iBRIGHT-GS: IMMEDIATE BRIGHT GAUSSIAN RECONSTRUCTION FROM LOW-LIGHT VIEWS

Anonymous authors
Paper under double-blind review

ABSTRACT

Training 3D Gaussian Splatting exclusively on low-light images anchors the learned appearance to dark observations, which leads to persistently dim novel-view renderings even when the geometry is fully recovered. To tackle this issue, we propose iBright-GS, a framework that learns bright-view Gaussian appearance representations solely from low-light multi-view inputs. It first stabilizes geometric optimization via raw low-light photometric constraints to preserve the intrinsic dark appearance distribution of inputs and avoid premature photometric drift that degrades geometric fidelity. A unified training-stage bright signal is then constructed by synergistically combining a curve- and LUT-based pseudo-bright target, a dark-render brightening branch, per-Gaussian local gain, global exposure lock, and multi-view consistency. After geometry converges to a stable state with negligible structural perturbation, the constructed bright signal guides the optimization of dedicated bright Gaussian appearance coefficients. The learned bright representation can be directly rasterized at inference without additional post-processing, view-specific adjustment, or reference guidance. Extensive evaluations on LOM and RealX3D benchmarks illustrate that iBright-GS achieves state-of-the-art average PSNR and SSIM performance, with substantially higher runtime efficiency than existing advanced offline low-light 3DGS baselines.

1 INTRODUCTION

The core challenge in low-light novel view synthesis lies in the inherent bias that low-light supervision imposes on the learned 3D representation. When all training views are under-exposed, the Gaussian color parameters are optimized toward dark radiance distributions, leaving even a geometrically perfect reconstruction incapable of producing bright, natural-looking novel views. This motivates our investigation into whether low-light multi-view inputs alone can yield an immediately renderable 3D Gaussian representation that directly outputs normal-brightness views.

This challenge carries critical significance for immediate 3D reconstruction. While i3DGS (Meuleman et al., 2026) extends 3D Gaussian Splatting (3DGS) to unordered, streaming online inputs, low contrast, sensor noise, under-exposure, and color deviation in raw captures still severely degrade feature matching reliability and appearance learning quality. Directly deploying an existing immediate reconstruction pipeline can recover scene structure, but it lacks a dedicated photometric mechanism to transform the dark appearance into a consistently bright, directly renderable representation.

While 2D image enhancement effectively improves per-frame visibility, existing low-light 3DGS methods have demonstrated that robust appearance modeling is a critical determinant of high-fidelity 3D reconstruction. For immediate reconstruction, the central requirement is cross-view stable bright appearance inside the Gaussian representation, under the low-latency constraints of unordered scene reconstruction.

In this paper, we propose **iBright-GS**, a method that retains the reconstruction skeleton of i3DGS while explicitly learning bright Gaussian appearance purely under low-light supervision. Specifically, iBright-GS first leverages low-light observations to anchor geometry and preserve dark appearance distribution. Subsequently, it constructs a training-stage bright signal jointly from low-light inputs and dark renderings, and uses this supervisory signal to optimize independent bright SH0 coefficients on the already stabilized shared geometry. At inference, the resulting bright Gaussian

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.

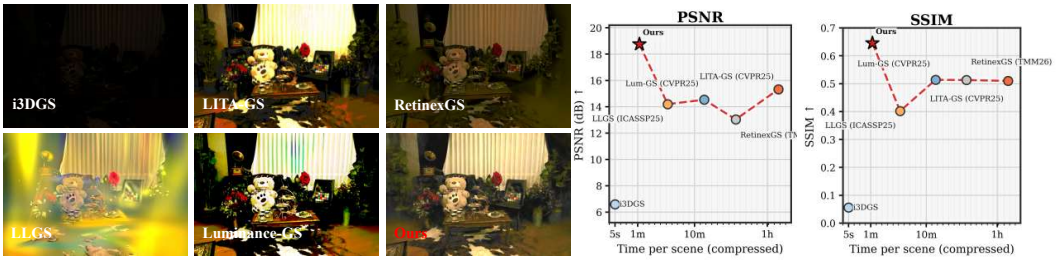

Figure 1: From low-light inputs, iBright-GS reconstructs bright Gaussian appearance for normal-brightness novel-view rendering. The left panel compares rendered novel views from recent low-light 3DGS methods, while the right panels summarize PSNR and SSIM under the same setting. iBright-GS produces brighter and more faithful appearance while maintaining the immediate reconstruction setting.

representation can be rasterized directly without extra processing. The contributions of this work are summarized as follows.

- ***Bright Gaussian reconstruction from low-light inputs.*** We investigate a practical setting where only low-light views are available for training, while all normal-exposure images are strictly reserved for evaluation, with the explicit goal of rendering physically plausible normal-brightness novel views directly from the optimized 3D Gaussians.
- ***Dark-anchored dual-appearance representation.*** We explicitly separate dark appearance from bright appearance on top of the shared, fixed scene geometry: the dark branch explains raw low-light observations and stabilizes the whole reconstruction process, while dedicated bright SH0 coefficients carry the full normal-brightness rendering capability.
- ***Low-light-driven bright appearance learning.*** We systematically integrate curve/LUT pseudo-bright targets, dark-render brightening branches, per-Gaussian local gain, global exposure lock, and multi-view consistency constraints into a coherent, stable training-time supervisory signal to guide the optimization of the target bright Gaussian appearance.

2 Related Work

3D Gaussian Splatting and Immediate Reconstruction. Neural Radiance Fields introduced continuous volumetric view synthesis from posed images (Mildenhall et al., 2020). Follow-ups accelerated training and rendering with explicit or compact representations, including PlenOctrees (Yu et al., 2021), DVGO (Sun et al., 2022), TensoRF (Chen et al., 2022), Instant-NGP (Müller et al., 2022), and Mip-NeRF 360 (Barron et al., 2022). 3DGS represents scenes with anisotropic Gaussians and differentiable rasterization for fast training and real-time novel view synthesis (Kerbl et al., 2023). Recent variants improve geometry with 2DGS (Huang et al., 2024), view-adaptive rendering with Scaffold-GS (Lu et al., 2024), and surface extraction with SuGaR (Guédon & Lepetit, 2024).

Standard 3DGS assumes known poses and a full image set. Pose-free methods such as BARF (Lin et al., 2021) and NoPe-NeRF (Bian et al., 2023) jointly optimize reconstruction and camera registration. i3DGS supports unordered online inputs via matching, pose estimation, depth guidance, keyframe selection, incremental densification, and appearance optimization (Meuleman et al., 2026). Yet both 3DGS and i3DGS still assume adequately lit inputs: low-light supervision drives Gaussian colors toward dark radiance and leaves renders dim. iBright-GS keeps this immediate reconstruction skeleton and studies how to form bright appearance inside it under low-light supervision.

Low-Light Image Enhancement. Low-light image enhancement (LLIE) recovers visibility from under-exposed, low-contrast, and noisy observations. Deep LLIE spans supervised raw enhancement, Retinex decomposition, zero-reference curves, unpaired enhancement, flows, unfolding networks, and self-calibrated illumination modeling (Chen et al., 2018; Wei et al., 2018; Guo et al., 2020; Jiang et al., 2021; Wang et al., 2022; Wu et al., 2022; Ma et al., 2022).

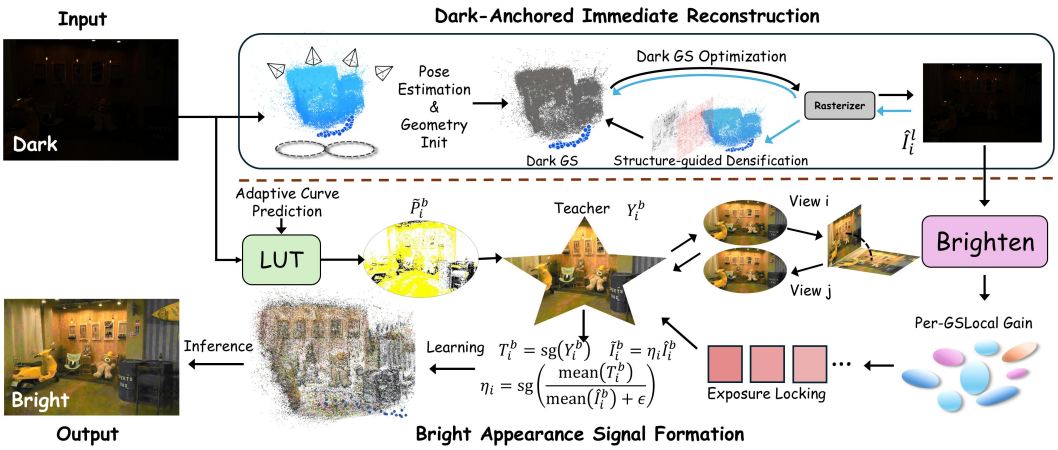

Figure 2: Overall pipeline of iBright-GS. Low-light inputs first drive pose estimation, geometry initialization, structural-gap-guided densification, and dark Gaussian optimization to obtain a dark-anchored reconstruction. Subsequently, iBright-GS constructs a bright training signal by integrating an adaptive curve/LUT prediction derived from the input and a brightened dark-render teacher. Exposure locking and per-Gaussian local gain further stabilize cross-view brightness consistency, and the constructed signal supervises the appearance learning of bright Gaussians to enable direct normal-brightness rendering.

In 3D pipelines, LLIE is usually applied either before reconstruction or after rendering. Input-side enhancement improves single-view visibility but often causes cross-view brightness drift and color inconsistency; post-render adjustment only edits the final 2D image and cannot yield a consistent bright 3D asset. iBright-GS therefore uses LLIE ideas only at training time, converting exposure adjustment, curve mapping, and local gain into supervision for bright Gaussian appearance.

Low-Light Neural Rendering and Low-Light 3DGS. Low-light neural rendering must model weak illumination, noise, exposure variation, and color shift inside the 3D representation. Prior work emphasizes raw/HDR modeling, dark-scene reconstruction, and structure recovery under low light (Mildenhall et al., 2022; Zhang et al., 2024; Guo et al., 2026).

Existing low-light 3DGS methods are mostly offline. Gaussian-in-the-Dark (Ye et al., 2024) and LE3D (Jin et al., 2024) target dark or HDR synthesis; LLGS (Gan et al., 2025), RetinexGS (Wang et al., 2026), Luminance-GS (Cui et al., 2025), LITA-GS (Zhou et al., 2025), and eRetinexGS (Yan et al., 2026) further add absorptance modulation, Retinex decomposition, view-adaptive curves, physical priors, or event streams. iBright-GS instead builds a bright Gaussian representation inside an immediate 3DGS pipeline from raw low-light multi-view inputs.

3 Methodology

3.1 Overview

Given a set of low-light multi-view inputs $\mathcal{I}^l = \{I_i^l\}_{i=1}^N$, where I_i^l denotes the i-*th* low-light view and N is the total number of training views, no normal-exposure images or paired reference supervision are introduced during the entire training phase; normal-exposure images are reserved exclusively for evaluation. The corresponding camera parameters are denoted as π_i. The proposed iBright-GS first reconstructs a structurally stable dark-anchored scene from the raw low-light observations, and then constructs a 3D-consistent bright signal Y_i^b to supervise the learning of bright Gaussian appearance.

3.2 Dark-Anchored Immediate Reconstruction

The dark-anchored reconstruction stage inherits the immediate reconstruction pipeline from i3DGS and operates directly on raw low-light images. It tightly couples pose estimation and geometry

optimization to the observed multi-view data. This stage stabilizes camera parameters, geometric structure, and dark appearance under low-light supervision, and provides a dependable shared geometry foundation for the subsequent bright appearance formation.

We further instantiate this dark-anchored paradigm into a dedicated dark rendering branch for explicit low-light modeling. The dark branch rasterizes a standard 3DGS model:

$$\hat{I}_i^l = \mathcal{R}(\mathcal{G}^d, \pi_i), \tag{1}$$

where $\mathcal{R}(\cdot)$ indicates the differentiable 3DGS rasterizer and $\mathcal{G}^d$ denotes the dark Gaussian representation. The low-light photometric loss is defined as follows:

$$\mathcal{L}_{low} = \lambda_s \left(1 - \mathrm{SSIM}(\hat{I}_i^l, I_i^l)\right) + (1 - \lambda_s)\|\hat{I}_i^l - I_i^l\|_1, \tag{2}$$

where λ_s serves as a balancing weight between the SSIM and L1 loss terms, and we define this combined formulation as the dark-anchoring loss term.

However, under low-light conditions, conventional densification is highly susceptible to interference from weak textures and dark noise. To address this issue, we develop a structure-gap-guided densification strategy that retains the immediate sampling and back-projection workflow, while scoring regions where the observed low-light structural information is absent from the current rendering. Let $\mathrm{Lap}(\cdot)$ denote the Laplacian response, then the densification score is formulated as follows,

$$S_{densify}(x) = \lambda_{img}\,\mathrm{Lap}(I_i^l)(x) + \lambda_{gap}\max\left(0, \mathrm{Lap}(I_i^l)(x) - \mathrm{Lap}(\hat{I}_i^l)(x)\right), \tag{3}$$

where the first term preserves the observed structural details, and the second term highlights the structural discrepancy between the input observation and the rendered output. The derived score guides the densification process toward boundaries and textures that remain missed in the current reconstruction.

3.3 Dual-Appearance Gaussian Representation

In general, standard 3DGS represents a scene as a collection of anisotropic Gaussians. Each Gaussian is defined by its center $\boldsymbol{\mu}_k$, covariance Σ_k, opacity α_k, and spherical-harmonic appearance coefficients, where the covariance is parameterized via scale and rotation as $\Sigma_k = R_k S_k S_k^\top R_k^\top$. Leveraging the robust shared geometry obtained from the prior dark-anchored stage, we adapt this base representation to support explicit dual-appearance modeling. Figure 3 summarizes the shared geometry and the two appearance branches. Concretely, as low-light supervision inherently drives these color parameters toward dark appearance learning, iBright-GS explicitly decouples and preserves separate dark and bright appearance representations on top of the shared underlying geometry:

$$\mathcal{G}^d = \{\boldsymbol{\mu}_k, \Sigma_k, \alpha_k, \mathbf{f}_{k,0}^d, \mathbf{f}_{k,>0}^d\}_{k=1}^K, \qquad \mathcal{G}^b = \{\boldsymbol{\mu}_k, \Sigma_k, \alpha_k, \mathbf{f}_{k,0}^b, \mathbf{f}_{k,>0}^d\}_{k=1}^K, \tag{4}$$

where K is the number of Gaussians. Both branches share geometry, opacity, and the higher-order SH coefficients $\mathbf{f}_{k,>0}^d$; the bright branch replaces only the dark SH0 $\mathbf{f}_{k,0}^d$ with an independent coefficient $\mathbf{f}_{k,0}^b$ that encodes normal-brightness appearance as a compact Gaussian attribute. Constraining the dark–bright gap to this low-order term mitigates view-dependent overfitting under low-light-only supervision.

With the shared geometry and dual appearance branches in place, the dark and bright renderings are

$$\hat{I}_i^l = \mathcal{R}(\mathcal{G}^d, \pi_i), \qquad \hat{I}_i^b = \mathcal{R}(\mathcal{G}^b, \pi_i). \tag{5}$$

We initialize the bright low-order color coefficient by applying a brightness enhancement operation directly on the pre-optimized dark low-order color. Let $C(\cdot)$ and $C^{-1}(\cdot)$ represent the forward and inverse SH0-RGB conversion mappings respectively:

$$\mathbf{c}_k^b = \mathrm{clip}\left(C(\mathbf{f}_{k,0}^d)^\gamma \odot \mathbf{a} + \mathbf{b}, 0, 1\right), \qquad \mathbf{f}_{k,0}^b = C^{-1}(\mathbf{c}_k^b). \tag{6}$$

This initialization step only provides a reasonable starting point, while the final high-quality bright appearance is fully optimized in the subsequent dedicated bright-appearance learning stage.

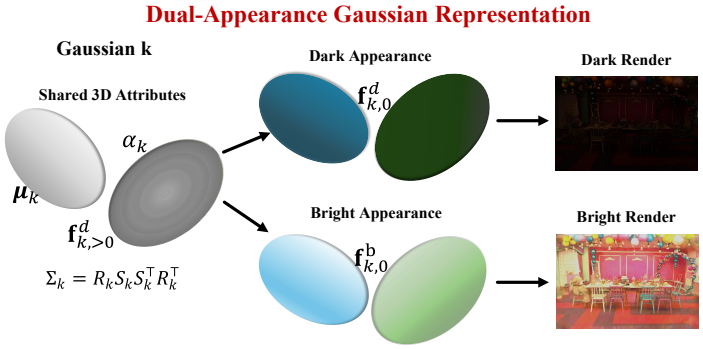

Figure 3: Dual-Appearance Gaussian Representation. The dark and bright branches fully share the Gaussian geometry, opacity, and all higher-order spherical harmonic coefficients, while the bright branch learns an exclusive independent SH0/DC color coefficient to support direct normal-brightness rendering.

3.4 Bright Appearance Signal Formation

Since no normal-exposure images are accessible throughout the training phase, our designed iBright-GS constructs a bright learning signal via two complementary sources: an image-space curve/LUT reference derived from the low-light input, and a 3D-constrained brightening basis extracted from the current dark rendering. We denote the adaptive curve parameters predicted from I_i^l be A_i:

$$P_i^b(x) = \mathrm{LUT}(I_i^l(x), A_i), \qquad \tilde{P}_i^b = \mathrm{clip}\left(P_i^b \cdot \frac{\tau}{q_{0.95}(I_i^l) + \epsilon}, 0, 1\right), \tag{7}$$

where x denotes the pixel location, $q_{0.95}(I_i^l)$ refers to the 95-*th* percentile brightness, and τ defines the target brightness scale. The dark-render branch further maps the current low-light rendering to the desired brightening basis, which anchors the learned signal to the appearance already modeled by the 3D reconstruction. The formula is defined as follows,

$$E_i^b = B(\hat{I}_i^l), \tag{8}$$

Furthermore, to address spatially non-uniform illumination, each Gaussian is assigned a bounded local gain,

$$G_k = 1 + \alpha_g \tanh(z_k), \qquad M_i^G = \mathcal{R}(\{\boldsymbol{\mu}_k, \Sigma_k, \alpha_k, G_k\}_{k=1}^K, \pi_i), \tag{9}$$

where z_k indicates an unconstrained learnable latent parameter for the k-*th* Gaussian, and M_i^G denotes the rasterized gain map. The gain-adjusted signal along with its corresponding regularizer are defined as follows,

$$E_i^G = \mathrm{clip}\left(E_i^b \odot M_i^G, 0, 1\right), \qquad \mathcal{L}_{gain} = \frac{1}{K}\sum_{k=1}^{K} \|z_k\|_2^2. \tag{10}$$

Afterward, the exposure lock module calibrates the global brightness to the precomputed curve/LUT reference,

$$\bar{L}_i = \mathrm{sg}\left(\mathrm{mean}(\tilde{P}_i^b)\right), \qquad Y_i^b = E_i^G \cdot \mathrm{sg}\left(\frac{\bar{L}_i}{\mathrm{mean}(E_i^G) + \epsilon}\right). \tag{11}$$

We further regularize this signal across different views. Specifically, given a pixel x in source view i and a neighboring view j, we use the current depth $D_i(x)$ to back-project x into 3D and re-project it into view j, obtaining its corresponding location x_j:

$$x_j = \pi_j\left(\pi_i^{-1}(x, D_i(x))\right), \tag{12}$$

and the multi-view consistency loss is formulated as follows,

$$\mathcal{L}_{mv} = \frac{1}{|\Omega_{ij}|}\sum_{x \in \Omega_{ij}} \left\| W_{i\to j}(Y_i^b)(x_j) - Y_j^b(x_j) \right\|_1, \tag{13}$$

Table 1: Quantitative comparison on LOM and RealX3D. Best and second-best results are marked in bold and underlined, respectively.

Method	LOM				RealX3D			
	PSNR↑	SSIM↑	LPIPS↓	Time	PSNR↑	SSIM↑	LPIPS↓	Time
3DGS	6.74	0.137	0.603	2:01	6.60	0.061	0.691	3:17
i3DGS	6.79	0.136	0.579	**0:04**	6.60	0.056	0.694	**0:07**
LLGS	17.74	0.783	0.330	4:02	12.35	0.493	0.609	6:09
Luminance-GS	18.69	0.784	0.310	6:54	14.54	<u>0.514</u>	0.489	18:27
LITA-GS	20.16	0.792	0.304	27:25	<u>15.33</u>	0.510	0.480	81:35
RetinexGS	<u>21.09</u>	**0.823**	**0.257**	12:06	13.04	0.513	<u>0.477</u>	20:01
iBright-GS	**23.63**	<u>0.819</u>	<u>0.288</u>	<u>0:22</u>	**18.77**	**0.646**	**0.475**	<u>1:31</u>

where Ω_{ij} contains valid, in-bounds, and unoccluded pixels. This operation suppresses cross-view brightness and color drift within the learned bright signal.

3.5 Bright Appearance Learning and Inference

After the dark-anchored reconstruction, the formation stage learns the curve reference, dark-render brightening branch, and per-Gaussian local gain, while using exposure lock and multi-view consistency to yield a 3D-consistent bright signal. Its overall optimization objective is as follows,

$$\mathcal{L}_{form} = \mathcal{L}_{low} + \lambda_{sig}\mathcal{L}_{sig} + \lambda_{spa}\mathcal{L}_{spa} + \lambda_{hist}\mathcal{L}_{hist} + \lambda_{lock}\mathcal{L}_{lock} + \lambda_{mv}\mathcal{L}_{mv} + \lambda_{gain}\mathcal{L}_{gain}, \quad (14)$$

where $\mathcal{L}_{low}$ preserves the dark-anchored reconstruction quality, $\mathcal{L}_{sig}$ aligns the learned brightening basis with the curve/LUT reference, $\mathcal{L}_{spa}$ and $\mathcal{L}_{hist}$ denote the spatial illumination and brightness-statistics regularizers, respectively, $\mathcal{L}_{lock}$ calibrates the global exposure level, $\mathcal{L}_{mv}$ enforces cross-view consistency, and $\mathcal{L}_{gain}$ regularizes the per-Gaussian gains. Concretely,

$$\mathcal{L}_{sig} = \lambda_e \left(1 - \mathrm{SSIM}(E_i^G, \mathrm{sg}(\tilde{P}_i^b))\right) + (1-\lambda_e)\|E_i^G - \mathrm{sg}(\tilde{P}_i^b)\|_1, \quad (15)$$

$$\mathcal{L}_{lock} = \left|\mathrm{mean}(E_i^G) - \mathrm{sg}(\mathrm{mean}(\tilde{P}_i^b))\right|. \quad (16)$$

The bright appearance learning stage freezes the geometry, camera poses, dark appearance, and signal-generation branches, and then optimizes only the bright SH0 coefficients. To convert the formed bright signal into a fixed supervision target, we block its gradients and define the target as $T_i^b = \mathrm{sg}(Y_i^b)$. We further normalize the bright rendering with a detached global scale factor η_i to align its overall brightness with the target.

$$T_i^b = \mathrm{sg}(Y_i^b), \qquad \tilde{I}_i^b = \eta_i \hat{I}_i^b, \qquad \eta_i = \mathrm{sg}\left(\frac{\mathrm{mean}(T_i^b)}{\mathrm{mean}(\hat{I}_i^b) + \epsilon}\right), \quad (17)$$

The bright SH0 coefficients are then optimized by minimizing the pixel-wise discrepancy between the normalized bright rendering and the fixed target:

$$\mathcal{L}_{learn} = \lambda_{learn}(t)\|\tilde{I}_i^b - T_i^b\|_1, \quad (18)$$

where $\lambda_{learn}(t)$ gradually increases as training progresses. During inference, the learned bright representation is directly rasterized as $I_{out} = \mathcal{R}(\mathcal{G}^b, \pi_{test})$.

4 Experiments

4.1 Experimental Setup

Datasets and Metrics. We conduct extensive experiments of our iBright-GS on LOM (Cui et al., 2024) and RealX3D (Liu et al., 2026) datasets. The LOM dataset comprises 5 real-world low-light scenes, while the RealX3D dataset includes 9 real low-light scenarios. Both benchmarks provide low-light / normal-exposure multi-view pairs. During the training phase, only the low-light views are utilized, and all normal-exposure images are strictly used for evaluation. For a fair comparison, we adopt the identical 8-split protocol, in which every 8-*th* image is withheld as the unseen test view. iBright-GS uses only the dataset images without COLMAP preprocessing; for baseline methods, we follow their official protocols and provide dataset images together with the dataset-provided COLMAP outputs. For evaluation, we benchmark all methods using four standard metrics: PSNR,



Figure 4: Qualitative comparison on RealX3D (*BlueHawaii*, *Sculpture*, *Cupcake*) and LOM (*sofa*). i3DGS remains dark; low-light 3DGS baselines brighten the view but often distort crops. iBright-GS raises brightness while keeping local structure because bright appearance is stored in the Gaussians.

SSIM (Wang et al., 2004), LPIPS (Zhang et al., 2018), and the measured wall-clock runtime per scene.

Implementation Details. We use the Adam optimizer for all iBright-GS parameters, and directly follow the official i3DGS configuration for the remaining reconstruction parameters. Each online keyframe uses 80 optimization iterations, followed by 80 finetuning epochs; the learning rates are set to 2×10^{-3} for the bright signal, 5×10^{-3} for bright SH0 coefficients, and 5×10^{-3} for per-Gaussian gain. We fix the random seed to 0 and keep this full configuration consistent across both the LOM and RealX3D datasets. Runtime is reported as single-scene wall-clock time on NVIDIA RTX 4090 GPUs.

4.2 Comparison with SOTAs

To validate the effectiveness of our proposed iBright-GS, we conduct a thorough comparative evaluation against six representative baselines, including 3DGS (Kerbl et al., 2023), i3DGS (Meuleman et al., 2026), LLGS (Gan et al., 2025), Luminance-GS (Cui et al., 2025), LITA-GS (Zhou et al., 2025), and RetinexGS (Wang et al., 2026). Among these, 3DGS and i3DGS directly fit the low-light input to measure the degradation without low-light appearance modeling; while the remaining four methods serve as low-light NVS or low-light 3DGS baselines.

Quantitative Analysis. Table 1 reports the quantitative averages on two datasets. Both vanilla 3DGS and i3DGS yield PSNR values around 6.7–6.8 dB on the LOM and 6.60 dB on RealX3D. Although these two methods are capable of recovering usable camera poses and coarse geometry, their learned spherical harmonic coefficients remain persistently anchored to the dark observations, which prevents the rendered held-out novel views from achieving normal exposure levels. This empirical observation directly validates our core representation claim: reliable geometry recovery does not inherently guarantee the successful recovery of a physically plausible bright appearance.

Table 2: Ablation studies of Ours and its four variants on LOM and RealX3D datasets.

Methods	LOM				RealX3D			
	PSNR↑	SSIM↑	LPIPS↓	ΔPSNR	PSNR↑	SSIM↑	LPIPS↓	ΔPSNR
w/o Bright SH0	6.81	0.139	0.527	-16.82	6.61	0.056	0.676	-12.16
w/o Freeze Geo	21.52	0.754	0.340	-2.11	18.57	0.637	0.479	-0.20
w/o Signal Sep	23.15	0.804	0.295	-0.48	16.79	0.599	0.506	-1.98
w/o Struct Dens.	23.30	0.806	0.302	-0.33	18.58	0.638	0.482	-0.19
Ours	23.63	0.819	0.288	0.00	18.77	0.646	0.475	0.00



Figure 5: Visual ablation on RealX3D *GearWorks*. From left to right: GT, *w/o* Bright SH0, *w/o* Freeze Geo, *w/o* Signal Sep, *w/o* Struct Dens., and Ours. *w/o* Bright SH0 collapses to dark fitting; *w/o* Freeze Geo smears structure; *w/o* Signal Sep softens detail; *w/o* Struct Dens. introduces mottled artifacts; Ours recovers bright appearance close to GT.

Existing low-light 3DGS methods partially narrow this performance gap, yet each introduces distinct practical trade-offs. As shown in Table 1, on the LOM dataset, RetinexGS obtains advanced performance, delivering 21.09 dB / 0.823 SSIM / 0.257 LPIPS with an average runtime of 12:06. In contrast, iBright-GS achieves a PSNR of 23.63 dB, marking a 2.54 dB improvement over RetinexGS and a 3.47 dB gain over LITA-GS, while completing per-scene optimization in only 0:22 and maintaining a competitive SSIM score of 0.819.

In addition, on the more challenging RealX3D dataset, iBright-GS achieves 18.77 dB PSNR, 0.646 SSIM, and 0.475 LPIPS in only 1.5 minutes per scene, outperforming all competing methods across three metrics. This further indicates that iBright-GS successfully decouples robust geometry recovery from dark-locked appearance fitting, enabling high-fidelity bright scene reconstruction even under extremely low-illumination inputs without compromising the efficiency advantages of the original Gaussian splatting paradigm. Appendix D visualizes the corresponding quality–speed tradeoff.

Qualitative Analysis. Fig. 4 compares i3DGS, four low-light 3DGS baselines, and iBright-GS on RealX3D scenes *BlueHawaii*, *Sculpture*, and *Cupcake*, and on the LOM scene *sofa*. From Fig. 4, the comparison shows that i3DGS confirms the quantitative dark-fitting result: structure is partly visible, but the final rendering stays underexposed and the zoomed crops lose color-chart patches, fabric grain, and text. LITA-GS, RetinexGS, LLGS, and Luminance-GS all elevate the global brightness of the rendered outputs, yet their crops still show color cast on the checkerboard and street sign in *BlueHawaii*, over-smoothed fabric texture on the sofa, noisy foliage surrounding *Sculpture*, and flickering, unstable table decorations in the *Cupcake* scene.

iBright-GS maintains a globally well-lit scene appearance, while its crops retain sharper edges and more coherent local albedo. This is consistent with storing normal-brightness appearance in dedicated SH0 coefficients on frozen geometry, rather than applying a view-wise enhancement map after rasterization. However, minor residual errors still remain concentrated in regions of extremely dark foliage and around complex, high-frequency highlight boundaries.

4.3 Ablation Analysis

To analyze the effectiveness of iBright-GS, we remove the modules that most directly affect the final bright Gaussians: the independent bright SH0 branch $\mathbf{f}^b_{k,0}$ (termed as *w/o* Bright SH0), geometry freezing during bright-appearance optimization (denoted as *w/o* Freeze Geo), the separation between training-time signal formation and bright SH0 learning (termed as *w/o* Signal Sep), and structure-gap-guided densification (known as *w/o* Struct Dens.). Fig. 5 shows a visual comparison on RealX3D *GearWorks*, and Table 2 reports the dataset averages.

Effectiveness Analysis. Table 2 shows that *w/o* Bright SH0 causes a severe degradation: PSNR falls to 6.81 on LOM (-16.82 dB) and 6.61 on RealX3D (-12.16 dB), which fully reverts the model to the dark-fitting regime observed in baseline 3DGS and i3DGS. Fig. 5 illustrates the same collapse

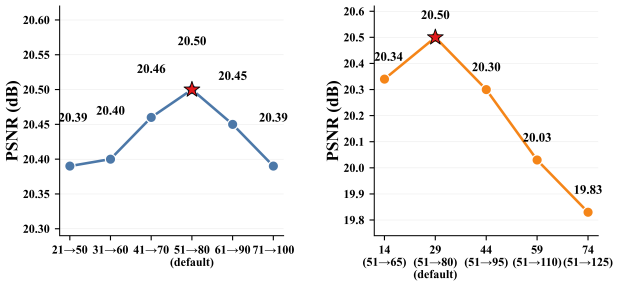

(a) Start step, fixed 29-step window (b) Window length, fixed start at 51

Figure 6: Sensitivity to the bright-appearance learning schedule on the pooled LOM+RealX3D average.

Table 3: Representative hyperparameter sensitivity on the pooled LOM+RealX3D average. Bold indicates the default setting.

Parameter	Value	PSNR↑	SSIM↑	LPIPS↓
w_{enh}	0.01	20.56	0.706	0.408
	1.0	**20.50**	**0.708**	**0.408**
	10	20.14	0.705	0.420
learning start	41 to 70	20.46	0.707	0.409
	51 to 80	**20.50**	**0.708**	**0.408**
	61 to 90	20.46	0.707	0.409
learning length	14 (51 to 65)	20.34	0.704	0.411
	29 (51 to 80)	**20.50**	**0.708**	**0.408**
	44 (51 to 95)	20.30	0.706	0.412

visually: without the bright SH0 branch, the render stays dark even when geometry is present. This is the main ablation supporting our representation claim: normal-brightness novel-view synthesis needs a dedicated Gaussian attribute, and sharing only dark SH0 is not enough.

The remaining factors govern how the bright attribute is optimized. As reported in Table 2, *w/o* Freeze Geo yields a -2.11 dB drop on LOM, while the degradation on RealX3D is only -0.20 dB. On LOM, the dark-anchored scaffold is already geometrically accurate; updating Gaussians while fitting the bright supervision smears edges and fine textures (SSIM 0.819→0.754, LPIPS 0.288→0.340). RealX3D's geometry contains inherent noise, so the absolute PSNR gap is smaller, but Fig. 5 still shows blurred bicycle spokes and softened surfaces when geometry remains unlocked.

w/o Signal Sep reveals a complementary pattern: it only causes a -0.48 dB drop on LOM, but a notable -1.98 dB degradation on RealX3D. When the dark-to-bright gap is large, jointly updating the teacher Y_i^b and the bright SH0 coefficients forces supervision to chase a moving target, which appears as haze and missing fine detail in Fig. 5.

w/o Struct Dens. yields a smaller but consistent drop of -0.33 dB on LOM and -0.19 dB on RealX3D; Fig. 5 shows mottled color noise on walls and floors, indicating that better low-light structural coverage assists the later bright branch. Overall, these ablations confirm our staged design: first recover dark-anchored geometry, then build a stable bright signal, and finally encode normal brightness into $\mathbf{f}_{k,0}^b$ on fully frozen geometry.

Sensitivity Analysis. We conduct a systematic evaluation of representative hyperparameters across 14 scenarios on two datasets, as reported in Fig. 6 and Table 3. Our default 51-to-80 training window with $w_{enh} = 1.0$ achieves 20.50 PSNR / 0.708 SSIM / 0.408 LPIPS. Loss weights stay stable around this default: when $w_{enh} \in \{0.01, 1.0\}$, the PSNR fluctuation remains within 0.06 dB, while $w_{enh} = 10$ degrades performance to 20.14 PSNR and raises LPIPS to 0.420. This confirms that over-weighting the image-space curve/LUT target diverts the learned bright teacher away from the 3D-constrained dark-render branch.

The learning schedule ablation further verifies our dark-then-bright staging design. With a fixed 29-step window, PSNR hits 20.46 at start step 41, peaks at 20.50 at step 51, and falls back to 20.46 at step 61. An early start updates bright SH0 before cameras and Gaussian density stabilize; a late start leaves too few steps for the bright branch. With the start step fixed at 51, a 14-step window yields 20.34 PSNR, while a 44-step window degrades to 20.30 PSNR. Short windows underfit the bright attribute, while long ones overfit the training-time teacher and raise LPIPS from 0.408 to 0.412. Fig. 6 visualizes this trend, with performance peaking exactly at our default setting.

5 Conclusion

We present iBright-GS, an efficient framework that recovers bright Gaussian appearance directly from low-light multi-view images. It learns normal-brightness appearance as a native renderable property within the 3D Gaussian representation, and outputs directly renderable bright Gaussians without any normal-exposure training images. Experiments on LOM and RealX3D datasets demonstrate strong reconstruction quality and significantly superior efficiency.

AI USE STATEMENT

Generative AI tools were used to assist with language polishing, organization, and experiment description editing. All method design, implementation, experiments, numerical results, and final claims were checked and confirmed by the authors. We take responsibility for the final content of this work, including text, claims, and artifacts produced with the aid of generative AI.

REFERENCES

Jonathan T. Barron, Ben Mildenhall, Dor Verbin, Pratul P. Srinivasan, and Peter Hedman. Mip-NeRF 360: Unbounded anti-aliased neural radiance fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5470–5479, 2022.

Wenjing Bian, Zirui Wang, Kejie Li, Jia-Wang Bian, and Victor Adrian Prisacariu. NoPe-NeRF: Optimising neural radiance field with no pose prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4160–4169, 2023.

Anpei Chen, Zexiang Xu, Andreas Geiger, Jingyi Yu, and Hao Su. TensoRF: Tensorial radiance fields. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 333–350, 2022.

Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 3291–3300, 2018.

Ziteng Cui, Lin Gu, Xiao Sun, Xianzheng Ma, Yu Qiao, and Tatsuya Harada. Aleth-NeRF: Illumination adaptive NeRF with concealing field assumption. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1435–1444, 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 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 26472–26482, 2025.

Jianwen Gan, Wenxin Li, Bo Zheng, Chengliang Wang, and Yingbo Wu. LLGS: Illuminating gaussian splatting via absorptance modulation. In *Proceedings of the IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5, 2025.

Antoine Guédon and Vincent Lepetit. SuGaR: Surface-aligned gaussian splatting for efficient 3D mesh reconstruction and high-quality mesh rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5354–5363, 2024.

Andrew Y. Guo, Anagh Malik, SaiKiran Tedla, Yutong Dai, Yiqian Qin, Zach Salehe, Benjamin Attal, Sotiris Nousias, Kiriakos N. Kutulakos, and David B. Lindell. Dark3R: Learning structure from motion in the dark. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2026.

Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1780–1789, 2020.

Binbin Huang, Zehao Yu, Anpei Chen, Andreas Geiger, and Shenghua Gao. 2D gaussian splatting for geometrically accurate radiance fields. In *ACM SIGGRAPH Conference Papers*, pp. 32:1–32:11, 2024.

Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. EnlightenGAN: Deep light enhancement without paired supervision. *IEEE Transactions on Image Processing*, 30:2340–2349, 2021.

Xin Jin, Pengyi Jiao, Zheng-Peng Duan, Xingchao Yang, Chun-Le Guo, Bo Ren, and Chongyi Li. 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.

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3D gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):139:1–139:14, 2023.

Chen-Hsuan Lin, Wei-Chiu Ma, Antonio Torralba, and Simon Lucey. BARF: Bundle-adjusting neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 5741–5751, 2021.

Shuhong Liu, Chenyu Bao, Ziteng Cui, Yun Liu, Xuangeng Chu, Lin Gu, Marcos V. Conde, Ryo Umagami, Tomohiro Hashimoto, Zijian Hu, Tianhan Xu, Yuan Gan, Yusuke Kurose, and Tatsuya Harada. RealX3D: A physically-degraded 3D benchmark for multi-view visual restoration and reconstruction. *International Journal of Computer Vision*, 2026. doi: 10.1007/s11263-026-02963-4.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 20654–20664, 2024.

Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5637–5646, 2022.

Andreas Meuleman, Linus Franke, Boris Zhestiankin, Camille Montemagni, and George Drettakis. Immediate 3D gaussian splat reconstruction of unordered input with global consistency. In *ACM SIGGRAPH Conference Papers*, 2026.

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 (ECCV)*, pp. 405–421, 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 (CVPR)*, 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.

Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5459–5469, 2022.

Xu Wang, Langren Xie, Zhiye Tang, Qiudan Zhang, Longhao Zou, You Yang, and Wenhui Wu. RetinexGS: Enhancing 3D gaussian splatting for low-light scenes via retinex-guided decomposition. *IEEE Transactions on Multimedia*, pp. 1–13, 2026. doi: 10.1109/TMM.2026.3705521.

Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex C. Kot. Low-light image enhancement with normalizing flow. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 2604–2612, 2022.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.

Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *Proceedings of the British Machine Vision Conference (BMVC)*, 2018.

Wenhui Wu, Jian Weng, Pingping Zhang, Xu Wang, Wenhan Yang, and Jianmin Jiang. URetinex-Net: Retinex-based deep unfolding network for low-light image enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5901–5910, 2022.

Haojie Yan, Zehao Chen, Yan Liu, Shi Gu, Peng Lin, De Ma, Huajin Tang, Qian Zheng, and Gang Pan. eRetinexGS: Retinex modeling for low-light scene enhancement via event streams and 3D gaussian splatting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8056–8066, 2026.

Chong Ye, Martin Danelljan, Christian Háne, Lin Wang, and Dengxin Liu. Gaussian in the dark: Real-time view synthesis from inconsistent dark images using gaussian splatting. *Computer Graphics Forum*, 43(7):e15213, 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 (ICCV)*, pp. 5752–5761, 2021.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 586–595, 2018.

Tianyi Zhang, Kaining Huang, Weiming Zhi, and Matthew Johnson-Roberson. DarkGS: Learning neural illumination and 3D gaussians relighting for robotic exploration in the dark. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2024.

Han Zhou, Wei Dong, and Jun Chen. LITA-GS: Illumination-agnostic novel view synthesis via reference-free 3D gaussian splatting and physical priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 21580–21589, 2025.

APPENDIX

A IMPLEMENTATION DETAILS

This appendix records implementation details omitted from the main text. The curve/LUT branch predicts view-dependent curve parameters from the low-light input and applies an adaptive stretch using the 95th-percentile brightness. The dark-render brightening branch takes the current dark rendering as input, so its signal remains tied to the reconstructed 3D scene. The per-Gaussian gain is parameterized by a scalar z_k and bounded by $G_k = 1 + \alpha_g \tanh(z_k)$, with a gain regularizer to avoid local amplification artifacts. The bright branch optimizes only the SH0/DC color coefficient $\mathbf{f}^b_{k,0}$ while sharing geometry, opacity, covariance, and higher-order SH terms with the dark branch. During the bright-appearance optimization stage, camera poses, geometry, dark appearance, signal-generation branches, and gain parameters are frozen; only the bright low-order color coefficients are optimized.

B TRAINING ALGORITHM

The training procedure consists of three stages:

1. **Dark-anchored reconstruction.** Process low-light views with the immediate reconstruction backbone, estimate cameras, initialize and densify Gaussians, and optimize the dark branch with $\mathcal{L}_{low}$.
2. **Bright-signal formation.** Construct the curve/LUT pseudo-bright target $\tilde{P}^b_i$, brighten the dark rendering to obtain E^b_i, rasterize the per-Gaussian gain map M^G_i, apply exposure lock to form Y^b_i, and regularize it with multi-view consistency.
3. **Bright appearance learning.** Freeze geometry, poses, dark appearance, and signal branches; use $T^b_i = \mathrm{sg}(Y^b_i)$ as the target; optimize the bright SH0 coefficients through $\mathcal{L}_{learn}$; render $\mathcal{G}^b$ for held-out test views.

This procedure keeps low-light training inputs and normal-exposure evaluation images separated throughout optimization.

Table 4: Per-scene quantitative results on LOM. Metrics are reported separately to avoid compressing multiple numbers into one cell.

Method	Metric	bike	buu	chair	shrub	sofa
3DGS	PSNR↑	6.36	7.37	5.98	7.91	6.10
	SSIM↑	0.070	0.238	0.144	0.030	0.201
	LPIPS↓	0.564	0.525	0.626	0.709	0.593
	Time	2:00	2:06	2:04	1:53	2:02
i3DGS	PSNR↑	6.39	7.51	5.92	7.99	6.16
	SSIM↑	0.073	0.247	0.131	0.029	0.199
	LPIPS↓	0.586	0.397	0.675	0.674	0.563
	Time	**0:03**	**0:04**	**0:04**	**0:05**	**0:03**
LLGS	PSNR↑	19.62	21.49	21.05	15.72	20.82
	SSIM↑	0.766	0.884	0.826	0.636	0.805
	LPIPS↓	0.371	0.219	0.383	**0.265**	0.414
	Time	3:33	3:47	3:32	5:49	3:29
Luminance-GS	PSNR↑	17.71	23.53	17.99	11.22	23.00
	SSIM↑	0.716	**0.908**	0.828	0.578	0.890
	LPIPS↓	0.464	**0.166**	0.333	0.272	0.317
	Time	7:58	6:32	6:36	7:02	6:22
LITA-GS	PSNR↑	22.20	19.79	21.04	18.72	19.03
	SSIM↑	0.774	0.865	0.843	0.630	0.850
	LPIPS↓	0.365	0.202	0.282	0.401	0.268
	Time	32:43	21:43	22:37	48:27	11:35
RetinexGS	PSNR↑	23.01	22.33	20.62	18.14	21.37
	SSIM↑	**0.810**	0.899	0.867	0.630	0.908
	LPIPS↓	**0.297**	0.178	**0.276**	0.284	**0.249**
	Time	7:51	10:13	8:24	26:15	7:45
iBright-GS	PSNR↑	**23.04**	**25.34**	**23.91**	**19.26**	**26.59**
	SSIM↑	0.776	0.900	**0.871**	**0.637**	**0.909**
	LPIPS↓	0.359	0.201	0.334	0.281	0.267
	Time	0:15	0:19	0:28	0:32	0:15

Table 5: Per-scene quantitative results on RealX3D. Metrics are reported separately to keep scene-level variation readable.

Method	Metric	Blue.	Choc.	Cup.	Gear.	Lab.	Milk.	Pop.	Scu.	Ujk.	Avg
3DGS	PSNR↑	7.31	8.02	4.37	6.76	7.45	5.65	7.46	6.04	6.31	6.60
	SSIM↑	0.046	0.085	0.062	0.088	0.046	0.065	0.060	0.023	0.073	0.061
	LPIPS↓	0.713	0.684	0.696	0.638	0.568	0.628	0.722	0.751	0.818	0.691
	Time	2:53	2:44	3:21	3:24	3:15	4:09	2:52	3:38	3:17	3:17
i3DGS	PSNR↑	7.30	8.10	4.39	6.61	7.40	5.88	7.44	6.02	6.29	6.60
	SSIM↑	0.041	0.085	0.060	0.077	0.033	0.079	0.056	0.022	0.055	0.056
	LPIPS↓	0.719	0.683	0.687	0.649	0.581	0.615	0.731	0.749	0.828	0.694
	Time	**0:05**	**0:08**	**0:09**	**0:04**	**0:07**	**0:05**	**0:04**	**0:11**	**0:07**	**0:07**
LLGS	PSNR↑	13.62	8.26	16.99	7.66	10.39	15.80	12.48	12.00	13.95	12.35
	SSIM↑	0.595	0.369	0.731	0.192	0.468	0.697	0.472	0.269	0.641	0.493
	LPIPS↓	0.618	0.784	0.382	0.612	0.623	0.581	0.643	0.684	0.555	0.609
	Time	4:29	5:11	5:53	7:52	6:21	6:15	5:08	9:05	5:11	6:09
Luminance-GS	PSNR↑	12.28	14.97	18.91	16.18	14.45	11.12	12.20	13.57	17.19	14.54
	SSIM↑	0.595	0.322	0.787	0.503	0.352	0.647	0.503	0.350	0.564	0.514
	LPIPS↓	0.535	0.551	0.319	0.549	0.485	0.483	0.443	**0.542**	0.493	0.489
	Time	19:08	15:02	11:58	17:27	25:55	12:34	19:22	30:38	14:00	18:27
LITA-GS	PSNR↑	17.73	16.91	12.87	10.87	16.88	12.39	18.81	12.57	18.93	15.33
	SSIM↑	0.625	0.476	0.619	0.345	0.548	0.527	0.542	0.267	0.637	0.510
	LPIPS↓	0.518	0.546	0.327	0.511	0.456	0.457	0.432	0.589	0.483	0.480
	Time	81:10	81:13	83:11	83:13	81:23	81:28	79:42	80:55	82:00	81:35
RetinexGS	PSNR↑	14.80	12.62	12.30	11.68	13.19	12.93	12.94	11.20	15.71	13.04
	SSIM↑	0.586	0.437	0.636	0.492	0.509	0.659	0.383	0.306	0.610	0.513
	LPIPS↓	0.492	0.573	**0.299**	**0.508**	**0.403**	**0.410**	0.578	0.574	**0.458**	0.477
	Time	15:32	16:52	20:11	20:41	25:15	19:47	14:01	28:09	19:41	20:01
iBright-GS	PSNR↑	**22.16**	**18.35**	**19.59**	**16.55**	**19.20**	**17.94**	**19.91**	**14.38**	**20.81**	**18.77**
	SSIM↑	**0.742**	**0.518**	**0.788**	**0.648**	**0.625**	**0.758**	**0.633**	**0.355**	**0.744**	**0.646**
	LPIPS↓	**0.467**	**0.522**	0.332	0.521	0.488	0.503	**0.413**	0.565	0.465	**0.475**
	Time	1:14	1:35	1:41	1:21	1:38	1:13	1:16	2:14	1:30	1:31

C PER-SCENE QUANTITATIVE RESULTS

Tables 4 and 5 provide per-scene results. The averages in the main text are computed from these scene-level measurements. We keep LOM and RealX3D as separate tables because the scene counts differ substantially.

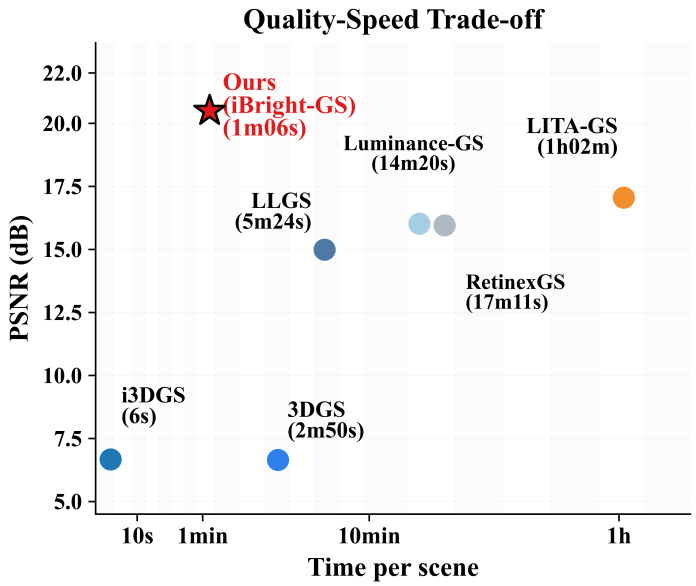

Figure 7: Quality–speed tradeoff averaged over LOM and RealX3D. Marker positions use per-scene wall-clock time (log scale) and average PSNR. iBright-GS lies near the immediate-reconstruction regime while leaving the dark-fitting cluster.

Table 6: Auxiliary ablation diagnostics on LOM and RealX3D. Each row removes one module and reports the corresponding diagnostic metric.

Removed module	Diagnostic	LOM			RealX3D		
		Full	Removed	Change	Full	Removed	Change
G_i	$\mathrm{PSNR}_{dark}\uparrow$	23.098	22.798	-0.300	19.880	19.803	-0.077
	$10^3 r_{sat}\downarrow$	0.006	0.348	+0.342	5.432	5.656	+0.224
exposure lock	$\mathrm{PSNR}_{dark}\uparrow$	23.098	23.109	+0.011	19.880	18.985	-0.895
	$10^3\sigma_{lum}\downarrow$	0.877	1.174	+0.297	9.863	9.553	-0.310
	$10^3 D_{lum}\downarrow$	1.175	1.520	+0.345	12.741	11.488	-1.253
	$10^3 r_{sat}\downarrow$	0.006	1.792	+1.786	5.432	18.118	+12.686
$\mathcal{L}_{mv}$	$10^3 D_{lum}\downarrow$	1.175	0.758	-0.417	12.741	12.656	-0.085
	$D_{hist}\downarrow$	0.313	0.319	+0.006	0.266	0.270	+0.004
structure-guided densification	$10^{-3} N_G$	88.274	69.594	-18.680	717.350	392.974	-324.376

D Quality–Speed Tradeoff

Figure 7 places dataset-average PSNR against per-scene wall-clock time. Offline low-light 3DGS methods occupy the slow high-quality or slow mid-quality region, while direct dark fitting (3DGS/i3DGS) stays near 6–7 dB. iBright-GS remains close to the immediate regime of i3DGS yet reaches the highest average PSNR among all compared methods, summarizing the quality–speed advantage reported in the main-text comparison table.

E Additional Ablations

The main text focuses on the modules that most directly affect final bright Gaussian rendering. This appendix reports auxiliary diagnostics for per-Gaussian gain, exposure lock, multi-view consistency, and structure-guided densification. We use PSNR_{dark} for the darkest 20% GT pixels, r_{sat} for the ratio of pixels brighter than 0.98, σ_{lum} for view-level luminance standard deviation, D_{lum} for pairwise view luminance difference, D_{hist} for pairwise luminance histogram distance, and N_G for the number of reconstructed Gaussians. Table 6 reports the same diagnostics on both LOM and RealX3D.

Table 7: Sensitivity to bright-signal loss weights on LOM and RealX3D. Bold marks the default configuration.

Parameter	Value	LOM			RealX3D		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
w_{enh}	0.01	23.77	0.808	0.286	18.77	0.642	0.477
	0.1	23.48	0.809	0.289	18.72	0.641	0.477
	1.0	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	10	22.95	0.805	0.296	18.58	0.642	0.490
w_{bright}	0.025	23.01	0.805	0.298	18.67	0.643	0.492
	0.25	23.03	0.806	0.297	18.65	0.645	0.490
	2.5	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	25	23.44	0.808	0.290	18.69	0.643	0.476
w_T	0.025	23.41	0.807	0.293	18.53	0.633	0.484
	0.25	23.45	0.808	0.294	18.70	0.641	0.477
	2.5	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	25	23.40	0.807	0.293	18.71	0.643	0.478

Table 8: Sensitivity to auxiliary regularization weights on LOM and RealX3D. Bold marks the default configuration.

Parameter	Value	LOM			RealX3D		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
w_{hist}	0.005	23.42	0.807	0.292	18.72	0.644	0.476
	0.05	23.35	0.809	0.292	18.66	0.636	0.476
	0.5	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	5	23.33	0.808	0.293	18.75	0.642	0.478
w_{spa}	0.0005	23.49	0.808	0.290	18.72	0.642	0.476
	0.005	23.27	0.808	0.292	18.75	0.643	0.477
	0.05	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	0.5	23.39	0.808	0.290	18.74	0.642	0.478
w_{gain}	0.0002	23.41	0.809	0.290	18.78	0.643	0.476
	0.002	23.37	0.808	0.292	18.74	0.642	0.476
	0.02	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	0.2	23.43	0.807	0.293	18.77	0.643	0.477
w_{mv}	0.0005	23.48	0.809	0.291	18.76	0.643	0.476
	0.005	23.49	0.809	0.288	18.72	0.643	0.476
	0.05	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	0.5	23.43	0.809	0.291	18.76	0.642	0.477

Table 6 shows complementary roles for the auxiliary modules. Removing G_i mainly hurts local dark-region fidelity and raises saturation on both datasets (LOM PSNR_{dark} $23.098 \rightarrow 22.798$; $10^3 r_{sat}$ $0.006 \rightarrow 0.348$), confirming that per-Gaussian gain is a mild but useful exposure corrector. Exposure lock is much more critical on RealX3D: without it, $10^3 r_{sat}$ jumps from 5.432 to 18.118 and PSNR_{dark} drops by 0.895 dB, whereas LOM mainly shows increased luminance variance and pairwise luminance drift. The multi-view term $\mathcal{L}_{mv}$ changes histogram distance only slightly on both datasets, so we treat it as a stabilizer rather than a primary accuracy source. Structure-guided densification mainly controls coverage, cutting the Gaussian count from 88k to 70k on LOM and from 717k to 393k on RealX3D when removed.

F Hyperparameter Sensitivity

This appendix reports sensitivity sweeps with LOM and RealX3D averages reported side by side. Bold marks the default setting. Learning-window rows use the recorded bright-learning intervals in the sensitivity logs; learning-length rows fix the start at step 51.

Across Tables 7–9, the default setting remains strongest on both datasets. For bright-signal weights, over-large w_{enh}=10 is the clearest failure mode: LOM PSNR falls to 22.95 and RealX3D LPIPS rises to 0.490, showing that an overly strong curve/LUT target pulls the teacher away from the 3D-bound dark render. Too-small w_{bright} or w_T also weakens both datasets, with RealX3D more sensitive to a weak teacher weight (w_T=0.025 drops RealX3D PSNR to 18.53 / SSIM 0.633). Auxiliary weights in Table 8 are comparatively flat around the default, and none of the tested settings overturn the ranking of the default recipe.

Table 9: Sensitivity to curve scale, bright-learning window, and learning length on LOM and RealX3D. Bold marks the default configuration.

Parameter	Value	LOM			RealX3D		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
curve scale	0.001	23.49	0.808	0.288	18.76	0.642	0.476
	0.01	23.26	0.807	0.292	18.80	0.644	0.476
	0.1	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	1.0	23.54	0.807	0.287	18.75	0.643	0.476
learning window	25 to 75	23.40	0.807	0.291	18.50	0.640	0.483
	50 to 100	23.32	0.807	0.291	18.44	0.643	0.483
	51 to 80	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	75 to 125	23.43	0.806	0.291	18.50	0.642	0.485
	100 to 150	23.37	0.807	0.290	18.39	0.639	0.485
learning length	14 (51 to 65)	23.45	0.808	0.290	18.61	0.639	0.479
	29 (51 to 80)	**23.63**	**0.819**	**0.288**	**18.77**	**0.646**	**0.475**
	44 (51 to 95)	23.37	0.808	0.291	18.59	0.643	0.480
	59 (51 to 110)	23.10	0.805	0.291	18.32	0.640	0.488
	74 (51 to 125)	23.07	0.803	0.292	18.03	0.637	0.499

The schedule results in Table 9 support a finite bright-learning stage after geometry is stable. Shifting the learning window earlier or later than 51–80 consistently lowers RealX3D PSNR (18.39–18.50 vs. 18.77), while LOM stays closer but still peaks at the default. With the start fixed at 51, shortening the window to 14 steps underfits the bright branch, and stretching it to 59/74 steps further hurts RealX3D (18.32 / 18.03 PSNR) and raises LPIPS, indicating overfitting to the training-time teacher. These trends match the main-text sensitivity figure and motivate keeping a bounded bright-appearance stage after the dark-anchored scaffold has settled.

G Evaluation Protocol and Reproducibility

All methods are evaluated with the same 8-split protocol: every eighth view is held out for testing, and the remaining low-light views are used for training. Normal-exposure images are used only for metric computation. Rendered images are resized to the GT resolution before PSNR, SSIM, and LPIPS are computed. LPIPS inputs are normalized consistently across methods. Running time is measured as wall-clock time per scene on the same server with 8 NVIDIA GeForce RTX 4090 GPUs. Different baselines require different geometry or pose inputs according to their public implementations; the main comparison therefore focuses on low-light appearance reconstruction quality under each method's standard setup.