

# AGLLDiff: Guiding Diffusion Models Towards Unsupervised Training-free Real-world Low-light Image Enhancement

Yunlong Lin<sup>1,2\*</sup>, Tian Ye<sup>3\*</sup>, Sixiang Chen<sup>3\*</sup>, Zhenqi Fu<sup>4</sup>, Yingying Wang<sup>1,2</sup>, Wenhao Chai<sup>5</sup>,  
Zhaohu Xing<sup>3</sup>, Wenxue Li<sup>3</sup>, Lei Zhu<sup>3,6</sup>, and Xinghao Ding<sup>1,2†</sup>

<sup>1</sup>Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, Xiamen University, China

<sup>2</sup>School of Informatics, Xiamen University, China

<sup>3</sup>The Hong Kong University of Science and Technology (Guangzhou), China

<sup>4</sup>Tsinghua University, China

<sup>5</sup>University of Washington, USA

<sup>6</sup>The Hong Kong University of Science and Technology, Hong Kong SAR, China  
dxh@xmu.edu.cn

## Abstract

Existing low-light image enhancement (LIE) methods have achieved noteworthy success in solving synthetic distortions, yet they often fall short in practical applications. The limitations arise from two inherent challenges in real-world LIE: 1) the collection of distorted/clean image pairs is often impractical and sometimes even unavailable, and 2) accurately modeling complex degradations presents a non-trivial problem. To overcome them, we propose the Attribute Guidance Diffusion framework (AGLLDiff), a training-free method for effective real-world LIE. Instead of specifically defining the degradation process, AGLLDiff shifts the paradigm and models the desired attributes, such as image exposure, structure and color of normal-light images. These attributes are readily available and impose no assumptions about the degradation process, which guides the diffusion sampling process to a reliable high-quality solution space. Extensive experiments demonstrate that our approach outperforms the current leading unsupervised LIE methods across benchmarks in terms of distortion-based and perceptual-based metrics, and it performs well even in sophisticated wild degradation.

## Introduction

Real-world low-light image enhancement (LIE) aims to ameliorate the quality and brightness of an image suffering from unknown degradation, such as low contrast, multiple artifacts, poor visibility, sensor noise, etc. Great improvement in enhancement quality has been witnessed over the past few years with the exploitation of generative priors (Fei et al. 2023; Jiang et al. 2021). For instance, Generative Adversarial Networks (GANs) (Jiang et al. 2021; Wang and Zhao 2024a) that are trained on extensive datasets of clean images and learn rich knowledge of real-world scenes have succeeded in LIE through GAN inversion. Compared to GANs, Denoising Diffusion Probabilistic Models (DDPMs) (Ho, Jain, and Abbeel 2020; Guo et al. 2024) yield

more high-fidelity and realistic details, thereby fostering a surge of interest in adapting diffusion models to LIE (Yin et al. 2023; Jiang et al. 2023).

Recent diffusion-based LIE methods can be roughly classified into two categories (He et al. 2024): **1) The common approaches** (Yin et al. 2023; Cai et al. 2023) are dedicated to accurately modeling degradation process via supervised learning, which show proficiency in synthetic degradation scenes but lack robustness to handle challenging unseen degradations. This inadaptability primarily stems from the inconsistency between the synthetic degradation of training data and the actual degradation in the real world. Enriching the synthetic data for model training would improve the models' generalizability, but it is obviously impractical to simulate every possible degradation in the real world. **2) The second ones** (Fei et al. 2023; Lin et al. 2024) strive to exploit the diffusion priors in the pre-trained diffusion models, which are effective in adapting to multiple degradations. Yet, despite their versatility, these methods are inevitably constrained in terms of generalizability, as they require prior knowledge of the specific degradation process in advance. In practice, degradations in the wild often include a mixture of multiple types, posing a challenge to accurately model them. In summary, two primary challenges are commonly encountered in real-world LIE: **i)** the collection of distorted/clean image pairs is often impractical and sometimes even unavailable, and **ii)** accurately modeling complex degradations presents a non-trivial problem.

To address the aforementioned challenges, we introduce **a novel training-free and unsupervised framework, named AGLLDiff**, for real-world low-light enhancement. In contrast to prior works that predefine the degradation process, **our approach models the desired attributes and incorporates this guidance within the diffusion generative process**. Concretely, we leverage a well-performing diffusion model (DM) (Dhariwal and Nichol 2021), which generates images through a stochastic iterative sampling process, and the attributes act as classifiers to constrain the generative process to a reliable high-quality (HQ) solution space. As shown

\*Equal contribution

†Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: **Motivation of our AGLLDiff.** (a) represents the data distribution of normal-light samples and degraded samples. It is evident that degradation significantly deviates from normal-light samples. (b) conceptually illustrate the geometries of the proposed attribute guidance sampling algorithm. It shows that, given the initial latent, which lies in the low-probability region, attribute guidance guides the latent to move towards its vicinal high-probability region. (c) presents that imposing gaussian noise on the degraded sample and its corresponding reference sample makes the distributions between them less distinguishable.

in Fig. 1, noisy images are degradation-irrelevant conditions for the DM generative process. Adding extra gaussian noise makes the degradation less distinguishable compared with its corresponding reference distribution. Since diffusion prior can serve as a natural image regularization, one could simply guide the sampling process with easily accessible attributes such as image exposure, structure and color of normal-light images. ***By constraining a reliable HQ solution space, the core of our philosophy is to bypass the difficulty of discerning the prior relationship between low-light and normal-light images, thus improving generalizability.*** DiffLLE (Yang et al. 2024b) shares a similar motivation by incorporating the diffusion model’s denoising prior into a low-light model through an unsupervised distillation training framework. In contrast to DiffLLE, our framework is fully training-free, achieving low-light enhancement through a dynamic guidance scheme and composite multi-attribute guidance, without relying on a specific physical degradation model. Our contributions can be summarized as follows:

- We introduce a novel paradigm, AGLLDiff, a training-free and unsupervised method that requires no degradation of prior knowledge but yields high fidelity and generality towards real-world low-light image enhancement.
- We demonstrate that AGLLDiff suffices to guide the pre-trained diffusion models to a reliable high-quality solution space through easily accessible attributes in the HQ image space.
- Comprehensive experiments reveal that our framework achieves both robustness and high quality on heavily degraded synthetic and real-world datasets.

## Related Work

**Low-light Image Enhancement.** To transform low-light images into visually satisfactory ones, numerous efforts have been made over the decades. The conventional approaches are first widely adopted (Ren et al. 2018; Hao et al.

2020; Huang, Cheng, and Chiu 2012). For example, Guo et al. (Li et al. 2018) suggested refining the initially estimated illumination map by incorporating a structural prior. Lee et al. (Lee, Lee, and Kim 2012) applied the layered difference representation of 2D histograms to amplify the gray-level differences between adjacent pixels. The enhancement performance of current conventional methods relies on tedious hand-crafted priors and is only applicable to specific scenarios. Recently, the paradigm of low-light image enhancement has gradually shifted to data-driven approaches based on deep learning (He et al. 2023a; Jin, Yang, and Tan 2022). For instance, Ma et al. (Ma et al. 2022) established a cascaded illumination estimation process to achieve fast and robust LIE in complex scenarios. Cai et al. (Cai et al. 2023) designed a novel one-stage Retinex-based framework for LIE. Additionally, in contexts where training images are scarce, the utility of unsupervised (Fu et al. 2023; Yang et al. 2023) and zero-shot learning approaches (Guo et al. 2020; Liang et al. 2023; Fei et al. 2023) becomes increasingly pronounced.

**Diffusion-Based Image Restoration and Low-Light Image Enhancement.** Diffusion Models has become increasingly influential in various field (Fei et al. 2023; Ho, Jain, and Abbeel 2020; Wang and Zhao 2024b; Wang, Xu, and Zhao 2024; Chen et al. 2025; Fan et al. 2024; Li et al. 2024; Fan et al. 2023), such as super resolution (Zhu et al. 2023; Xia et al. 2024), multi-modal fusion (Lin et al. 2023; Wang et al. 2023, 2024; Wang and Zhao 2023), de-hazing (Ye et al. 2022; Jin et al. 2023), image enhancement (Lin et al. 2024; He et al. 2023b). These methods could be broadly categorized into supervised and unsupervised paradigms. Supervised-based IR solutions usually rely on large-scale, pre-collected paired datasets to train their models with great success. Jiang et al. (Jiang et al. 2023) employed wavelet transformation to decrease the input size and a high-frequency restoration module to maintain the details. Yin et al. (Yin et al. 2023) directly utilized the color map as an extra conditional control to preserve the color information. A major challenge is that they implicitly assume training and testing data should be identically distributed. As a result, these methods often deteriorate seriously in performance when testing cases deviate from the pre-assumed distribution. Another prevailing research line is unsupervised-based IR approaches. They adopt a zero-shot approach to leverage a pre-trained diffusion model for restoration without the need for task-specific training. As an early attempt, Kawar et al. (Kawar et al. 2022) hypothesized the linear degradation model and relied on the desirable property of linear formulation to sample from posterior distribution. Fei et al. (Fei et al. 2023) applied a simultaneous estimation of degradation model to address blind degradation. Yang et al. (Yang et al. 2024a) introduced a partial guidance mechanism for blind face restoration, wherein intermediate outputs of the diffusion model are constrained by a classifier to perform photo restoration. Previous diffusion-based image restoration methods explicitly leveraged a degradation model by solving a maximum posterior problem or a posterior sampling problem to generate solutions. However, for many practical image enhancement problems, the underlying degradation model may not be available. ***In this work,***

we propose to model the desired attributes of normal-light images. Such a strategy is independent of the degradation process, circumventing the difficulty of modeling the degradation process.

## Methodology

### Preliminaries of Diffusion Models

The diffusion models (Ho, Jain, and Abbeel 2020; Guo et al. 2024; Chai et al. 2023) belong to a category of generative models that operate by incrementally incorporating Gaussian noise into training data and subsequently acquiring the denoiser to restore the data distribution  $p(\mathbf{x})$  by reversing the process of noise injection.

The forward process  $q(\mathbf{x}_t | \mathbf{x}_{t-1})$  transforms an initial image  $\mathbf{x}_0$  into Gaussian noise  $\mathbf{x}_T \sim \mathcal{N}(0, 1)$  over  $T$  iterations. The following equation can express the process of each iteration in the diffusion:

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1 - \beta_t} \mathbf{x}_{t-1}, \beta_t \mathbf{I}), \quad (1)$$

where  $\mathbf{x}_t$  denotes the noisy image at time-step  $t$ ,  $\beta_t$  is the pre-determined scaling factor, and  $\mathcal{N}$  represents the Gaussian distribution. Under the reparameterization trick,  $\mathbf{x}_t$  can be written as:

$$\mathbf{x}_t = \sqrt{\bar{\alpha}_t} \mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad (2)$$

where  $\bar{\alpha}_t := \prod_{i=1}^t (1 - \beta_i)$  and  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ . Then  $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$  if  $T$  is big enough, usually  $T = 1000$ .

The reverse generative process of the inference stage, starting from a Gaussian random noise map  $\mathbf{x}_T \sim \mathcal{N}(0, \mathbf{I})$  and iteratively performing the denoising step until it attains a high-quality output  $\mathbf{x}_0$ :

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \Sigma_\theta \mathbf{I}), \quad (3)$$

where variance  $\Sigma_\theta \mathbf{I}$  can be either time-dependent constants or learnable parameters. The mean value  $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$  is generally parameterized by a network  $\epsilon_\theta(\mathbf{x}_t, t)$ :

$$\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{\bar{\alpha}_t}} \left( \mathbf{x}_t - \frac{\beta_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(\mathbf{x}_t, t) \right). \quad (4)$$

In practice, one can also directly approximate  $\hat{\mathbf{x}}_0$  from  $\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)$ :

$$\hat{\mathbf{x}}_0 = \frac{\mathbf{x}_t}{\sqrt{\bar{\alpha}_t}} - \frac{\sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(\mathbf{x}_t, t)}{\sqrt{\bar{\alpha}_t}}. \quad (5)$$

**Classifier Guidance.** The classifier guidance is employed to direct an unconditional diffusion model towards achieving conditional generation. Here,  $\mathbf{y}$  represents the target and  $p_\phi(\mathbf{y} | \mathbf{x})$  symbolizes a classifier, the conditional distribution is formulated to resemble a Gaussian distribution akin to its unconditional counterpart, but with the mean shifted by  $\Sigma_\theta(\mathbf{x}_t, t) \mathbf{g}$  (Dhariwal and Nichol 2021):

$$p_{\theta, \phi}(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{y}) \approx \mathcal{N}(\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) + \Sigma_\theta(\mathbf{x}_t, t) \mathbf{g}, \Sigma_\theta(\mathbf{x}_t, t)), \quad (6)$$

where  $\mathbf{g} := \nabla_{\mathbf{x}} \log p_\phi(\mathbf{y} | \mathbf{x})|_{\mathbf{x}=\boldsymbol{\mu}_\theta(\mathbf{x}_t, t)}$ . The gradient  $\mathbf{g}$  serves as a guidance that leads the unconditional sampling distribution towards the condition target  $\mathbf{y}$ .

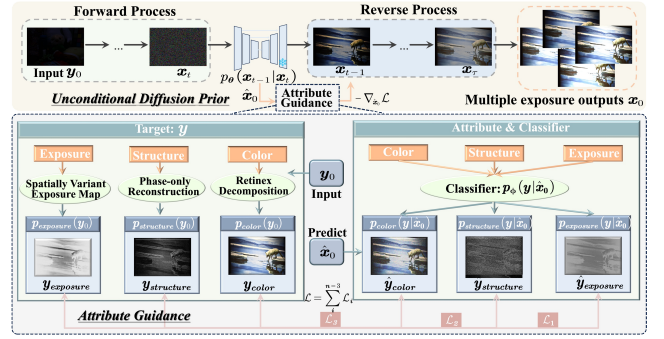


Figure 2: The overall framework of our proposed AGLLDiff.

### Overview of the AGLLDiff Framework

Our core motivation is to model the desired attributes of normal-light images and apply them to guide the diffusion generative process into a reliable high-quality (HQ) space. Such a design is agnostic to the degradation process and bypasses the difficulty of modeling the degradation process, making it more suitable for real-world LIE. The enhanced image should simultaneously satisfy: i) it is faithful to the degraded image, and ii) it conforms to the model distribution of pre-trained diffusion models that incorporate a vast repository of prior knowledge about HQ natural images. The overview of AGLLDiff is summarized in Fig. 2 and Algorithm 1. Given a degraded low-light image  $\mathbf{y}_0$  in the wild domain, the diffusion forward process adds a few steps of slight Gaussian noise to the  $\mathbf{y}_0$ , aiming to narrow the distribution between the degraded image and its potential counterpart, i.e., the HQ image. After obtaining the noisy image  $\mathbf{x}_t$ , we implement the reverse generative process through a pre-trained diffusion denoiser and attribute guidance to generate the enhancement result  $\mathbf{x}_0$ . The inherent attributes of normal-light samples, such as image exposure, structure and color, can be readily derived from their degraded counterparts. Further elaboration on the pivotal components of our method, including attributes, classifiers, and targets, is provided in subsequent sections.

#### Attribute Guidance

Our attribute guidance eschews any assumptions about degradation. Instead, with diffusion prior acting as a regularization, we provide guidance only on the desired attributes of HQ images. The key to AGLLDiff is to construct proper guidance on the generative process.

**Attribute and Classifier.** The initial step of AGLLDiff is to determine the desired attributes that the normal-light output possesses. Each of these attributes corresponds to a specific classifier  $\log p_\phi(\mathbf{y} | \mathbf{x}_0)$ , and the intermediate outputs  $\mathbf{x}_t$  are updated by back-propagating the gradient computed on the loss between the classifier output and the target  $\mathbf{y}$ . Through this mechanism, enhanced results can be obtained via an iterative refinement. The significance of attributes in achieving the desired final result cannot be overstated. This raises a fundamental question: *What attributes are possessed by HQ images? Through comprehensive ob-*

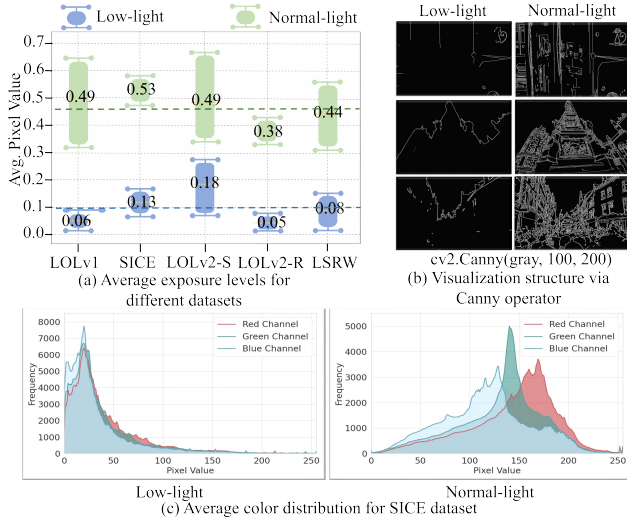


Figure 3: **Statistics and Visualization.** (a) The average exposure values of the low- and normal-light subsets of the five LIE datasets. (b) Visualization of the structure of low- and normal-light images by the Canny operator. (c) Histogram of the color distribution of the low- and normal-light images.

*ervation and statistical analysis, we conclude that the following three fundamental attributes are typically found in HQ images: 1) well exposure, 2) clear structure, and 3) vivid colors.* As presented in Fig. 3, we analyzed the average exposure levels and color distributions of five LIE datasets, visualizing the structure of low- and normal-light images. The observations are as follows: 1) there exists a significant discrepancy in average exposure values, with low-light images at around 0.1 and normal-light images at approximately 0.46, 2) normal-light images have clearer and richer structures compared to low-light images, and 3) the color distribution of images in normal light is both more colorful and homogeneous than in low-light images.

**Exposure Guidance Formulation.** To guide the output exposure toward that of normal-light images, we utilize the spatially variant exposure map (Li et al. 2022) to constrain the exposure of the output  $\hat{x}_0$ . The loss is formulated as follows:

$$\mathcal{L}_1 = \|\text{Mean}(\hat{x}_0) - \text{Mean}(E)\|_2^2, \quad (7)$$

where  $E$  denotes the spatially variant exposure map. Specifically, the exposure map is set with non-uniform exposure values in different regions, e.g., the underexposed region is assigned a large exposure value while the overexposed region is allocated a small exposure value. To achieve the spatially variant exposure map, we first obtain the luminance channel  $Y$  by color space decomposition (e.g., YCbCr, YUV) and its average value  $Y_{avg}$  of a low-light image  $y_0$ . Then we calculate the spatially variant condition exposure map by:

$$E = A \times \text{Norm}(Y - Y_{avg}) + B, \quad (8)$$

where  $B$  is the base exposure value,  $A$  is the adjustment amplitude, and Norm operation normalizes its input to the




Figure 4: (a) Visualization of the spatially variant exposure maps. (b) Visualization of the phase-only reconstruction image (PCI) in the spatial domain. (c) Visualization of the Retinex decomposition.

range of  $[-1, 1]$ . Based on extensive statistics in Fig. 3(a),  $B$  and  $A$  are empirically set to 0.46 and 0.25, respectively. Fig. 4(a) presents three outcomes of our method utilizing spatially variant exposure maps, where underexposed areas receive higher exposure values and well-exposed or overexposed areas receive lower ones. Such a spatially variant setting enables precise exposure adjustments and allows for generating multiple results at varying exposure levels by adjusting the exposure values.

**Structure Guidance Formulation.** For constraining the structure of the output faithful to the degraded image, we minimize the phase error between the degraded image  $y_0$  and the output  $\hat{x}_0$ :

$$\mathcal{L}_2 = \|\mathcal{P}(\hat{x}_0) - \mathcal{P}(y_0)\|_2^2, \quad (9)$$

where  $\mathcal{P}(\cdot)$  indicates the phase in the Fourier domain. In Fig. 4(b), we perform the inverse discrete Fourier transform to obtain phase-only reconstruction images in the spatial domain. As observed, the phase-only reconstruction versions of low-light and normal-light exhibit structural consistency. This is because most illumination information is expressed as amplitudes, and structural information is revealed in phases (Lin et al. 2024). We find that this simple phase constraint is sufficient to produce reliable results.

**Color Guidance Formulation.** According to the Retinex theory (Wei et al. 2018), a low-light image can be decomposed into illumination  $L$  and reflectance  $R$ . As shown in Fig. 4(c), the reflectance map  $R$  represents the physical properties of the objects, which contain abundant color information. Therefore, we could guide the color synthesis process with the reflectance map  $R$ . Equivalently, the loss is formulated as follows:

$$\mathcal{L}_3 = \|\mathcal{F}(\hat{x}_0) - \mathcal{F}(y_0)\|_2^2, \quad (10)$$

where  $\mathcal{F}(\cdot)$  denotes the pre-trained Retinex-based decomposition network (Fu et al. 2023), termed RNet. It takes a low-light image and generates the reflectance map  $R$  and illumination map  $L$ . RNet learns adaptive physical-based constraints from low-light image pairs in a self-supervised manner, significantly reducing the dependence on hand-crafted priors. Such an effective constraint ensures that our method can generalize well to various exposure scenes, which is why we chose it.

Method	Type	LOLv1			SICE			LOLv2-synthetic		
		PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
SDD (Hao et al. 2020)	T	13.34	0.63	0.74	15.34	0.73	0.26	16.46	0.73	0.35
LECARM (Ren et al. 2018)	T	14.40	0.54	0.32	18.59	0.78	0.26	17.44	0.76	0.37
MBLLEN (Lv et al. 2018)	S	15.25	0.70	0.32	18.41	0.73	0.31	18.16	0.80	0.28
RetinexNet (Wei et al. 2018)	S	17.60	0.64	0.38	19.57	0.78	0.27	17.41	0.67	0.34
DiffLL (Jiang et al. 2023)	S	26.19	0.85	0.11	21.33	0.84	0.22	29.46	0.92	0.09
PyDiff (Zhou, Yang, and Yang 2023)	S	27.56	0.87	0.10	21.18	0.83	0.23	26.13	0.92	0.08
CUE (Zheng et al. 2023)	S	22.67	0.79	0.20	20.06	0.82	0.24	24.47	0.90	0.12
Retinexformer (Cai et al. 2023)	S	25.15	0.84	0.13	22.32	0.85	0.20	25.66	0.95	0.05
EnlightenGAN (Jiang et al. 2021)	U	17.48	0.65	0.32	18.73	0.82	0.23	16.79	0.76	0.31
RUAS (Liu et al. 2021)	U	16.40	0.49	0.27	13.21	0.72	0.43	16.31	0.65	0.38
SCI (Ma et al. 2022)	U	14.78	0.52	0.33	15.94	0.78	0.45	18.07	0.77	0.27
PairLIE (Fu et al. 2023)	U	19.46	0.73	0.24	21.23	0.83	0.22	19.12	0.77	0.23
NeRCo (Yang et al. 2023)	U	19.81	0.73	0.24	20.73	0.82	0.23	19.14	0.74	0.26
ZeroDCE (Guo et al. 2020)	U	14.86	0.55	0.33	18.67	0.80	0.26	17.75	0.83	0.16
ZeroDCE++ (Li, Guo, and Loy 2021)	U	15.32	0.56	0.33	18.65	0.81	0.27	17.55	0.83	0.18
RRDNet (Zhu et al. 2020)	U	11.38	0.51	0.36	13.27	0.68	0.32	14.85	0.65	0.24
CLIP-LIT (Liang et al. 2023)	U	12.39	0.49	0.38	13.70	0.73	0.30	16.18	0.79	0.20
LightenDiffusion (Jiang et al. 2024)	U	20.44	0.81	0.19	20.59	0.83	0.22	20.44	0.84	0.16
GDP (Fei et al. 2023)	U	15.83	0.61	0.34	14.12	0.67	0.31	13.21	0.49	0.36
AGLLDiff (Ours)	U	21.81	0.84	0.15	22.12	0.84	0.21	21.11	0.87	0.13

Table 1: Quantitative comparison on LOLv1, SICE and LOLv2-synthetic. ‘‘T’’, ‘‘S’’ and ‘‘U’’ represent ‘‘Traditional’’, ‘‘Supervised’’ and ‘‘Unsupervised’’ methods, respectively. The best results of ‘‘S’’ and ‘‘U’’ are marked in blue and orange, respectively.

### Algorithm 1: Sampling with attribute guidance

**Require:** A pre-trained diffusion model ( $\mu_\theta(s_t, t), \Sigma_\theta(x_t, t)$ ), classifier  $p_\theta(y|x_0)$ , target  $y$ , gradient scale  $s$ , the number of gradient steps  $N$  and the iteration steps  $\omega$  of adding and removing noise.

**Input:** A low-light image  $y_0$

**Output:** Output image  $x_0$

$x_\omega \leftarrow \sqrt{\alpha_t}y_0 + \sqrt{1 - \alpha_t}\epsilon$

**for**  $t = \omega$  **to** 1 **do**

$\mu_t, \Sigma \leftarrow \mu_\theta(x_t, t), \Sigma_\theta(x_t, t)$

$\hat{x}_0 \leftarrow \frac{1}{\sqrt{\alpha_t}}x_t - \frac{\sqrt{1-\alpha_t}}{\alpha_t}\epsilon_\theta(x_t, t)$

$\hat{s} = \frac{\|x_t - x_{t-1}\|_2}{\|\nabla_{\hat{x}_0}\mathcal{L}\|_2} \cdot s$  ▷ Dynamic guidance scale

$\hat{N} \leftarrow \max\left(1, \frac{\|x_t - x_{t-1}\|_2}{\|\nabla_{\hat{x}_0}\mathcal{L}\|_2} \cdot N\right)$  ▷ Dynamic gradient steps

**repeat**

$x_t \leftarrow \text{sample from } \mathcal{N}(\mu_t - \hat{s}\Sigma\nabla_{\hat{x}_0} \log p_\theta(y|\hat{x}_0), \Sigma)$

$\hat{x}_0 \leftarrow \frac{1}{\sqrt{\alpha_t}}x_t - \frac{\sqrt{1-\alpha_t}}{\alpha_t}\epsilon_\theta(x_t, t)$

**until**  $\hat{N} - 1$  times

$x_{t-1} \leftarrow \mathcal{N}(\mu_t - \hat{s}\Sigma\nabla_{\hat{x}_0} \log p_\theta(y|\hat{x}_0), \Sigma)$

**end for**

**return**  $\hat{x}_0$

Our attribute guidance controls only the attributes of HQ outputs, and therefore composing the classifiers and summing the loss corresponding to each attribute can easily guide the diffusion model to generate HQ results:

$$\mathcal{L} = \lambda_1\mathcal{L}_1 + \lambda_2\mathcal{L}_2 + \lambda_3\mathcal{L}_3, \quad (11)$$

where  $\lambda_1$ ,  $\lambda_2$  and  $\lambda_3$  are constants controlling the relative importance of the different losses, which are empirically set to 1000, 10 and 0.03 in all experiments, respectively.

**Dynamic Guidance Scheme.** As illustrated in Fig. 2, given the desired attributes, we construct the corresponding classifier and apply classifier guidance during the generative process. Similarly to the external classifier gradient guidance in (Dhariwal and Nichol 2021), we evaluate the anti-

gradient  $-\nabla_{\hat{x}_0}\mathcal{L}$  to bring the attribute-guidance to the generative process. *However, our observations indicate that the traditional guidance scheme often leads to suboptimal outcomes.* Concretely, the traditional guidance scheme, which adopts a constant gradient scale  $s$ , often falls short in guiding the output towards the target value. Additionally, it executes merely one single gradient step per denoising step, which may not sufficiently steer the output towards the intended target, especially when early-phase denoising process intermediate outputs are significantly affected by noise. Such limitations are particularly adverse within LIE tasks that demand high similarity to the target. To mitigate this issue, we introduce a dynamic guidance scheme that consists of two distinct components, i.e., the dynamic guidance scale  $\hat{s}$  and the dynamic gradient steps  $\hat{N}$  at each denoising step. Specifically, we calculate the  $\hat{s}$  and  $\hat{N}$  based on the magnitude change of the intermediate image (Voynov, Aberman, and Cohen-Or 2023; Chefer et al. 2023):

$$\begin{aligned} \hat{s} &= \frac{\|x_t - x_{t-1}\|_2}{\|\nabla_{\hat{x}_0}\mathcal{L}\|_2} \cdot s \text{ and} \\ \hat{N} &= \max\left(1, \frac{\|x_t - x_{t-1}\|_2}{\|\nabla_{\hat{x}_0}\mathcal{L}\|_2} \cdot N\right), \end{aligned} \quad (12)$$

where  $x_{t-1} \sim \mathcal{N}(\mu_\theta, \Sigma_\theta)$ ,  $s$  and  $N$  are empirically set to 1.8 and 3 in all experiments, respectively. Such a dynamic guidance scheme affords users the flexibility to adjust the strength of guidance for attributes as per their unique requirements, thereby improving overall controllability.

## Experiments

### Implementation and Datasets

**Inference Requirements.** The pre-trained diffusion model we employ is a  $256 \times 256$  denoising network trained on the ImageNet dataset provided by (Dhariwal and Nichol 2021). The total number of iteration steps is defaulted to 1000. We select the final 10 steps to implement the noise addition and

Method	DICM			MEF			LIME			NPE			VV		
	NIQE↓	BRI↓	PI↓	NIQE↓	BRI↓	PI↓	NIQE↓	BRI↓	PI↓	NIQE↓	BRI↓	PI↓	NIQE↓	BRI↓	PI↓
SDD(T)	4.64	31.74	4.18	4.52	38.90	4.12	4.58	29.75	3.84	4.64	37.10	3.72	3.62	23.46	3.42
LECARM (T)	4.24	28.70	4.34	4.54	33.60	4.47	4.92	31.64	4.12	9.61	38.70	5.92	3.68	23.66	3.31
MBLLEN (S)	4.54	36.18	4.15	5.03	38.75	4.38	4.70	32.87	3.84	4.13	30.72	3.48	4.68	43.49	5.06
RetinexNet (S)	4.19	23.42	3.14	4.56	35.91	3.91	5.54	36.58	4.19	4.76	33.51	3.16	5.34	46.80	5.18
DiffLL (S)	<b>3.70</b>	18.08	<b>3.13</b>	3.46	23.27	<b>2.99</b>	<b>3.60</b>	19.44	3.06	<b>3.46</b>	<b>14.97</b>	<b>2.52</b>	<b>2.75</b>	<b>18.53</b>	<b>3.03</b>
PyDiff (S)	3.98	29.79	3.55	4.12	29.19	3.65	4.58	32.82	3.93	3.66	26.60	2.82	3.74	31.23	4.04
CUE (S)	3.76	17.87	3.32	3.63	25.81	3.21	3.83	16.90	3.08	3.53	19.82	2.75	3.54	21.83	3.96
Retinexformer (S)	3.72	<b>17.22</b>	3.08	<b>3.44</b>	<b>22.01</b>	3.07	3.86	<b>17.41</b>	<b>2.96</b>	3.39	20.51	2.56	2.97	<b>16.24</b>	3.35
EnlightenGAN (U)	3.89	29.23	3.41	3.56	25.31	3.25	4.21	25.34	3.34	3.66	24.10	2.95	3.63	27.79	3.82
RUAS (U)	6.58	45.06	5.22	5.40	41.68	4.57	5.36	31.62	4.34	7.12	46.89	5.61	4.86	34.03	4.24
SCI (U)	4.12	31.64	3.74	3.63	24.57	3.23	4.38	24.85	3.32	4.12	27.31	3.41	5.13	21.45	3.49
PairLIE (U)	4.13	28.59	3.69	4.18	29.54	3.10	4.51	25.21	3.26	4.17	26.22	3.03	3.66	25.88	3.69
NeRCo (U)	3.86	28.34	3.45	3.53	23.22	2.97	3.68	24.49	2.98	3.56	25.21	2.80	3.70	32.18	3.12
ZeroDCE (U)	3.91	24.05	3.13	3.51	26.63	3.13	4.34	26.48	3.21	3.80	22.34	2.92	4.12	25.13	3.33
ZeroDCE++ (U)	3.87	23.40	3.21	3.49	23.59	3.12	4.29	27.24	3.23	3.81	22.60	2.95	3.96	26.11	3.32
RRDNet (U)	3.81	23.95	3.21	3.55	25.92	3.35	4.35	35.23	3.81	3.69	25.51	3.24	3.65	29.37	3.28
CLIP-LIT (U)	3.71	25.78	3.24	3.57	27.72	3.22	3.99	24.41	3.07	3.62	24.14	2.74	3.33	28.66	3.09
LightenDiffusion (U)	3.70	26.59	3.32	3.57	23.92	3.12	4.25	20.61	3.27	3.56	23.83	2.74	3.59	25.28	2.92
GDP (U)	4.08	30.11	3.58	4.10	28.94	3.31	4.65	27.05	3.61	3.72	25.38	3.09	3.61	28.29	3.17
AGLLDiff (U)	<b>3.57</b>	<b>19.13</b>	<b>3.07</b>	<b>3.44</b>	<b>23.21</b>	<b>3.11</b>	<b>3.64</b>	<b>19.13</b>	<b>2.98</b>	<b>3.50</b>	<b>15.13</b>	<b>2.58</b>	<b>3.50</b>	<b>21.13</b>	<b>2.77</b>

Table 2: Quantitative comparison on DICM, MEF, LIME, NPE and VV. “T”, “S” and “U” represent “Traditional”, “Supervised”, “Unsupervised” methods, respectively. The best results of “S” and “U” are marked in blue and orange, respectively. BRI denotes BRISQUE.

attribute guidance. The inference process is carried out on the NVIDIA RTX 3090 GPU.

**Testing Datasets.** We construct one synthetic dataset and seven real-world datasets for testing. The LOLv1 (Wei et al. 2018) dataset is composed of 500 low-light and normal-light image pairs and divided into 485 training pairs and 15 testing pairs. The LOLv2-synthetic (Wei et al. 2018) dataset is officially divided into two parts, i.e., 900 low-light images for training and 100 low-light images for testing. The SICE benchmark collects 224 normal-light images and 783 low-light images. Each normal-light image corresponds to 2~4 low-light images. We adopt the first 50 normal-light images and the corresponding 150 low-light images for testing, and the rest (633 low-light images) for training. Moreover, we further assess our method on five commonly used real-world unpaired benchmarks: LIME (Li et al. 2018), NPE (Wang et al. 2013), MEF (Ma, Zeng, and Wang 2015), DICM (Lee, Lee, and Kim 2012), and VV (Vonikakis, Kouskouridas, and Gasteratos 2017). Notably, we only utilize the testing sets to evaluate our approach.

**Metrics.** For the paired datasets, we adopt two distortion-based metrics: PSNR and SSIM to evaluate the performance of the proposed method, and also the perceptual-based metric LPIPS to measure the visual quality of the enhanced results. For the other five unpaired datasets, we use three non-reference perceptual-based metrics: NIQE, BRISQUE, and PI for evaluation.

### Comparison with the State-of-the-Art

For a more comprehensive analysis, we compare our proposed AGLLDiff with three categories of existing state-of-the-art methods, including: 1) traditional methods SDD (Hao et al. 2020) and LECARM (Ren et al. 2018), 2) supervised approaches MBLLEN (Lv et al. 2018), RetinexNet (Wei et al. 2018), DiffLL (Jiang et al. 2023), PyDiff (Zhou, Yang, and Yang 2023), CUE (Zheng et al. 2023) and Retinexformer (Cai et al. 2023), and 3) unsupervised methods EnlightenGAN (Jiang et al. 2021), RUAS (Liu

et al. 2021), SCI (Ma et al. 2022), PairLIE (Fu et al. 2023), NeRCo (Yang et al. 2023), ZeroDCE (Guo et al. 2020), ZeroDCE++ (Li, Guo, and Loy 2021), RRDNet (Zhu et al. 2020), CLIP-LIT (Liang et al. 2023), LightenDiffusion (Jiang et al. 2024) and GDP (Fei et al. 2023). Note that the results of all those methods are reproduced by using the official codes with recommended parameters.

**Quantitative Comparisons.** Tables 1 and 2 report the quantitative performance of three paired datasets and five unpaired datasets, respectively. The best results of supervised and unsupervised methods are highlighted in blue and orange, respectively. Compared with recent competitive unsupervised approaches, AGLLDiff achieves the most advanced quantitative performance in terms of distortion-based and perceptual-based metrics across all benchmarks. Note that AGLLDiff even outperforms partial supervised approaches, which confirms the superiority of our solution.

**Visual Comparisons.** For a more comprehensive comparison, we further provide visual comparisons with leading algorithms in Figs 5 and 6. Our observations are twofold: 1) the proposed method distinctly surpasses other approaches in delivering aesthetically superior enhancements in terms of brightness, color fidelity, contrast, and natural appearance, especially under extremely low-light conditions where others falter; and 2) despite supervised approaches like Retinexformer, DiffLL, and CUE exhibiting notable efficacy on LOLv1, SICE and LOLv2-synthetic datasets, their generalization capabilities may be limited as supervised learning is sensitive to the data distribution. For more visual comparisons, please refer to the supplementary material.

### Ablation study

To assess the impact of our approach’s key components. We conduct several ablation studies on the LOLv1 dataset.

**Impact of the attributes.** We undertake ablation studies to verify the effectiveness of the mentioned attributes in Sec. . Concretely, we have tested the following three variations over the original setting: (i) without the exposure at-




Figure 5: Visual comparisons on SICE.




Figure 6: Visual comparisons on MEF.




Figure 7: Visual and quantitative results of ablation studies on LOL. The full model achieves the best performance.

tribute guidance. (ii) without the structure attribute guidance. (iii) without the color attribute guidance. Results are shown in Fig. 7. We have the following observations: 1) The removal of exposure attribute guidance limits users’ control over exposure adjustments. 2) The lack of structure attribute guidance leads to blurring in the structure. 3) The absence of color attribute guidance causes severe color distortions, and the objective measures degrade significantly. Our full model yields clear and natural outputs, validating the efficacy of the introduced attributes.

**Effectiveness of Dynamic Guidance Scale.** The effectiveness of the dynamic guidance scale  $\hat{s}$  is evaluated quantitatively and qualitatively. As illustrated in Fig. 8, without the dynamic guidance scale, although plausible results can be generated, the guidance scale must be manually adjusted for specific scenes, and the fidelity and clarity of the content cannot be guaranteed. In contrast, with the dynamic guidance scale  $\hat{s}$  replacing the constant guidance scale  $\hat{s}$ , high quality and fidelity results can be robustly delivered. The results highlight the indispensable role of our dynamic guidance scale in ensuring high fidelity to the target during the guidance process.

**Effectiveness of Dynamic Gradient Steps.** The dynamic gradient steps  $\hat{N}$  serve to adaptively adjust the strength of guiding the output toward the intended target. As depicted in Fig. 9(a), employing constant gradient steps yields sub-




Figure 8: Ablation studies on the dynamic gradient scale  $\hat{s}$ .




Figure 9: Ablation studies on dynamic gradient steps  $\hat{N}$  (a-b) and different iteration steps  $\omega$ .

optimal results, either under- or over-enhancement, with artifacts and noise. Conversely, in Fig. 9(b), artifacts and noise are progressively removed and finer details are generated. During the early phases of the denoising process, the  $\hat{N}$  is larger, while in the later stages,  $\hat{N}$  is smaller. Such a phenomenon suggests that the intermediate outputs are laden with noise in the early phases, and hence the gradient step should be increased to effectively steer the outputs towards the intended target. Whereas in the later phases, the gradient step should be decreased to produce refined results.

## Conclusion and Limitations

This manuscript introduces an Attribute Guidance Diffusion (AGLLDiff) framework to alleviate the challenges in real-world low-light image enhancement (LIE). AGLLDiff innovatively focuses on modeling desired high-quality image attributes such as exposure, structure and color, which do not depend on specific assumptions about the degradation process. This attribute-based guidance facilitates the diffusion sampling process towards achieving high-quality image recovery. Despite outstanding quantitative and qualitative performance achieved in eight challenging LIE benchmarks, there remain areas for improvement, such as accelerating sampling via advanced techniques and exploring more underlying high-quality attributes. Furthermore, future work will extend the application of this framework to various restoration challenges.

## Acknowledgments

The work was supported in part by the National Natural Science Foundation of China under Grant 82172033, U19B2031, 61971369, 52105126, 82272071, 62271430, and the Fundamental Research Funds for the Central Universities 20720230104.

## References

- Cai, Y.; Bian, H.; Lin, J.; Wang, H.; Timofte, R.; and Zhang, Y. 2023. Retinexformer: One-stage Retinex-based Transformer for Low-light Image Enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 12504–12513.
- Chai, W.; Guo, X.; Wang, G.; and Lu, Y. 2023. Stablevideo: Text-driven consistency-aware diffusion video editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 23040–23050.
- Chefer, H.; Alaluf, Y.; Vinker, Y.; Wolf, L.; and Cohen-Or, D. 2023. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *ACM Transactions on Graphics (TOG)*, 42(4): 1–10.
- Chen, S.; Ye, T.; Zhang, K.; Xing, Z.; Lin, Y.; and Zhu, L. 2025. Teaching Tailored to Talent: Adverse Weather Restoration via Prompt Pool and Depth-Anything Constraint. In *European Conference on Computer Vision*, 95–115. Springer.
- Dhariwal, P.; and Nichol, A. 2021. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34: 8780–8794.
- Fan, Z.; Cong, W.; Wen, K.; Wang, K.; Zhang, J.; Ding, X.; Xu, D.; Ivanovic, B.; Pavone, M.; Pavlakos, G.; Wang, Z.; and Wang, Y. 2024. InstantSplat: Unbounded Sparse-view Pose-free Gaussian Splatting in 40 Seconds. arXiv:2403.20309.
- Fan, Z.; Wang, K.; Wen, K.; Zhu, Z.; Xu, D.; and Wang, Z. 2023. LightGaussian: Unbounded 3D Gaussian Compression with 15x Reduction and 200+ FPS. arXiv:2311.17245.
- Fei, B.; Lyu, Z.; Pan, L.; Zhang, J.; Yang, W.; Luo, T.; Zhang, B.; and Dai, B. 2023. Generative Diffusion Prior for Unified Image Restoration and Enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9935–9946.
- Fu, Z.; Yang, Y.; Tu, X.; Huang, Y.; Ding, X.; and Ma, K.-K. 2023. 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*, 22252–22261.
- Guo, C.; Li, C.; Guo, J.; Loy, C. C.; Hou, J.; Kwong, S.; and Cong, R. 2020. Zero-Reference Deep Curve Estimation for Low-Light Image Enhancement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1777–1786.
- Guo, J.; Chai, W.; Deng, J.; Huang, H.-W.; Ye, T.; Xu, Y.; Zhang, J.; Hwang, J.-N.; and Wang, G. 2024. Versat2i: Improving text-to-image models with versatile reward. arXiv preprint arXiv:2403.18493.
- Hao, S.; Han, X.; Guo, Y.; Xu, X.; and Wang, M. 2020. Low-Light Image Enhancement With Semi-Decoupled Decomposition. *IEEE Transactions on Multimedia*, 22(12): 3025–3038.
- He, C.; Fang, C.; Zhang, Y.; Ye, T.; Li, K.; Tang, L.; Guo, Z.; Li, X.; and Farsiu, S. 2023a. Reti-diff: Illumination degradation image restoration with retinex-based latent diffusion model. arXiv preprint arXiv:2311.11638.
- He, C.; Li, K.; Xu, G.; Yan, J.; Tang, L.; Zhang, Y.; Wang, Y.; and Li, X. 2023b. Hqg-net: Unpaired medical image enhancement with high-quality guidance. *IEEE Transactions on Neural Networks and Learning Systems*.
- He, C.; Shen, Y.; Fang, C.; Xiao, F.; Tang, L.; Zhang, Y.; Zuo, W.; Guo, Z.; and Li, X. 2024. Diffusion Models in Low-Level Vision: A Survey. arXiv preprint arXiv:2406.11138.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Huang, S.-C.; Cheng, F.-C.; and Chiu, Y.-S. 2012. Efficient contrast enhancement using adaptive gamma correction with weighting distribution. *IEEE transactions on image processing*, 22(3): 1032–1041.
- Jiang, H.; Luo, A.; Fan, H.; Han, S.; and Liu, S. 2023. Low-light image enhancement with wavelet-based diffusion models. *ACM Transactions on Graphics (TOG)*, 42(6): 1–14.
- Jiang, H.; Luo, A.; Liu, X.; Han, S.; and Liu, S. 2024. LightenDiffusion: Unsupervised Low-Light Image Enhancement with Latent-Retinex Diffusion Models. In *European Conference on Computer Vision*.
- Jiang, Y.; Gong, X.; Liu, D.; Cheng, Y.; Fang, C.; Shen, X.; Yang, J.; Zhou, P.; and Wang, Z. 2021. Enlightengan: Deep light enhancement without paired supervision. *IEEE transactions on image processing*, 30: 2340–2349.
- Jin, Y.; Lin, B.; Yan, W.; Yuan, Y.; Ye, W.; and Tan, R. T. 2023. Enhancing visibility in nighttime haze images using guided apsf and gradient adaptive convolution. In *Proceedings of the 31st ACM International Conference on Multimedia*, 2446–2457.
- Jin, Y.; Yang, W.; and Tan, R. T. 2022. Unsupervised night image enhancement: When layer decomposition meets light-effects suppression. In *European Conference on Computer Vision*, 404–421. Springer.
- Kawar, B.; Elad, M.; Ermon, S.; and Song, J. 2022. Denoising diffusion restoration models. *Advances in Neural Information Processing Systems*, 35: 23593–23606.
- Lee, C.; Lee, C.; and Kim, C.-S. 2012. Contrast enhancement based on layered difference representation. In *2012 19th IEEE international conference on image processing*, 965–968. IEEE.
- Li, C.; Guo, C.; Feng, R.; Zhou, S.; and Loy, C. C. 2022. CuDi: Curve Distillation for Efficient and Controllable Exposure Adjustment. arXiv preprint arXiv:2207.14273.
- Li, C.; Guo, C.; and Loy, C. C. 2021. Learning to enhance low-light image via zero-reference deep curve estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(8): 4225–4238.
- Li, C.; Liu, H.; Liu, Y.; Feng, B. Y.; Li, W.; Liu, X.; Chen, Z.; Shao, J.; and Yuan, Y. 2024. Endora: Video generation models as endoscopy simulators. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, 230–240. Springer Nature Switzerland Cham.
- Li, M.; Liu, J.; Yang, W.; Sun, X.; and Guo, Z. 2018. Structure-Revealing Low-Light Image Enhancement Via

- Robust Retinex Model. *IEEE Transactions on Image Processing*, 27(6): 2828–2841.
- Liang, Z.; Li, C.; Zhou, S.; Feng, R.; and Loy, C. C. 2023. Iterative prompt learning for unsupervised backlit image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 8094–8103.
- Lin, Y.; Fu, Z.; Meng, G.; Wang, Y.; Dong, Y.; Fan, L.; Yu, H.; and Ding, X. 2023. Domain-irrelevant Feature Learning for Generalizable Pan-sharpening. In *Proceedings of the 31st ACM International Conference on Multimedia*, 3287–3296.
- Lin, Y.; Fu, Z.; Wen, K.; Ye, T.; Chen, S.; Meng, G.; Wang, Y.; Huang, Y.; Tu, X.; and Ding, X. 2024. Unsupervised Low-light Image Enhancement with Lookup Tables and Diffusion Priors. *arXiv preprint arXiv:2409.18899*.
- Liu, R.; Ma, L.; Zhang, J.; Fan, X.; and Luo, Z. 2021. Retinex-Inspired Unrolling With Cooperative Prior Architecture Search for Low-Light Image Enhancement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10561–10570.
- Lv, F.; Lu, F.; Wu, J.; and Lim, C. 2018. MBLLEN: Low-light image/video enhancement using CNNs. In *British Machine Vision Conference (BMVC)*, 1–13.
- Ma, K.; Zeng, K.; and Wang, Z. 2015. Perceptual quality assessment for multi-exposure image fusion. *IEEE Transactions on Image Processing*, 24(11): 3345–3356.
- Ma, L.; Ma, T.; Liu, R.; Fan, X.; and Luo, Z. 2022. Toward Fast, Flexible, and Robust Low-Light Image Enhancement. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5637–5646.
- Ren, Y.; Ying, Z.; Li, T. H.; and Li, G. 2018. LECARM: Low-light image enhancement using the camera response model. *IEEE Transactions on Circuits and Systems for Video Technology*, 29(4): 968–981.
- Vonikakis, V.; Kouskouridas, R.; and Gasteratos, A. 2017. On the evaluation of illumination compensation algorithms. *Multimedia Tools and Applications*, 77: 9211–9231.
- Voynov, A.; Aberman, K.; and Cohen-Or, D. 2023. Sketch-guided text-to-image diffusion models. In *ACM SIGGRAPH 2023 Conference Proceedings*, 1–11.
- Wang, S.; Zheng, J.; Hu, H.-M.; and Li, B. 2013. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE transactions on image processing*, 22(9): 3538–3548.
- Wang, Y.; He, X.; Dong, Y.; Lin, Y.; Huang, Y.; and Ding, X. 2024. Cross-Modality Interaction Network for Pan-sharpening. *IEEE Transactions on Geoscience and Remote Sensing*.
- Wang, Y.; Lin, Y.; Meng, G.; Fu, Z.; Dong, Y.; Fan, L.; Yu, H.; Ding, X.; and Huang, Y. 2023. Learning high-frequency feature enhancement and alignment for pan-sharpening. In *Proceedings of the 31st ACM International Conference on Multimedia*, 358–367.
- Wang, Y.; Xu, J.; and Zhao, Y. 2024. Wavelet encoding network for inertial signal enhancement via feature supervision. *IEEE Transactions on Industrial Informatics*.
- Wang, Y.; and Zhao, Y. 2023. Handwriting recognition under natural writing habits based on a low-cost inertial sensor. *IEEE Sensors Journal*.
- Wang, Y.; and Zhao, Y. 2024a. Scale and direction guided gan for inertial sensor signal enhancement. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24, International Joint Conferences on Artificial Intelligence Organization*, 5126–5134.
- Wang, Y.; and Zhao, Y. 2024b. Wavelet Dynamic Selection Network for Inertial Sensor Signal Enhancement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 15680–15688.
- Wei, C.; Wang, W.; Yang, W.; and Liu, J. 2018. Deep retinex decomposition for low-light enhancement. In *British Machine Vision Conference (BMVC)*, 1–12.
- Xia, J.; Yang, Z.; Li, S.; Zhang, S.; Fu, Y.; Gündüz, D.; and Li, X. 2024. Blind Super-Resolution Via Meta-Learning and Markov Chain Monte Carlo Simulation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Yang, P.; Zhou, S.; Tao, Q.; and Loy, C. C. 2024a. PGDiff: Guiding Diffusion Models for Versatile Face Restoration via Partial Guidance. *Advances in Neural Information Processing Systems*, 36.
- Yang, S.; Ding, M.; Wu, Y.; Li, Z.; and Zhang, J. 2023. Implicit neural representation for cooperative low-light image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12918–12927.
- Yang, S.; Zhang, X.; Wang, Y.; Yu, J.; Wang, Y.; and Zhang, J. 2024b. DiffLLE: Diffusion-based Domain Calibration for Weak Supervised Low-light Image Enhancement. *International Journal of Computer Vision*, 1–20.
- Ye, T.; Zhang, Y.; Jiang, M.; Chen, L.; Liu, Y.; Chen, S.; and Chen, E. 2022. Perceiving and Modeling Density for Image Dehazing. In *European Conference on Computer Vision*, 130–145. Springer.
- Yin, Y.; Xu, D.; Tan, C.; Liu, P.; Zhao, Y.; and Wei, Y. 2023. CLE Diffusion: Controllable Light Enhancement Diffusion Model. In *Proceedings of the 31st ACM International Conference on Multimedia*, 8145–8156.
- Zheng, N.; Zhou, M.; Dong, Y.; Rui, X.; Huang, J.; Li, C.; and Zhao, F. 2023. Empowering Low-Light Image Enhancer through Customized Learnable Priors. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 12559–12569.
- Zhou, D.; Yang, Z.; and Yang, Y. 2023. Pyramid Diffusion Models For Low-light Image Enhancement. *arXiv preprint arXiv:2305.10028*.
- Zhu, A.; Zhang, L.; Shen, Y.; Ma, Y.; Zhao, S.; and Zhou, Y. 2020. Zero-shot restoration of underexposed images via robust retinex decomposition. In *2020 IEEE International Conference on Multimedia and Expo (ICME)*, 1–6. IEEE.
- Zhu, Y.; Zhang, K.; Liang, J.; Cao, J.; Wen, B.; Timofte, R.; and Van Gool, L. 2023. Denoising Diffusion Models for Plug-and-Play Image Restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1219–1229.