

Enhancing Real-Time Face Recognition Robustness against Low Lighting through Dynamic Feature Enhancement

Muhammad Rival^{a*}, Dadang Iskandar Mulyana^b

^{a,b} Sekolah Tinggi Ilmu Komputer Cipta Karya Informatika, East Jakarta City, Special Capital Region of Jakarta, Indonesia.

ABSTRACT

Low-light conditions continue to challenge real-time face recognition because dim illumination often produces noisy and low-contrast images that weaken the quality of feature extraction. This study investigates how different preprocessing strategies influence the performance of a dlib-ResNet-based recognition system under such conditions. Two reference dataset sizes—33 and 1000 images—were used to observe how reference variation affects embedding stability. Enhancement was applied either offline to the reference dataset or in real time to incoming video frames, and both approaches were also tested in combination. The experiments show that offline preprocessing provides the most reliable improvement. Enhancing reference images raised the F1-Score by 7.28% (small dataset) and 7.50% (large dataset) without reducing processing speed, indicating that clearer embeddings at registration contribute to more stable matching. Real-time preprocessing, however, produced inconsistent results. While slight gains appeared in specific cases, the added computation and occasional distortion of facial structure reduced accuracy in other scenarios. The combined method produced the weakest performance, with the large dataset showing a 33.71% decline, suggesting that excessive modification disrupts structural consistency between reference and test images. Overall, the results highlight the importance of maintaining coherent facial features rather than applying aggressive adjustments to every frame. Offline enhancement is the most practical strategy for low-light deployments, whereas real-time enhancement should be used selectively. Future work may explore adaptive illumination adjustment capable of tuning enhancement parameters automatically to match varying lighting conditions.

ARTICLE HISTORY

Received 18 October 2025
Accepted 28 October 2025
Published 30 November 2025

KEYWORDS

Low-Light Recognition;
Preprocessing; Dlib-Resnet;
Face Embedding; Real-Time
Performance.

1. Introduction

If a contemporary biometric system is to be implemented in security setups, automatic attendance systems, access management, device authentication, and monitoring, it has to carry out immediate face acknowledgment. Biometric verification has the upside of giving more certainty in identifying an individual than password or PIN methods that can easily be compromised through theft or sharing (Chauhan *et al.*, 2025; Juluri, 2024). The typical solution to this issue is based on finding the face and extracting distinguishing features for comparison with reference data. Deep learning has largely contributed toward improving recognition reliability under well-lit and

controlled environments (Datta *et al.*, 2016; Gong *et al.*, 2000). Several real-world conditions have blocked recognition performance such as illumination change, pose variation, expression difference aging, and partial occlusion that disrupt stability in feature representation. Among these factors bad illumination always remains one of the most disturbing. Cameras operating under dim conditions usually produce images characterized by heavy noise with low contrast and blurry structural patterns. These distortions degrade the quality of features extracted by deep learning models and reduce their discriminability in the embedding space (Fang & Zhou, 2023; Goos *et al.*, 2014). Recognition systems do not gain any accuracy when using surveillance cameras or low-cost imaging devices without infrared support.

Low-Light Image Enhancement techniques apply Histogram equalization or contrast stretching to solve low light issues. Most of the LLE techniques are not applicable for face recognition. These enhancement algorithms are usually aesthetic and enhance over the images and noise content, distort gradients, create texture artifacts that downstream feature extraction will obstruct. Recognition-based enhancements have already proved that deep models which know about lighting conditions align better with the recognition strategy than aesthetic visual appearance alone for the effectiveness of recognition-based enhancements. Such studies bring out very clearly that improving performance on recognition under low illumination conditions needs preprocessing that enhances structural cues without adding misleading artifacts. This study uses a dlib-ResNet model characterized by an embedding dimension of 128 where the strength of this embedding depends fully on facial features being captured during the detection phase. Important structures such as edges become poorly defined when frames fall under low illumination so now local textures and facial contours generate embeddings with no representational clarity; even worse in real deployment since lighting variation can happen between successive frames. Therefore, preprocessing methods to enhance facial features before embedding are required for stabilizing recognition results over different illumination conditions. Offline enhancement of reference images guarantees consistent improvement since reference embeddings do not change after being generated once while real-time enhancement helps only under some conditions and often at a very high computational cost.

Recent works on illumination-aware enhancement use learning-based frameworks and self-calibrating neural networks to dynamically tune the illumination parameters for structural fidelity preservation while attaining noise reduction (Wang & Zhang, 2024; Zhu *et al.*, 2024). These models work well for their intended purposes but cannot be applied in lightweight real-time systems due to complexity and hardware requirements. The heavy computational burden from advanced illumination models coupled with enhancement method limitations makes it urgent to find a practical solution that will improve structural features without any additional system requirements. This paper is an extension of previous works on image enhancement, adaptive illumination correction, and low-light face recognition by exploring preprocessing techniques meant to boost recognition performance on facial features under poor lighting conditions to see which one between offline or real-time application of these approaches brings such an accuracy improvement at a practically deployable responsiveness level.

2. Methodology

The study adopts an experimental design using a private single-subject facial dataset developed to examine how dataset size influences the performance of a dlib-ResNet-based face recognition system. Two dataset variants were prepared to represent different

operational conditions: a small set consisting of 33 enrollment images, reflecting scenarios where registration data are limited, and a large set containing 1000 images, representing high-load deployments. All enrollment images were collected through a custom web interface and subsequently converted into 128-dimensional embeddings using the dlib-ResNet model. This transformation converts visual facial patterns into numerical representations that the system can compare with high precision during verification or identification. To ensure that the evaluation reflects real lighting challenges, two types of video inputs were used. The first consists of genuine low-illumination recordings captured directly under dim conditions, preserving natural noise patterns and degradation. The second consists of simulated low-illumination videos recorded under normal lighting and digitally altered to resemble dim scenes. Using both video types makes it possible to examine the system's behavior under authentic and controlled conditions, enabling a more reliable assessment of how lighting affects recognition accuracy.

The system workflow begins with real-time frame acquisition, followed by facial detection. Depending on the testing scenario, a preprocessing step is applied either to the reference dataset or to the incoming video frames. After preprocessing, feature extraction is performed using the dlib-ResNet model to generate a 128-dimensional embedding for each detected face. These embeddings are then matched against stored references to determine identity. Two primary preprocessing strategies were evaluated: offline dataset preprocessing, where enrollment images undergo Dynamic Feature Enhancement using CLAHE and Gamma Correction before embedding generation, and real-time preprocessing, where enhancement is applied directly to each incoming frame. To measure the effectiveness of these strategies, four testing scenarios were defined: a Baseline scenario with no preprocessing, an offline-preprocessing-only scenario, a real-time-preprocessing-only scenario, and a combined scenario that integrates both methods. Each scenario was tested across the two dataset sizes and in two modes—Benchmark Mode, which runs without visualization to measure raw computational performance, and Video Mode, which includes visual output to simulate real-world operation. System performance was quantified using True Positive (TP), True Negative (TN), False Positive (FP), and False Negative (FN) values. These metrics were used to compute Accuracy, Precision, Recall, and F1-Score using the following equations:

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}$$

$$Precision = \frac{TP}{TP + FP}$$

$$Recall = \frac{TP}{TP + FN}$$

$$F1 - Score = 2 \times \frac{Precision \times Recall}{Precision + Recall}$$

This structured evaluation enables a clear assessment of how offline and real-time preprocessing influence recognition accuracy and system responsiveness under low-illumination conditions.

3. Results

The evaluation was conducted to compare the performance of the face recognition system under different preprocessing strategies, dataset sizes, and low-illumination

conditions. The results show clear variations in recognition accuracy and processing speed depending on the treatment applied. To observe the system’s behavior when only a small number of reference images are available, testing was first performed using a dataset of 33 images. As shown in Table 1, the baseline configuration produced an F1-Score of 0.6441 with an average processing speed of 25.45 FPS. Applying offline preprocessing to the dataset improved the F1-Score to 0.6910, a relative increase of 7.28%, although the processing speed declined slightly. Real-time preprocessing produced the weakest outcome, reducing accuracy to 0.5700. When both offline and real-time preprocessing were combined, the system showed only a small improvement compared with the real-time-only scenario but still performed below the baseline.

Table 1. Results for the Small Dataset (33 Images)

Scenario (Benchmark Mode)	F1-Score	Δ F1-Score (%)	Average FPS	Δ FPS (%)
Baseline	0.6441	–	25.45	–
Dataset Preprocessing	0.6910	+7.28%	24.49	-3.77%
Video Preprocessing	0.5700	-11.51%	24.49	-3.78%
Both Preprocessing	0.6264	-2.76%	23.37	-8.19%

Evaluation of the larger dataset containing 1000 images revealed a different pattern, particularly when both preprocessing methods were applied simultaneously. As shown in Table 2, the baseline F1-Score was 0.8822, already higher than the small dataset due to richer reference variation. Offline preprocessing increased accuracy by 7.50% without reducing system speed. Real-time preprocessing offered only a minor improvement but resulted in lower FPS. The most notable outcome occurred in the combined-preprocessing scenario, where accuracy fell sharply to 0.5848, a decrease of more than 33%. This significant drop indicates that applying enhancement to both reference and test images can distort their structural consistency and reduce embedding stability.

Table 2. Results for the Large Dataset (1000 Images)

Scenario (Benchmark Mode)	F1-Score	Δ F1-Score (%)	Average FPS	Δ FPS (%)
Baseline	0.8822	–	24.16	–
Dataset Preprocessing	0.9484	+7.50%	24.26	+0.41%
Video Preprocessing	0.9014	+2.17%	23.62	-2.25%
Both Preprocessing	0.5848	-33.71%	23.66	-2.07%

The results demonstrate that offline preprocessing consistently strengthens recognition performance, as it enhances reference images before embedding, allowing the model to build more stable identity representations. This improvement is reflected in higher Precision and Recall values across both dataset sizes. In contrast, real-time preprocessing introduces additional computational load for every incoming frame, which results in lower FPS. For small datasets, this approach also decreases accuracy, indicating that enhancements applied immediately before recognition can introduce noise or distort facial structure. The larger dataset shows greater baseline accuracy due to the availability of more varied reference embeddings, although this advantage comes with slightly reduced FPS because of the heavier matching workload. Offline preprocessing still provides a notable improvement, reinforcing its scalability across dataset sizes. The combined preprocessing approach, however, produces the strongest negative impact—especially in the large dataset—suggesting that excessive manipulation of both reference and input images creates inconsistencies that disrupt the model’s ability to generate reliable embeddings. Overall, the findings underline a clear trade-off between accuracy and processing speed in low-illumination environments. Offline preprocessing offers the

most balanced improvement, while real-time enhancement must be used cautiously due to its computational demands and its tendency to reduce recognition stability.

4. Discussion

The findings of this study demonstrate that preprocessing strategies influence recognition reliability in distinct ways when applied to a dlib-ResNet–based face recognition system operating under low-illumination conditions. The most consistent improvement was observed when enhancement was applied offline to the reference dataset. This strategy increased the F1-Score by 7.28% in the small dataset and 7.50% in the large dataset, indicating that strengthening reference embeddings is more advantageous than modifying incoming frames during runtime. This outcome aligns with prior studies emphasizing the importance of high-quality reference representations in embedding-based models, where discriminative clarity improves the stability of vector matching (Datta *et al.*, 2016; Gong *et al.*, 2000; Fang & Zhou, 2023). Offline preprocessing using Dynamic Feature Enhancement produced more distinct facial structures without introducing additional computational load, enabling the system to maintain a stable processing rate of approximately 24 FPS. Real-time preprocessing, in contrast, showed inconsistent behavior. Although enhancement theoretically improves the visibility of facial structure, the experiments revealed that applying such adjustments to each incoming frame often reduced accuracy, especially in the small dataset.

Similar challenges have been reported in studies where enhancement algorithms such as CLAHE and traditional histogram-based methods amplify background noise or generate artificial gradients that obscure identity-specific features (Liu *et al.*, 2021; Merikapudi *et al.*, 2020). The computational overhead of processing every frame also contributed to a reduction in system responsiveness, which is critical for time-sensitive applications such as surveillance or access control (Chauhan *et al.*, 2025; Juluri, 2024). Even in the large dataset, where richer reference variation typically improves recognition robustness, real-time enhancement produced only a marginal F1-Score increase (+2.17%), confirming that its benefits are limited. A notable outcome of this study is the substantial accuracy drop observed when offline and real-time preprocessing were combined, particularly in the large dataset, where accuracy declined by 33.71%. This phenomenon suggests the presence of over-enhancement, a condition in which both reference and test images undergo aggressive contrast or brightness transformations. Excessive manipulation distorts fine-grained facial structures that deep models rely on for generating stable embeddings, creating discrepancies between the processed reference images and incoming video frames. Similar concerns have been noted in illumination-focused research, where mismatched enhancement across datasets leads to reduced feature consistency and degraded model performance (Fan *et al.*, 2024; Zhou *et al.*, 2024; Dharmadinata & Kusuma, 2023).

These results reinforce the importance of maintaining visual coherence between reference data and test inputs rather than optimizing each independently. Dataset size also played a significant role in shaping system behavior. The large dataset produced a far stronger baseline F1-Score (0.8822) compared with the small dataset (0.6441), reflecting the benefit of diverse reference images in building more representative embeddings. However, larger datasets also increased matching workload, which marginally reduced FPS—a behavior that aligns with findings from research on computational efficiency in large-scale recognition systems (Goos *et al.*, 2014; Hadi, 2025). Despite this additional workload, offline preprocessing remained effective in both dataset sizes, demonstrating its scalability across different deployment scenarios. The broader implications of these observations highlight the limitations of fixed-parameter

enhancement techniques such as CLAHE and Gamma Correction. Illumination conditions in real deployments fluctuate widely, and static parameters are unlikely to suit all lighting variations, a challenge also emphasized in recent illumination-adaptive studies (Wang & Zhang, 2024; Zhu *et al.*, 2024; Masyitoh *et al.*, 2024). The inconsistencies produced by fixed enhancements provide a strong rationale for developing adaptive preprocessing algorithms capable of adjusting contrast, brightness, or gamma levels based on the lighting characteristics of each frame. Such adaptability would reduce the risk of over-enhancement and improve recognition stability without adding excessive computational burden.

Overall, the findings emphasize that maintaining consistent feature integrity is more valuable than applying aggressive visual corrections. Enhancing the reference dataset offline offers the most reliable improvement, while real-time enhancement requires careful consideration due to its unpredictable effects. Future development of adaptive illumination correction methods may further improve recognition performance by balancing accuracy and computational efficiency in complex low-illumination environments.

5. Conclusion

The study demonstrates that improving the quality of reference images through offline preprocessing provides the most reliable way to strengthen face recognition performance under low-illumination conditions. Applying Dynamic Feature Enhancement to the reference dataset consistently increased the F1-Score—by 7.50% in the large dataset and 7.28% in the small dataset—while maintaining real-time processing speed. These findings indicate that enhancing the clarity of reference embeddings contributes more directly to recognition stability than modifying incoming frames during system operation. In contrast, real-time preprocessing produced inconsistent results and introduced a clear trade-off between accuracy and computational efficiency. While marginal accuracy gains were observed in specific cases, processing each incoming frame added computational overhead that reduced the frame rate. The most critical issue emerged when both offline and real-time preprocessing were used together, resulting in substantial accuracy degradation due to excessive manipulation of facial features. This outcome highlights the risk of introducing structural inconsistencies when enhancement is applied to both reference and test images.

Based on these observations, systems designed for low-light environments should prioritize the use of reference datasets that have undergone offline enhancement, particularly in applications involving large volumes of facial data. Real-time preprocessing should be applied selectively and only in scenarios where lighting conditions fluctuate drastically and cannot be compensated by the reference data alone. Future work may focus on developing adaptive preprocessing algorithms capable of adjusting enhancement parameters dynamically according to the lighting characteristics of each frame, reducing the likelihood of over-enhancement. Further investigations involving hardware acceleration (such as GPU-based optimization) and evaluations under more complex conditions—such as strong backlighting or partial facial occlusion—will help validate the system's robustness in practical deployments.

References

Chauhan, D., Singh Thakur, R., Choubey, U., Tripathi, S., Rajput, V., & Pagare, S. (2025). *E Crowd Vision: Real-time face recognition and crowd counting in dynamic environments*. Retrieved from <https://www.researchgate.net/publication/388069204>

- Datta, A. K., Datta, M., & Banerjee, P. K. (2016). *Face detection and recognition: Theory and practice*.
- Dharmadinata, O. J., & Kusuma, G. P. (2023). Improving face recognition in low illumination condition using combination of image enhancement and face recognition methods. *Journal of Theoretical and Applied Information Technology*, 28(4).
- Fan, Y., et al. (2024). Low-FaceNet: Face recognition-driven low-light image enhancement. *IEEE Transactions on Instrumentation and Measurement*, 73, 1–13. <https://doi.org/10.1109/TIM.2024.3372230>
- Fang, Z., & Zhou, Z. (2023). Studies advanced in robust face recognition under complex light intensity. In *Proceedings of the International Conference on...* (pp. 1005–1012). https://doi.org/10.2991/978-94-6463-300-9_101
- Gong, S., McKenna, S. J., & Psarrou, A. (2000). *Dynamic vision: From images to face recognition*.
- Goos, G., et al. (2014). *Face and facial expression recognition from real world videos*. Retrieved from <http://www.springer.com/series/7412>
- Hadi, A. (2025). *Enhancing face recognition in surveillance systems using local binary pattern and PCA-based feature selection*. Retrieved from www.scienticreview.com
- Juluri, S. (2024). Smart attendance monitoring system using multimodal biometrics. *Sigma Journal of Engineering and Natural Sciences*, 147–? <https://doi.org/10.14744/sigma.2024.00030>
- Li, J.-B., Chu, S.-C., & Pan, J.-S. (n.d.). *Kernel learning algorithms for face recognition*.
- Liu, J., Xu, D., Yang, W., Fan, M., & Huang, H. (2021). Benchmarking low-light image enhancement and beyond. *International Journal of Computer Vision*, 129(4), 1153–1184. <https://doi.org/10.1007/s11263-020-01418-8>
- Masyitoh, S. L., Ma'Ruf, K., & Setiawan, R. J. (2024). Local binary pattern and principal component analysis for low-light face recognition. In *Proceedings of the International Conference on Electrical Engineering, Computer Science and Informatics (EECSI)* (pp. 211–216). IEEE. <https://doi.org/10.1109/EECSI63442.2024.10776119>
- Merikapudi, S., Math, S., Nandini, C., & Rafi, M. (2020). Face recognition using CNN trained with histogram equalization-based image enhancement scheme.
- Mulyana, D. I., & Edi. (2023). Penerapan face recognition dengan algoritma Viola-Jones dalam sistem presensi kehadiran siswa dan guru pada Sekolah IDN Boarding School Jonggol. *Jurnal Indonesia: Manajemen Informatika dan Komunikasi*, 4(3), 1749–1757. <https://doi.org/10.35870/jimik.v4i3.398>
- Mulyana, D. I., et al. (2023). Penerapan face recognition dengan algoritma Haar Cascade untuk sistem absensi pada Yayasan Pusat Pengembangan Anak Jakarta.
- Pavani, S.-K. (2010). *Methods for face detection and adaptive face recognition*.

- Wang, H., & Zhang, P. (2024). Research on HLS implementation of self-calibrated illumination neural network: Low-light image enhancement. In *IEEE Information Technology, Networking, Electronic and Automation Control Conference (ITNEC)* (pp. 800–804). IEEE. <https://doi.org/10.1109/ITNEC60942.2024.10732922>
- Zhou, X., Gao, Z., Gong, H., & Li, S. (2024). DeFFace: Deep face recognition unlocked by illumination attributes. *Electronics*, *13*(22), 4566. <https://doi.org/10.3390/electronics13224566>
- Zhu, S. (2024). Enhancing facial recognition: A review of deep learning approaches and future perspectives. *Applied and Computational Engineering*, *110*(1), 137–145. <https://doi.org/10.54254/2755-2721/110/2024MELB0107>
- Zhu, Z., Yang, X., Lu, R., Shen, T., Zhang, T., & Wang, S. (2024). Ghost imaging in the dark: A multi-illumination estimation network for low-light image enhancement. *IEEE Transactions on Circuits and Systems for Video Technology*. <https://doi.org/10.1109/TCSVT.2024.3472278>