Future Technology
Open Access Journal
ISSN 2832-0379

February 2027| Volume 06 | Issue 01 | Pages 01-13

Journal homepage: https://fupubco.com/futech

https://doi.org/10.55670/fpll.futech.6.1.1

Future Publishing LLC

Article

Perception-guided Retinex-inspired framework for low-light facial image enhancement and robust face analysis

Mohammed S. H. Al-Tamimi[1]*, Roaa S. Abed Alah[2], Sarah Riyadh Adnan[3]

[1]Computer Science Department, College of Science, University of Baghdad, Baghdad, Iraq
[2]Department of Intelligent Medical Systems, University of Information Technology and Communications, Baghdad, Iraq
[3]High Institute for Infertility Diagnosis Assisted Reproductive Techniques, Al-Nahrain University, Baghdad, Iraq

ARTICLE INFO

Article history:
Received 01 May 2026
Received in revised form
12 October 2026
Accepted 02 November 2026

Keywords:
Low-light image enhancement, Face detection, Retinex-inspired enhancement,
Perception-guided enhancement, DARK FACE, EfficientNet-B0

*Corresponding author
Email address:
mohammed.s@sc.uobaghdad.edu.iq

DOI: 10.55670/fpll.futech.6.1.2

ABSTRACT

Low-light facial images often suffer from reduced visibility, noise, color distortion, and loss of structural information, which can degrade subsequent face-analysis performance. This study proposes a perception-guided framework for single-image low-light facial enhancement based on Retinex-inspired illumination–reflectance decomposition and spatially adaptive enhancement. Local contrast and luminance-masking cues are incorporated through a Perception Contrast Optimization Module (PCOM) to regulate enhancement intensity across differently illuminated regions. Illumination correction, perception-guided adaptive fusion, and edge-preserving post-processing are subsequently applied to improve facial visibility while maintaining relevant structural information. Experiments were conducted on a held-out set of 900 DARK FACE images using no-reference image-quality assessment and downstream face-analysis tasks. Enhancement reduced the mean NIQE score from 4.8652 to 4.7828. Using a fixed YOLOv8-Face detector, AP@0.50 increased from 16.80% to 26.22%, while mAP@0.50:0.95 increased from 4.99% to 7.82%. In binary Face/Non-Face classification, ImageNet-pretrained EfficientNet-B0 achieved 95.22% accuracy and 99.08% recall on enhanced samples, compared with 95.69% accuracy and 91.74% recall on the corresponding original samples. The results demonstrate that the proposed enhancement substantially improves downstream face detectability and Face recall under challenging low-light conditions. However, a trade-off is observed in some perceptual-quality and precision measures.

1. Introduction

Face detection and analysis under low-light conditions remain challenging tasks in computer vision, with important applications in intelligent surveillance, mobile photography, human-computer interaction, and automated visual analysis [1]. Although deep-learning-based methods have made substantial progress in facial analysis, their performance still depends heavily on the input image's visual quality [2]. In real-world nighttime environments, insufficient and non-uniform illumination can substantially reduce facial visibility and obscure discriminative structural information. Low-light images may additionally suffer from sensor noise, color distortion, reduced contrast, and loss of fine details, which can negatively affect subsequent face analysis [3]. Consequently, improving the visual quality of low-light facial images before downstream processing represents an important research problem [4]. Traditional enhancement techniques, including histogram equalization and Retinex-based methods, attempt to improve visibility by redistributing image intensity or separating illumination from reflectance [1]. However, under severe and spatially non-uniform illumination, conventional enhancement may amplify noise, introduce halo artifacts, distort colors, or suppress important structural details [5]. More recent learning-based approaches offer greater restoration capability, but balancing illumination correction, structural preservation, perceptual quality, and computational efficiency remains challenging [6,7]. Retinex theory provides a useful representation of a low-light image through illumination and reflectance components. The reflectance representation primarily captures structural and texture information, whereas illumination characterizes the spatial distribution of lighting. These complementary properties have led to the development of specific processing techniques rather than processing both images the same way

[8]. Furthermore, stimuli related to human visual perception from the Human Visual System (HVS), such as local contrast sensitivity and luminance masking, may offer insight into adapting the level of image enhancement across varying spatial light levels. In this study, these HVS concepts are used as perceptual approximations rather than a complete physiological model of human vision. Despite recent advances, several limitations remain in low-light facial image enhancement. Existing approaches may improve global brightness while simultaneously amplifying noise or suppressing local facial details, whereas aggressive local enhancement may introduce artifacts that adversely affect downstream face analysis [9]. Furthermore, improvements in perceptual image quality do not necessarily translate directly into improved face-detection or classification performance. This creates a need for an enhancement strategy that considers both illumination correction and preservation of task-relevant facial information, and that is evaluated using both image-quality measures and downstream face-analysis performance [4].

To address these limitations, this study proposes a perception-guided framework for single-image low-light facial enhancement. The framework uses a Retinex-inspired decomposition to obtain illumination and reflectance representations [10]. The reflectance pathway is associated with local structural and texture preservation, whereas the illumination pathway represents global and spatially varying lighting characteristics [11]. PCOM incorporates local contrast and luminance-masking cues to construct a spatially adaptive enhancement weight. This weight regulates the contribution of the corrected representation during reconstruction, helping to control enhancement intensity across differently illuminated image regions. Each image is processed independently; therefore, the proposed method does not employ temporal modeling, video tracking, identity association, or inter-frame information. The experimental evaluation is conducted on the DARK FACE dataset using full-resolution low-light images [12]. Because the main annotated DARK FACE set does not provide a corresponding normal-light ground-truth image for each low-light image, no-reference image-quality measures are used for perceptual evaluation [13]. The effect of enhancement on downstream facial information is further investigated using YOLOv8-Face detection and ImageNet-pretrained EfficientNet-B0 [14] binary Face/Non-Face classification. In addition, ablation experiments are performed to examine the contribution of the principal enhancement components, and computational runtime is measured on full-resolution images.

2. Specific objectives of this study

- To develop a single-image low-light facial enhancement framework based on Retinex-inspired decomposition and perception-guided adaptive processing.
- To improve illumination while preserving local facial structures and texture information under challenging low-light conditions.
- To incorporate local contrast and luminance-masking cues through PCOM to spatially regulate enhancement intensity.
- To evaluate enhancement quality using no-reference image-quality metrics on the DARK FACE dataset.
- To assess the effect of enhancement on downstream face analysis using YOLOv8-Face detection and EfficientNet-B0 Face/Non-Face classification.
- To investigate the contribution of the principal enhancement components through ablation experiments and evaluate computational runtime on full-resolution images.

3. Problem formulation

Let $I \in \mathbb{R}^{H\times W\times 3}$ denote a low-light *RGB* image, where H and W represent the image height and width, respectively. The objective of the proposed framework is to generate an enhanced image $\hat{I}$ with improved illumination and contrast while preserving color consistency and relevant structural information. The enhancement process is generally formulated as:

$$\hat{I} = F(I) \tag{1}$$

where *F(.)* denotes the proposed perception-guided enhancement process.

Following the Retinex image-formation principle, the observed low-light image is modeled as the interaction between reflectance and illumination components with an additional noise term [10]:

$$I = R \circledast L + N \tag{2}$$

where R denotes the reflectance component containing structural and texture information, L represents the spatially varying illumination component, N accounts for sensor and environmental noise, and $\circledast$ denotes element-wise multiplication. Under low-light conditions, insufficient and non-uniform illumination can obscure structural information and increase the influence of noise [15].

To address these degradations, the proposed framework estimates the illumination component from the image luminance using Gaussian smoothing. It derives the reflectance representation by normalizing the input image by the estimated illumination. It then applies illumination correction using a gamma transformation, followed by reconstruction of the corrected representation. In parallel, the PCOM uses local contrast and luminance-masking cues to generate a spatially adaptive enhancement weight. This weight regulates the fusion between the corrected representation and the original low-light image. Finally, global gamma adjustment and bilateral filtering are applied to produce the enhanced output $\hat{I}$.

3.1 Single-image enhancement stage

Given a low-light RGB image *I,* the proposed framework processes each image independently to generate an enhanced output $\hat{I}$. The enhancement procedure combines Retinex-inspired illumination reflectance decomposition, local contrast estimation, luminance masking, gamma-based illumination correction, and perception-guided adaptive fusion [10]. The resulting image is subsequently refined using mild global gamma adjustment and bilateral filtering to improve illumination while preserving relevant structural information.

3.2 Downstream face analysis

Enhanced images are subsequently evaluated using fixed downstream face-analysis models. YOLOv8-Face is used to assess face-detection performance on the DARK FACE bounding-box annotations, while EfficientNet-B0 [14] is used for binary Face/Non-Face classification. The same evaluation protocol is applied to the original and enhanced images to quantify the effect of low-light enhancement on downstream facial analysis.

4. Deep learning-based detectors

Deep learning-based detectors mainly employ Convolutional Neural Networks (CNNs) to automatically learn high-level feature representations from input images and significantly outperform traditional detectors, which rely on handcrafted features [16]. The Region-based Convolutional Neural Network (R-CNN), fast R-CNN, and Faster R-CNN are some of the advanced CNN architectures that are widely used for face detection applications. The major drawback of the conventional region-based methods is that they lack any class-specific bounding box regression module and hence fail to do well in optimally localizing the parts of the object along with the classifiers [17]. These computational constraints have been tackled by single-stage detectors such as You Only Look Once (YOLO) and Single Shot Multibox Detector (SSD), which redefine the detection problem as a single regression problem [18]. These architectures enable real-time inference due to the single pass of the image through the feed-forward architecture. Despite these advances, the performance of deep learning-based face detectors can degrade considerably when input images are captured under severe low-light conditions [19]. Poor or variable lighting may cause facial features to be poorly contrasted, edges to be blurred, color to be distorted, and sensor noise to be increased. This makes it harder to identify and examine features on the face [20]. These effects are particularly important for small or poorly illuminated faces, where discriminative visual information is already limited. Therefore, improving the visual quality of individual low-light images before downstream facial analysis is an important consideration [21]. The proposed hybrid framework addresses this problem by enhancing illumination while preserving local facial structures, edges, and texture information. The proposed method operates independently on each low-light image, consistent with the image-based experimental protocol adopted using the DARK FACE dataset.

5. Challenges in low-light face image analysis

Face detection and analysis in low-light images remain challenging because facial appearance is strongly affected by variations in illumination intensity, color, and direction [22]. Abundant illumination is essential for face detection, yet notoriously unreliable in real-world settings, varying in intensity, color, and direction. Images are often taken under unbalanced or unnatural lighting conditions, including dim environments, shadows in the object region, and backlight contexts [23]. Performance often plummets in the presence of severe illumination bias and low resolution owing to large or small distances from the camera [19]. However, as HVS has evolved to address these problems, face detection systems also improve when images are captured under substantially different lighting conditions. Similarly, face detection systems are easily fooled when detecting faces in extreme poses. Although classic frontal face datasets provide realistic training images, they ignore variations that naturally occur within a face set. Agents may also appear in extreme pose situations in many real-life scenarios, such as in crowd scenes [24]. The localization of faces under partial occlusion is also challenging, particularly when shadows, external objects, or extreme pose variations obscure facial regions [25]. Real-world faces may appear at substantially different scales, while changes in viewing angle can alter facial appearance from frontal to profile views. In addition, very small or low-resolution facial regions may contain insufficient visual information for reliable detection and analysis [26].

These challenges further motivate the development of enhancement approaches that preserve facial structures while improving visibility under difficult illumination conditions.

6. Literature review

The studies on low-light image enhancement have significantly evolved in recent years, especially because they need to address the problem of noise enhancement, loss of structures, and disjunction of perception in the dark [8]. The Retinex theory has been one of the major paradigms, giving a physically interpretable breakdown of the illumination and reflectance [27]. Despite this conceptual superiority, several studies indicate that classical implementations of Retinex are prone to poor performance in extreme noise or spatially varying illuminant conditions and that more adaptive and learning approaches are needed [28]. Recent work has explored how to incorporate Retinex priors into contemporary deep architectures. An example is Bai et al. (2024) [5], who proposed a Retinex Mamba that uses the state-space modeling features of Mamba architectures to learn a more powerful feature representation for illumination-reflectance decomposition. They found it more stable in complex lighting conditions. In the latter, Chen et al. [29] proposed a spatial redundancy and channel redundancy compression scheme by introducing the concept of joint attention mechanisms. They suggested that it could be possible to increase the structural fidelity while reducing the computational load, thanks to the spatial redundancy and channel redundancy. Furthermore, Jiang et al. (2024) [30] proposed an attention-based Retinex approach that integrates a CBAM with local adaptive enhancement to improve brightness, contrast, and detail preservation while reducing noise and color distortion. Whereas the multi-frequency feature separation method proposed by Liang et al. (2023) [31] pointed out that multi-frequency feature separation is critical in the global brightness correction and fine detail preservation trade-off, particularly for highly degraded images.

Beyond enhancement architectures, recent studies have increasingly incorporated perceptual considerations into low-light image enhancement. Liang et al. (2025) [32] emphasized that evaluating low-light image enhancement is highly subjective and proposed integrating human aesthetic preferences into the enhancement process. These findings suggest that perceptual constraints can complement pixel-level and frequency-domain optimization in low-light image enhancement, particularly when the objective is to improve visual naturalness and perceptual quality [33]. Previous approaches have made significant progress in low-light image restoration through illumination estimation, noise reduction, and structure reconstruction [8]. Nevertheless, some approaches struggle to achieve high-quality results that preserve image detail and perceptual realism while also addressing spatially varying illumination, color distortion, noise, and image structure [33]. This research exploits this gap by proposing a perception-aware module. One pathway focuses on image reconstruction and structural/surface refinement with an emphasis on illumination. The second pathway uses priors based on HVS properties to control the degree of image enhancement while maintaining the realism of the image content. Low-light image enhancement and restoration for robust face detection under challenging illumination conditions remain important research topics, particularly for applications such as surveillance and other vision-based systems [23]. Ding and Hu (2020) [34] proposed

a two-stage framework for extreme low-light face processing, consisting of a Face Location Network (FLNet) for face localization and a Face Enhancement Network (FE-Net) for progressive face restoration using semantic information from facial landmarks. Recent deep-learning approaches have explored both adapting face detectors to low-light conditions and using multi-exposure representations to improve face detection under challenging illumination [23]. Each study's data is summarized in Table 1.

7. Methodology

7.1 Processing operations

This section presents the overall processing pipeline of the proposed perception-guided low-light enhancement framework, as illustrated in Figure 1. The framework is designed for single-image low-light facial enhancement and combines Retinex-inspired illumination–reflectance decomposition [38] with perception-guided adaptive processing.

Table1. Summary of related studies on Low-Light image enhancement and face detection

RF	Aim	Limitation
[34]	The paper focuses on detecting and enhancing human faces in extreme low-light images using a Face Location Network and a Face Enhancement Network. Still, it does not specifically address dual-branch perception-guided learning or the YouTube Database.	• The context provided does not mention any specific limitations of the study. • There is no information regarding potential drawbacks or challenges faced by the authors in their approach. • The abstract focuses on the methodology and results without addressing limitations. • The authors do not discuss any constraints related to the dataset or experimental conditions. • Overall, the context lacks details on limitations from the authors' perspective.
SID-Face Variant (Learning to See Faces in the Dark) + 300 x Amplification gain: up to >20% increase in detection rate over sRGB pipelines.		
[35]	The paper does not address enhancing low-light image restoration via dual-branch perception-guided learning using the YouTube Faces Database. It focuses on a joint high-low adaptation framework for low-light face detection, utilizing existing normal-light data.	• Huge gap between normal and low-light conditions. • Existing methods are not robust for dark face detection. • Existing methods struggle with normal-to-low-light adaptation. • Data collection for low light is time-consuming and laborious.
HLA-Face (Joint High-Low Adaptation) + mAP = 44.4%		
[36]	The paper does not specifically address enhancing low-light image restoration or dual-branch perception-guided learning for robust face detection in dynamic lighting conditions using the YouTube Faces Database. It focuses on various methods for facial image processing and analysis.	• The paper indicates that many problems in facial image processing and analysis remain unsolved despite significant progress in face detection and recognition. • It highlights that current systems lack the robustness required for reliable deployment in real-world scenarios. • The challenges faced in face detection include outdoor illumination, pose variation with large rotation angles, low image quality, low resolution, occlusion, and background changes. • Face recognition algorithms are still a major area of research, particularly those that can cope with a variety of aspects such as viewpoint, background lighting, occlusion, facial expressions, and aging. • Research should be carried out on enhancing face detection and recognition, performance analysis, and real-time biometric system implementation.
Special Issue Overview (Facial Image Processing and Analysis) + 30% to 40% performance drop under extreme directional illumination.		
[37]	This paper uses a deep neural network approach to improve face detection in low light using the MMOD algorithm. The paper does not apply dual-branch perception-guided learning, nor does it use the YouTube Faces Database in either implementation or results.	• The context provided does not mention any specific limitations of the study. • There is no information regarding potential drawbacks or challenges faced by the proposed method. • The abstract focuses on the effectiveness and applications of the low-light face detection system without discussing limitations. • The absence of limitations in the context suggests that the authors may not have highlighted any in the provided text. • Therefore, no limitations can be identified based on the available information.
The developed approach uses the Zero-DCE super-resolution enhancement method and the MMOD (Max-Margin Object Detection) algorithm for face feature detection and illumination extraction without the use of reference images. The approach exhibits a real time capability for nighttime surveillance applications, low computational delay, and achieves 96.0% detection rate.		
[23]	The study proposes REGDet, which combines Recurrent Exposure Generation (REG) and Multi-Exposure Detection (MED) to improve face detection under low-light conditions and non-uniform illumination.	• Limited photons and unevenly distributed noise in low-light images. • Non-uniform illumination and noise issues affecting face detection performance.
The REGDet framework consists of two components: the REG module and an adaptive Feature Pyramid Network, which gradually restore and detect hidden facial features. It was evaluated on the DARK FACE dataset and achieved a record 48.3% mAP (scaling up to 52.1%), which is significantly better than the conventional low-light pipelines.		

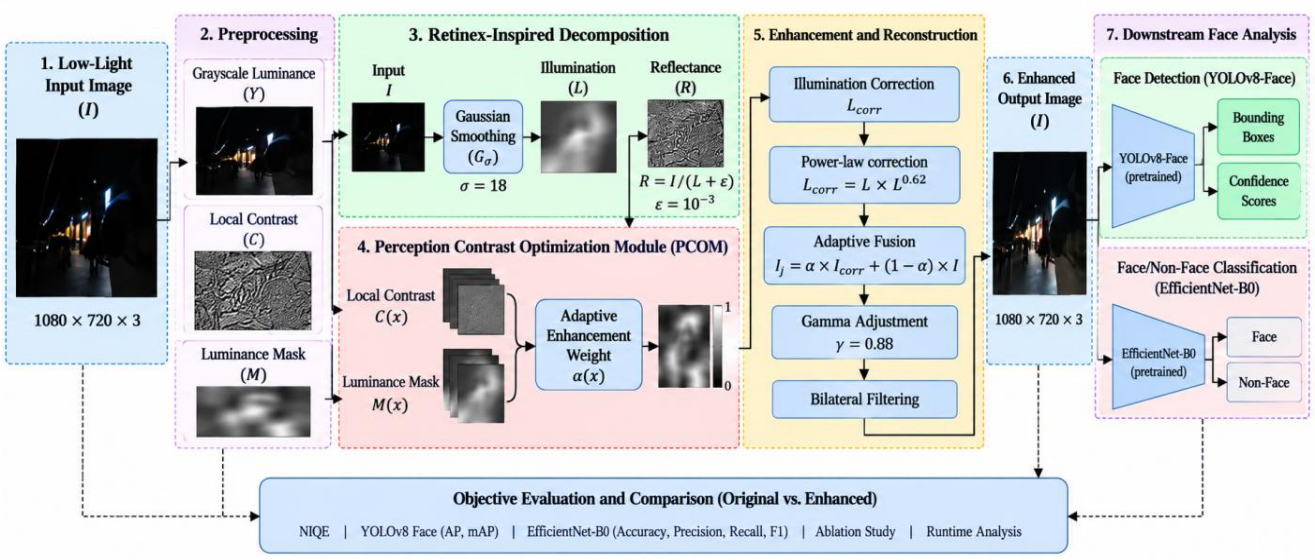


Figure 1. Proposed perception-guided illumination-reflectance enhancement followed by face analysis

First, the luminance component is extracted from the input image, and Gaussian smoothing is used to estimate the illumination representation. The reflectance representation is then obtained by normalizing the input image with respect to the estimated illumination [38]. Illumination correction is performed using gamma adjustment, followed by reconstruction of the corrected image. In parallel, the Perception Contrast Optimization Module (PCOM) derives local contrast and luminance-masking cues to generate a spatially adaptive enhancement weight. The parameter weight determines how the restored image is combined with the low-quality image.

Finally, bilateral filtering and global gamma correction are used to generate the final result that retains desired image features and minimizes any remaining noise [39]. The system takes a single low-light image I and first converts it to a grayscale luminance image Y. It then estimates the illumination term L from Y using Gaussian smoothing. The reflectance component R is subsequently obtained by normalizing the original image with respect to the estimated illumination [38]. This Retinex-inspired decomposition separates spatial illumination variations from the structural information required for subsequent enhancement [10]. The estimated illumination component is corrected using a gamma transformation to improve the visibility of poorly illuminated regions [11]. The corrected illumination is combined with the reflectance image to form an intermediate corrected image Icorr, which retains the structural content of the low-light image while correcting for illumination effects. In addition, the PCOM extracts local image contrast C(x) and luminance masking M(x) from the luminance. Using these perceptual inputs, the method creates an adaptive enhancement factor that adjusts the enhancement level to local image content. This factor is spatially smoothed and restricted to the range [0.25,0.85] to avoid over- or under-enhancement. The corrected image and low-light image are then adaptively fused using the perceptual factors. This allows more of the corrected image to be used in areas that need the most enhancement and the low-light image in areas where aggressive enhancement would create artifacts. A bilateral filter is then used to reduce any remaining noise while retaining edges [40].

A gentle gamma correction also increases visibility. The resulting image $\hat{I}$ represents the final enhanced output of the proposed framework.

7.2 Retinex-Inspired decomposition and perception-guided fusion

Given a normalized low-light input image I, its luminance representation Y is first obtained by grayscale conversion. As illustrated in Figure 1, the first stage of the proposed framework performs a Retinex-inspired decomposition, in which the illumination component is estimated using Gaussian smoothing, while the reflectance component is obtained by normalizing the input image with respect to the estimated illumination [38]:

$$L = G_{\sigma=18}(Y), \quad R = clip(\frac{I}{L+\epsilon}, 0, 1.6) \tag{3}$$

where $G_\sigma(.)$ denotes Gaussian smoothing with standard deviation σ, L represents the estimated illumination, R denotes the reflectance representation, and ϵ is a small constant used to avoid numerical instability. The illumination values are constrained to the range *[0.025,1.0]* before reflectance estimation.

7.3 Perception-Guided enhancement and adaptive fusion

As illustrated in Figure 1, the proposed framework processes the illumination and reflectance representations obtained through Retinex-inspired decomposition. The illumination component is estimated from the grayscale luminance representation using Gaussian smoothing, while the reflectance component is obtained by normalizing the input image with respect to the estimated illumination [38]. The illumination representation is subsequently corrected using gamma transformation and recombined with the reflectance representation to reconstruct an intermediate corrected image [11]. In parallel, local contrast and luminance-masking cues are derived through the PCOM to generate a spatially adaptive enhancement weight. This weight regulates fusion between the corrected representation and the original low-light image based on local illumination and contrast characteristics. Finally, global gamma adjustment and bilateral filtering are applied to

obtain the enhanced output while preserving relevant structural information [39]. The perception-sensitive module uses local contrast and luminance-masking information from the input image to produce an adaptive factor α(x) that controls the amount of enhancement at each image location. Illumination correction is performed using:

$$L_{corr}(x) = L(x)^{\gamma}, \qquad \gamma = 0.62 \tag{4}$$

The corrected illumination component $L_{corr}(x)$ is obtained by raising the original illumination component $L(x)$ to the power of the gamma correction parameter [11] γ where $\gamma = 0.62$.

The corrected image is reconstructed as:

$$I_{corr}(x) = clip(R(x) \circledast L_{corr}(x), 0,1) \tag{5}$$

where $I_{corr}(x)$ denotes the reconstructed corrected image at spatial location *(x)*, *(R(x))* represents the reflectance component, and $L_{corr}(x)$ denotes the corrected illumination component. The operator $\circledast$ represents element-wise multiplication between the reflectance and illumination components, while the clipping operation constrains the resulting pixel values to the normalized intensity range [0,1]. The final adaptive fusion between the corrected representation and the original low-light image is then defined as:

$$L_{enh}(x) = \alpha(x) I_{corr}(x) + [1 - \alpha(x)] I(x) \tag{6}$$

where $L_{enh}(x)$ denotes the final enhanced image at spatial location *(x)*, $I_{corr}(x)$ represents the corrected image, and *I(x)* denotes the original low-light input image. The adaptive weight $\alpha(x)$ controls the contribution of the corrected image at each spatial location, while $[1 - \alpha(x)]$ determines the contribution of the original image. The value of $\alpha(x)$ is constrained to the range [0.25,0.85].

The complete sequence of the proposed perception-guided low-light enhancement procedure is summarized in Algorithm 1, which presents the main processing steps, including Retinex-inspired decomposition, dual-branch feature processing, perception-guided adaptive weighting, illumination correction, image reconstruction, adaptive fusion, gamma adjustment, and bilateral filtering [11, 38, 39].

Algorithm 1: Perception-Guided Low-Light Enhancement

1: Normalize I to the range [0,1]
2: Y ← Grayscale(I)
3: L ← GaussianBlur(Y, σ = 18)
   L ← clip(L, 0.025, 1.0)
4: R ← I / L
   R ← clip(R, 0, 1.6)
5: μ ← GaussianBlur(Y, σ = 4)
   v ← max(GaussianBlur(Y^2, σ = 4) − μ^2, 0)
6: C ← √v
   C ← C / (max(C) + ε)
7: M ← clip(1 − Y, 0, 1)
8: α ← 0.25C + 0.75M
   α ← GaussianBlur(α, σ = 2)
   α ← clip(α, 0.25, 0.85)
9: Lcorr ← L^0.62
10: Icorr ← R ⊙ Lcorr
    Icorr ← clip(Icorr, 0, 1)
11: Ienh ← α ⊙ Icorr + (1 − α) ⊙ I
12: Î ← clip(Ienh, 0, 1)^0.88
13: Convert Î to 8-bit representation
14: Î ← BilateralFilter(Î, d = 5, σ color = 18, σ space = 18)
15: Return Î

8. Experimental setup in datasets and preprocessing

For evaluation of the proposed low-light face image enhancement method and its application to subsequent face recognition tasks, the DARK FACE dataset [12] was employed. This dataset consists of 6,000 real outdoor low-light images taken at night in various scenarios. All images are annotated with a face bounding box. Each image is 1080×720 pixels. In addition, there are a number of unlabeled low-light images also available in the dataset. There is also a subset of low light/normal light image pairs available for image restoration experiments. In this paper, we use the following features of the DARK FACE dataset, as shown in Table 2.

Table 2. Summary of the DARK FACE dataset characteristics

Property	Details
Dataset	DARK FACE
Main Task	Low-light face detection and enhancement evaluation
Image Type	Real-world nighttime / low-light images
Annotated Images Used	6,000 images
Training Set	4,200 images
Validation Set	900 images
Held-Out Test Set	900 images
Training Face Annotations	35,377 faces
Validation Face Annotations	7,612 faces
Test Face Annotations	7,410 faces
Image Resolution	1080 × 720 pixels
Annotation Type	Face bounding boxes
Enhancement Input	Full-resolution low-light images
Quality Evaluation	NIQE and BRISQUE
Detection Evaluation	YOLOv8-Face
Classification Evaluation	EfficientNet-B0 Face/Non-Face

For the experiments conducted in this study, the 6,000 annotated low-light images were used as the principal dataset because the downstream evaluation focuses on facial information under challenging illumination conditions. The bounding-box annotations were retained for face localization and detector-based evaluation. The annotated images used in this study were divided into 4,200 training images, 900 validation images, and 900 held-out test images. The corresponding numbers of annotated faces were 35,377, 7,612, and 7,410, respectively. The enhancement procedure was applied to the full-resolution (1080×720) images so that the original spatial coordinates of the face annotations remained unchanged.

Because the principal annotated DARK FACE set does not provide a normal-light reference image for each low-light image [12], full-reference image-quality metrics were not computed on the 6,000-image annotated set. Instead, the enhancement performance on this set was assessed using no-reference image-quality metrics, including NIQE [41] and BRISQUE [42], together with downstream face detection and face/non-face classification experiments. This protocol avoids treating the main DARK FACE benchmark as a paired low-light/normal-light dataset and provides an evaluation consistent with its low-light face-analysis setting. The experiments were conducted using the annotated low-light images from the DARK FACE dataset. The 6,000 images were divided into 4,200 training images, 900 validation images, and 900 held-out test images. The corresponding numbers of annotated faces were 35,377, 7,612, and 7,410, respectively.

For the enhancement and face-detection experiments, the proposed enhancement procedure was applied directly to the full-resolution (1080×720) images rather than to randomly extracted (256×256) patches. This strategy preserves small facial regions and maintains the original spatial correspondence of the DARK FACE bounding-box annotations. Because the main annotated DARK FACE set does not provide a corresponding normal-light reference for each low-light image, these images were not treated as paired samples. Consequently, evaluation on the held-out set was performed using no-reference image-quality metrics and downstream face-analysis tasks. YOLOv8-Face was applied to the original and enhanced full-resolution images using the same detection protocol to quantify the effect of enhancement on face-detection performance. EfficientNet-B0 was additionally employed for a separate binary Face/Non-Face classification experiment rather than for facial identity recognition. The configuration of the EfficientNet-B0 model used for this classification experiment is summarized in Table 3 [14].

Table 3. EfficientNet-B0 configuration used for Face/Non-Face classification

Component	Description
Model	EfficientNet-B0
Input Size	224×224×3
Pretraining	ImageNet pretrained weights
Classification Task	Binary Face/Non-Face classification
Main Building Block	Mobile Inverted Bottleneck Convolution (MBConv)
Convolution	Depthwise separable convolution within MBConv blocks
Feature Attention	Squeeze-and-Excitation (SE) mechanism
Feature Extraction	Hierarchical convolutional feature extraction
Pooling	Global Average Pooling
Classification Head	Fully connected layer with 2 output classes
Output Classes	Face and Non-Face
Total Parameters	5,288,548
Model Variant	EfficientNet-B0

Face samples were generated from the DARK FACE bounding-box annotations, while non-face samples were extracted from background regions that did not overlap annotated faces. A balanced dataset of 1,384 samples was constructed, consisting of 692 Face and 692 Non-Face crops. The samples were divided into 968 training, 207 validation, and 209 test samples. An ImageNet-pretrained EfficientNet-B0 was used as the classification backbone. The same held-out test samples were evaluated in their original and enhanced forms to determine whether enhancement affected the preservation and detectability of facial information. Classification performance was assessed using accuracy, precision, recall, F1-score, and confusion matrices.

9. Experimental results

The proposed low-light facial image enhancement framework was evaluated using images from the DARK FACE dataset. The main experimental evaluation was conducted on a held-out set of 900 low-light images that was not used during the development of the enhancement procedure. Since the main labeled DARK FACE set does not provide a corresponding normal-light reference image for every low-light image, the evaluation focused on no-reference image-quality assessment and downstream face-analysis performance. NIQE and BRISQUE were used to assess perceptual image quality, while a fixed YOLOv8-Face detector was employed to quantify face-detection performance before and after enhancement. In addition, ImageNet-pretrained EfficientNet-B0 was evaluated on a clearly defined binary face-presence classification task. Ablation experiments were subsequently conducted to investigate the contribution of the principal enhancement components.

9.1 No-Reference image quality assessment

To evaluate enhancement quality without requiring paired normal-light reference images, NIQE and BRISQUE were calculated on the same 900-image held-out test set. Lower values for both metrics indicate better no-reference perceptual quality. The results of the no-reference image-quality assessment are presented in Table 4. As shown in Table 4, the mean NIQE score decreased from 4.8652 for the original low-light images to 4.7828 after enhancement. The corresponding standard deviation also decreased substantially from 1.6292 to 0.7524, indicating more consistent NIQE-based quality across the enhanced images. In contrast, the mean BRISQUE score increased from 10.4981 to 17.6841, although its standard deviation decreased from 29.2551 to 18.7534. Therefore, the two no-reference measures do not indicate uniform improvement. Rather, the results suggest a perceptual-quality trade-off, emphasizing the importance of evaluating the proposed enhancement in conjunction with its intended downstream face-analysis task rather than relying on a single image-quality metric.

Table 4. No-Reference image-quality assessment on the 900-image Held-Out test set

Image Condition	NIQE ↓	BRISQUE ↓
Original Low-Light	4.8652 ± 1.6292	10.4981 ± 29.2551
Enhanced	4.7828 ± 0.7524	17.6841 ± 18.7534

A representative visual comparison between the original low-light and enhanced DARK FACE images is presented in Figure 2, providing qualitative evidence of the visual changes produced by the proposed enhancement procedure. Figure 2 presents representative qualitative examples from the held-out test images. The enhanced outputs exhibit increased visibility in poorly illuminated regions while preserving the overall spatial structure required for subsequent face analysis. These qualitative observations are interpreted together with the quantitative no-reference metrics and downstream detection results rather than as evidence of uniform perceptual improvement.

9.2 Face detection performance before and after enhancement

The practical objective of low-light facial enhancement is not only to improve visual appearance but also to make facial information more accessible to downstream computer-vision systems. Therefore, YOLOv8-Face was used as a fixed face detector. The same detector configuration and confidence threshold were applied to both the original and enhanced versions of the 900 held-out images, and detections were matched against the DARK FACE ground-truth bounding boxes using an IoU threshold of 0.50. The YOLOv8-Face detection performance on the original and enhanced images is presented in Table 5.



Figure 2. Representative original low-light and enhanced DARK FACE images

The enhancement produced a substantial increase in face-detection sensitivity. The number of true-positive detections increased from 848 to 1,418, while false negatives decreased from 6,562 to 5,992. Consequently, recall increased from 11.44% to 19.14%, and the F1-score increased from 20.86% to 30.77%. AP@0.50 improved from 16.80% to 26.22%, corresponding to an absolute gain of 9.42 percentage points. Similarly, mAP@0.50:0.95 increased from 4.99% to 7.82%. This improvement was accompanied by an increase in false positives from 197 to 388, resulting in a moderate reduction in precision from 81.15% to 78.52%. Thus, the enhancement primarily improves the recovery of difficult low-light faces, with a trade-off in the form of additional false-positive detections. The observed improvements in AP@0.50, mAP@0.50:0.95, recall, and F1-score indicate that the enhanced images improved the detection performance of the fixed face detector under the evaluated conditions. Representative YOLOv8-Face detection results on the original and enhanced low-light images are presented in Figure 3, providing a qualitative illustration of the improved detectability of facial regions after enhancement. Figure 3 illustrates representative examples of face detections before and after enhancement. The visualization should use the same test images and detector configuration for both conditions, allowing the additional recovered faces after enhancement to be directly observed.

Table 5. YOLOv8-Face detection performance on original and enhanced images

Metric	Original Low-Light	Enhanced	Change
True Positives	848	1,418	+570
False Positives	197	388	+191
False Negatives	6,562	5,992	−570
Precision	81.15%	78.52%	−2.77 pp
Recall	11.44%	19.14%	+7.69 pp
F1-score	20.86%	30.77%	+10.71pp
AP@0.50	16.80%	26.22%	+9.42 pp
mAP@0.50:0.95	4.99%	7.82%	+2.83 pp

9.3 EfficientNet-B0 binary face-presence classification

To provide a well-defined downstream classification experiment, EfficientNet-B0 was evaluated for binary face-presence classification, rather than facial identity recognition. The two classes were defined as Face and Non-Face. Positive face samples were generated from the ground-truth facial bounding boxes provided by DARK FACE, while negative samples were obtained from background regions that did not overlap the annotated facial regions. A balanced dataset containing 1,384 image crops was constructed, comprising 692 Face and 692 Non-Face samples. The dataset was divided into 968 training samples, 207 validation samples, and 209 held-out test samples. An ImageNet-pretrained EfficientNet-B0 was used. Its feature-extraction backbone was frozen, and the original 1,000-class output layer was replaced with a two-class classification layer. Only the classification head was optimized for the binary task. Input crops were resized to (224×224) pixels and normalized using the ImageNet normalization parameters. During five training epochs, validation accuracy increased from 93.72% after the first epoch to 96.62% and subsequently remained stable. The final trained classifier was then evaluated on the independent 209-sample test partition. The EfficientNet-B0 face-presence classification results on the held-out test set are presented in Table 6. On the original low-light samples, EfficientNet-B0 achieved an accuracy of 95.69%, precision of 100.00%, recall of 91.74%, and F1-score of 95.69%. Following enhancement, accuracy and F1-score remained nearly unchanged at 95.22% and 95.58%, respectively. However, recall increased markedly from 91.74% to 99.08%, indicating that substantially fewer facial samples were missed after enhancement. Precision decreased to 92.31%, reflecting an increased number of background samples classified as faces.

Table 6. EfficientNet-B0 face-presence classification on the held-out test set

Metric	Original Low-Light	Enhanced	Change
Accuracy	95.69%	95.22%	−0.47 pp
Precision	100.00%	92.31%	−7.69 pp
Recall	91.74%	99.08%	+7.34 pp
F1-score	95.69%	95.58%	−0.11 pp



Figure 3. Representative YOLOv8-face detection results on original and enhanced Low-Light images

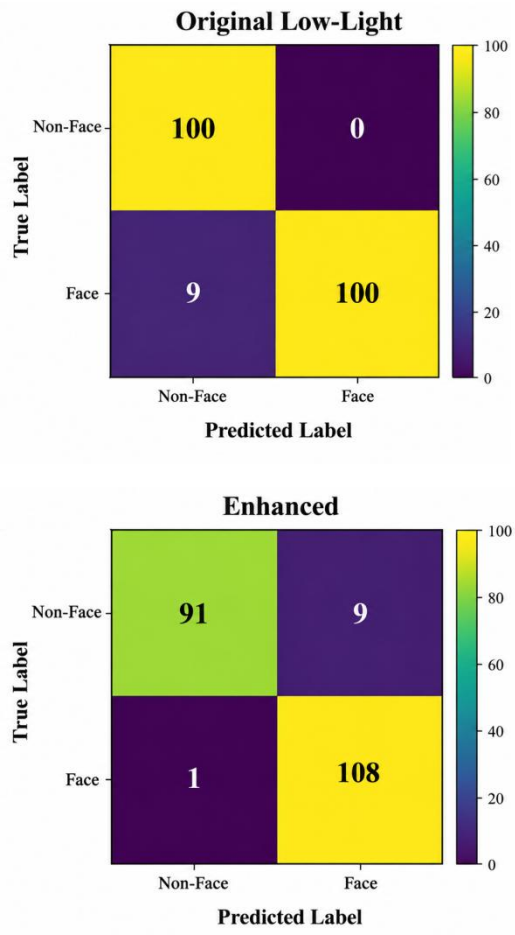


Figure 4. Confusion matrices of EfficientNet-B0 for original and enhanced face-presence classification

The confusion matrices provide further insight into this behavior. For the original low-light test samples, the classifier produced 100 true negatives, 0 false positives, 9 false negatives, and 100 true positives. After enhancement, the same 209 test samples yielded 91 true negatives, 9 false positives, 1 false negative, and 108 true positives. Thus, enhancement reduced false-negative face classifications from 9 to 1 and increased true-positive classifications from 100 to 108, while introducing 9 false-positive predictions. These findings are consistent with the YOLOv8-Face experiment: enhancement increases sensitivity to low-light facial information while introducing a moderate precision trade-off. The corresponding confusion matrices for the original and enhanced face-presence classification results are presented in Figure 4, providing a visual comparison of the classification outcomes before and after enhancement. Figure 4 compares the confusion matrices obtained from the same 209 held-out samples. The reduction in false-negative facial classifications from nine to one shows EfficientNet-B0's increased sensitivity to facial information after enhancement.

9.4 Ablation Study

An ablation study was conducted on a fixed subset of 100 test images to investigate the contribution of the principal components of the enhancement procedure. Three variants were compared with the complete method: removal of HVS guidance, removal of adaptive gating, and removal of Retinex decomposition. All variants were evaluated using the same images and YOLOv8-Face configuration. The results of the ablation study on the 100 held-out DARK FACE images are presented in Table 7. Removing HVS guidance decreased the detection F1-score from 33.45% to 32.71%, while removing adaptive gating reduced it to 32.95%. These results suggest that HVS-guided processing and spatially adaptive fusion contribute to downstream face-detection performance under the evaluated experimental conditions.

Table 7. Ablation study on 100 Held-Out DARK FACE images

Variant	NIQE ↓	BRISQUE ↓	Precision (%)	Recall (%)	F1-score (%)
Full Method	4.7942	17.8619	80.77	21.09	33.45
w/o HVS Guidance	4.7221	16.4080	80.35	20.54	32.71
w/o Adaptive Gating	4.8048	18.1609	79.83	20.76	32.95
w/o Retinex	4.5814	15.1185	83.12	21.43	34.07

However, removal of the Retinex decomposition produced the strongest performance within this 100-image ablation subset, achieving an F1-score of 34.07%, NIQE of 4.5814, and BRISQUE of 15.1185. This finding indicates that the contribution of the Retinex stage is not uniformly positive under the current configuration. Rather than claiming that every individual component independently improves all evaluation measures, the ablation results reveal a trade-off between decomposition-based enhancement and downstream/perceptual performance. Further optimization of the Retinex stage therefore represents an important direction for future refinement of the framework.

9.5 Computational performance

Computational runtime was measured using 100 DARK FACE images at their original resolution of (1080×720) pixels. The same enhancement procedure used in the preceding experiments was evaluated without changing the image resolution. The computational runtime of the enhancement procedure is reported in Table 8.

Table 8. Computational runtime of the enhancement procedure

Measure	Result
Image Resolution	1080×720
Number of Images	100
Mean Runtime	149.50 ms/image
Standard Deviation	4.41 ms
Processing Speed	6.69 FPS

The proposed enhancement procedure required an average processing time of 149.50 ms per (1080×720) image, with a standard deviation of 4.41 ms. These measurements provide an empirical estimate of the computational cost of the implemented enhancement procedure under the tested execution environment. The reported runtime reflects independent processing of individual low-light images and demonstrates the computational requirements of the proposed method at full image resolution.

9.6 Overall performance analysis

The combined experiments demonstrate that the effect of low-light enhancement cannot be characterized by a single metric. NIQE showed a modest improvement after enhancement, whereas BRISQUE increased. More importantly for the intended facial-analysis application, YOLOv8-Face exhibited substantial gains in recall, F1-score, AP@0.50, and mAP@0.50:0.95. EfficientNet-B0 produced a similar pattern: face-class recall increased from 91.74% to 99.08%, while overall accuracy and F1-score remained approximately stable and precision decreased. These findings indicate that the principal benefit of the enhancement procedure is increased facial visibility and downstream sensitivity under difficult low-light conditions, rather than uniform improvement across every perceptual-quality measure. The ablation experiments further indicate positive contributions from HVS guidance and adaptive gating to detection F1-score, while also identifying the Retinex stage as a component requiring further optimization. Therefore, the experimental evidence supports a more specific conclusion: the enhancement procedure improves the recoverability of low-light facial information for downstream analysis, but this improvement involves measurable perceptual and precision trade-offs.

10. Conclusion

This study presented a perception-guided Retinex-inspired framework for enhancing facial information in real-world low-light images. The proposed approach combines illumination–reflectance decomposition with local contrast and luminance-guided adaptive enhancement to improve poorly illuminated regions while preserving relevant structural information. The experimental protocol was conducted on the DARK FACE dataset using full-resolution (1080×720) images, thereby maintaining the original spatial information and face bounding-box annotations. The effectiveness of the enhancement process was evaluated from both image-quality and downstream face-analysis perspectives. No-reference NIQE and BRISQUE metrics were employed because the main annotated DARK FACE set does not provide corresponding normal-light ground-truth images. Face-detection experiments further demonstrated that enhancement increased the number of detected faces under challenging illumination conditions. In the 100-image evaluation subset, the proposed full method achieved an F1-score of 33.45%, with 189 true-positive face detections, compared with the weaker detection performance observed on the original low-light images. The ablation experiments further demonstrated the influence of the perception-guided components on enhancement and downstream detection behavior. A separate EfficientNet-B0 Face/Non-Face classification experiment was conducted using a balanced set of 1,384 facial and non-facial crops. On the held-out test set of 209 samples, the original images achieved 95.69% accuracy and 95.69% F1-score, whereas the enhanced images achieved 95.22% accuracy and 95.58% F1-score. Although overall classification accuracy remained nearly unchanged, enhancement substantially increased Face recall from 91.74% to 99.08%, reducing false negatives from nine to one. This indicates that the enhancement process can improve the visibility of difficult facial regions, although the associated reduction in precision highlights a trade-off between recovering poorly visible faces and introducing additional false-positive responses. The proposed enhancement method achieved an average runtime of 149.50 ms for each (1080×720) image at 6.69 FPS. In general, the experiments demonstrate that low-light image enhancement with facial perception guidance can enhance face recognition and visibility with minimal impact on the succeeding classification task. More tests with other state-of-the-art low-light image enhancement algorithms, standard face-detection mAP metrics, and larger independent datasets are needed. Additionally, more efficient implementations and learned adaptive enhancement techniques will be explored to help find an optimal compromise between image quality, computational complexity, false-positive reduction, face restoration, and face detection accuracy. There are a few limitations to this experiment. First, testing was only performed on the DARK FACE dataset. Second, the primary manually annotated dataset in DARK FACE does not have normal-light ground-truth images, which prevents full-reference image quality measurement. Third, while face visibility and recognition were enhanced, there was also an increase in the number of false positives. This presents a compromise between enhancing visibility of faces that are not well visible and the accuracy of the face detection. Additional testing on independent low-light face datasets with a broader range of lighting conditions is needed to provide evidence of the generalization of this technique.

Ethical issue
The authors are aware of and comply with best practices in publication ethics, specifically with regard to authorship (avoidance of guest authorship), dual submission, manipulation of figures, competing interests, and compliance with policies on research ethics. The authors adhere to publication requirements that the submitted work is original and has not been published elsewhere.

Data availability statement
The manuscript contains all the data. However, more data will be available upon request from the authors. The DARK FACE used in this work is made publicly available for research use. Experiments were performed on the publicly available DARK FACE dataset, along with annotations for the low-light images and face bounding boxes in the images. The publicly available DARK FACE distribution used in this work provides access to the dataset. None of the data used in this work were independently collected or private.

Conflict of interest
The authors declare no potential conflict of interest.

References

[1] M. M. Tatana, M. S. T., and Maswanganyi, R. C., "Low-Light Image and Video Enhancement for More Robust Computer Vision Tasks: A Review," Journal of Imaging, vol. 11, no. 4, pp. 11040125, 2025. DOI: https://doi.org/10.3390/jimaging11040125.

[2] Y. Fan, Y. Wang, D. Liang, Y. Chen, H. Xie, F. L. Wang, J. L., and Wei, M., "Low-FaceNet: Face Recognition-Driven Low-Light Image Enhancement," IEEE Transactions on Instrumentation and Measurement, vol. 73, no. 5019413, pp. 5019413, 2024. DOI: https://doi.org/10.1109/TIM.2024.3372230.

[3] Q. Zhao, G. Li, B. He, R. S., "Deep Learning for Low-Light Vision: A Comprehensive Survey," IEEE Transactions on Neural Networks and Learning Systems, vol. 36, no. 9, pp. 3566647, 2025. DOI: https://doi.org/10.1109/TNNLS.2025.3566647.

[4] Liu, F., and Fan, L., "A review of advancements in low-light image enhancement using deep learning," Neurocomputing, vol. 652, pp. 131052, 2025. DOI: https://doi.org/10.1016/j.neucom.2025.131052.

[5] Bai, J., Yin, Y., He, Q., Li, Y., and Zhang, X., RetinexMamba: Retinex-Based Mamba for Low-Light Image Enhancement, 2024. DOI: https://doi.org/https://doi.org/10.1007/978-981-96-6596-9_30.

[6] Tian, Z., Qu, P., Li, J., Sun, Y., Li, G., Liang, Z., Zhang, W., "A Survey of Deep Learning-Based Low-Light Image Enhancement," Sensors, vol. 23, no. 18, pp. 23187763, 2023. DOI: https://doi.org/https://doi.org/10.3390/s23187763.

[7] Liu, M., Li, X., and Fang, Y., "Low-Light Image Enhancement via Dual Information-Based Networks," Electronics, vol. 13, no. 18, pp. 13183713, 2024. DOI: https://doi.org/https://doi.org/10.3390/electronics13183713.

[8] J. Zhan, E. S. G., and M. S. Sunar, "Low-light image enhancement: A comprehensive review on methods, datasets and evaluation metrics," Journal of King Saud University - Computer and Information Sciences, vol. 36, no. 10, pp. 102234, 2024. DOI: https://doi.org/https://doi.org/10.1016/j.jksuci.2024.102234.

[9] Y. Feng, S. Hou, H. Lin, Y. Zhu, P. Wu, W. Dong, J. Sun, Q. Y., and Zhang, Y., "DiffLight: Integrating Content and Detail for Low-light Image Enhancement," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops, 2024, pp. 6152.

[10] Jiang, Y., Zhu, J., Li, L., Ma, H., "A Joint Network for Low-Light Image Enhancement Based on Retinex," Cognitive Computation, vol. 16, pp. 12559, 2024. DOI: https://doi.org/https://doi.org/10.1007/s12559-024-10347-4.

[11] R. Weng, Y. Zhang, H. Wu, W. Wang, D. W., "The Retinex enhancement algorithm for low-light intensity image based on improved illumination map," IET Image Processing, vol. 18, no. 12, pp. 13180, 2024. DOI: https://doi.org/https://doi.org/10.1049/ipr2.13180.

[12] Yang, W., Yuan, Y., Ren, W., Liu, J., Scheirer, W.J., Wang, Z., Zhang, T., Zhong, Q., Xie, D., Pu, S., Zheng, Y., "Advancing image understanding in poor visibility environments: A collective benchmark study," IEEE Transactions on Image Processing, vol. 29, pp. 5752, 2020. DOI: https://doi.org/10.1109/TIP.2020.2981922.

[13] Wang, Xiaofeng, Liang Huang, Mingxuan Li, Chengshan Han, Xin Liu, T. N., "Fast, Zero-Reference Low-Light Image Enhancement with Camera Response Model," Sensors, vol. 24, no. 15, pp. 5019, 2024. DOI: https://doi.org/https://doi.org/10.3390/s24155019.

[14] M. Tan, Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," In Proceedings of the 36th International Conference on Machine Learning (ICML), 2019, pp. 6105–6114.

[15] J. Wu, B. Ding, B. Zhang, J. D., "A Retinex-based network for image enhancement in low-light environments," PLOS ONE, vol. 19, no. 5, pp. e303696, 2024. DOI: https://doi.org/https://doi.org/10.1371/journal.pone.0303696.

[16] X. Zhao, L. Wang, Y. Zhang, X. Han, M. Deveci, M. P., "A review of convolutional neural networks in computer vision," Artificial Intelligence Review, vol. 57, pp. 99, 2024. DOI: https://doi.org/https://doi.org/10.1007/s10462-024-10721-6.

[17] J. Cao, B. Peng, M. Gao, H. Hao, X. L. H. M., "Object Detection Based on CNN and Vision- Transformer: A Survey," IET Computer Vision, vol. 19, no. 1, pp. e70028, 2025. DOI: https://doi.org/https://doi.org/10.1049/cvi2.70028.

[18] O. E. Olorunshola, P. O. Jemitola, A. K. A., "Comparative Study of Some Deep Learning Object Detection Algorithms: R-CNN, FAST R-CNN, FASTER R-CNN, SSD, and YOLO," Nile Journal of Engineering and Applied Science, vol. 1, no. 1, pp. 70–80, 2023. DOI: https://doi.org/https://doi.org/10.5455/NJEAS.150264.

[19] J. Xu, Y. Qian, K. Wang, C. Liu, J. L., "Low-light face detection method based on the cross fusion of high- and low-frequency channel features," CAAI Transactions on Intelligent Systems, vol. 19, no. 2, pp. 472–481, 2024. DOI:

https://doi.org/https://doi.org/10.11992/tis.202208034.

[20] Kim M, Jung Y, L. S., "Attention Retinex Network (A4R-Net) for face detection under low-light environment," ICT Express, vol. 10, no. 6, pp. 1206–12011, 2024. DOI: https://doi.org/https://doi.org/10.1016/j.icte.2024.09.009.

[21] Z. Li, Y. Bai, X. Wang, L. Huang, C. Lin, S. L., "Low illumination face detection based on image enhancement," Journal of Computer Applications, vol. 44, no. 8, pp. 2588–2594, 2024.

[22] Li, P., "Multi-model fusion of DSFD and Zero-DCE for low-light face detection," Applied and Computational Engineering, vol. 86, pp. 20241584, 2024. DOI: https://doi.org/https://doi.org/10.54254/2755-2721/86/20241584.

[23] Liang, J., Wang, J., Quan, Y., Chen, T., Liu, J., Ling, H., and Xu, Y., "Recurrent exposure generation for low-light face detection," IEEE Transactions on Multimedia, vol. 24, pp. 1609–1621, 2021.

[24] F. Kasim, T. E. Boult, R. Mora, B. Biesseck, R. Ribeiro, J. Schlueter, T. Repák, R. Vareto, D. Menotti, W. R. Schwartz, M. G., "Watchlist Challenge: 3rd Open-set Face Detection and Identification," Proceedings of the 2024 IEEE International Joint Conference on Biometrics (IJCB), 2024, pp. 1–10. DOI: https://doi.org/https://doi.org/10.1109/IJCB62174.2024.10744535.

[25] A. Alashbi, A. H. H. M. Mohamed, A. A. El-Saleh, I. Shayea, M. S. Sunar, Z. R. Alqahtani, F. Saeed, B. S., "Human face localization and detection in highly occluded unconstrained environments," Engineering Science and Technology, an International Journal, vol. 61, pp. 101893, 2025. DOI: https://doi.org/https://doi.org/10.1016/j.jestch.2024.101893.

[26] L. Liu, G. Wang, Q. M., "ADYOLOv5-Face: An Enhanced YOLO-Based Face Detector for Small Target Faces," Electronics, vol. 13, no. 21, pp. 4184, 2024. DOI: https://doi.org/https://doi.org/10.3390/electronics13214184.

[27] Zhang Y, Jiang J, Wang Z, Zhang Q, Jiang Y, Liu J, H. Z., "Low-light image enhancement method based on retinex theory and dual-tree complex wavelet transform," Journal of King Saud University Computer and Information Sciences, vol. 37, no. 5, pp. 83, 2025. DOI: https://doi.org/https://doi.org/10.1007/s44443-025-00102-6.

[28] Y. Sun, Y. Jin, X. Chen, Y. Xu, X. Yan, Z. L., "Unsupervised detail and color restorer for Retinex-based low-light image enhancement," Engineering Applications of Artificial Intelligence, vol. 153, pp. 110867, 2025. DOI: https://doi.org/https://doi.org/10.1016/j.engappai.2025.110867.

[29] J. Chen, Z. Xiao, X. Qin, D. L., "Retinex-Based Low-Light Image Enhancement via Spatial-Channel Redundancy Compression and Joint Attention," Electronics, vol. 14, no. 11, pp. 2212, 2025. DOI: https://doi.org/https://doi.org/10.3390/electronics14112212.

[30] S. Jiang, Y. Shi, Y. Zhang, Y. Z., "An Improved Retinex-Based Approach Based on Attention Mechanisms for Low-Light Image Enhancement," Electronics, vol. 13, no. 18, pp. 3645, 2024. DOI: https://doi.org/https://doi.org/10.3390/electronics13183645.

[31] X. Liang, X. Chen, K. Ren, X. Miao, Z. Chen, Y. J., "Low-light image enhancement via adaptive frequency decomposition network," Scientific Reports, vol. 13, no. 1, pp. 14107, 2023. DOI: https://doi.org/https://doi.org/10.1038/s41598-023-40899-8.

[32] D. Liang, Y. Gao, L. Li, Z. Xu, S.-J. Huang, S. C., "Aesthetics-Guided Low-Light Enhancement," IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 47, no. 7, pp. 5866–5883, 2025. DOI: https://doi.org/https://doi.org/10.1109/TPAMI.2025.3554639.

[33] B. Chen, L. Zhu, H. Zhu, W. Yang, L. Song, S. W., "Gap-Closing Matters: Perceptual Quality Evaluation and Optimization of Low-Light Image Enhancement," IEEE Transactions on Multimedia, vol. 26, pp. 3430–3443, 2024. DOI: https://doi.org/https://doi.org/10.1109/TMM.2023.3312851.

[34] Ding, X., and Hu, R., "Learning to see faces in the dark," 2020 IEEE International Conference on Multimedia and Expo (ICME), 2020, pp. 33.

[35] Wang, W., Yang, W., and Liu, J., "Hla-face: Joint high-low adaptation for low light face detection," Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, 2021, pp. 16195–16204.

[36] Grgic, M., Shan, S., Lukac, R., Wechsler, H., and Bartlett, M. S., "Facial Image Processing and Analysis," Int J Patt Recog & Artif Intell, vol. 23, no. 3, pp. 355–358, 2009.

[37] Ananthi, G., Balaji, T., Pugalenthi, M., and Rajkumar, M., "Low Light Face Detection System," Recent Advancements in Multimedia Data Processing and Security: Issues, Challenges, and Techniques, 2023.

[38] S. Munaf, A. Bharathi, A. N. J., "FPGA-based low-light image enhancement using Retinex algorithm and coarse-grained reconfigurable architecture," vol. 14, no. 1, pp. 28770, 2024. DOI: https://doi.org/https://doi.org/10.1038/s41598-024-80339-9.

[39] G. Wenjuan, C. Ding, J. Wei, Y. Yin, X. L., "Low-light image enhancement based on the fusion of Bilateral filter MSR and AutoMSRCR," Optics and Precision Engineering, vol. 31, no. 24, pp. 3606–3617, 2023. DOI: https://doi.org/https://doi.org/10.37188/OPE.20233124.3606.

[40] J. Wang, S. Huang, Z. Huo, S. Zhao, Y. Q., "Bilateral enhancement network with signal-to-noise ratio fusion for lightweight generalizable low-light image enhancement," Scientific Reports, vol. 14, no. 1, pp. 29832, 2024. DOI: https://doi.org/https://doi.org/10.1038/s41598-024-81706-2.

[41] A. Mittal, R. Soundararajan, A. C. B., "Making a 'Completely Blind' Image Quality Analyzer," IEEE Signal Processing Letters, vol. 20, no. 3, pp. 209–212, 2013. DOI: https://doi.org/https://doi.org/10.1109/LSP.2012.2227726.

[42] Mittal, A. K. Moorthy, A. C. B., "No-Reference Image Quality Assessment in the Spatial Domain," IEEE Transactions on Image Processing, vol. 21, no. 12, pp. 4695–4708, 2012. DOI: https://doi.org/https://doi.org/10.1109/TIP.2012.2214050.



This article is an open-access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).