputs across (*top half*) synthetic and (*bottom half*) real scenes.

ever, these CRFs generalize poorly to unseen exposures, resulting in overexposed or underexposed renderings and color artifacts. This sharp gap between observed- and unseen-exposure performance indicates overfitting to the training exposure, consistent with our earlier discussions in Section 1 and observations in Figure 5. SeHDR mitigates this overfitting and the resulting exposure errors and color artifacts by heuristically fixing a CRF for each scene. However, it still suffers from limited scene adaptability and exhibits a narrow response range due to inaccurate CRF preselection. In contrast, PriorGS produces substantially more accurate novel-view LDR re-exposures at unseen exposure levels while maintaining high fidelity at the observed training exposure, consistently outperforming all baselines. This behavior aligns with our earlier observations in Figure 5 and highlights the intended balance of our CRF formulation: the prior regularizes its behavior beyond the supervised exposure, while its learnable parameters exhibit excellent scene adaptability. Together, these results validate the need for a learnable CRF prior in single-exposure HDR-NVS and demonstrate the effectiveness of our proposed model.



Figure 10: Further qualitative re-exposed LDR-NVS results. Using single-exposure LDR views for training, we render a novel view at all five exposures, including the training exposure. The training exposure is (*top half*) Δt_3 for a synthetic scene and (*bottom half*) Δt_2 for a real scene.

Qualitative Analysis of Learned CRF. In Figure 11, we compare our learned CRF curve with the MLP-based CRF of GaussHDR (Liu et al., 2025) for the synthetic *bear* scene. Both methods are trained under the same single-exposure setting, using LDR views at exposure $\Delta t_1 = 0.125s$. We additionally present their novel-view LDR re-exposures at the training exposure $\Delta t_1 = 0.125s$, and the test exposure, $\Delta t_3 = 2.0s$. Both methods faithfully reproduce the novel-view LDR re-exposure at the training exposure, where the target intensities fall within the output range supervised by training observations. Beyond this range, however, the flexible MLP-based CRF curve exhibits pronounced divergence between its RGB responses, early saturation of the blue channel, and local violations of monotonicity. Without explicit priors to regularize the response inside this unsupervised region, its flexibility permits implausible response learning. Since rendering at the test exposure involves LDR intensities beyond the supervised region, MLP-based re-exposure leads to substantial brightness and color distortions with a blue tint. Our proposed CRF, however, leverages its DoRF-derived prior

to maintain monotonicity and a more plausible response beyond the supervised range, producing substantially more faithful LDR renderings at the test exposure. We note that the differences between the learned CRFs within the supervised range is due to different HDR reconstructions by the two methods, as each CRF maps its respective reconstruction to the target exposures. We mark the approximate supervision boundary in Figure 11 at the 99.5-th percentile of pixel intensities across all input LDR views. These results validate the need for CRF regularization with a learnable prior that ensures monotonicity and promotes plausible response learning.

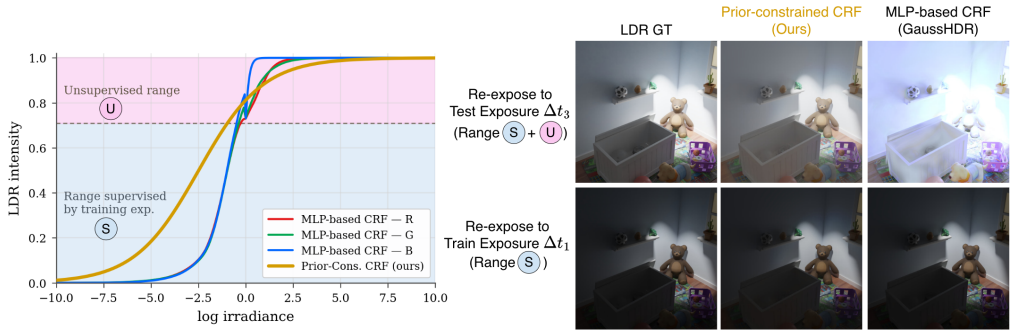

Figure 11: Learned MLP-based and Prior-Constrained CRF curves and novel-view LDR re-exposures. Both CRFs faithfully reproduce the LDR view at training exposure, where target intensities fall within their CRFs' supervised range. Beyond this range, MLP-based CRF curve exhibits RGB divergence, early saturation, and nonmonotonic behavior, resulting in poor rendering at test exposure that involves mapping to CRF's unsupervised range. Our prior-constrained CRF maintains a plausible, monotonic response in this range, producing more faithful test-exposure renderings.

A.6 Extension of 3D HDR Appearance Editing

Discussion. In Section 4.2, we discussed that without a 3D gain field, our 3D HDR editing capability cannot be achieved by scaling Gaussians' HDR radiance, as this does not expand the dynamic range (DR) and is only equivalent to an exposure shift. Figure 12 illustrates this point: Scaling 3D HDR radiance with ρ only shifts the log-luminance histogram and leaves the DR unchanged. In contrast, increasing ρ for 3D gain field scaling substantially expands the DR and changes the histogram. The histograms show the log-luminance distribution of HDR view 1 from the *desk* scene shown in Figure 6, restricted to values between the 1st and 99th percentiles.

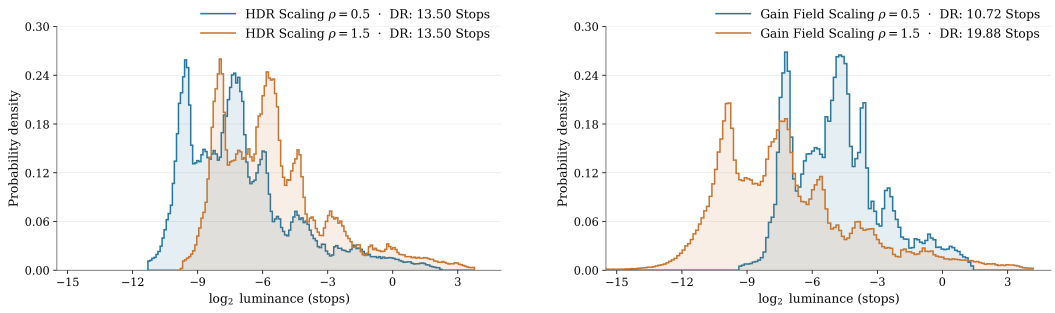

(a) Log-luminance histogram; HDR scaling

(b) Log-luminance histogram; Gain Field scaling

Figure 12: Changes in an HDR view's log-luminance histogram by varying ρ from 0.5 to 1.5 when ρ scales (a) 3D HDR radiance and (b) the 3D gain field. Increasing ρ for 3D HDR radiance scaling only shifts the histogram (equivalent to an exposure shift) and leaves the dynamic range (DR) unchanged. For 3D gain-field scaling, it instead changes the histogram and substantially expands the DR.

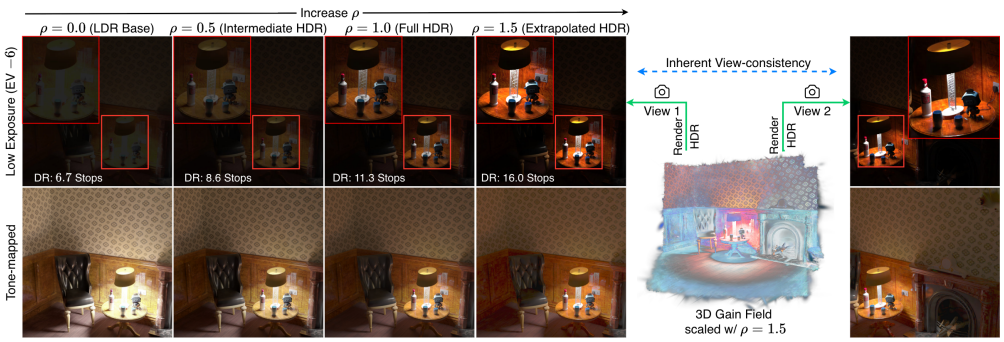

Figure 13: Further example of 3D HDR appearance editing via gain field scaling with ρ. We sweep ρ and show an HDR view at a low-exposure (*top row*) and after a fixed tone-mapping (*bottom row*). Increasing ρ expands dynamic range (DR) and enhances HDR appearance and chromatic contrast. Since ρ operates directly on 3D Gaussians, the edits are inherently view-consistent.

Another Example of Global Editing. Figure 13 complements Figure 6 by demonstrating 3D HDR appearance editing on another scene using the same ρ-scaling of the 3D gain field. As observed before, increasing ρ expands the DR and enhances the HDR appearance and chromatic contrast.

Extension to Localized Editing. We extend our 3D HDR appearance editing application to support localized edits. Specifically, instead of varying ρ globally across all 3D Gaussians (i.e. scaling the gain field globally with ρ), we adjust it only for a manually selected subset of Gaussians, enabling targeted control over the HDR appearance of a specific scene region or object. Figure 14 presents the results. We demonstrate different local adjustments to ρ across three scenes, each chosen to achieve a distinct HDR appearance in the desired region. For each scene, we show 2D visualizations of the 3D gain field before and after local editing. We also present HDR renderings from two viewpoints before and after the edits, displayed at a selected exposure. We observe that ρ adjustments achieve the desired local appearance changes. Furthermore, as with global editing, the local changes are inherently view-consistent as ρ adjustments are applied directly on 3D Gaussians.

A.7 Further Ablations

We provide qualitative evidence beyond those presented in Figure 7 for the ablations reported in Table 2 of main paper. Furthermore, we carry out an ablation on LDR-base supervision.

LDR-to-HDR Prior & Closed-Loop Reconstruction. Figure 15 provides further qualitative ablations of our LDR-to-HDR reconstruction prior and closed-loop reconstruction. Removing the LDR-to-HDR reconstruction prior yields significantly inaccurate HDR radiance recovery, while disabling the reconstruction loop introduces artifacts and degrades HDR quality. This further confirms the need for learned LDR-to-HDR reconstruction priors for high-fidelity HDR-NVS from single-exposure views and demonstrates the HDR quality gains from closed-loop reconstruction.

Base-Gain HDR Representation. Figure 16 presents additional qualitative ablations, where we replace our base-gain HDR representation with direct HDR encoding. We note that in our implementation of direct HDR encoding, each Gaussian's color SH coefficients encode log-radiance. This log-space formulation follows prior works (Cai et al., 2024; Liu et al., 2025) and, in our experiments, improves training stability and HDR reconstruction quality compared with linear encoding. We also increase the SH degree of the direct HDR encoding variant to match the color representation capacity of our base–gain representation and ensure fair comparison. As shown in Figure 16, cross-view inconsistencies among independently generated pseudo-HDR targets cause severe color and geometry artifacts. However, our base-gain HDR representation promotes robustness to those inconsistencies and achieves superior performance in HDR reconstruction. Additionally, in Figure 17, we present further qualitative ablations for replacing the gain field with MLPs that predict each Gaussian's HDR radiance from its LDR-base color $\boldsymbol{c}$, position $\boldsymbol{\mu}$, and viewpoint direction $\boldsymbol{d}$ ($\mathrm{MLP}(\boldsymbol{c}, \boldsymbol{\mu}, \boldsymbol{d})$) or base color $\boldsymbol{c}$ alone ($\mathrm{MLP}(\boldsymbol{c})$). Both MLP variants exhibit poor performance in reconstructing high-frequency textures in challenging dark regions. For the variant conditioned on

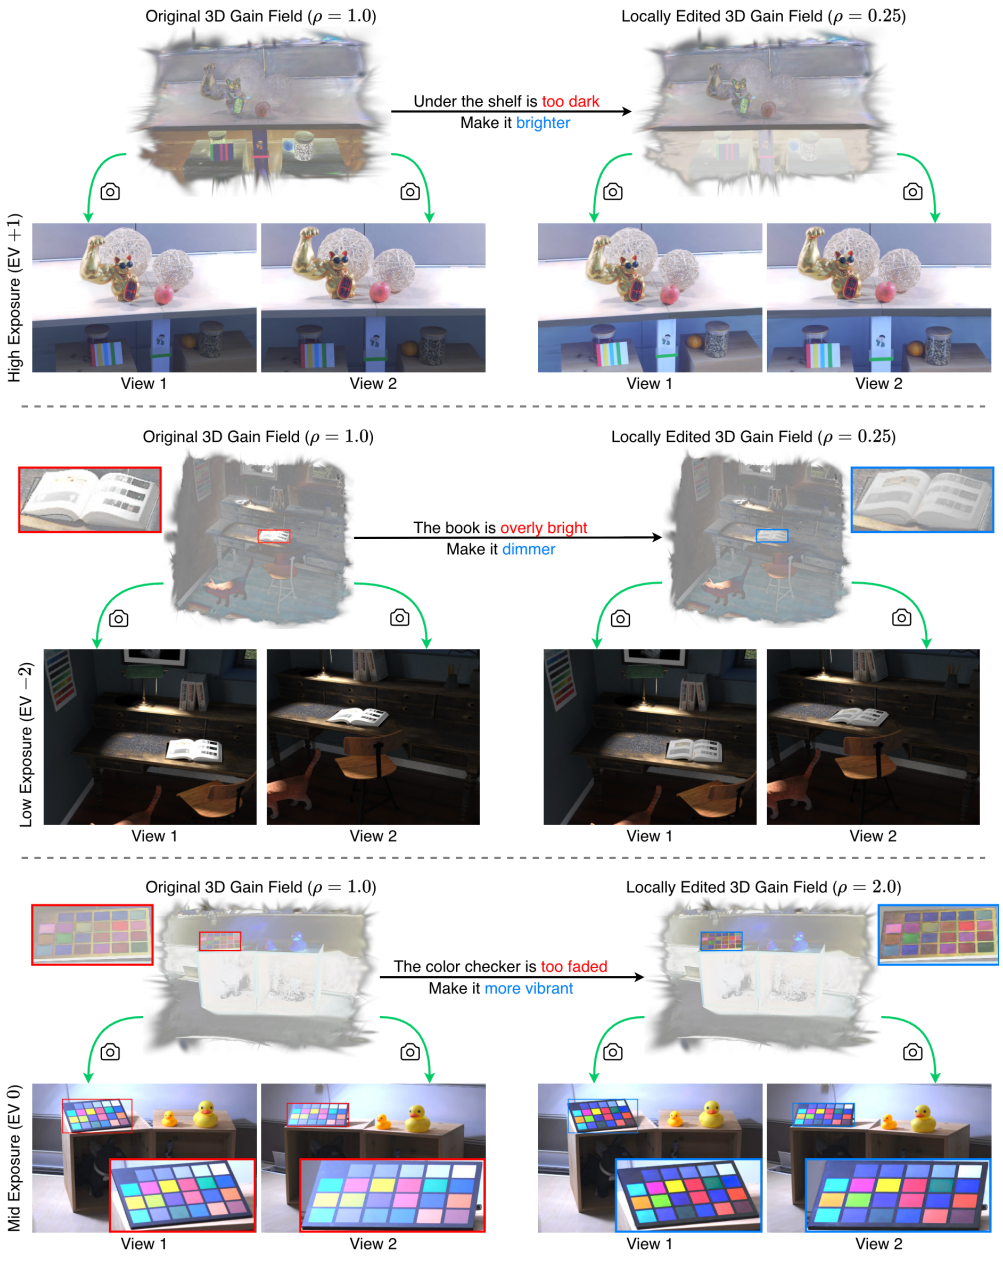

Figure 14: Localized 3D HDR appearance editing by adjusting ρ for a manually selected subset of 3D Gaussians across three scenes: (*top*) brightening the region beneath a shelf, (*middle*) dimming an overly bright book, and (*bottom*) enhancing a color checker's vividness. We show before-and-after 2D visualizations of the 3D gain field and HDR renderings from two viewpoints, displayed at a selected exposure. As ρ operates directly in 3D, the edits are inherently view-consistent.

LDR-base color, position, and viewing direction, this behavior is consistent with the low-frequency bias of MLPs operating on raw positions and view directions, which hinders reconstruction of sharp surface boundaries and detailed textures (Mildenhall et al., 2021). The color-only variant, on the other hand, predicts identical HDR radiance for identical base colors. However, Gaussians with the same base color may have different HDR radiance values, and this variant lacks local texture and broader context to distinguish these cases. Our base-gain representation, however, overcomes all discussed limitations and recovers high-frequency textures with higher fidelity. Finally, we note that

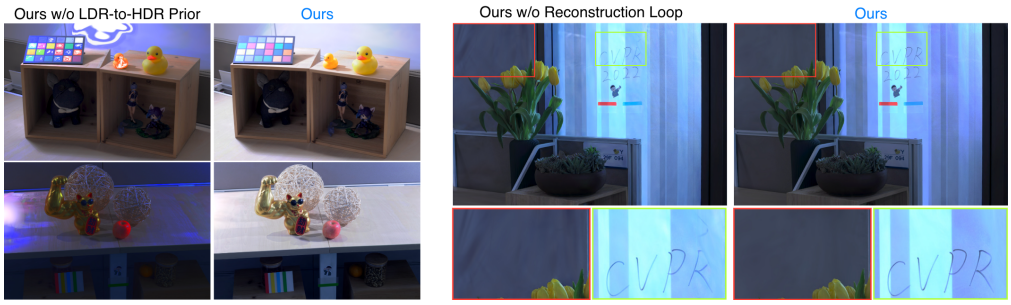

Figure 15: Further qualitative ablations on HDR views for (*left*) removing LDR-to-HDR reconstruction priors and (*right*) disabling the reconstruction loop. Removing the LDR-to-HDR reconstruction prior results in inaccurate HDR radiance recovery, while disabling the reconstruction loop degrades HDR reconstruction quality. All HDR views are tone-mapped identically for visualization.

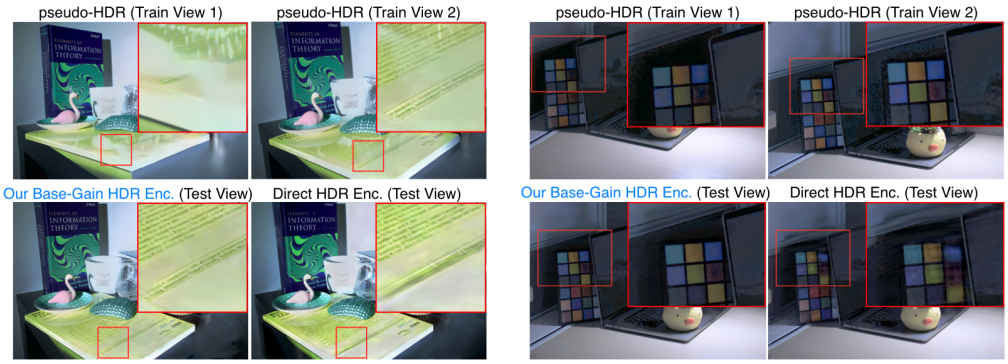

Figure 16: Further qualitative comparison (ablation) of our base-gain HDR representation against direct HDR encoding. (*top row*) Cross-view inconsistencies in pseudo-HDR targets (*bottom row*) produce severe artifacts with direct HDR encoding, while our base–gain representation promotes robustness to such inconsistencies and achieves higher quality HDR reconstruction.

direct HDR encoding and MLP variants preclude 3D HDR editing application enabled by the gain field, which further highlights the benefits of our base-gain representation.

Prior-Constrained CRF Modeling. In Figure 18, we ablate our proposed CRF with (1) fixed gamma encoding $\gamma = 2.2$; (2) an MLP applied on 2D log-irradiance obtained by exposing HDR renderings; and (3) an MLP applied on 3D log-irradiance obtained by exposing per-Gaussian HDR radiance. Following the recommendation of Cai et al. (2024) and Liu et al. (2025), we feed MLPs log-transformed irradiance rather than linear irradiance to improve training stability. The MLP-based CRFs closely fit the training exposure but generalize poorly to the unseen exposure, producing overexposure and severe color artifacts. Fixed gamma encoding also fails to reproduce the unseen exposure as it cannot adapt to individual scenes. Our CRF, however, achieves substantially higher fidelity than all alternatives at the unseen exposure: its strong prior regularizes its behavior beyond the observed exposure, while its learnable parameters enable scene-specific adaptation.

LDR-Base Supervision. In Figure 19, we qualitatively ablate LDR-base supervision by disabling its loss term in Equation 7. We observe that absence of LDR-base supervision severely impairs scene reconstruction, causing failure in texture recovery around the scene. This highlights the necessity of LDR-base supervision for robust HDR scene reconstruction. We further note that without LDR-base supervision HDR radiance is encoded arbitrarily across the LDR base and gain field, which lose their intended roles as LDR appearance and its HDR extension. This loss of a meaningful base–gain decomposition further precludes the 3D HDR editing application enabled by our base-gain representation and forgos joint reconstruction of the 3D LDR-base and HDR scenes.

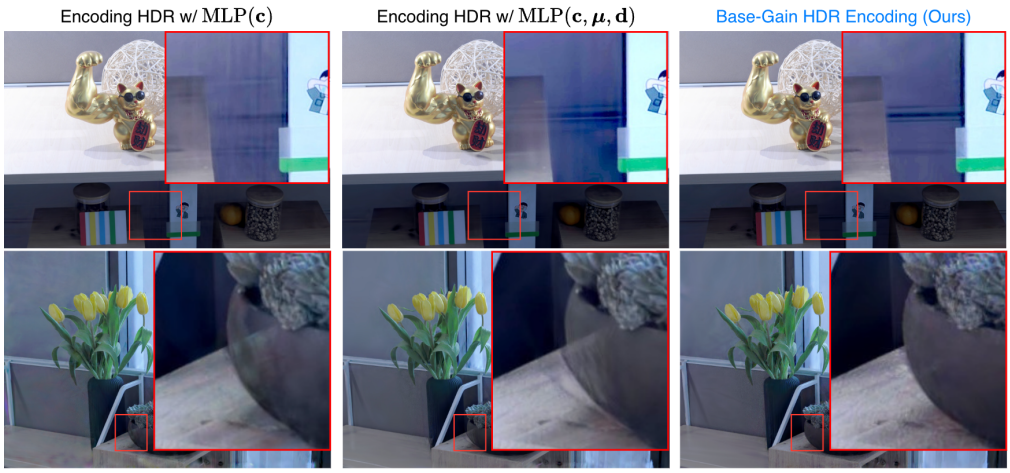

Figure 17: Further qualitative ablations on HDR views for replacing our gain field with two MLP variants: one predicts each Gaussian's HDR radiance from its LDR-base color $\boldsymbol{c}$ alone; the other also uses position $\boldsymbol{\mu}$ and direction $\boldsymbol{d}$. Our base-gain representation outperforms both MLP variants in reconstructing high-frequency details in regions with challenging lighting.

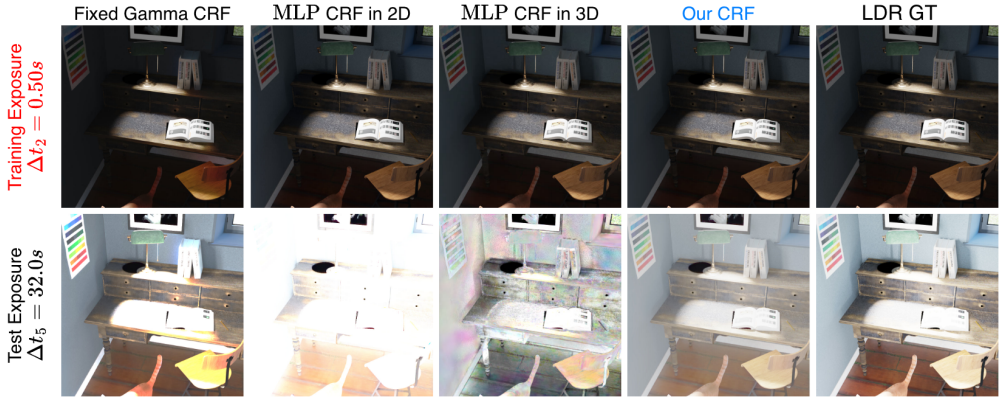

Figure 18: Qualitative ablations on LDR views for replacing our proposed CRF with MLPs in 2D and 3D and with fixed gamma CRF. All configurations reproduce the training exposure, yet our CRF achieves much higher reconstruction fidelity at the test exposure.

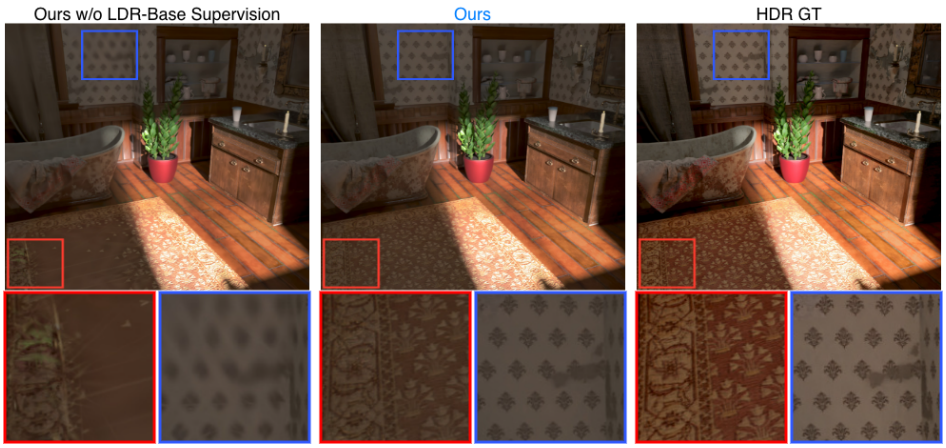

Figure 19: Qualitative ablation of LDR-base supervision on HDR views. Removing this supervision severely impairs scene reconstruction, causing failure in texture recovery all across the scene.