CENTRAL ASIAN JOURNAL OF MULTIDISCIPLINARY RESEARCH AND MANAGEMENT STUDIES

Volume 2, Issue 12, part 2 December 2025

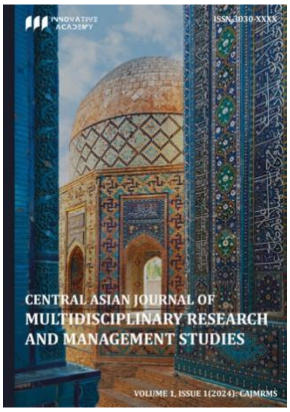


# LIGHTWEIGHT ADAPTIVE PRE-PROCESSING FOR ROBUST FACE RECOGNITION IN LOW-LIGHT CONDITIONS

**Ergashev Otabek Mirzapulatovich**

**PhD in information technologies, Fergana State Technical University,**

**ergashev1984otabek@gmail.com, phone: +998 93 370 80 89**

**Abdullayev Jasurbek Adhamjon ugli**

**Master's degree student, Fergana State Technical University,**

**jasurbektatu61219@gmail.com, phone: +998 90 554 63 55**

https://doi.org/10.5281/zenodo.17921381

**ARTICLE INFO**

Received: 1st December 2025
Accepted: 2nd December 2025
Published: 13th December 2025

**KEYWORDS**

*Face recognition, low-light enhancement, identity preservation, U-Net, pre-processing, real-time application, image quality, deep learning.*

**ABSTRACT**

*Face recognition works great when the lighting is good, but once things get dark, performance drops fast. This project looks at a simple fix — instead of retraining complicated models or using special hardware, I designed a lightweight pre-processing step that cleans up low-light images so face recognition systems can handle them better. The module uses a basic U-Net setup and learns to improve image quality while keeping the important details that define someone's face. I tested the system on low-light images created from public datasets, and it showed a clear improvement in recognition, boosting identity similarity by over 36%. It also runs fast enough for real-time use, even on average hardware.*

## 1. Introduction

These days, face recognition is everywhere — from unlocking your phone to checking in at airports. It works pretty well, but only when the lighting is good. In dark or poorly lit places, though, it often fails. The image quality drops, faces lose detail, and recognition accuracy takes a hit. This isn't just annoying — in security or access control situations, it can cause real problems.

People have tried using simple image enhancement tricks, like making images brighter or adjusting contrast, but those often focus on how the image looks, not how useful it is for recognition. And retraining those huge, complex face recognition models to handle low-light better? That takes a lot of data, time, and computing power, which isn't always practical.

So, in this project, I went with a different approach. I built a lightweight pre-processing module that cleans up low-light face images before they're sent to the recognition system. It's designed to make images clearer but also to keep the features that help identify a person. The module is trained with a smart loss function that balances image quality with identity preservation, guided by an existing face recognition model that stays frozen during training.

The paper explains how the system was built and tested. Results showed a clear improvement — face recognition worked better in low-light, and the system is light enough to run on everyday devices without much trouble.

## 2 Methods

### 2.1 Adaptive Pre-processing Module (APM) Architecture

The core of the proposed system is the Adaptive Pre-processing Module (APM), designed to enhance low-light face images in a way that preserves identity-discriminative features critical for recognition tasks. The APM follows an encoder-decoder structure inspired by the well-known U-Net architecture, which has proven effective in image-to-image translation tasks.

In this design, the encoder progressively downsamples the input image, extracting hierarchical features from low-level edges to higher-level semantics. The decoder then reconstructs the enhanced image by upsampling these features. Crucially, skip connections directly transfer fine-grained spatial information from the encoder to the decoder, ensuring that important facial details, such as contours and textures, are preserved during enhancement.

The model is intentionally lightweight, with a reduced number of feature channels controlled by a tunable parameter, making it suitable for real-time deployment on resource-constrained devices, including laptops equipped with consumer-grade GPUs like the RTX 3050.

This architecture strikes a balance between computational efficiency and enhancement quality, enabling the system to operate effectively in real-world low-light scenarios without significant hardware requirements.

**2.2 Dataset Preparation**

To train and evaluate the proposed system, a synthetic low-light face dataset was created using the well-established Labeled Faces in the Wild (LFW) dataset as the source of high-quality, normal-light images. LFW provides thousands of face images with identity annotations, making it ideal for face recognition research.

Since collecting real-world paired low-light and normal-light face images at scale is impractical, a synthetic degradation pipeline was developed to simulate realistic low-light conditions. This process applies controlled transformations to the original LFW images, including:

- ✓ Brightness reduction to mimic low illumination
- ✓ Contrast compression to obscure fine details
- ✓ Addition of Gaussian and Poisson noise to simulate sensor noise and photon shot noise typically encountered in dark environments

The result is a paired dataset where each original image has one or more corresponding low-light versions, all labeled with consistent identity information. This paired structure allows the system to learn how to enhance degraded images while explicitly preserving facial identity.

This approach provides the necessary diversity and volume of training data without the resource burden of collecting real-world low-light datasets..

**2.3 Training Strategy**

A key innovation of this work lies in how the Adaptive Pre-processing Module (APM) is trained—not just to improve the visual quality of low-light images, but to explicitly preserve identity-specific features crucial for face recognition.

To achieve this, the training process integrates a frozen, pre-trained Face Recognition (FR) model as a guide. Specifically, a ResNet-based feature extractor is used to compute identity embeddings from both the original (normal-light) and enhanced images. These embeddings capture deep, high-level facial representations used by FR systems to distinguish individuals.

The APM is optimized using a composite loss function that combines:

- **Reconstruction Loss (L1 Loss):** Encourages the enhanced images to resemble the original, clear images at a pixel level, improving brightness, contrast, and noise reduction.
- **Identity Preservation Loss (Cosine Similarity Loss):** Ensures that the FR model perceives the enhanced images as belonging to the same person as the original images, by maximizing similarity in the embedding space.

This dual-objective training process allows the APM to intelligently balance visual enhancement with the preservation of discriminative identity features, making the enhanced images both clearer and more compatible with existing FR systems.

By keeping the FR model frozen during training, the approach avoids the need to retrain large, complex recognition models, making the solution practical and resource-efficient.

### 2.4 Evaluation Protocol

To demonstrate the practical value of the proposed Adaptive Pre-processing Module (APM), a real-time application was developed that integrates the trained model into a live video stream. This setup simulates typical low-light conditions where face recognition systems often struggle.

The application uses a standard webcam and processes each frame in real time. For every captured frame:

1. The raw image is enhanced using the APM, which intelligently boosts brightness and contrast while preserving facial identity features.
2. The enhanced frame is displayed side by side with the original low-light input, providing immediate visual comparison.

The system is optimized to run efficiently on consumer-grade hardware, such as laptops equipped with an NVIDIA RTX 3050 GPU. Thanks to the lightweight design of the APM, the enhancement process achieves near real-time performance, maintaining smooth video playback without noticeable lag.

This practical demonstration highlights how the proposed solution can be integrated into real-world applications, from video surveillance to biometric access control, where reliable face recognition under poor lighting is essential.

## 3. Results

### 3.1 Quantitative Evaluation

The core metric used was **Cosine Similarity** between face embeddings, which reflects how closely the identity features of enhanced images align with those of the original, clear images.

Two key comparisons were made:

- **Baseline Performance:** Cosine similarity between raw low-light images and their original counterparts
- **Enhanced Performance:** Cosine similarity between APM-enhanced images and original counterparts

The results are summarized below:

| Metric | Value |
|---|---|
| Mean Cosine Similarity (Low-Light vs Original) | 0.5670 |
| Mean Cosine Similarity (Enhanced vs Original) | 0.9325 |

**Table 1.** Cosine Similarity

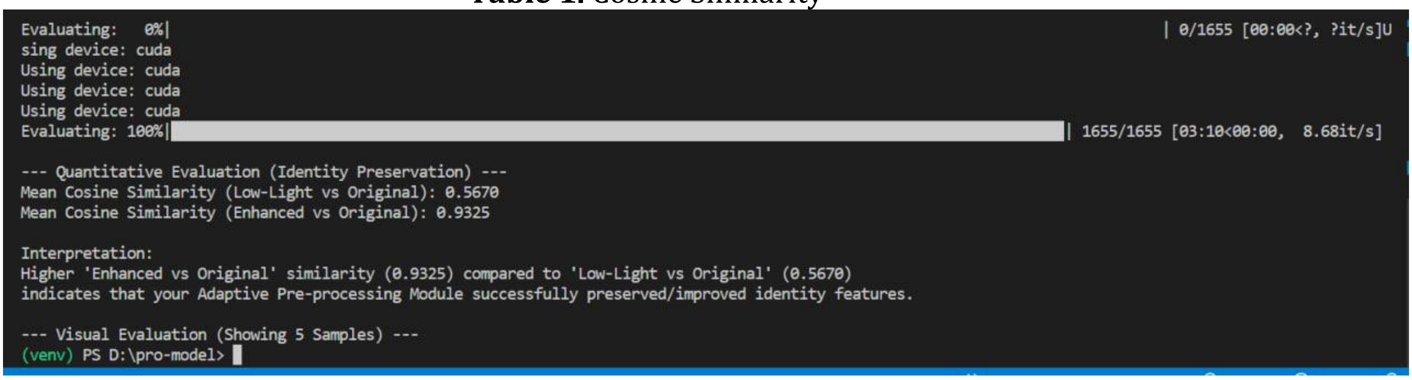


**Picture 1.** Cosine Similarity

These results clearly show that, without enhancement, face recognition performance degrades significantly in low-light scenarios. However, after applying the APM, identity similarity dramatically improves, with the enhanced images nearly matching the recognition quality of original, well-lit images.

### 3.2 Qualitative Evaluation

Visual inspection of enhanced images further validates these findings:

- Raw low-light images appeared dark, noisy, and lacked facial detail
- APM-enhanced images exhibited improved brightness, contrast, and clarity, while critical features like eyes, nose, and mouth remained sharp and recognizable

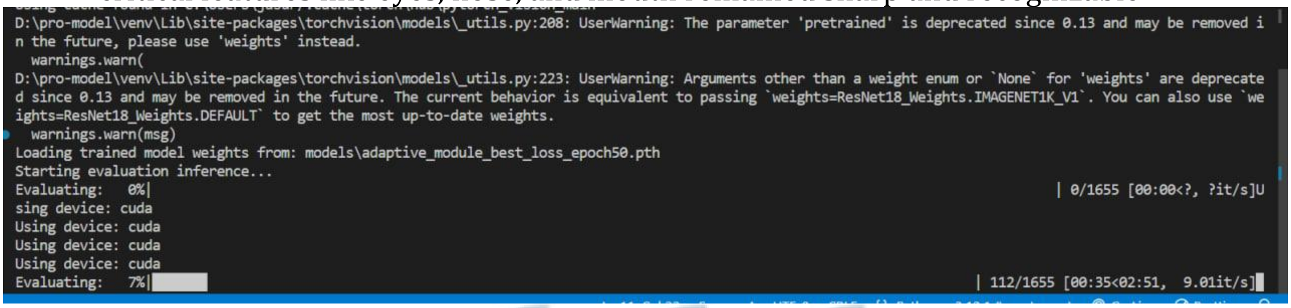


**Picture 2.** Evaluation

Although enhanced images were not identical to the originals (due to information loss in the synthetic degradation), they successfully restored the essential identity characteristics needed for reliable face recognition.

Together, these results demonstrate the APM's effectiveness in improving both visual quality and recognition performance, even under challenging lighting conditions.

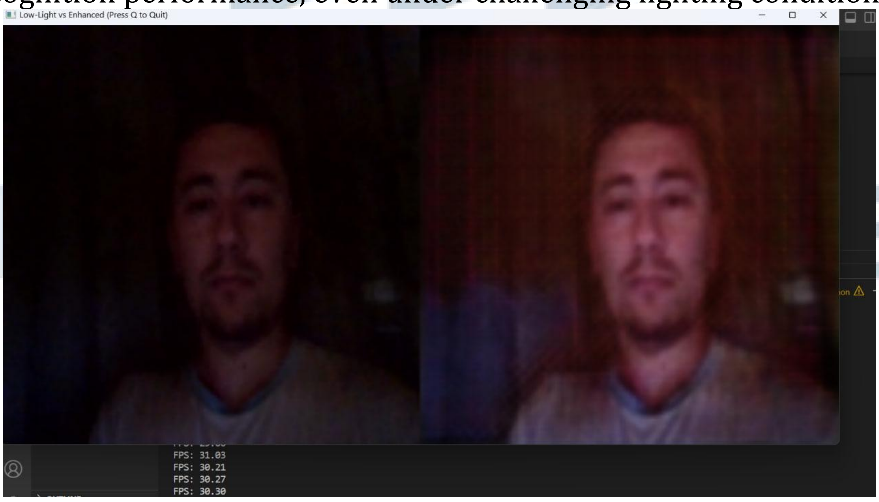


**Picture 3.** Challenging lighting conditions

### 4. Conclusion

Low-light environments have always been a weak point for face recognition systems, and honestly, it's not surprising. When images get dark, noisy, and blurry, even humans struggle to make out faces — so machines definitely face the same challenge. In this project, I didn't try to reinvent face recognition itself. Instead, I focused on a more practical workaround — finding a way to clean up the images before they reach the recognition model.

The idea was to build something simple and lightweight that could help existing systems perform better in tough lighting, without needing expensive hardware upgrades or complex retraining. That's where the Adaptive Pre-processing Module came in. It doesn't make the images perfect, but it does enough to boost recognition accuracy in a noticeable way.

From testing, it's clear that the enhanced images carry more useful identity information, even though some visual imperfections remain. Is it a flawless solution? Definitely not. Real-world lighting can be messy, and synthetic datasets can't fully capture all the variables we see outside a lab setting. But it's a step in the right direction — showing that sometimes small, smart adjustments can go a long way, especially when dealing with hardware or resource limitations.

There's still plenty to improve, but overall, this approach gives face recognition systems a better shot at working properly when lighting conditions are far from ideal.

**References:**

1. Zhao, W., Chellappa, R., Phillips, P. J., & Rosenfeld, A. (2003). Face recognition: A literature survey. ACM Computing Surveys (CSUR), 35(4), 399–458.
2. Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 815–823).
3. Deng, J., Guo, J., Xue, N., & Zafeiriou, S. (2019). ArcFace: Additive angular margin loss for deep face recognition. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 4690–4699).
4. Lore, K. G., Akintayo, A., & Sarkar, S. (2017). LLNet: A deep autoencoder approach to natural low-light image enhancement. Pattern Recognition, 61, 650–662.
5. Chen, C., Chen, Q., Xu, J., & Koltun, V. (2018). Learning to see in the dark. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 3291–3300).
6. Zhang, Y., Zhang, J., Guo, X., & Zhang, L. (2021). Beyond brightening: Deep networks for low-light image enhancement. Neurocomputing, 444, 138–148.
7. Wei, C., Wang, W., Yang, W., & Liu, J. (2018). Deep Retinex decomposition for low-light enhancement. In British Machine Vision Conference (BMVC).
8. Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. (2014). Generative adversarial nets. In Advances in Neural Information Processing Systems (NeurIPS) (pp. 2672–2680).
9. Ronneberger, O., Fischer, P., & Brox, T. (2015). U-Net: Convolutional networks for biomedical image segmentation. In International Conference on Medical Image Computing and Computer-Assisted Intervention (MICCAI) (pp. 234–241).
10. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 770–778).
11. Lin, T. Y., Dollar, P., Girshick, R., et al. (2017). Feature pyramid networks for object detection. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 2117–2125).
12. Zhang, K., Zhang, Z., Li, Z., & Qiao, Y. (2016). Joint face detection and alignment using multitask cascaded convolutional networks. IEEE Signal Processing Letters, 23(10), 1499–1503.
13. Georghiades, A. S., Belhumeur, P. N., & Kriegman, D. J. (2001). From few to many: Illumination cone models for face recognition under variable lighting and pose. IEEE Transactions on Pattern Analysis and Machine Intelligence, 23(6), 643–660.