

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 LEARNING TO ENHANCE LOW-LIGHT IMAGES WITH RELIABLE ATTENTION AND REINFORCED DISTRIBUTION ALIGNMENT

Anonymous authors

Paper under double-blind review

ABSTRACT

Low-light image enhancement (LLIE) methods have recently adopted the HVI color space, which alleviates the entanglement between luminance and color and improves color fidelity through chrominance polarization and intensity compression. However, existing approaches may suffer from error accumulation during the interaction between luminance and chrominance components, and the lack of fine-grained modeling of color distribution can lead to unsatisfactory enhancement results. To address these challenges, we propose a novel low-light image enhancement framework, Learning to Enhance Low-Light Images with Reliable Attention and Reinforced Distribution Alignment. Specifically, we introduce two key modules: the Reliable Cross Attention (RCA) module, which aggregates luminance and chrominance features with reliable queries, and the Reinforced Distribution Alignment (RDA) module, which robustly fits the color distribution in a more fine-grained manner. These designs significantly improve the quality of enhanced images under low-light conditions. Extensive experiments on multiple benchmark datasets demonstrate that our method achieves state-of-the-art performance compared with existing approaches.

1 INTRODUCTION

Images captured by imaging sensors in low-light conditions often suffer from significant noise. To address this issue, the task of low-light image enhancement has emerged, focusing on improving brightness, contrast, and the visibility of details in dark environments. In addition to its standalone benefits, low-light enhancement serves as an important foundation for various downstream vision tasks, such as object detection Zou et al. (2023), tracking Feichtenhofer et al. (2017), and image matching Cheng et al. (2025). However, in the conventional RGB color space, color and luminance are closely intertwined, which can lead to color distortions or unnatural brightness after enhancement. To alleviate these problems, multiple distinct categories of methods have been developed.

Traditional sRGB-based methods Wang et al. (2022a) often cause color shifts and distortions due to the coupling of luminance and chromaticity. To address this, some approaches Guo & Hu (2023) convert images to the HSV color space for more precise luminance enhancement. However, HSV introduces new issues such as red channel discontinuities and black-plane noise, resulting in visual artifacts and color distortions. To tackle these challenges, CIDNet Yan et al. (2025) introduces the HVI color space, which is specifically designed for low-light image enhancement. The HVI color space polarizes the hue-saturation plane to reduce red-channel discontinuities and employs a learnable intensity compression function to adaptively suppress low-luminance areas, effectively minimizing black noise artifacts. This design significantly enhances both color fidelity and perceived naturalness in low-light conditions. Nevertheless, residual noise and incomplete decoupling between luminance and chromaticity still introduce errors during feature interaction, which degrade the overall naturalness and smoothness of enhanced images. Moreover, due to the challenges in precisely modeling the distribution in HV space, global color distortions remain a common issue.

From our discussion on HVI-based methods, we have identified two main challenges that need to be addressed to achieve more accurate and robust low-light image enhancement. **(1) How to efficiently aggregate luminance and chromaticity features while avoiding noise amplification.** Previous ap-

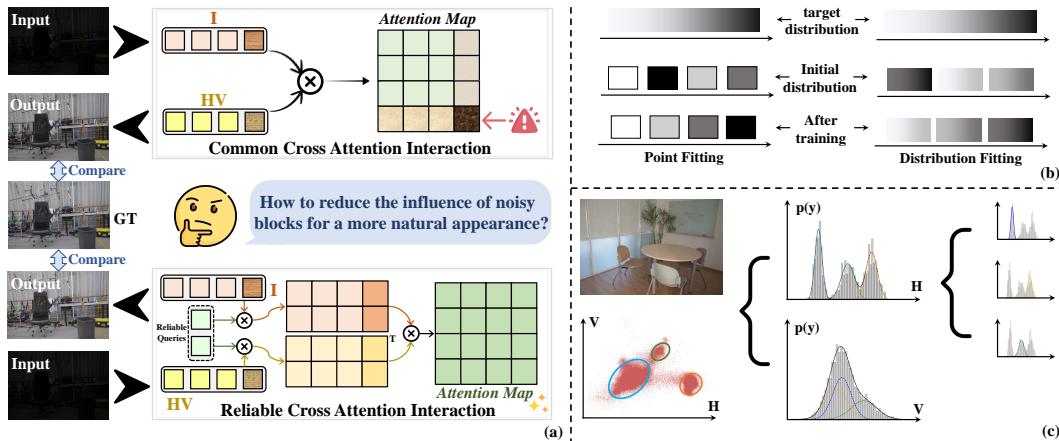


Figure 1: (a) Visualization comparing our method with traditional interactions. (b) Comparison between point-based and distribution-based modeling, showing that distribution modeling provides a more fine-grained fit to the target distribution. (c) Visualization of distribution modeling, where clustering and Gaussian mixture modeling effectively capture the distributions.

proaches often employ conventional transformer-based interactions, which tend to amplify noise when low-quality regions from luminance and color spaces are aggregated. This leads to noticeable degradation in overall image quality. To address this issue (see Figure 1(a)), we argue that leveraging high-quality features to guide aggregation can better align the enhancement process with real-scene information, thereby improving enhancement quality while suppressing the interference of noisy or invalid features in dark regions. **(2) How to precisely constrain the distribution of chromaticity features to achieve more natural color rendering.** Existing methods commonly adopt point-based constraints to approximate the feature distribution, but such strategies suffer from a lack of global consistency and insufficient modeling of the overall distribution (Figure 1(b)). To overcome this limitation, we propose to approximate the feature distribution with multiple fitted distributions (Figure 1(c)), enabling a more fine-grained characterization of the joint statistics between luminance and chromaticity. Furthermore, we introduce a reinforcement learning mechanism to dynamically adjust the distribution fitting process, allowing the model to adaptively select optimal distribution parameters. This not only preserves global consistency but also enhances the naturalness and robustness of the enhanced results.

To address the above challenges, we propose Learning to Enhance Low-Light Images with Reliable Attention and Reinforced Distribution Alignment, which introduces two key modules: the Reliable Cross Attention (RCA) and Reinforced Distribution Alignment (RDA). RCA employs reliable queries to bridge luminance and chromaticity features, replacing traditional transformer interactions. This design suppresses noise from redundant features, improves the naturalness of enhanced images, and reduces computational complexity for higher efficiency. RDA targets the chromaticity components (H and V), applying unsupervised clustering and Gaussian modeling to capture multiple distributions. A reinforcement learning mechanism then dynamically adjusts distribution parameters for better alignment with global chromaticity statistics. As a result, our model achieves more natural and stable enhancement, especially under low-light conditions.

In summary, our contributions are threefold:

1. We propose a novel framework, Learning to Enhance Low-Light Images with Reliable Attention and Reinforced Distribution Alignment, achieving state-of-the-art performance in low-light image enhancement.
2. We propose the Reliable Cross Attention (RCA) module to adaptively aggregate luminance and chromaticity features, suppress noise, and enhance illumination balance and color fidelity. Additionally, we introduce the Reinforced Distribution Alignment (RDA) module, which models multiple chromaticity distributions using Gaussian clustering and dynamically refines them through reinforcement learning, leading to clearer and more natural results.
3. Extensive experiments and ablation studies on ten benchmark datasets validate the superiority of our method.

108

2 RELATED WORKS

109
110 In this section, we review existing approaches for low-light image enhancement.111 **Traditional Methods.** Early low-light enhancement methods were largely heuristic and did not
112 require training data. Histogram equalization Pizer et al. (1987) and gamma correction Rahman
113 et al. (2016) improve contrast and brightness by redistributing pixel intensities, but often ignore
114 scene illumination, leading to over-enhanced or washed-out results. Retinex-based approaches Land
115 & McCann (1971); Rahman et al. (2004) decompose an image into illumination and reflectance
116 components and refine the illumination with structural priors. Although more physically motivated,
117 they rely on idealized assumptions and are prone to noise amplification and color distortion in real-
118 world conditions.119 **Learning-Based Methods.** Deep learning has transformed low-light enhancement into a data-
120 driven task. RetinexNet Wei et al. (2018) and KinD Zhang et al. (2019) embed Retinex decom-
121 position into CNNs, but remain sensitive to illumination estimation, often amplifying noise or shifting
122 colors. ZeroDCE Guo et al. (2020) and RUAS Liu et al. (2021) avoid explicit decomposition by
123 learning pixel-adaptive curves or structural priors, but may introduce artifacts or unstable chromi-
124 nance. Flow-based methods such as LLFlow Wang et al. (2022a) deliver high-fidelity restora-
125 tion via normalizing flows, but incur heavy computational costs and require paired data. GAN-
126 based approaches like EnlightenGAN Jiang et al. (2021) enhance perceptual realism through ad-
127 versarial training, though sometimes at the expense of unnatural textures. UFormer Wang et al.
128 (2022b) introduces a Transformer-based U-shaped architecture with a Locally-enhanced Window
129 Transformer and multi-scale restoration modulator, achieving top performance in image restora-
130 tion. Restormer Zamir et al. (2022a) presents an efficient Transformer model with a multi-Dconv
131 head attention mechanism and multi-scale design, excelling in image deraining, deblurring, and
132 denoising. MIRNet Zamir et al. (2022b) uses multi-scale residual blocks and non-local atten-
133 tion to preserve spatial details and context, achieving state-of-the-art results in image denoising,
134 super-resolution, and enhancement. SNR-aware networks Xu et al. (2022) integrate noise priors
135 to reduce artifacts but still struggle with color inconsistency. Transformer-based models, e.g., LL-
136 Former Wang et al. (2023) and RetinexFormer Cai et al. (2023), capture long-range dependencies
137 but lack explicit channel-level alignment. Bread Guo & Hu (2023) mitigates noise–color entangle-
138 ment in YCbCr space. GSAD Hou et al. (2023) employs a global structure-aware diffusion process,
139 and QuadPrior Wang et al. (2024) introduces physical priors to constrain illumination enhancement.
140 However, these methods often face issues such as overexposure, color shifts, or high computational
141 cost. RetinexMamba Bai et al. (2024) combines traditional Retinex theory with deep learning to
142 improve illumination estimation and noise suppression for low-light enhancement. More recently,
143 CIDNet Yan et al. (2025) explored the HVI color space to alleviate red discontinuities and black
144 noise. Our method builds upon existing low-light image enhancement techniques, addressing key
145 challenges in feature aggregation and chromaticity distribution modeling to achieve superior results.146

3 METHOD

147 The proposed method is illustrated in Figure 2. The input image is first mapped to the HVI color
148 space to separate luminance from chromaticity, and then passed through the Reliable Cross Atten-
149 tion (RCA) module. This module aggregates luminance and chromaticity features using reliable
150 queries, effectively suppressing noise and enhancing naturalness. Next, the Reinforced Distribution
151 Alignment (RDA) module models chromaticity distributions through unsupervised clustering and
152 Gaussian mixture modeling, while reinforcement learning dynamically adjusts parameters for better
153 global alignment. Finally, the enhanced HVI representation is mapped back to the RGB color space.
154 The role of the HVI transformation is explained in the Appendix A.1.155

3.1 RELIABLE CROSS ATTENTION MODULE

156
157 To establish robust and reliable interactions between the luminance and chromaticity features and
158 mitigate noise amplification during cross-domain feature aggregation, we propose the Reliable Cross
159 Attention (RCA) module. The RCA module employs a set of learnable query vectors $Q_0 \in \mathbb{R}^{M \times d}$
160 (where M denotes the number of query vectors) initialized close to zero to extract a compact col-
161 lection of high confidence descriptors denoted as R_Q from the joint intensity-chromaticity repre-
162 sentation via cross attention. The input feature maps of the intensity and chrominance branches are

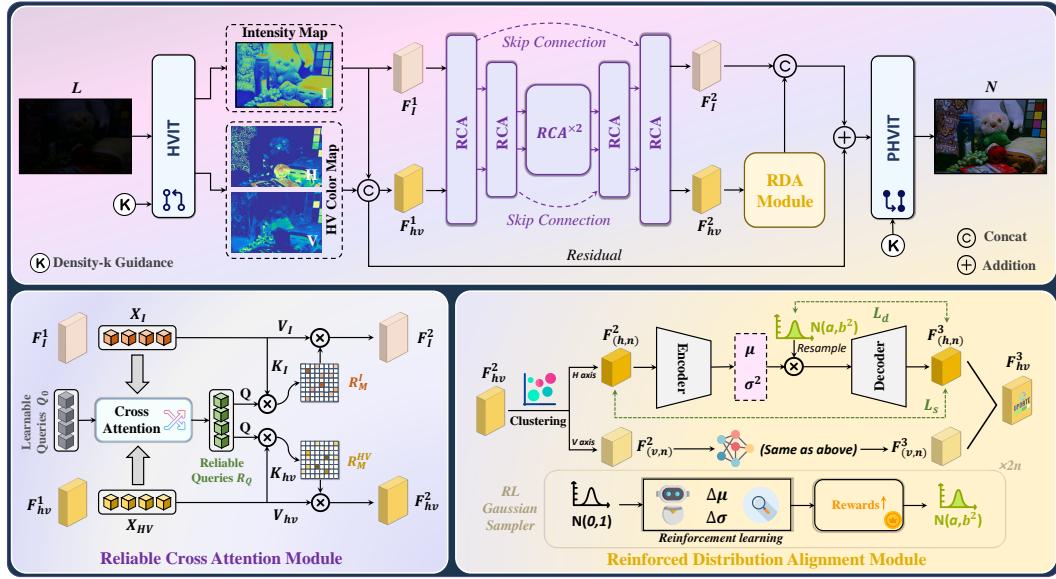


Figure 2: The overall pipeline of our method process begins by transforming the input into the HVI space. The Reliable Cross Attention (RCA) module aggregates luminance and chromaticity features with reliable queries, suppressing noise and improving naturalness. The Reinforced Distribution Alignment (RDA) module models chromaticity distributions via clustering and Gaussian modeling, with reinforcement learning dynamically refining parameters for better global alignment. The final enhanced image is reconstructed in the sRGB color space.

denoted by $F_I^1 \in \mathbb{R}^{C \times H \times W}$ and $F_{hv}^1 \in \mathbb{R}^{C \times H \times W}$, respectively, which are flattened along spatial dimensions to obtain token matrices $X_I \in \mathbb{R}^{N \times C}$ and $X_{hv} \in \mathbb{R}^{N \times C}$ with $N = H \cdot W$.

RCA operates in two stages: extracting reliable queries from the joint intensity-chromaticity representation, and guiding each branch with per-branch reliability maps that modulate the value tensors before residual fusion. In the first stage, branch-specific linear projections produce keys and values,

$$\begin{aligned} K_I &= X_I W_K^I, & V_I &= X_I W_V^I, \\ K_{hv} &= X_{hv} W_K^{hv}, & V_{hv} &= X_{hv} W_V^{hv}, \end{aligned} \quad (1)$$

where $W_K^I, W_V^I, W_K^{hv}, W_V^{hv} \in \mathbb{R}^{C \times d}$ are learned projection matrices and d is the scaling factor to stabilize gradients. Through concatenation, joint key/value matrices are formed as:

$$K_{\text{joint}} = \text{Concat}(K_I, K_{hv}) \in \mathbb{R}^{N' \times d}, \quad V_{\text{joint}} = \text{Concat}(V_I, V_{hv}) \in \mathbb{R}^{N' \times d}, \quad (2)$$

when $N' = 2N$. The empty queries Q_0 probe the joint matrix via scaled dot-product cross-attention to produce a compact set of reliable query descriptors $R \in \mathbb{R}^{M \times d}$:

$$A_{\text{joint}} = \text{softmax}\left(\frac{Q_0 K_{\text{joint}}^\top}{\sqrt{d}}\right) \in \mathbb{R}^{M \times N'}, \quad R_Q = A_{\text{joint}} V_{\text{joint}} \in \mathbb{R}^{M \times d}. \quad (3)$$

In the second stage, the reliable queries R_Q act as queries against each branch key to yield per-query compatibility maps for the intensity and HV branches:

$$R_M^I = \text{softmax}\left(\frac{R_Q K_I^\top}{\sqrt{d}}\right) \in \mathbb{R}^{M \times N}, \quad R_M^{HV} = \text{softmax}\left(\frac{R_Q K_{hv}^\top}{\sqrt{d}}\right) \in \mathbb{R}^{M \times N}. \quad (4)$$

where $x \in \{1, \dots, N\}$, R_M^I and R_M^{HV} denote the per-query compatibility maps in the I and HV branches, respectively. Finally, these per-branch reliability maps multiplicatively modulate the branch value tensors element-wise,

$$\tilde{F}_I^2 = R_M^I V_I(x), \quad \tilde{F}_{hv}^2 = R_M^{HV} V_{hv}(x), \quad (5)$$

where F_I^2 and F_{hv}^2 denote the intermediate feature representations of the intensity and HV branches, respectively. The modulated responses are projected back into the channel space and reshaped to the original spatial resolution, then fused with input features via residual addition. Subsequently, LayerNorm is applied to stabilize the feature distribution and enable reliable optimization, this step is omitted from Figure 2 for visual clarity.

Overall, the RCA module selectively aggregates reliable information across intensity and chromaticity branches while suppressing noise amplification, thereby improving enhancement quality and reducing computational complexity from $O(N^2)$ to $O(MN)$ where M is significantly smaller than N , compared to conventional global attention mechanisms.

3.2 REINFORCED DISTRIBUTION ALIGNMENT MODULE

After obtaining high-quality luminance and chrominance components through the RCA module, we further consider the distribution modeling of the chrominance branch. Previous methods lack targeted constraints in this aspect, and common approaches that approximate the distribution via point-wise constraints often suffer from insufficient global consistency and inadequate modeling of overall statistics. To address this issue, we design the proposed RDA module, which models the distribution as a Gaussian Mixture Model (GMM) Rasmussen (1999) through unsupervised clustering Xie et al. (2016). Furthermore, reinforcement learning is employed to drive the policy that adaptively adjusts the posterior distribution along the H and V directions, thereby fitting an optimal chrominance distribution for low-light image enhancement.

Formally, given the chrominance components $F_H, F_V \in \mathbb{R}^{n \times d}$, we first perform deep embedded clustering Xie et al. (2016); Guo et al. (2018) to partition the features into K groups. Each group is then parameterized by a Gaussian mixture model (GMM):

$$p(\mathbf{x}) = \sum_{k=1}^K \pi_k \mathcal{N}(\mathbf{x} | \mu_k, \Sigma_k), \quad (6)$$

where π_k denotes the mixture weight, and μ_k, Σ_k are the mean and covariance of the k -th Gaussian component. This formulation allows us to capture fine-grained statistics compared to single-point constraints. Given clustered features $F_{h,n}^2$ and $F_{v,n}^2$, the encoder predicts the mean and variance (μ, σ^2) . By adopting the reparameterization trick

$$\mathbf{z} = \mu + \sigma \odot \epsilon, \quad \epsilon \sim \mathcal{N}(a, b). \quad (7)$$

The training objective of the RDA module consists of two components: the reconstruction loss and the KL divergence regularization, which are designed to ensure visual consistency of the enhanced results and alignment of feature distributions. Given the decoder outputs $F_{h,n}^3, F_{v,n}^3$ and the original input features $F_{h,n}^2, F_{v,n}^2$, we require them to remain consistent in the feature space. Accordingly, the reconstruction loss is defined as:

$$\mathcal{L}_s = \mathbb{E}_{q_\phi(\mathbf{z} | F^2)} [\|F_{h,n}^3 - F_{h,n}^2\|_2^2 + \|F_{v,n}^3 - F_{v,n}^2\|_2^2], \quad (8)$$

where $q_\phi(\mathbf{z} | F^2)$ denotes the posterior distribution parameterized by the encoder. This term constrains the decoder to faithfully recover the input features under low-light conditions, thereby enhancing detail fidelity and luminance consistency. To prevent the posterior distribution from straying too far from the prior, which is defined by the Gaussian Mixture Model (GMM), we introduce a KL divergence regularization term:

$$\mathcal{L}_d = D_{\text{KL}}(q_\phi(\mathbf{z} | F_{h,n}^2) \| p(\mathbf{z})) + D_{\text{KL}}(q_\phi(\mathbf{z} | F_{v,n}^2) \| p(\mathbf{z})), \quad (9)$$

where $p(\mathbf{z})$ represents the mixture prior distribution obtained via unsupervised clustering and Gaussian modeling. This regularization encourages the encoder outputs to better align with the global statistics, effectively suppressing instability of chrominance components and avoiding undesired color shifts or distribution collapse during enhancement.

In conventional frameworks, the posterior distribution parameters (μ, σ) are directly predicted by the encoder and regularized to align with the prior. However, such static alignment is insufficient in low-light scenarios, where chrominance distributions exhibit large variations and noise amplification becomes severe. To overcome this limitation, we introduce a reinforcement learning (RL) mechanism into the RDA module, enabling dynamic adjustment of Gaussian parameters guided by enhancement quality feedback. We regard (μ, σ) as adjustable parameters and employ an RL policy network to output correction terms $(\Delta\mu, \Delta\sigma)$, yielding the updated distribution:

$$\mu' = \mu \pm \Delta\mu, \quad \sigma' = \sigma \pm \Delta\sigma, \quad (10)$$

where $(\Delta\mu, \Delta\sigma)$ are actions sampled from the policy $\pi_\theta(a | s)$ given the current state s . The state encodes the current distribution fitting quality (e.g., KL divergence and histogram statistics), allowing the policy to flexibly adapt the distribution shape under different conditions. To directly link policy optimization with enhancement performance, we design the reward r as:

$$r = \frac{1}{\|F_h^3 - F_{hv,gt}^3\|_1}, \quad (11)$$

where F_h^3 represents the enhanced chrominance features, $F_{hv,gt}^3$ represents the ground truth chrominance features, and $\|\cdot\|_1$ denotes the L1 norm. The inverse of the L_1 loss encourages the policy to reduce the difference between the enhanced and ground truth features, thereby improving the enhancement quality.

270 We adopt the REINFORCE algorithm to optimize the policy network. The objective is defined as the expected
 271 cumulative reward:

$$\mathcal{J}(\theta) = \mathbb{E}_{a \sim \pi_\theta(\cdot|s)}[r], \quad (12)$$

273 where r represents the reward, which is the feedback signal given for a particular action, and π_θ is the policy
 274 network that determines the probability distribution of actions a given the state s . The parameter θ denotes the
 275 parameters of the policy network. The gradients of this objective are estimated as:

$$\nabla_\theta \mathcal{J}(\theta) \approx \frac{1}{N} \sum_{i=1}^N r_i \nabla_\theta \log \pi_\theta(a_i | s_i), \quad (13)$$

276 where N denotes the number of sampled actions, a_i is the action taken in the i -th sample, and s_i is the corre-
 277 sponding state for the action. This optimization procedure enables the policy to iteratively refine the distribution
 278 parameters according to the feedback received from the image quality, thus enhancing the model's stability and
 279 adaptability in low-light conditions.

282 In order to directly link policy optimization with the enhancement performance, we design the RL loss \mathcal{L}_r as:

$$\mathcal{L}_r = -\mathbb{E}_{a \sim \pi_\theta(\cdot|s)}[r \cdot \log \pi_\theta(a | s)], \quad (14)$$

284 which encourages the policy to minimize the discrepancy between the enhanced chrominance features and
 285 the ground truth, improving the enhancement quality by reducing the L_1 loss between the predicted and true
 286 chrominance values.

288 3.3 LOSS FUCTION

289 To constrain the training of the proposed framework, we employ a comprehensive loss that combines the pri-
 290 mary reconstruction loss in both the RGB and HVI spaces with the VCF and CDA losses. Specifically, let I_{out}
 291 and I_{gt} represent the enhanced and ground-truth images in the RGB domain, and let $I_{\text{out}}^{\text{HVI}}$ and $I_{\text{gt}}^{\text{HVI}}$ represent
 292 their counterparts in the HVI color space. The reconstruction loss is defined as:

$$\mathcal{L}_t = \|I_{\text{out}} - I_{\text{gt}}\|_1 + \lambda_1 \left\| I_{\text{out}}^{\text{HVI}} - I_{\text{gt}}^{\text{HVI}} \right\|_1, \quad (15)$$

294 where $\|\cdot\|_1$ denotes the ℓ_1 -norm. The total loss function is then formulated as:

$$\mathcal{L} = \mathcal{L}_t + \lambda_2 \mathcal{L}_s + \lambda_3 \mathcal{L}_d + \lambda_4 \mathcal{L}_r, \quad (16)$$

295 where $\lambda_1, \lambda_2, \lambda_3$ and λ_4 are weighting coefficients.

298 4 EXPERIMENTS

300 4.1 DATASETS AND SETTINGS

301 **Datasets.** To validate the effectiveness of the proposed method, we conduct experiments on seven LLIE
 302 benchmark datasets, including three paired datasets: LOLv1 Wei et al. (2018), LOLv2 Yang et al. (2021), and
 303 SICE Cai et al. (2018), and four unpaired datasets, including DICM Lv et al. (2018), LIME Guo et al. (2016),
 304 MEF Ma et al. (2015), NPE Wang et al. (2013), and VV Vonikakis et al. (2018). The LOLv1 dataset has 485
 305 paired training images and 15 for testing. LOLv2 consists of two subsets: LOLv2-Real (689 training, 100
 306 testing) and LOLv2-Synthetic (900 training, 100 testing). The SICE dataset includes 589 paired low-light and
 307 well-exposed images, with 100 randomly selected for testing and the rest for training and validation. For SID,
 308 we convert raw images to sRGB without gamma correction, resulting in extremely dark images. We crop the
 309 training images into 256×256 patches and train for 1,000 epochs with a batch size of 4.

310 **Experiment Settings.** We implement the proposed method using PyTorch and train all models on a single
 311 NVIDIA RTX 3090 GPU. The optimizer is Adam Kingma & Ba (2014) with parameters $\beta_1 = 0.9$ and $\beta_2 =$
 312 0.999. The initial learning rate is set to 1×10^{-4} and is gradually reduced to 1×10^{-7} using a cosine annealing
 313 schedule Loshchilov & Hutter (2016). During training, the batch size is consistently set to 8 and input images
 314 are cropped into 400×400 patches for all datasets except the LOLv2-Synthetic subset, for which full-resolution
 315 images are used without cropping. The $\lambda_1, \lambda_2, \lambda_3$ and λ_4 set to 1, 1, 0.5 and 0.5, respectively.

316 **Evaluation Metrics.** Following our baseline Yan et al. (2025), for paired datasets, we adopt Peak Signal-
 317 to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) Wang et al. (2004) as distortion-based metrics
 318 to evaluate reconstruction fidelity. To further assess the perceptual quality of the enhanced results, we re-
 319 port the Learned Perceptual Image Patch Similarity (LPIPS) Zhang et al. (2018), computed using a pretrained
 320 AlexNet Krizhevsky et al. (2012a). For *unpaired* datasets, we employ two no-reference image quality assess-
 321 ment metrics, BRISQUE Krizhevsky et al. (2012b) and NIQE Mittal et al. (2012), to evaluate perceptual real-
 322 ism. Moreover, to provide a comprehensive comparison, our method is benchmarked against 11 state-of-the-art
 323 supervised learning methods, including RetinexNet Wei et al. (2018), KinD Zhang et al. (2019), LLFlow Wang
 324 et al. (2022a), EnlightenGAN Jiang et al. (2021), SNR-Aware Xu et al. (2022), Bread Guo & Hu (2023), Pair-
 325 LIE Fu et al. (2023), LLFormer Wang et al. (2023), RetinexFormer Cai et al. (2023), GSAD Hou et al. (2023)
 326 and CIDNet Yan et al. (2025), as well as 3 unsupervised learning methods, such as ZeroDCE Guo et al. (2020),
 327 RUAS Liu et al. (2021), QuadPrior Wang et al. (2024), RetinexMamba Bai et al. (2024), UFormer Wang et al.
 328 (2022b), Restormer Zamir et al. (2022a) and MIRNet Zamir et al. (2022b) across all datasets.

324 Table 1: Quantitative results of PSNR \uparrow /SSIM \uparrow /LPIPS \downarrow on LOL (v1 and v2) datasets. Best performance
 325 in **purple**, second best in **cyan**.

326 327 Methods	328 329 330 331 332 333 334 335 336 337 338 339 340 LOLv1			328 329 330 331 332 333 334 335 336 337 338 339 340 LOLv2-Real			328 329 330 331 332 333 334 335 336 337 338 339 340 LOLv2-Synthetic		
	328 329 330 331 332 333 334 335 336 337 338 339 340 PSNR \uparrow	328 329 330 331 332 333 334 335 336 337 338 339 340 SSIM \uparrow	328 329 330 331 332 333 334 335 336 337 338 339 340 LPIPS \downarrow	328 329 330 331 332 333 334 335 336 337 338 339 340 PSNR \uparrow	328 329 330 331 332 333 334 335 336 337 338 339 340 SSIM \uparrow	328 329 330 331 332 333 334 335 336 337 338 339 340 LPIPS \downarrow	328 329 330 331 332 333 334 335 336 337 338 339 340 PSNR \uparrow	328 329 330 331 332 333 334 335 336 337 338 339 340 SSIM \uparrow	328 329 330 331 332 333 334 335 336 337 338 339 340 LPIPS \downarrow
RetinexNet Wei et al. (2018)	18.915	0.427	0.470	16.097	0.401	0.543	17.137	0.762	0.255
Kind Zhang et al. (2019)	23.018	0.843	0.156	17.544	0.669	0.375	18.320	0.796	0.252
ZeroDCE Guo et al. (2020)	21.880	0.640	0.335	16.059	0.580	0.313	17.712	0.815	0.169
RUAS Liu et al. (2021)	18.654	0.518	0.270	15.326	0.488	0.176	13.765	0.638	0.305
EnlightenGAN Jiang et al. (2021)	20.003	0.691	0.317	18.230	0.617	0.309	16.570	0.734	0.220
LLFlow Wang et al. (2022a)	24.998	0.871	0.117	17.433	0.831	0.315	24.870	0.919	0.067
UFormer Wang et al. (2022b)	19.610	0.755	0.197	19.410	0.657	0.194	19.660	0.871	0.075
Restormer Zamir et al. (2022a)	22.430	0.823	0.184	19.940	0.827	0.183	21.410	0.830	0.062
MIRNet Zamir et al. (2022b)	24.140	0.830	0.154	20.020	0.820	0.175	21.940	0.876	0.058
SNR-Aware Xu et al. (2022)	26.716	0.851	0.152	21.480	0.849	0.163	24.140	0.928	0.056
Bread Guo & Hu (2023)	25.299	0.847	0.155	20.830	0.847	0.174	17.630	0.919	0.091
PairLIE Fu et al. (2023)	23.526	0.755	0.248	19.855	0.778	0.317	19.074	0.794	0.230
LLFormer Wang et al. (2023)	25.758	0.823	0.167	20.056	0.792	0.211	24.038	0.909	0.066
RetinexFormer Cai et al. (2023)	27.140	0.850	0.129	22.794	0.840	0.171	25.670	0.930	0.059
GSAD Hou et al. (2023)	27.605	0.876	0.092	20.153	0.846	0.113	24.472	0.929	0.051
QuadPrior Wang et al. (2024)	22.849	0.800	0.201	20.592	0.811	0.202	16.108	0.758	0.114
RetinexMamba Bai et al. (2024)	24.030	0.827	0.146	22.450	0.844	0.174	25.890	0.935	0.054
CIDNet Yan et al. (2025)	28.201	0.889	0.079	24.111	0.871	0.108	25.705	0.942	0.045
Ours	29.123	0.880	0.075	24.892	0.875	0.110	26.512	0.947	0.040

4.2 RESULTS ON PAIRED DATASETS

We evaluate our proposed method on three widely-used low-light image enhancement benchmarks: LOLv1, LOLv2-Real, and LOLv2-Synthetic. As illustrated in Table 1, our approach consistently achieves the best performance across PSNR, SSIM, and LPIPS metrics. In contrast, existing methods exhibit notable weaknesses: RUAS and LLFlow often produce over-smoothed or distorted textures, PairLIE and GSAD suffer from unstable color rendition with visible hue shifts, while CIDNet fails to adequately suppress residual noise, resulting in unnatural tone mapping.

Table 2: Complexity of Different Methods.

	ZeroDCE	RUAS	LLFlow	EnlightenGAN	SNR-Aware	Bread	PairLIE	LLFormer
Params/M	0.075	0.003	17.42	114.35	4.01	2.02	0.33	24.55
FLOPs/G	4.83	0.83	358.4	61.01	26.35	19.85	20.81	22.52

Benefiting from the Reliable Cross Attention (RCA) module and Reinforced Distribution Alignment (RDA) module, our method effectively suppresses noisy or inconsistent features and ensures better alignment between luminance and chromaticity statistics. This enables more balanced illumination and natural color reproduction, particularly in extremely dark regions. Quantitatively, our approach achieves a PSNR of 29.123 on LOLv1, outperforming the previous state-of-the-art by 0.771 dB, and reaches 24.892 on LOLv2-Real with an SSIM improvement of 0.022 over the second best. On LOLv2-Synthetic, our method attains 26.512 PSNR and 0.947 SSIM, both ranking first. These results highlight the effectiveness and robustness of our design, which delivers state-of-the-art visual quality while maintaining competitive computational efficiency. Furthermore, compared with CIDNet, the proposed method introduces only a marginal increase in parameters (+0.28M) and FLOPs (+2.96G2), yet consistently outperforms it across multiple benchmark metrics, which validates the efficiency of our design and demonstrates its favorable trade-off between complexity and performance.

Figure 3 presents a qualitative comparison of the enhancement results on the LOL (v1 and v2) and SICE datasets, showcasing the performance of our method in comparison to several state-of-the-art approaches, including RUAS, LLFlow, PairLIE, GSAD, EnlightenGAN, RetinexFormer, and CIDNet. As seen in the figure, our method demonstrates superior performance in enhancing low-light images, effectively improving both brightness and contrast while preserving details in the images. Our approach outperforms the competing meth-

Table 3: Quantitative result on SID, SICE and the five unpaired datasets (DICM, LIME, MEF, NPE, and VV). The top-ranking score is in **Bold**.

370 371 Methods	372 373 374 375 376 377 SICE		372 373 374 375 376 377 SID		372 373 374 375 376 377 Unpaired	
	372 373 374 375 376 377 PSNR \uparrow	372 373 374 375 376 377 SSIM \uparrow	372 373 374 375 376 377 PSNR \uparrow	372 373 374 375 376 377 SSIM \uparrow	372 373 374 375 376 377 BRISQUE \downarrow	372 373 374 375 376 377 NIQE \downarrow
RetinexNet	12.424	0.613	15.695	0.395	23.286	4.558
ZeroDCE	12.452	0.639	14.087	0.090	26.343	4.763
RUAS	8.656	0.494	12.622	0.081	26.372	4.800
LLFlow	12.737	0.617	16.226	0.367	26.087	4.221
CIDNet	13.435	0.642	22.904	0.676	23.521	3.523
Ours	16.195	0.714	23.116	0.727	22.894	3.417

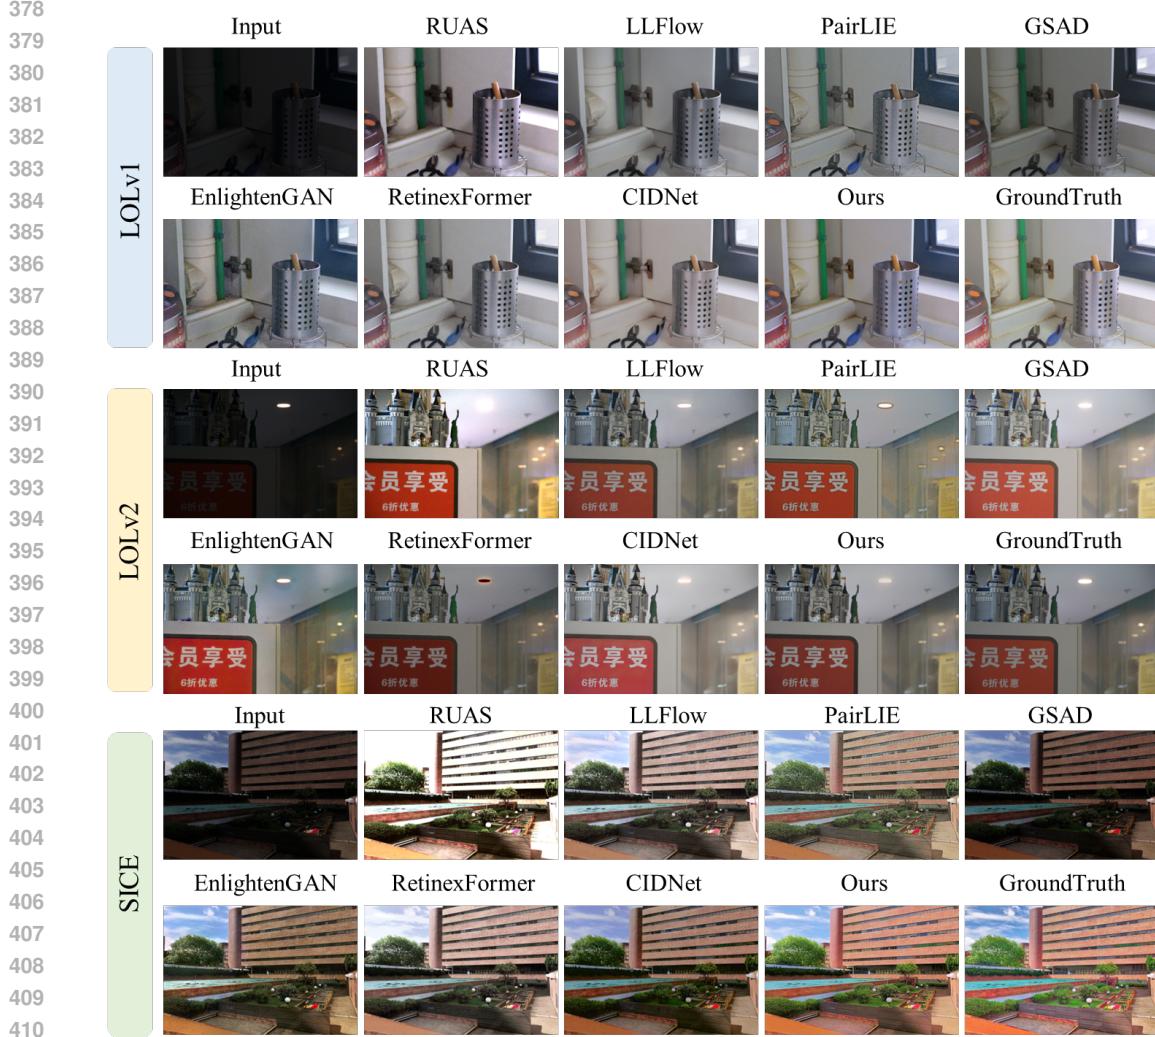


Figure 3: Qualitative comparison of enhancement results on the LOL (v1 and v2) and SICE dataset, generated by various methods.

ods in terms of naturalness and color fidelity, particularly in challenging low-light scenarios. The enhanced images show clearer details and reduced noise, providing a more visually pleasing result compared to the other methods, as highlighted in the visual comparison. The effectiveness of our method is particularly evident in the SICE dataset, where it significantly enhances the image quality without introducing noticeable artifacts.

4.3 RESULTS ON UNPAIRED DATASETS

We conduct comprehensive evaluations on unpaired datasets. For the unpaired datasets (DICM, LIME, MEF, NPE, and VV), we report two widely used no-reference perceptual quality measures, BRISQUE and NIQE,



Figure 4: Qualitative comparison of enhancement results on the unpaired dataset, generated by various methods.

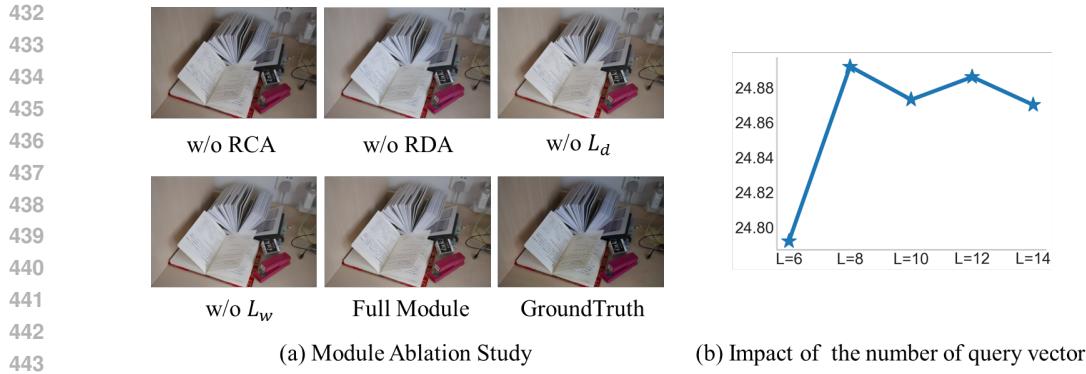


Figure 5: Ablation study on the LOLv2-real dataset.

to evaluate visual realism. As shown in Table 3, our method consistently achieves the best results across all metrics. Specifically, it outperforms existing supervised methods such as RetinexNet, LLFlow, and CIDNet, as well as unsupervised approaches including ZeroDCE and RUAS. On the five unpaired datasets, our model also yields the lowest BRISQUE and NIQE scores, indicating superior perceptual quality. These results demonstrate that our method not only preserves structural fidelity in paired scenarios but also generalizes effectively to real-world unpaired conditions.

4.4 ABLATION STUDIES

To validate the effectiveness of each component in our proposed framework, we perform a series of ablation experiments on the LOLv2-real dataset. As summarized in Table 4, the model’s performance is assessed using PSNR and SSIM metrics. In the ablation studies, we sequentially remove individual modules to analyze their contributions. The removal of the RCA results in a noticeable decline in performance, underscoring the importance of adaptively aggregating luminance and chromaticity features. Additionally, excluding the RDA module or our design loss L_d or L_s leads to further performance degradation, highlighting the critical role of precise chromaticity distribution modeling in achieving more natural and accurate color rendering.

Table 4: Ablation studies of modules, w means with and w/o means without.

exp.	RCA	RDA	L_d	L_s	PSNR↑	SSIM↑
1	w/o	w/o	w/o	w/o	24.111	0.871
2	w	w/o	w/o	w/o	24.459	0.872
3	w	w	w/o	w/o	24.681	0.874
4	w	w	w	w/o	24.792	0.875
5	w	w	w	w	24.892	0.875

Moreover, as demonstrated in Figure 5(a), removing the RDA module causes significant degradation in color fidelity, as it eliminates the mechanism for aligning chromaticity distributions. The absence of the RCA module leads to substantial deterioration in both luminance and chrominance quality, highlighting its critical role in suppressing noise and maintaining balanced illumination and accurate color restoration during enhancement.

In the ablation study shown in Figure 5(b), we evaluate the impact of the number of query vectors, L , on performance. Increasing L from 6 to 8 leads to significant improvement, with diminishing returns as L exceeds 8. This suggests that $L = 8$ offers a good balance between performance and computational efficiency.

4.5 CONCLUSION

In this work, we introduce a novel low-light image enhancement framework, Learning to Enhance Low-Light Images with Reliable Attention and Reinforced Distribution Alignment, which effectively addresses the challenges of noise amplification and color distortion in low-light conditions. Our framework combines the Reliable Cross Attention (RCA) module and the Reinforced Distribution Alignment (RDA) module to improve the interaction between luminance and chromaticity features while preserving color fidelity and naturalness. The RCA module suppresses redundant features and enhances the efficiency of feature aggregation, while the RDA module refines chromaticity distributions through unsupervised clustering and reinforcement learning. Extensive experiments on ten benchmark datasets demonstrate that our method achieves state-of-the-art performance, offering superior visual quality and strong generalization across various lighting conditions. Our results confirm that VCR is an effective and efficient solution for low-light image enhancement, pushing the boundaries of image quality under challenging conditions.

486
487
ETHICS STATEMENT488
489
490
491
492
493
494
495
This research adheres to the ICLR Code of Ethics. We ensure that no ethical violations have occurred during
the research process. All datasets used comply with publicly available privacy policies, and we have ensured
the security and privacy of the data during collection and use. There are no conflicts of interest or funding issues
in this research. All methods and applications used in this research follow principles of fairness and objectivity
to ensure the integrity and transparency of the research.496
497
REPRODUCIBILITY STATEMENT
498
499
500501
502
503
504
505
All improvements in this research are based on open-source code and datasets. We provide comprehensive
experimental details and algorithm descriptions, including the models, datasets, and training processes used. All
relevant source code and datasets will be made open-source. We encourage readers to use the same experimental
setups and parameters to reproduce our results and validate the theories and algorithms presented in this work,
ensuring the reproducibility of the research and supporting the validation of the results.506
507
REFERENCES
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

Jiesong Bai, Yuhao Yin, Qiyuan He, Yuanxian Li, and Xiaofeng Zhang. Retinexmamba: Retinex-based mamba for low-light image enhancement. In *International Conference on Neural Information Processing*, pp. 427–442. Springer, 2024.

Jianrui Cai, Shuhang Gu, and Lei Zhang. Learning a deep single image contrast enhancer from multi-exposure images. *IEEE Transactions on Image Processing*, 27(4):2049–2062, 2018.

Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 12504–12513, 2023.

Zhixin Cheng, Jiacheng Deng, Xinjun Li, Xiaotian Yin, Bohao Liao, Baoqun Yin, Wenfei Yang, and Tianzhu Zhang. Ca-i2p: Channel-adaptive registration network with global optimal selection. *arXiv preprint arXiv:2506.21364*, 2025.

Christoph Feichtenhofer, Axel Pinz, and Andrew Zisserman. Detect to track and track to detect. In *Proceedings of the IEEE international conference on computer vision*, pp. 3038–3046, 2017.

Zhenqi Fu, Yan Yang, Xiaotong Tu, Yue Huang, Xinghao Ding, and Kai-Kuang 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*, pp. 22252–22261, 2023.

Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 1780–1789, 2020.

Xiaojie Guo and Qiming Hu. Low-light image enhancement via breaking down the darkness. *International Journal of Computer Vision*, 131(1):48–66, 2023.

Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE Transactions on image processing*, 26(2):982–993, 2016.

Xifeng Guo, En Zhu, Xinwang Liu, and Jianping Yin. Deep embedded clustering with data augmentation. In *Asian conference on machine learning*, pp. 550–565. PMLR, 2018.

Jinhui Hou, Zhiyu Zhu, Junhui Hou, Hui Liu, Huanqiang Zeng, and Hui Yuan. Global structure-aware diffusion process for low-light image enhancement. *Advances in Neural Information Processing Systems*, 36:79734–79747, 2023.

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

Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012a.

540 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural
 541 networks. *Advances in neural information processing systems*, 25, 2012b.

542

543 Edwin H Land and John J McCann. Lightness and retinex theory. *Journal of the Optical society of America*,
 544 61(1):1–11, 1971.

545

546 Risheng Liu, Long Ma, Jiaao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative
 547 prior architecture search for low-light image enhancement. In *Proceedings of the IEEE/CVF conference on*
 548 *computer vision and pattern recognition*, pp. 10561–10570, 2021.

549

550 Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint*
 551 *arXiv:1608.03983*, 2016.

552

553 Feifan Lv, Feng Lu, Jianhua Wu, and Chongsoon Lim. Mblen: Low-light image/video enhancement using
 554 cnns. In *Bmvc*, volume 220, pp. 4. Northumbria University, 2018.

555

556 Kede Ma, Kai Zeng, and Zhou Wang. Perceptual quality assessment for multi-exposure image fusion. *IEEE*
 557 *Transactions on Image Processing*, 24(11):3345–3356, 2015.

558

559 Anish Mittal, Rajiv Soundararajan, and Alan C Bovik. Making a “completely blind” image quality analyzer.
 560 *IEEE Signal processing letters*, 20(3):209–212, 2012.

561

562 Stephen M Pizer, E Philip Amburn, John D Austin, Robert Cromartie, Ari Geselowitz, Trey Greer, Bart ter
 563 Haar Romeny, John B Zimmerman, and Karel Zuiderveld. Adaptive histogram equalization and its varia-
 564 tions. *Computer vision, graphics, and image processing*, 39(3):355–368, 1987.

565

566 Shanto Rahman, Md Mostafijur Rahman, Mohammad Abdullah-Al-Wadud, Golam Dastegir Al-Quaderi, and
 567 Mohammad Shoyaib. An adaptive gamma correction for image enhancement. *EURASIP Journal on Image*
 568 *and Video Processing*, 2016(1):35, 2016.

569

570 Zia-ur Rahman, Daniel J Jobson, and Glenn A Woodell. Retinex processing for automatic image enhancement.
 571 *Journal of Electronic imaging*, 13(1):100–110, 2004.

572

573 Carl Rasmussen. The infinite gaussian mixture model. *Advances in neural information processing systems*, 12,
 574 1999.

575

576 Vassilios Vonikakis, Rigas Kouskouridas, and Antonios Gasteratos. On the evaluation of illumination compen-
 577 sation algorithms. *Multimedia Tools and Applications*, 77(8):9211–9231, 2018.

578

579 Shuhang Wang, Jin Zheng, Hai-Miao Hu, and Bo Li. Naturalness preserved enhancement algorithm for non-
 580 uniform illumination images. *IEEE transactions on image processing*, 22(9):3538–3548, 2013.

581

582 Tao Wang, Kaihao Zhang, Tianrun Shen, Wenhan Luo, Bjorn Stenger, and Tong Lu. Ultra-high-definition
 583 low-light image enhancement: A benchmark and transformer-based method. In *Proceedings of the AAAI*
 584 *conference on artificial intelligence*, volume 37, pp. 2654–2662, 2023.

585

586 Wenjing Wang, Huan Yang, Jianlong Fu, and Jiaying Liu. Zero-reference low-light enhancement via physical
 587 quadruple priors. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*,
 588 pp. 26057–26066, 2024.

589

590 Yufei Wang, Renjie Wan, Wenhan Yang, Haoliang Li, Lap-Pui Chau, and Alex Kot. Low-light image enhance-
 591 ment with normalizing flow. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36,
 592 pp. 2604–2612, 2022a.

593

594 Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer:
 595 A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF conference on*
 596 *computer vision and pattern recognition*, pp. 17683–17693, 2022b.

597

598 Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error
 599 visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

600

601 Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhance-
 602 ment. *arXiv preprint arXiv:1808.04560*, 2018.

603

604 Junyuan Xie, Ross Girshick, and Ali Farhadi. Unsupervised deep embedding for clustering analysis. In *Inter-
 605 national conference on machine learning*, pp. 478–487. PMLR, 2016.

606

607 Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *Pro-
 608 ceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 17714–17724, 2022.

594 Qingsen Yan, Yixu Feng, Cheng Zhang, Guansong Pang, Kangbiao Shi, Peng Wu, Wei Dong, Jinqiu Sun, and
 595 Yanning Zhang. Hvi: A new color space for low-light image enhancement. In *Proceedings of the Computer*
 596 *Vision and Pattern Recognition Conference*, pp. 5678–5687, 2025.

597 Wenhan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient regularized
 598 deep retinex network for robust low-light image enhancement. *IEEE Transactions on Image Processing*, 30:
 599 2072–2086, 2021.

600 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan
 601 Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the*
 602 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 5728–5739, 2022a.

603 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang,
 604 and Ling Shao. Learning enriched features for fast image restoration and enhancement. *IEEE transactions*
 605 *on pattern analysis and machine intelligence*, 45(2):1934–1948, 2022b.

606 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness
 607 of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and*
 608 *pattern recognition*, pp. 586–595, 2018.

609 Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer.
 610 In *Proceedings of the 27th ACM international conference on multimedia*, pp. 1632–1640, 2019.

611 Zhengxia Zou, Keyan Chen, Zhenwei Shi, Yuhong Guo, and Jieping Ye. Object detection in 20 years: A survey.
 612 *Proceedings of the IEEE*, 111(3):257–276, 2023.

615 A APPENDIX

617 A.1 HVI COLOR SPACE

618 In the standard sRGB color space, image brightness and chromatic information are tightly coupled across the
 619 three color channels, which may disrupt the perceived illumination or color balance of the entire image when
 620 making adjustments to any individual channel. Although the HSV color space separates intensity from chro-
 621 maticity, it inadvertently amplifies noise in regions of extreme red and near-black areas, producing pronounced
 622 “red-discontinuity” and “black-plane” artifacts during enhancement. To address the above limitations, the HVI
 623 color space has been proposed to alleviate inherent color noise, which is composed of three channels: I_{\max} ,
 624 \hat{H} , and \hat{V} , designed to mitigate the artifacts introduced by the HSV representation. Here, $C_k(x)$ denotes a
 625 learnable intensity collapse function that remaps the maximum intensity $I_{\max}(x)$ for stabilizing low-light re-
 626 sponds. The parameter k , termed *density-k*, controls the density of black-plane points in HVT/PHVIT, thereby
 627 balancing noise suppression and detail preservation. According to the Max-RGB, for each individual pixel x ,
 628 the intensity map of image I can be estimated:

$$629 I_{\max}(x) = \max_{c \in \{R, G, B\}} I_c(x). \quad (17)$$

630 Meanwhile, according to the sRGB-HSV transformation, the saturation s of the image can be obtained:

$$631 s = \begin{cases} 0, & I_{\max} = 0 \\ \frac{I_{\max} - \min(I_c)}{I_{\max}}, & I_{\max} \neq 0 \end{cases} \quad (18)$$

632 and the hue h of the image is formulated as follows:

$$633 h = \begin{cases} 0, & s = 0 \\ \left(\frac{I_G - I_B}{I_{\max} - \min(I_c)} \right) \bmod 6, & I_{\max} = I_R \\ 2 + \frac{I_B - I_R}{I_{\max} - \min(I_c)}, & I_{\max} = I_G \\ 4 + \frac{I_R - I_G}{I_{\max} - \min(I_c)}, & I_{\max} = I_B \end{cases} \quad (19)$$

634 where s and h correspond to any pixel in the saturation map $S(x)$ and hue map $H(x)$, respectively. Moreover,
 635 corresponding to HVT in Figure 2, the horizontal chromaticity component $\hat{H}(x)$ and the vertical component
 636 $\hat{V}(x)$ are constructed by polarizing the hue angle from HSV into Cartesian space, defined as:

$$637 \begin{cases} \hat{H}(x) = C_k(x) \cdot S(x) \cdot \cos\left(\frac{\pi H(x)}{3}\right), \\ \hat{V}(x) = C_k(x) \cdot S(x) \cdot \sin\left(\frac{\pi H(x)}{3}\right), \end{cases} \quad (20)$$

648 where $C_k(x)$ is a learnable intensity collapse function defined as:
 649

$$650 \quad C_k(x) = k \cdot \sqrt{\sin\left(\frac{\pi I_{\max}(x)}{2}\right) + \varepsilon}, \quad (21)$$

651 with k as a trainable parameter and ε as a small constant (set to 10^{-8}) for training stability. Moreover, as shown
 652 in Fig. 2, the Perceptual-inverse HVI Transformation (PHVIT) is performed to convert the HVI space back to
 653 HSV. The hue $H(x)$, saturation $S(x)$, and value $V(x)$ are estimated as:
 654

$$655 \quad \begin{cases} H(x) = \frac{1}{2\pi} \cdot \arctan\left(\frac{\hat{v}(x)}{\hat{h}(x)}\right) \bmod 1, \\ 656 \quad S(x) = \alpha_S \cdot \sqrt{\hat{h}^2(x) + \hat{v}^2(x)}, \\ 657 \quad V(x) = \alpha_I \cdot I_{\max}(x), \end{cases} \quad (22)$$

658 where α_S and α_I are linear scaling parameters that control the output image's saturation and brightness, re-
 659 spectively. The normalized intermediate chromaticity coordinates are computed as:
 660

$$661 \quad \begin{cases} \hat{h}(x) = \frac{\hat{H}(x)}{C_k(x) + \varepsilon}, \\ 662 \quad \hat{v}(x) = \frac{\hat{V}(x)}{C_k(x) + \varepsilon}. \end{cases} \quad (23)$$

665 A.2 BLIND/REFERENCELESS IMAGE SPATIAL QUALITY EVALUATOR (BRISQUE)

666 BRISQUE Krizhevsky et al. (2012b) is a blind image quality assessment method that leverages natural scene
 667 statistics (NSS) in the spatial domain. It normalizes local luminance values and characterizes their distribution
 668 through an asymmetric generalized Gaussian distribution (AGGD). From this distribution, descriptive statistics
 669 such as shape and variance are derived. A support vector regression model, trained with subjective quality
 670 annotations, maps these statistics to a perceptual quality score. Lower BRISQUE values correspond to higher
 671 visual quality. Since it operates without reference images, BRISQUE is particularly suitable for evaluating
 672 real-world or unpaired data.

674 A.3 NATURALNESS IMAGE QUALITY EVALUATOR (NIQE)

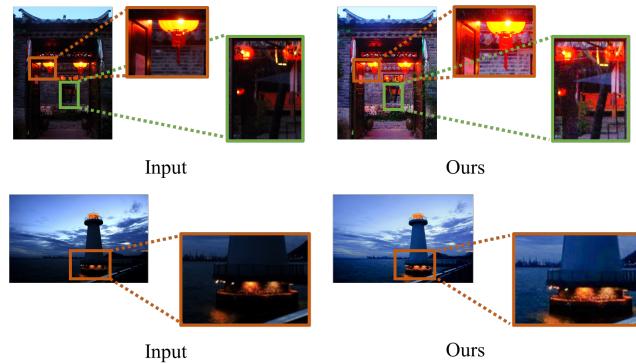
675 NIQE Mittal et al. (2012) is another no-reference quality assessment approach, designed to capture deviations
 676 from the statistical regularities of natural images. It first constructs a multivariate Gaussian model using NSS-
 677 based features (e.g., mean-subtracted contrast-normalized coefficients and local pixel correlations) extracted
 678 from pristine natural images. For a test image, the same features are computed, and the Mahalanobis distance
 679 to the Gaussian model is used as the quality score:

$$680 \quad \text{NIQE}(I) = \sqrt{(f - \mu)^\top \Sigma^{-1} (f - \mu)},$$

681 where f represents the feature vector of the test image, and μ, Σ denote the mean and covariance estimated
 682 from natural data. A lower NIQE value implies stronger alignment with natural image statistics and thus better
 683 perceptual quality.

685 A.4 FAILURE CASES

686 Figure 6 presents typical failure cases on the unpaired dataset (DICM). In extreme scenarios, our method
 687 occasionally struggles with inadequate brightness restoration and insufficient noise suppression. We plan to
 688



700 Figure 6: Failure cases.
 701

702 address these limitations in future work by integrating explicit noise priors, enforcing temporal consistency in
703 video sequences, and introducing locally adaptive color temperature adjustment.
704

705 **A.5 USE OF LARGE MODELS**
706

707 In this work, large language models are employed solely for language polishing and improving the readability
708 of the manuscript. They are not involved in problem formulation, algorithm design, model implementation, or
709 experimental analysis. All technical contributions and experimental results are independently developed and
710 verified by the authors.
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755