

---

# Luminance-Aware Statistical Quantization: Unsupervised Hierarchical Learning for Illumination Enhancement

---

Derong Kong<sup>1</sup> Zhixiong Yang<sup>1</sup> Shengxi Li<sup>2</sup> Shuaifeng Zhi<sup>1</sup>  
Li Liu<sup>1</sup> Zhen Liu<sup>1</sup> Jingyuan Xia<sup>1†</sup>

<sup>1</sup>College of Electronic Science and Technology, National University of Defense Technology

<sup>2</sup>College of Electronic and Information Engineering, Beihang University

## Abstract

Low-light image enhancement (LLIE) faces persistent challenges in balancing reconstruction fidelity with cross-scenario generalization. While existing methods predominantly focus on deterministic pixel-level mappings between paired low/normal-light images, they often neglect the continuous physical process of luminance transitions in real-world environments, leading to performance drop when normal-light references are unavailable. Inspired by empirical analysis of natural luminance dynamics revealing power-law distributed intensity transitions, this paper introduces Luminance-Aware Statistical Quantification (LASQ), a novel framework that reformulates LLIE as a statistical sampling process over hierarchical luminance distributions. Our LASQ re-conceptualizes luminance transition as a power-law distribution in intensity coordinate space that can be approximated by stratified power functions, therefore, replacing deterministic mappings with probabilistic sampling over continuous luminance layers. A diffusion forward process is designed to autonomously discover optimal transition paths between luminance layers, achieving unsupervised distribution emulation without normal-light references. In this way, it considerably improves the performance in practical situations, enabling more adaptable and versatile light restoration. This framework is also readily applicable to cases with normal-light references, where it achieves superior performance on domain-specific datasets alongside better generalization-ability across non-reference datasets. The code is available at: <https://github.com/XYLGroup/LASQ>.

## 1 Introduction

In low-light environments, images frequently experience degradations like reduced visibility and heightened noise, which hinder subsequent vision-related tasks (1; 2; 3; 4). Low-light image enhancement (LLIE) aims to reconstruct perceptually natural scenes by establishing mappings between low-light and normal-light distributions. However, this problem is fundamentally ill-posed, as natural luminance transitions follow continuous physical processes governed by scene radiance and sensor responses, rather than discrete pixel-level correspondences.

While recent deep learning methods—whether supervised (5; 6; 7; 8; 9) or unsupervised (10; 11; 12)—attempt to model light variations through paired or unpaired training, they inherently overfit to static relationships between low/normal-light domains. Supervised methods rely on pixel-level correspondences in paired data, forcing models to prioritize localized correlations over the physics of gradual luminance evolution. Unpaired approaches, though avoiding direct pairing, still depend

---

Derong Kong and Zhixiong Yang contributed equally to this work (<sup>†</sup> Corresponding author: Jingyuan Xia).

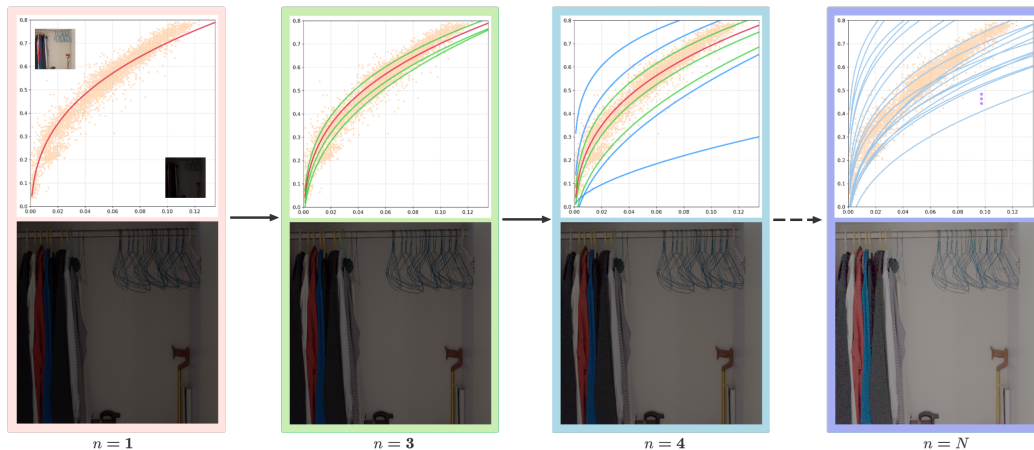


Figure 1: The physics-driven regularity of luminance intensity evolution.

on pseudo-references derived from empirical gamma corrections (10), inheriting prior biases. Both paradigms oversimplify the inherently context-dependent and continuous nature of luminance dynamics, resulting in constrained generalization: models excel in domain-specific scenarios with reference normal-light samples but struggle to adapt to unseen environments or sensor-specific degradations (13; 14). This underscores the necessity for a paradigm shift toward learning luminance transitions from the intrinsic continuity of real-world illumination processes.

This work is motivated by an empirical revelation: natural luminance transitions between low-light and normal-light conditions inherently adhere to power-law density distributions across hierarchical intensity coordinates, as visualized in Fig. 1. Unlike the black-box mappings learned by existing methods, these distributions reveal a physics-driven regularity—pixel intensities evolve along stratified luminance layers governed by cumulative power functions (Fig. 1). Specifically, each layer corresponds to a distinct power-law parameter that dictates how localized or global the luminance adaptation should be. For instance, a single power function approximates uniform luminance adjustment across the entire image (e.g., global gamma correction), while multiple overlapping functions capture spatially varying transitions, mimicking the interplay of scene radiance and sensor responses. By parameterizing these layers through variable-density sampling, where the number of power functions determines the granularity of luminance adaptation, we bridge the gap between pixel-level fidelity and cross-scenario generalization. Denser sampling (more functions) prioritizes localized intensity corrections resembling pixel-wise mappings, whereas sparser sampling (fewer functions) enforces smoother, physically consistent transitions across regions. This hierarchical decomposition fundamentally redefines LLIE: instead of deterministic low-to-normal mappings, we model LLIE as a statistically driven process that progressively traverses luminance layers, emulating the continuous and context-aware nature of real-world illumination dynamics.

In this instance, we propose a luminance-aware statistical framework named Luminance-Aware Statistical Quantization (LASQ) that translates the hierarchical power-law distributions of natural illumination into an adaptive statistical sampling process. We first formulate a scale-adaptive luminance intensity estimation function governed by power-law exponents, where both the base and exponent are dynamically computed from localized intensity statistics across adjustable regions of the luminance map. This function is derived from power-law regularity observed in natural scenes, enabling seamless computation for pixel-level corrections to global adjustments luminance adaptation operators. On the basis of this, we formulate a distribution space for luminance adaptation operators that spans granularity levels—ranging from coarse, scene-wide adjustments to fine-grained, region-specific refinements. A Markov Chain Monte Carlo (MCMC) sampling strategy is then designed to progressively explore this space, initiating from global equilibrium states and iteratively introducing spatially varying layers to simulate the continuum of real-world luminance transitions. The sampled operators follow a Gaussian-like distribution, where high-probability candidates correspond to

physically plausible global adaptations, while low-probability ones represent localized refinements. This naturally reflects the rarity of extreme, pixel-level corrections in natural illumination transitions.

We note that this sampling mechanism is embedded into the forward process of a diffusion model, which learns to traverse luminance layers in an unsupervised manner. By aligning the diffusion trajectory with the hierarchical granularity of luminance adjustments, our framework emulates the gradual, across-scene robust propagation of light in real environments, achieving fidelity-generalization equilibrium through statistically grounded layer-wise enhancement. The proposed LASQ framework thereby attains an optimal balance between local reconstruction precision and global robustness across diverse scenarios, eliminating the need for normal-light reference acquisition. Extensive experiments validate that LASQ, when integrated with a vanilla diffusion model, achieves state-of-the-art performance on non-reference datasets while attaining comparable performance to reference-dependent methods on normal-light benchmark datasets. Furthermore, LASQ exhibits versatile compatibility: it seamlessly adapts to scenarios where normal-light references are available, delivering superior domain-specific enhancement alongside unparalleled cross-dataset generalization capabilities.

Our main contributions are summarized as follows:

- We propose LASQ that fundamentally redefines LLIE by establishing the first physics-aware statistical model grounded in hierarchical power-law luminance transitions. This innovation bridges the gap between physical regularity modeling and data-driven learning paradigms, shifting the LLIE paradigm from deterministic pixel-wise mappings to stochastic processes governed by natural illumination statistics.
- We establish a statistical sampling on hierarchical luminance adaptation operator to emulate illumination transitions, where multi-scale power-law distributions are systematically parameterized and sampled via adaptive MCMC strategies. This enables automatic adaptation from global equilibrium adjustments to localized refinements based on scene-based brightness characteristics.
- We introduce a diffusion-driven learning architecture that systematically incorporates physical illumination priors through progressive luminance layer traversal during the forward diffusion process. This design enables unsupervised hierarchical enhancement while achieving dual-mode compatibility with both reference-based and reference-free scenarios, thereby eliminating dependency on paired reference data.
- Comprehensive experiments show that LASQ achieves i) superior performance on non-reference datasets without any reference guidance, ii) attains comparable performance to reference-based methods on reference-available benchmarks even when references are withheld, and iii) outperforms existing reference-based approaches when references are utilized.

## 2 Related Work

### 2.1 LLIE via Pixel-Level Consistency

Contemporary approaches focusing on pixel-level consistency can be categorized into three evolutionary stages (5; 15; 16; 17; 18; 19; 20; 21). Early works like LLNet (22) and Retinex-Net (23) leveraged paired datasets to train CNNs with pixel-wise losses, achieving precise local corrections but suffering from domain overfitting. Methods like EnlightenGAN (24) introduced cycle-consistency constraints, while Zero-DCE (25) used non-paired training with empirical illumination curves, both inheriting biases from heuristic priors (2). Recent diffusion models (10) enhanced flexibility through noise-to-clean transitions, and FeatEnHancer (1) proposed hierarchical feature fusion to bridge pixel-level and semantic gaps. These methods improved generalization but remained black-box mappings (3). While progressively reducing dependency on strict pixel correspondences, these methods universally prioritize pixel-wise fidelity over the continuous, context-dependent nature of luminance transitions, resulting in constrained generalization when facing unseen scenarios or sensor-specific degradations (13).

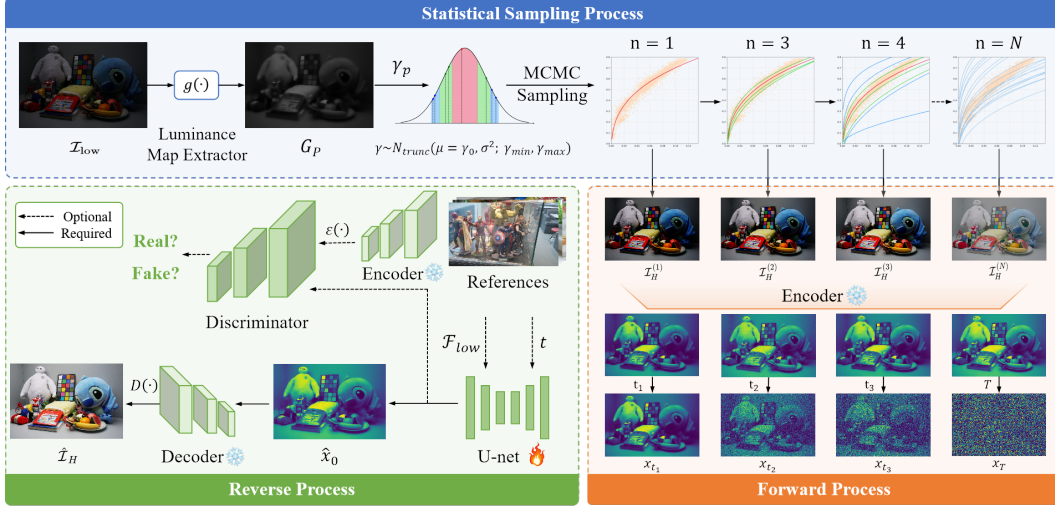


Figure 2: The framework of our LASQ.

## 2.2 LLIE with Illumination Priors

To alleviate these issues, several studies (26; 27; 25) have incorporated illumination-aware correction into deep neural networks, framing LLIE as a curve-estimation problem through different gamma correction samples. Specifically, (27) introduced learnable gamma transforms for illumination adjustment, yet enforced uniform corrections across all pixels. Methods like KinD (26) decomposed images into illumination-reflectance components, but relied on manually designed reflectance priors that oversimplified real-world radiance interactions, and its reliance on downstream pre-training objectives (e.g., perceptual losses) further introduces external biases that impair adaptability to diverse degradation patterns. More latest works like LightenDiffusion (10) integrated Retinex theory into diffusion steps, while GPP-LLIE (28) embedded illumination gradients as diffusion guidance. Though enhancing physical plausibility, these methods still imposed global correction strategies through rigid equation constraints.

Current physics-aware approaches generally ignore the hierarchical power-law distributions that dictate natural luminance changes. These methods, which depend on global gamma-like adjustments or empirical reflectance models, do not adequately represent the layered luminance levels where both localized and global adjustments interact, thereby restricting adaptability to sensor-specific issues and real-world lighting variations.

## 3 Methodology

### 3.1 Notation

The overall framework of our LASQ is illustrated in Fig. 2. At its core, we denote a low-light image by  $\mathcal{I}_L \in [0, 1]^{H \times W}$  and its normal-light counterpart by  $\mathcal{I}_N \in [0, 1]^{H \times W}$ , and for each pixel index  $i \in \{1, \dots, I\}$  with  $I = H \times W$ , we write the luminance pair  $s_i = (\mathcal{I}_L^{(i)}, \mathcal{I}_N^{(i)})$ . Besides,  $\mathcal{P}$  indicates the image region. We denote the luminance map by  $G_{\mathcal{P}}$  and define  $\gamma_{\mathcal{P}}$  as our hierarchical luminance adaptation operator.  $\mathcal{H}$  and  $\mathcal{F}$  represent the hierarchical enhanced image set and latent representations, respectively. Let  $\mathcal{E}$  and  $\mathcal{D}$  denote the encoder and the decoder.  $G_{\theta}$  and  $\mathcal{D}_{\phi}$  express the generator and the discriminator.

### 3.2 Hierarchical Luminance Modeling

**Luminance Variation Coordinate System:** In preparation for our statistical modeling of low-light image enhancement, we present a unified two-dimensional coordinate system designed to map the

relationship between "normal-light" and "low-light" luminance intensities. Let  $\mathcal{I}_L \in [0, 1]^{H \times W}$  denote an observed low-light image and  $\mathcal{I}_N \in [0, 1]^{H \times W}$  be its normal-light counterpart. For each element located at  $(x_i, y_i)$ , we write  $\mathcal{I}_L^{(i)} = \mathcal{I}_L(x_i, y_i)$  and  $\mathcal{I}_N^{(i)} = \mathcal{I}_N(x_i, y_i)$ ,  $i = 1, \dots, I$ , where  $I = H \times W$ . By treating each pair  $(\mathcal{I}_L^{(i)}, \mathcal{I}_N^{(i)})$  as a point  $s_i$  in the plane, we form the Luminance Variation (LV) coordinate system:

$$\left\{ s_i = \left( \mathcal{I}_L^{(i)}, \mathcal{I}_N^{(i)} \right) \right\}_{i=1}^I. \quad (1)$$

Empirical observation reveals that, within the normalized range  $[0, 1]$ , the low-light intensities exhibit a heavy-tailed power-law distribution, which can be approximated by a set of power-law functions

$$ax^\kappa, x \in [0, 1], \kappa > 0, a > 0, \quad (2)$$

where  $\kappa$  denotes the illumination change level by this power-law curve. As depicted in the fourth column of Fig. 1, the asymmetric distribution of underexposed elements can be approximated by a set of sampled power-law curves, providing a rigorous foundation for adjusting the low-light distribution to a specified target using mapping or sampling methods.

**Statistical Sampling Process:** Utilizing the LV coordinate system, we initially introduce the regional luminance scalar  $G_{\mathcal{P}}$  for any designated image area  $\mathcal{P} \subseteq [1, H] \times [1, W]$ . This scalar encapsulates the characteristic distribution of luminance within the region, and its precise formulation is provided in the Appendix. Utilizing Eq. (2), we proceed to compute our hierarchical luminance adaptation operator (LAO)  $\gamma_{\mathcal{P}}$  as follows:

$$\gamma_{\mathcal{P}} = (\alpha + G_{\mathcal{P}})^{\beta_{\mathcal{P}}}, \quad \beta_{\mathcal{P}} = 2G_{\mathcal{P}} - 1 + \eta \frac{\sigma_{G_{\mathcal{P}}}^2}{\sigma_{G_{\mathcal{P}}}^2 + \delta}, \quad (3)$$

where  $\alpha \in (0, 1]$ ,  $\eta, \delta$  are hyper-parameters that control adjustment strength and contrast gain. Empirical analysis reveals distinct luminance adaptation patterns: single LAO exhibits uniform global luminance modulation, while multi-LAO configurations enable region-specific refinement. This phenomenon emerges because curves in the central regime of the power-law distribution (highlighted in red in Fig. 2) demonstrate universal traversal across all LAO set cardinalities, whereas boundary regions (depicted in blue) are exclusively accessible to high-density LAO sets implementing fine-grained adjustments. We consequently model this operator distribution through a symmetrically truncated Gaussian distribution (Fig. 2, top center) as follows:

$$p(\gamma) \propto \exp\left(-\frac{1}{2}(\gamma - \gamma_0)^2 / \sigma^2\right), \quad \gamma \in [\gamma_{\min}, \gamma_{\max}], \quad (4)$$

where  $\gamma_{\min} = \min_{i,j} \gamma(i,j)$ ,  $\gamma_{\max} = \max_{i,j} \gamma(i,j)$  and  $\gamma_0 = \frac{1}{HW} \sum_{i=1}^H \sum_{j=1}^W \gamma(i,j)$ . The bounded distribution can be formally expressed as  $\gamma \sim \mathcal{N}_{\text{trunc}}(\mu = \gamma_0, \sigma^2; \gamma_{\min}, \gamma_{\max})$ .

Building upon the symmetrically truncated Gaussian distribution  $p(\gamma)$ , we devise a hierarchical Markov Chain Monte Carlo (MCMC) sampling scheme to generate LAO sets  $\Gamma = \{\Gamma_n\}_{n=1}^N$ , where each iteration  $n$  produces  $2^{n-1}$  distinct LAO configurations via adaptive chain transitions, referring to  $\Gamma_n = \left\{ \gamma_{\mathcal{P},z}^{(n)} \right\}_{z=1}^{2^{n-1}}$ . The MCMC process at the  $n$ -th iteration is given by:

$$p(\mathcal{I}_H^{(n)}) = \int p(\mathcal{I}_H^{(n)} | \Gamma_n) p(\Gamma_n) d\Gamma_n \approx \sum_{z=1}^{2^{n-1}} p(\mathcal{I}_H^{(n)} | \gamma_{\mathcal{P},z}^{(n)}) p(\gamma_{\mathcal{P},z}^{(n)}), \quad (5)$$

derived from the continuous formulation through discrete sampling approximation. Each trial constructs a Markov chain defined by the transition kernel:

$$q\left(\gamma_{\mathcal{P},z}^{(n)} \mid \gamma_{\mathcal{P},z-1}^{(n)}\right) = \mathcal{N}_{\text{trunc}}\left(\gamma_{\mathcal{P},z}^{(n)} \mid \gamma_{\mathcal{P},z-1}^{(n)}, \lambda^2; \gamma_{\min}, \gamma_{\max}\right), \quad (6)$$

where the step size  $\lambda$  adaptively balances exploration-exploitation trade-offs across hierarchy levels.

The dynamically partitioned grid strategy ensures progressive refinement: at iteration  $n$ , the image is divided into  $m_n \times w_n$  non-overlapping patches ( $m_n = 2^{\lceil (n-1)/2 \rceil}$ ,  $w_n = 2^{\lfloor (n-1)/2 \rfloor}$ ). This induces hierarchical luminance-corrected images  $\mathcal{H} = \{\mathcal{I}_H^{(n)}\}_{n=1}^N$ , where each  $\mathcal{I}_H^{(n)}$  encapsulates  $2^{n-1}$  locally optimized gamma correction patterns. Crucially, every MCMC trial synthesizes a self-consistent LAO set that traverses luminance hierarchies through state-dependent transitions, enabling coarse-to-fine representation learning where global brightness constraints guide local refinements and vice versa.

### 3.3 Hierarchically-Guided Diffusion

**Forward Process with Hierarchical Guidance:** The sampled set  $\mathcal{H} = \{\mathcal{I}_H^{(n)}\}_{n=1}^N$  employs stochastic learning via diffusion transitions, exploiting the Markov property—each  $\mathcal{I}_H^{(n)}$  relies only on its forerunner—to adaptively direct noise injection. The low-light image  $\mathcal{I}_L$  and  $\mathcal{H}$  are encoded together using  $\mathcal{E}(\cdot)$ , incorporating  $k$  residual blocks and max-pooling for latent features  $\mathcal{F}_L \in \mathbb{R}^{\frac{H}{2^k} \times \frac{W}{2^k} \times C}$  and  $\{\mathcal{F}_H^{(n)}\}_{n=1}^N \in \mathbb{R}^{N \times \frac{H}{2^k} \times \frac{W}{2^k} \times C}$ . We align the  $T$ -step diffusion with  $\{\mathcal{F}_H^{(n)}\}_{n=1}^N$  using a temporal mapping  $\psi : \{1, \dots, T\} \rightarrow \{1, \dots, N\}$ ,  $N \leq T$ , by  $\psi(t) = \lfloor t \cdot N/T \rfloor$ , such that:

$$T = \sum_{n=1}^N |T_n|, \quad t \in T_n \Rightarrow x_0 = \mathcal{F}_H^{(\psi(t))}. \quad (7)$$

The forward diffusion adds Gaussian noise progressively:

$$q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I), \quad (8)$$

where  $t \in \{1, \dots, T\}$  denotes the diffusion timestep,  $x_t$  represents the random variable at timestep  $t$ , and  $\beta_t$  denotes the noise variance. Accordingly, for each temporal interval  $T_n$ , the corresponding spatial variant  $\mathcal{F}_H^{(\psi(t))}$  is utilized as the illumination normalization reference, thereby maintaining luminance-consistent forward sampling. By incrementally incorporating spatial luminance awareness from coarse to fine scales into the diffusion forward trajectory, the model acquires a multi-level representation of illumination dynamics. This hierarchical perception facilitates adaptive noise scheduling and enhances robustness across a wide range of lighting conditions.

**Hierarchically-Guided Diffusion Denoising:** During the reverse training phase, the denoising network  $\epsilon_\theta(x_t, t, \mathcal{F}_L)$  is trained to achieve:

$$x_{t-1} = \frac{1}{\sqrt{1 - \beta_t}}(x_t - \beta_t \epsilon_\theta(x_t, t, \mathcal{F}_L)) + \sigma_t b, \quad b \sim \mathcal{N}(0, I), \quad (9)$$

where  $\sigma$  denotes the standard deviation. Based on the variational lower bound of the forward process, we minimize the mean squared error between the true noise and the network prediction, leading to the simplified noise prediction objective:

$$\mathcal{L}_d = \sum_{t=1}^T \mathbb{E}_{q(x_0, \mathcal{F}_L)} \left[ \|\epsilon - \epsilon_\theta(x_t, t, \mathcal{F}_L)\|^2 \right], \quad (10)$$

where  $\epsilon$  represents the actual injected noise. To ensure overall image smoothness and preserve fine details while minimizing generation artifacts, we employ the global label  $\mathcal{F}_H^{(\psi(0))}$  to weakly guide the reverse diffusion process:

$$\mathcal{L}_g = \left\| \mathcal{D}(\hat{x}_0) - \mathcal{D}(\mathcal{F}_H^{(\psi(0))}) \right\|_1, \quad (11)$$

where the  $\mathcal{D}(\cdot)$  denotes the decoder. During inference, under the guidance of the low-light input  $\mathcal{F}_L$ , we employ the diffusion model’s implicit sampling strategy (29) for reverse denoising. The model utilizes its learned distribution to fit the optimal illumination-enhanced feature representation  $\hat{\mathcal{F}}_N$ , which is then decoded to yield the final output  $\hat{\mathcal{I}}_N = \mathcal{D}(\hat{\mathcal{F}}_N)$ .

The LASQ framework can integrate effortlessly with normal-light references. In this setup, an optional adversarial discriminator  $\mathcal{D}_\phi$  complements the LASQ, creating a hybrid diffusion-GAN model. The generator’s training involves a combined loss:

$$\mathcal{L}_{\text{total}} = \lambda_d \mathcal{L}_d + \lambda_g \mathcal{L}_g + \underbrace{\lambda_{\text{GAN}} \mathbb{E}_{\mathcal{I}_L} [-\log \mathcal{D}_\phi(G_\theta(\mathcal{I}_L))]}_{\text{adversarial penalty}}, \quad (12)$$

where adversarial training refines high-level textures while preserving the physical grounding from diffusion priors. These results are presented with LSAQ++ in the simulation. Implementation details for reference-augmented training are provided in the Appendix.



Figure 3: Qualitative comparison of our method and competitive methods on the LOLv1 and LSRW test sets. "LASQ++" denotes the incorporation of unpaired normal-light references.

## 4 Experiments

### 4.1 Experimental Settings

All experiments are carried out on a cluster of four NVIDIA A800 GPUs under Python 3.9 and PyTorch 2.0, with a fixed batch size of 16. We employ the Adam optimizer (30), setting the denoising diffusion process learning rate to  $2 \times 10^{-5}$ , while using a sampling ratio  $k = 3$ . The hyperparameters  $\lambda_d$ ,  $\lambda_g$  and  $\lambda_{GAN}$  (if activated) are set to 0.9, 0.005 and 0.7 respectively. Noise estimation during diffusion training is performed using the U-Net (31) architecture with  $T = 1000$  time steps.

### 4.2 Datasets and Metrics

We validated our approach using both paired and unpaired low-light benchmarks. For the paired evaluation, we used the LOLv1 (22) and LSRW (32) test sets—each comprising matched low- and normal-illumination image pairs—and reported restoration fidelity via PSNR and SSIM (33), alongside the full-reference perceptual score LPIPS (34). To assess performance in the absence of ground-truth references, we then tested on the unpaired LIME (35), DICM (36), NPE (37), and VV (38) collections, measuring perceptual quality with the no-reference NIQE (39) and PI (40) metrics. Our comparative study encompasses six supervised approaches (RetinexNet (22), KinD++ (23), LCDPNet (21), URetinexNet (41), SMG (17) and PyDiff (42)) and six unsupervised approaches (Zero-DCE (25), EnlightenGAN (24), SCI (19), PairLIE (43), SCL-LLE (44), LightenDiffusion (10) and NeRCO (18)).

### 4.3 Qualitative Comparison

As shown in Fig. 3, LASQ achieves enhanced local brightness adaptation and superior detail fidelity comparable to supervised methods URetinexNet (41) and KinD++ (23) on ground truth-annotated datasets, while demonstrating improved domain-adaptive color reproduction through integration with unpaired normal-light images in LASQ++. By contrast, existing methods exhibit distinct limitations: SCI (19) and SCL-LLE (44) suffer from persistent underexposure, whereas EnlightenGAN (24) produces blurred structural details, and NeRCO (18) tends to generate localized over-exposure artifacts. Furthermore, Fig. 4 confirms LASQ’s exceptional performance in real-world scenarios through its complete avoidance of local overexposure, noise amplification, and artifacts that persistently affect other methods: EnlightenGAN (24), NeRCO (18) and PairLIE (43) notably show severe localized overexposure and lens flare artifacts, and even supervised approaches URetinexNet (41) and KinD++ (23) still struggle to fully suppress these issues. Crucially, LASQ maintains natural



Figure 4: Qualitative comparison of our method and competitive methods on the LIME, and VV datasets. More results will be provided in the Appendix.

scene characteristics without compromising detail fidelity or color consistency, thereby demonstrating unprecedented cross-scenario generalization capability across both constrained laboratory settings and unconstrained environmental conditions. This comprehensive evaluation systematically validates LASQ’s technical superiority in terms of adaptive illumination control, artifact suppression, and domain transfer effectiveness. More results will be provided in the Appendix.

#### 4.4 Quantitative Comparison

The quantitative evaluation results across diverse datasets are summarized in Table 1, where LASQ demonstrates performance parity with leading supervised techniques on LOLv1 (22) and LSRW (32) while achieving state-of-the-art results among unsupervised methods through integration of unpaired normal-light images. Notably, on datasets DICM (36), NPE (37), and VV (38), LASQ outperforms existing approaches across most perceptual metrics, thereby confirming its intrinsic generalization prowess without domain-specific adaptation. Although normal-light reference integration improves color fidelity, its tendency toward overfitting partially counteracts the model’s inherent generalization capacity, resulting in metric degradation in LASQ++. Crucially, this performance highlights LASQ’s fundamental advantage in balancing domain adaptation with cross-scenario robustness, whereas LASQ++ prioritizes target-domain color accuracy at the expense of slight generalization capability. The extended experimental results, including comprehensive qualitative analyses and quantitative evaluations, are provided in the Appendix.

#### 4.5 Computational Cost

We show the computational complexity metrics in Table 2 (NVIDIA A800, LOLv1 dataset). The coarse-to-fine MCMC sampling mechanism is only used during training and is embedded into the forward diffusion process. It guides the model to traverse luminance layers in a hierarchical manner, enabling structured learning of light propagation. During inference, our model only performs the denoising step conditioned on the low-light input within the diffusion model, which is significantly more efficient.

We compare LASQ against both early non-diffusion-based methods (e.g., EnlightenGAN, KinD++), and recent diffusion-based approaches (e.g., WCDM, LightenDiffusion). While the early methods are lightweight, their performance lags far behind diffusion-based models across all key metrics. Existing diffusion models, although significantly more effective, tend to suffer from high computational cost due to deep architectures and iterative sampling. LASQ, while maintaining the performance

Table 1: The quantitative comparison results of partial experiments, with the best-performing results marked in red and the second-best in blue. The notations "SL" and "UL" respectively represent supervised and unsupervised learning approaches. "LASQ++" denotes the incorporation of unpaired normal-light references

Type	Method	LOLv1			LSRW			DICM		NPE		VV	
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	NIQE↓	PI↓	NIQE↓	PI↓	NIQE↓	PI↓
SL	RetinexNet	16.774	0.462	0.390	15.609	0.414	0.393	4.487	3.242	4.732	3.219	5.881	3.727
	KinD++	17.752	0.758	0.198	16.085	0.394	0.366	4.027	3.999	4.005	3.144	3.586	2.773
	LCDPNet	14.506	0.575	0.312	15.689	0.474	0.344	4.110	3.250	4.126	3.127	5.039	3.347
	URetinexNet	19.842	<b>0.824</b>	<b>0.128</b>	18.271	0.518	<b>0.295</b>	4.774	3.565	4.028	3.153	3.851	2.891
	SMG	<b>23.814</b>	0.809	0.144	17.579	0.538	0.456	6.224	4.228	5.300	3.627	5.752	3.757
	PyDiff	<b>23.275</b>	<b>0.859</b>	<b>0.108</b>	17.264	0.510	0.335	4.499	3.792	4.082	3.268	4.360	3.678
UL	Zero-DCE	14.861	0.562	0.330	15.867	0.443	0.315	3.951	3.149	3.826	2.918	5.080	3.307
	EnGAN	17.606	0.653	0.319	17.106	0.463	0.322	3.832	3.256	3.775	2.953	3.689	2.749
	SCI	14.784	0.525	0.333	15.242	0.419	0.321	4.519	3.700	4.124	3.534	5.312	3.648
	PairLIE	19.514	0.731	0.254	17.602	0.501	0.323	4.282	3.469	4.661	3.543	3.373	2.734
	SCL-LLE	10.754	0.506	0.382	13.110	0.310	0.396	5.129	3.809	4.873	3.692	5.513	4.316
	NeRCo	19.738	0.740	0.239	17.844	0.535	0.371	4.107	3.345	3.902	3.037	3.765	3.094
	LigDiff	20.453	0.803	0.192	<b>18.555</b>	0.539	0.311	3.724	3.144	3.618	2.879	2.941	<b>2.558</b>
	LASQ	20.375	0.814	0.191	18.137	<b>0.547</b>	<b>0.308</b>	<b>3.715</b>	<b>3.128</b>	<b>3.571</b>	<b>2.764</b>	<b>2.777</b>	<b>2.623</b>
	LASQ++	20.481	0.807	0.205	<b>18.584</b>	<b>0.540</b>	0.316	<b>3.723</b>	<b>3.137</b>	<b>3.601</b>	<b>2.789</b>	<b>2.850</b>	2.691

Table 2: Comparison of computational efficiency and resource usage.

Method	FLOPs (G)	Params (M)	Inference Time (ms)	Memory Usage (MB)
EnlightenGAN	16.45	8.64	70.16	241.48
KinD++	17.49	8.27	4279.70	372.19
NeRCo	184.20	23.30	354.77	2320.87
PairLIE	81.84	0.34	900.70	3499.79
SCI	0.13	-	50.14	20.01
SCL-LLE	19.01	0.08	60.59	324.31
URetinex	81.35	0.34	129.70	568.48
WCDM	374.47	22.92	206.66	6017.86
LightenDiffusion	367.99	27.83	257.94	8049.95
LASQ	219.75	24.08	213.89	6496.68

advantages of diffusion models, achieves inference efficiency comparable to non-diffusion-based methods. This makes it highly suitable for real-world deployment. Considering the substantial performance gain without reliance on reference images, the moderate computational overhead of LASQ is a practical and acceptable trade-off.

## 5 Ablation

### 5.1 Fixed Luminance Adjustment

We replace our adaptive MCMC-based luminance adaptation with static, hand-crafted functions (e.g., global gamma correction or fixed tone curves) drawn from prior LLIE methods (45; 46; 47; 48). Embedding these into the forward diffusion degrades PSNR, SSIM, and LPIPS (Table 3), showing reduced feature richness and generative fidelity (Fig. 5). In contrast, our physics-driven, multi-scale operator—via power-law stratification and adaptive sampling—better balances global consistency and local detail, yielding stronger cross-scenario generalization and perceptual quality.

### 5.2 Limited Hierarchy

We keep adaptive MCMC sampling but limit the operator to two layers: a global adjustment and per-pixel correction, omitting mid-level power-law strata. This two-layer variant still beats the fixed baseline, yet its PSNR, SSIM, and perceptual metrics fall short of the full LASQ (Table 3), confirming that intermediate layers are essential for smoothly refining illumination and preserving both quantitative performance and visual fidelity.

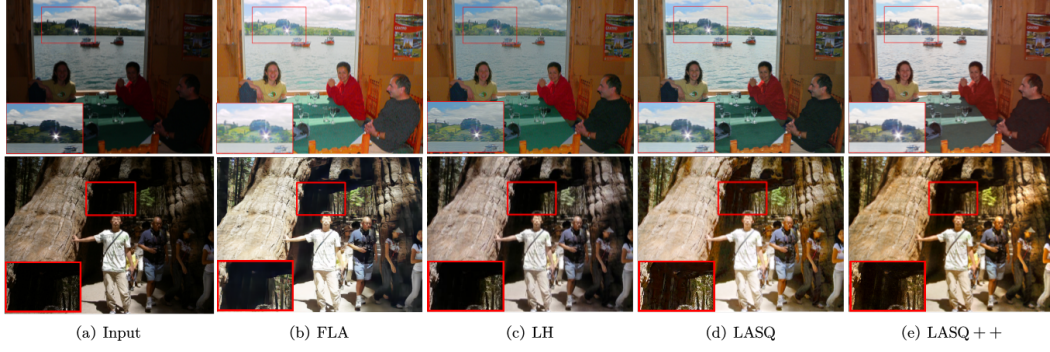


Figure 5: The qualitative results of ablation studies.

Table 3: Quantitative results of ablation studies. The "FLA" and "LH" respectively represent the ablation of Fixed Luminance Adjustment and Limited Hierarchy.

Method	LOLv1			LSRW			DICM		NPE		VV	
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	NIQE $\downarrow$	PI $\downarrow$	NIQE $\downarrow$	PI $\downarrow$	NIQE $\downarrow$	PI $\downarrow$
FLA	16.741	0.715	0.273	15.490	0.508	0.399	4.265	3.529	3.937	3.114	3.683	3.007
LH	19.139	0.792	0.243	18.026	0.522	0.333	3.759	3.396	3.648	2.996	3.006	2.730
LASQ	20.375	0.814	0.191	18.137	0.547	0.308	3.715	3.128	3.571	2.764	2.777	2.623
LASQ++	20.481	0.807	0.205	18.584	0.540	0.316	3.723	3.137	3.601	2.789	2.850	2.691

### 5.3 Hyperparameter Sensitivity

As illustrated in the table, we systematically varied key hyperparameters including  $\alpha$ ,  $\eta$ ,  $\lambda_d$ , and  $\lambda_g$  over a range of values ( $\beta_P$  is determined by  $\eta$ ). The results show that while performance slightly fluctuates with different settings, the overall impact on metrics remains moderate. For instance, varying between 0.05 and 0.6 only causes a minor PSNR change (within 0.3 dB) and negligible shifts in perceptual scores. Similarly, other hyperparameters demonstrate a stable trend without sharp degradation. These experiments demonstrate that our method is not overly sensitive to hyperparameter selection, and maintains consistently strong performance across a broad range of settings—highlighting its robustness, practical stability, and generalization potential.

Table 4: Ablation study on key hyperparameters. The best results for are highlighted in bold.

Param	Value	PSNR $\uparrow$	LPIPS $\downarrow$	SSIM $\uparrow$
$\alpha$	0.05 / 0.15 / 0.3 / 0.6	17.81 / <b>18.10</b> / 17.92 / 17.84	<b>0.319</b> / 0.322 / 0.320 / 0.324	0.512 / <b>0.543</b> / 0.530 / 0.519
$\eta$	0.1 / 1.0 / 3.0 / 6.0	17.85 / <b>18.35</b> / 18.17 / 17.95	0.335 / <b>0.321</b> / 0.324 / 0.329	0.537 / 0.543 / <b>0.546</b> / 0.540
$\lambda_d$	0.1 / 1.0 / 10 / 20	17.82 / <b>18.04</b> / 17.85 / 17.87	0.324 / <b>0.315</b> / 0.318 / 0.323	0.545 / <b>0.553</b> / 0.549 / 0.531
$\lambda_g$	0.001 / 0.005 / 0.01 / 0.1	17.76 / <b>18.22</b> / 18.16 / 17.88	0.312 / 0.310 / <b>0.309</b> / 0.311	0.536 / <b>0.548</b> / 0.547 / 0.540

## 6 Conclusion

The proposed LASQ framework reorients the LLIE challenge by merging illumination continuity physics with deep-learning, moving beyond conventional pixel-level methods. It redefines LLIE as a continuous stochastic task using adaptive MCMC, capturing real illumination through layered luminance analysis. Transitioning from reliance on paired data to unsupervised layer exploration, it improves generalization and can support reference contexts. It balances overall illumination with local detail, effectively resolving fidelity vs. adaptability issues, while avoiding gamma-related biases and overfitting of supervised methods. These advancements underscore a broader insight: LLIE must evolve from pixel-wise function approximation to spatiotemporal reconstruction of light dynamics. Future work should explore dynamic power-law parameterization for time-varying scenes and hardware-software co-design to align statistical priors with sensor-specific noise profiles. Narrowing these gaps will enhance computational imaging systems to better mimic biological vision in low light.

## Acknowledgments and Disclosure of Funding

This work is supported by the National Natural Science Foundation of China under Grant 62576350, 62131020, 62376283 and 62531026.

## References

- [1] L. Ye and Z. Ma, “Llod: a object detection method under low-light condition by feature enhancement and fusion,” in *2023 4th international seminar on artificial intelligence, networking and information technology (AINIT)*. IEEE, 2023, pp. 659–662.
- [2] Z. Li, X. Li, Y. Niu, C. Rong, and Y. Wang, “Infrared and visible light fusion for object detection with low-light enhancement,” in *2024 IEEE 7th International Conference on Information Systems and Computer Aided Education (ICISCAE)*. IEEE, 2024, pp. 120–124.
- [3] L. Xiong, X. Feng, L. Zhang, Q. Sun, and H. Ren, “Semantic segmentation algorithm of nighttime road surface water based on improved low-light image enhancement,” in *2024 5th International Conference on Intelligent Computing and Human-Computer Interaction (ICHCI)*. IEEE, 2024, pp. 104–108.
- [4] Z. Lu, H. Sun, L. Lei, Y. Xu, Y. Sun, and G. Kuang, “Diffdual-ad: Diffusion-based dual-stage adversarial defense framework in remote sensing with denoiser constraint,” *IEEE Transactions on Geoscience and Remote Sensing*, 2025.
- [5] X. Yi, H. Xu, H. Zhang, L. Tang, and J. Ma, “Diff-retinex: Rethinking low-light image enhancement with a generative diffusion model,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 12 302–12 311.
- [6] Y. Cai, H. Bian, J. Lin, H. Wang, R. Timofte, and Y. Zhang, “Retinexformer: One-stage retinex-based transformer for low-light image enhancement,” in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 12 504–12 513.
- [7] F. Lv, F. Lu, J. Wu, and C. Lim, “Mbllen: Low-light image/video enhancement using cnns.” in *Bmvc*, vol. 220, no. 1. Northumbria University, 2018, p. 4.
- [8] L. Shen, Z. Yue, F. Feng, Q. Chen, S. Liu, and J. Ma, “Msr-net: Low-light image enhancement using deep convolutional network,” *arXiv preprint arXiv:1711.02488*, 2017.
- [9] F. Yam and Z. Hassan, “Innovative advances in led technology,” *Microelectronics Journal*, vol. 36, no. 2, pp. 129–137, 2005.
- [10] H. Jiang, A. Luo, X. Liu, S. Han, and S. Liu, “Lightendiffusion: Unsupervised low-light image enhancement with latent-retinex diffusion models,” in *European Conference on Computer Vision*. Springer, 2024, pp. 161–179.
- [11] S. Yang, X. Zhang, Y. Wang, J. Yu, Y. Wang, and J. Zhang, “Diffle: Diffusion-guided domain calibration for unsupervised low-light image enhancement,” *arXiv preprint arXiv:2308.09279*, 2023.
- [12] Y. Shi, D. Liu, L. Zhang, Y. Tian, X. Xia, and X. Fu, “Zero-ig: zero-shot illumination-guided joint denoising and adaptive enhancement for low-light images,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2024, pp. 3015–3024.
- [13] S. Tu, W. Yang, and B. Fei, “Taming generative diffusion prior for universal blind image restoration,” *Advances in Neural Information Processing Systems*, vol. 37, pp. 21 172–21 206, 2024.
- [14] J. Hou, Z. Zhu, J. Hou, H. Liu, H. Zeng, and H. Yuan, “Global structure-aware diffusion process for low-light image enhancement,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 79 734–79 747, 2023.
- [15] Y. Wu, C. Pan, G. Wang, Y. Yang, J. Wei, C. Li, and H. T. Shen, “Learning semantic-aware knowledge guidance for low-light image enhancement,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 1662–1671.

- [16] X. Xu, R. Wang, C.-W. Fu, and J. Jia, "Snr-aware low-light image enhancement," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 17 714–17 724.
- [17] X. Xu, R. Wang, and J. Lu, "Low-light image enhancement via structure modeling and guidance," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2023, pp. 9893–9903.
- [18] S. Yang, M. Ding, Y. Wu, Z. Li, and J. Zhang, "Implicit neural representation for cooperative low-light image enhancement," in *Proceedings of the IEEE/CVF international conference on computer vision*, 2023, pp. 12 918–12 927.
- [19] L. Ma, T. Ma, R. Liu, X. Fan, and Z. Luo, "Toward fast, flexible, and robust low-light image enhancement," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 5637–5646.
- [20] R. Liu, L. Ma, J. Zhang, X. Fan, and Z. Luo, "Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2021, pp. 10 561–10 570.
- [21] H. Wang, K. Xu, and R. W. Lau, "Local color distributions prior for image enhancement," in *European conference on computer vision*. Springer, 2022, pp. 343–359.
- [22] C. Wei, W. Wang, W. Yang, and J. Liu, "Deep retinex decomposition for low-light enhancement," *arXiv preprint arXiv:1808.04560*, 2018.
- [23] Y. Zhang, X. Guo, J. Ma, W. Liu, and J. Zhang, "Beyond brightening low-light images," *International Journal of Computer Vision*, vol. 129, pp. 1013–1037, 2021.
- [24] Y. Jiang, X. Gong, D. Liu, Y. Cheng, C. Fang, X. Shen, J. Yang, P. Zhou, and Z. Wang, "Enlightengan: Deep light enhancement without paired supervision," *IEEE transactions on image processing*, vol. 30, pp. 2340–2349, 2021.
- [25] C. Guo, C. Li, J. Guo, C. C. Loy, J. Hou, S. Kwong, and R. Cong, "Zero-reference deep curve estimation for low-light image enhancement," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 1780–1789.
- [26] W. Wang, R. Luo, W. Yang, and J. Liu, "Unsupervised illumination adaptation for low-light vision," *IEEE Transactions on Pattern Analysis & Machine Intelligence*, no. 01, pp. 1–15, 2024.
- [27] Y. Wang, Z. Liu, J. Liu, S. Xu, and S. Liu, "Low-light image enhancement with illumination-aware gamma correction and complete image modelling network," in *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023, pp. 13 128–13 137.
- [28] H. Zhou, W. Dong, X. Liu, Y. Zhang, G. Zhai, and J. Chen, "Low-light image enhancement via generative perceptual priors," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 39, no. 10, 2025, pp. 10 752–10 760.
- [29] J. Song, C. Meng, and S. Ermon, "Denoising diffusion implicit models," *arXiv preprint arXiv:2010.02502*, 2020.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [31] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical image computing and computer-assisted intervention–MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III* 18. Springer, 2015, pp. 234–241.
- [32] J. Hai, Z. Xuan, R. Yang, Y. Hao, F. Zou, F. Lin, and S. Han, "R2rnet: Low-light image enhancement via real-low to real-normal network," *Journal of Visual Communication and Image Representation*, vol. 90, p. 103712, 2023.

- [33] Z. 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*, vol. 13, no. 4, pp. 600–612, 2004.
- [34] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 586–595.
- [35] X. Guo, Y. Li, and H. Ling, "Lime: Low-light image enhancement via illumination map estimation," *IEEE Transactions on image processing*, vol. 26, no. 2, pp. 982–993, 2016.
- [36] C. Lee, C. Lee, and C.-S. Kim, "Contrast enhancement based on layered difference representation of 2d histograms," *IEEE transactions on image processing*, vol. 22, no. 12, pp. 5372–5384, 2013.
- [37] S. Wang, J. Zheng, H.-M. Hu, and B. Li, "Naturalness preserved enhancement algorithm for non-uniform illumination images," *IEEE transactions on image processing*, vol. 22, no. 9, pp. 3538–3548, 2013.
- [38] V. Vonikakis, R. Kouskouridas, and A. Gasteratos, "On the evaluation of illumination compensation algorithms," *Multimedia Tools and Applications*, vol. 77, pp. 9211–9231, 2018.
- [39] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a "completely blind" image quality analyzer," *IEEE Signal processing letters*, vol. 20, no. 3, pp. 209–212, 2012.
- [40] Y. Blau, R. Mechrez, R. Timofte, T. Michaeli, and L. Zelnik-Manor, "The 2018 pirm challenge on perceptual image super-resolution," in *Proceedings of the European conference on computer vision (ECCV) workshops*, 2018, pp. 0–0.
- [41] W. Wu, J. Weng, P. Zhang, X. Wang, W. Yang, and J. 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*, 2022, pp. 5901–5910.
- [42] D. Zhou, Z. Yang, and Y. Yang, "Pyramid diffusion models for low-light image enhancement," *arXiv preprint arXiv:2305.10028*, 2023.
- [43] Z. Fu, Y. Yang, X. Tu, Y. Huang, X. Ding, and K.-K. Ma, "Learning a simple low-light image enhancer from paired low-light instances," in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2023, pp. 22 252–22 261.
- [44] D. Liang, L. Li, M. Wei, S. Yang, L. Zhang, W. Yang, Y. Du, and H. Zhou, "Semantically contrastive learning for low-light image enhancement," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 36, no. 2, 2022, pp. 1555–1563.
- [45] S.-C. Huang, F.-C. Cheng, and Y.-S. Chiu, "Efficient contrast enhancement using adaptive gamma correction with weighting distribution," *IEEE transactions on image processing*, vol. 22, no. 3, pp. 1032–1041, 2012.
- [46] F. Kallel, M. Sahnoun, A. Ben Hamida, and K. Chtourou, "Ct scan contrast enhancement using singular value decomposition and adaptive gamma correction," *Signal, Image and Video Processing*, vol. 12, pp. 905–913, 2018.
- [47] Z. Huang, T. Zhang, Q. Li, and H. Fang, "Adaptive gamma correction based on cumulative histogram for enhancing near-infrared images," *Infrared Physics & Technology*, vol. 79, pp. 205–215, 2016.
- [48] W. Wang, X. Wu, X. Yuan, and Z. Gao, "An experiment-based review of low-light image enhancement methods," *Ieee Access*, vol. 8, pp. 87 884–87 917, 2020.

## NeurIPS Paper Checklist

### 1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: In the abstract and introduction, we propose Luminance-Aware Statistical Quantization (LASQ) as a physics-informed statistical framework for low-light image enhancement. This framework is theoretically grounded in hierarchical power-law distributions of natural illumination (as formalized in Sec. 3) and experimentally validated through comprehensive evaluations of both reference-based and reference-free scenarios (detailed in Sec. 4).

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

### 2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: The primary limitations of LASQ lie in its relatively high computational complexity and the requirement for manual hyperparameter design. Detailed discussions of these limitations are provided in the Appendix.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

### 3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification: The variables of the proposed methodology have been formally defined in Sec. 3, while extended mathematical derivations and rigorous proofs will be documented in the Appendix.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

### 4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [\[Yes\]](#)

Justification: In implementation details, we give the parameter settings and network architecture of our LASQ in Sec. 4.1. Besides, the hyperparameter-tuning provides the choice of our hyperparameters in the Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
  - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
  - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
  - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).

- (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

## 5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: We will provide open access to our data and code upon acceptance.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

## 6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [Yes]

Justification: We provide the explicit experimental details. In Sec. 4.1 and Appendix, we provide the details of datasets, including the number, data splits. Besides, we set the Adam optimizer to train our LASQ.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

## 7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [Yes]

Justification: We run the experiments for 5 times and report the average value.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

#### 8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Our LASQ is trained on a cluster of four NVIDIA A800 GPUs. The detailed information about computer resources are given in Appendix.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

#### 9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines?>

Answer: [Yes]

Justification: Our research conducted in the paper conforms, in every respect, with the NeurIPS Code of Ethics.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

#### 10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [Yes]

Justification: The impact of our LASQ has been discussed in the Appendix.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

#### 11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: Our paper does not pose such risks.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

#### 12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification: Our paper does not use the existing above assets.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.

- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, [paperswithcode.com/datasets](https://paperswithcode.com/datasets) has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

### 13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: This paper does not release new assets.

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

### 14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: This paper not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

### 15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification: This paper does not involve crowdsourcing nor research with human subjects.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

#### 16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The LLMs is not the core methods in our research.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

## A The Detailed LASQ Workflow

**Luminance Variation Coordinate System:** To systematically characterize low-light degradation, we establish the Luminance Variation (LV) coordinate system that geometrically encodes pixel-wise illumination relationships. Let  $\mathcal{I}_L \in [0, 1]^{H \times W}$  and  $\mathcal{I}_N \in [0, 1]^{H \times W}$  denote paired low/normal-light images. For each pixel  $i$ , we define its luminance state as:

$$\{s_i = (\mathcal{I}_L^{(i)}, \mathcal{I}_N^{(i)})\}_{i=1}^I. \quad (13)$$

**Power Law like Transformation:** As visualized in Fig. 6 (a), these coordinate points  $s_i$  constitute a geometric manifold where the horizontal/vertical axes respectively represent low/normal-light intensity spaces. This formulation provides two critical physical insights: 1) Each  $s_i$  encodes a pixel-specific luminance attenuation pattern; 2) The points' spatial distribution reflects global illumination degradation statistics. The transition from Fig. 6 (b) to (d) reveals our hierarchical modeling strategy. For an individual  $s_i$ , we derive its local transformation strategy through:

$$y = ax^\kappa, x \in [0, 1], \kappa > 0, a > 0, \quad (14)$$

where  $\kappa$  quantifies the exposure compensation magnitude at pixel  $i$  (Fig. 6 (b)). We note that  $\kappa$  exhibits spatial correlation with scene content—lower  $\kappa$  values (steeper curves) correspond to regions requiring aggressive enhancement (e.g., shadows), while higher  $\kappa$  preserves highlights.

Extending to multiple pixels (Fig. 6(c)), we observe heterogeneous curve families governed by parameter distributions  $\{\kappa_p\}$ . Each curve represents a distinct luminance adjustment strategy:

- Highlight preservation ( $\kappa_p \rightarrow 1$ ): Linear mapping maintaining high-intensity features.
- Mid-tone enhancement ( $0.5 < \kappa_p < 1$ ): Concave curves amplifying moderate intensities.
- Shadow recovery ( $\kappa_p < 0.5$ ): Convex curves dramatically boosting dark regions.

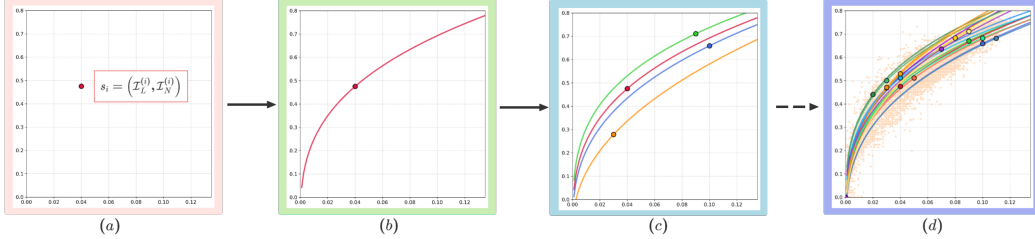


Figure 6: Luminance Variation Coordinate and the progressive modeling on full-domain luminance intensity elements via power law curves. (a) Luminance variation coordinate with pink background, where each sample  $s_i$  is geometrically determined by the horizontal low-light intensity  $\mathcal{I}_L^{(i)}$  and vertical normal-light intensity  $\mathcal{I}_N^{(i)}$ , marked by a central red dot; (b) Single-sample modeling (green background) demonstrating that the luminance transformation of individual elements follows a red-colored power law curve; (c) Multi-sample extension (cyan background) showing diversified mapping relationships through color-coded power law curves corresponding to different elements; (d) Full-domain fitting (purple background) achieved by dense overlapping curves that comprehensively cover the coordinate space.

Full-domain fitting in Fig. 6 (d) employs dense overlapping curves to approximate the entire  $\{s_i\}_{i=1}^I$  distribution. While achieving theoretical completeness through

$$\bigcup_{i=1}^I \{(x, ax^{\kappa_i}) \mid x \in [0, 1]\} \rightarrow \text{LV coordinate coverage}, \quad (15)$$

this over-parameterized paradigm introduces critical limitations. First, optimizing curves at the pixel level ignores spatial coherence, leading to artifacts as adjacent pixels apply divergent enhancement

strategies. Second, estimating  $\kappa_p$  becomes problematic in low-intensity areas ( $\mathcal{I}_L^{(i)} < 0.1$ ), where slight intensity shifts cause significant curve fluctuations. Lastly, focusing heavily on the individual  $\{\kappa_p\}$  parameters greatly restricts practical implementation due to poor generalization. Therefore, we statistically formulate an MCMC-sampling-based physics-aware  $\{\kappa_p\}$  prediction to preserve the LV system’s statistical advantages while mitigating overfitting.

**Hierarchical Sampling of Luminance Adaptation Operators:** Given a low-light RGB image  $\mathcal{I}_L \in \mathbb{R}^{H \times W \times 3}$ , we first convert it into the YUV color space and perform guided filtering to compute the smoothed illumination map  $G(i, j)$ , where  $(i, j)$  denotes pixel coordinates. The local linear coefficients in the guided filtering process can be formulated as:

$$a(i, j) = \frac{\sigma_L^2(i, j)}{\sigma_L^2(i, j) + \epsilon}, \quad b(i, j) = \mu_L(i, j) - a(i, j) \cdot \mu_L(i, j), \quad (16)$$

where  $\mu_L(i, j)$  represents the local window mean and  $\sigma_L^2(i, j)$  denotes the local window variance. Through coefficient smoothing,  $G(i, j)$  is obtained as:

$$G(i, j) = \bar{a}(i, j) \cdot L(i, j) + \bar{b}(i, j), \quad (17)$$

where  $\bar{a}$  and  $\bar{b}$  denote the mean-filtered results of  $a$  and  $b$ , respectively. This luminance map  $G(i, j)$  is then aggregated over arbitrary image regions  $\mathcal{P} \subseteq [1, H] \times [1, W]$  to obtain the regional scalar  $G_{\mathcal{P}}$ , which quantifies local luminance characteristics and serves as the input to compute the LAO  $\gamma_{\mathcal{P}}$  via the parameterized transformation introduced in the main text **Section 3.2 Statistical Sampling Process**.

**Coarse-to-Fine Hierarchy Gaussian distribution** As illustrated in Fig.7, LAOs are sampled from a truncated Gaussian distribution  $\mathcal{N}_{\text{trunc}}(\mu = \gamma_0, \sigma^2; \gamma_{\min}, \gamma_{\max})$ . Commencing from a globally balanced state, spatially-varying layers are iteratively introduced to simulate the continuity of real-world luminance transformations. The sampling operator employs a truncated Gaussian distribution wherein high-probability candidates correspond to physically plausible global adaptations, while low-probability candidates represent local refinements. This characteristic inherently reflects the natural illumination transitions where extreme pixel-level modifications rarely occur. The transition between states is governed by a kernel that ensures smooth evolution across iterations:

$$q\left(\gamma_{\mathcal{P},z}^{(n)} \mid \gamma_{\mathcal{P},z-1}^{(n)}\right) = \mathcal{N}_{\text{trunc}}\left(\gamma_{\mathcal{P},z}^{(n)} \mid \gamma_{\mathcal{P},z-1}^{(n)}, \lambda^2; \gamma_{\min}, \gamma_{\max}\right), \quad (18)$$

This hierarchical sampling yields a sequence of LAO sets  $\Gamma = \{\Gamma_n\}_{n=1}^N$ , with each  $\Gamma_n = \left\{\gamma_{\mathcal{P},z}^{(n)}\right\}_{z=1}^{2^{n-1}}$  representing the spatially adaptive enhancement parameters. When applied, these sets give rise to a family of luminance-enhanced images  $\mathcal{H} = \{\mathcal{I}_H^{(n)}\}_{n=1}^N$ , capturing a coarse-to-fine progression of illumination correction. This structure enables flexible exploration of global and local enhancement effects through state-dependent sampling transitions across the luminance field.

**Hierarchically-Guided Diffusion:** The forward process employs a hierarchically-guided diffusion framework to progressively enhance low-light images through semantically-aligned noise injection. By encoding the input  $\mathcal{I}_L$  and its luminance-enhanced variants  $\mathcal{H}$  into latent features  $\mathcal{F}_L$  and  $\left\{\mathcal{F}_H^{(n)}\right\}_{n=1}^N$ , the method establishes temporal correspondence between diffusion steps and multi-scale guidance via a mapping function  $\psi(t)$ . This alignment dynamically steers the noise trajectory, where each transition step incorporates hierarchical semantic shifts derived from  $\mathcal{F}_H^{(\psi(t))}$  to modulate the degradation path. Throughout the forward pass, coarse-to-fine illumination characteristics are gradually imprinted into the diffusion process. During reverse denoising, the network  $\epsilon_{\theta}$  iteratively removes noise while synergizing hierarchical cues from original low-light features  $\mathcal{F}_L$ . The denoising trajectory is regularized by a dual objective: a noise prediction loss ensures faithful reconstruction aligned with hierarchical semantics, while a global constraint  $\mathcal{L}_g$  enforces the structural consistency.

**Adversarial Discriminator for Iterative Refinement:** To optionally further enhance realism and perceptual fidelity—only when a normal-light reference is available—we introduce a discriminator  $\mathcal{D}_{\phi}$  alongside our diffusion-based generator  $G_{\theta}$ . Given an unpaired sample  $\mathcal{I}_{\text{normal}} \sim p_{\text{normal}}$  (if available) from the normal-light distribution, the discriminator learns to distinguish between true normal-light images and our enhanced outputs  $\hat{\mathcal{I}}_N = G_{\theta}(\mathcal{I}_L)$ . We therefore cast the overall framework as a hybrid

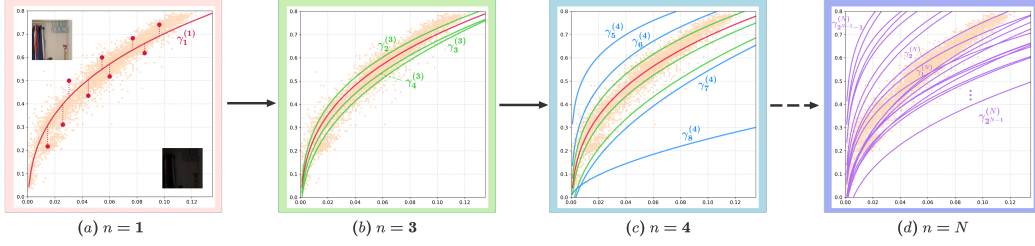


Figure 7: The hierarchical MCMC sampling scheme to generate LAO sets. (a) Initial state ( $n = 1$ ) with a uniform LAO operator  $\gamma_1^{(1)}$  (red curve) globally adjusting pixel intensities via orthogonal projections to fit shallow orange sample points (approximate power-law distribution); (b) Intermediate sampling ( $n = 3$ ) generates three green curves ( $\gamma_{i=2,3,4}^{(3)}$ ) near the red curve, grounding on the Markov chain posterior; (c) Increased sampling times ( $n = 4$ ) introduces external blue curves ( $\gamma_{i=5,6,7,8}^{(4)}$ ), which entail power law functions ( $\gamma_{i=5,8}^{(4)}$ ) that are outlying of the main distribution area; (d) Converged state ( $n = N$ ) achieves pixel-level granularity through dense overlapping curves, dynamically fitting the orange distribution.

diffusion-GAN, optimizing both the variational diffusion objective and an adversarial objective in an alternating fashion.

Concretely, at each training iteration  $t$ , we perform the discriminator update as:

$$\begin{aligned} \phi &\leftarrow \phi - \eta_\phi \nabla_\phi \mathcal{L}_{\text{adv}}(\mathcal{D}_\phi; G_\theta), \\ \mathcal{L}_{\text{adv}} &= -\mathbb{E}_{\mathcal{I}_{\text{normal}} \sim p_{\text{normal}}} [\log \mathcal{D}_\phi(\mathcal{I}_{\text{normal}})] - \mathbb{E}_{\mathcal{I}_L} [\log(1 - \mathcal{D}_\phi(G_\theta(\mathcal{I}_L)))] \end{aligned} \quad (19)$$

and the generator update can be formulate as:

$$\begin{aligned} \theta &\leftarrow \theta - \eta_\theta \nabla_\theta (\mathcal{L}_{\text{total}} + \lambda_{\text{GAN}} \mathcal{L}_{\text{G-adv}}), \\ \mathcal{L}_{\text{G-adv}} &= -\mathbb{E}_{\mathcal{I}_L} [\log \mathcal{D}_\phi(G_\theta(\mathcal{I}_L))], \end{aligned} \quad (20)$$

where  $\eta_\phi, \eta_\theta$  are learning rates, and  $\lambda_{\text{GAN}}$  balances adversarial and diffusion losses. Therefore, the full generator loss becomes:

$$\mathcal{L}_{\text{total}} = \lambda_d \mathcal{L}_d + \lambda_g \mathcal{L}_g + \underbrace{\lambda_{\text{GAN}} \mathbb{E}_{\mathcal{I}_L} [-\log \mathcal{D}_\phi(G_\theta(\mathcal{I}_L))]}_{\text{adversarial penalty}}. \quad (21)$$

By unifying physically grounded diffusion priors with adversarial discrimination, our hybrid framework yields outputs that not only preserve structural details and smooth lighting transitions but also exhibit the sharp textures and natural contrasts characteristic of true normal-light images.

**Inference Stage:** During the inference phase, the model synthesizes an optimal enhanced representation  $\hat{\mathcal{F}}_N$  that balances global illumination correction with local texture fidelity through implicit latent sampling guided by  $\mathcal{F}_L$ , starting from Gaussian random noise. This representation is subsequently decoded into the final output  $\hat{\mathcal{I}}_N$  through the learned mapping function.

## B The Algorithmic Framework

Algorithm 1 summarizes the full LASQ (Luminance-Adaptation Sampling and Hierarchically-Guided Diffusion) workflow. It first extracts a latent representation via an encoder, then performs multi-scale luminance-adaptive sampling to generate a hierarchy of intermediate images. These are used to guide both the forward and reverse diffusion processes, and finally a decoder reconstructs the enhanced high-quality output, optionally applying adversarial training for further refinement.

---

**Algorithm 1** The LASQ pipeline.

---

```
1: Input: Low-light image  $\mathcal{I}_L$ 
2: Hierarchical LAO sampling:
3:   Initialize  $\Gamma = \{\}, \mathcal{I}_H^{(0)} = \{\}$ 
4:   Compute mean luminance  $G_p$ , then
5:    $\gamma_P = (\alpha + G_P)^{\beta_P}$ 
6:   for  $n = 1 \dots N$  do
7:     Sample  $\gamma \sim \mathcal{N}_{\text{trunc}}(\gamma_0, \sigma^2, \gamma_{\min}, \gamma_{\max})$ 
8:      $\Gamma_n = \left\{ \gamma_{\mathcal{P},z}^{(n)} \right\}_{z=1}^{2^{n-1}}$ ; Update  $\mathcal{I}_H^{(n)}$ 
9:   end for
10:  Collect  $\mathcal{H} = \{\mathcal{I}_H^{(n)}\}_{n=1}^N$ 
11: Hierarchically-Guided diffusion:
12:   $\mathcal{F}_L \leftarrow \mathcal{E}(\mathcal{I}_L); \left\{ \mathcal{F}_H^{(n)} \right\}_{n=1}^N \leftarrow \mathcal{E}(\mathcal{H})$ 
13:  Define  $\psi(t) = \lceil tN/T \rceil$ , set  $x_0 = \mathcal{F}_H^{(0)}$ 
14:  while not converged do
15:    for  $t = 1 \dots T$  do
16:       $x_0 = \mathcal{F}_H^{(\psi(t))}$ 
17:       $x_t \sim \mathcal{N}(\sqrt{1 - \beta_t}x_{t-1}, \beta_t I)$ 
18:      Perform gradient descent steps on  $\nabla_{\theta} \|\epsilon - \epsilon_{\theta}(x_t, t, \mathcal{F}_L)\|^2$ 
19:    end for
20:    for  $t = T \dots 1$  do
21:      Predict  $\epsilon_{\theta}(\hat{x}_t, t, \mathcal{F}_L)$ 
22:       $\hat{x}_{t-1} = \frac{1}{\sqrt{1 - \beta_t}}(\hat{x}_t - \beta_t \epsilon_{\theta}) + \sigma_t b$ 
23:    end for
24:     $\hat{\mathcal{I}}_N = \mathcal{D}(\hat{x}_0)$ 
25:    Perform gradient descent steps on  $\nabla_{\theta} \left\| \hat{\mathcal{I}}_N - \mathcal{D}(\mathcal{F}_H^{(\psi(0))}) \right\|_1$ 
26:    Optional:  $\nabla_{\theta} [-\log \mathcal{D}_{\phi}(G_{\theta}(\mathcal{I}_L))]; \nabla_{\phi} [\log \mathcal{D}_{\phi}(\mathcal{I}_{\text{normal}}) + \log(1 - \mathcal{D}_{\phi}(G_{\theta}(\mathcal{I}_L)))]$ 
27:  end while
28: Return  $\hat{\mathcal{I}}_N$ 
```

---

## C Supplementary Experiments and Extended Results

**Visual Comparison:** As demonstrated in the supplemental visual comparisons (Fig. 8-10), LASQ exhibits remarkable stability across both controlled laboratory settings and challenging real-world scenarios. Figures 8 and 9 systematically visualize LASQ’s consistency in ground truth-annotated scenes, where it maintains precise alignment with reference images in terms of color temperature, dynamic range distribution, and fine-grained texture preservation. The visual trajectories across multiple test cases confirm that LASQ effectively decouples scene semantics from lighting interference, avoiding the common pitfalls of oversaturation or detail loss observed in supervised baselines. Fig. 10 further underscores LASQ’s robustness under extreme conditions, specifically in nocturnal environments and localized overexposure scenarios—settings that are notoriously difficult for conventional methods. In night-time scenes with extremely low ambient illumination, LASQ demonstrates a pronounced ability to preserve structural integrity and suppress noise amplification, resulting in outputs that maintain visual coherence and readability. Unlike comparative methods that often introduce flare artifacts, chromatic aberrations, or abrupt luminance shifts under high-contrast lighting, LASQ adaptively regulates brightness and minimizes highlight clipping, ensuring that no important spatial detail is lost. Particularly in overexposed regions—such as headlights, reflective surfaces, or artificially illuminated zones—LASQ retains nuanced intensity gradients and avoids flattening or unnatural glow, thereby delivering smooth tonal transitions that remain faithful to human perceptual expectations. These results confirm that LASQ not only generalizes well across diverse lighting conditions but also excels where existing techniques tend to break down.

**Quantitative Results:** To provide a more comprehensive evaluation of LASQ and LASQ++ under realistic, unconstrained conditions, we present in Table 5 a quantitative comparison on five widely used no-reference low-light image enhancement benchmarks: DICM, NPE, VV, LIME, and MEF.

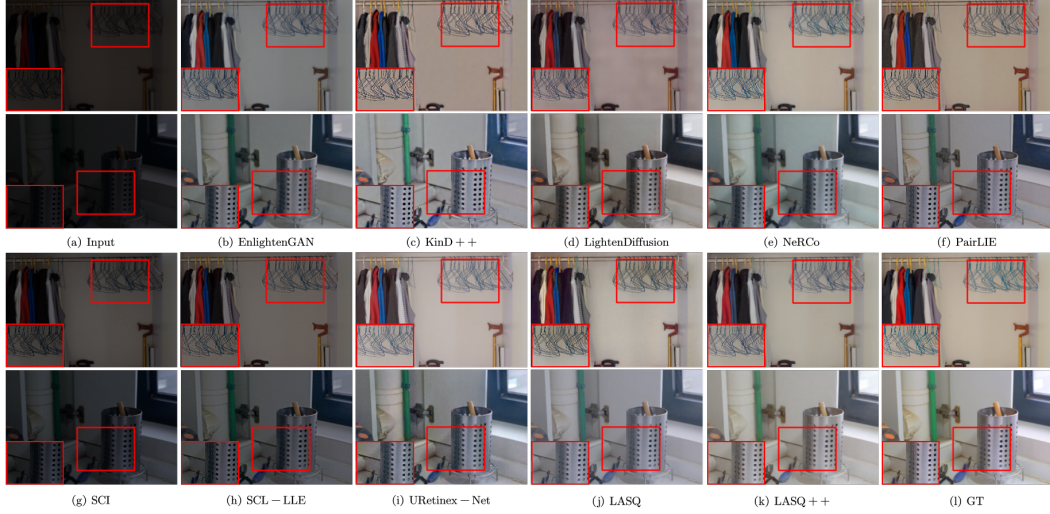


Figure 8: The supplemental qualitative comparison of our method and competitive methods on the LOLv1 test sets. "LASQ++" denotes the incorporation of unpaired normal-light references.

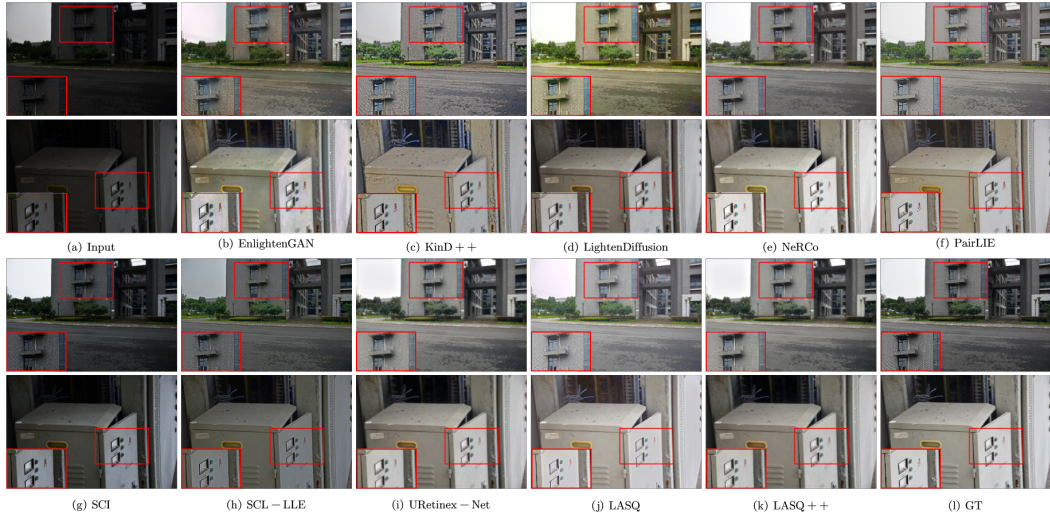


Figure 9: The supplemental qualitative comparison of our method and competitive methods on the LSRW test sets. "LASQ++" denotes the incorporation of unpaired normal-light references.

These datasets lack paired ground-truth (GT) normal-light references, making them a robust testbed for assessing generalization and real-world applicability. We evaluate performance using two standard no-reference metrics—NIQE and PI—and compare LASQ and LASQ++ against a broad range of state-of-the-art supervised (SL) and unsupervised (UL) methods. As shown, LASQ consistently achieves the best or second-best results across almost all datasets and metrics, outperforming all other unsupervised methods and remaining competitive with even several fully supervised approaches. This highlights the strong generalization capability of LASQ, which requires no paired data and yet adapts effectively to diverse lighting conditions. LASQ++, which incorporates unpaired normal-light references during training, pushes perceptual quality even further. While this slightly reduces robustness in the most challenging cross-domain settings, it reinforces the complementary strengths of our two models: LASQ is ideal for broad deployment due to its high robustness and adaptability, whereas LASQ++ excels when enhanced fidelity is needed in target-specific domains.



Figure 10: Visual method comparison in challenging scenarios. "LASQ++" denotes the incorporation of unpaired normal-light references.

## D Computation Resources and Hyper-parameter Tuning

All experiments were conducted on a high-performance computing node equipped with four NVIDIA A100 80GB GPUs interconnected via NVLink. The system specifications are as follows:

- **OS:** Ubuntu 22.04 LTS with Linux 5.15 kernel
- **CPU:** Dual AMD EPYC 7763 64-Core @ 2.45GHz (128 cores/256 threads)
- **GPU Interconnect:** NVLink 3.0 (600GB/s bisectional bandwidth)
- **Memory:** 1TB DDR4 ECC @ 3200MHz
- **Storage:** 16TB NVMe SSD RAID (3.5GB/s sustained read)
- **Accelerators:** 4×NVIDIA A100 80GB (FP32: 19.5 TFLOPS, FP16: 312 TFLOPS)

The hyperparameter configuration for the proposed framework is comprehensively summarized as follows. All experiments are conducted on four NVIDIA A800 GPUs utilizing Python 3.9 and PyTorch 2.0 with a fixed batch size of 16, employing the Adam optimizer with a denoising diffusion learning rate of  $2 \times 10^{-5}$  and a sampling ratio  $k = 3$ . The loss weighting coefficients  $\lambda_d$ ,  $\lambda_g$ , and  $\lambda_{GAN}$  are empirically set to 0.9, 0.005, and 0.7 respectively when adversarial training is activated. The diffusion process operates over  $T=1000$  time steps with a U-Net-based noise estimation architecture. In the luminance adaptation framework, the power-law adjustment parameters  $\alpha$ ,  $\eta$ , and  $\delta$  governing local contrast enhancement are initialized to 2, 0.1, and 0.01 respectively, while the MCMC sampling process employs an adaptive step size  $\lambda = 0.2$  to balance exploration-exploitation dynamics across hierarchy levels. The temporal mapping function  $\psi(t)$  synchronizes  $N = 100$  hierarchical guidance levels with the  $T$ -step diffusion through linear interpolation. The truncated Gaussian distribution for LAO sampling is bounded by  $\gamma_{\min}$  and  $\gamma_{\max}$  derived dynamically from image statistics, ensuring physically plausible luminance adjustments. For adversarial training augmentation, the discriminator  $D_\phi$  maintains architectural hyperparameters aligned with LSGAN conventions (including convolutional layer configurations and spectral normalization usage), though implemented with standard binary cross-entropy objectives rather than least-squares formulations. The network depth and feature channel dimensions are adaptively scaled according to the spatial resolution of input pairs from the generator's decomposition process, while preserving LSGAN's fundamental design principles in layer progression and discriminator receptive field structure.

Table 5: Supplementary quantitative comparisons on the no-reference image datasets from the main text, with the best-performing results marked in red and the second-best in blue. The notations "SL" and "UL" respectively represent supervised and unsupervised learning approaches. "LASQ++" denotes the incorporation of unpaired normal-light references.

Type	Method	DICM		NPE		VV		LIME		MEF	
		NIQE↓	PI↓	NIQE↓	PI↓	NIQE↓	PI↓	NIQE↓	PI↓	NIQE↓	PI↓
SL	RetinexNet	4.487	3.242	4.732	3.219	5.881	3.727	4.802	3.522	4.152	3.411
	KinD++	4.027	3.999	4.005	3.144	3.586	2.773	4.035	3.217	3.874	3.285
	LCDPNet	4.110	3.250	4.126	3.127	5.039	3.347	4.128	3.332	3.912	3.398
	URetinexNet	4.774	3.565	4.028	3.153	3.851	2.891	3.987	3.104	3.721	3.185
	SMG	6.224	4.228	5.300	3.627	5.752	3.757	5.312	3.615	5.028	3.804
	PyDiff	4.499	3.792	4.082	3.268	4.360	3.678	4.412	3.685	4.228	3.572
UL	Zero-DCE	3.951	3.149	3.826	2.918	5.080	3.307	3.625	3.512	3.608	3.217
	EnlightenGAN	3.832	3.256	3.775	2.953	3.689	2.749	3.427	3.424	3.524	3.108
	SCI	4.519	3.700	4.124	3.534	5.312	3.648	4.032	3.518	3.892	3.415
	PairLIE	4.282	3.469	4.661	3.543	3.373	2.734	3.782	3.215	3.412	3.028
	SCL-LLE	5.129	3.809	4.873	3.692	5.513	4.316	5.104	4.302	4.872	4.115
	NeRCo	4.107	3.345	3.902	3.037	3.765	3.094	3.712	3.078	3.328	3.112
	LightenDiffusion	3.724	3.144	3.618	2.879	2.941	2.558	3.218	3.128	3.305	3.024
	LASQ	3.715	3.128	3.571	2.764	2.777	2.623	3.152	3.002	3.294	3.001
	LASQ++	3.723	3.137	3.601	2.789	2.850	2.691	3.167	3.046	3.309	3.013

Table 6: Supplementary ablation results.

Method	Pre-trained Dataset			VV		MEF		NPE		DICM		LIME	
	LOLv1	LSRW	MEF	NIQE↓	PI↓	NIQE↓	PI↓	NIQE↓	PI↓	NIQE↓	PI↓	NIQE↓	PI↓
LASQ	✓			2.777	2.623	3.294	3.001	3.571	2.764	3.715	3.128	3.152	3.002
		✓		2.801	2.637	3.307	3.019	3.584	2.762	3.750	3.160	3.191	3.222
			✓	2.882	2.703	3.287	3.010	3.627	2.758	3.744	3.177	3.248	3.335

## E Ablation

As shown in Table 6, we present extend ablation studies to systematically evaluate the proposed framework. LASQ is trained on three datasets (LOLv1, LSRW, and MEF with manually curated training splits augmented via random cropping), followed by cross-dataset evaluations on MEF and VV benchmarks. Empirical results demonstrate LASQ’s superior generalization capability in unseen scenarios, achieving consistent reconstruction quality under diverse illumination conditions. This evidence confirms that LASQ avoids overfitting to region-specific normal-light patterns and exhibits reduced sensitivity to pretraining data selection. Two principal advantages are emphasized: (1) Cross-domain adaptation through layered luminance modeling enables high-fidelity visual outputs without paired supervision; (2) Effective suppression of noise and artifacts under extreme low-light conditions, substantially improving perceptual quality. These advancements highlight LASQ’s potential to redefine low-light image enhancement paradigms by harmonizing physical priors with data-driven learning.

## F Limitations

The main limitations of our approach stem from its hierarchical sampling depth  $N$  and diffusion step count  $T$ : as  $N$  grows, the number of local correction patches doubles at each level, and as  $T$  increases, each diffusion step must incorporate the corresponding hierarchical features, so jointly large  $N$  and  $T$  lead to exponential increases in computation and memory demands. Moreover, the final enhancement quality is quite sensitive to the choice of  $N$  (too few layers underfits, too many layers overfits and amplifies noise) and  $T$  (insufficient steps yield coarse results, excessive steps waste resources), as well as to key hyperparameters in the  $\gamma$ -calculation (e.g. the contrast gain terms  $\eta$  and  $\delta$ ), which must be manually tuned to balance enhancement strength against artifact suppression across different scenes.

## **G   Impacts**

Our LASQ advances low-light imaging for critical applications like target detection, self-driving and medical diagnostics, where reliable scene interpretation under poor illumination is essential. Our LASQ mimics natural light adaptation to improve nighttime safety. Its unsupervised design removes the need for paired data, expanding access to low-light enhancement for applications lacking annotated datasets.