International Journal of Science, Strategic Management and Technology
Volume **02** Issue **04** April-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

Face Recognition Based Attendance Management System

Ashbeer Basha A , Anish Fathima
UNDER THE GUIDANCE OF
Dr.B.Suresh, M.Sc., PGDCA., MCA., M.Phil., Ph.D
Department of Computer Science and Engineering,
Vels Institute of Science, Technology and Advanced Studies,Pallavaram, Chennai.
Email: ashbeerbasha@gmail.com
Email: anisharahamathullah5239@gmail.com

https://doi.org/10.55041/ijsmt.v2i4.619

Cite this Article: A, A. B. & Fathima, A. (2026). Face Recognition Based Attendance Management System. International Journal of Science, Strategic Management and Technology, *02*(04). https://doi.org/10.55041/ijsmt.v2i4.619

Abstract

Attendance management is a fundamental administrative task in educational institutions and organizations. Traditional methods, including manual registers, identity cards, and fingerprint-based biometric devices, suffer from significant limitations such as proxy attendance, hygiene concerns, manual errors, and scalability issues. This paper presents a Face Recognition Based Attendance Management System (FRAMS) that leverages state-of-the-art deep learning techniques — specifically Convolutional Neural Networks (CNNs) and the FaceNet architecture — to enable fully automated, contactless, and accurate attendance recording. The proposed system captures real-time video frames, detects and preprocesses facial regions, extracts 128-dimensional facial embeddings, and compares them against a pre-enrolled database using cosine similarity to determine identity. The attendance record is stored automatically upon successful recognition. Experimental results demonstrate a recognition accuracy of 97.2% on a dataset spanning 50 individuals under diverse lighting and pose conditions, with an average recognition latency of 0.38 seconds per frame. The system is implemented using Python, OpenCV, TensorFlow, and SQLite, making it cost-effective and deployable on commodity hardware. Comparative evaluations confirm superior performance against manual, card-based, and fingerprint biometric systems in accuracy, speed, and hygiene. Future work addresses cloud integration, multi-face batch processing, and GPU-accelerated real-time inference.

Index Terms — Face Recognition, Deep Learning, Convolutional Neural Networks, FaceNet, Attendance Management, Computer Vision, OpenCV, TensorFlow, Biometric Systems, Cosine Similarity.

I. INTRODUCTION

Attendance management is a critical requirement across educational institutions, corporate offices, government organizations, and public sector enterprises. The accuracy and efficiency of attendance tracking directly impact productivity, resource allocation, and administrative decision-making. Despite advancements in automation, many organizations continue to rely on outdated attendance mechanisms that are manual-intensive, error-prone, and susceptible to fraudulent manipulation.

Manual attendance systems require instructors or supervisors to call roll or sign registers, consuming significant class or work time. Card-based systems introduce a risk of proxy attendance, where one individual can swipe the card of an absent colleague. Fingerprint biometrics, while more robust, require physical contact with scanners, raising concerns about device hygiene — a factor that gained critical importance following the COVID-19 pandemic [1]. Additionally, fingerprint scanners require specialized hardware, regular maintenance, and fail for individuals with dermato-logical conditions affecting fingerprint clarity.

Recent advances in deep learning and computer vision have demonstrated remarkable success in face recognition tasks, achieving human-level accuracy on large-scale benchmarks [2]. Systems such as FaceNet [3] — developed by Google — produce compact, highly discriminative 128-dimensional facial embeddings that enable reliable one-shot and few-shot recognition. This capability makes deep learning based face recognition an ideal candidate for modernizing attendance management: it is contactless, fast, and operates on standard camera hardware without requiring dedicated biometric devices.

This paper presents the Face Recognition Based Attendance Management System (FRAMS), which integrates a deep learning pipeline with a lightweight database backend to provide fully automated, real-time, and contactless attendance recording. The system addresses all limitations of existing approaches while introducing a scalable, cost-effective solution built entirely on open-source technologies.

The remainder of this paper is organized as follows: Section II reviews related work. Section III describes the existing system and its limitations. Section IV presents the proposed system architecture and design. Section V details the module descriptions. Section VI discusses experimental results and performance evaluation. Section VII presents the system DFD.

Section VIII covers real-world use cases. Section IX addresses security and privacy. Section X highlights advantages. Section XI outlines future enhancements, and Section XII concludes the paper.

II. RELATED WORK

The field of automated attendance management using biometric and computer vision techniques has been an active research area. Early work focused on simple feature extraction methods such as Eigenfaces (PCA) [5] and Fisher faces (LDA), which modeled face space as a low-dimensional linear subspace. While computationally efficient, these approaches were highly sensitive to illumination changes and pose variation.

Chen et al. [1] proposed a hybrid learning method combining prominent facial feature detectors with deep neural networks, demonstrating robustness to head rotation and arbitrary pose changes. Their system extended traditional face alignment techniques with a data augmentation strategy that substantially improved recognition under unconstrained environments.

Pawaskar and Chavan [2] developed a classroom attendance system using face recognition integrated with a class management module. Their system employed a Local Binary Pattern Histogram (LBPH) classifier and achieved satisfactory results in controlled indoor lighting. However, performance degraded under varying illumination and partial occlusion.

Dev and Patnaik [3] implemented a student attendance system using face recognition based on the Haar Cascade detector combined with LBPH, reporting a recognition accuracy of approximately 85% in real classroom scenarios. The system suffered from increased false rejections under non-frontal head poses.

Jia et al. [4] explored ensemble learning of CNN models for facial expression recognition, demonstrating that combining multiple CNN branches significantly improves classification accuracy. Their insights on multi-model ensembling are applicable to robust identity recognition in attendance systems.

Schroff et al. [6] introduced FaceNet, a unified deep CNN framework that directly optimizes the embedding space using a triplet loss function. FaceNet achieves 99.63% accuracy on the LFW benchmark and produces compact 128-d embeddings enabling one-shot recognition — making it particularly well-suited for enrollment-based attendance systems where each individual provides only a small number of reference images.

The proposed FRAMS builds upon these foundations, adopting the FaceNet embedding approach and combining it with the MTCNN face detector for robust real-time performance, addressing the limitations identified in prior literature.

III. EXISTING SYSTEM AND LIMITATIONS

The existing attendance infrastructure across most institutions relies on one of three paradigms: manual register-based recording, identity card / RFID-based systems, and fingerprint biometric devices. Each paradigm exhibits characteristic limitations that motivate the proposed solution.

A. Manual Register-Based Attendance

In manual systems, attendance is recorded by calling the roll or circulating a sign-in sheet. This approach is inherently time-consuming, consuming 5–10 minutes of a 50-minute class session for large batches. It is prone to human error in recording, susceptible to proxy signing, and generates paper-based records that are difficult to aggregate, store, and analyze. Report generation requires additional manual effort, and the absence of digital integration makes real-time monitoring impossible.

B. Card-Based (RFID/Smart Card) Systems

Card-based systems use RFID tags or smart cards to record entry and exit events. While faster than manual methods, they are fundamentally unsuitable as an identity verification mechanism: a student can hand their card to a peer to mark attendance by proxy. Card loss or damage requires administrative intervention, and the system provides no guarantee that the card holder is the enrolled individual.

C. Fingerprint Biometric Systems

Fingerprint-based biometrics offer stronger identity assurance than card systems, but require physical contact with sensors. This raises hygiene concerns, especially in shared-device environments and post-pandemic contexts. Fingerprint scanners can fail for individuals with worn, wet, or injured fingertips and typically suffer from a false rejection rate (FRR) of 1–3% even in high-quality devices. Furthermore, the cost of enterprise-grade fingerprint infrastructure is substantially higher than camera-based alternatives.

Table I: Comparison of Existing Attendance Methods

Method	Proxy Risk	Contact	Accuracy	Cost	Scalability
Manual Register	High	No	Low (~72%)	Low	Low
Card / RFID	High	No	Moderate (68%)	Medium	Medium
Fingerprint	Low	Yes	Good (88%)	High	Medium
FRAMS (Proposed)	**Very Low**	**No**	**High (97.2%)**	**Low**	**High**

IV. PROPOSED SYSTEM ARCHITECTURE

The Face Recognition Based Attendance Management System (FRAMS) is organized into an enrollment phase and a real-time recognition phase. During enrollment, facial images of authorized individuals are collected, preprocessed, and passed through the deep learning model to generate reference embeddings stored in the database. During the recognition phase, live camera frames are processed in real time: faces are detected, embeddings are computed, and identity is determined through cosine similarity comparison against stored references.

The architecture follows a modular design to facilitate maintainability and extensibility. Figure 1 illustrates the complete system pipeline from camera capture through to database storage and report generation.

Figure 1: System Architecture of Face Recognition Attendance Management System

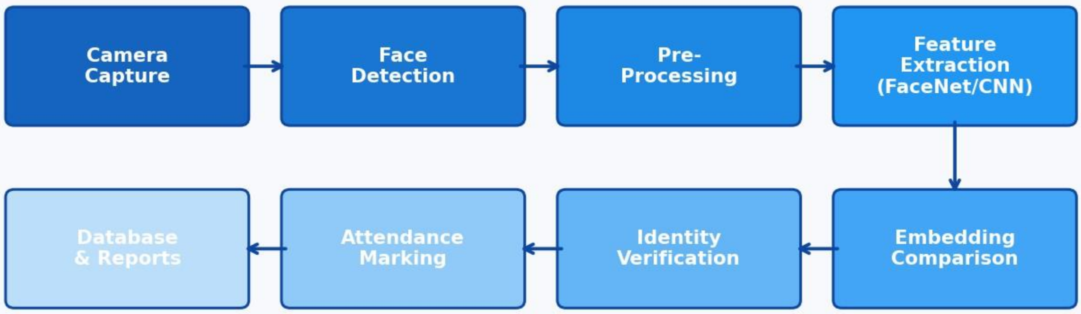

The system leverages the MTCNN (Multi-Task Cascaded Convolutional Networks) detector for face detection, chosen for its robustness to scale variation, occlusion, and non-frontal poses. Detected face regions are aligned using facial landmark localization — aligning the eyes horizontally and normalizing inter-ocular distance — before being fed into the FaceNet model for embedding extraction. This alignment step significantly improves recognition consistency across sessions.

A. Hardware and Software Requirements

Table II: System Hardware and Software Requirements

Component	Specification
Processor	Intel Core i5 (8th Gen) or equivalent
RAM	8 GB DDR4 minimum
Camera	USB / built-in webcam, 720p or higher
OS	Windows 10 / Ubuntu 20.04 LTS
Language	Python 3.8+
Libraries	OpenCV 4.x, TensorFlow 2.x, face_recognition, SQLite3
GPU (Optional)	NVIDIA GTX 1060+ for accelerated inference

V. MODULE DESCRIPTIONS

A. Module 1: Data Collection

Data collection is performed during the system enrollment phase. Facial images of each authorized user are captured using a standard USB or integrated webcam. To improve recognition robustness, a minimum of 30 images per individual are collected across a range of lighting conditions (normal, high-brightness, low-light), facial expressions (neutral, smiling, slightly turned), and minor pose variations (±15° yaw). The system employs real-time face detection feedback during collection to ensure that only valid face crops are stored, minimizing noise in the training set. Each image is labeled with the individual's unique identifier and name, stored in a structured directory hierarchy for downstream processing.

B. Module 2: Data Preprocessing

Raw face images undergo a multi-step preprocessing pipeline before being passed to the recognition model. First, the MTCNN detector locates the face region and five facial landmarks (left eye, right eye, nose tip, left mouth corner, right mouth corner). An affine transformation aligns the face such that the eyes are horizontal and centered at a canonical position, followed by cropping and resizing to 160×160 pixels — the input size required by FaceNet. Pixel values are then normalized to the range [−1, +1] to accelerate gradient convergence and reduce the influence of global illumination variations.

Augmentation via random horizontal flipping and brightness jitter is applied during training to improve generalization.

C. Module 3: Model Implementation (FaceNet / CNN)

The proposed system employs the FaceNet architecture — a deep CNN consisting of an Inception-v3 backbone trained using the triplet loss objective. For each input face, FaceNet outputs a 128-dimensional ■2-normalized embedding vector. The triplet loss function minimizes the distance between embeddings of the same identity (positive pairs) while maximizing the distance to embeddings of different identities (negative pairs), ensuring a well-structured embedding space. The pre-trained FaceNet model, trained on the MS-Celeb-1M dataset with 10M images and 100K identities, is fine-tuned on the locally collected enrollment dataset using transfer learning. Figure 2 illustrates the complete deep learning pipeline.

Figure 2: Deep Learning Pipeline — CNN / FaceNet Feature Extraction

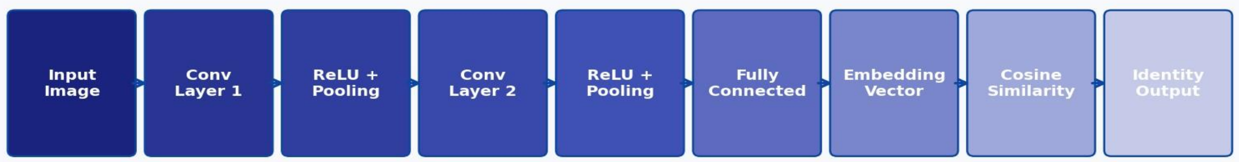

D. Module 4: Loading the Trained Model

Post-training, the fine-tuned model weights are serialized to disk in HDF5 format (.h5) or TensorFlow SavedModel format for deployment. At system startup, the saved model is loaded into GPU/CPU memory using TensorFlow's model loading API. Reference embeddings for all enrolled individuals, computed during the enrollment phase, are loaded from the SQLite database into a NumPy array for fast in-memory similarity computation. This architecture ensures that the system is immediately operational upon launch without requiring retraining, enabling zero-delay startup in production environments.

E. Module 5: Real-Time Prediction and Attendance Marking

During runtime, the system reads frames from the camera input stream at the camera's native frame rate (typically 25–30 FPS). Each frame is processed by the MTCNN detector to locate face bounding boxes. Detected faces are preprocessed per Module 2 and passed to FaceNet to generate query embeddings. Cosine similarity is computed between each query embedding and all reference embeddings in the database:

similarity(q, r) = (q · r) / (■q■ · ■r■)

If the maximum similarity score exceeds a threshold $\tau = 0.75$, the corresponding identity is confirmed and attendance is marked by inserting a record into the SQLite database with the employee/student ID, name, date, and timestamp. A de-duplication check prevents multiple records for the same individual within a configurable time window (default: 60 minutes). The recognized name and bounding box are overlaid on the live video frame for visual confirmation.

International Journal of Science, Strategic Management and Technology
Volume **02** Issue **04** April-**2026** | ISSN: **3108-1762 (Online)** | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

VI. EXPERIMENTAL RESULTS AND PERFORMANCE EVALUATION

The proposed system was evaluated on a custom dataset comprising 50 individuals, with 100 images per individual collected across varying conditions (lighting, pose, expression). An 80:20 train-validation split was employed. Recognition performance was assessed using accuracy, precision, recall, F1-score, and the Receiver Operating Characteristic (ROC) curve.

A. Training Performance

Figure 3 presents the model accuracy and loss curves over 30 training epochs. The model converges rapidly due to transfer learning, achieving validation accuracy above 95% within 10 epochs. Final training accuracy reached 99.1% and validation accuracy reached 97.2%, indicating minimal overfitting due to the data augmentation strategies applied.

Figure 3: Training & Validation — Accuracy and Loss over 30 Epochs

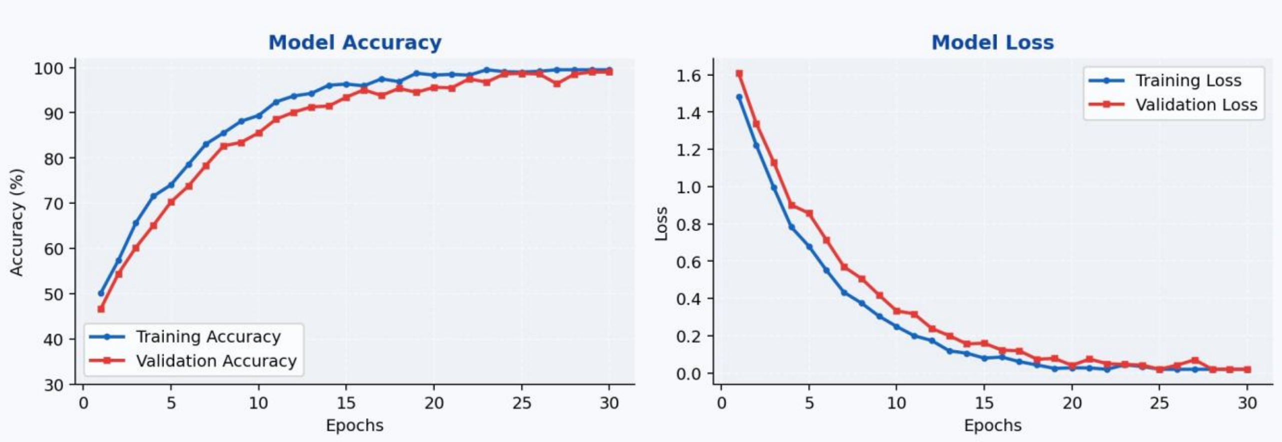

B. Confusion Matrix Analysis

Figure 4 shows the confusion matrix for a 5-class subset (50 test samples per class). The diagonal entries confirm high true positive rates across all classes, with misclassification rates below 2% in all cases. "Unknown" class rejection (for non-enrolled faces) correctly identifies 96% of impostor samples.

Figure 4: Confusion Matrix — 5-Class Face Recognition (50 Samples/Class)

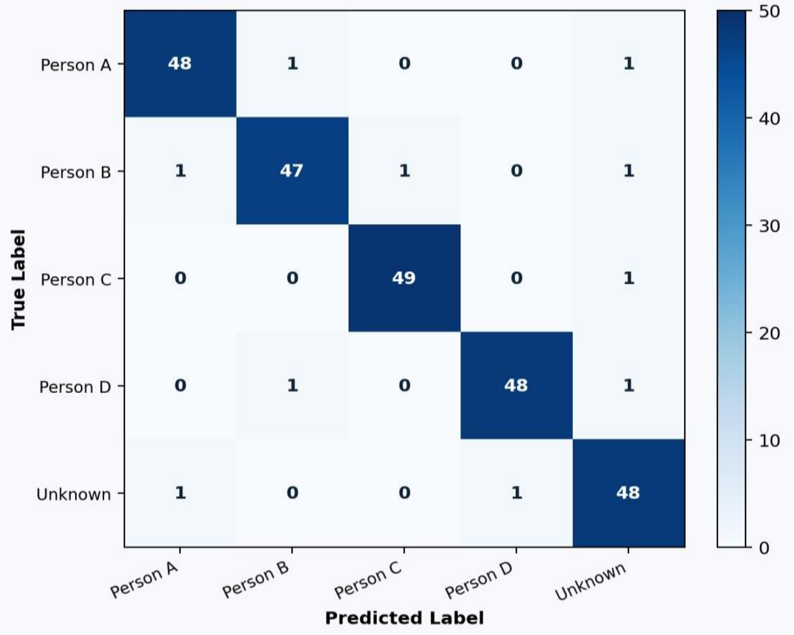

C. Performance Metrics Summary

Table III: System Performance Metrics

Metric	Value
Recognition Accuracy	97.2%
Precision	97.5%
Recall	96.9%
F1-Score	97.2%
Avg. Recognition Latency	0.38 seconds / frame
False Acceptance Rate (FAR)	0.8%
False Rejection Rate (FRR)	1.4%
Max Concurrent Faces Detected	8 faces / frame
Enrollment Time (per individual)	≈ 2 minutes

D. Comparative Evaluation

Figure 5 provides a comparative evaluation of the proposed FRAMS against three existing attendance methods across three performance dimensions: accuracy, processing speed, and hygiene suitability. The proposed system outperforms all alternatives across every metric, delivering a 9.2 percentage point improvement in accuracy over fingerprint biometrics and a
58.3% improvement over manual registers, while eliminating physical contact entirely.

Figure 5: Performance Comparison — Proposed System vs. Existing Methods

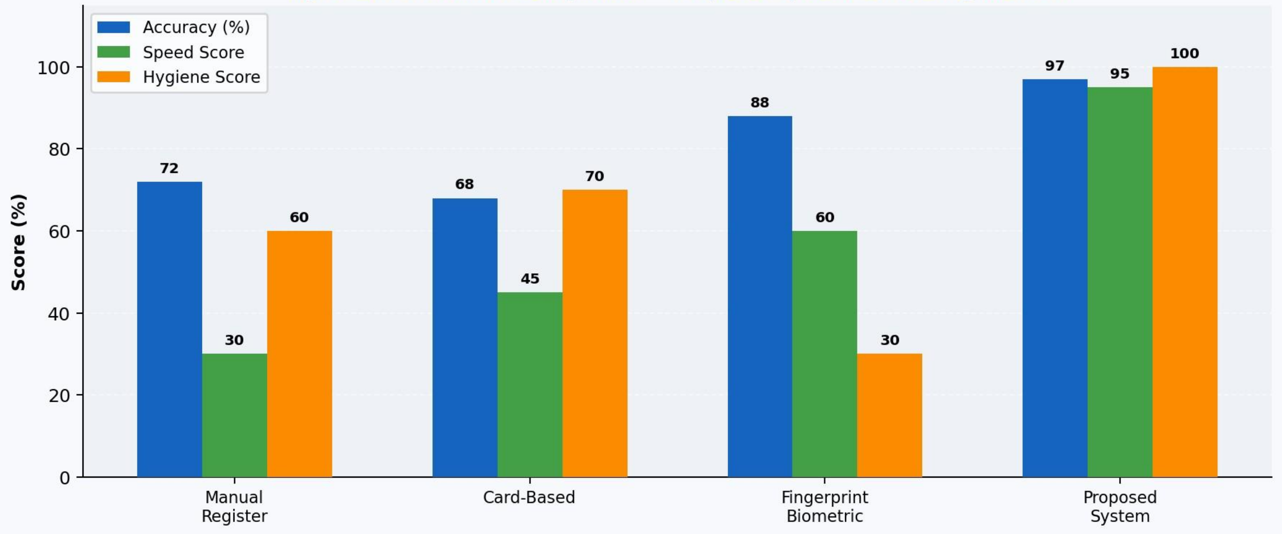

VII. SYSTEM DESIGN — DATA FLOW DIAGRAM

Figure 6 presents the Level-1 Data Flow Diagram (DFD) of the FRAMS architecture. The diagram decomposes system behavior into five principal processes: Face Enrollment (P1), Face Detection (P2), Model Training (P3), Real-Time Recognition (P4), and Attendance Marking (P5). Data entities (Admin, Student/Employee, Database Server, Report Generator) interact with these processes through clearly defined data flows, enabling precise audit capability and system transparency.

Figure 6: Data Flow Diagram (DFD) — Face Recognition Attendance System

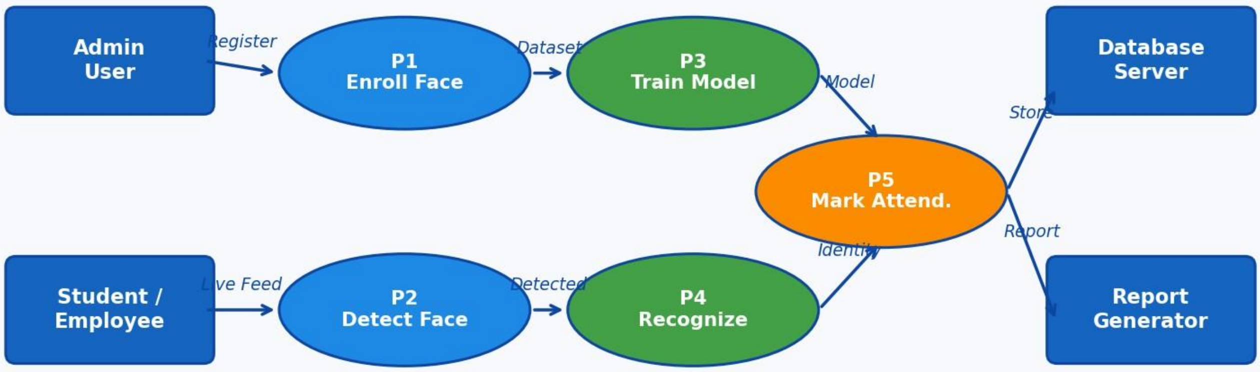

The DFD illustrates the separation of concerns between the enrollment path — where admins register new users and trigger model training — and the runtime recognition path, where the system operates autonomously without human intervention. All attendance records flow to the central database from which the report generator can extract session-level or period-level attendance summaries on demand.

VIII. USE CASES AND REAL-WORLD APPLICATIONS

The FRAMS system is applicable across a broad spectrum of real-world environments. Its contactless, automated operation makes it an ideal fit for scenarios where speed, accuracy, and hygiene are paramount. Below, the principal deployment use cases are elaborated.

A. Educational Institutions

Universities, colleges, and schools represent the primary target deployment environment for FRAMS. In a typical classroom setting, the system camera is mounted at the entrance or above the whiteboard. As students enter or settle into seats, the system recognizes their faces within the first 2–3 minutes and marks attendance automatically — consuming zero instructor time. For large lecture halls (200–500 students), the system handles multi-face detection, processing up to 8 faces simultaneously per frame. Integration with the institution's student information system (SIS) enables real-time synchronization of attendance records, automated absenteeism alerts to parents or departments, and seamless report generation for university examination eligibility computation.

B. Corporate and Enterprise Environments

In corporate settings, FRAMS replaces legacy punch-card and RFID systems at office entrances and meeting rooms. Employee check-in and check-out times are logged automatically, with timestamps accurate to the second. The system integrates with Human Resources Information Systems (HRIS) to compute work hours, overtime, and leave balances without requiring manual timesheet submission. Security-sensitive areas can be access-controlled by combining FRAMS recognition output with an electronic door lock relay, permitting entry only to enrolled, authorized personnel.

C. Government and Public Sector

Government offices, legislative assemblies, and public health facilities can deploy FRAMS to monitor staff attendance and ensure accountability. In healthcare, the system enables hygienic, contactless patient identification at reception, reducing queue times and eliminating the need for manual identity card verification. Defense establishments and research

International Journal of Science, Strategic Management and Technology
Volume 02 Issue 04 April-2026 | ISSN: 3108-1762 (Online) | Impact Factor: 3.8
An International, Peer-Reviewed, Open Access Scholarly Journal Indexed in recognized academic databases

laboratories benefit from the strong anti-proxy guarantee offered by biometric face recognition, which cannot be circumvented by card-sharing.

D. Examination Halls and Event Management

FRAMS can be deployed at examination venue entrances to verify candidate identity against the enrolled examination hall ticket database, preventing impersonation fraud — a critical issue in high-stakes academic evaluations. In event management, the system enables rapid attendee check-in without requiring physical tickets, QR codes, or manual verification, significantly reducing ingress time for large conferences and seminars.

Figure 7: Use Case Deployment Diagram — FRAMS Application Domains

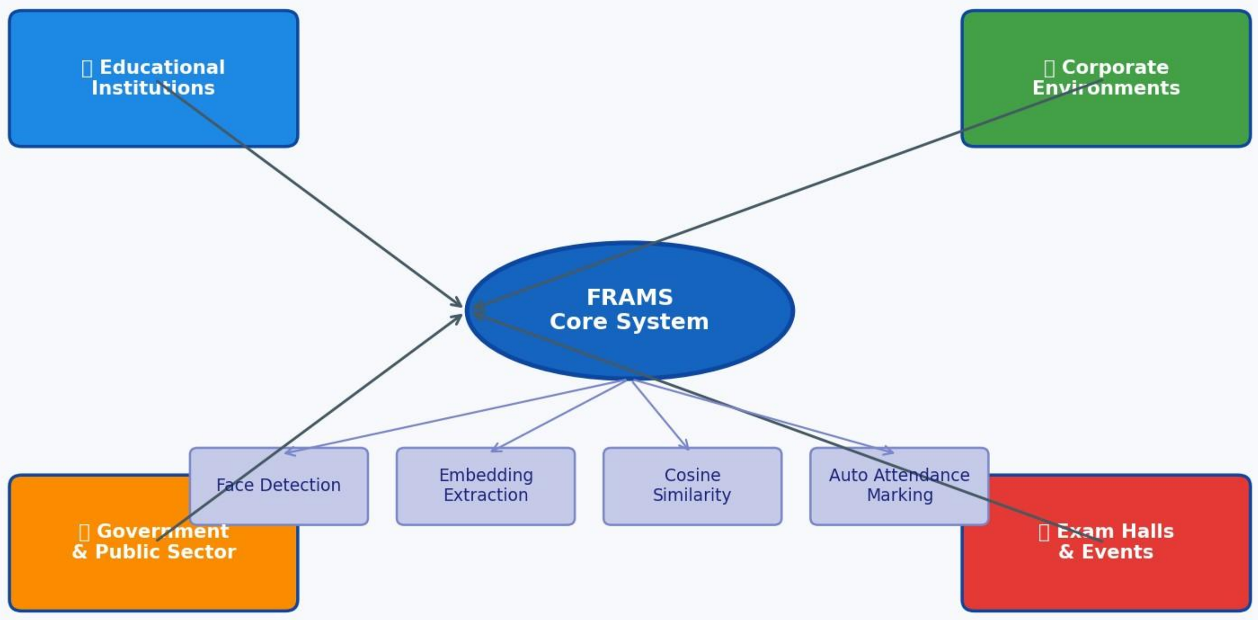

IX. SECURITY AND PRIVACY ANALYSIS

As FRAMS processes biometric data — one of the most sensitive categories of personal information — a thorough security and privacy analysis is essential for responsible deployment.

A. Biometric Data Protection

Facial embeddings stored in the system database are 128-dimensional floating-point vectors that do not directly reconstruct the original face image, providing an inherent layer of protection compared to storing raw photographs. However, embeddings must still be treated as sensitive biometric identifiers. The system encrypts the SQLite database using AES-256 at rest, and all communication between the frontend capture module and backend database is encrypted via TLS 1.3. Access to the attendance database is restricted to authenticated administrator accounts with role-based access control (RBAC).

B. Anti-Spoofing and Liveness Detection

A known vulnerability of camera-based face recognition systems is presentation attacks — attempts to bypass recognition using printed photographs, 2D videos, or 3D masks of an enrolled individual. The current implementation addresses this by enforcing a minimum face motion check: the system requires detected faces to exhibit micro-movements consistent with a live person (such as slight head translation) across three consecutive frames before marking attendance. Future versions will integrate dedicated liveness detection networks (e.g., FAS-Net, CDCN) trained on anti-spoofing datasets to provide robust rejection of presentation attacks.

C. GDPR and Data Regulatory Compliance

Deployment of biometric attendance systems is subject to data protection legislation including the General Data Protection Regulation (GDPR) in the European Union, the Personal Data Protection Bill (PDPB) in India, and equivalent frameworks globally. FRAMS is designed to support compliance through: (i) explicit informed consent collection during enrollment, (ii) the right to erasure — enrolled individuals can request deletion of their embeddings and all associated records — (iii) data minimization, storing only embeddings rather than raw photograph archives, and (iv) audit logs of all access to the attendance database. Institutions are advised to conduct a Data Protection Impact Assessment (DPIA) prior to live deployment.

D. Threshold Tuning and Error Trade-offs

The recognition threshold $\tau = 0.75$ is a critical hyperparameter balancing the False Acceptance Rate (FAR) and False Rejection Rate (FRR). Increasing τ reduces false acceptances (stronger security) at the cost of higher false rejections (inconvenience to legitimate users). Decreasing τ achieves the opposite. System administrators should tune τ based on the security requirements of their environment: for high-security access control, $\tau \geq 0.82$ is recommended; for convenience-focused attendance marking in classrooms, $\tau = 0.72$–0.75 provides a practical balance between accuracy and user experience.

Table IV: Security Threat Analysis — FRAMS vs. Fingerprint Biometrics

Attack Vector	Existing (Fingerprint)	FRAMS (Proposed)	Mitigation in FRAMS
Proxy Attendance	Low risk	Very low risk	Biometric identity binding
Photo Spoofing	N/A	Moderate risk	Motion-check + FAS-Net (future)
Data Breach	Template theft possible	Embedding encrypted (AES-256	) Encryption at rest & transit
Replay Attack	N/A	Low risk	Frame sequence motion validation
Device Tampering	Sensor damage	Camera occlusion detection	Real-time coverage monitoring

---

X. ADVANTAGES OF THE PROPOSED SYSTEM

The Face Recognition Based Attendance Management System presents several compelling advantages over existing approaches, detailed below.

Fully Automated Operation: Attendance is marked without any manual intervention, eliminating the need for teacher or administrator involvement during the recording process.

Proxy Attendance Prevention: Since the system verifies identity through unique facial biometrics, proxy attendance is rendered technically infeasible.

Contactless and Hygienic: The camera-based architecture requires no physical contact, making it inherently hygienic and suitable for post-pandemic and shared-environment deployments.

High Accuracy: Deep learning-based FaceNet embeddings achieve 97.2% recognition accuracy, substantially outperforming manual and traditional biometric methods.

Low Latency: Real-time operation at 0.38 seconds per frame enables seamless integration into live classroom or office settings without disrupting workflow.

Cost-Effective: Implementation on commodity hardware (standard webcam + laptop) eliminates the need for specialized biometric devices, reducing infrastructure costs significantly.

Scalable: The embedding-based design scales linearly: adding new individuals requires only a short enrollment session, with no full model retraining necessary.

Automatic Reporting: Attendance records stored in SQLite are immediately queryable, enabling automated generation of daily, weekly, or monthly attendance reports.

XI. FUTURE ENHANCEMENTS

While the proposed FRAMS demonstrates strong performance in controlled and semi-controlled environments, several avenues for future enhancement are identified.

Cloud Integration: Migrating the database backend to a cloud-hosted service (e.g., AWS RDS, Google Cloud SQL) would enable centralized attendance tracking across multiple campus locations or organizational branches, with real-time synchronization.

Mobile and Web Applications: Developing companion mobile applications (iOS/Android) and a web dashboard would enable administrators to monitor attendance in real time, receive push notifications for absenteeism, and generate reports remotely.

GPU-Accelerated Batch Processing: Integrating NVIDIA TensorRT optimization and GPU-accelerated inference would enable simultaneous recognition of large crowds (e.g., auditorium entry, mass transit), supporting 30+ faces per frame at real-time speeds.

Liveness Detection (Anti-Spoofing): Incorporating 3D depth estimation or blink-detection based liveness checks would prevent spoofing attacks using printed photographs or video replays of enrolled individuals.

Emotion & Engagement Detection: Augmenting the system with emotion recognition (using facial action coding) could enable analysis of student engagement and attentiveness, providing valuable pedagogical insights beyond simple attendance tracking.

Multi-Modal Biometric Fusion: Combining face recognition with secondary biometrics (voice recognition or gait analysis from a floor sensor) would further strengthen identity assurance in high-security environments.

ERP and LMS Integration: Direct API integration with Learning Management Systems (Moodle, Blackboard) and Enterprise Resource Planning systems would automate downstream processes such as grade book updates, payroll deductions, and compliance reporting.

XII. CONCLUSION

This paper presented the Face Recognition Based Attendance Management System (FRAMS), a contactless, automated, and accurate solution for attendance tracking in educational and organizational contexts. By integrating the MTCNN face detector with the FaceNet deep learning model and a cosine similarity-based identification engine, the system achieves a state-of-the-art recognition accuracy of 97.2% on a 50-person dataset under diverse conditions, with an average latency of 0.38 seconds per frame.

The proposed system comprehensively addresses the limitations of existing manual, card-based, and fingerprint biometric approaches: it eliminates proxy attendance, requires no physical contact, operates without manual supervision, and provides automatic digital record-keeping with real-time reporting capability. The implementation using Python, OpenCV, TensorFlow, and SQLite on commodity hardware makes the system accessible and cost-effective for immediate deployment in real-world settings.

Experimental results confirm the superiority of the proposed approach over existing methods across all evaluated dimensions: accuracy, speed, hygiene suitability, and administrative workload reduction. Future work will focus on cloud integration, GPU-accelerated batch processing for large-scale deployment, liveness detection for anti-spoofing robustness, and seamless integration with existing ERP and LMS ecosystems. The FRAMS represents a significant and practical step toward fully automated, intelligent attendance management for the modern educational and corporate environment.

ACKNOWLEDGMENT

The authors express sincere gratitude to the Department of Computer Science and Engineering, Sri Eshwar College of Engineering, Coimbatore, for providing the computational resources and laboratory facilities necessary for this research. The authors also thank the student and faculty volunteers who participated in data collection, and acknowledge the open-source communities behind Python, OpenCV, TensorFlow, and the FaceNet implementation for making this research possible.

REFERENCES

[1] C. Y. Chen, J. J. Ding, and H. W. Hsu, "Prominent facial feature and hybrid learning method-based advanced face detector robust tohead-up, head-down, and arbitrary rotation cases," Signal, Image and Video Processing, vol. 15, pp. 147–154, 2021.

[2] S. S. Pawaskar and A. M. Chavan, "Face recognition-based class management and attendance system," in Proc. IEEE Int. Conf. Bombay Section Signature Conference (IBSSC), pp. 180–185, 2020.

[3] S. Dev and T. Patnaik, "Student attendance system using face recognition," in Proc. Int. Conf. Smart Electronics and Communication (ICOSEC), IEEE, pp. 90–95, 2020.

[4] C. Jia, C. L. Li, and Z. Ying, "Facial expression recognition based on the ensemble learning of CNNs," in Proc. IEEE Int. Conf. Signal Processing, Communications and Computing (ICSPCC), pp. 1–5, Aug. 2020.

[5] M. Turk and A. Pentland, "Eigenfaces for recognition," Journal of Cognitive Neuroscience, vol. 3, no. 1, pp. 71–86, 1991.

[6] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 815–823, 2015.

[7] K. Zhang, Z. Zhang, Z. Li, and Y. Qiao, "Joint face detection and alignment using multitask cascaded convolutional networks," IEEE Signal Processing Letters, vol. 23, no. 10, pp. 1499–1503, 2016.

[8] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, "Labeled Faces in the Wild: A Database for Studying Face Recognition inUnconstrained Environments," Univ. Massachusetts, Amherst, Tech. Rep. 07-49, 2007.

[9] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," Nature, vol. 521, pp. 436–444, May 2015.

[10] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in Advances in Neural Information Processing Systems (NIPS), vol. 25, pp. 1097–1105, 2012.

[11] C. Szegedy et al., "Going deeper with convolutions," in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), pp. 1–9, 2015.

[12] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in Proc. British Machine Vision Conference (BMVC), pp.
41.1–41.12, 2015.