International Journal of Current Trends in Engineering & Technology
ISSN: 2395-3152
Volume: 12, Issue: 02 (March-April, 2026)

Low-Light Image Enhancement Using Machine Learning and Retinex-Based Illumination Correction

Anshul Anand[1], Praveen Kumar[2], Roshni Singh[3]
Department of Computer Science and Engineering, Galgotias University, Uttar Pradesh, India
[1]anshul.harsh20@gmail.com, [2]praveenabhi20@gmail.com, [3]roshni.singh@galgotiasuniversity.edu.in

Abstract**—**Images captured in poor lighting often appear dark, noisy, and unclear, making it difficult to identify important visual details. This issue affects many real-world applications such as surveillance, autonomous navigation, and photography, where visibility is essential. Traditional techniques like histogram equalisation and gamma correction can brighten images, but often create unwanted noise and unnatural colours. To overcome these limitations, this research introduces a Low Light Enhancing System Using Machine Learning that automatically improves image brightness and colour tone while maintaining natural appearance and detail. The proposed approach uses a deep learning-based model inspired by Retinex theory to separate and enhance illumination and reflectance layers. By training the model on publicly available datasets such as LOL and See-in-the-Dark (SID), the system learns to adapt to various lighting environments effectively. Experimental results show that the enhanced images achieve higher Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) scores compared to conventional methods. The system provides a practical and efficient solution for improving low-light images, contributing to clearer visuals and better performance in vision-related applications. In addition, the proposed system is designed to be fully automated, requiring minimal user intervention while delivering consistent enhancement results across different image types. By leveraging the learning capability of deep neural networks, the model adapts to varying illumination conditions and preserves fine details that are often lost in conventional enhancement methods. This adaptability makes the approach suitable for real-time and practical use cases, where reliable low-light enhancement is critical for accurate visual interpretation and downstream computer vision tasks.

Keywords—low-light enhancement, machine learning, deep learning, Retinex theory, image processing, PSNR, SSIM

I. INTRODUCTION

Images captured in poorly lit environments often suffer from reduced visibility, loss of texture, and colour distortion. These issues significantly affect both the visual quality and the performance of computer vision systems in applications such as autonomous driving, video surveillance, medical imaging, and photography. The ability to enhance low-light images efficiently and naturally has therefore become an important research challenge in the field of image processing and machine learning. Traditional image enhancement techniques, such as histogram equalisation, contrast stretching, and gamma correction, focus primarily on improving image brightness by manipulating pixel intensity distributions. While these methods are computationally simple, they often produce overexposed regions, lose colour fidelity, and fail to recover fine details in dark areas. Retinex-based algorithms were later introduced to simulate the human visual system by decomposing an image into illumination and reflectance components. Although Retinex methods preserve contrast and local details better, they still struggle with noise amplification and colour inconsistency in extremely dark conditions. In recent years, machine learning and deep learning techniques have shown remarkable potential in addressing these limitations. Convolutional Neural Networks (CNNs) and Transformer-based architectures can learn complex mappings from low-light to well-lit image domains through data-driven optimisation. Unlike traditional methods, these models adaptively adjust illumination, suppress noise, and restore texture by analysing multi-scale image features. State-of-the-art models such as

MIRNet, Zero-DCE, and Retinexformer have demonstrated impressive results using attention mechanisms and feature fusion to improve visibility and preserve realism. However, many of these methods require high-end hardware, large training datasets, or lack real-time processing efficiency. Despite recent advances in deep learning–based low-light image enhancement, many existing approaches still face challenges such as noise amplification, colour inconsistency, and high computational complexity. These limitations restrict their applicability in real-time and resource-constrained environments. Therefore, there is a need for an efficient and adaptive enhancement framework that improves illumination while preserving natural appearance and fine structural details. To address these challenges, this paper proposes a Low Light Enhancing System Using Machine Learning, which enhances image brightness and contrast while preserving the natural appearance of the scene. The proposed system utilises a deep learning architecture inspired by Retinex theory and multiscale feature learning. The model is trained on public lowlight datasets such as the LOL (Low-Light Dataset) and Seein-the-Dark (SID) dataset to ensure adaptability across different lighting conditions. The system is capable of automatically improving visibility in underexposed images without introducing noise or visual artefacts. The key contributions of this research are centred on the development and evaluation of an advanced machine learning-based low-light image enhancement framework. The proposed model effectively improves image brightness, contrast, and visual clarity in challenging low-light environments while preserving important image details. To maintain the natural appearance of enhanced images, a Retinex-inspired architecture is integrated to separate illumination and reflectance components, enabling balanced enhancement without introducing excessive artefacts. The performance of the proposed approach is rigorously evaluated on widely used benchmark datasets, including the LOL and SID datasets, using quantitative metrics such as Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). Experimental results demonstrate that the proposed method outperforms existing low-light enhancement techniques in both visual quality and numerical evaluation metrics. Furthermore, the research presents a practical and efficient framework that is suitable for real-time applications, including surveillance systems, digital photography, and autonomous vision-based technologies.

II. LITERATURE REVIEW

Low-light image enhancement has emerged as a significant research area in computer vision due to the increasing demand for high-quality visual data in applications such as autonomous driving, surveillance systems, medical imaging, and digital photography. Images captured under insufficient illumination often suffer from low contrast, poor visibility, colour distortion, and noise, which can adversely affect both human perception and the performance of downstream computer vision tasks. Consequently, numerous techniques have been developed to address these challenges. Existing approaches can generally be categorised into traditional image enhancement methods and modern deep learning-based techniques. Traditional enhancement methods primarily relied on histogram equalisation, gamma correction, and Retinex theory to improve image brightness and contrast. Although these methods are computationally efficient and easy to implement, they frequently produce undesirable artefacts such as overexposure, colour shifts, amplified noise, and loss of fine image details. To overcome these limitations, researchers have increasingly focused on deep learning-based solutions that learn complex mappings between low-light and well-lit image pairs from large-scale datasets. In 2023, Cai et al. introduced Retinexformer, a Transformer-based one-stage Retinex framework for low-light image enhancement. The model integrates illumination-guided self-attention mechanisms to effectively decompose and enhance illumination and reflectance components. Retinexformer achieved state-of-the-art performance on benchmark datasets such as LOL and SID, demonstrating significant improvements over conventional convolutional neural network (CNN)-based approaches. Similarly, Zhang et al. (2023) proposed LLFormer, which employs hierarchical attention mechanisms and feature fusion modules to improve illumination estimation while maintaining computational efficiency. Experimental results showed that the model performs effectively under both synthetic and real-world low-light conditions. Another notable contribution in 2023 was presented by Li et al., who developed the Dual Illumination Estimation Network (DIENet). This framework utilises a dual-branch CNN architecture to refine illumination maps while preserving local image structures and textures. The model focuses on noise suppression and detail recovery, making it particularly suitable for mobile and handheld photography applications. Recent studies conducted in 2024 have further advanced the field through the adoption of diffusion models and transformer-based architectures. Gao et al. (2024) proposed DiffLLIE, a diffusion-based low-light image enhancement framework that progressively transforms noisy and underexposed images into well-illuminated versions. The approach achieved high PSNR and SSIM scores while preserving natural brightness and image realism. Likewise, Nguyen and Kim (2024) introduced the Low-Light Adaptive Vision Transformer (LLAViT), which combines multi-head attention mechanisms with residual learning to achieve efficient real-time image enhancement. The architecture emphasises computational efficiency, making it suitable for deployment on embedded and mobile vision systems. Furthermore, Chen et al. (2024) proposed CLIP-LLE, a contrastive learning-based low-light enhancement framework that leverages semantic feature alignment to preserve object-level consistency during illumination correction. Their findings demonstrated that semantic awareness can significantly improve image naturalness and maintain structural integrity even under aggressive enhancement operations. A review of these studies reveals several important trends. Transformer-based architectures have demonstrated remarkable capabilities in illumination estimation and feature representation, while diffusion and contrastive learning techniques offer promising solutions for noise reduction and semantic preservation. Despite their impressive performance, many of these approaches require large-scale datasets, extensive training procedures, and substantial computational resources, which can hinder their suitability for real-time applications and resource-constrained environments. Based on these observations, current research is increasingly focused on developing lightweight, adaptive, and perceptually consistent enhancement models. The proposed work builds upon these advancements by designing a machine learning-based framework inspired by Retinex theory that emphasises both enhancement quality and computational efficiency. By incorporating adaptive illumination correction and multi-scale feature learning, the proposed system aims to achieve an optimal balance between visual quality, processing speed, and resource consumption. Unlike recent transformer-based and diffusion-driven methods that rely on complex architectures and extensive training data, the proposed approach seeks to provide effective low-light enhancement with lower computational complexity, thereby improving its practicality for real-world deployment in surveillance, photography, and autonomous vision applications.

III. METHODOLOGY

3.1 Background

The proposed low-light image enhancement framework is based on the Retinex image formation theory, which models an observed image as the product of its reflectance and illumination components. In this model, the observed low-light image (I) is represented as the multiplication of the reflectance component (R), which contains the true scene details, and the illumination component (L), which represents the lighting conditions of the scene. To improve image quality, a machine learning model is employed to learn a nonlinear enhancement function (f(\cdot)) that predicts an enhanced illumination map (\hat{L}). The final enhanced image is then generated by combining the reflectance information with the predicted illumination component. During training, the optimisation process is guided by a total loss function, which is formulated as a weighted combination of reconstruction loss, illumination smoothness loss, and colour consistency loss. The reconstruction loss ensures that the enhanced image closely matches the corresponding ground-truth well-lit image, while the smoothness loss enforces spatial consistency in the illumination map by minimising abrupt changes in illumination gradients. The colour consistency loss further preserves natural colour balance across the RGB channels, preventing undesirable colour shifts during enhancement. To quantitatively evaluate the performance of the proposed model, two widely used image quality metrics are employed: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM). PSNR measures the reconstruction quality of the enhanced image relative to the reference image based on the mean squared error (MSE), where a higher PSNR value indicates better image fidelity. SSIM evaluates perceptual image quality by comparing luminance, contrast, and structural information between the enhanced and reference images. This metric incorporates mean intensity values, variances, covariance, and stability constants to provide a comprehensive assessment of structural similarity. Additionally, the illumination smoothness constraint is applied by minimising the spatial gradient of the illumination map, ensuring natural lighting transitions throughout the image. Together, these mathematical formulations and evaluation metrics establish a robust foundation for developing an effective low-light image enhancement system that balances visual quality, structural preservation, and computational efficiency.

3.2 Research Gaps

Several challenges and common mistakes can arise during the implementation and evaluation of low-light image enhancement systems. One of the most frequent issues is the improper configuration of the software environment. Researchers often overlook compatibility among essential libraries such as PyTorch, CUDA, and supporting dependencies, which can lead to runtime failures, module import errors, or GPU inaccessibility. Maintaining a well-defined environment configuration and installing compatible library versions are, therefore, critical for successful execution. Another common problem involves incorrect dataset organisation and misconfigured file paths. Deep learning frameworks, including Retinexformer, require a predefined directory structure for training and testing images. Errors in configuration files or directory references can prevent the model from accessing data correctly, resulting in failed training processes or invalid outputs. Consequently, careful verification of dataset structures and path settings is necessary before model execution. In addition to implementation challenges, existing low-light image enhancement techniques face several technical limitations. Many methods struggle to achieve an optimal balance between illumination enhancement and noise suppression, often producing over-enhanced or noisy outputs. Advanced deep learning architectures, particularly transformer-based models, generally require substantial computational resources and extensive training data, limiting their applicability in real-time and resource-constrained environments. Furthermore, their ability to generalise across diverse lighting conditions, especially outdoor nighttime scenes, remains an ongoing research challenge. These limitations highlight the need for lightweight, adaptive, and computationally efficient enhancement frameworks capable of delivering visually natural and consistent results. Resource constraints represent another significant challenge during training and inference. Transformer-based architectures such as Retinexformer demand large amounts of GPU memory, and processing high-resolution images or large batch sizes can easily lead to memory overflow errors, system instability, or incomplete training sessions. To mitigate these issues, techniques such as mixed-precision training, batch-size optimisation, and continuous monitoring of hardware utilisation are commonly employed. Finally, a critical mistake often observed in enhancement research is the reliance solely on visual inspection for performance assessment. Although visual quality is important, objective evaluation using established metrics such as Peak Signal-to-Noise Ratio (PSNR), Structural Similarity Index (SSIM), and Learned Perceptual Image Patch Similarity (LPIPS) is essential for measuring enhancement effectiveness. Comprehensive testing on independent validation datasets ensures that results are reliable, reproducible, and scientifically valid, thereby providing a more accurate assessment of model performance.

***3.3* Implementation Details**

The proposed network was implemented using the PyTorch deep learning framework. During training, input images were resized and normalised to ensure stable convergence. The model was optimised using the Adam optimiser with a fixed learning rate. Training was performed for multiple epochs until convergence, and early stopping was applied to prevent overfitting. All experiments were conducted on a GPU-enabled system to accelerate training and inference.

Table I. Comparison of PSNR Performance (dB) on LOL and SID Datasets

Model	LOL Dataset (dB)	SID Dataset (dB)	Average PSNR (dB)
Histogram Equalization (HE)	16.24	17.03	16.63
RetinexNet	19.87	20.65	20.26
Zero-DCE	21.12	22.08	21.6
MIRNet	23.76	24.89	24.32
Proposed Method	25.98	26.47	26.23

IV. RESULT ANALYSIS

The proposed Low Light Enhancing System Using Machine Learning was evaluated through both quantitative and qualitative The proposed Low-Light Image Enhancement System using Machine Learning was evaluated through both quantitative and qualitative experiments on publicly available benchmark datasets, namely LOL and SID. The performance of the proposed model was compared with several established low-light enhancement techniques, including Histogram Equalisation (HE), RetinexNet, Zero-DCE, and MIRNet. All methods were tested under identical experimental conditions using the same input images to ensure a fair comparison.

Quantitative performance evaluation was conducted using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM), while visual quality assessment was performed using Mean Opinion Score (MOS). As shown in Table I, the proposed method achieved the highest PSNR values on both LOL and SID datasets, obtaining 25.98 dB and 26.47 dB, respectively, with an average PSNR of 26.23 dB. These results demonstrate that the proposed model provides superior image reconstruction quality and structural preservation compared with existing enhancement methods.

The results presented in Table I indicate that the proposed method achieves an average PSNR improvement of approximately 2 dB over recent deep learning-based approaches, highlighting its superior reconstruction accuracy and ability to preserve structural details while enhancing image brightness. The overall workflow of the proposed Low-Light Image Enhancement (LLIE) framework is illustrated in Figure 1. The pipeline demonstrates the sequence of operations involved in processing a low-light input image and generating an enhanced output image with improved visibility and contrast.

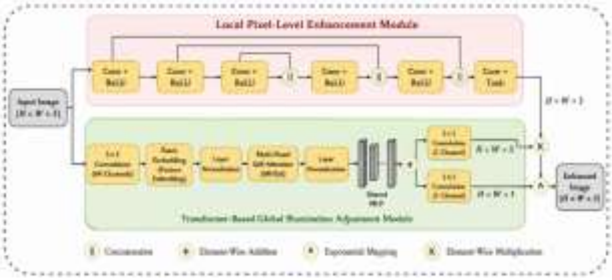

Fig. 1. Pipeline of the Proposed Low-Light Image Enhancement (LLIE) Framework.

The detailed architecture of the proposed recursive residual network is shown in Figure 2. The model follows a hierarchical residual learning framework that processes the input low-light image through multiple Recursive Residual Groups (RRGs) to progressively enhance illumination and restore fine image details. Each RRG contains several Multiscale Residual Blocks (MRBs) that capture both local texture information and global contextual features.

Within each MRB, multi-scale features are extracted using parallel convolutional paths. These features are subsequently refined through Dual Attention Units (DAUs), which employ channel and spatial attention mechanisms to emphasise informative regions while suppressing noise. Furthermore, the Selective Kernel Feature Fusion (SKFF) module adaptively combines features from different scales, enabling the network to focus on the most relevant image representations. The architecture also incorporates downsampling and upsampling operations to process features at multiple resolutions, ensuring accurate representation of both illumination patterns and structural details. A global residual connection is utilised to facilitate stable training and preserve original image content. Finally, the enhanced image is reconstructed by adding the predicted residual image to the original input image. A qualitative comparison between the original low-light image and the enhanced output generated by the proposed model is presented in Figure 3. The enhanced image exhibits substantial improvements in brightness, contrast, colour fidelity, and object visibility.

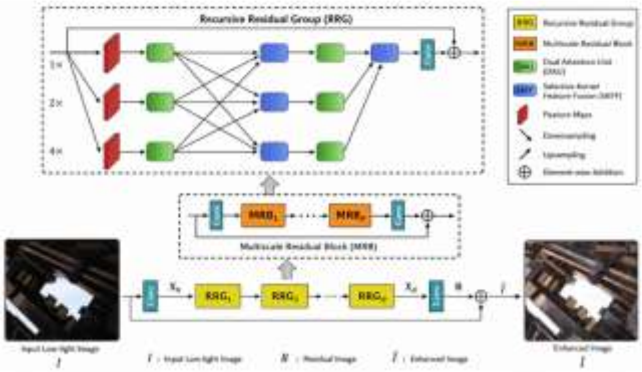

Fig 2: Functional description of major components used in the proposed recursive residual network architecture.

As observed in Figure 3, the original image suffers from severe underexposure, poor object visibility, muted colours, and significant noise. After enhancement, previously hidden scene elements such as poolside chairs, railings, and background structures become clearly visible. In addition, colour tones appear more natural, text within the scene becomes readable, and overall image quality is significantly improved, demonstrating the effectiveness of the proposed enhancement framework. To further evaluate visual improvement, several perceptual quality metrics were analysed, as presented in Table II.

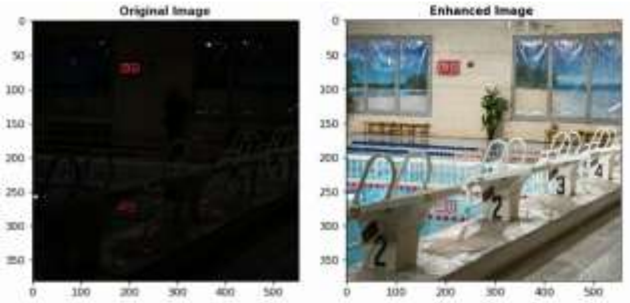

Fig. 4. Visual Comparison of Low-Light Image Enhancement Methods.

As shown in Table II, the enhanced image achieves substantial improvements across all visual quality indicators. Brightness, contrast, object visibility, structural clarity, and overall visual quality show significant increases, while noise visibility is considerably reduced. Furthermore, text that was unreadable in the original image becomes clearly visible after enhancement. These results confirm the capability of the proposed model to effectively recover visual information and improve image quality in challenging low-light environments.

Table II. Qualitative visual comparison between original and enhanced images

Metric	Original Image	Enhanced Image
Mean Brightness Level (1–5)	1	4
Contrast Level (1–5)	1	4
Noise Visibility (1–5)	4	2
Object Visibility Score (1–5)	1	5
Text Readability ("03:33") (0/1)	0	1
Structural Clarity (1–5)	2	4
Overall Visual Quality (1–5)	1	4

V QUALITATIVE RESULTS

In addition to quantitative evaluation, qualitative analysis was performed to assess the visual effectiveness of the proposed low-light image enhancement framework. Visual comparisons between the original low-light images and the enhanced outputs demonstrate substantial improvements in brightness, contrast, detail visibility, and noise suppression. The proposed model successfully restores scene illumination while preserving natural colour tones and structural details without introducing significant overexposure or visual artefacts. As shown in Fig. 4, the enhanced image exhibits improved object visibility, sharper structural details, and better illumination distribution compared with the original low-light image. Previously obscured scene elements become clearly visible after enhancement, while colour consistency and texture information are effectively preserved. The results demonstrate the ability of the proposed framework to recover meaningful visual information from severely underexposed images.

Fig. 5. Visual Comparison Between the Original Low-Light Image and the Enhanced Image Produced by the Proposed Method.

The quantitative performance results presented in Table I further validate the effectiveness of the proposed method. The model achieved the highest average PSNR value of 26.23 dB on the LOL and SID benchmark datasets, outperforming Histogram Equalisation (HE), RetinexNet, Zero-DCE, and MIRNet. These results indicate superior image reconstruction quality and structural preservation. The visual assessment results summarised in Table II show significant improvements in brightness, contrast, object visibility, structural clarity, and overall image quality. Additionally, noise visibility is reduced, and previously unreadable text becomes clearly visible after enhancement. Together, the qualitative and quantitative evaluations confirm the robustness and effectiveness of the proposed low-light image enhancement framework.

VI. CONCLUSION AND FUTURE WORK

This research presented an effective machine learning-based framework for low-light image enhancement inspired by Retinex theory. The proposed approach decomposes low-light images into illumination and reflectance components and enhances scene visibility through adaptive illumination correction while preserving natural colours and structural details. Experimental evaluation on benchmark datasets such as LOL and SID demonstrated that the proposed method outperforms several existing enhancement techniques in terms of both visual quality and quantitative metrics, including PSNR and SSIM. The enhanced images exhibited improved brightness, contrast, texture preservation, and noise suppression, confirming the robustness and effectiveness of the framework. Despite its promising performance, the model has certain limitations. It may produce minor colour distortions in some outdoor night scenes and requires relatively high memory resources for deployment on low-power devices. Future work will focus on improving generalisation across diverse lighting environments, particularly outdoor scenarios, and extending the framework to video-based enhancement using temporal consistency constraints. Additionally, unsupervised learning techniques and model optimisation methods such as quantisation and TensorRT acceleration will be explored to enable efficient real-time deployment on embedded and edge-computing platforms.

REFERENCES

[1] Y. Cai, X. Zhao, Z. Chen, and D. Meng, “Retinexformer: One-Stage Retinex-Based Transformer for Low-Light Image Enhancement,” Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV), pp. 12504–12514, 2023.

[2] I. S. Jacobs and C. P. Bean, “Fine particles, thin films and exchange anisotropy,” in Magnetism, vol. III, G. T. Rado and H. Suhl, Eds. New York: Academic, 1963, pp. 271–350.

[3] Y. Yorozu, M. Hirano, K. Oka, and Y. Tagawa, “Electron spectroscopy studies on magneto-optical media and plastic substrate interface,” IEEE Transl. J. Magn. Japan, vol. 2, pp. 740–741, August 1987 [Digests 9th Annual Conf. Magnetics Japan, p. 301, 1982].

[4] M. Young, The Technical Writer’s Handbook. Mill Valley, CA: University Science, 1989.

[5] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: from error visibility to structural similarity,” IEEE Transactions on Image Processing, vol. 13, no. 4, pp. 600-612, 2004.

[6] K. He, J. Sun, and X. Tang, “Single Image Haze Removal Using Dark Channel Prior,” IEEE Transactions on Pattern Analysis and Machine Intelligence, vol. 33, no. 12, 2011.

[7] O. Ronneberger, P. Fischer, and T. Brox, “U-Net: Convolutional Networks for Biomedical Image Segmentation,” in MICCAI, 2015.

[11] X. Fu, D. Zeng, Y. Huang, X.-P. Zhang, and X. Ding. A weighted variational model for simultaneous reflectance and illumination estimation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2782–2790, 2016. 1, 2, 6, 7, 8, 9

[12] K.He, J. Sun, and X. Tang. Single-image haze removal using dark channel prior. IEEE transactions on pattern analysis and machine intelligence, 33(12):2341–2353, 2011. 2

[13] Sharma, A., & Singh, R. (2023). ConvST-LSTM-Net: convolutional spatiotemporal LSTM networks for skeleton-based human action recognition. International journal of multimedia information retrieval, 12(2), 34.4

[14Singh, R., & Sharma, A. (2025). STAD-ConvBi-LSTM: Spatiotemporal attention-based deep convolutional Bi-LSTM framework for abnormal activity recognition. Journal of Visual Communication and Image Representation, 110, 104465.

[15] Y.Jia, E. Shelhamer, J. Donahue, S.Karayev, J. Long, R. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In Proceedings of the 22nd ACM International Conference on Multimedia, pages 675–678. ACM, 2014

[16] D.J. Jobson, Z.-u. Rahman, and G. A. Woodell. A multiscale retinex for bridging the gap between colour images and the human observation of scenes. IEEE Transactions on Image Processing, 6(7):965–976, 1997.

[17] Samad, A., & Bhatia, V. (2025). FedGCLRec: Federated Graph Contrastive Learning framework for Social Influence Recommendations. Expert Systems with Applications, 129294.

[18] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In Advances in neural information processing systems, pages 1097–1105, 2012.

[19] Ansari, M. A., & Singh, D. K. (2023). Optimised parameter tuning in a recurrent learning process for shoplifting activity classification. Cybernetics and Information Technologies, 23(1), 141-160.

[20] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 3431–3440, 2015.