

Low-Light Image Enhancement with Illumination-Aware Gamma Correction and Complete Image Modelling Network

Yinglong Wang^{1*} Zhen Liu^{2*} Jianzhuang Liu³ Songcen Xu⁴ Shuaicheng Liu^{5,2†}

¹Meituan Inc. ²Megvii Technology ³Shenzhen Institute of Advanced Technology

⁴Huawei Noah's Ark Lab ⁵University of Electronic Science and Technology of China

ylwanguestc@gmail.com, liuzhen03@megvii.com, jz.liu@siat.ac.cn

xusongcen@huawei.com, liushuaicheng@uestc.edu.cn

Abstract

This paper presents a novel network structure with illumination-aware gamma correction and complete image modelling to solve the low-light image enhancement problem. Low-light environments usually lead to less informative large-scale dark areas, directly learning deep representations from low-light images is insensitive to recovering normal illumination. We propose to integrate the effectiveness of gamma correction with the strong modelling capacities of deep networks, which enables the correction factor gamma to be learned in a coarse to elaborate manner via adaptively perceiving the deviated illumination. Because exponential operation introduces high computational complexity, we propose to use Taylor Series to approximate gamma correction, accelerating the training and inference speed. Dark areas usually occupy large scales in low-light images, common local modelling structures, e.g., CNN, SwinIR, are thus insufficient to recover accurate illumination across whole low-light images. We propose a novel Transformer block to completely simulate the dependencies of all pixels across images via a local-to-global hierarchical attention mechanism, so that dark areas could be inferred by borrowing the information from far informative regions in a highly effective manner. Extensive experiments on several benchmark datasets demonstrate that our approach outperforms state-of-the-art methods.

1. Introduction

Images captured in low-light environments usually contain large scales of dark areas. They are of low contrast and intensities, which submerge the useful image contents, making them invisible to humans as well as damaging the performances of numerous computer vision algorithms

*Equal contribution

†Corresponding author




Figure 1. This figure shows a challenging low-light image, whose middle areas are completely swallowed by the dark. Our IAGC obtains higher illumination and recovers more image details than SOTA methods. This is attributed to illumination-aware gamma correction which effectively enhances illumination as well as COMO-ViT which recovers better image details by completely modelling the dependencies among all pixels of images.

[12, 21, 25, 44]. Many attempts have been made to solve the Low-Light Image Enhancement (LLIE) task. Besides the early methods, e.g., histogram equalization [33], Retinex-based optimization methods [8, 31, 34, 23, 10], CNNs are introduced to learn illumination-recovering deep representations from low-light images. Many CNN-based methods [49, 51, 50, 41, 40] solve LLIE by combining Retinex theory [14], which decomposes an image \mathbf{I} into reflectance \mathbf{R} and illumination \mathbf{L} :

$$\mathbf{I} = \mathbf{R} \odot \mathbf{L} \quad (1)$$

where \odot is point-wise multiplication. There still exist some other deep networks learning a mapping from low-light images to normal-light ones, e.g., the supervised methods [2, 27, 15, 45, 18, 53, 4], and unsupervised methods [7, 13, 28].

However, the convolutional mechanism is usually limited by its inductive biases [30]. The locality is unable to

handle long-range dark areas well, and the spatial invariance regards dark and bright areas in the same manner, so that some useful information in the bright areas cannot be properly used. Therefore, CNN-based methods are insensitive to learning effective illumination-recovering deep representations, leading to still low illumination and inaccurate color recovery for some images.

Transformers are introduced to enhance low-light images by modelling longer-range context dependencies [52, 43]. An effective paradigm is to combine self-attention and convolution in a complementary way. However, these methods usually downsample an image before calculating self-attention to reduce computation. The dependencies among some pixels are inevitably ignored, losing some useful information and leading to unsatisfactory recovery results for some low-light images.

Allowing for the insensitivity of CNNs to handling low-light images, we propose to make the best use of the effectiveness of Gamma Correction to illumination enhancement. Existing methods usually neglect the differences between different image contents and even different images by using single empirical gamma values [8, 10, 23, 31, 34]. We propose to adaptively learn correction factors for different image contents according to their unique brightness by introducing an Illumination-Aware Gamma Correction network (IAGC). IAGC integrates the effectiveness of gamma correction with the strong modelling capacities of deep neural networks, so that more effective gamma can be evaluated via perceiving different image contents.

IAGC adopts a coarse-to-fine strategy to learn correction factors. Coarsely, we globally recover the brightness of low-light images by learning a whole correction factor with a Global Gamma Correction Module (GGCM) perceiving global illumination. Finely, we further build a Local Gamma Correction Module (LGCM) to locally explore pixel dependencies to elaborately enhance illumination according to different local image contents. An example is shown in Fig. 1 to illustrate our improvement on visual quality over existing state-of-the-art LLIE works.

In our work, gamma is learnable, gamma correction thus introduces exponential operation in both forward and backward propagation, bringing high computational complexity. We propose to use Taylor Series to approximate gamma correction so that exponential operations are avoided in both forward and backward propagation.

Low-light images usually contain large-scale dark areas, which are ineffective to be enhanced by CNNs due to their locality bias. Similarly, calculating attention patchwisely [24] or on a downsampled image [43, 52] is also verified to be not enough to handle such large dark areas. We propose to completely model the dependencies of all pixels in an image, so that information from far bright regions can be used to recover dark areas. To reach such a

goal, we propose a Completely Modelling Vision Transformer (COMO-ViT) block to describe images via a local-to-global hierarchical self-attention mechanism. Extensive experiments in later sections verify the stronger modelling capacity of our COMO-ViT for modelling long-range image dependencies than state-of-the-art Transformer blocks.

The main contributions of this paper are summarized as:

- We are the first to propose utilizing learnable gamma correction to adaptively solve the insensitivity of existing models to extract effective illumination-recovering deep representation. To reach this goal, two illumination-aware gamma correction modules (GGCM and LGCM) are introduced to predict global and local adaptive correction factors. Moreover, Taylor Series are used to approximate exponential operation to improve running efficiency.
- We propose to completely express images by exploring extensive dependencies of all pixels. A novel vision Transformer (COMO-ViT) block is thus introduced to connect all pixels via a local-to-global hierarchical self-attention. COMO-ViT is verified to be able to extract highly-effective deep features and keep strong modelling capacities.
- We introduce an effective network architecture, i.e., IAGC, for low-light image enhancement by combining GGCM, LGCM, and COMO-ViT block together. Extensive experiments show that we obtain better quantitative evaluation and hue recovery than SOTA methods compared with ground truth.

2. Related Work

In this section, we comprehensively review existing works on low-light image enhancement and analyze their limitations in producing satisfactory results.

Conventional Methods. Many conventional LLIE methods are based on Retinex theory, which formulates an image as a multiplication of reflectance and illumination (Eq. (1)). Some methods estimate the reflectance and illumination via a variational framework, the estimated illumination is further adjusted to restore the original low-light image [6, 16, 31]. Other Retinex-based methods usually optimize an energy function derived from the Maximum-a-Posteriori (MAP) framework to enhance low-light images, in which some heuristic priors are designed to constrain the properties of reflectance or illumination [8, 10, 23, 34]. For example, Gaussian total variation in [10], structure-aware regularization in [8] and the additional noise term in [23]. However, these hand-craft priors are usually inaccurate and limited when applied to different cases, leading to apparent color deviation.

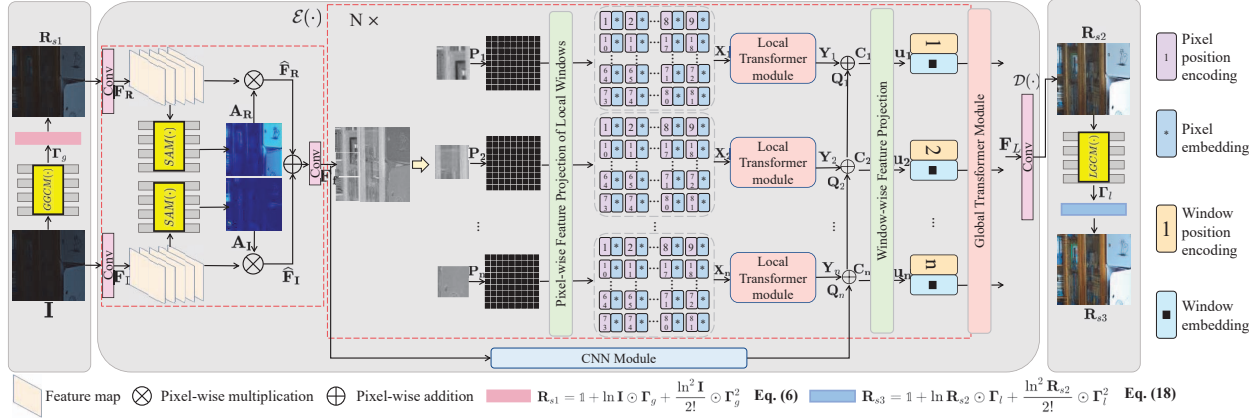


Figure 2. Overview of our IAGC framework. We adopt a three-stage illumination enhancement strategy. Firstly, we coarsely enhance \mathbf{I} via a global gamma correction module (GGCM). Secondly, an encoder-decoder structure is used to further enhance illumination and remove noise. Finally, we introduce another local gamma correction module (LGCM) to elaborately tune the illumination.

CNN-based Learning Methods. CNN has pushed forward the solution of LLIE by large margins [20]. Some methods integrate Retinex theory with CNNs [40, 49, 51, 50, 41, 40]. For example, Wei *et al.* propose an end-to-end Retinex-Net, in which a decomposition module and an illumination adjustment module are used to learn reflectance and illumination [40]. Wu *et al.* propose URetinex-Net which unfolds an optimization problem into several learnable network modules [41]. Some methods design effective networks to predict normal-light images directly from low-light ones via a supervised learning strategy [2, 27, 15, 45, 18, 53, 4]. For instance, [45] learns a 3D lookup table, and [53] uses color consistency to constrain network training. There are still some unsupervised methods to solve LLIE [7, 13, 28]. One famous work is [28], in which a self-calibrated illumination learning framework is proposed for fast, flexible, and robust low-light image enhancement. Because CNNs are usually limited by inductive bias, the illumination still cannot be well recovered for some images.

Transformer-based Methods. Transformer-based methods usually combine self-attention and convolution to extract long- and short-range dependencies, so that better performance can be obtained [43, 52, 37]. However, the long-range context dependencies are usually obtained by building patch-level relationships, and these methods downsample images before calculating self-attention to reduce computation. Hence, they may lose some useful information, which is adverse to recover high-quality images.

In this paper, we propose a local-to-global hierarchical self-attention to completely express images. We would like to point out that the work [9] aiming at high-level tasks proposes a TNT structure to explore local attention to assist global attention. Our work is different from it in the following aspects: 1) TNT still focuses on patch-level attention,

pixel-level dependencies are still not explored, i.e., it does not completely model images. While ours completely express images via pixel-level attention. 2) Local attention in TNT just provides auxiliary information for the patch embedding of global attention. While our local attention models pixel dependencies locally and the global attention extends pixel dependencies to the whole image. 3) TNT adopts a bypass structure to realize auxiliary information enhancement. While our local attentions are directly embedded as the input for global attention.

3. Our Method

We first overview the proposed IAGC in Sec. 3.1. Then, we present its architecture with illumination-aware gamma correction modules (GGCM and LGCM) and the completely modelling Transformer (COMO-ViT) block in Sec. 3.2. Finally, we introduce loss functions used to train IAGC in Sec. 3.3.

3.1. Overview

The overall architecture of IAGC is illustrated in Fig. 2. Given a low-light image \mathbf{I} , we propose to enhance its illumination via a three-stage coarse-to-fine strategy. In the first stage, we introduce a global gamma correction module (GGCM) to extract illumination features and learn a whole correction factor Γ_g to coarsely enhance \mathbf{I} :

$$\Gamma_g = \text{GGCM}(\mathbf{I}), \quad \mathbf{R}_{s1} = \mathbf{I}^{\Gamma_g}, \quad (2)$$

where \mathbf{R}_{s1} is the enhanced result. Compared with \mathbf{I} , \mathbf{R}_{s1} is of higher brightness and many informative contents emerge from the dark areas. Experiments also show that GGCM enables attention to capture more informative image contents, as shown in Fig. 3, providing more knowledge for the second stage. Therefore, our GGCM is effective to solve the in-

sensitivity of existing methods to enhance illumination and recover the swallowed contents in the low-light image \mathbf{I} .

In the second stage, we fuse the information in \mathbf{I} and \mathbf{R}_{s1} and build a backbone encoder $\mathcal{E}(\cdot)$ to learn illumination-recovered representations \mathbf{F} from both of them, and then further enhance brightness and recover image details via a decoder $\mathcal{D}(\cdot)$. Because low-light images usually contain large-scale dark areas and only a small part may possess relatively higher illumination, it is important to completely model pixel-level dependencies of images to infer the contents in the dark based on the informative bright areas. We introduce a novel vision Transformer block COMO-ViT in $\mathcal{E}(\cdot)$ to connect all pixels via a local-to-global self-attention. Due to the strong modelling capacity of COMO-ViT, the noise which may appear in the illumination enhancement [43] is also implicitly removed during this stage, as shown in Fig. 4. We formulate this stage as:

$$\mathbf{F} = \mathcal{E}(\mathbf{I}, \mathbf{R}_{s1}), \quad \mathbf{R}_{s2} = \mathcal{D}(\mathbf{F}), \quad (3)$$

where \mathbf{R}_{s2} is the enhanced result of the second stage.

In order to obtain better hue recovery compared with ground truth, we introduce a local gamma correction module (LGCM) to further perceive the illumination gaps between \mathbf{R}_{s2} and the ground truth. The pixel-wise correction factor Γ_l is predicted to enhance illumination in a more elaborate manner, which is formulated as:

$$\Gamma_l = \text{LGCM}(\mathbf{R}_{s2}), \quad \mathbf{R}_{s3} = \mathbf{R}_{s2}^{\Gamma_l}, \quad (4)$$

where \mathbf{R}_{s3} is our final low-light enhancement result.

Different from existing models learning illumination-recovered representation directly from the less informative low-light images, we introduce GGCM and LGCM to perceive illumination differences and adaptively predict correction factors, so that the degraded illumination and image details are effectively recovered. Complementarily, our COMO-ViT explores better pixel dependencies of an image via a local-to-global hierarchical self-attention. Hence, our IAGC makes better use of the bright pixels to recover every dark pixel than existing Transformer-based models which only develop patch-level dependencies [43, 52].

3.2. Architecture of IAGC

In this section, we give the detailed structure of our IAGC and meanwhile show some intermediate results to illustrate the effectiveness of our key designs.

GGCM. GGCM takes \mathbf{I} as input, perceives its illumination, and predicts a whole correction factor Γ_g via exploring global deep features. Its architecture is formulated as:

$$\Gamma_g = \phi(\text{FC}(\text{Avg-Pool}(\text{Conv}(\mathbf{I})))), \quad \mathbf{R}_{s1} = \mathbf{I}^{\Gamma_g}, \quad (5)$$

where $\text{Conv}(\cdot)$ is a convolutional layer projecting \mathbf{I} into feature space, $\text{Avg-Pool}(\cdot)$ is an average pooling layer which




Figure 3. This figure shows that GGCM enables SAM to capture more informative image contents, which provides more knowledge for feature learning at the next stage. The learned correction factor $\Gamma_g = 0.5976$ for this example.

extracts global features by squeezing spatial dimensions, $\text{FC}(\cdot)$ is a fully-connected layer for fusing channel information, and $\phi(\cdot)$ is a Sigmoid layer to predict Γ_g . For simplicity, we omit the operation expanding a scalar to a matrix Γ_g . Note that Eq. (5) introduces exponential operation with high complexity during both forward and backward propagation, we utilize Taylor Series to approximate it:

$$\mathbf{R}_{s1} = \mathbb{1} + \ln \mathbf{I} \odot \Gamma_g + \frac{\ln^2 \mathbf{I}}{2!} \odot \Gamma_g^2 \quad (6)$$

where $\mathbb{1}$ is all-1 matrix. Fig. 3 visualizes an example of \mathbf{I} and corresponding \mathbf{R}_{s1} , showing that more image details emerge in \mathbf{R}_{s1} , especially at the middle regions.

We then fuse \mathbf{I} and \mathbf{R}_{s1} to supply more knowledge for the second stage by projecting them into feature space:

$$\mathbf{F}_I = \text{Conv}(\mathbf{I}), \quad \mathbf{F}_R = \text{Conv}(\mathbf{R}_{s1}). \quad (7)$$

Because bright pixels are more informative and play a more important role in illumination enhancement than dark ones, we adopt a Spatial Attention Module (SAM) to highlight bright areas by large weights and constrain dark pixels to reduce their negative effect by smaller weights:

$$\begin{aligned} \mathbf{A}_I &= \phi(\text{Conv}(\mathbf{F}_I)), \quad \widehat{\mathbf{F}}_I = \mathbf{A}_I \odot \mathbf{F}_I, \\ \mathbf{A}_R &= \phi(\text{Conv}(\mathbf{F}_R)), \quad \widehat{\mathbf{F}}_R = \mathbf{A}_R \odot \mathbf{F}_R, \end{aligned} \quad (8)$$

where \mathbf{A}_I and \mathbf{A}_R are the spatial attention map of \mathbf{F}_I and \mathbf{F}_R , respectively. In Fig. 3, we visualize \mathbf{A}_I and \mathbf{A}_R , observing that \mathbf{A}_R captures more informative image contents, which further illustrates the effectiveness of our GGCM. Finally, the fused deep feature \mathbf{F}_f is obtained as:

$$\mathbf{F}_f = \text{Conv}(\widehat{\mathbf{F}}_I + \widehat{\mathbf{F}}_R). \quad (9)$$

Epochs	Optimizer	Batch size	Learning rate	LR decay	Weight decay	Drop path	Embedding dim	Head	L	Window size	Patch size
300	Adam[17]	8	4e-4	cosine	1e-7	0.1	15	5	2	16	512

Table 1. Default training and network hyper-parameters used in our method, unless stated otherwise.

COMO-ViT. Given the input feature $\mathbf{F}_{l-1} \in \mathbb{R}^{H \times W \times c}$ of COMO-ViT, we conduct two branches of operations. In the first branch, we uniformly split it into n non-overlapping windows $\mathcal{P} = [\mathbf{P}^1, \mathbf{P}^2, \dots, \mathbf{P}^n] \in \mathbb{R}^{n \times w \times w \times c}$, where (w, w) is the window resolution. SNR [43] and STAR [52] downsample images, losing local structures and some important pixel-level information. Instead, the proposed COMO-ViT completely models the dependencies among all pixels of an image via a local-to-global hierarchical self-attention. Locally, each pixel in a window \mathbf{P}^i is regarded as an individual, we thus reshape \mathbf{P}^i as follows:

$$\mathbf{P}^i \rightarrow [p^{i,1}, p^{i,2}, \dots, p^{i,m}], \quad (10)$$

where $p^{i,j} \in \mathbb{R}^{1 \times 1 \times c}$, $m = w^2$ is the number of pixels in \mathbf{P}^i . With a linear projection, we then transform the pixels into a sequence of pixel embeddings $\mathbf{X}^i = [x^{i,1}, x^{i,2}, \dots, x^{i,m}]$, where $x^{i,j} \in \mathbb{R}^{c1}$ is the j -th pixel embedding, $c1$ is the embedding dimension. For \mathbf{X}^i , we utilize a local Transformer module to extract deep features as follows:

$$\begin{aligned} \mathbf{Y}'^i &= \mathbf{X}^i + \text{MSA}(\text{LN}(\mathbf{X}^i)), \\ \mathbf{Y}^i &= \mathbf{Y}'^i + \text{MLP}(\text{LN}(\mathbf{Y}'^i)), \end{aligned} \quad (11)$$

where \mathbf{Y}^i is the feature learned by the local Transformer module, $\text{MSA}(\cdot)$ is the Multi-head Self-Attention [35], $\text{LN}(\cdot)$ is layer normalization [1] for stable training and faster convergence, $\text{MLP}(\cdot)$ is multi-layer perceptron for feature transformation at channel dimension and non-linearity. In such a process, we adopt 1D learnable location embedding to encode the spatial information of pixels.

To complement the non-overlapping window attention, in the second branch which is parallel with local attention, we use a CNN module to model local pixel dependencies in \mathbf{F}_{l-1} via an overlapped sliding kernel to recover image details, in which a SE block [11] is used to explore channel relationship to boost representative power:

$$\mathbf{F}' = \text{Conv}(\text{LN}(\mathbf{F}_{l-1})), \quad \mathbf{F}_{conv} = \mathbf{F}' \odot \text{SE}(\mathbf{F}'). \quad (12)$$

\mathbf{F}_{conv} is then split into n non-overlapping windows $\mathcal{Q} = [\mathbf{Q}^1, \mathbf{Q}^2, \dots, \mathbf{Q}^n] \in \mathbb{R}^{n \times w \times w \times c}$, and each \mathbf{Q}^i is reshaped:

$$\mathbf{Q}^i \rightarrow [q^{i,1}, q^{i,2}, \dots, q^{i,m}]. \quad (13)$$

We combine the features from both branches as:

$$\mathcal{C} = [\mathbf{C}^1, \mathbf{C}^2, \dots, \mathbf{C}^n], \quad \mathbf{C}^i = \mathbf{Q}^i + \mathbf{Y}^i \quad (14)$$

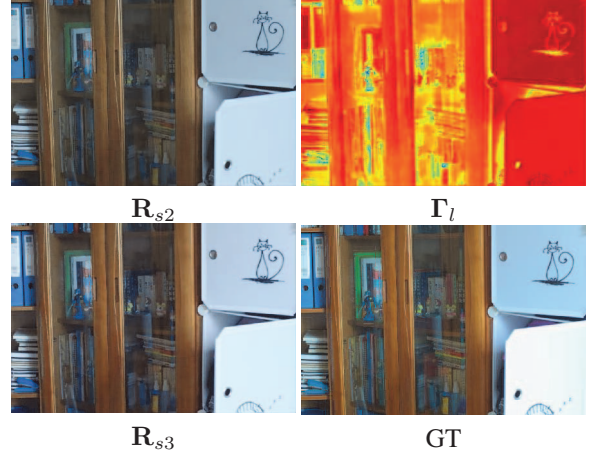


Figure 4. We observe that \mathbf{R}_{s3} keeps higher illumination than \mathbf{R}_{s2} , illustrating the effectiveness of our LGCM. Γ_l is corresponding pixel-wise local gamma map.

Global pixel dependencies are explored by calculating window attention via a global attention module. Firstly, \mathcal{C} is transformed into a sequence of window embedding:

$$\mathbf{U} = [u^1, u^2, \dots, u^n], \quad u^i = \text{FC}(\text{Vec}(\mathbf{C}^i)), \quad (15)$$

where $\text{Vec}(\cdot)$ is vectorization operation. Then, we utilize a global Transformer module to explore inter-window dependencies, obtaining the feature $\mathbf{F}_l \in \mathbb{R}^{H \times W \times c}$, where $l \in \{1, 2, \dots, L\}$, and L is the COMO-ViT number. When $l = 1$, \mathbf{F}_{l-1} is the fused feature \mathbf{F}_f in Eq. (9). The result of the second stage is obtained by decoding \mathbf{F}_L with a convolutional layer ($\mathcal{D}(\cdot)$):

$$\mathbf{R}_{s2} = \text{Conv}(\mathbf{F}_L). \quad (16)$$

We visually show \mathbf{R}_{s2} in Fig. 4, observing that the illumination is enhanced and image details are also recovered. Especially, the noise is well removed by our COMO-ViT.

LGCM. LGCM takes \mathbf{R}_{s2} as input, and learns local deep features to perceive illumination gap between \mathbf{R}_{s2} and ground truth and elaborately enhances illumination to reduce local color deviation:

$$\Gamma_l = \phi(\text{Conv}_3(\mathbf{R}_{s2})), \quad \mathbf{R}_{s3} = \mathbf{R}_{s2}^{\Gamma_l}. \quad (17)$$

Similarly, we also apply Taylor Series to approximate exponential operation as follows:

$$\mathbf{R}_{s3} = \mathbb{1} + \ln \mathbf{R}_{s2} \odot \Gamma_l + \frac{\ln^2 \mathbf{R}_{s2}}{2!} \odot \Gamma_l^2. \quad (18)$$

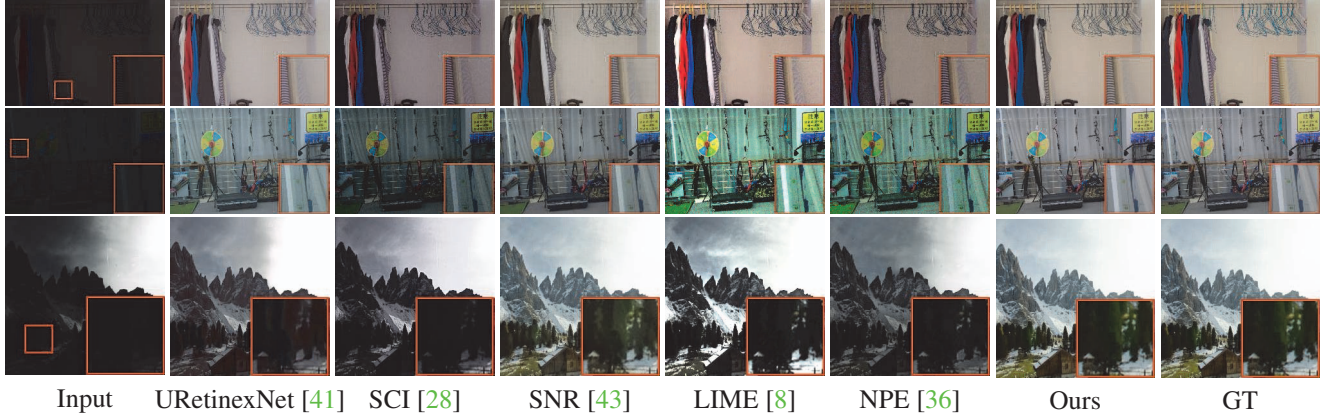


Figure 5. Visual quality comparison with SOTA methods on the LOL-v1 (the first row), LOL-v2-real (the second row), and LOL-v2-synthetic (the third row) datasets, respectively. Our method recovers better illumination and preserves more image details.

Methods	SNR [43]	URetinex-Net[41]	DeepLPF [29]	SCI [28]	LIME [8]	MF [5]	NPE [36]	SRIE [6]	SDD [10]	CDEF [19]	LPNet [22]
PSNR	24.49	21.32	15.28	15.80	16.76	16.96	16.96	11.86	13.34	16.33	21.46
SSIM	0.840	0.835	0.473	0.527	0.444	0.505	0.481	0.493	0.635	0.583	0.802
Methods	DRBN [46]	KinD [51]	RUAS [26]	FIDE [42]	EG [13]	MS-RDN [47]	Retinex-Net[40]	MIRNet [48]	IPT [3]	Uformer [39]	Ours
PSNR	20.13	20.87	18.23	18.27	17.48	17.20	16.77	24.14	16.27	16.36	24.53
SSIM	0.830	0.800	0.720	0.665	0.650	0.640	0.560	0.830	0.504	0.507	0.842

Table 2. Quantitative comparisons with state-of-the-art methods on the LOL-v1 dataset.

Γ_l and \mathbf{R}_3 are shown in Fig. 4. We observe that the illumination is enhanced further, obtaining better visual quality.

3.3. Loss Functions

Our IAGC adopts a three-stage strategy to enhance low-light images. In each stage, we adopt the same loss functions to constrain the network training. The details are:

$$\mathcal{L} = \sum_{i=1}^3 (\sqrt{\|\mathbf{R}_{si} - \mathbf{G}\|^2 + \epsilon^2} + \|\nabla \mathbf{R}_{si} - \nabla \mathbf{G}\|^2). \quad (19)$$

The first term is the Charbonnier loss constraining network in the spatial domain. The second term is a gradient loss to supervise the network in the gradient domain, assisting the Charbonnier loss to recover more image details. \mathbf{R}_{si} is the result of stage i , $\epsilon = 10^{-3}$, \mathbf{G} is ground truth, ∇ denotes the combination of the horizontal and vertical gradients.

4. Experiments

4.1. Datasets and Experimental Setting

Datasets. We evaluate our method on the widely-used LOL (including v1 and v2) datasets. LOL-v1 [40] contains 485 real low-/normal-light pairs for training and 15 pairs

for testing. LOL-v2 [47] is composed of two sets. LOL-v2-real contains 689 pairs for training and 100 pairs for testing. The majority of the low-light images are collected by changing exposure time and ISO while fixing other camera configurations. LOL-v2-synthetic is collected by analyzing the illumination distribution so that synthetic images match the property of real dark photography. This set contains 900 synthetic pairs for training and 100 pairs for testing.

Experimental Setting. Our IAGC adopts an end-to-end training strategy to optimize the network parameters of 3 stages simultaneously. Our model is implemented with PyTorch [32] and trained on an NVIDIA v100NV32 GPU. Other hyper-parameters of training and network structure are listed in Table 1 for clarity.

4.2. Comparison with Current Methods

We compare the results of the proposed IAGC with a rich collection of SOTA methods to make comprehensive quantitative comparisons. For qualitative comparison, we select some very recent or robust methods to visually show their performance, including SNR [43], URetinex-Net [41], DeepLPF [29], SCI [28], LIME [8], MF [5], NPE [36], SRIE [6], SDD [10], CDEF [19], LPNet [22], DRBN [46], KinD [51], RUAS [26], FIDE [42], EG [13], MS-RDN [47], Retinex-Net [40], MIRNet [48]. Moreover, following [43],

Methods	SNR [43]	URetinex -Net[41]	DeepLPF [29]	SCI [28]	LIME [8]	MF [5]	NPE [36]	SRIE [6]	SDD [10]	CDEF [19]	LPNet [22]
PSNR	21.36	21.22	14.10	17.79	15.24	18.73	17.33	17.34	16.64	19.75	17.80
SSIM	0.842	0.860	0.480	0.568	0.470	0.559	0.461	0.686	0.677	0.627	0.792
Methods	DRBN [46]	KinD [51]	RUAS [26]	FIDE [42]	EG [13]	MS-RDN [47]	Retinex -Net[40]	MIRNet [48]	IPT [3]	Uformer [39]	Ours
PSNR	20.29	14.74	18.37	16.85	18.23	20.06	15.47	20.02	19.80	18.82	22.20
SSIM	0.831	0.641	0.723	0.678	0.617	0.816	0.567	0.820	0.813	0.771	0.863

Table 3. Quantitative comparisons with state-of-the-art methods on the LOL-v2-real dataset.

Methods	SNR [43]	URetinex -Net[41]	DeepLPF [29]	SCI [28]	LIME [8]	MF [5]	NPE [36]	SRIE [6]	SDD [10]	CDEF [19]	LPNet [22]
PSNR	23.84	18.75	16.02	17.73	16.88	17.50	16.59	14.50	16.46	18.56	19.51
SSIM	0.905	0.829	0.587	0.763	0.776	0.751	0.778	0.616	0.728	0.842	0.846
Methods	DRBN [46]	KinD [51]	RUAS [26]	FIDE [42]	EG [13]	MS-RDN [47]	Retinex -Net[40]	MIRNet [48]	IPT [3]	Uformer [39]	Ours
PSNR	23.22	13.29	16.55	15.20	16.57	22.05	17.13	21.94	18.30	19.66	25.58
SSIM	0.927	0.578	0.652	0.612	0.734	0.905	0.798	0.876	0.811	0.871	0.940

Table 4. Quantitative comparisons with state-of-the-art methods on the LOL-v2-synthetic dataset.

another two classical architecture IPT [3] and Uformer [39] are also selected to make more extensive comparisons.

Quantitative Comparisons. We use PSNR and SSIM [38] as quantitative metrics. Table 2, 3 and 4 show that our method obtains consistent best values. Especially on the LOL-v2-real and LOL-v2-synthetic datasets, our method surpasses the second best SNR [43] by large margins. Note that these numbers are obtained either by running their respective codes or from their respective papers.

Visual Quality Comparisons. Due to space limitation, we select the very recent URetinexNet [41], SCI [28], SNR [43] and two classic LIME [8] and NPE [36] to visually make comparison with our method. Fig. 5 shows three low-light images respectively from LOL-v1 (the first line), LOL-v2-real (the second line), and LOL-v2-synthetic (the third line). We observe that LIME and NPE recover good image details and color hue for the first image, but apparent noise remains. They do not produce good results on the second and third images. SCI may produce slight noise, e.g., in the first and second images. Moreover, this method does not recover the illumination well. URetinexNet obtains good performance on the first and second images but does not perform well on the third synthetic image. SNR performs well on all these three images. Especially, this method handles the noise well. By comparison, our method enhances better illumination and preserves more details.

4.3. Ablation Studies

Effect of Gamma Correction Modules. Illumination-aware gamma correction modules, i.e., GGCM and LGCM

Datasets	\mathcal{G}_1		\mathcal{G}_2		\mathcal{G}_3	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LOL-v1	23.18	0.821	23.96	0.838	24.53	0.842
LOL-v2-real	21.32	0.834	21.87	0.851	22.20	0.863
LOL-v2-syn	24.88	0.912	25.26	0.925	25.58	0.940

Table 5. Quantitative results of our ablation study on the proposed gamma correction modules GGCM and LGCM.



Figure 6. Qualitative comparisons of our ablation study on the proposed GGCM and LGCM.

play important roles in illumination enhancement. GGCM is to perceive the whole illumination and adaptively predict a global gamma correction factor to enhance brightness. LGCM is aware of illumination differences by exploring local pixel dependencies and predicts gamma correction factors pixel-wisely. We verify their effect by ablating them separately. \mathcal{G}_1 is the model ablating GGCM, \mathcal{G}_2 is the one removing LGCM, \mathcal{G}_3 is our whole network. As shown in Table 5, SSIM does not have too large differences for these three cases. The PSNR of \mathcal{G}_1 decreases more largely than that of \mathcal{G}_2 , illustrating that GGCM plays a more important role in illumination enhancement than LGCM. Fig. 6 visually shows similar results, when ablating GGCM, the hue

Datasets	\mathcal{A}_1	\mathcal{A}_2	\mathcal{A}_3	\mathcal{A}_4
LOL-v1	23.72	24.13	24.21	24.53
	0.797	0.816	0.829	0.842
LOL-v2-real	21.54	21.93	22.04	22.20
	0.811	0.834	0.851	0.863
LOL-v2-syn	24.89	25.27	25.31	25.58
	0.902	0.917	0.923	0.940

Table 6. Quantitative results (PSNR/SSIM) of our ablation study on the proposed local-to-global self-attention mechanism.




Figure 7. Qualitative comparisons of our ablation study on the proposed local-to-global self-attention mechanism.

deviates more apparently from that of ground truth.

Effect of Local-to-Global Self-Attention. COMO-ViT is designed to completely model the dependencies of all pixels via a local-to-global hierarchical self-attention mechanism. The local attention is to explore pixel relationships in a window. The global attention connects different windows so that all pixels across the whole image could be combined together and completely modeled. Here, we test their effect on illumination enhancement. We define \mathcal{A}_1 as the model ablating local attention, which degrades to a patch-embedding Transformer. \mathcal{A}_2 is the model removing global self-attention, which acts like SwinIR [24]. \mathcal{A}_3 combines local and global attention as in TNT [9]. \mathcal{A}_4 is our whole network. Table 6 shows that performances of \mathcal{A}_1 and \mathcal{A}_2 both decrease. But \mathcal{A}_1 possesses lower performance, illustrating that local self-attention plays a more important role. The performance of \mathcal{A}_4 is higher than that of \mathcal{A}_3 , which verifies that our local-to-global hierarchical self-attention mechanism is more effective than the bypass architecture in TNT. The reason is that our COMO-ViT better expresses images via pixel-level self-attention, the dark areas can be better inferred by using far informative bright areas. The visual results are in Fig. 7.

Effect of Window Size. In COMO-ViT, the input image is split into $w \times w$ windows, the default w is 16. Because the value of w also influences the embedding dimension of our global self-attention, here we test the effect of other values of w on the model performance. Table 7 and Fig. 8 both

Win size (w)	4		16		32	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
LOL-v1	24.46	0.836	24.53	0.842	24.37	0.830
LOL-v2-real	22.14	0.857	22.20	0.863	22.23	0.851
LOL-v2-syn	25.44	0.931	25.58	0.940	25.39	0.935

Table 7. Quantitative results of different window sizes.




Figure 8. Visual comparison results for different window sizes.

show that w just has slight influence on the performance, we thus set $w = 16$ for its efficiency (0.735M parameters for $w = 4$; 0.326M for $w = 16$; 1.016M for $w = 32$).

Limitations. Although IAGC achieves better performance than previous methods, there are still some open problems that IAGC does not solve. There exists slight local color deviation when encountering very challenging low-light cases, e.g., the color gaps at the upper right corner in Fig. 7. This can be understood as contrast and hue are seriously damaged and accurate contextual relationships are not available to recover complete image color. This is a limitation of all existing methods, including IAGC.

5. Conclusion

In this paper, we propose a novel Illumination-Aware Gamma Correction (IAGC) network for low-light image enhancement. In particular, we combine the effectiveness of gamma correction with the strong modelling capacities of networks and introduce two modules to predict the global and local gamma correction factors via perceiving the global (GGCM) and local (LGCM) illumination. Moreover, in order to solve the large-scale dark areas in low-light images, we design a novel Transformer block (COMO-ViT) to completely model the dependencies of all pixels across images through a local-to-global self-attention mechanism. The local self-attention explores pixel-level dependencies based on local windows, while the global self-attention further expands pixel dependencies to the whole image. Compared with conventional Transformers which either only focus on local information but discard global dependencies or emphasize global relationships but neglect local information, our IAGC recovers better illumination and preserves more image details. Extensive experiments have demonstrated the effectiveness of the proposed IAGC architecture.

Acknowledgements This work was supported by Sichuan Science and Technology Program of China under grant No.2023NSFSC0462.

References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization, 2016. arXiv:1607.06450.
- [2] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun. Learning to see in the dark. In *CVPR*, 2018.
- [3] Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *CVPR*, 2021.
- [4] Xingbo Dong, Wanyan Xu, Zhihui Miao, Lan Ma, Chao Zhang, Jiewen Yang, Zhe Jin, Andrew Beng Jin Teoh, and Jiajun Shen. Abandoning the bayer-filter to see in the dark. In *CVPR*, 2022.
- [5] Xueyang Fu, Delu Zeng, Yue Huang, Yinghao Liao, Xinghao Ding, and John Paisley. A fusion-based enhancing method for weakly illuminated images. *SP*, 129:82–96, 2016.
- [6] Xueyang Fu, Delu Zeng, Yue Huang, Xiao-Ping Zhang, and Xinghao Ding. A weighted variational model for simultaneous reflectance and illumination estimation. In *CVPR*, 2016.
- [7] Chunle Guo, Chongyi Li, Jichang Guo, Chen Change Loy, Junhui Hou, Sam Kwong, and Runmin Cong. Zero-reference deep curve estimation for low-light image enhancement. In *CVPR*, 2020.
- [8] Xiaojie Guo, Yu Li, and Haibin Ling. Lime: Low-light image enhancement via illumination map estimation. *IEEE TIP*, 26(2):982–993, 2016.
- [9] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. In *NeurIPS*, 2021.
- [10] Shijie Hao, Xu Han, Yanrong Guo, Xin Xu, and Meng Wang. Low-light image enhancement with semi-decoupled decomposition. *IEEE TMM*, 22(12):3025–3038, 2020.
- [11] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018.
- [12] Yukun Huang, Zheng-Jun Zha, Xueyang Fu, Richang Hong, and Liang Li. Real-world person re-identification via degradation invariance learning. In *CVPR*, 2020.
- [13] Yifan Jiang, Xinyu Gong, Ding Liu, Yu Cheng, Chen Fang, Xiaohui Shen, Jianchao Yang, Pan Zhou, and Zhangyang Wang. Enlighten: Deep light enhancement without paired supervision. *IEEE TIP*, 30:2340–249, 2021.
- [14] Daniel J Jobson, Zia-ur Rahman, and Glenn A Woodell. Properties and performance of a center/surround retinex. *IEEE TIP*, 6(3):451–462, 1997.
- [15] Hanul Kim, Su-Min Choi, Chang-Su Kim, and Yeong Jun Koh. Representative color transform for image enhancement. In *ICCV*, 2021.
- [16] Ron Kimmel, Michael Elad, Doron Shaked, Renato Keshet, and Irwin Sobel. A variational framework for retinex. *IJCV*, 52(1):7–23, 2003.
- [17] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [18] Jaihyun Koh, Jangho Lee, and Sungroh Yoon. Bnucd: A two-branched deep neural network for restoring images from under-display cameras. In *CVPR*, 2022.
- [19] Xiaozhou Lei, Zixiang Fei, Wenju Zhou, Huiyu Zhou, and Minrui Fei. Low-light image enhancement using the cell vibration model. *IEEE TMM*, pages 1–1, 2022.
- [20] Chongyi Li, Chunle Guo, Ling-Hao Han, Jun Jiang, Ming-Ming Cheng, Jinwei Gu, and Chen Change Loy. Low-light image and video enhancement using deep learning: a survey. *IEEE TPAMI*, pages 1–1, 2021.
- [21] Chengxi Li, Xiangyu Qu, Abhiram Gnanasambandam, Omar A. Elgendy, Jiaju Ma, and Stanley H. Chan. Photon-limited object detection using non-local feature matching and knowledge distillation. In *ICCVW*, 2021.
- [22] Jiaqian Li, Juncheng Li, Faming Fang, Fang Li, and Guixu Zhang. Luminance-aware pyramid network for low-light image enhancement. *IEEE TMM*, 23:3153–3165, 2020.
- [23] Mading Li, Jiaying Liu, Wenhan Yang, Xiaoyan Sun, and Zongming Guo. Structure-revealing low-light image enhancement via robust retinex model. *IEEE TIP*, 27(6):2828–2841, 2018.
- [24] Jingyun Liang, Jie Zhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *ICCVW*, 2021.
- [25] Jiaying Liu, Dejjia Xu, Wenhan Yang, Minhao Fan, and Haofeng Huang. Benchmarking low-light image enhancement and beyond. *IJCV*, 129(4):1153–1184, 2021.
- [26] Risheng Liu, Long Ma, Jiao Zhang, Xin Fan, and Zhongxuan Luo. Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. In *CVPR*, 2021.
- [27] Kin Gwn Lore, Adedotun Akintayo, and Soumik Sarkar. L1-net: A deep autoencoder approach to natural low-light image enhancement. *PR*, 61:650–662, 2017.
- [28] Long Ma, Tengyu Ma, Risheng Liu, Xin Fan, and Zhongxuan Luo. Toward fast, flexible, and robust low-light image enhancement. In *CVPR*, 2022.
- [29] Sean Moran, Pierre Marza, Steven McDonagh, Sarah Parisot, and Gregory Slabaugh. Deeplpf: Deep local parametric filters for image enhancement. In *CVPR*, 2020.
- [30] Muhammad Muzammal Naseer, Kanchana Ranasinghe, Salman H Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing properties of vision transformers. *NeurIPS*, 34:23296–23308, 2021.
- [31] Michael K Ng and Wei Wang. A total variation model for retinex. *SIAM JIS*, 4(1):345–365, 2011.
- [32] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, and et al. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019.
- [33] Stephen M Pizer. Contrast-limited adaptive histogram equalization: Speed and effectiveness. In *VBC*, 1990.
- [34] Xutong Ren, Wenhan Yang, Wen-Huang Cheng, and Jiaying Liu. Lr3m: Robust low-light enhancement via low-rank regularized retinex model. *IEEE TIP*, 29:5862–5876, 2020.
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017.

- [36] Shuhang Wang, Jin Zheng, Hai-Miao Hu, and Bo Li. Naturalness preserved enhancement algorithm for non-uniform illumination images. *IEEE TIP*, 22(9):3538–3548, 2013.
- [37] Yinglong Wang, Chao Ma, and Jianzhuang Liu. Smartasign: Learning a smart knowledge assignment strategy for deraining and desnowing. In *CVPR*, 2023.
- [38] Zhou Wang, A.C. Bovik, H.R. Sheikh, and E.P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE TIP*, 13(4):600–612, 2004.
- [39] Zhendong Wang, Xiaodong Cun, Jianmin Bao, and Jianzhuang Liu. Uformer: A general u-shaped transformer for image restoration. In *CVPR*, 2022.
- [40] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *BMVC*, 2018.
- [41] Wenhui Wu, Jian Weng, Pingping Zhang, Xu Wang, Wenhan Yang, and Jianmin Jiang. Uretinex-net: Retinex-based deep unfolding network for low-light image enhancement. In *CVPR*, 2022.
- [42] Ke Xu, Xin Yang, Baocai Yin, and Rynson WH Lau. Learning to restore low-light images via decomposition-and-enhancement. In *CVPR*, 2020.
- [43] Xiaogang Xu, Ruixing Wang, Chi-Wing Fu, and Jiaya Jia. Snr-aware low-light image enhancement. In *CVPR*, 2022.
- [44] Xin Xu, Shiqin Wang, Zheng Wang, Xiaolong Zhang, and Ruimin Hu. Exploring image enhancement for salient object detection in low light images. *ACM TMCCA*, 17(1s):1–19, 2021.
- [45] Canqian Yang, Meiguang Jin, Xu Jia, Yi Xu, and Ying Chen. Adaint: Learning adaptive intervals for 3d lookup tables on real-time image enhancement. In *CVPR*, 2022.
- [46] Wenhan Yang, Shiqi Wang, Yuming Fang, Yue Wang, and Jiaying Liu. Band representation-based semi-supervised low-light image enhancement: bridging the gap between signal fidelity and perceptual quality. *IEEE TIP*, 30:3461–3473, 2021.
- [47] Wenhan Yang, Wenjing Wang, Haofeng Huang, Shiqi Wang, and Jiaying Liu. Sparse gradient regularized deep retinex network for robust low-light image enhancement. *IEEE TIP*, 30:2072–2086, 2021.
- [48] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhancement. In *ECCV*, 2020.
- [49] Yu Zhang, Xiaoguang Di, Bin Zhang, and Chunhui Wang. Self-supervised image enhancement network: Training with low light images only, 2020. arXiv:2002.11300.
- [50] Yonghua Zhang, Xiaojie Guo, Jiayi Ma, Wei Liu, and Jiawan Zhang. Beyond brightening low-light images. *IJCV*, 129(4):1013–1037, 2021.
- [51] Yonghua Zhang, Jiawan Zhang, and Xiaojie Guo. Kindling the darkness: A practical low-light image enhancer. In *ACM MM*, 2019.
- [52] Zhaoyang Zhang, Yitong Jiang, Jun Jiang, Xiaogang Wang, Ping Luo, and Jinwei Gu. Star: A structure-aware lightweight transformer for real-time image enhancement. In *ICCV*, 2021.
- [53] Zhao Zhang, Huan Zheng, Richang Hong, Mingliang Xu, Shuicheng Yan, and Meng Wang. Deep color consistent network for low-light image enhancement. In *CVPR*, 2022.