



Author: Bruce A. Liboy Maldonado  
 Advisor: Dr. Jeffrey Duffany  
 Master in Computer Science

## Abstract

Low-Light Image Enhancement (LLIE) plays a crucial role in photography, surveillance, autonomous systems, and scientific imaging. Traditional enhancement techniques often struggle to recover fine details and may introduce unwanted color distortions. In this project, a lightweight U-Net model will be developed for low-light image enhancement. Transfer learning will be employed by first pre-training the model on a synthetic low-light dataset, followed by fine-tuning on real paired low-light images. Using the LOw-Light (LOL) dataset on Kaggle, model performance will be evaluated before and after transfer learning using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) image quality metrics. The lightweight design of the model makes it well-suited for deployment on edge devices and mobile platforms.

## Introduction

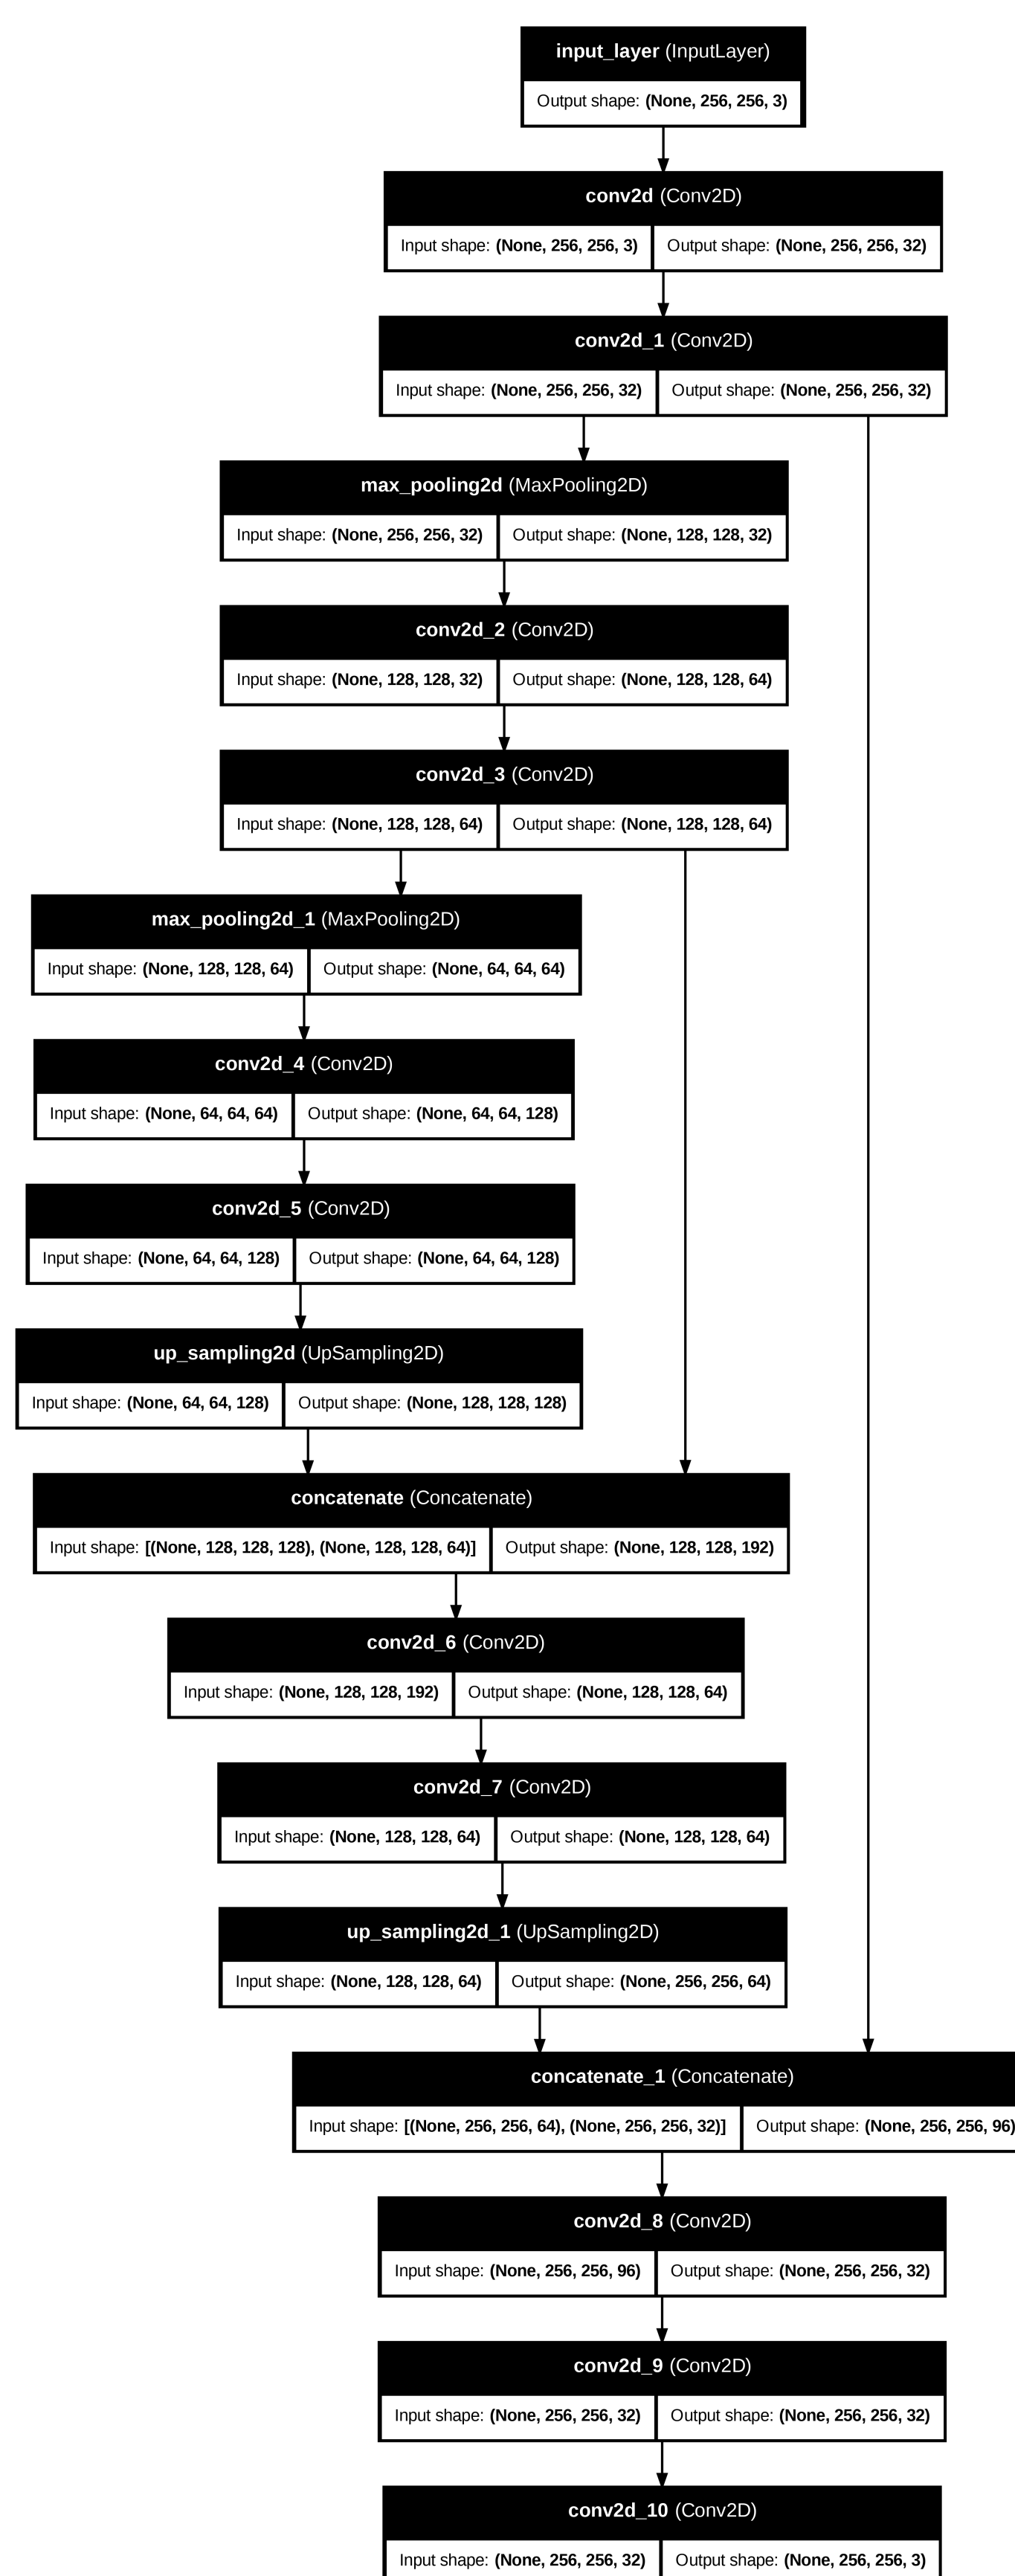
This research presents a series of experiments using a U-Net architecture to perform low-light image enhancement by generating visually augmented images with improved illumination. The overarching goal is to develop a lightweight, efficient model capable of processing raw low-light camera footage and producing well-lit images or video in real time on constrained edge computing devices for mission-critical, in situ applications such as security and exploration. U-Net was chosen due to its simplicity and computational efficiency, making it suitable for low-power environments. The experiments leverage the LOL dataset, which consists of 500 paired low-light and well-lit images (485 for training and 15 for testing), primarily indoor scenes with a resolution of 400×600, originally introduced in the *Deep Retinex Decomposition for Low-Light Enhancement* paper. The study adopts a comparative framework evaluating U-Net models with and without fine-tuning and benchmarks their performance against classical histogram equalization to assess the impact of fine-tuning and the relative advantages of machine learning-based approaches over traditional methods, with detailed training and evaluation procedures described in the Methodology section.

## Background

Low-light image enhancement (LLIE) seeks to improve visibility and perceptual quality in images captured under insufficient illumination, a common challenge in applications such as surveillance, autonomous driving, and medical imaging. Classical methods like Histogram Equalization (HE) and its variants enhance contrast through heuristic pixel-level adjustments but often introduce artifacts, amplify noise, and struggle with complex lighting conditions. Deep learning has reframed LLIE as a supervised image-to-image translation task, where convolutional neural networks—particularly U-Net-based architectures—learn non-linear mappings from low-light to well-lit images using paired datasets such as LOL, generally achieving superior structural and perceptual quality. However, these methods rely on costly paired data and are difficult to evaluate due to the subjective nature of visual quality, typically using metrics such as PSNR and SSIM. In this work, a U-Net-based LLIE model is evaluated with and without fine-tuning and compared against a classical histogram equalization baseline to assess the benefits of learned representations and task-specific refinement.

## Methodology

The methodology consists of data loading, preprocessing, model training, and evaluation using the LOL dataset within a Kaggle environment. Image pairs were loaded, split into training, validation, and testing sets, and processed through a TensorFlow pipeline that resized images to a uniform resolution, normalized pixel values, shuffled samples, and optimized input throughput through caching, batching, and prefetching. In addition to using real low-light images, synthetic low-light data were generated from well-lit images via random darkening, brightness reduction, and noise injection to support pre-training. A U-Net architecture was employed for image-to-image enhancement, leveraging encoder-decoder blocks with skip connections to preserve spatial detail. Two models were trained: one using only synthetic to-ground-truth pairs and another incorporating an additional fine-tuning stage with real low-light images. Training was performed using the Adam optimizer and MAE loss over fixed epochs. Model performance was evaluated on a held-out test set using PSNR to assess pixel-level fidelity and SSIM to measure perceptual and structural similarity, providing a balanced assessment of enhancement quality.



## Results and Discussion

Overall, the base model effectively enhances brightness while preserving color balance and avoiding extreme highlights or shadows, but its outputs suffer from softness and low contrast that reduce visual appeal. The fine-tuned model overcomes these issues by improving sharpness, edge definition, and color richness without losing brightness, resulting in more visually pleasing and structurally faithful images. In contrast, histogram equalization lags behind learning-based methods for natural images, as it degrades color fidelity and amplifies noise, producing desaturated and unnatural results; however, its strong contrast stretching makes it well suited for high-contrast, structure-focused tasks like text or document images, where edge clarity matters more than color realism.

Metric	Base Model	Fine-tuned Model	Hist. Eq.
PSNR Mean	14.291	16.251	15.134
PSNR Std. Dev.	3.611	2.917	2.403
SSIM Mean	0.709	0.776	0.463
SSIM Std. Dev.	0.133	0.094	0.112

TABLE I  
 PERFORMANCE COMPARISON OF LLIE METHODS




Fig. 2. A comparative 2X2 figure that depicts the output difference of each approach VS the ground truth (upper right).




Fig. 3. A comparative 2X2 figure that depicts how Histogram Equalization degrades color saturation.

## Conclusions

The results indicate that the baseline U-Net model provides a strong starting point, while fine-tuning via transfer learning yields clear improvements in both PSNR and SSIM, demonstrating enhanced pixel fidelity and perceptual quality. The fine-tuned model outperforms the non-fine-tuned version and shows more consistent results, whereas classical histogram equalization, despite reasonable PSNR, performs poorly in structural similarity. Overall, the findings confirm that learning-based LLIE methods—especially those incorporating fine-tuning—surpass classical techniques in both quantitative performance and visual quality.

## Future Work

Future work includes exploring alternative and more optimized model architectures, experimenting with different training configurations, and evaluating additional image quality metrics such as LPIPS, NIQE, BRISQUE, and error-based measures to capture perceptual and statistical aspects of enhancement. Expanding the dataset with diverse images and augmentation pipelines may further improve performance, while investigating training strategies such as varying epochs, pruning, distillation, and quantization could reduce model size for edge deployment. Finally, testing alternative loss functions and implementing the model for real-time, on-device enhancement represent promising directions to further improve both efficiency and output quality in lightweight LLIE systems.

## Acknowledgements

The author gratefully acknowledges Dr. Jeffrey Duffany for his guidance and support throughout this work, Alexandra Artusio-Glimpse for providing an Overleaf template aligned with IEEE Xplore formatting for the CPEM 2024 conference, and the Kaggle team for enabling easy access to the LOL dataset through their Jupyter notebook environment.

## References

- [1] Soumik Rakshit. LOL Dataset. Kaggle. Accessed November 29, 2025. <https://www.kaggle.com/datasets/soumikrakshit/lol-dataset>.
- [2] Liang, X., Chen, X., Ren, K. et al. Low-light image enhancement via adaptive frequency decomposition network. *Sci Rep* 13, 14107 (2023). <https://doi.org/10.1038/s41598-023-40899-8>
- [3] Nguyen, T. et al. (2026). Towards a Real-Time and Compact Low-Light Image Enhancement Method Using Deep Learning on ESP32-S3-EYE. In: Quan, T.T., Sombatheera, C., Pham, HA., Tran, N.T. (eds) *Multi-disciplinary Trends in Artificial Intelligence*. MIWAI 2025. *Lecture Notes in Computer Science*(), vol 16355. Springer, Singapore. [https://doi.org/10.1007/978-981-95-4963-4\\_4](https://doi.org/10.1007/978-981-95-4963-4_4)
- [4] Tahia Binte Chowdhury; Md. Khaliluzzaman; Md. Jashim Uddin; Karnish Khanam; Md. Monirul Islam. Low Light Enhancer: A Low Light Image Enhancement Model Based on U-Net using Smartphone. 2022 International Conference on Innovations in Science, Engineering and Technology (ICISSET), 26-27 Feb. 2022