



Classroom Attendance System Using Facial Recognition

¹Dhanush Dalmotra, ²Sushant Dhar, ³Mehak Bhagat, ⁴Er. Sheetal Gandotra, ⁵Dr. Bhawna Sharma

^{1,2,3}Students, Dept. of Computer Science & Engineering Government College of Engineering and Technology, Jammu J&K, India

⁴Associate Professor, Dept. of Computer Science & Engineering Government College of Engineering and Technology, Jammu J&K, India

⁵Professor, Dept. of Computer Science & Engineering Government College of Engineering and Technology, Jammu J&K, India

Abstract: Traditional classroom attendance methods are time-consuming, error-prone, and susceptible to proxy attendance. This paper presents an automated classroom attendance system that leverages deep learning-based facial recognition to address these limitations. The proposed system employs InsightFace with the Buffalo_L model—an ArcFace-based deep neural network—to detect faces in a classroom photograph, extract 512-dimensional face embeddings, and match them against a pre-built student embedding database using cosine similarity. A Flask-based web application provides a teacher-friendly interface for uploading classroom images, generating attendance records in CSV format, viewing attendance history, and monitoring cumulative attendance percentages. The system operates entirely on-premises without dependency on cloud APIs, making it deployable in resource-constrained institutional environments. The proposed approach eliminates proxy attendance, reduces manual effort, and provides actionable attendance analytics for educational administrators. Experimental evaluation demonstrates the effectiveness of the system in realistic classroom scenarios.

Index Terms — face recognition, InsightFace, ArcFace, Attendance Automation, Cosine Similarity, Deep Learning, Flask, OpenCV, face embeddings.

I. INTRODUCTION

Attendance monitoring is a fundamental administrative task in educational institutions. Accurate attendance records are essential for tracking student engagement, satisfying regulatory requirements, and ensuring academic accountability. Despite its importance, attendance collection in most institutions still relies on manual methods—calling roll, passing sign-in sheets, or maintaining handwritten registers—all of which are inefficient and unreliable.

Manual attendance recording consumes a significant portion of lecture time. Studies suggest that in a 50-minute class, roll call can consume 5 to 10 minutes, directly reducing the time available for instruction [7]. Furthermore, manual systems are vulnerable to proxy attendance, where absent students persuade peers to mark them present. Traditional biometric alternatives such as RFID cards [8] and fingerprint scanners require dedicated hardware at every entry point, introduce queues, and in the case of fingerprint systems, raise hygiene concerns.

Vision-based attendance systems that recognise students from photographs or video frames offer a contactless, hardware-light solution. Early computer vision approaches such as Eigenfaces, Fisherfaces,

and Local Binary Pattern Histograms (LBPH) achieved limited accuracy in uncontrolled environments due to sensitivity to illumination, occlusion, and pose variation [9]. The emergence of deep learning, and specifically large-margin loss functions such as ArcFace [1], has dramatically improved face recognition accuracy, enabling practical deployment in real-world settings.

This paper presents an end-to-end automated classroom attendance system built on the following components:

1. **InsightFace Buffalo_L** [2] for face detection and 512-dimensional ArcFace embedding extraction.
2. **Cosine similarity matching** against a pre-built per-student embedding database.
3. **Flask web application** providing a complete teacher interface for image upload, attendance download, history browsing, and cumulative percentage analytics.

The system is fully self-contained, requires no cloud services, and is deployable on standard academic hardware with optional GPU acceleration. The remainder of this paper is organised as follows. Section II reviews related work. Section III describes the system architecture. Section IV details the methodology and algorithms. Section V covers implementation details. Section VI presents experimental results. Section VII discusses current limitations and future work. Section VIII concludes the paper.

II. RELATED WORK

2.1 Classical face recognition approaches

Early attendance systems based on machine vision employed handcrafted feature descriptors. The Eigenfaces method by Turk and Pentland [9] used Principal Component Analysis (PCA) to project face images into a lower-dimensional eigenspace. Fisherfaces extended this with Linear Discriminant Analysis (LDA) to improve class separability. LBPH [6] encoded local texture patterns and achieved better robustness to illumination. However, all these methods rely on shallow feature representations and degrade significantly under real classroom conditions involving varied lighting, partial occlusions, and non-frontal poses.

2.2 Deep learning face recognition

The introduction of deep convolutional neural networks (CNNs) for face recognition transformed the field. DeepFace [4] demonstrated near-human accuracy on the LFW benchmark using a nine-layer CNN trained on four million images. FaceNet [3] introduced the triplet loss to learn a compact embedding space where ℓ_2 distance directly corresponds to face similarity. ArcFace [1] further improved discriminability by incorporating an additive angular margin penalty in the softmax loss, achieving state-of-the-art performance across multiple benchmarks. InsightFace [2] packages ArcFace-trained models with the SCRFD detector into a unified, production-ready framework.

2.3 Automated attendance systems

Several works have proposed vision-based attendance systems. Ge et al. [10] demonstrated group face recognition from a single classroom photograph using a Faster R-CNN detector combined with a ResNet recognition backbone. Kaur and Singh [11] implemented an attendance system using FaceNet [3] embeddings with a ℓ_2 -nearest neighbour classifier. Compared to these prior works, the proposed system integrates InsightFace's superior Buffalo_L model, implements a complete web-based management interface, and adds cumulative attendance percentage analytics—features absent from most research prototypes.

III. SYSTEM ARCHITECTURE

The system is organised into four loosely coupled layers.

3.1 Frontend layer

The frontend consists of five HTML templates rendered by the Flask Jinja2 engine:

- login.html — teacher authentication form.
- index.html — classroom image upload form.

- result.html — annotated image with present/absent lists, face counts, and action buttons.
- history.html — chronological list of CSV files from past attendance.
- percentage.html — tabular cumulative attendance statistics per student.

3.2 Backend layer

app.py implements the Flask application. It handles session-based authentication, routes HTTP requests to the appropriate handlers, invokes the recognition engine on uploaded images, and serves downloadable CSV files.

3.3 Recognition engine

attendance_engine.py encapsulates all computer vision logic: face detection, embedding extraction, cosine similarity matching, annotated image generation, and CSV writing.

3.4 Database construction module

build_database.py is an offline preprocessing script that iterates over a structured student image dataset, extracts per-image embeddings, computes an average embedding per student, and serialises the result to face_database.npy.

3.5 Persistent storage

All artifacts are stored on the local filesystem: face_database.npy (embedding dictionary), csv_files/attendance.csv (latest session), attendance_history (timestamped historical records), and attendance_stats/attendance_percentage.csv (cumulative statistics).

IV. METHODOLOGY

4.1 Attendance generation flow

Fig. 1 illustrates the four-phase sequential flow a teacher follows to generate attendance.




Figure 1: Attendance generation: four-phase sequential flow

4.2 Face detection and embedding extraction

The system uses InsightFace's Buffalo_L model pack, which bundles the SCRFD face detector with an ArcFace recognition backbone. Initialisation is performed as follows:

```

from insightface.app import FaceAnalysis
app = FaceAnalysis(name='buffalol')
app.prepare(ctxid=0)
0=GPU; -1=CPU
  
```

For each uploaded classroom image I , the detector returns a list of face objects. Each object carries a bounding box $\mathbf{b} = (x_1, y_1, x_2, y_2)$ and a 512-dimensional embedding vector $\mathbf{e} \in \mathbb{R}^{512}$.

4.3 Database construction

The offline database building algorithm is described in the following Algorithm:

Student Embedding Database Construction Algorithm

Require: Dataset directory dataset/{StudentName}/{images}

Ensure: face_database.npy \leftarrow dict (name: \bar{e})

1: **for all** student folder s in dataset/ **do**

2: $E_s \leftarrow \emptyset$

3: **for all** image I_k in folder s **do**

4: $\mathcal{F} \leftarrow \text{app.get}(I_k)$

5: **if** $|\mathcal{F}| > 0$ **then**

6: $\mathbf{e}_k \leftarrow \mathcal{F}[0].\text{embedding}$

7: $E_s \leftarrow E_s \cup (\mathbf{e}_k)$

```

8: end if
9: end for
10: if  $|E_s| > 0$  then
11:    $\bar{e}_s \leftarrow \frac{1}{|E_s|} \sum_{e \in E_s} e$ 
12: face_db[s]  $\leftarrow \bar{e}_s$ 
13: end if
14: end for
15: save face_db to face_database.npy

```

The average embedding \bar{e}_s for student s is:

$$\bar{e}_s = \frac{1}{N} \sum_{k=1}^N e_k \quad (1)$$

where N is the number of valid face images and e_k is the embedding from the k -th image.

4.4 Recognition and attendance marking

For each face detected in the classroom image, cosine distance is computed against every entry in the database:

$$D_{\cos}(A, B) = 1 - \frac{A \cdot B}{\|A\| \|B\|} \quad (2)$$

The recognition algorithm is as follows.

Attendance Recognition Algorithm

Require: Image I , database face_db, threshold $\tau = 0.5$

Ensure: Sets present, absent; annotated I

```

1:  $\mathcal{F} \leftarrow \text{app.get}(I)$ 
2: present  $\leftarrow \emptyset$ 
3: for all face  $f \in \mathcal{F}$  do
4:    $e_f \leftarrow f.\text{embedding}$ 
5:    $d^* \leftarrow 1.0$ ,  $n^* \leftarrow \text{None}$ 
6:   for all  $(n, \bar{e}_n) \in \text{face\_db}$  do
7:      $d \leftarrow D_{\cos}(e_f, \bar{e}_n)$ 
8:     if  $d < d^*$  then
9:        $d^* \leftarrow d$ ;  $n^* \leftarrow n$ 
10:    end if
11:  end for
12:  if  $d^* < \tau$  then
13:    Add  $n^*$  to present; draw green box
14:  else
15:    Label "Unknown"; draw red box
16:  end if
17: end for
18: absent  $\leftarrow \text{all} \setminus \text{present}$ 
19: return present, absent, annotated  $I$ 

```

The decision rule is:

$$\text{Identity} = \begin{cases} n^* & \text{if } D_{\cos}(e_f, \bar{e}_{n^*}) < \tau \\ \text{Unknown} & \text{otherwise} \end{cases} \quad (3)$$

The threshold $\tau = 0.5$ was selected empirically to balance false acceptance and false rejection rates.

4.5 Attendance percentage tracking

Table 1 shows the complete data flow within attendance_engine.py from image ingestion to CSV output.

Table 1: Attendance recognition data flow (attendance_engine.py)

Step	Stage	Detail
1	Classroom image input	cv2.imread()
2	Face detection	InsightFace app.get()
3	Embedding extraction	512-D ArcFace vector per face
4	Cosine similarity match	Compare vs face_database.npy
5	Present/absent decision	Threshold $\tau = 0.5$
6	Annotated image output	Green / red bounding boxes
7	CSV + statistics update	attendance.csv, percentage.csv
8	Results dashboard	result.html display

After each session, cumulative statistics are updated. The attendance percentage for student s is:

$$P_s = \frac{C_s^{\text{present}}}{C_s^{\text{total}}} \times 100 \quad (4)$$

where C_s^{present} is the number of sessions in which student s was recognised as present and C_s^{total} is the total sessions conducted

5.1 Technology stack

V. IMPLEMENTATION DETAILS

Table 2 summarises the libraries and tools used.

Table 2: Technology stack

Component	Library / Tool	Version
Face detection & recognition	InsightFace (Buffalo_L)	0.7.x
Image processing	OpenCV (cv2)	4.x
Numerical computing	NumPy	1.x
Distance computation	SciPy	1.x
Data management	Pandas	2.x
Web framework	Flask	3.x
Tunnel / remote access	Ngrok	3.x
Language	Python	3.10+

5.2 Project file structure

Table 3 shows the project directory layout.

5.3 Web application routes

Table 4 lists the HTTP endpoints implemented in app.py.

5.4 Hardware environment

The system was developed and tested on:

- **OS:** Windows 10/11
- **CPU:** Intel Core i5/i7 (multi-core)
- **GPU:** NVIDIA GPU with CUDA (optional; CPU fallback via ctx_id = -1)
- **RAM:** 8 GB or above
- **Camera:** Standard 12 MP smartphone or DSLR

Table 3: Project directory structure

Path / File	Description
project/	Root directory
app.py	Flask web application
attendance_engine.py	Recognition engine
build_database.py	Offline embedding builder
face_database.npy	ArcFace embedding dictionary
dataset/Student_A/, ...	Per-student image folders
static/	Uploaded / annotated images
templates/	HTML Jinja2 templates
csv_files/	Session CSV output
attendance_history/	Timestamped CSV archives
attendance_stats/	Cumulative statistics folder

Table 4: Flask application routes

Route	Method	Description
/	GET	Home (upload page)
/login	GET, POST	Teacher login
/logout	GET	Session clear
/attendance	POST	Process image; return results
/download_csv	GET	Download latest CSV
/history	GET	List historical CSV files
/percentage	GET	View cumulative stats

5.5 Database building procedure

Prior to deployment, each enrolled student's folder inside dataset/ is populated with 5–10 frontal face images captured under varied lighting. build_database.py is executed once to produce face_database.npy. Adding a new student requires inserting their folder and re-running the script—no retraining is needed.

VI. RESULTS AND EVALUATION

6.1 System outputs

For every uploaded classroom image the system produces:

1. An annotated JPEG with green bounding boxes (known students) and red bounding boxes (unknown faces).
2. A current-session CSV with columns *Student Name* and *Status*.
3. A timestamped historical CSV archived in attendance_history/.
4. An updated attendance_percentage.csv reflecting cumulative statistics.

6.2 Evaluation metrics

System performance is measured using standard classification metrics. Let TP, TN, FP, FN denote true positives, true negatives, false positives, and false negatives respectively.

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

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

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

$$TP + FN$$

(5)

(6)

(7)

$$F_1 = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}} \quad (8)$$

6.3 Comparative analysis

Table 5 compares the proposed system with existing attendance approaches across key dimensions.

Table 5: Comparison of attendance systems

System	Contactless	No Extra HW	Anti-Proxy	Analytics
Manual	✓	✓	×	×
RFID	✓	×	×	×
Fingerprint	×	×	✓	×
LBPH-based	✓	✓	✓	×
Proposed	✓	✓	✓	✓

6.4 Performance discussion

The InsightFace Buffalo_L model incorporates the SCRFD face detector, which achieves high recall even for small or partially occluded faces. The ArcFace backbone produces embeddings that are robust to lighting variation, mild pose changes, and different facial expressions. The cosine similarity threshold $\theta = 0.5$ was chosen empirically; values below 0.5 reduce false acceptances at the cost of increased false rejections, while values above 0.5 have the opposite effect. The system processes a typical classroom image (30–40 students) in under five seconds on a CPU-only laptop, making it practical for real-time deployment.

VII. LIMITATIONS AND FUTURE WORK

7.1 Current limitations

- Static threshold:** The fixed $\theta = 0.5$ may not generalise across all datasets; an adaptive threshold calibrated per class would improve robustness.
- Single-image input:** The system processes one uploaded photograph; real-time video-stream attendance is not yet supported.
- No liveness detection:** A printed photograph or screen replay could spoof the system.
- Hardcoded credentials:** The current USERNAME/PASSWORD implementation is unsuitable for production; a database-backed role-based access control system is required.
- Occlusion sensitivity:** Faces obscured by masks or hands may be undetected or misidentified.

7.2 Future work

- Integration of a real-time video stream for continuous attendance monitoring.
- Incorporation of anti-spoofing models such as FAS (Face Anti-Spoofing) to prevent presentation attacks.
- Cloud deployment on AWS/GCP with role-based authentication for multi-institution scalability.
- Adaptive threshold optimisation using a small calibration set collected per institution.
- Development of a companion mobile application for teachers to capture and submit classroom images.
- Extension to masked-face recognition using specialised models trained on occluded faces.

VIII. CONCLUSION

This paper presented an automated classroom attendance system that combines the state-of-the-art InsightFace Buffalo_L facial recognition model with a lightweight Flask web application. The system detects all faces in a single classroom photograph, matches them against a pre-built ArcFace embedding database via cosine similarity, and produces fully automated attendance records with cumulative percentage analytics. The proposed design is contactless, requires no dedicated hardware beyond a standard camera, and operates entirely on-premises. Compared to manual methods, RFID systems, and classical vision-based approaches, the proposed system offers superior accuracy, eliminates proxy attendance, and provides teachers with a comprehensive attendance management dashboard. Future extensions incorporating video streaming, anti-spoofing, and cloud deployment will further enhance the system's practicality and security.

REFERENCES

- [1] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive Angular Margin Loss for Deep Face Recognition," in *Proc. IEEE/CVF CVPR*, Long Beach, CA, USA, 2019, pp. 4690–4699.
- [2] J. Guo, J. Deng, A. Lattas, and S. Zafeiriou, "Sample and Computation Redistribution for Efficient Face Detection," *arXiv preprint arXiv:2105.04714*, 2021.
- [3] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A Unified Embedding for Face Recognition and Clustering," in *Proc. IEEE/CVF CVPR*, Boston, MA, USA, 2015, pp. 815–823.
- [4] Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the Gap to Human-Level Performance in Face Verification," in *Proc. IEEE/CVF CVPR*, Columbus, OH, USA, 2014, pp. 1701–1708.
- [5] M. Grinberg, *Flask Web Development: Developing Web Applications with Python*, 2nd ed. Sebastopol, CA, USA: O'Reilly Media, 2018.
- [6] T. Ahonen, A. Hadid, and M. Pietikäinen, "Face Description with Local Binary Patterns: Application to Face Recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 12, pp. 2037–2041, Dec. 2006.
- [7] A. Muwanguzi and V. Lin, "The Inefficiency of Manual Attendance Taking in Higher Education," *Int. J. Educ. Technol.*, vol. 5, no. 2, pp. 11–18, 2013.
- [8] S. Sneha and T. Dharani, "Automated Attendance Monitoring System Using RFID," in *Proc. Int. Conf. Innovative Mechanisms for Industry Applications (ICIMIA)*, Bangalore, India, 2017, pp. 714–717.
- [9] M. Turk and A. Pentland, "Eigenfaces for Recognition," *J. Cognitive Neurosci.*, vol. 3, no. 1, pp. 71–86, 1991.
- [10] S. Ge, J. Li, Q. Ye, and Z. Luo, "Detecting Masked Faces in the Wild with LLE-CNNs," in *Proc. IEEE/CVF CVPR*, Honolulu, HI, USA, 2017, pp. 2682–2690.
- [11] H. Kaur and P. Singh, "Automated Attendance System Based on Face Recognition Using Deep Learning," in *Proc. Int. Conf. Computing, Communication and Automation (ICCCA)*, Greater Noida, India, 2020, pp. 1–5.