Advances in Consumer Research
https://acr-journal.com/
Volume-3 | Issue-2 | Feb 2026
Original Researcher Article

Smart-Campus: A Privacy-Preserving Distributed Edge-AI Framework for Hierarchical Face-Based Attendance and Natural Language Analytics

Dr. R. Poornachandran[1], Dr. R. Rameshbabu[2], S. Santhiya[3]
[1]Assistant Professor , Department Of Electronics and Communication Engineering, V.S.B Engineering College Karur – 639111
Email ID : rrameshbabu15@gmail.com
[2]Assistant Professor , Department Of Electronics And Communication Engineering, V.S.B Engineering College Karur – 639111
Email ID :surenecevsb@gmail.com
[3]M.E (Applied electronics) VSB Engineering College, Karur Karur,
Email ID :santhiyakabi@gmail.com

ABSTRACT
Automated attendance systems have become an integral component of modern smart campus infrastructures. However, most existing solutions rely on centralized process- ing pipelines that introduce high latency, network dependency, and significant privacy risks due to continuous transmission of sensitive biometric data. To address these challenges, this paper proposes a distributed edge-AI framework for real-time attendance monitoring, where facial recognition is performed locally at the classroom level and only essential metadata is synchronized with a central server.
The proposed system introduces a hierarchical recognition strategy that prioritizes staff identification over student matching, thereby reducing cross-identity misclassification during simulta- neous entries. Furthermore, a natural language-driven analytics interface is integrated to enable non-technical administrators to query attendance records conversationally. A real-time prototype is implemented using open-source tools and evaluated under practical conditions. Experimental results demonstrate low- latency recognition, reliable identity segregation, and improved usability.
The proposed framework bridges the gap between intelligent edge computing and real-world campus management, offering a scalable, privacy-aware, and interactive attendance solution.
Keywords: Terms—Edge Computing, Face Recognition, Smart Campus, Attendance Automation, Distributed Systems, Privacy- Preserving AI, Natural Language Interfaces

1. INTRODUCTION:

Educational institutions increasingly rely on automated at- tendance systems to enhance administrative efficiency and minimize manual intervention. Traditional approaches such as RFID cards, fingerprint scanners, and manual roll calls suffer from multiple limitations, including susceptibility to proxy attendance, hygiene concerns, device dependency, and high operational overhead.

Recent advancements in computer vision and deep learn- ing have enabled face recognition-based attendance systems, offering contactless and user-friendly interaction. However, most existing implementations adopt centralized cloud-based architectures, where captured facial data is transmitted to remote servers for processing. This design introduces critical drawbacks such as increased latency, vulnerability to network failures, scalability constraints, and severe privacy risks asso- ciated with transmitting and storing sensitive biometric data.

Edge computing has emerged as a promising paradigm to mitigate these limitations by relocating computational in- telligence closer to the data source. By executing inference locally, edge-based systems significantly reduce response time, enhance robustness, and prevent raw data exposure. This paradigm is particularly suitable for smart campus environ- ments, where real-time responsiveness and data confidentiality are crucial.

In this work, we propose a distributed edge-AI attendance framework that performs real-time face recognition at the classroom level. Unlike conventional systems, the proposed framework incorporates a hierarchical recognition mechanism that prioritizes staff identification before student matching, preventing cross-category misclassification. Additionally, the system integrates a natural language-based analytics mod- ule that allows administrators to retrieve attendance insights through conversational queries.

The proposed system is implemented as a real-time proto- type and evaluated under practical operating conditions. The results demonstrate the feasibility of deploying intelligent, privacy-aware attendance systems using lightweight edge de- vices.

Contributions

The main contributions of this paper are summarized as follows:

How to cite : Dr. R. Poornachandran, Dr. R. Rameshbabu, S. Santhiya , Smart-Campus: A Privacy-Preserving Distributed Edge-AI Framework for Hierarchical Face-Based Attendance and Natural Language Analytics. *Advances in Consumer Research*. 2026;3(2): 319-323

We propose a distributed edge-based attendance frame- work that eliminates dependency on centralized biometric processing and enhances privacy.

We introduce a hierarchical identity recognition strategy that prioritizes staff detection before student matching, preventing cross-identity misclassification.

We integrate a natural language-driven analytics interface that enables intuitive, conversational access to attendance records.

We implement a real-time prototype to validate the feasibility of the proposed architecture under practical conditions.

We present a comprehensive experimental evaluation fo- cusing on latency, usability, and system reliability.

stem Overview

The proposed attendance framework is designed for smart campus environments, where multiple classrooms operate as independent yet interconnected units. Each classroom is treated as an autonomous edge node responsible for real-time face detection, recognition, and preliminary decision-making. Unlike centralized architectures, the proposed system per- forms all biometric processing locally. Only essential atten- dance metadata, such as identity, timestamp, and class ID, is synchronized with the central server. This design significantly reduces network dependency and prevents exposure of raw

facial data.

The system supports two user categories: staff and students. Staff identities are globally visible across all edge nodes, while student identities are restricted to their assigned classrooms. This selective visibility enables efficient matching and prevents unnecessary memory usage.

In addition, an AI-driven natural language interface is integrated into the administrative dashboard, allowing non- technical users to retrieve attendance insights through con- versational queries.

2. METHODOLOGY

The proposed attendance framework follows a modular pipeline consisting of face acquisition, preprocessing, feature extraction, hierarchical matching, attendance logging, and an- alytics.

Initially, live video frames are captured through a camera attached to the edge node. Each frame undergoes preprocess- ing to normalize illumination and reduce noise. Face detection is then performed, followed by feature extraction using a deep embedding model that produces compact numerical represen- tations.

The extracted features are compared against two identity pools: a global staff pool and a class-specific student pool. A prioritized matching strategy is employed, where staff matching is always attempted first. If no staff match is found, student matching is performed.

Upon successful identification, attendance is logged locally and synchronized with the central server. The updated records are visualized on the administrative dashboard and can be queried using natural language commands.

System Architecture

The overall architecture of the proposed framework is illustrated in Fig. 1. The system follows a distributed edge computing paradigm, where each classroom operates as an independent processing unit.

The sensing layer consists of cameras deployed at classroom entry points. These cameras continuously capture live video streams and forward frames to the edge intelligence layer. The edge layer performs face detection, feature encoding, and hierarchical identity matching.

Each edge node maintains a lightweight local cache contain- ing class-specific student encodings and a global pool of staff

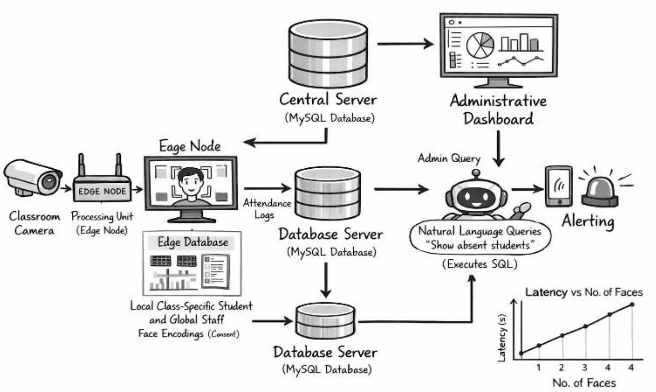

Fig. 1. Proposed distributed edge-AI architecture for smart campus attendance monitoring.

encodings. This selective loading strategy reduces memory overhead and improves recognition speed.

The central server aggregates attendance logs from multiple edge nodes and stores them in a relational database. The application layer consists of a web-based dashboard and a natural language interface, enabling real-time monitoring and intuitive interaction.

Algorithm Description

The attendance logging mechanism is governed by a hier- archical decision algorithm designed to ensure accuracy, effi- ciency, and administrative correctness. The complete workflow is illustrated in Fig. 2.

Initially, the system verifies whether the current time falls within predefined operational hours. This prevents unneces- sary computation during inactive periods. Once activated, the camera module captures live frames, which are passed to the face detection module.

Detected faces are encoded into fixed-length feature vectors. These encodings are first compared against the staff identity pool. If a match is found, attendance is immediately logged as staff, and no further comparisons are performed for that face. If no staff match is found, the system proceeds to compare the encoding with the class-specific student pool. If a valid match is identified, student attendance is logged. If no match is found in either pool, the face is classified as unknown and

ignored.

This prioritized matching strategy significantly reduces false positives and prevents staff-student cross-identification.

Implementation

To validate the proposed framework, a real-time prototype was developed using open-source tools and libraries. A stan- dard laptop equipped with a webcam was used to simulate a classroom-level edge node.

The system is implemented in Python and utilizes OpenCV for video acquisition and frame preprocessing. Face detection is performed using a HOG-based detector, followed by feature

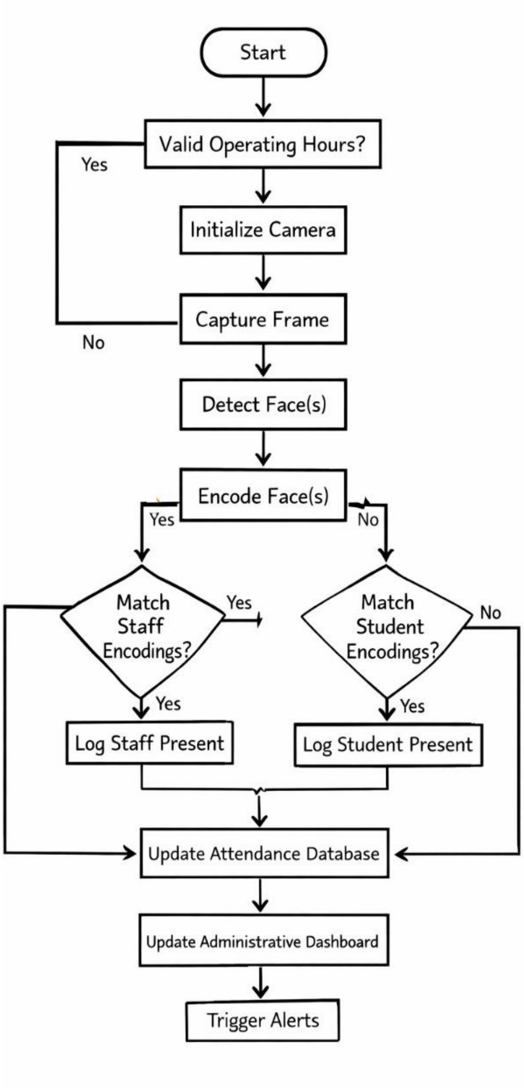

Fig. 2. Operational workflow of the hierarchical recognition and attendance logging process.

extraction using a deep embedding model that generates 128- dimensional face encodings.

The edge node maintains an in-memory cache of class-specific student encodings and a global pool of staff encodings. This selective loading mechanism reduces memory usage and accelerates identity matching. Once a face is identified, attendance is logged locally and synchronized with the central database.

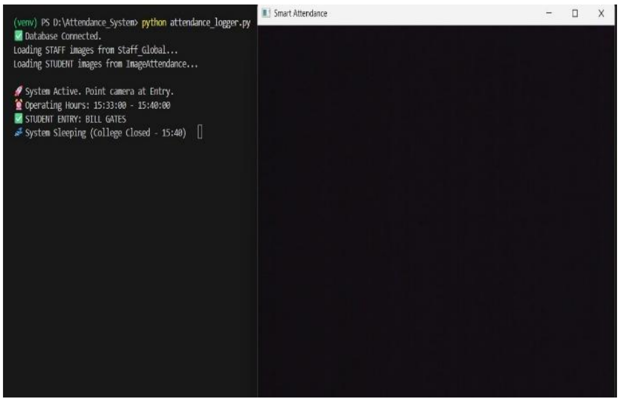

Fig. 3 illustrates the experimental edge node prototype executing real-time recognition tasks.

Fig. 4 presents a sample recognition output, showing de- tected faces, identity labels, and category-based classification.

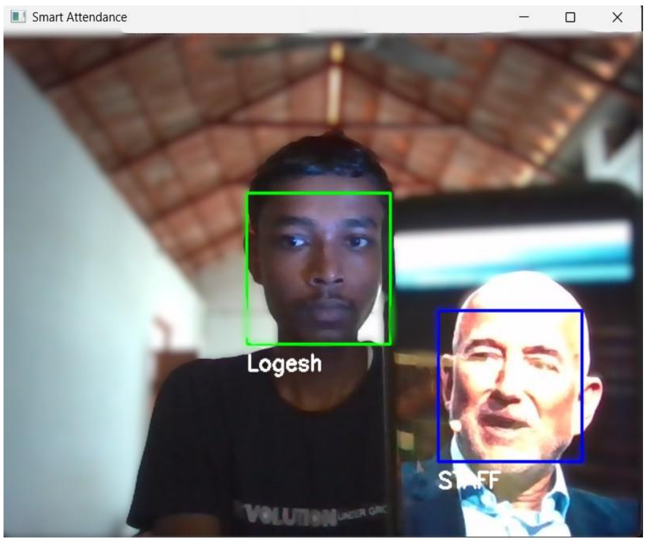

A web-based dashboard was implemented using Streamlit, allowing administrators to visualize attendance trends, access individual reports, and interact with the AI assistant. The nat- ural language interface translates user queries into structured database operations.

Fig. 3. Experimental edge node prototype implemented on a standard laptop.

Fig. 4. Sample recognition output showing detected faces, identity labels, and classification results.

Experimental Setup

The proposed framework was evaluated using a real-time prototype deployed on a standard laptop configured to emulate a classroom-level edge node. The device was equipped with an integrated webcam, and all recognition tasks were performed locally to validate the feasibility of edge-based deployment.

A controlled dataset consisting of registered student and staff identities was created for evaluation. Each identity was enrolled using multiple facial images captured under varying illumination and pose conditions to enhance robustness. The generated embeddings were stored locally based on the as- signed classroom configuration.

The system was tested under real-world conditions, includ- ing partial occlusions, variable lighting, and multi-face scenar- ios. Key performance indicators such as recognition latency, identity segregation accuracy, and system responsiveness were recorded.

Network connectivity was enabled only for metadata syn- chronization with the central server, while all biometric com- putations were retained at the edge level. This setup reflects realistic deployment conditions for privacy-preserving smart campus systems.

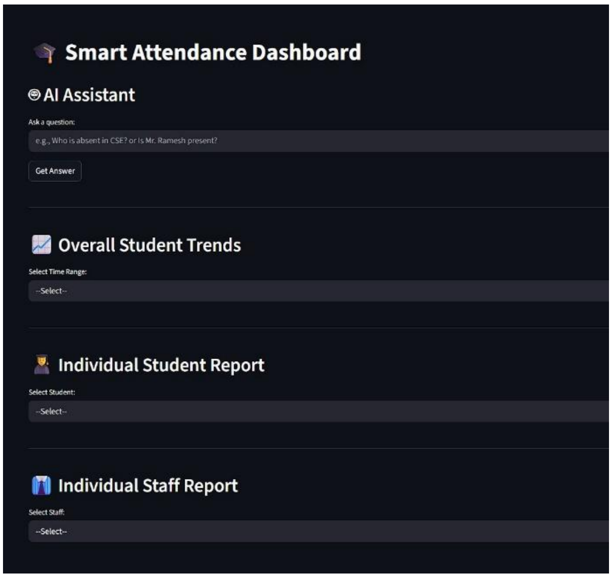

Fig. 5. Web-based administrative dashboard with real-time visualization and AI-assisted querying.

3. RESULTS AND DISCUSSION

The proposed distributed edge-AI framework demonstrated reliable real-time performance under practical operating con- ditions. The hierarchical matching strategy successfully elim- inated staff-student cross-identification, which is a common limitation in conventional systems.

Fig. 5 shows the web-based administrative dashboard, which provides real-time visualization of attendance statistics, indi- vidual reports, and natural language-driven query support. This interface significantly reduces the operational complexity for non-technical users.

To analyze system scalability, recognition latency was mea- sured by increasing the number of faces present in the frame. As shown in Fig. 6, the latency increases linearly while main- taining sub-second response time, confirming the suitability of the proposed approach for real-time classroom environments. The selective identity loading mechanism further improved matching speed by restricting comparisons to relevant identity subsets. Moreover, local processing of biometric data mini-

mized network dependency and enhanced privacy.

Overall, the experimental results validate the effectiveness of the proposed system in achieving low-latency operation, accurate identity segregation, and improved usability.

Deployment

To demonstrate real-world usability, the trained recognition pipeline was integrated into an interactive web application us- ing Streamlit. The interface enables administrators to monitor attendance records, visualize trends, and retrieve reports in real time.

The system performs preprocessing, feature extraction, and identity matching locally at the edge node, while metadata

Fig. 6. Latency performance of the proposed system under increasing number of detected faces.

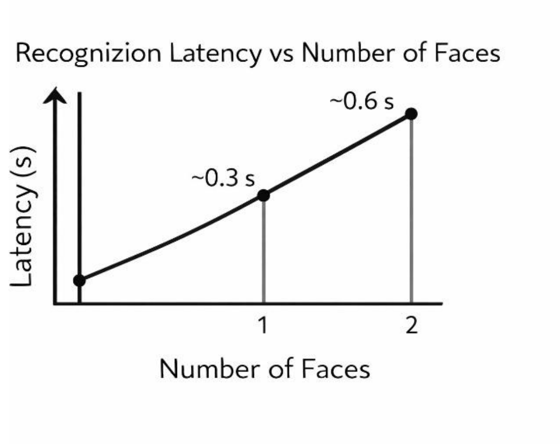

synchronization is handled in the background. The natural language interface allows users to issue queries such as ”Who is absent today?” or ”Is a particular staff member present?” without requiring technical knowledge.

This deployment strategy ensures accessibility, ease of use, and seamless integration into existing campus workflows.

4. LIMITATIONS

Although the proposed system demonstrates strong per- formance, certain limitations must be acknowledged. The current prototype has been evaluated on a limited number of identities, which may not fully represent large-scale campus environments.

The framework primarily focuses on entry-point-based at- tendance logging and does not track continuous presence throughout the day. Additionally, the system has been eval- uated using frontal face views, and extreme pose variations may impact recognition accuracy.

Finally, deep learning-based models require higher com- putational resources compared to classical approaches, which may affect deployment on ultra-low-power edge devices.

5. CONCLUSION

This paper presented a privacy-preserving distributed edge- AI framework for automated attendance monitoring in smart campus environments. By relocating biometric processing from centralized servers to localized edge nodes, the proposed system significantly reduces latency, enhances scalability, and improves data confidentiality.

The introduction of a hierarchical recognition strategy en- sures accurate identity segregation between staff and students, while the natural language-based analytics module transforms the system into an intelligent administrative assistant.

A real-time prototype was implemented and evaluated under practical conditions, demonstrating low-latency performance, reliable identity classification, and improved usability. The re- sults confirm the feasibility of deploying intelligent attendance systems using edge computing paradigms.

6. FUTURE WORK

Future work will focus on extending the framework to include exit-point monitoring for continuous presence tracking and bunking detection. Federated learning mechanisms can be explored to enable collaborative model improvement across multiple edge nodes without sharing raw biometric data.

Further optimization will target deployment on low-power embedded devices such as Raspberry Pi and NVIDIA Jetson platforms. Additionally, multimodal biometric fusion, incorpo- rating gait or voice recognition, can be investigated to enhance robustness under challenging conditions.

REFERENCES

[1] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, “DeepFace: Closing the Gap to Human-Level Performance in Face Verification,” in 2014 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2014, pp. 1701–1708.

[2] S. Teerapittayanon, B. McDanel, and H. T. Kung, “Distributed Deep Neural Networks over the Cloud, the Edge and End Devices,” in 2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS), 2017, pp. 328–339.

[3] M. Satyanarayanan, “The Emergence of Edge Computing,” Computer, vol. 50, no. 1, pp. 30–39, Jan. 2017.

[4] A. Graves and J. Schmidhuber, “Framewise Phoneme Classification with Bidirectional LSTM Networks,” in Proceedings of the International Joint Conference on Neural Networks (IJCNN), vol. 4, 2005, pp. 2047–2052.

[5] A. Vaswani et al., “Attention Is All You Need,” in Advances in Neural Information Processing Systems (NeurIPS), vol. 30, 2017, pp. 5998–6008

.