Employing Multiple Priors in Retinex-Based Low-Light Image Enhancement

Weipeng Yang[1] , Hongxia Gao[†,1,2] , Tongtong Liu[1] , Jianliang Ma[1,3] , Wenbin Zou[1] , Shasha Huang[1]

[1]School of Automation Science and Engineering, South China University of Technology, Guangzhou, China
[2]Research Center for Brain-Computer Interface, Pazhou Laboratory, Guangzhou, China
[3]KUKA Robotics Guangdong Co., Ltd., Foshan, China

Abstract
In the field of low-light image enhancement, images captured under low illumination suffer from severe noise and artifacts, which are often exacerbated during the enhancement process. Our method, grounded in the Retinex theory, tackles this challenge by recognizing that the illuminance component predominantly contains low-frequency image information, whereas the reflectance component encompasses high-frequency details, including noise. To effectively suppress noise in the reflectance without compromising detail, our method uniquely amalgamates global, local, and non-local priors. It utilizes the tensor train rank for capturing global features along with two plug-and-play denoisers: a convolutional neural network and a Color Block-Matching 3D filter (CBM3D), to preserve local details and non-local self-similarity. Furthermore, we employ Proximal Alternating Minimization (PAM) and the Alternating Direction Mthod of Multipliers (ADMM) algorithms to effectively separate the reflectance and illuminance components in the optimization process. Extensive experiments show that our model achieves superior or competitive results in both visual quality and quantitative metrics when compared with state-of-the-art methods. Our code is available at https://github.com/YangWeipengscut/GLON-Retinex.

CCS Concepts
• Computing methodologies → Image processing;

1. Introduction

In the realm of low-light image enhancement (LLIE), the quest for optimal methods that proficiently balance noise suppression with detail preservation remains paramount. In the domain of optimization-based methods, algorithms based on the Retinex theory [Lan77] aim to decompose the observed image S into illuminance L and reflectance R, according to the Retinex theory formula $S = L \circ R$, where $\circ$ denotes element-wise multiplication. The illuminance L represents the influence of light source intensity and color on the object's color, which is typically piecewise smooth. while the reflectance R represents the object's inherent detail information and also contains noise present in low-light images. Subsequently, the enhancement result is obtained by applying gamma correction to the illuminance and multiplying it by the denoised reflectance component [LLY*18, LL22, YGZ*23].

When it comes to noise reduction in low-light images, there are typically two approaches. One involves denoising the image after enhancement, but this method struggles to eliminate noise that has already been amplified during the enhancement process [GLL16]. The other, and more commonly employed approach, is to leverage prior knowledge about the image and noise to construct noise removal regularizers pertaining to the reflectance and establish a unified optimization objective function. This allows for the simultaneous estimation of the illuminance and reflectance, and such methods are referred to as variational Retinex-based methods [XHR*20].

In the construction of the regularizers for the reflectance, the most commonly employed approach is the use of singular local prior knowledge from the image. Methods such as WVM [FZH*16b], JieP [CXG*17], SRRR [LLY*18], STAR [XHR*20], and PnPRetinex [LL22] implement total variation (TV), relative total variation, or exponential relative total variation to formulate the regularizers for the reflectance. Furthermore, the BTRetinex [YGZ*23] method pioneered the use of dictionary learning, a type of local prior, for noise removal in the reflectance. Local priors are geared towards preserving details in smaller regions and maintaining edges and textures, yet they might result in inconsistencies and are less effective against uniformly distributed noise. Therefore, in addition to utilizing TV regularizers, the LR3M [RYCL20] method innovatively employs non-local self-similarity as the regularizers for the reflectance. Non-local priors exploit the similarity in distant regions, demonstrating superior performance in noise suppression. However, the use of non-local self-similarity priors often faces

---

† Corresponding author

Figure 1: *Close-up comparative analysis of enhanced results on the LSRW [HXY*23] dataset. Our proposed method integrating global-local-nonlocal priors excels in brightness enhancement, noise suppression, structural information preservation, and color fidelity. **Zoom in for a better view**.*

challenges like low computational efficiency, substantial memory consumption, and a tendency to cause loss of image details.

In addition, the global prior, such as inherent low-rankness, is often neglected by existing variational Retinex-based LLIE methods. However, natural images can typically be well-approximated with low-rank structures [ZHZ*19]. Leveraging this trait can effectively eliminate noise that doesn't adhere to this low-rank feature from the reflectance. Furthermore, current variational Retinex-based LLIE methods often process RGB images channel-by-channel or convert them to the HSV color space, focusing solely on the V channel. This approach, including matrix vectorization, typically leads to the loss of spatial and spectral relationships [YGZ*23]. Tensors, which extend vectors and matrices into higher dimensions, offer a superior way to represent high-dimensional data. By using tensors, the inherent high-dimensional nature of data is preserved, along with its contiguous structural information, as highlighted in [CYZ17]. This is particularly relevant as most real-world data are characterized by multidimensionality that goes beyond two orders. For example, RGB images are third-order tensors with dimensions of height, width, and color channel, while videos are fourth-order tensors, adding time as an extra dimension.

The remarkable advancements of deep learning technologies in the field of image processing have propelled the application of learning-based methods in LLIE, such as URetinex [WWZ*22], ZeroDCE [GLG*20], and SCI [MML*22]. Particularly, with the growing prominence of flow and diffusion models, methods such as LLFlow [WWY*22] (based on flow models) and diffusion model-based DIFFLL [JLF*23] and GSAD [HZH*23] have been introduced. To date, DIFFLL and GSAD have achieved leading results on paired datasets like LOL-v1/2 [WWYL18, YWH*21]. However, these LLIE methods often face challenges with generalization or require extensive training. Their success depends on accurate physical modeling and diverse training data, which may not align well with the vast range of real-world low-light conditions. Consequently, this can limit their effectiveness on unpaired datasets,

potentially resulting in color distortion, overexposure, noise amplification, or artifacts in the output.

Building on the need to address these challenges, our approach aims to overcome the limitations of existing LLIE methods, offering four key contributions.

1) **Integration of Global-Local-Nonlocal Priors:** We adopted a strategy that integrates global tensor low-rankness, based on tensor train rank [ZYM*21], with local and nonlocal self-similarity priors within a plug-and-play ADMM [CWE16] framework. Utilizing denoisers CBDNet [GYZ*19] for local detail preservation and CBM3D [DFKE07] for nonlocal similarity, our approach blends model-based and learning-based strengths. This integration effectively maintains detail preservation while reducing noise, enhancing performance in diverse low-light environments.

2) **Adaptive Visual Data Tensorization:** To apply tensor train rank priors to RGB images, we transform them into higher-dimensional tensors using Visual Data Tensorization (VDT) [YZGC19]. To overcome the limitations of traditional VDT, which necessitates equal factor counts in an image's height and width dimensions, we designed an adaptive VDT method. This innovative technique enables tensorization of images with arbitrary dimensions, thereby resolving the constraints that previously limited the processing of certain image sizes.

3) **Innovative Algorithm Design:** In order to tackle the complexities of multivariate optimization and decouple the illumination and reflectance components, we designed a stable solution algorithm employing Proximal Alternating Minimization (PAM) [ABS11] and Alternating Direction Method of Multipliers (ADMM) [LLS11].

4) **Superior Experimental Performance:** We designed numerous experiments, with results revealing that our approach surpasses existing LLIE methods in several key aspects such as brightness enhancement, noise reduction, and detail retention under diverse low-light conditions.
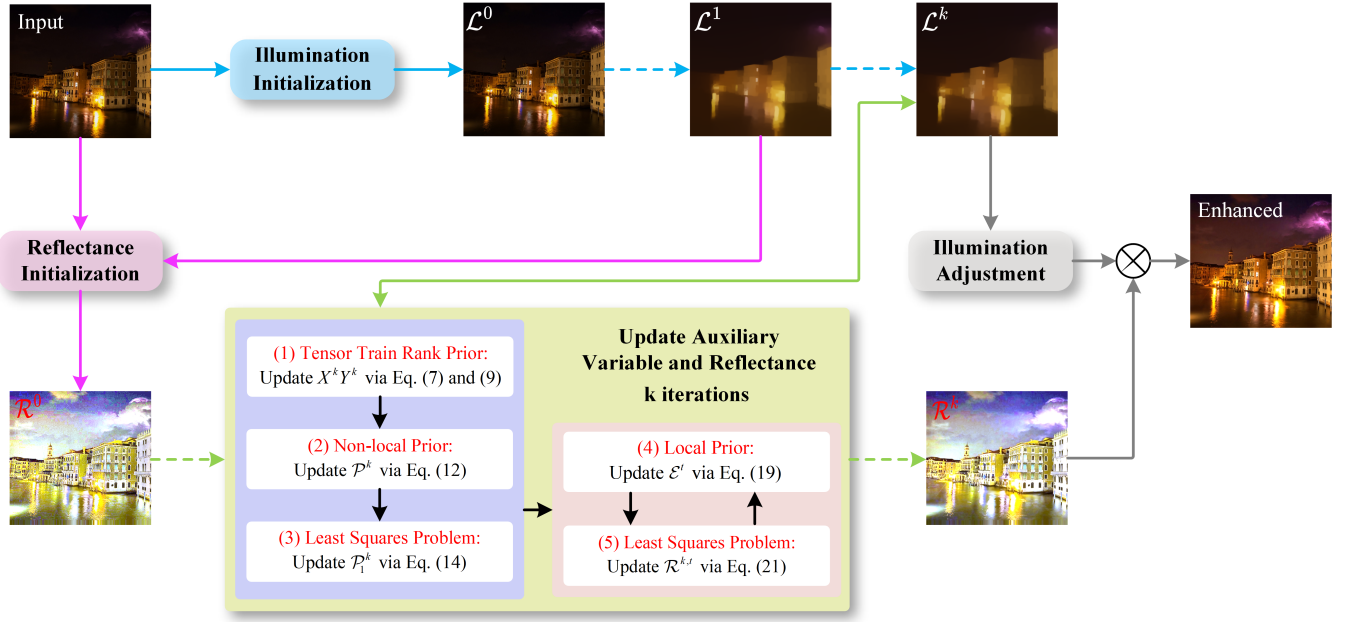
Figure 2: *Schematic diagram of our proposed model's framework. Initially, a low-light input image is used as the initial value for illuminance $\mathcal{L}^0$. After computing $\mathcal{L}^1$, the initial reflectance $\mathcal{R}^0$ is obtained by element-wise division of the low-light input image by $\mathcal{L}^1$. This value is then processed through the auxiliary variable and reflectance update module. The iteration stops when the maximum number of iterations K is reached or when the iteration error is less than δ. The final enhanced image is obtained by element-wise multiplication of the corrected illuminance and reflectance values.*

2. Our method

2.1. Adaptive Visual Data Tensorization

In this paper, an *N*-order tensor is denoted by an Euler script letter, for instance, $\mathcal{Z} \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_N}$. A matrix, representing a 2-order tensor, is denoted by a capital letter, such as *Z*.

Traditional VDT requires equal factorization of the spatial dimensions *m* and *n* in a third-order visual data tensor $\mathcal{Z} \in \mathbb{R}^{m \times n \times p}$, a constraint often impractical for images with incompatible dimensions. Adaptive VDT addresses this by padding image dimensions with edge pixels, allowing for uniform factorization.

The adaptive process can be explained as follows: Given the original dimensions *m* and *n* of an image, the image is augmented to dimensions $\tilde{m}$ and $\tilde{n}$ through edge pixel replication until $\tilde{m} = \prod_{d=1}^{\tilde{q}} \tilde{m}_d$ and $\tilde{n} = \prod_{d=1}^{\tilde{q}} \tilde{n}_d$ are achieved, with $\tilde{q}$ representing the new equal number of factors for both $\tilde{m}$ and $\tilde{n}$. We then permute and reshape these dimensions for tensorization, transforming $\tilde{m} \times \tilde{n} \rightarrow \tilde{m}_1 \times \tilde{n}_1 \times \tilde{m}_2 \times \tilde{n}_2 \times \ldots \times \tilde{m}_{\tilde{q}} \times \tilde{n}_{\tilde{q}}$. The tensor is then permuted and reshaped into a higher-order form $\mathcal{Z}'$, with dimensions $\tilde{m}_1 \tilde{n}_1 \times \ldots \times \tilde{m}_{\tilde{q}} \tilde{n}_{\tilde{q}} \times p$.

The adaptive approach allows VDT to be used with images of any dimension, establishing it as a robust tool in tensor-based image processing.

2.2. Our method

Different from previous methods which required an RGB to HSV color space conversion, our **G**obal-**LO**cal-**N**on-local **Retinex** (GLON-Retinex) model directly processes low-light images in the RGB space. The optimization objective function of our variational Retinex-based model for LLIE is as follows:

$$
\underset{\mathcal{L}, \mathcal{R}, X, Y}{\arg\min} \|\mathcal{S} - \mathcal{L} \circ \mathcal{R}\|_F^2 + \lambda_1 \sum_{d=1}^{3mn} mBTV(\mathcal{L})_d
$$
$$
+ \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R})_{[i]} - X_i Y_i \right\|_F^2 + \Phi_L(\mathcal{R}) + \lambda_3 \Phi_N(\mathcal{R}), \tag{1}
$$

where λ_1, λ_2 and λ_3 are penalty parameters, $\mathcal{S}, \mathcal{L}, \mathcal{R} \in \mathbb{R}^{m \times n \times 3}$ signify the low-light image, illuminance, and reflectance, respectively. *mBTV* is the modified Bilateral Total Variation [YGZ*23] used to obtain an illuminance that has clear image boundary while also being sufficiently piecewise smooth. *d* is pixel indices. In our model, the prior for the reflectance $\mathcal{R}$ is thoughtfully formulated by merging a global prior, which is based on tensor train rank, indicated by the third term, with two plug-and-play priors, namely $\Phi_L(\mathcal{R})$ using CBDNet and $\Phi_N(\mathcal{R})$ using CBM3D. $\mathcal{K}()$ denotes the VDT operator and *l* is the order of the high-dimensional tensor augmented by applying VDT to $\mathcal{R}$. $\mathcal{K}(\mathcal{R}) \in \mathbb{R}^{I_1 \times I_2 \times \cdots \times I_l}$ is the tensor augmented by VDT. $\mathcal{K}(\mathcal{R})_{[i]}$ denotes the tensor's canonical matricization [Ose11] operation. α_i are positive weight-

ing parameters that adhere to the constraint $\sum_{i=1}^{l} \alpha_i = 1$. The sets $X = (X_1, X_2, \ldots, X_{l-1})$ and $Y = (Y_1, Y_2, \ldots, Y_{l-1})$ constitute factor matrices. For details on tensor train rank and canonical matricization, see [Ose11].

2.3. Solution for Our Model

Utilizing the half quadratic splitting (HQS) [GR92] technique, we introduce auxiliary variables $\mathcal{P}$ and $\mathcal{P}_1$, thereby reformulating Eq. (1) into the following problem:

$$
\begin{aligned}
\underset{\mathcal{L},\mathcal{R},\mathcal{P},\mathcal{P}_1,X,Y}{\arg\min} \; & \|\mathcal{S} - \mathcal{L} \circ \mathcal{P}_1\|_F^2 + \lambda_1 \sum_{d=1}^{3mn} mBTV(\mathcal{L})_d \\
& + \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R})_{[i]} - X_i Y_i \right\|_F^2 + \Phi_L(\mathcal{R}) + \lambda_3 \Phi_N(\mathcal{P}) \\
& + \frac{\beta}{2} \|\mathcal{P} - \mathcal{R}\|_F^2 + \frac{\beta_1}{2} \|\mathcal{P}_1 - \mathcal{R}\|_F^2,
\end{aligned} \tag{2}
$$

where β and β_1 is a penalty parameter. Figure 2 presents a simplified flowchart of the solution process for our model.

Leveraging the Proximal Alternating Minimization (PAM) algorithm [ABS11], we design an algorithm to decompose problem (2) into a series of more manageable subproblems. Within this alternating minimization framework, we systematically update the variables $\mathcal{L}, \mathcal{R}, X, Y, \mathcal{P}$, and $\mathcal{P}_1$, in an iterative manner as follows:

$$
\begin{cases}
\mathcal{L}^{k+1} = \underset{\mathcal{L}}{\arg\min} \quad Z(\mathcal{L}, X^k, Y^k, \mathcal{P}^k, \mathcal{P}_1^k, \mathcal{R}^k) \\
\qquad\qquad + \frac{\rho}{2} \left\| \mathcal{L} - \mathcal{L}^k \right\|_F^2 \\
X^{k+1} = \underset{X}{\arg\min} \quad Z(\mathcal{L}^{k+1}, X, Y^k, \mathcal{P}^k, \mathcal{P}_1^k, \mathcal{R}^k) \\
\qquad\qquad + \frac{\rho}{2} \left\| X - X^k \right\|_F^2 \\
Y^{k+1} = \underset{Y}{\arg\min} \quad Z(\mathcal{L}^{k+1}, X^{k+1}, Y, \mathcal{P}^k, \mathcal{P}_1^k, \mathcal{R}^k) \\
\qquad\qquad + \frac{\rho}{2} \left\| Y - Y^k \right\|_F^2 \\
\mathcal{P}^{k+1} = \underset{\mathcal{P}}{\arg\min} \quad Z(\mathcal{L}^{k+1}, X^{k+1}, Y^{k+1}, \mathcal{P}, \mathcal{P}_1^k, \mathcal{R}^k) \\
\qquad\qquad + \frac{\rho}{2} \left\| \mathcal{P} - \mathcal{P}^k \right\|_F^2 \\
\mathcal{P}_1^{k+1} = \underset{\mathcal{P}_1}{\arg\min} \quad Z(\mathcal{L}^{k+1}, X^{k+1}, Y^{k+1}, \mathcal{P}^{k+1}, \mathcal{P}_1, \mathcal{R}^k) \\
\qquad\qquad + \frac{\rho}{2} \left\| \mathcal{P}_1 - \mathcal{P}_1^k \right\|_F^2 \\
\mathcal{R}^{k+1} = \underset{\mathcal{R}}{\arg\min} \quad Z(\mathcal{L}^{k+1}, X^{k+1}, Y^{k+1}, \mathcal{P}^{k+1}, \mathcal{P}_1^{k+1}, \mathcal{R}) \\
\qquad\qquad + \frac{\rho}{2} \left\| \mathcal{R} - \mathcal{R}^k \right\|_F^2
\end{cases} \tag{3}
$$

Here, $Z(\mathcal{L}, \mathcal{R}, \mathcal{P}, \mathcal{P}_1, X, Y)$ denotes the objective function in Eq. (2), with ρ as a positive constant and k indicating the iteration count. Within this alternating minimization framework, we systematically update variables $\mathcal{L}, X, Y, \mathcal{P}, \mathcal{P}_1$, and $\mathcal{R}$ in an iterative manner.

(1) Update $\mathcal{L}$: The subproblem for L is as follows:

$$
\begin{aligned}
\mathcal{L}^{k+1} = \underset{\mathcal{L}}{\arg\min} \; & \left\| \mathcal{S} - \mathcal{L} \circ \mathcal{P}_1^k \right\|_F^2 + \lambda_1 \sum_{d=1}^{3mn} mBTV(\mathcal{L})_d \\
& + \frac{\rho}{2} \left\| \mathcal{L} - \mathcal{L}^k \right\|_F^2 \\
= \underset{\mathcal{L}}{\arg\min} & \mathcal{L}^T \left(\mathcal{P}_{1,diag}^k \right)^T \mathcal{P}_{1,diag}^k \mathcal{L} + \rho \mathcal{L}^T \mathcal{L}/2 \\
& + \lambda_1 \mathcal{L}^T \left(\mathcal{C}_h^T \mathcal{W}_h \mathcal{C}_h + \mathcal{C}_v^T \mathcal{W}_v \mathcal{C}_v \right) \mathcal{L} \\
& - 2 \mathcal{S}^T \left(\mathcal{P}_{1,diag}^k \right)^T \mathcal{L} - \rho \left(\mathcal{L}^k \right)^T \mathcal{L}
\end{aligned} \tag{4}
$$

where $\mathcal{C}_h, \mathcal{C}_v$ are Toeplitz matrices from discrete gradient operators with forward difference. The weight $\mathcal{W}_h, \mathcal{W}_v$ are computed according to the methods described in [YGZ*23]. $\mathcal{P}_{1,diag}^k$ is obtained by diagonalizing $\mathcal{P}_1^k$. Subsequently, finding the solution for $\mathcal{L}$ can be reformulated into solving a sequence of linear equation systems:

$$
\mathcal{L}^{k+1} = \mathcal{Q}^{-1} \left(\left(\mathcal{P}_{1,diag}^k \right)^T \mathcal{S} + \rho \mathcal{L}^k / 2 \right) \tag{5}
$$

where $\mathcal{Q} = \left(\mathcal{P}_{1,diag}^k \right)^T \mathcal{P}_{1,diag}^k + \rho \mathcal{I}/2 + \lambda_1 \mathcal{C}_h^T \mathcal{W}_h \mathcal{C}_h + \lambda_1 \mathcal{C}_v^T \mathcal{W}_v \mathcal{C}_v$.

(2) Update X and Y: The minimization processes for X_i and Y_i are decoupled, allowing us to break them down into $l - 1$ independent subproblems. At the $k + 1$ iteration, we update X_i and Y_i by resolving the ensuing optimization problems.

For X_i:

$$
\begin{aligned}
X^{k+1} = \underset{X}{\arg\min} \, & \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R}^k)_{[i]} - X_i Y_i \right\|_F^2 \\
& + \frac{\rho}{2} \left\| X - X^k \right\|_F^2
\end{aligned} \tag{6}
$$

which has the closed-form solution:

$$
\begin{aligned}
X^{k+1} = & \left(\lambda_2 \alpha_i \mathcal{K}(\mathcal{R}^k)_{[i]} \left(Y_i^k \right)^T + \rho X_i^k \right) \\
& \left(\lambda_2 \alpha_i Y_i^k \left(Y_i^k \right)^T + \rho I \right)^{-1}
\end{aligned} \tag{7}
$$

For Y_i:

$$
\begin{aligned}
Y^{k+1} = \underset{Y}{\arg\min} \, & \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R}^k)_{[i]} - X_i^{k+1} Y_i \right\|_F^2 \\
& + \frac{\rho}{2} \left\| Y - Y^k \right\|_F^2
\end{aligned} \tag{8}
$$

which has the closed-form solution:

$$
\begin{aligned}
Y_i^{k+1} = & \left(\lambda_2 \alpha_i \left(X_i^{k+1} \right)^T X_i^{k+1} + \rho I \right)^{-1} \\
& \left(\lambda_2 \alpha_i \left(X_i^{k+1} \right)^T \mathcal{K}(\mathcal{R}_{[i]}^k) + \rho Y_i^k \right)
\end{aligned} \tag{9}
$$

where $I \in \mathbb{R}^{r_i \times r_i}$ is an identity matrix.

(3) Update $\mathcal{P}$: The subproblem for $\mathcal{P}$ is as follows:

$$\mathcal{P}^{k+1} = \arg\min_{\mathcal{P}} \lambda_3 \Phi_N(\mathcal{P}) + \frac{\beta}{2} \left\| \mathcal{P} - \mathcal{R}^k \right\|_F^2 \\ + \frac{\rho}{2} \left\| \mathcal{P} - \mathcal{P}^k \right\|_F^2 \quad (10)$$

Specifically, the computation of P^{k+1} involves:

$$\mathcal{P}^{k+1} = \arg\min_{\mathcal{P}} \lambda_3 \Phi_N(\mathcal{P}) + \frac{\beta+\rho}{2} \left\| \mathcal{P} - \frac{\beta\mathcal{R}^k + \rho\mathcal{P}^k}{\beta+\rho} \right\|_F^2 \quad (11)$$

Incorporating a denoiser with input $\frac{\beta\mathcal{R}^k + \rho\mathcal{P}^k}{\beta+\rho}$, we arrive at the updated $\mathcal{P}$:

$$\mathcal{P}^{k+1} = D_N \left(\left(\beta\mathcal{R}^k + \rho\mathcal{P}^k \right) / (\beta+\rho), \sigma_2 \right) \quad (12)$$

where $\sigma_2 = \sqrt{\lambda_3/(\beta+\rho)}$ and D_N represents the denoiser, capturing the nonlocal prior within the Plug-and-Play ADMM [CWE16] framework. Here, $(\beta\mathcal{R}^k + \rho\mathcal{P}^k)/(\beta+\rho)$ is regarded as the observed tensor, and the output $\mathcal{P}^{k+1}$ constitutes the refined estimate following the denoising process enabled by the applied prior.

(4) Update $\mathcal{P}_1$: The subproblem for $\mathcal{P}_1$ is as follows:

$$\mathcal{P}_1^{k+1} = \arg\min_{\mathcal{P}_1} \left\| \mathcal{S} - \mathcal{L}^{k+1} \circ \mathcal{P}_1 \right\|_F^2 + \frac{\beta_1}{2} \left\| \mathcal{P}_1 - \mathcal{R}^k \right\|_F^2 \\ + \frac{\rho}{2} \left\| \mathcal{P}_1 - \mathcal{P}_1^k \right\|_F^2 \quad (13)$$

This problem is characterized as a least squares problem, and the closed-form solution is expressed as:

$$\mathcal{P}_1^{k+1} = \left(2\left(\mathcal{L}_{diag}^{k+1} \right)^T \mathcal{L}_{diag}^{k+1} + (\beta_1 + \rho)\mathcal{I} \right)^{-1} \\ \left(2\left(\mathcal{L}_{diag}^{k+1} \right)^T \mathcal{S} + \beta_1 \mathcal{R}^k + \rho\mathcal{P}_1^k \right) \quad (14)$$

We applies the preconditioned conjugate gradient (PCG) method [BBC*94] to efficiently solve this linear systems problem.

(5) Update $\mathcal{R}$: The subproblem for $\mathcal{R}$ is as follows:

$$\mathcal{R}_1^{k+1} = \arg\min_{\mathcal{R}} \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R})_{[i]} - X_i^{k+1} Y_i^{k+1} \right\|_F^2 + \Phi_L(\mathcal{R}) \\ + \frac{\beta_1}{2} \left\| \mathcal{P}_1 - \mathcal{R} \right\|_F^2 + \frac{\beta}{2} \left\| \mathcal{P} - \mathcal{R} \right\|_F^2 + \frac{\rho}{2} \left\| \mathcal{R} - \mathcal{R}^k \right\|_F^2 \quad (15)$$

The resolution of Eq.(15) is approached iteratively using the Alternating Direction Method of Multipliers (ADMM) [LLS11]. By introducing the auxiliary variable $\mathcal{E}$, Eq.(15) is reformulated as follows:

$$\mathcal{R}_1^{k+1} = \arg\min_{\mathcal{R}} \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R})_{[i]} - X_i^{k+1} Y_i^{k+1} \right\|_F^2 + \Phi_L(\mathcal{E}) \\ + \frac{\beta_1}{2} \left\| \mathcal{P}_1 - \mathcal{R} \right\|_F^2 + \frac{\beta}{2} \left\| \mathcal{P} - \mathcal{R} \right\|_F^2 + \frac{\rho}{2} \left\| \mathcal{R} - \mathcal{R}^k \right\|_F^2 \\ s.t. \quad \mathcal{E} = \mathcal{R} \quad (16)$$

The augmented Lagrangian function corresponding to Eq. (6) is

expressed as:

$$L(\mathcal{R}, \mathcal{E}, \mathcal{J}) = \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R})_{[i]} - X_i^{k+1} Y_i^{k+1} \right\|_F^2 \\ + \Phi_L(\mathcal{E}) + \frac{\beta_1}{2} \left\| \mathcal{P}_1^{k+1} - \mathcal{R} \right\|_F^2 + \frac{\beta}{2} \left\| \mathcal{P}^{k+1} - \mathcal{R} \right\|_F^2 \\ + \frac{\rho}{2} \left\| \mathcal{R} - \mathcal{R}^k \right\|_F^2 + \frac{\theta}{2} \left\| \mathcal{E} - \mathcal{R} + \frac{\mathcal{J}}{\theta} \right\|_F^2 \quad (17)$$

where $\mathcal{J}$ is Lagrange multiplier, ρ is the penalty parameter. The solution for $\mathcal{R}$ and $\mathcal{E}$ is then obtained by iteratively solving the following two subproblems.

1) *Solving the Subproblem of $\mathcal{E}$*: The subproblem for $\mathcal{E}$ is

$$\mathcal{E}^{t+1} = \arg\min_{\mathcal{E}} \Phi_L(\mathcal{E}) + \frac{\theta}{2} \left\| \mathcal{E} - \mathcal{R}^t + \frac{\mathcal{J}^t}{\theta} \right\|_F^2 \quad (18)$$

where t is the iteration number during the iterative solution of $\mathcal{R}$ and $\mathcal{E}$. Based on the Plug-and-Play ADMM framework, we can view this problem as a denoising step, the solution for $\mathcal{E}$ is

$$\mathcal{E}^{t+1} = D_L \left(\mathcal{R}^t - \mathcal{J}^t/\theta \right) \quad (19)$$

where $\mathcal{R}^t - \mathcal{J}^t/\theta$ is used as the input to the denoiser, denoted by D_L, which effectively captures the nonlocal prior.

2) *Solving the Subproblem of $\mathcal{R}$*: The subproblem for $\mathcal{R}$ is

$$\mathcal{R}^{k+1,t+1} = \arg\min_{\mathcal{R}} \lambda_2 \sum_{i=1}^{l-1} \frac{\alpha_i}{2} \left\| \mathcal{K}(\mathcal{R})_{[i]} - X_i^{k+1} Y_i^{k+1} \right\|_F^2 \\ + \frac{\beta_1}{2} \left\| \mathcal{P}_1^{k+1} - \mathcal{R} \right\|_F^2 + \frac{\beta}{2} \left\| \mathcal{P}^{k+1} - \mathcal{R} \right\|_F^2 \\ + \frac{\rho}{2} \left\| \mathcal{R} - \mathcal{R}^k \right\|_F^2 + \frac{\theta}{2} \left\| \mathcal{E}^{t+1} - \mathcal{R} + \frac{\mathcal{J}}{\theta} \right\|_F^2 \quad (20)$$

This problem is also a least squares problem, and its closed-form solution is

$$\mathcal{R}^{k+1,t+1} = \mathcal{M}/(\lambda_2 + \theta + \beta_1 + \beta + \rho). \quad (21)$$

where $\mathcal{M} = \lambda_2 \sum_{i=1}^{l-1} \alpha_i \mathcal{K}(X_i^{k+1} Y_i^{k+1})_{[i]}^{-1} + \theta\mathcal{E}^{t+1} + \mathcal{J}^t + \beta_1\mathcal{P}_1 + \beta\mathcal{P} + \rho\mathcal{R}^{k,t}$. Operator $\mathcal{K}()_{[i]}^{-1}$ represents the inverse of the VDT and canonical matricization, used to reconstruct the original three-dimensional RGB image.

3) *Update multiplier $\mathcal{J}$*: The Lagrangian multiplier is updated as follows

$$\mathcal{J}^{t+1} = \mathcal{J}^t + \theta \left(\mathcal{E}^{t+1} - \mathcal{R}^{k+1,t+1} \right). \quad (22)$$

It is noteworthy that our method employs a rank-increasing strategy tailored for the Tensor Train (TT) rank of the reflectance component, which is critical for capturing the global structure of the tensor. This adaptive strategy adjusts the TT-rank based on the size of the unfolding matrices, with alterations being triggered upon meeting the specific norm ratio criterion described in [XHYS13].

2.4. Illuminance Adjustment

Upon acquiring estimates for illuminance $\mathcal{L}$ and reflectance $\mathcal{R}$, we apply gamma correction to $\mathcal{L}$, as outlined in [CXG*17, LL22], to adjust the estimated illuminance. The final enhanced image is formulated as

$$\hat{\mathcal{S}} = \left(\frac{\mathcal{L}}{\Omega}\right)^{\frac{1}{\gamma}} \circ \mathcal{R}, \tag{23}$$

where γ is empirically set to 2.2. In gamma correction, Ω is typically set to the maximum possible value of the image, which is 255. However, to enhance the brightness of the results more effectively and prevent over-enhancement, we experientially employed the following strategy for setting the value of Ω:

$$\Omega = \begin{cases} \min(\max(\mathcal{L}), 255)) & \max(\mathcal{L}) \geq 140 \\ \max(\max(L), 100) & else \end{cases} \tag{24}$$

Algorithm 1 outlines the procedure for LLIE with our GLON-Retinex model.

Algorithm 1 Low-Light Image Enhancement Based on the GLON-Retinex Model in Eq. (1)

Input: Observed image $\mathcal{S}$, parameters λ_1, λ_2 and λ_3, stopping criteria δ.

1: Initialize $k = 0$, $\mathcal{L}^0 = \mathcal{S}$, β, β_1, and θ;
2: **while** $\left\|\mathcal{R}^{k+1} - \mathcal{R}^k\right\|_F^2 / \left\|\mathcal{R}^k\right\|_F^2 \geqslant \delta$ and $k \leqslant 10$ **do**
3:     Update $\mathcal{L}^{k+1}$ via Eq. (5) and initialize $\mathcal{R}^0 = \mathcal{S}/\mathcal{L}^1$.
4:     **for** $i = 1$ to $l - 1$ **do**
5:        Update X_i^{k+1} via Eq. (7).
6:        Update Y_i^{k+1} via Eq. (9).
7:     **end for**
8:     Update $\mathcal{P}^{k+1}$ via Eq. (12).
9:     Update $\mathcal{P}_1^{k+1}$ via Eq. (14).
10:     **while** $\left\|\mathcal{R}^{t+1} - \mathcal{R}^t\right\|_F^2 / \left\|\mathcal{R}^t\right\|_F^2 \geqslant \delta$ and $t \leqslant 15$ **do**
11:        Update $\mathcal{R}^{t+1}$ via Eq. (21).
12:        Update $\mathcal{E}^{t+1}$ via Eq. (19).
13:        Update $\mathcal{J}^{t+1}$ via Eq. (22).
14:        Update the parameter $\theta^{t+1} = \min\left(1.2\theta^t, 10^6\right)$.
15:     **end while**
16: **end while**

Output: The enhanced image $\hat{\mathcal{S}}$ via Eq. (23).

---

3. Experiments

In this section, we evaluate the performance of our model against state-of-the-art LLIE methods through both subjective and objective comparisons. Experiments were conducted using MATLAB 2022b on a PC equipped with AMD Ryzen 5 3600 6-Core Processor and NVIDIA RTX 2070 GPU. For additional visual results, please refer to the supplementary materials provided with our paper.

Datasets and Evaluation Metrics: We utilized a dataset named PG660 [YGZ*23], comprising 660 unpaired images collected from datasets MF [FZH*16a], LIME [GLL16], SIRE [WBSS04], DICM [LLK13], NPE [WZHL13], MEF [MZW15], and ExDark [LC19] to test the generalization performance of various methods on real-world low-light images. This dataset includes low-light images under a range of lighting conditions and noise levels that one might encounter in real-world scenarios. In PG660 dataset, we employed Natural Image Quality Evaluator (NIQE) [MSB12], No-reference Image Quality Metric for Contrast distortion (NIQMC) [GLZ*16], and Visual Information Fidelity (VIF) [SB06] metrics to assess the proximity between the enhanced results and the characteristics of natural images, contrast, and visual information fidelity, respectively. A lower value for the NIQE indicates higher perceived image quality, whereas the opposite is true for the latter two. Additionally, for a more precise characterization of natural image features, we utilized 424 ground truth nighttime images from the SID [CCXK18] dataset and 2560 high-resolution images from the Flickr2K [TAVG*17] dataset to compute the distribution of natural images, which was then used in the calculation of the NIQE metric.

Furthermore, we evaluated noise suppression and detail retention capabilities of various methods using paired datasets LOL-v1 [WWYL18], LOL-v2 [YWH*21], and LSRW [HXY*23], which have 15, 100, and 50 paired images, respectively, employing PSNR, SSIM [WBSS04], and LPIPS [ZIE*18] for assessment. PSNR (Peak Signal-to-Noise Ratio) measures the ratio between the maximum possible power of a signal and the power of corrupting noise, where a higher value indicates better image quality. SSIM (Structural Similarity Index Measure) assesses visual impact of three characteristics of an image: luminance, contrast, and structure, with higher values indicating more similarity to the original image. LPIPS (Learned Perceptual Image Patch Similarity) quantifies perceptual differences between images using deep learning features, where lower values indicate lesser perceptual differences.

Parameters Setting: In the experiments conducted on the non-paired dataset PG660, the parameters of our model were set as $\lambda_1 = 0.05, \lambda_2 = 0.01, \lambda_3 = 0.0015, \beta = 0.5, \beta_1 = 6, \theta = 0.13, \rho = 0.07, K = 5, \delta = 2e - 6$. Due to the extremely low brightness of most low-light images in the paired datasets LOL-v1, LOL-v2, and LSRW, these images exhibit significant noise. Consequently, we adjusted the parameters $\lambda_3 = 0.1, \beta = 15, \theta = 0.05$, to enhance noise suppression more effectively.

3.1. Experimental Results on Unpaired Dataset PG660

The comparative methods are shown in Table 1, where methods URetinex, GSAD, and DIFFLL are supervised and also included to assess their generalization on unpaired datasets. Table 1, along with Figure 3 and 4, display quantitative and visual results of these comparisons. From Table 1, it's clear that our method significantly outperforms others in the NIQE, indicating a closer statistical resemblance to high-quality images in the SID and Flickr2K datasets. Besides LR3M and GSAD, other methods demonstrated lower NIQE scores than the original low-light images, indicating that while brightness was enhanced, overall image quality deteriorated. This deterioration is characterized by various issues such as overexposure in SCI and URetinex; color distortion in LLFlow and ZeroDCE; noise amplification in WVM, JieP, STAR, PnPRetinex, BTRetinex, and ZeroDCE; loss of detail in SRRR; and the intro-
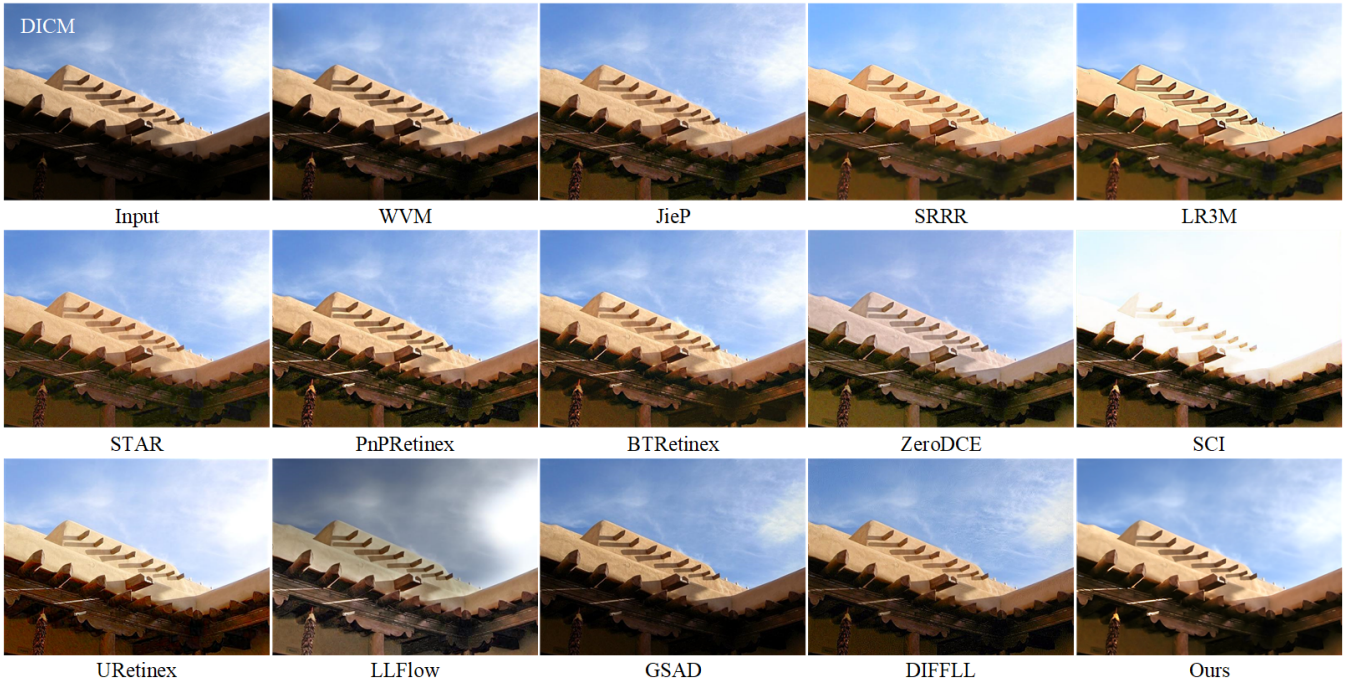
Figure 3: *Global Comparative analysis of enhanced results on the DICM [LLK13] dataset: The visual quality of the results from our proposed method is notably impressive.* **Zoom in for a better view**.

Table 1: *Comparison of NIQE, NIQMC, and VIF metrics for enhanced results by various methods on unpaired dataset PG660. Bold numbers indicate the best results, while underlined numbers represent the second-best results.*

Metric	NIQE↓	NIQMC↑	VIF↑
Input	3.46	4.14	1.00
WVM [FZH*16b]	3.74	4.68	1.82
JieP [CXG*17]	3.84	4.73	1.82
SRRR [LLY*18]	4.04	4.92	1.85
LR3M [RYCL20]	3.45	4.84	1.85
STAR [XHR*20]	3.79	4.81	1.75
PnPRetinex [LL22]	4.14	5.00	2.94
BTRetinex [YGZ*23]	4.03	4.94	2.34
ZeroDCE [GLG*20]	4.17	4.93	2.06
SCI [MML*22]	4.41	**5.28**	**3.06**
URetinex [WWZ*22]	3.94	<u>5.17</u>	2.34
LLFlow [WWY*22]	3.52	5.30	2.75
GSAD [HZH*23]	<u>3.26</u>	4.75	1.63
DIFFLL [JLF*23]	3.87	5.06	1.87
Ours	**3.072**	5.14	<u>2.99</u>

duction of artifacts in URetinex, LLFlow, and DIFFLL, as illustrated in Figure 3 and 4.

On the contrast metric NIQMC and the visual fidelity metric VIF, our results achieved third and second places, respectively. Although

the SCI and URetinex methods may attain superior outcomes on these metrics, it is noteworthy that excessively high NIQMC and VIF values are often caused by over-enhancement, as seen in the SCI and URetinex methods (shown in Figure 3).

Overall, our method yields competitive contrast and visual fidelity on non-paired datasets while maintaining optimal perceptual quality, as indicated by the best NIQE scores achieved.

3.2. Experimental Results on Paired Datasets

To comprehensively evaluate the capability of various unsupervised methods in enhancing image brightness while suppressing noise and preserving details, we conducted experiments on the paired datasets LOL-v1, LOL-v2, and LSRW. The methods involved in the comparison are listed in Table 2.

Table 2, Figure 1, and Figure 5 collectively present detailed quantitative comparisons and visual results. According to the data in Table 2, our method demonstrates the best performance in both PSNR and SSIM metrics on LOL-v1 and LSRW datasets, with PSNR values exceeding the second-best method by 3.38dB and 1.19dB, respectively. Additionally, on the LOL-v2 dataset, our method also leads in the PSNR metric, outperforming the second-best method by 1.21dB, and is closely competitive in the SSIM metric with the top method SRRR. Although the SRRR method achieves the highest SSIM values at this time, it typically smooths out a significant amount of image details, as illustrated in Figure 1 and 5.
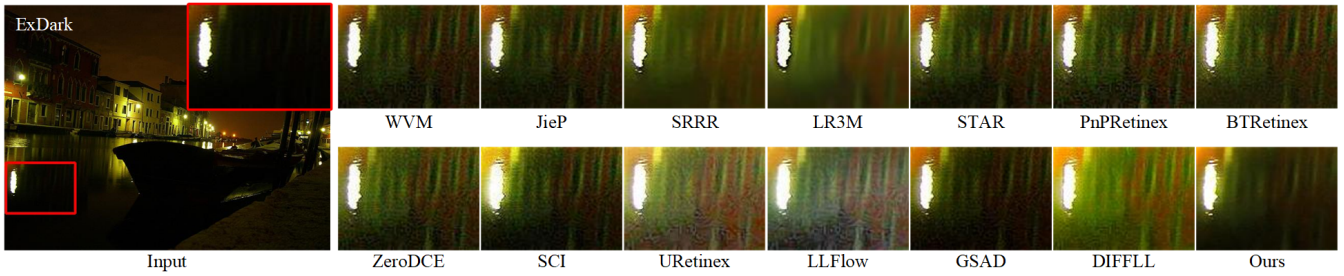
Figure 4: *Close-Up comparative analysis of enhanced results on the ExDark [LC19] dataset.* **Zoom in for a better view**.
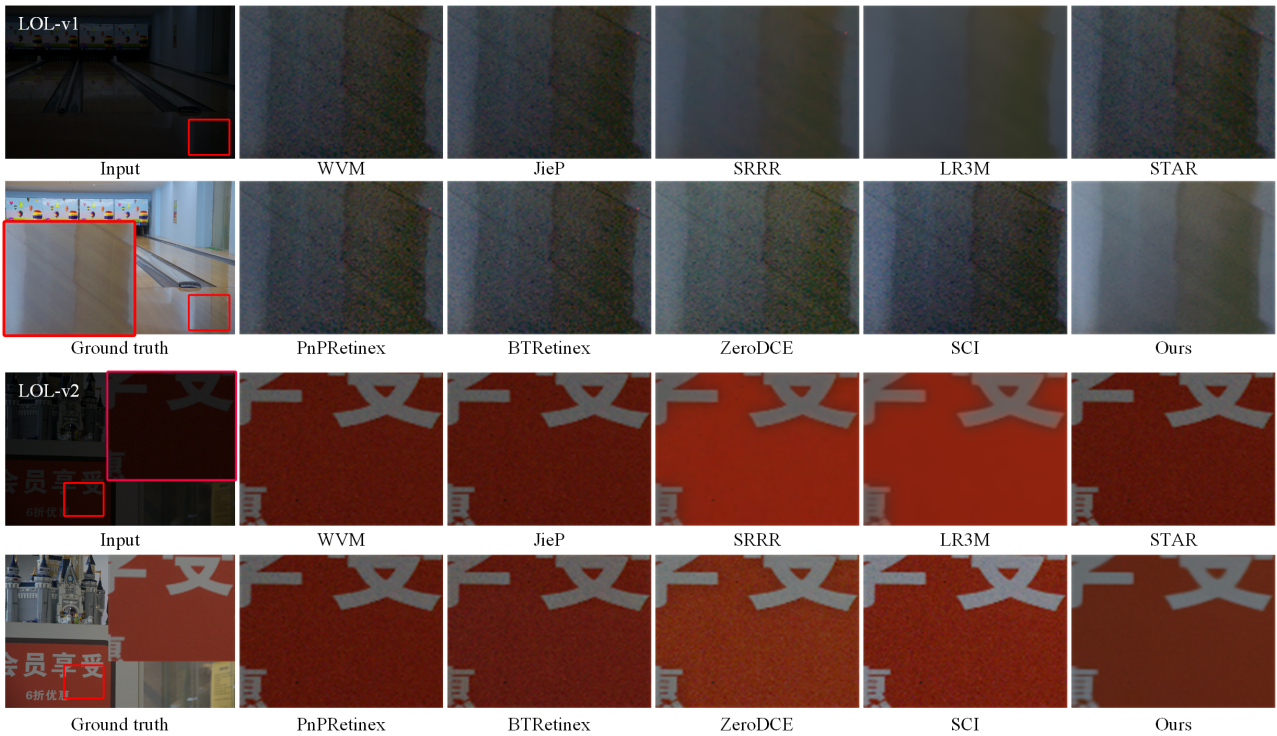
Figure 5: *Close-Up comparative analysis of enhanced results on the LOL-v1 [WWYL18] and LOL-v2 [YWH*21] datasets.* **Zoom in for a better view**.

Regarding the LPIPS metric, our method ranks as the best on both the LOL-v1 and LOL-v2 datasets, with a particularly significant lead on the LOL-v1 dataset. Although not the top performer on the LSRW dataset, our score of 0.342 is very close to the best score of 0.318 by ZeroDCE. Overall, these outcomes confirm that our method is highly competitive in retaining perceptual quality, as corroborated by the visualizations presented in our paper.

In Table 3, we present the performance metrics of two supervised methods, DIFFLL and GSAD, on paired datasets. Since these methods were trained on the LOL-v1 and LOL-v2 datasets, we can naturally expect them to achieve excellent results on the PSNR, SSIM, and LPIPS metrics for these datasets, significantly outperforming the unsupervised methods mentioned earlier. However,

their lack of training on the LSRW dataset has resulted in a drastic decline in their performance metrics on this dataset, potentially even falling below the levels achieved by the unsupervised methods we proposed.

In Figure 6, we present a visual comparison of the supervised enhancement methods DIFFLL and GSAD applied to paired datasets. As depicted in Figure 6, both methods generally achieve good results on the LOL-v1 and LOL-v2 datasets. However, in the LSRW dataset, the DIFFLL method not only produces significant artifacts but also suffers from severe color distortion. Additionally, the GSAD method results in comparatively darker images with noticeable noise.

These results highlight that our method not only excels at

Table 2: *Comparison of PSNR, SSIM, and LPIPS metrics for **Unsupervised** Methods on Paired Datasets.*

Dataset	LOL-v1			LOL-v2			LSRW		
Metric	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
WVM [FZH*16b]	11.855	0.498	0.517	14.451	0.542	0.312	13.349	0.467	0.342
JieP [CXG*17]	12.047	0.512	0.495	14.719	0.558	0.280	13.782	0.485	0.327
SRRR [LLY*18]	13.877	0.658	0.446	17.345	0.686	0.254	15.023	0.556	0.445
LR3M [RYCL20]	13.900	0.646	0.487	17.423	0.658	0.315	15.079	0.557	0.472
STAR [XHR*20]	12.637	0.538	0.490	15.576	0.572	0.276	14.336	0.497	0.341
PnPRetinex [LL22]	13.072	0.557	0.347	16.148	0.583	0.317	14.352	0.481	0.340
BTRetinex [YGZ*23]	13.455	0.571	0.524	16.642	0.6200	0.300	14.527	0.504	0.341
ZeroDCE [GLG*20]	14.861	0.585	0.335	18.059	0.603	0.313	15.876	0.497	**0.318**
SCI [MML*22]	14.784	0.523	0.522	17.304	0.555	0.308	15.163	0.443	0.336
Ours	**18.244**	**0.717**	**0.223**	**19.274**	0.665	**0.248**	**17.070**	**0.578**	0.342

Table 3: *Results of PSNR, SSIM, and LPIPS metrics for **Supervised** Methods on Paired Datasets.*

Dataset	LOL-v1			LOL-v2			LSRW		
Metric	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
DIFFLL [JLF*23]	26.318	0.849	0.118	28.881	0.895	0.010	17.206	0.451	0.450
GSAD [HZH*23]	27.569	0.872	0.092	28.789	0.883	0.095	14.810	0.486	0.334

Input   DIFFLL   GSAD   Ground truth
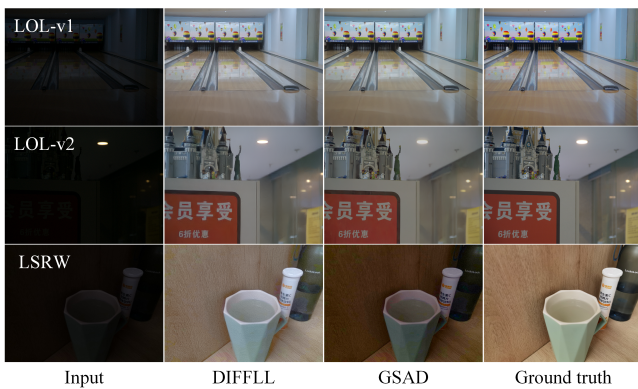
Figure 6: *Visual Comparison of Supervised Methods DIFFLL [JLF*23] and GSAD [HZH*23] on Paired Datasets. **Zoom in for a better view**.*

Table 4: *Results of average PSNR, SSIM, and LPIPS metrics from ablation experiments on paired dataset LOL-v1.*

Global prior: TT-rank	✗	✔	✗	✗	✔
Local prior: CBDNet	✗	✗	✔	✗	✔
Non-local prior: CBM3D	✗	✗	✗	✔	✔
Avg PSNR↑	17.599	17.803	17.794	18.178	**18.267**
Avg SSIM↑	0.6202	0.6423	0.6416	0.7133	**0.7203**
Avg LPIPS↓	0.3444	0.3260	0.3041	0.2067	**0.2017**

Specifically, the synergistic integration of global, local, and non-local constraints delivers the most substantial improvements.

3.4. Limitation

One challenge with our optimization-based iterative algorithm is its lack of computational efficiency. In particular, the *mBTV* and CBM3D modules require several seconds per iteration for each image, which substantially extends the processing time. For instance, processing an image of $600 \times 400 \times 3$ resolution from the LOL-v1 dataset takes about 55 seconds with our method. While this is significantly faster than LR3M's 108 seconds, it still lags behind PnPRetinex's swift 4-second processing time.

Additionally, our model faces the issue of parameter tuning, which demands adjustments based on both the physical interpretation of the model and empirical observations—a notably time-intensive task. We have provided comprehensive guidelines for parameter selection on our code page. Looking ahead, we aim to enhance the model's adaptability by integrating advanced parameter adjustment strategies, thus refining its overall utility.

4. CONCLUSION

Our method, rooted in Retinex theory and incorporating global, local, and non-local priors, demonstrates outstanding performance in

noise suppression, consistently achieving the highest or competitive PSNR values, but also effectively preserves details. This dual capability demonstrates our method's strong competitiveness in both PSNR and SSIM metrics.

3.3. Ablation Study

To assess the effectiveness of each component in our model, we conducted ablation studies using the paired LOL-v1 dataset with specific parameter settings ($\lambda_1 = 0.05, \lambda_2 = 2, \lambda_3 = 0.01, \beta = 15, \beta_1 = 6, \theta = 0.05$). The findings are detailed in Table 4.

The results from these studies clearly demonstrate that each module significantly enhances the performance of our model.

Low-Light Image Enhancement (LLIE). By combining tensor train rank with two advanced plug-and-play denoisers, our approach not only effectively suppresses noise but also preserves critical image details. Although there is a compromise in processing speed, the method's capability to markedly improve visual quality and retain essential details establishes it as a highly competitive and effective solution in the field of LLIE.

Acknowledgments

References

[ABS11] ATTOUCH H., BOLTE J., SVAITER B.: Convergence of descent methods for semi-algebraic and tame problems: Proximal algorithms, forward-backward splitting, and regularized gauss-seidel methods. Mathematical Programming 137 (01 2011). doi:10.1007/s10107-011-0484-9. 2, 4

[BBC*94] BARRETT R., BERRY M., CHAN T. F., DEMMEL J., DONATO J., DONGARRA J., EIJKHOUT V., POZO R., ROMINE C., VAN DER VORST H.: Templates for the solution of linear systems: building blocks for iterative methods. SIAM, 1994. 5

[CCXK18] CHEN C., CHEN Q., XU J., KOLTUN V.: Learning to see in the dark. In Proceedings of the IEEE conference on computer vision and pattern recognition (2018), pp. 3291–3300. 6

[CWE16] CHAN S. H., WANG X., ELGENDY O. A.: Plug-and-play admm for image restoration: Fixed-point convergence and applications. IEEE Transactions on Computational Imaging 3, 1 (2016), 84–98. 2, 5

[CXG*17] CAI B., XU X., GUO K., JIA K., HU B., TAO D.: A joint intrinsic-extrinsic prior model for retinex. In Proceedings of the IEEE international conference on computer vision (2017), pp. 4000–4009. 1, 6, 7, 9

[CYZ17] CHANG Y., YAN L., ZHONG S.: Hyper-laplacian regularized unidirectional low-rank tensor recovery for multispectral image denoising. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2017), pp. 4260–4268. 2

[DFKE07] DABOV K., FOI A., KATKOVNIK V., EGIAZARIAN K.: Color image denoising via sparse 3d collaborative filtering with grouping constraint in luminance-chrominance space. In 2007 IEEE international conference on image processing (2007), vol. 1, IEEE, pp. I–313. 2

[FZH*16a] FU X., ZENG D., HUANG Y., LIAO Y., DING X., PAISLEY J.: A fusion-based enhancing method for weakly illuminated images. Signal Processing 129 (2016), 82–96. 6

[FZH*16b] FU X., ZENG D., HUANG Y., ZHANG X.-P., DING X.: A weighted variational model for simultaneous reflectance and illumination estimation. In Proceedings of the IEEE conference on computer vision and pattern recognition (2016), pp. 2782–2790. 1, 7, 9

[GLG*20] GUO C., LI C., GUO J., LOY C. C., HOU J., KWONG S., CONG R.: 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. 2, 7, 9

[GLL16] GUO X., LI Y., LING H.: Lime: Low-light image enhancement via illumination map estimation. IEEE Transactions on image processing 26, 2 (2016), 982–993. 1, 6

[GLZ*16] GU K., LIN W., ZHAI G., YANG X., ZHANG W., CHEN C. W.: No-reference quality metric of contrast-distorted images based on information maximization. IEEE transactions on cybernetics 47, 12 (2016), 4559–4565. 6

[GR92] GEMAN D., REYNOLDS G.: Constrained restoration and the recovery of discontinuities. IEEE Transactions on Pattern Analysis and Machine Intelligence 14, 3 (1992), 367–383. doi:10.1109/34.120331. 4

[GYZ*19] GUO S., YAN Z., ZHANG K., ZUO W., ZHANG L.: Toward convolutional blind denoising of real photographs. In Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (2019), pp. 1712–1722. 2

[HXY*23] HAI J., XUAN Z., YANG R., HAO Y., ZOU F., LIN F., HAN S.: R2rnet: Low-light image enhancement via real-low to real-normal network. Journal of Visual Communication and Image Representation 90 (2023), 103712. 2, 6

[HZH*23] HOU J., ZHU Z., HOU J., LIU H., ZENG H., YUAN H.: Global structure-aware diffusion process for low-light image enhancement. arXiv preprint arXiv:2310.17577 (2023). 2, 7, 9

[JLF*23] JIANG H., LUO A., FAN H., HAN S., LIU S.: Low-light image enhancement with wavelet-based diffusion models. ACM Transactions on Graphics (TOG) 42, 6 (2023), 1–14. 2, 7, 9

[Lan77] LAND E. H.: The retinex theory of color vision. Scientific american 237, 6 (1977), 108–129. 1

[LC19] LOH Y. P., CHAN C. S.: Getting to know low-light images with the exclusively dark dataset. Computer Vision and Image Understanding 178 (2019), 30–42. 6, 8

[LL22] LIN Y.-H., LU Y.-C.: Low-light enhancement using a plug-and-play retinex model with shrinkage mapping for illumination estimation. IEEE Transactions on Image Processing 31 (2022), 4897–4908. 1, 6, 7, 9

[LLK13] LEE C., LEE C., KIM C.-S.: Contrast enhancement based on layered difference representation of 2d histograms. IEEE transactions on image processing 22, 12 (2013), 5372–5384. 6, 7

[LLS11] LIN Z., LIU R., SU Z.: Linearized alternating direction method with adaptive penalty for low-rank representation. Advances in neural information processing systems 24 (2011). 2, 5

[LLY*18] LI M., LIU J., YANG W., SUN X., GUO Z.: Structure-revealing low-light image enhancement via robust retinex model. IEEE Transactions on Image Processing 27, 6 (2018), 2828–2841. 1, 7, 9

[MML*22] MA L., MA T., LIU R., FAN X., LUO Z.: 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. 2, 7, 9

[MSB12] MITTAL A., SOUNDARARAJAN R., BOVIK A. C.: Making a "completely blind" image quality analyzer. IEEE Signal processing letters 20, 3 (2012), 209–212. 6

[MZW15] MA K., ZENG K., WANG Z.: Perceptual quality assessment for multi-exposure image fusion. IEEE Transactions on Image Processing 24, 11 (2015), 3345–3356. 6

[Ose11] OSELEDETS I. V.: Tensor-train decomposition. SIAM Journal on Scientific Computing 33, 5 (2011), 2295–2317. 3, 4

[RYCL20] REN X., YANG W., CHENG W.-H., LIU J.: Lr3m: Robust low-light enhancement via low-rank regularized retinex model. IEEE Transactions on Image Processing 29 (2020), 5862–5876. 1, 7, 9

[SB06] SHEIKH H. R., BOVIK A. C.: Image information and visual quality. IEEE Transactions on image processing 15, 2 (2006), 430–444. 6

[TAVG*17] TIMOFTE R., AGUSTSSON E., VAN GOOL L., YANG M.-H., ZHANG L.: Ntire 2017 challenge on single image super-resolution: Methods and results. In Proceedings of the IEEE conference on computer vision and pattern recognition workshops (2017), pp. 114–125. 6

[WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: from error visibility to structural similarity. IEEE transactions on image processing 13, 4 (2004), 600–612. 6

[WWY*22] WANG Y., WAN R., YANG W., LI H., CHAU L.-P., KOT A.: Low-light image enhancement with normalizing flow. In Proceedings of the AAAI conference on artificial intelligence (2022), pp. 2604–2612. 2, 7

[WWYL18] WEI C., WANG W., YANG W., LIU J.: Deep retinex decomposition for low-light enhancement. arXiv preprint arXiv:1808.04560 (2018). 2, 6, 8

[WWZ*22] WU W., WENG J., ZHANG P., WANG X., YANG W., JIANG J.: 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. 2, 7

[WZHL13] WANG S., ZHENG J., HU H.-M., LI B.: Naturalness preserved enhancement algorithm for non-uniform illumination images. IEEE transactions on image processing 22, 9 (2013), 3538–3548. 6

[XHR*20] XU J., HOU Y., REN D., LIU L., ZHU F., YU M., WANG H., SHAO L.: Star: A structure and texture aware retinex model. IEEE Transactions on Image Processing 29 (2020), 5022–5037. 1, 7, 9

[XHYS13] XU Y., HAO R., YIN W., SU Z.: Parallel matrix factorization for low-rank tensor completion. Inverse Problems and Imaging 9 (12 2013). doi:10.3934/ipi.2015.9.601. 5

[YGZ*23] YANG W., GAO H., ZOU W., HUANG S., CHEN H., MA J.: Enhancing low-light images: A variation-based retinex with modified bilateral total variation and tensor sparse coding. In Computer Graphics Forum (2023), Wiley Online Library, p. e14960. 1, 2, 3, 4, 6, 7, 9

[YWH*21] YANG W., WANG W., HUANG H., WANG S., LIU J.: Sparse gradient regularized deep retinex network for robust low-light image enhancement. IEEE Transactions on Image Processing 30 (2021), 2072–2086. 2, 6, 8

[YZGC19] YUAN L., ZHAO Q., GUI L., CAO J.: High-order tensor completion via gradient-based optimization under tensor train format. Signal Processing: Image Communication 73 (2019), 53–61. 2

[ZHZ*19] ZHENG Y.-B., HUANG T.-Z., ZHAO X.-L., JIANG T.-X., MA T.-H., JI T.-Y.: Mixed noise removal in hyperspectral image via low-fibered-rank regularization. IEEE Transactions on Geoscience and Remote Sensing 58, 1 (2019), 734–749. 2

[ZIE*18] ZHANG R., ISOLA P., EFROS A. A., SHECHTMAN E., WANG O.: 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. 6

[ZYM*21] ZHAO X.-L., YANG J.-H., MA T.-H., JIANG T.-X., NG M. K., HUANG T.-Z.: Tensor completion via complementary global, local, and nonlocal priors. IEEE Transactions on Image Processing 31 (2021), 984–999. 2