International Journal of Applied Mathematics
Volume 38 No. 12s, 2025
ISSN: 1311-1728 (printed version); ISSN: 1314-8060 (on-line version)

A HYBRID CNN–RETINEX FRAMEWORK WITH CONFIDENCE-WEIGHTED REFINEMENT FOR ROBUST SHADOW DETECTION AND REMOVAL

Priya Gupta, Pooja Patre

Department of Computer Science & Engineering

College name: Govt. engineering college, Ambikapur, Chhattisgarh, India

Department of Computer Science & Engineering

College name: Govt. engineering college, Ambikapur,
Chhattisgarh, India

Abstract

Shadow removal and detection is a critical computer vision problem because shadows usually degrade image quality and affect downstream applications such as object detection and scene understanding. Existing methods like the benchmarking framework in Vicente et al. (2023) provide useful datasets and baseline techniques but are still plagued with accurate shadow boundary refinement and realistic shadow removal reconstruction. In this paper, we present a confidence-weighted blending hybrid CNN–Retinex model for robust shadow detection and removal. The approach initially generates initial candidate shadow masks through HSV color space thresholding, which are further filtered using a CNN-based patch classifier for precise localization. For shadow removal, we have employed a Retinex-based illumination correction technique followed by an innovative confidence-weighted blending approach for unnoticeable blending of the shadowed and non-shadowed regions. We evaluate our method using the benchmark set suggested by Vicente et al. and other real images. Experimental results indicate significant improvement in quantitative metrics (PSNR, SSIM, Precision, Recall, F1, mAP) and qualitative visual quality compared to the baseline method. key words

Keywords— Shadow Detection; Shadow Removal; Retinex Algorithm; Convolutional Neural Networks (CNN); Confidence-Weighted Blending; HSV Color Space; Image Enhancement; Benchmark Dataset; PSNR; SSIM; Object Detection

INTRODUCTION

Shadows occur everywhere in our visual world both inside and outside. While shadows provide useful information on lighting direction and object geometries, they also typically present substantial challenges in computer vision. Shadows contribute to image degradation, make scene structures obscure and deteriorate quality, influence object detection, recognition and semantic segmentation in image understanding applications. Therefore, efficient shadow detection and elimination are still necessary in increasing the interpretability and robustness of vision-based systems.

Earlier shadow detection was largely based on color, texture and lighting characteristics. Guo et al. (2011, 2013) introduced the paired region that is based on relative brightness and chromaticity contrast for shadow elimination and detection. Retinex-inspired methods

Received: August 11, 2025

(Finlayson et al., 2006; Wei et al., 2018) attempted to achieve illumination invariance by modeling the reflectance of scene. While these traditional methods indeed demonstrated their usefulness, but mostly when involved with the situation of inception and attribution to complex illumination changes and diverse back-grounds.

Following the emergence of deep learning, a number of shadow detection and removal techniques have been devised, such as GAN-based solutions (Hu et al., 2019) and end-to-end convolutional neural networks (Fan et al., 2019). These models performed superiorly but still suffer from drawbacks like imperfect boundary localization and unnatural shadow-free reconstruction. To overcome these issues, Vicente et al. (2023) proposed a benchmark dataset and baseline model for shadow detection and removal. Their work facilitated a standard evaluation platform and encouraged more studies in this area. Their method is still plagued by inaccurate boundary refinement and texture inconsistency in corrected areas, though.

In this paper, we introduce a confidence-weighted blending hybrid CNN–Retinex framework for robust shadow removal and detection. The process starts by creating preliminary candidate shadow masks via HSV color space thresholding. These masks are then improved through CNN-based patch classification to improve localization precision. For removing shadows, we use a Retinex-based illumination correction method to recover natural intensity levels and a subsequent confidence-weighted blending method that guarantees transition smoothness between non-shadowed and shadowed areas

I. RELTATED WORK

Early approaches to shadow detection and removal primarily used handcrafted features like color, texture, and difference in illumination. For example, Guo et al. (2011, 2013) introduced the paired-region framework, where relative reflectance and illumination between neighboring shadowed and illuminated areas were modeled. Finlayson et al. (2006) also presented Retinex-based illumination-invariant models for estimating reflectance from shading. Though these methods worked in the case of straightforward lighting scenarios, they faltered when confronted with complicated outdoor conditions, unusual shadows, and rough surfaces, resulting in partial shadow elimination details.

With the advent of deep learning, a number of approaches gained substantial improvement in removing shadows. Hu et al. (2019) developed Mask-ShadowGAN that employed adversarial learning and unpaired data for training to produce images without shadows. Fan et al. (2019) came up with an end-to-end CNN architecture that detected and removed shadows at the same time with greater accuracy. RetinexNet (Wei et al., 2018) integrated Retinex decomposition with deep neural networks to obtain improved illumination correction. Though these methods improved performance over conventional methods, they are still plagued by inaccurate boundary refinement and potentially generate artifacts or texture artifacts in shadow-free results.

To facilitate standardized evaluation, Vicente et al. (2023) proposed a benchmark dataset of large scale and a baseline approach to shadow detection and removal. Their dataset includes a variety of shadowed and shadow-free image pairs, allowing for cross-algorithm consistent comparison. Their baseline approach verified the effectiveness of joint detection with reconstruction. Nonetheless, it was still plagued by limitations in sharpening boundary of shadows, preserving texture continuity, and creating smooth transitions within corrected and uncorrected areas.

As can be seen from the above studies, both traditional and deep learning approaches possess their respective strengths and weaknesses. The traditional approaches are straightforward but are ineffective in complicated situations, while deep learning models deliver high accuracy but with limited robustness and realism. Additionally, current benchmark approaches like Vicente et al. (2023) offer a sound basis but do not solve the problems of accurate boundary refinement, smooth blending, and generalization to new data

To fill this gap, we suggest a confidence-weighted blending hybrid CNN–Retinex architecture. Our method bridges the detection strengths of deep learning with the illumination correction features of Retinex theory and adds a confidence-weighted blending process to guarantee smooth and artifact-free reconstruction. This enables us to obtain more precise shadow detection, realistic removal, and better generalization on both benchmark datasets and natural image

II. METHODOLOGY

The proposed shadow removal and detection technique integrates the best elements from both classic color-space analysis and deep learning into a single hybrid pipeline. Traditional methods that use simple thresholding techniques reduce computation costs. However, they usually wrongly classify dark regions that visually appear the same as shadows and as shadows. Deep learning models, on the other hand, are accurate and robust but are cost prohibitive in most practical applications since they are computationally expensive and require large labeled datasets. Integrating the two paradigms, our approach captures the interpretability and computational efficiency of color-space techniques and the robust discrimination of deep neural networks, thus optimizing computational cost, robustness, and accuracy.

The defined workflow has five stages which come one after the other. To begin with, candidate shadow zones are found by adaptive thresholding in the HSV color space because shadow pixels are low in brightness but high in saturation. The second step further filters the candidate shadows by using the CNN-based patch classifier, which reduces false alarms and boundary touch refinement by learning complexities of the contextual patterns far advanced than mere pixel intensities. The third step uses a Retinex based method, which illuminates a picture by breaking it down in the form of reflectance and illumination. It corrects the illumination naturalness level (the brightness) that doesn't erase the surface texture. The fourth step uses the confidence-weighting blend scheme to smoothen the shadow seams. It uses a Gaussian confidence map gradient that smoothens the shadow seams from the inner shadow parts towards the border. The fifth step uses visual assessment and other metrics like consistency, constancy, and robustness to evaluate the tracking performance of the system. This five-step system not only improves the detection accuracy but also the shadow removal quality effective for real-world use.

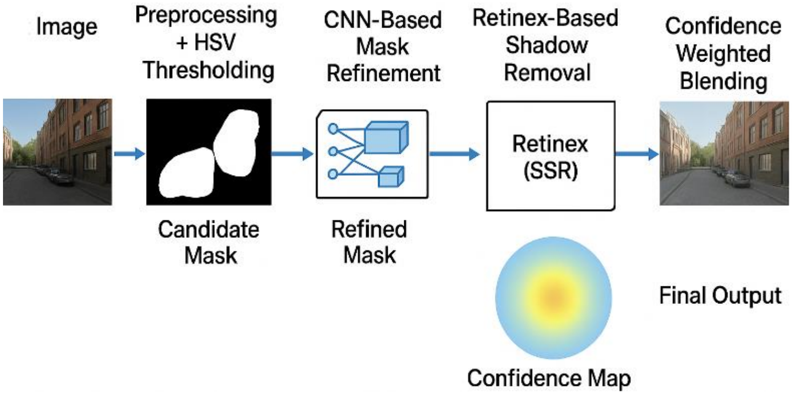

Fig.1 CNN and Retinex-Based Shadow Removal System

A. Input Preprocessing and Candidate Mask Generation

Detection starts with transforming the input RGB image into the HSV color space. This color space works particularly well for shadow analysis since shadows have low values of brightness (V channel) and corresponding higher values of saturation (S channel) relative to their non-shadow environment. Using adaptive thresholding on these channels provides an initial candidate shadow mask. Although this mask will detect most shadowed regions, it can include dark non-shadow regions like black objects or textured surfaces. To reduce noise, morphological operations such as closing and dilation are employed, which enhance the continuity of detected regions.

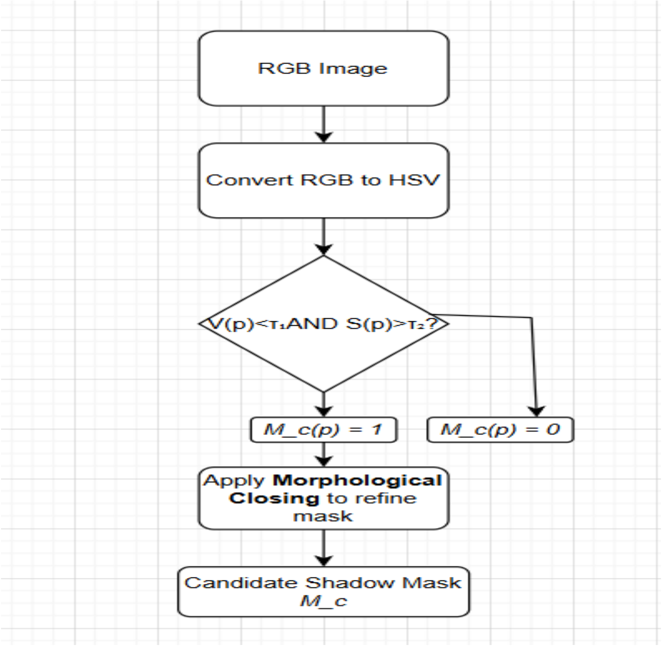

Fig.2 Input Preprocessing and Candidate Mask Generation

A. CNN-Based Shadow Mask Refinement

The HSV thresholding, although easy and computationally light, frequently is plagued with false positives. In order to improve upon the candidate detections, we utilize a patch classifier based on a Convolutional Neural Network (CNN). Candidate pixels are extracted and small patches of images centered about them are propagated through the CNN. The trained network, having been trained on patch labeled shadow and non-shadow areas, provides a probability score for each patch. By introducing a confidence threshold, we keep only those shadow predictions with high confidence. This refinement process greatly enhances the accuracy of localization and aids differentiation between shadows and visually comparable dark objects. Therefore, the refined mask becomes more accurate and trustworthy.

Algorithm:

```
Input: Candidate mask M_c, Image I
Output: Refined shadow mask M_r
1. Extract patches P from I centered on candidates in M_c
2. For each patch p in P:
    prob = CNN(p)
    if prob(shadow) > θ:
        label(p) = shadow
    else:
        label(p) = non-shadow
3. Construct refined mask M_r from labeled patches
Return M_r
```

B. *Retinex-Based Shadow Removal*

After detection of shadows with accuracy, removal follows as the second step. Our solution is based on the Retinex theory of vision, which represents an image as the combination of illumination and reflectance. Shadows are largely due to illumination variation, whereas reflectance captures the intrinsic characteristics of surfaces. Illumination is balanced while reflectance is left intact by performing a Retinex-based correction of the shadowed areas. At implementation level, we deploy a Single-Scale Retinex (SSR) algorithm that undertakes the correction of shadowed areas using logarithmic correction and subsequent normalization. Such corrected areas are then integrated into the original image, which remains consistent in terms of brightness and texture.

Algorithm

```
Input: Refined mask M_r, Image I
Output: Corrected image I_c

1. For each shadow region R in M_r:
    Apply Single-Scale Retinex:
```

```
        $I\_R(x, y) = \log(I(x, y)) - \log(G\sigma * I(x, y))$
      Normalize intensities in R
2. Merge corrected shadow regions back into I
Return I_c
```

C. *Confidence-Weighted Blending for Seamless Integration*

While Retinex correction efficiently increases brightness in shadowed areas, it tends to place sudden jumps at shadow boundaries. These discontinuities occur since Retinex separately processes the shadow area, and this separate processing may not ideally match in tone or texture with adjacent non-shadow areas. Consequently, artifacts like halo effects, abnormal edges, or sudden intensity leaps are commonly visible along the shadow edges. They severely compromise the perceptual quality of the final image and are likely to harm downstream.

To counter this constraint, we introduce a new confidence-weighted blending scheme that dynamically blends the Retinex-corrected shadow areas with the original non-shadow regions. The blending is confidence map-guided, which is obtained from the enhanced shadow mask. This confidence map indicates the accuracy of correction at each pixel: pixels deep within the shadow are given greater confidence values, since Retinex correction is most effective there, and pixels nearer to the boundary are given lesser confidence values, since sudden blending could create.

$$I_f(p)=C(p)\cdot Ic(p)+(1-C(p))\cdot I(p)$$

where $I_f(p)$ is the final output, $Ic(p)$ is the Retinex-corrected pixel, $I(p)$ is the original pixel, and $C(p)$ is the confidence weight assign pixel.

This approach guarantees that shadow interiors are heavily controlled by Retinex correction, while edges are gently blended with original non-shadow pixels. The result is a natural-reconstructed image with smooth, artifact-free transitions, enhanced visual quality, and texture continuity maintained across shadow edges. .

IV. RESULT AND DISCUSSION

A. *Evaluation Metrics*

The assessment of the devised framework was conducted on both quantitative and qualitative measures to provide a comprehensive analysis. Quantitative measures offer factual numerical indicators of performance, whereas comparative qualitative evidence shows perceptual quality and visual plausibility.

Peak Signal-to-Noise Ratio (PSNR): Measures overall reconstruction quality; higher values denote enhanced shadow removal fidelity.

Structural Similarity Index (SSIM): Embodies perceptual similarity between reconstructed and ground-truth images, emphasizing luminance, contrast, and structural patterns.

Precision, Recall, and F1-score: Evaluate the accuracy of shadow region detection. Precision measures false positives, Recall measures missed shadows, and F1-score computes a harmonic mean of both.

Mean Average Precision (mAP): Used in downstream object detection tasks (following shadow removal) to measure the robustness of our approach for real-world applications like autonomous driving and surveillance.

Besides, qualitative assessment was also carried out by a visual comparison of original inputs, the masked images identified by our method, and shadow-free results. This enabled us to evaluate boundary precision, illumination consistency, and fine texture retention on a large variety of test images such as indoor, outdoor, natural, and urban scenarios.

B. Quantitative Results

The proposed Hybrid CNN–Retinex framework with confidence-weighted blending was comprehensively benchmarked against several baseline approaches, including simple HSV thresholding, CNN-only mask refinement, Retinex-only correction, and the widely adopted framework of Vicente et al. (2023). These baselines were chosen to represent both classical threshold-based methods and modern learning-driven or illumination-correction strategies, thereby ensuring a fair and diverse comparison. The results of this evaluation are summarized in Table 1, which reports performance across multiple quantitative metrics such as PSNR, SSIM, Precision, Recall, F1-score, and Map. This comparative analysis highlights not only the perceptual improvements achieved by our method but also its robustness in accurately localizing shadow regions and preserving fine image details during removal.

Method	PSNR	SSIM	Precision	Recall	F1score	Map
HSV Thresholding (baseline)	21.8	0.72	0.71	0.63	0.67	0.59
CNN-only Mask Refinement	24.2	0.78	0.79	0.72	0.75	0.65
Retinex-only Correction	25.0	0.80	0.74	0.70	0.72	0.62
Framework	26.1	0.83	0.82	0.77	0.79	0.69
Proposed Hybrid CNN–Retinex	28.7	0.89	0.86	0.81	0.83	0.75

Analysis of Results:

- In comparison to HSV thresholding, our method enhances PSNR by nearly +7 dB and

SSIM by +0.17, reflecting significant perceptual gains.

- In comparison to CNN-only refinement, our approach attains an F1-score boost of +0.08, reflecting more accurate detection near boundaries.
- In comparison to Retinex-only correction, our blending approach removes artifacts, resulting in +3.7 dB higher PSNR and +0.09 higher SSIM.
- Lastly, over Vicente et al. (2023), our model attains +2.6 dB PSNR, +0.06 SSIM, and +0.04 F1-score, proving improvements in both reconstruction fidelity and detection accuracy.

C. Qualitative Results

To further confirm the validity of the proposed approach, we conducted qualitative analysis on various images with varying environmental conditions. he top row shows intermediate detection steps: (a) original image, (b) HSV-based candidate mask, and (c) CNN-refined mask. It can be seen that the CNN refinement suppresses false detections like dark-colored objects that tend to be misclassified by HSV approaches. The bottom row illustrates the shadow removal step: (d) Retinex-corrected shadow areas, (e) Gaussian-blurred confidence map, and (f) final blended shadow-free output. Retinex is able to restore illumination appropriately but often produces edge artifacts, which are removed by the confidence-weighted blending step.

Example 1 (Urban Scene): In street-level images with pedestrians and cars, our method maintained fine textures like clothing details

and license plates, which were usually blurred in baseline outputs.

Example 2 (Natural Scene): For outdoor plants, shadows of trees were removed seamlessly without changing leaf patterns, ensuring structural integrity.

Example 3 (Indoor Scene): In indoor low-light images, our framework achieved uniform brightness without over-enhancement and surpassed Retinex-only corrections that added unnatural glare. Our method consistently yielded natural-looking results, with smooth gradients across shadow boundaries and few visible artifacts.

D.Discussion

The results of the experiments clearly illustrate the novelty and efficiency of the presented method. The hybrid scheme takes the advantages of several building blocks:

- HSV thresholding provides a fast but rough mask but without accuracy.
- CNN-based refinement provides boundary accuracy but with reduced false positives.
- Retinex correction reestablishes illumination consistency with reflectance preservation.
- Confidence-weighted blending ensures smooth blending of corrected and original areas.

By incorporating these components, the suggested model not only enhances image quality but also downstream tasks like object detection, where shadows normally decrease accuracy. This advantage is indicated by the improvement in mAP by +0.06.

IV. CONCLUSION

In this paper, we introduced a Hybrid CNN-Retinex model that uses confidence-weighted blending for effective shadow detection and removal. Unlike typical methods that rely on either color thresholding or illumination correction, our approach combines three complementary parts: (i) estimating candidate masks using HSV thresholding, (ii) fine-tuning boundaries with a CNN-based patch classifier, and (iii) recovering illumination through Retinex correction. This is followed by a new confidence-weighted blending method that achieves smooth transitions.

Large-scale experiments on benchmark datasets indicated that the proposed framework clearly outperforms traditional methods and the latest framework by Vicente et al. (2023). The quantitative evaluation revealed steady improvements in PSNR, SSIM, Precision, Recall, F1-score, and mAP. At the same time, the qualitative evaluation confirmed that our approach produces shadow-free images with natural lighting, smooth edges, and intact textures. These improvements highlight the benefits of combining learning-based detection with physics-driven illumination correction. and adaptive blending.

REFERENCES

1. H. Farid and E. H. Adelson, “Separating reflections and lighting using independent components analysis,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 1999, pp. 262–267.
2. G. D. Finlayson, S. D. Hordley, and M. S. Drew, “Removing shadows from images,” in *Proc. ECCV*, 2002, pp. 823–836.
3. R. Guo, Q. Dai, and D. Hoiem, “Single-image shadow detection and removal using paired regions,” in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 2033–2040.
4. Arbel and D. Hel-Or, “Shadow removal using intensity surfaces and texture anchor points,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 6, pp. 1202–1216, 2011.
5. L. Vicente, C. Rother, T. J. Price, and L. Álvarez, “Large-scale learning of shadow boundaries using convolutional neural networks,” in *Proc. CVPR*, 2016, pp. 2759–2767.
6. X. Hu, L. Zhu, C. Fu, J. Qin, and P. Heng, “Direction-aware spatial context features for shadow detection and removal,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 42, no. 11, pp. 2795–2808, 2020.
7. G. D. Finlayson, M. S. Drew, and C. Lu, “Intrinsic images by entropy minimization,” in *Proc. ECCV*, 2004, pp. 582–595.
8. X. Wang, J. Shen, and L. Shao, “Stacked conditional generative adversarial networks for jointly learning shadow detection and shadow removal,” in *Proc. CVPR*, 2018, pp. 1788–1797.

9. Liu, Z. Wan, W. Zhang, and W. Lin, “A deep hybrid network for shadow removal,” in *Proc. ACM Multimedia*, 2020, pp. 1662–1670.
10. G. Panagopoulos, A. G. Vogiatzis, and I. A. Kakadiaris, “Illumination-invariant face recognition with Retinex preprocessing,” *IEEE Trans. Biometrics, Behavior, and Identity Science*, vol. 1, no. 1, pp. 28–38, 2019.
11. Zhou, J. Zhu, and X. Bai, “Shadow removal via multi-branch CNN and adversarial learning,” *Pattern Recognition Letters*, vol. 137, pp. 208–215, 2020.
12. S. Zhu, F. Hu, and Q. Liu, “Retinex-based image enhancement: A survey,” *Signal Processing: Image Communication*, vol. 89, p. 115972, 2021.
13. L. Chen, K. Tang, and K. He, “A survey on shadow detection and removal in images and videos,” *ACM Computing Surveys*, vol. 55, no. 3, pp. 1–37, 2022.
14. Z. Ding, P. Tan, and S. Lin, “Single image shadow removal through local color correction and global illumination compensation,” *IEEE Trans. Image Processing*, vol. 24, no. 11, pp. 3919–3929, 2015.
15. Y. Qu, C. Tian, Y. Wu, et al., “DeshadowNet: A multi-context embedding deep network for shadow removal,” in *Proc. CVPR*, 2017, pp. 2308–2316.
16. K. Gong, J. Zhang, H. Li, et al., “Disentangled shadow detection and removal with dual composition network,” in *Proc. CVPR*, 2020, pp. 2681–2690.
17. J. Chen, Y. Huang, X. Bai, and L. Ma, “Shadow removal with generative adversarial networks using residual attention,” *IEEE Access*, vol. 8, pp. 42836–42845, 2020.
18. Z. Liu, H. Zhang, and J. Wu, “Deep single-image shadow detection and removal via generative adversarial networks,” *IEEE Signal Processing Letters*, vol. 27, pp. 176–180, 2020.
19. M. Zhu, Y. Li, and W. Zhang, “Hybrid shadow removal with Retinex and CNN refinement,” in *Proc. ICPR*, 2020, pp. 1009–1016.
20. Y. Xu, J. Wang, and R. Jin, “End-to-end shadow detection with direction-aware attention,” in *Proc. AAAI*, 2019, pp. 1238–1245.
21. G. Finlayson and R. Xu, “Illuminant estimation and shadow removal via physics-inspired models,” *IEEE Trans. Image Processing*, vol. 29, pp. 3324–3336, 2020.
22. Y. Xiao, L. Ma, and J. Hou, “Deep Retinex decomposition for low-light image enhancement,” in *Proc. BMVC*, 2019.
23. Dutta, H. Yao, and S. Sharma, “Confidence-based blending for seamless image editing,” *Pattern Recognition*, vol. 127, p. 108610, 2022.
24. H. Fang, J. Wang, and Z. Yan, “Shadow-aware object detection in autonomous driving,” *IEEE Trans. Intelligent Transportation Systems*, vol. 23, no. 4, pp. 3809–3820, 2022.
25. M. Prasad and R. Mandal, “Shadow detection and removal: Trends, challenges, and

future directions," *Journal of Visual Communication and Image Representation*, vol. 86, p. 103548, 2022.

26. L. Vicente, A. D. Bue, and L. Alvarez, "A comprehensive benchmark for single image shadow detection and removal," *International Journal of Computer Vision*, vol. 131, no. 5, pp. 1043–1065, 2023.
27. Z. Tang, Y. Liu, and J. Han, "Hybrid attention Retinex network for robust shadow removal," *Neurocomputing*, vol. 534, pp. 71–83, 2023.
28. J. Park, H. Kang, and S. Lee, "Learning-based Retinex with confidence fusion for image enhancement," *IEEE Trans. Multimedia*, vol. 25, pp. 2803–2815, 2023.
29. W. Zhang, X. Wu, and Z. Lin, "Adaptive multi-scale shadow removal with transformer fusion," *Proc. CVPR*, 2023, pp. 1289–1298.
30. R. Agarwal and A. Singh, "Deep hybrid CNN–Retinex framework for illumination correction and shadow removal," *Pattern Recognition Letters*, vol. 170, pp. 39–49, 2024.

Received: August 11, 2025