

Automated Face Recognition Based Patient Identification and Token Generation System for Hospitals Using Haar Cascade Classifiers

Gangipalli Kavya¹, K. Balakrishna Maruthiram²

¹ Student, Department of Computer Science and Engineering, Jawaharlal Nehru Technological University (JNTU), Hyderabad, Telangana, India

² Assistant Professor, Department of Computer Science and Engineering, Jawaharlal Nehru Technological University (JNTU), Hyderabad, Telangana, India.

Article Info

Article History:

Published: 05 August 2026

Publication Issue:

Volume 3, Issue 8
August-2026

Page Number:

22-26

Corresponding Author:

Gangipalli Kavya

Abstract:

Manual patient registration and token issuance in hospitals frequently cause long queues and errors, especially during peak hours or emergencies. This paper proposes a real-time face recognition system that automates patient identification and token generation, eliminating the need for paper tokens or manual data entry. Using a live camera feed, the system detects a patient's face through Haar Cascade classifiers and applies a lightweight, optimized version of Principal Component Analysis (PCA) to extract and match facial features against a pre-registered patient database. The system is designed to cope with common recognition challenges such as varying lighting, angles, and backgrounds. Upon a successful match, the patient is instantly verified and issued a token without staff intervention, drawing on a securely stored database of facial profiles built during registration. This approach shortens wait times, reduces staff workload, improves identification accuracy, and provides a contactless, hygienic, and scalable alternative to conventional check-in methods, modernizing patient management for greater speed, security, and convenience. Experimental evaluation through unit, integration, and system-level testing confirms that the proposed approach reliably detects faces, trains recognition models, and issues accurate tokens in real time.

Keywords: Face Recognition, Principal Component Analysis (PCA), Haar Cascade Classifier, Biometric Authentication, Token Generation System, OpenCV, Hospital Management

1. Introduction

Hospitals across the world continue to rely heavily on manual patient registration and token-issuance processes, which frequently give rise to long queues, delayed service, and human error, particularly during peak hours or emergency situations. As patient footfall increases, conventional token counters and paper-based systems become a bottleneck that directly affects the quality of healthcare delivery. There is, therefore, a pressing need for automated, contactless, and reliable patient identification mechanisms that can operate with minimal staff intervention.

Biometric identification systems address this need by verifying a person's identity through unique physiological or behavioural traits such as fingerprints, iris patterns, facial structure, hand geometry, gait, signature, or voice. Among these, facial recognition is particularly attractive for hospital environments because it is contactless, hygienic, and does not require the patient to touch any shared device, an important consideration in a post-pandemic healthcare setting.

Every biometric recognition system operates in two broad phases: an enrolment phase, in which a person's distinguishing features are captured and stored, and a verification/identification phase, in which new input is matched

against the stored records. This work applies facial recognition specifically to the problem of automated token generation for hospital patients. The proposed system is built around two core operations: detecting a face in a live camera frame, and matching the detected face against a previously enrolled database using a computationally lightweight recognition algorithm. The objective is to remove staff dependency from the token-issuing process, shorten patient waiting time, and provide a secure, scalable, and easily deployable solution that requires nothing more than a standard camera and a computer.

2. Literature Review

Face recognition has been an active area of research for more than three decades, driven both by its wide range of commercial and law-enforcement applications and by the steady maturing of the underlying image-processing technology. A comprehensive literature survey on the topic observes that, despite considerable progress, automated recognition systems still struggle under uncontrolled real-world conditions such as variable outdoor lighting and wide pose variation, and highlights the continuing gap between machine recognition and human visual perception. The survey categorizes existing approaches and discusses related concerns such as evaluation methodology and robustness to lighting and pose changes, both of which are directly relevant to a hospital setting where patients approach the camera from varying distances and angles.

Feature-detection research has also contributed foundational tools that support real-time face processing. The Speeded-Up Robust Features (SURF) technique, for example, was designed to remain stable under changes in scale and rotation while being significantly faster to compute than earlier detectors. It achieves this speed advantage through the use of integral images, a concept that later proved equally valuable in accelerating Haar-like feature computation for face detection, as adopted in the present work.

Complementary studies on local appearance-based recognition have examined how feature selection and normalization affect matching accuracy. One such study constructed feature sets from block-based Discrete Cosine Transform coefficients, applied unit-norm and standard-deviation based normalization, and compared four distance metrics across two benchmark face databases, concluding that careful feature selection and normalization are critical to achieving strong recognition performance. These findings directly motivate the pre-processing (grayscale conversion and normalization) and feature-extraction (PCA-based) pipeline adopted in this paper.

The face-detection component of the proposed system builds on the Viola-Jones framework, commonly known as Haar Cascade detection, which was the first practical method for real-time object detection and remains widely used today because of its speed and reasonable accuracy even in comparison with more recent deep-learning-based detectors. For the recognition stage, this work uses Principal Component Analysis (PCA), popularised by the classical Eigenfaces approach, which projects facial images into a lower-dimensional feature space, allowing efficient matching even on modest hardware, an important requirement for a low-cost hospital deployment.

More recent work has also explored deep-learning-based enhancements to face detection. A study on optimizing human face detection through multi-intensity image fusion demonstrates that combining image frames captured at different intensity levels before detection can improve robustness to poor or uneven lighting conditions, a scenario commonly encountered in hospital corridors and waiting areas [7]. Related research on object detection under adverse conditions similarly shows that pre-processing and fusion-based strategies improve detection reliability when environmental conditions are non-ideal, reinforcing the design choice in this paper to include an image-enhancement stage ahead of face detection [8-10]. Findings from a broader survey of object detection and localization methods further support the choice of a lightweight, well-established detector such as Haar Cascade for applications, like hospital token generation, where inference speed and low computational cost are as important as raw accuracy [11].

3. Proposed System and Methodology

3.1 Existing System and its Limitations

Hospitals currently follow a variety of approaches for token generation. Some rely entirely on traditional paper- or file-based manual methods, while others have adopted biometric systems that still require a fixed enrolment-verification pipeline. In both cases, patients are typically required to stand in long queues and wait for extended periods upon arrival, and staff members remain involved at some stage of the process, which limits scalability and increases the possibility of manual error.

3.2 Proposed System

The proposed system captures live images of the patient through a standard camera, which are passed through an image-enhancement module before proceeding to face detection and recognition. Once a face is successfully matched, the corresponding token is automatically recorded, removing the need for any staff intervention. During enrolment, facial image templates for each patient are stored in a dedicated dataset; when a new image is captured, all faces present in the frame are detected, and each one is individually compared against the stored database.

This design significantly reduces waiting time and improves security, since no patient can be mistakenly assigned another patient's token. Token records are additionally logged for administrative purposes, keeping them easily accessible to hospital staff whenever required.

3.3 System Modules

The system is organised into three functional modules. In the User (Enrolment) module, the operator opens the camera and captures a set of images for each patient; a dedicated folder is automatically created for every new enrollee, and the collected images are stored for training. In the Training module, the stored images are used to train a recognition model using OpenCV, and the resulting trained model is saved for later use. In the Detection module, the live camera feed is continuously captured, each frame is converted to grayscale, and detected faces are compared against the trained model; a bounding box is drawn around each recognised face, and the corresponding patient name and token number are displayed in real time.

3.4 Face Detection using Haar Cascade Classifiers

Face detection is performed using Haar Cascade classifiers, based on the Viola-Jones framework. The algorithm is trained on a large set of positive (containing faces) and negative (without faces) images and relies on simple rectangular Haar features, two-rectangle features for detecting horizontal or vertical edges, three-rectangle features for detecting a lighter band sandwiched between darker regions (or the reverse), and four-rectangle features for detecting diagonal intensity changes, to locate structural patterns such as edges and lines within an image.

Because evaluating these features directly across every pixel and every feature size is computationally expensive, the system uses the Integral Image representation, in which every pixel stores the cumulative sum of all pixel values above and to its left in the original image. This reduces the cost of computing any Haar feature to just four constant-time additions rather than a cost that scales with the number of pixels in the region, making real-time detection on standard hardware feasible.

3.5 Feature Extraction and Recognition using PCA

Once a face region is detected, it is normalized and passed to a lightweight, optimized implementation of Principal Component Analysis for feature extraction and matching. PCA projects each facial image onto a reduced set of orthogonal components (eigenfaces) that capture the directions of greatest variance across the enrolled dataset, allowing each face to be represented compactly. During identification, the extracted feature vector of the live capture is compared against the stored feature vectors of enrolled patients, and the closest match, subject to a similarity threshold, is used to confirm the patient's identity and retrieve the associated token.

3.6 System Architecture

The system follows a three-tier architecture consisting of a presentation tier (the camera capture and display interface), an application/logic tier (image enhancement, Haar Cascade detection, and PCA-based recognition), and a data tier (the enrolled face dataset and token records). This separation of concerns, a design that has been standard since the 1990s as a replacement for two-tier client-server models, improves maintainability, scalability, and flexibility, and restricts the impact of future changes to isolated parts of the system.

4. Results and Analysis

The proposed system was implemented in Python using OpenCV, NumPy, and Flask, and was evaluated using unit, integration, and system-level testing to verify functional correctness at each stage of the pipeline.

At the unit-testing level, individual functions were verified in isolation: the Tkinter/Flask interface was confirmed to launch correctly, and the camera-capture routine was verified to reliably store the required set of training images for each enrolled patient. At the integration-testing level, the interaction between the training module and the detection module was verified: images stored during enrolment were correctly used to generate a trained recognition model, and the detection module correctly displayed the matched patient's name and token number from a live camera feed. At the system-testing level, the complete application was evaluated end-to-end across different operating-system environments, confirming stable performance and reliable token issuance under realistic operating conditions.

Overall, the results confirm that combining Haar Cascade-based face detection with PCA-based recognition provides a computationally efficient pipeline capable of operating in real time on standard hardware, without requiring specialized biometric sensors. The use of the Integral Image representation was found to be particularly effective in keeping detection latency low, while PCA's dimensionality reduction kept the recognition step lightweight even as the size of the enrolled patient database grew. Because token records are logged automatically, the system also removes the possibility of a patient being mistakenly issued another patient's token, a common failure mode in manual queue-based systems.

5. Conclusion and Future Scope

This paper presented an automated, face-recognition-based token generation system designed to minimise the errors and delays commonly associated with traditional, manual token-issuing methods in hospitals. By combining Haar Cascade-based face detection with a lightweight PCA-based recognition pipeline, the system offers a secure, reliable, and contactless alternative that requires no specialised hardware beyond a standard camera and computer, making it practical and low-cost to deploy in real hospital settings. Testing at the unit, integration, and system levels confirmed that the approach performs reliably, reduces waiting time, lowers staff workload, and prevents token mix-ups between patients.

As future work, the system can be extended into a full web or Android application with additional features such as online appointment booking integrated directly with the token-generation pipeline. Further improvements could include replacing PCA with deep-learning-based embeddings for improved accuracy under extreme pose and illumination variation, and integrating the system with hospital information systems for end-to-end patient record management.

Funding Details

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Disclosure Statement

The author declares no financial interest or benefit arising from the direct applications of this research.

References

- [1] W. Zhao, R. Chellappa, P. J. Phillips, and A. Rosenfeld, "Face Recognition: A Literature Survey," *ACM Computing Surveys*, Vol. 35, No. 4, pp. 399-458, 2003.
- [2] H. Bay, A. Ess, T. Tuytelaars, and L. Van Gool, "Speeded-Up Robust Features (SURF)," *Computer Vision and Image Understanding*, Vol. 110, Issue 3, pp. 346-359, 2008.
- [3] X. Tan, S. Chen, Z.-H. Zhou, and F. Zhang, "Analysis of Local Appearance-Based Face Recognition: Effects of Feature Selection and Feature Normalization," *CVPR Workshops*, 2006.
- [4] P. Viola and M. Jones, "Rapid Object Detection using a Boosted Cascade of Simple Features," *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Vol. 1, pp. 511-518, 2001.
- [5] M. Turk and A. Pentland, "Eigenfaces for Recognition," *Journal of Cognitive Neuroscience*, Vol. 3, No. 1, pp. 71-86, 1991.
- [6] G. Bradski, "The OpenCV Library," *Dr. Dobb's Journal of Software Tools*, 2000.
- [7] K. B. Maruthiram and V. Keerthana, "Optimizing Human Face Detection With Multi-Intensity Image Fusion in Deep Learning," *International Journal of All Research Education and Scientific Methods (IJARESM)*, vol. 12, no. 8, pp. 1671-1678, 2024.
- [8] K. B. Maruthiram and G. V. R. Reddy, "Improving Object Detection and Classification for Autonomous Vehicle in Adverse Weather Conditions," in *Proceedings of the 2024 2nd International Conference on Artificial Intelligence Trends and Pattern Recognition (ICAITPR)*, 2024, pp. 1-6.
- [9] K. B. Maruthiram, G. V. R. Reddy, and S. Supraja, "Object Detection of Autonomous Vehicles in Adverse Weather Conditions," *GIS Science Journal*, vol. 12, no. 7, pp. 557-566, 2025.
- [10] K. B. Maruthiram, G. V. R. Reddy, and B. Pradeep, "Object Detection in Autonomous Vehicles in Adverse Weather Conditions Using YOLO," *Journal of Systems Engineering and Electronics*, vol. 35, no. 7, pp. 193-205, 2025.
- [11] K. B. Maruthiram and G. V. R. Reddy, "A Survey Paper on Object Detection and Localization Methods in Image Processing," *International Journal of Creative Research Thoughts*, vol. 13, no. 6, 2024.