

A Frequency-Aware Self-Supervised Learning for Ultra-Wide-Field Image Enhancement

Weicheng Liao^{1,2†}, Zan Chen^{1†}, Jianyang Xie³, Yalin Zheng³, Yuhui Ma^{2(✉)},
and Yitian Zhao^{2,4(✉)}

¹ College of Information Engineering, Zhejiang University of Technology, Hangzhou, China

² Ningbo Institute of Materials Technology and Engineering, Chinese Academy of Sciences, Ningbo, China
{mayuhui, yitian.zhao}@nimte.ac.cn

³ Department of Eye and Vision Science, University of Liverpool, Liverpool, UK

⁴ School of Biomedical Engineering, ShanghaiTech University, Shanghai, China

Abstract. Ultra-Wide-Field (UWF) retinal imaging has revolutionized retinal diagnostics by providing a comprehensive view of the retina. However, it often suffers from quality-degrading factors such as blurring and uneven illumination, which obscure fine details and mask pathological information. While numerous retinal image enhancement methods have been proposed for other fundus imageries, they often fail to address the unique requirements in UWF, particularly the need to preserve pathological details. In this paper, we propose a novel frequency-aware self-supervised learning method for UWF image enhancement. It incorporates frequency-decoupled image deblurring and Retinex-guided illumination compensation modules. An asymmetric channel integration operation is introduced in the former module, so as to combine global and local views by leveraging high- and low-frequency information, ensuring the preservation of fine and broader structural details. In addition, a color preservation unit is proposed in the latter Retinex-based module, to provide multi-scale spatial and frequency information, enabling accurate illumination estimation and correction. Experimental results demonstrate that the proposed work not only enhances visualization quality but also improves disease diagnosis performance by restoring and correcting fine local details and uneven intensity. To the best of our knowledge, this work is the first attempt for UWF image enhancement, offering a robust and clinically valuable tool for improving retinal disease management.

Keywords: Ultra-wide-field · deblur · enhancement · frequency-aware.

1 Introduction

Ultra-Wide-Field (UWF) retinal imaging represents a significant advancement in retinal diagnostics, offering a broader and more detailed view of the retina

[†] These authors contributed equally.




Fig. 1: Illustrations of UWF imaging. (a-b) A color fundus image and an UWF image acquired from the same eye. (c) Sample low-quality factors observed in UWF image: vasculature blurring and uneven illumination mask the lesion.

compared with conventional imaging approaches. Fig. 1 (a-b) illustrates this by comparing a retinal color fundus image with a UWF image, highlighting the enhanced ability of UWF imaging to capture a more comprehensive retina, improving the diagnosis and management of various retinal diseases.

In recent years, numerous image enhancement methods have been proposed [4, 19–21] to improve image readability. However, the majority of these methods are primarily designed for natural images and fail to adequately address the specific requirements of medical imaging, particularly retinal imaging, where preserving micro pathological details is critical. Numerous retinal image enhancement methods have been introduced that learn general priors from extensive datasets comprising low- and high-quality images [3, 6]. However, acquiring such paired datasets poses significant challenges in medical imaging. Consequently, researchers [14, 11, 15, 8] have synthesized low-quality images by degrading high-quality ones, thereby constructing paired datasets for training. Nevertheless, these approaches rely heavily on the availability of extensive collections of high-quality images free from quality-degrading factors.

The UWF enhancement is more complicated by several low-quality factors, with blurring and uneven illumination being the most significant. Blurring, as shown in Fig. 1 (c), often obscures fine local details, reducing the diagnostic utility of the images. Similarly, uneven illumination, depicted in Fig. 1 (c), particularly in the peripheral retinal regions, the poor contrast may mask critical pathological information. The aforementioned issues collectively hinder the identification of important structural features, making it more difficult for clinicians to accurately and confidently interpret the images. UWF images are acquired through red and green laser scanning, which means they do not follow the "gray world assumption" and are thus prone to color distortion during light compensation [18]. In addition to blur and uneven illumination, the preservation of pathological details, peripheral distortions obtained by unsupervised approach, and the lack of paired datasets for training. Overcoming these challenges is essential to ensure that enhanced UWF images provide reliable and clinically useful information for diagnosing and managing retinal diseases.

In this work, we introduce a frequency-aware self-supervised learning method, which is specifically designed for UWF image enhancement. The method consists of a FREquency-Decoupled deblurring (FRED) module and a Retinex-guided Illumination CompEnsation (RICE) module. In the FRED module, an asymmetric channel integration operation is introduced, to effectively combine global and local views by leveraging separated high- and low-frequency information. This ensures the preservation of both fine and broader structural features. In the RICE module, a color preservation unit is proposed to provide multi-scale spatial and frequency information, enabling accurate illumination estimation and correction. Together, these innovations allow the model to restore fine local details and correct intensity inhomogeneities while preserving structural integrity.

The main contributions of this paper are summarized as follows: **1)** To the best of our knowledge, this work represents the first dedicated effort for deblurring and illumination compensation in UWF image enhancement. The experimental findings demonstrate that the work not only enhances the visual quality but also improves the performance of the subsequent disease diagnosis. **2)** The asymmetric channel integration unit effectively combines global and local views by leveraging separated high- and low-frequency information, to ensure the preservation of both fine details and broader structural features. **3)** The color preservation unit is able to provide multi-scale spatial and frequency information, enabling accurate illumination estimation and correction.

2 Proposed Method

Fig. 2 (a) illustrates the overview of the proposed framework, where the image is processed in two stages. In the first stage, the FRED module is proposed to perform deblurring on UWF images, which aims to align the image domain with clinical acquisition conditions of UWF fundus cameras for clearer details. In the second stage, the RICE module is proposed for further illumination compensation based on the deblurring results, which aims at better quality of UWF images in terms of clarity and uniform intensity.

2.1 Frequency-decoupled image deblurring (FRED)

In the deblurring stage, a random blurring degradation operation [14] is firstly applied to high-quality images with clear structures to generate paired blur/clear samples. These paired samples are then used to train the model, enabling it to learn an effective restoration mapping and effectively deblur UWF images while preserving fine structural details. In order to prevent mutual interference between varying frequencies, we use average pooling separation (APS) [22] to decompose the UWF image into high- and low-frequency features, which are processed separately through dual frequency streams with encoder-decoder structure. Specifically, the generated blurry image is firstly fed into the APS to obtain the low-frequency features via downsampling by an average pooling layer and upsampling

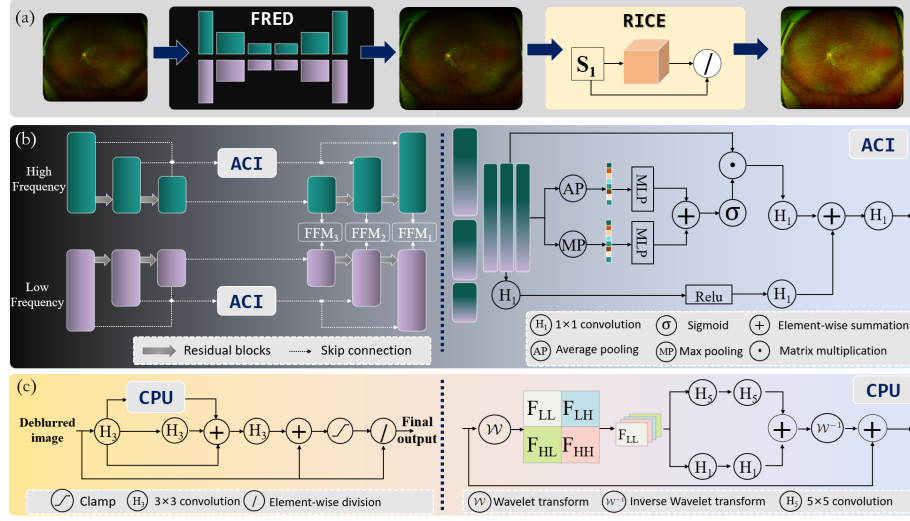


Fig. 2: Overview of the proposed framework and its main components. (a) Overall architecture: the low-quality images are sequentially passed through the deblurring and illumination compensation modules. (b-c) Architectures of FRED and RICE.

by bilinear interpolation successively. Then the high-frequency features are obtained by subtraction of the low-frequency features from the input blurry image. Next, both high- and low-frequency features are downsampled to multiple scales and then fed into different encoder layers of dual frequency streams respectively, which have the same encoder-decoder structure. In order to effectively aggregate global and local information in both high- and low-frequency features, an Asymmetric Channel Integration (ACI) unit is proposed for cross-scale latent feature fusion in the skip connections of each frequency stream, as shown in Fig. 2 (b).

Asymmetric Channel Integration (ACI) The proposed ACI unit aims to achieve more comprehensive fusion of global and local features from different encoder layers and further enhance feature representations in the decoder. Specifically, the features from different encoder layers are firstly rescaled to corresponding decoder feature size and concatenated to obtain the fusion feature $f_s \in \mathbb{R}^{H \times W \times C}$. Then f_s is input into global information extraction branch and local information extraction branch, respectively. Global information extraction branch adopts channel-wise attention learning to refine f_s . In this branch, max and average pooling operations followed by a Multi Layer Perceptron (MLP) are separately applied to f_s to produce two feature vectors $f_s^m, f_s^a \in \mathbb{R}^C$. Then the two feature vectors are fused via summation and Sigmoid to generate the final attention weight $W_s^g \in \mathbb{R}^C$, which is multiplied with f_s in a channel-wise manner followed by a 1×1 convolutional layer to obtain the global information f_g . For local information extraction branch, two 1×1 convolutional layers with

ReLU is applied to f_s to obtain the local information f_l . Finally, a 1×1 convolutional layer is applied to summation of f_g and f_l to calculate the output feature f_{ACI} of the ACI unit, which will be integrated into the corresponding decoder layer. The whole process can be formulated as:

$$\begin{aligned} f_s^a &= M(AP(f_s)), f_s^m = M(MP(f_s)), \\ W_s^g &= \sigma(f_s^a + f_s^m), f_g = H_1(W_s^g \times f_s), \\ f_l &= H_1(\delta(H_1(f_s))), f_{ACI} = H_1(f_g + f_l), \end{aligned} \quad (1)$$

where M , AP , MP , δ , σ and H_1 represent MLP, average pooling, max pooling, ReLU, Sigmoid and 1×1 convolution, respectively.

The FRED module further integrates features of each decoder layer separately in high- and low-frequency streams through Frequency Fusion module (FFM) [22] to output multi-scale deblurring results, as illustrated in Fig. 2 (b). Each FFM is composed of convolutional blocks and channel attention mechanisms, which facilitates the integration of high- and low-frequency information. The output of FFM on the last decoder layer is the final deblurring result. The FRED module adopted the deep supervision strategy for training and is optimized through multi-scale content loss L_{cont} , multi-scale frequency reconstruction loss L_{MSFR} [2] and perceptual loss L_{per} [24], which can be formulated as:

$$L_{deblur} = L_{cont} + \beta L_{MSFR} + \gamma L_{per} \quad (2)$$

where β and γ are weighting factors of multi-scale frequency reconstruction loss and perceptual loss, respectively.

2.2 Retinex-based illumination compensation (RICE)

Inspired by SCI [10], we propose a Retinex-based module called RICE for further illumination compensation of the deblurred UWF images, as shown in Fig. 2 (c). According to the Retinex theory [10], a given image \mathbf{I} can be decomposed into two components, respectively the reflectance \mathbf{R} and illumination \mathbf{L} : $\mathbf{I} = \mathbf{R} \otimes \mathbf{L}$. Similarly, the enhanced image \mathbf{I}' can also be expressed as $\mathbf{I}' = \mathbf{R}' \otimes \mathbf{L}'$. For the enhancement task that requires only illumination compensation, we can assume that the reflectance should remain constant after enhancement, i.e., $\mathbf{R}' = \mathbf{R}$. Therefore, the enhanced image \mathbf{I}' can be further formulated as $\mathbf{I}' = \mathbf{R}' \otimes \mathbf{L} \otimes (\mathbf{L}' \oslash \mathbf{L}) = \mathbf{R} \otimes \mathbf{L} \otimes (\mathbf{L} \oslash \mathbf{L}') = \mathbf{I} \oslash \mathbf{r}$, where $\mathbf{r} = \mathbf{L} \oslash \mathbf{L}'$ represents the ratio of illumination compensation. The proposed RICE aims to precisely estimate the ratio \mathbf{r} to achieve the effective illumination compensation.

However, due to the optical nature of ultra-wide angle fundus imaging, which is likely to raise color distortion for Retinex-based methods during illumination compensation [18]. To address this issue, we propose a Color Preservation Unit (CPU) that incorporates a two-dimensional discrete wavelet transform to ensure color fidelity, as shown in Fig. 2 (c). This unit utilizes the multi-scale frequency separation capability of the wavelet transform to decouple chromaticity and luminance, leveraging the low-frequency sub-bands to focus on light estimation and suppressing the interference of high-frequency chromaticity information.

Color Preservation Unit (CPU) The discrete wavelet transform \mathcal{W} involved in the proposed CPU utilizes four filters defined as:

$$f^{LL} = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, f^{LH} = \begin{pmatrix} -1 & -1 \\ 1 & 1 \end{pmatrix}, f^{HL} = \begin{pmatrix} -1 & 1 \\ -1 & 1 \end{pmatrix}, f^{HH} = \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}. \quad (3)$$

Through the discrete wavelet transform \mathcal{W} , the feature $F \in \mathbb{R}^{H \times W \times C}$ of the deblurred UWF image extracted by a 3×3 convolutional layer is decomposed into four frequency sub-bands: $F^{LL}, F^{LH}, F^{HL}, F^{HH} \in \mathbb{R}^{\frac{H}{2} \times \frac{W}{2} \times C}$, where F^{LL} mainly contains the global features and approximate contours of the image, F^{LH}, F^{HL} and F^{HH} capture the details and edge variations of the image. Each frequency sub-band is input to two parallel branches respectively composed of two successive 1×1 and 5×5 convolutional layers, then the output feature maps of the two branches are fused via an element-wise summation. Finally, the discrete wavelet inverse transform \mathcal{W}^{-1} is applied to these processed frequency sub-bands, then a shortcut connection of the feature F is introduced to produce the output of CPU. Let set $\tilde{F} = \{F^{LL}, F^{LH}, F^{HL}, F^{HH}\}$, and the whole process can be defined as:

$$\begin{aligned} \tilde{F} &= \mathcal{W}(F), \\ F_w &= \{H_5^2(F^*) + H_1^2(F^*) \mid F^* \in \tilde{F}\}, \\ F_{wt} &= \mathcal{W}^{-1}(F_w) + F, \end{aligned} \quad (4)$$

where H_1^2 and H_5^2 denote two successive 1×1 and 5×5 convolutional layers, respectively, F_w and F_{wt} denote the set of the processed frequency sub-bands and the output of CPU, respectively. For effective optimization of the RICE module, we adopt a joint loss function containing fidelity loss L_f , smoothness loss L_s [10] and exposure control loss L_{exp} [5], which can be formulated as:

$$L_{ic} = \alpha L_f + L_s + L_{exp}, \quad (5)$$

where α is the parameter to control the weight of fidelity loss.

3 Experimental Results

Dataset description In this study, we first construct a UWF image enhancement dataset, which involves both healthy controls and patients with diabetic retinopathy (DR). The dataset contains 834 UWF images captured with the Optos 200Tx device, with a resolution of 3900×3072 pixels and a 200° field of view. In our experiments, 400 images with varying exposure levels were selected for training, while 434 images with different levels of blur and illumination were used to evaluate the performance of the image enhancement method. All image acquisitions adhered to the relevant regulatory approvals and patient consent, ensuring compliance with ethical standards and privacy protection protocols.

Experimental settings Our method was implemented in Python with PyTorch library. The experiments were carried out on two NVIDIA GPUs (GeForce RTX 4090, 24 GB). All training images were randomly cropped to a size of 256×256 . Adam optimization was adopted with the initial learning rate of 0.0001 and 0.0003 for training the FRED and RICE modules, respectively. The weighted parameters in the loss function were set as: $\beta = 0.1$, $\gamma = 0.01$, $\alpha = 1.5$.




Fig. 3: The visualization results by using different methods.

3.1 Image enhancement performances

We first present the performance evaluation of the proposed method. Six state-of-the-art methods are used for comparisons, and the NIQE [13], BRISQUE [12] and PIQE [17] are used as evaluation metrics. Fig. 3 demonstrates the enhancement results of different methods. It can be observed that several compared methods produce undesirable effects. For example, CLAHE [25] amplifies background noise and thus lead to image quality degradation, while both Zero-DCE++ [7] and URRN-Net [1] generate noticeable light purple artifacts in the peripheral regions, as shown in the bottom-right patch. Deblurring methods such as MRDNet [23] can effectively enhance structural information, but hardly deal with illumination non-uniformity. RUAS [9] improves overall brightness while maintaining color consistency, but suffers from overexposure in the optic nerve region, leading to an unnatural appearance. Although ZeroIG [16] effectively enhances overall brightness, it might cause some blurring in fine details, such as capillaries and leopard-like patterns, which can affect the clarity of these regions. In contrast, our method outperforms the others by preserving the continuity of microvessels and maintaining natural color tones. Table 1 further demonstrate the superiority of our proposed method in quantitative analysis. Our method yields the lowest NIQE, BRISQUE and PIQE scores with large margin when compared with the competing methods.

3.2 Ablation study

In order to investigate the contributions of components in our method, we conducted ablation studies, as illustrated in the bottom four rows of Table 1.

Effectiveness of two stages (FRED and RICE) The results demonstrate a 21.3% increase in PIQE after removing the FRED module, indicating that intrinsic blur in UWF images has a detrimental effect on illumination compensation, thereby degrading overall image quality. This emphasises the critical role of FRED in eliminating blur, which are essential for image quality enhancement. Furthermore, the absence of the RICE module resulted in a 28% increase

Table 1: Image enhancement performance in terms of *quality assessment* (the first three metrics) and *disease grading* tasks (the latter two metrics).

| Method | NIQE ↓ | BRISQUE ↓ | PIQE ↓ | ACC(%) ↑ | F1(%) ↑ |
|----------------|-------------------|--------------------|-------------------|--------------|--------------|
| Original | 5.87 ±0.91 | 35.25 ±8.87 | 16.32 ±7.49 | 57.13 | 48.24 |
| CLAHE [25] | 4.11 ±0.90 | 22.69 ±8.59 | 10.41 ±5.13 | 61.21 | 53.06 |
| MRDNet [23] | 4.86 ±0.61 | 29.45 ±10.80 | 15.85 ±7.18 | 62.74 | 53.16 |
| Zero-DCE++ [7] | 6.80 ±0.53 | 41.12 ±15.12 | 11.75 ±7.52 | 47.81 | 25.86 |
| URRN-Net [1] | 7.03 ±0.89 | 43.19 ±16.68 | 18.84 ±5.88 | 46.62 | 26.91 |
| RUAS [9] | 5.82 ±1.89 | 24.51 ±15.4 | 9.21 ±6.43 | 39.95 | 33.62 |
| ZeroIG [16] | 4.99 ±1.39 | 24.17 ±10.48 | 8.67 ±6.82 | 47.03 | 33.89 |
| Ours | 4.09 ±0.51 | 19.15 ±6.02 | 8.26 ±4.22 | 69.49 | 64.07 |
| w/o FRED | 4.39 ±1.06 | 23.09 ±8.02 | 10.02 ±5.60 | 66.86 | 59.47 |
| w/o RICE | 4.86 ±0.55 | 24.51 ±9.65 | 14.59 ±7.06 | 62.36 | 51.96 |
| w/o ACI | 4.19 ±0.91 | 20.72 ±9.10 | 8.39 ±4.15 | 67.39 | 61.48 |
| w/o CPU | 4.31 ±0.56 | 20.67 ±7.46 | 8.49 ±4.51 | 68.71 | 62.93 |

in BRISQUE, confirming the vital role of RICE in ensuring accurate illumination compensation. The absence of RICE led to a failure to balance the overall lighting effect, also leading to a significant degradation in image quality.

Effectiveness of two units (ACI and CPU) We further validate the effectiveness of ACI in the FRED stage. Disabling ACI resulted in an 8.2% increase in BRISQUE, indicating the effectiveness of its multi-scale feature fusion within different frequency bands on UWF image deblurring. This might be because the framework without ACI only relies on skip connections for feature fusion, which is less effective in capturing global and local information. Similarly, disabling the CPU in the RICE stage led to a 5.4% increase in NIQE, confirming its significance in maintaining color consistency. The absence of CPU can not effectively extract frequency sub-bands for accurate illumination estimation, thus underperforming in preserving color fidelity during illumination compensation.

3.3 Effectiveness of enhancement on disease diagnosis

To further validate the effectiveness of image enhancement on disease diagnosis, we first collected an additional 902 UWF images (444 normal, 195 mild, 103 moderate, 79 severe non-proliferative DR and 81 proliferative DR), which were split into training and test sets at a 7:3 ratio. Then we trained a ResNet34 model for DR grading using the training set, and employed the well-trained model in the test images enhanced by different methods for comparisons. The right two columns of Table 1 report the DR grading performance on images enhanced by different methods. Our method improves ACC and F1-score by 12.3% and 15.8% respectively when compared with the original images, and also outperforms other enhancement methods. This significant improvement is primarily attributed to the FRED and RICE modules, which enhance structural clarity and illumination uniformity, respectively. The findings demonstrate the advantages of our method on improving the visibility of peripheral retinal pathology, thereby facilitating disease diagnosis.

4 Conclusions

In this study, we propose a novel frequency-aware self-supervised learning framework for UWF image enhancement. The primary advantage of our framework is that it can learn to achieve image deblurring and illumination compensation simultaneously via two specially designed modules. First, the FRED module leverages asymmetric channel integration to preserve structural details by combining global and local information in high- and low-frequency information. Second, the RICE module employs a color preservation unit based on frequency decoupling by discrete wavelet transform to mitigate color distortions while achieving illumination uniformity. Experimental results demonstrate that our method outperforms other approaches, indicating the potential of the framework in improving image quality for higher accuracy of clinical diagnosis.

Acknowledgments. This study was funded by the National Natural Science Foundation of China (62402478, 62422122, 62272444) and China Postdoctoral Science Foundation (2023M743629)

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Chen, X., Yu, Y.: An unsupervised low-light image enhancement method for improving v-slam localization in uneven low-light construction sites. *Automation in Construction* **162**, 105404 (2024)
2. Cho, S.J., Ji, S.W., Hong, J.P., Jung, S.W., Ko, S.J.: Rethinking coarse-to-fine approach in single image deblurring. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 4641–4650 (October 2021)
3. Deng, Z., Cai, Y., Chen, L., Gong, Z., Bao, Q., Yao, X., Fang, D., Yang, W., Zhang, S., Ma, L.: Rformer: Transformer-based generative adversarial network for real fundus image restoration on a new clinical benchmark. *IEEE Journal of Biomedical and Health Informatics* **26**(9), 4645–4655 (2022)
4. Dong, J., Pan, J., Yang, Z., Tang, J.: Multi-scale residual low-pass filter network for image deblurring. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 12345–12354 (2023)
5. Guo, C., Li, C., Guo, J., Loy, C.C., Hou, J., Kwong, S., Cong, R.: Zero-reference deep curve estimation for low-light image enhancement. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2020)
6. Lee, K.G., Song, S.J., Lee, S., Yu, H.G., Kim, D.I., Lee, K.M.: A deep learning-based framework for retinal fundus image enhancement. *PLoS One* **18**(3), e0282416 (2023)
7. Li, C., Guo, C., Loy, C.C.: Learning to enhance low-light image via zero-reference deep curve estimation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **44**(8), 4225–4238 (2022)

8. Liu, H., Li, H., Fu, H., Xiao, R., Gao, Y., Hu, Y., Liu, J.: Degradation-invariant enhancement of fundus images via pyramid constraint network. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 507–516. Springer (2022)
9. Liu, R., Ma, L., Zhang, J., Fan, X., Luo, Z.: Retinex-inspired unrolling with cooperative prior architecture search for low-light image enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10561–10570 (June 2021)
10. Ma, L., Ma, T., Liu, R., Fan, X., Luo, Z.: Toward fast, flexible, and robust low-light image enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5637–5646 (June 2022)
11. Ma, Y., Liu, Y., Cheng, J., Zheng, Y., Ghahremani, M., Chen, H., Liu, J., Zhao, Y.: Cycle structure and illumination constrained gan for medical image enhancement. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2020: 23rd International Conference, Lima, Peru, October 4–8, 2020, Proceedings, Part II 23. pp. 667–677. Springer (2020)
12. Mittal, A., Moorthy, A.K., Bovik, A.C.: No-reference image quality assessment in the spatial domain. *IEEE Transactions on image processing* **21**(12), 4695–4708 (2012)
13. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters* **20**(3), 209–212 (2012)
14. Pérez, A.D., Perdomo, O., Rios, H., Rodríguez, F., González, F.A.: A conditional generative adversarial network-based method for eye fundus image quality enhancement. In: International Workshop on Ophthalmic Medical Image Analysis. pp. 185–194. Springer (2020)
15. Shen, Z., Fu, H., Shen, J., Shao, L.: Modeling and enhancing low-quality retinal fundus images. *IEEE Transactions on Medical Imaging* **40**(3), 996–1006 (2021)
16. Shi, Y., Liu, D., Zhang, L., Tian, Y., Xia, X., Fu, X.: Zero-ig: Zero-shot illumination-guided joint denoising and adaptive enhancement for low-light images. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3015–3024 (2024)
17. Venkatanath, N., Praneeth, D., Bh, M.C., Channappayya, S.S., Medasani, S.S.: Blind image quality evaluation using perception based features. In: 2015 twenty first national conference on communications (NCC). pp. 1–6. IEEE (2015)
18. Wang, W., Li, B., Zheng, J., Xian, S., Wang, J.: A fast multi-scale retinex algorithm for color image enhancement. In: 2008 International Conference on Wavelet Analysis and Pattern Recognition. vol. 1, pp. 80–85 (2008)
19. Wang, W., Yang, H., Fu, J., Liu, J.: Zero-reference low-light enhancement via physical quadruple priors. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 26057–26066 (2024)
20. Wu, Y., Pan, C., Wang, G., Yang, Y., Wei, J., Li, C., Shen, H.T.: Learning semantic-aware knowledge guidance for low-light image enhancement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1662–1671 (2023)
21. Xu, X., Wang, R., Lu, J.: Low-light image enhancement via structure modeling and guidance. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 9893–9903 (2023)
22. Yang, K., Hu, T., Dai, K., Chen, G., Cao, Y., Dong, W., Wu, P., Zhang, Y., Yan, Q.: Crnet: A detail-preserving network for unified image restoration and enhancement task. *arXiv preprint arXiv:2404.14132* (2024)

23. Zhang, B., Sun, J., Sun, F., Wang, F., Zhu, B.: Image deblurring method based on self-attention and residual wavelet transform. *Expert Systems with Applications* **244**, 123005 (2024)
24. Zhang, K., Ren, W., Luo, W., Lai, W.S., Stenger, B., Yang, M.H., Li, H.: Deep image deblurring: A survey. *International Journal of Computer Vision* **130**(9), 2103–2130 (2022)
25. Zuiderveld, K.: Contrast limited adaptive histogram equalization. In: *Graphics gems IV*. pp. 474–485. Academic Press Professional, Inc. (1994)