International Journal of Engineering & Extended Technologies Research (IJEETR)

IJEETR

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

Real-Time Student Attendance Management System using Deep Learning-Based Facial Recognition

Dr. G. Kishor Kumar M. Tech, Ph.D.[1], G.Swarnalatha[2], S.Nagendra[3], V.Babu Saheb[4]

Professor & Head, Department of CSE & Business Systems, , RGM College of Engineering and Technology, Nandyal, Andhra Pradesh, India[1]

B.Tech Students, Department of CSE & Business Systems, RGM College of Engineering and Technology, Nandyal, Andhra Pradesh, India[2,3,4]

Publication History: Received: 25.02.2026; Revised: 20.03.2026; Accepted: 25.03. 2026; Published: 28.03.2026.

ABSTRACT: Student attendance management is an undertaking, which consumes a great deal of resources, and whose results can still be marred by proxy fraud and clerical error. This proposal is an AttendAI, full-stack, real-time, automated attendance management system, a blend of facial recognition, which is deep learning-associated and a Flask REST API on a MongoDB Atlas cloud database. Using the InsightFace ArcFace pipeline, multiple students can be recognized and identified right in real-time on live webcam images so that the teacher makes them available through a web browser interface. Poor attendance journalizing is only written on a legit timetable slot which gives a schedule integrity to the server level. Offline enrollment averages per-student embeddings on a dataset of reference images, and enhances resistance to pose and expression change. To get the highest recognition rate a cosine-similarity parameter (τ) = 0.35 produces a highest recognition rate of 96% with a per-frame CPU processing latency of 120178ms and a maximum number of a total of 6 students simultaneously. The entire system encompasses the complete attendance cycle - selecting the classes and recognizing them in real time, going through a teacher review verification step, and permanent storage in the database, along with student, faculty, and timetable-management multiple views.

KEYWORDS: ArcFace, attendance automation, biometric identification, deep learning, facial recognition, Flask, InsightFace, MongoDB Atlas, proxy prevention, real-time systems.

I. INTRODUCTION

Proper documentation of attendance is one of the basic needs of administration in institutions of higher learning. Although it is a routine process, this access to faculty members has a regular time cost: Roll calls cost several minutes of lecture time and the paperwork resulting has to be manually input into institutional systems two-step process that further amplifies potential error rates. In addition to its inefficiency, conventional methods are also prone to proxy attendance where a physically absent student will have someone impersonate him/her. Such fraud can only be detected manually, either by having a human continually scrutinize the classroom or by verifying them biometrically - neither of which would be feasible at scale without software.

Biometrics solutions offer a philosophical remedy. Face recognition, compared to other modalities, has a very advantageous place to be implemented in classrooms: it is completely non-contact, does not require any special equipment on the side of the student, can be passively used by a single fixed camera, and will not cause any physical friction in the classroom. A student who walks into a classroom and sits is identified without any conscious effort on his part and all compliance weight is taken off her.

Several early face-recognition attendance system implementations, based on handcrafted features Local Binary Patterns (LBP), Eigenfaces based on PCA projection, or Fisher discriminant analysis, set the conceptual framework but always failed in practice when subjected to real-world circumstances; changing illumination as the day progresses, off-axis views taken when students stare up at the board or their notes or reference books, partial occlusions These constraints inspired a change towards end-to-end learned representations.

Performance ceiling was greatly reached when large-margin loss functions were introduced to deep face recognition. Builds upon ArcFace [2], which defines recognition as classification on a unit hypersphere with an additive angular margin penalty, yields the embeddings in 512 dimensions in which the cosine distance yields identity similarity in a clean way. Combined with RetinaFace Multi-Scale detector and InsightFace open-source toolkit [3], these improvements enable the deployment of an effective multi-face recognition system with a standard CPU-only setup without special infrastructure.

On this basis, we introduce our concept AttendAI which is a complete production-bound attendance automation. The main findings of this work are: (i) integrated real-time face recognition with timetable-gated access control/end-to-end system architecture; (ii) mean-embedding enrolment plan of converting many reference images into just one robust gallery vector/student; (iii) Laplacian-variance blur pre-filter discarding and dead frames before neural inference process; and (iv) multi-view teacher dashboard

The rest of the paper is formatted as follows. Section II discusses related work. The background and motivation are detailed in Section III. Part IV includes the system architecture and design. Section V describes the face recognition engine. In section VI, there is a database design. Section VII describes the implementation and web interfaces. Section VIII presents the results and discussion of the experiments. Section IX has a conclusion and provides future work directions.

II. LITERATURE SURVEY

Automated attendance systems are closely related in the history of computer vision overall and machine learning in particular. The previous one employed Viola-Jones Haar cascade detectors [1] to find the faces, and allow verification of the identity using Eigenface or Fisherface classifier. In spite of being able to compute cheaply and run on much less hardware, they were very vulnerable to performance decreases with respect to illumination change and non-frontal poses and so were not useful in spontaneous classroom applications.

Illumination sensitivity was partly attended to by using the Gabor wavelets and Local Binary Pattern histograms to work on local texture statistics as opposed to global pixel intensities [5]. Description Excellent attendance systems are known to have a satisfactory accuracy under semi-laboratory conditions and could not be extrapolated to unrestricted environments in which the target could move around freely and the light conditions slowly changed throughout the day.

The Deep Learning world gave birth to paradigm shift. Results of Taigman et al. [7], using a nine-layer convolutional network, trained with a 3D refinement step demonstrated an approximation of human verification accuracy, that the results of learned representations could outperform human-crafted features (areas) significantly. Subsequent studies of loss geometry design like contrastive loss, triplet loss, SphereFace [8], CosFace, and others, began refinements on the embedding geometry, with ArcFace [2], which deploys a fixed angular margin in the angular space of the embedding hypersphere. ArcFace embeddings exhibit intra-class clustering and inter-class separation and this shows that the thresholding based on cosine-distance is an efficient rule to make identity decisions.

At the detection end, RetinaFace [3] suggested an addition to multi-task learning architecture to simultaneously predict face bounding boxes, facial landmarks and 3D face mesh in a single forward pass that offers scale- and rotation-invariant robust detection. The InsightFace toolkit is a collection of both arcface recognition and RetinaFace detection in one and easy to operate Python library enabling practical applications of both to be made without the burden of intermediary expertise in model execution.

These developments have seen different levels take advantage of them based on the attendance specific systems. To build an initial prototype Neto et al. [4] created the prototype basing on LBP features that have been obtained via a web interface showing real-world relevance of enrollment management and session-level deduplication. Recent studies [6] have indicated the CNN-based recognition of smart classrooms using IoT sensors, but on individual students, rather than multi-face recognition simultaneously. The cloud-based deployments have been explored as far as the question of scalability goes but otherwise in most of the instances have not adopted the schedule-enforcement feature that prevents instances of fraud out-of-class attendance.

The under development fulfils these gaps by adding a state of the art deep recognition with fully implemented application stack: timetable-gated REST endpoint endpoints, an online cloud Atlas MongoDB storage, a mean-embedding enrollment pipeline, a blur-based quality filter and a full teacher facing web dashboard - it can be deployed

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

on commodity hardware and does not require any GPU infrastructure.

III. BACKGROUND AND MOTIVATION

A. The School Problem of Absenteism.

It is the legal and administrative responsibility of academic institutions to keep proper attendance registers of each course section. These reports control scholarship eligibility, decision and report on progression and regulatory compliance. In practice attendance information is usually gathered by roll calls at the beginning of each lecture, creating a per-session list which needs to be transcribed, stored and summed up over the semester.

This workflow as a system is weak in three aspects. First, it consumes lecture time in the same fraction as does 5 or 7 minutes per session 60 student section can consume a slice of 5-7 minutes of lecture time in the same proportion that it occupies lecture time (a 60 student section can take 5-7 minutes in a section). Second, manual data entry is a procedure that instifles data entry defects which are difficult to reverse audit in the back part. Third, and most crucially, it does not provide any way in which we might be confident that the student answering is looking at us in the face: a banal verbal assurance can always be made trivially.

B. Why Face Recognition

Face recognition may be regarded as one of the most appropriate biometric modalities due to its potential to be passive, nondestructively gathered and with the least effort to implement in practice as the classroom scenario requires the combination of a variety of modalities: fingerprint, iris, voice, gait and face. Fingerprints scanners have to be touched individually and iris scanned, voice recognition is sensitive to other voices. Conversely, face recognition merely requires any commercial camera that has an unobstructed view of the classroom, students do not need to take any action in person, and can be scaled to simultaneous identification of an entire classroom.

Deep learning based face recognition algorithms are already state of the art with greater than 99.8 percent verification accuracy on idealized benchmarks [2], meaning that in the real world (and under moderately challenging conditions) there are almost no errors. The remaining problems, which comprise the following extremes of illumination, extremely low resolution, extreme occlusion and adversarial spoofing, can be resolved by design choices (the camera placement, selecting a filter to remove shadows, liveness detection) and not by intrinsic limitations of the algorithm.

C. Motivation for AttendAI

Existing open attendance systems (open source) are either not reliable in going to production (single-face processing, not persistent on clouds) or not a good fit to operate with a large number of rack-mountable servers. AttendAI is motivated by the wish to achieve: (i) accurate enough to be deployed to an unattended classroom; (ii) fast enough to run live video stream on a CPU; (iii) secure enough to thwart out-of-ignored fraud; and (iv) complete enough to take over the entire manual workflow without creating parallel systems.

IV. SYSTEM ARCHITECTURE

A. Three-Tier Design

AttendAI is a typical, three-tier web application structure. Presentation tier is a web-based interface, which receives the webframe of the web cam through HTML5 MediaDevices API and displays the attendance feedback real-time. The application layer is a Flask REST server which accepts the frames via the HTTPS stream, coordinates face recognition, enforces access control based on the time schedule, and manages the state of sessions. The MongoDB Atlas data tier is a fully managed cloud platform database that offers horizontal scalability, automatic failover and geographic replication.

This isolation of responsibilities enables scaled operation of each level: the browser client can be utilized by the current device of the teacher, a Flask server can be deployed in any VM with a cloud provider or any on-premises computer, and data persistence can be done by MongoDB Atlas without the application server needing to operate data infrastructure. As can be seen in Figure 1, the teacher dashboard will be the starting point of all the attendance sessions.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

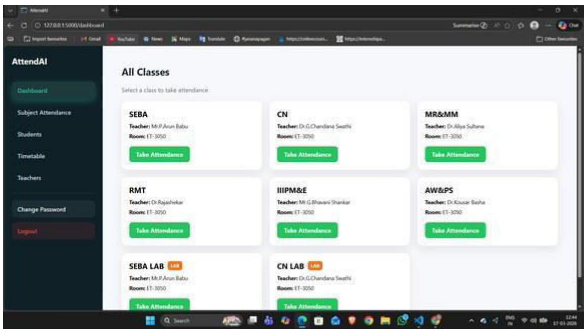

Fig.1. AttendAI Teacher Dashboard — Class Selection

B. Request Processing Lifecycle

Upon opening a class and initiating a live session, the browser will go into an endless capture loop: each 2 seconds, a JPEG image will be captured off the live camera stream, it is client-side compressed to save on transmission bandwidth and sent as raw bytes over to the /process_frame endpoint. The server unpacks the binary payload with OpenCV, feeds the resultant NumPy array into the InsightFace recognition engine, and is returned a list of detection objects - which are each a bounding box, matched identity string, and cosine similarity score.

The server has a per-session Python set in application memory (CURRENT_SESSION_PRESENT). Every new student found horde the name of which is not in this set is inserted and recorded in the JSON response, and later frames which again identify the same student will not cause further writes to the database. By clicking Stop, the teacher sets the passing of the session set to a confirmation screen in which individual entries may be unselected manually prior to receipt. Finalization cycles through the verified list and features a single attendance document per student in MongoDB and clears the session set.

C. Time Based Access Control.

An important security aspect of AttendAI is a mandatory attendance marking only within designated time frames. All calls to /process_frame call get active timetable (subject, department) first, asking the timetable collection to find a document, where day field = current week-day and start and end time interval contain the current system clock time. In case of the non-presence of such document, the endpoint sends a frame-destroying reply at once, with the status 403 Forbidden and does not even bother with recognition.

This design makes it impossible to record an attendance beyond the scheduled class session even when an accidental teacher leaves a window of a session open or an attacker revisits recorded frames. At query time, the MongoDB is loaded with the timetable data, thus a schedule change is immediately put into effect without a restart of a server. The weekly timetable (Figure 2) includes the type of slots painted in color.

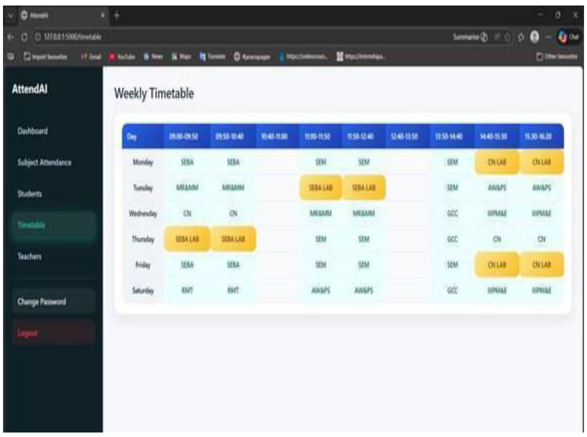

Fig.2. Weekly Timetable — Regular (Teal) and Lab (Amber) Slots

V. FACE REECOGNITION ENGINE

A. InsightFace Pipeline

The recognition subsystem is itself located in the file face engine.py and is launched when the server starts starting with InsightFace FaceAnalysis class and CPUExecutionProvider backend. This selection of execution providers suggests that the system can take advantage of mere normal processor hardware, not using a discrete GPU, and deployment to any modern laptop or cloud VM is likely to be free. The pipeline is made of three consecutive steps where each new frame passes through:

Stage 1 - Detection: retinaface detects axis-aligned bounding box and five face landmark features (eye center, nose tip, and mouth corners) of every face that it can detect in the input frame by scanning it with multi-scale convolutional filters. Detecting a large range of different scales of faces, the detector has a fairly good accuracy with faces that are as small as 20 priority of face scales and 20 pixels per pixel in the input resolution.

Stage 2 — Alignment: All face crops as seen are warped by another similarity transform (rotation, uniform scaling, translation) that forms map equivalents of the five predicted landmarks to a canonical template on a 112x112 pixel grid. This normalization is essential: it eliminates the variability in in-plane rotation and makes sure that the next feature extractor will receive a constant-oriented input irrespective of the pose of the head of the student.

Stage 3 -Embedding: The aligned crop is embedded to an ArcFace ResNet-50 backbone which produces a 512-dimensional floating-point embedding. Before matching, this vector is L2-normalized to unit length thus limiting any embedding to a unit hypersphere and making the matching operation a dot product that is equivalent to cosine similarity.

B. Student Enrollment Procedure

Enrollment is done as an offline batch operation before deploying the system. One of the reference collections of face pictures is organized as a per-student directory: each of the subdirectories is labeled by the identifier of the student and contains 5 to 20 pictures under various conditions: in various lighting conditions, slightly different positioning and with or without glasses.

The enrollment script identifies all faces in an image and chooses largest detected face (in terms of bounding box area) to avoid spurious detection of background persons, the embedding of a face is obtained and L2- normalized. The mean of the embeddings obtained must be considered once all the images of a specific student have been averaged which will be one 512-dimensional mean vegetable. This is an average value of the vectors which is re-normalized again to a unit length and is saved in the MongoDB Students collection as a record of a student.

The mean-embedding strategy is one of the design options. An individual snapshot of the student captures the student in a single situation; across a series of snapshots, averaging gives the student a centroid in embedding space which is more abundant of within-class variance. Such an approach empirically, has a large reducing power in pupils whose looks vary by session.

C. Cosine Similarity Identity Matching

The query embedding q is compared with the entire gallery at inference time by comparing the dot product scores(q, r_i) = q · r_i to every gallery vector r i (unit-normalized). The gallery entry that has the top score is provisionally identified as corresponding identity. When this maximum score is greater than the threshold τ, then the detection is registered as a known student, otherwise the face is considered Unknown and is not marked as attending.

Figure 3 indicates recognition accuracy with respect to τ tested on the test dataset. Accuracy peaks at 96% for $\tau = 0.35$. At lower thresholds, low-quality embeddings of half-occluded or blurred faces are mistakenly recognized (false acceptances); at higher thresholds, true matches of students at non-frontal angles are below the threshold (false rejections). It uses the value, $\tau = 0.35$ as the operating point, which achieves the highest balanced accuracy in both types of errors.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

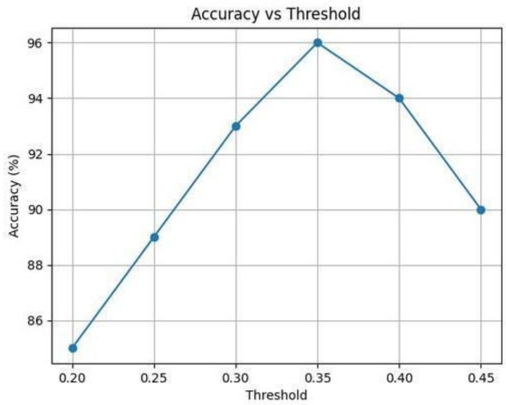

Fig.3. Recognition Accuracy vs. Cosine Similarity Threshold τ

D. Laplacian Variance Blur Pre-Filter

Live video streams always include frames that have been degraded due to motion blur especially when a student is about to sit down, turn his/her head or when a student is moving seats. Embedding a blurred face crop to Arcface extractor results in an inaccurate embedding, which would lead to false identifications. To do this at a cheaper cost than neural inference on each crop, we use a lightweight pre-filter: the Laplacian operation is convolved with the grayscale image of each face crop, and the gradient magnitude image variance is calculated. A low variance means spatial smoothness (in accord with blur). Silently discarding crops with a Laplacian variance less than 60 prior to embedding extraction is a gating step effectively allowing only the costly neural inference on inputs that pass the quality test to be performed.

VI. DATABASE DESIGN

A. MongoDB Collection Schema

AttendAI deploys MongoDB Atlas with 5 collections that constitute the overall data model. The columns in the Teachers collection include: teacher_id, hashed password (PBKDF2-SHA256, using Werkzeug) name, username, and the department. The Students table is where student-id, name as well as the 512 length floating-point embedding array is stored. The Classes collection links every subject code with the name of a teacher, class name, and a room name. The Timetable collection captures the day of the week, commencement time (HH:MM), the conclusiveness of time (HH:MM), subject, department and room spent by each slot. The Markattendance collection will record the entire record of each attendance instance: student id, student name, subject, date (YYYY-MM-DD), day of week, clock time, start and finish time of a schedule slot, and a status field with a value of Present.

The uniqueness constraint is at the database level enforced by a compound index on the primary key (student_id, subject, date) of gallery collection Markattendance: although the application layer might write the same student on the same date with the same subject, MongoDB has rejected all but the first such write, which offers a defense-in-depth guarantee against duplication of the attendance record.

P. Waveform aggregation pipeline.

A multi-stage MongoDB array of aggregation creates the student report. The initial stage takes all the Markattendance records and aggregates them around the composite key student-id / subject, and counts the number of attended sessions per subject per student. The second step re-groups by student id, and shoves the per subject counts into an array and adds them to get a final figure of sessions attended. The third phase predicts the end result, and to calculate the overall attendance percentage, (total attend / TOTAL classes) / 100) will be calculated and rounded off to two decimal points. The outcome documents are translated into the Students view in Figure 4 with percent badges less than 75% as the red ones to indicate students who will be targeted by academic punishments via attendance.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

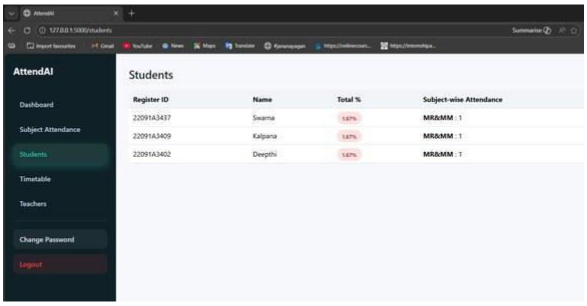

Fig 4. Student Report — Register ID, Attendance % and Subject Breakdown

The report by the teachers relies on a similar pipeline which counts the number of unique tuples of (subject, date) pairs of teacher of their subjects, as a measure of the classes taught. This view is displayed in figure 5.

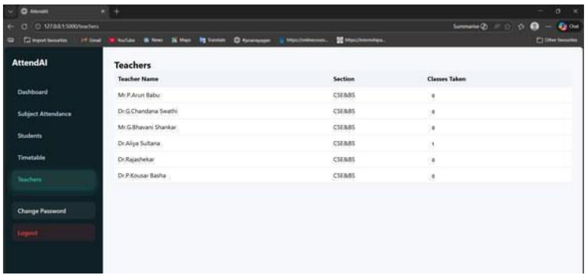

Fig 5. Teacher Report — Classes Conducted per Faculty

VII. IMPLEMENTATION

A. Technology Stack

The Python 3.10 server-side application is built on the Flask micro-framework, serving Python 3.10, as an HTTP routing and session management framework. This is because PyMongo allows connectivity to a database and it automatically resolves to the topology of the cluster by using SRV DNS resolution using the official MongoDB Atlas connection string. OpenCV 4.8 is utilized in activities of computer vision: frame decoding, color conversion and blur scoring. InsightFace library version 0.7 is a face detector and face recognition prototype which provides pre-trained retinaface and arcface models. The web frontend is coded in HTML5 and CSS3 with bespoke properties and vanilla JavaScript to operate the camera capture loop, and to submit AJAX frames.

B. Authentication and Security.

Teacher accounts are secured with Werkzeug, by creating a pass word hash, and checking password hash (by Invoking generate password hash and password hash checking functions) based on PBKDF2-SHA256 and with a randomly generated per-password salt. Never are passwords stored in plaintext. The decorated routes all have a login required decorator with failed routes redirecting to the login page. On the server side, cookies are signed with a secret key thereby deterring the client-side fraud. Additional access control provided by the time schedule enforcement mechanism listed in Section IV-C can serve to provide an additional access control layer to the attendance process.

C. System Flow

The entire credit of the way how AttendAI works is as follows. When it is authenticated by a teacher, the dashboard displaying all the classes set up is shown and the target class is selected which directs to the live attendance view. The browser begins to save and queue up the frames; the server recognizes students pushing names back to the UI, as they fall in. At the end of the session, the teacher also can click on Stop, audits the confirmation list, and can remove the falsely identified names, and finalize. The server will leave a record of the attendance of the registered students and send to the dashboard. This entire process, which involves the onset of the session up to its conclusion, will typically take 3-5 minutes in an example of a classroom with thirty students.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

VIII. EXPERIMENTAL RESULTS AND DISCUSSION

A. Experimental Setup

It was implemented and tested in an actual classroom setting at the Department of Electronics and Telecommunication Engineering. The test data included students taking a variety of courses. Enrollment reference images were taken in controlled (frontal, neutral expression, good light) and varied (side views, glasses, days training) conditions. The tests were carried out in natural fluorescent light conditions (indoor) with a standard 1080p USB web camera that was placed at a height of around 2 meters, and a distance of about 3-4 meters before the 1st seat of the row.

B. Method Comparison

The test dataset has also been run through two baseline methods: LBP histogram matching and nearest-neighbor classification to place AttendAI into perspective, and an Eigenface projection score based on the PCA and cosine distance. The recognition accuracy is shown in Figure 6. LBP has 75% accuracy, which indicates that it is sensitive to the difference in illumination. PCA-based Eigenfaces are considered to be enhanced to 82% because they can carry out the discriminative activity in a lower-dimensional subspace. AttendAI ArcFace index has an accuracy of 95 percent, 13 percentage points better than PCA and 20 points better than LBP, which justifies the use of deep learned representations in this deployment scenario.

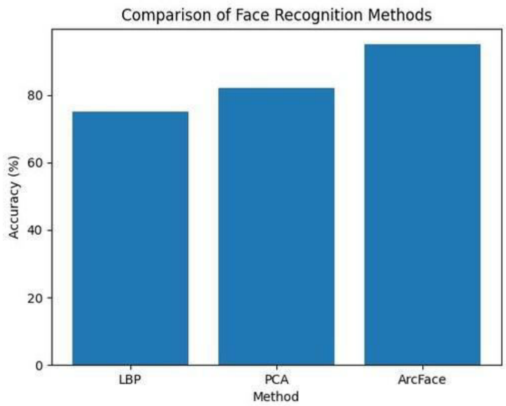

Fig.6. Recognition Accuracy: LBP (75%) vs. PCA (82%) vs. ArcFace (95%)

C. Threshold Sensitivity

Figure 3 (Section V-C) indicated that the accuracy is maximized at τ =0.35, and increases outwards in both directions- smaller thresholds allow marginal matches and larger thresholds differentiate valid ones. This can be observed to be the case with the geometric interpretation of the ArcFace embedding space: real pairs will cluster around cosine similarity of 0.40 or higher whereas impostor pairs will not without clustering to the lower half and the overlap area will center around 0.35, which is the case in this dataset. The system is calibrated every time the institution is deployed in new settings (more students, other geometry of the cameras) because τ is expected to change.

D. Processing Latency

Figure 7 shows the curve of per-frame processing time vs. 20 consecutive test frames on a typical consumer laptop (Intel Core i5, 8 GB RAM, no discrete graphics card). Response times span between 120 ms and 178 ms with an average response time of about 152 ms along with a headroom of over 92 percent per sample period. The system can easily run within real-time constraints even on CPU-only hardware and has over 92 percent latency budget headroom at a 2 seconds server-side sample rate.

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

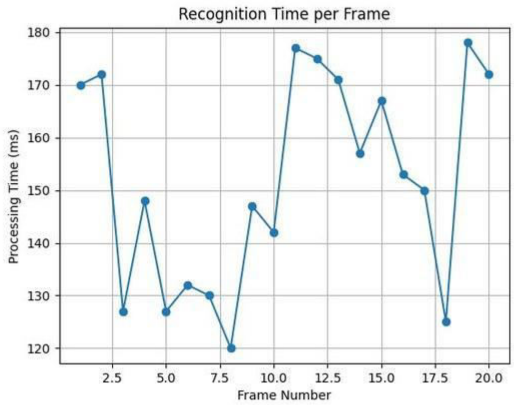

Fig.7. Per-Frame Recognition Processing Time (ms) — CPU Only

The difference in processing duration (range: 58 ms), has been largely explained by the quantity of faces observed at one frame per process: frames with 5-6 students attract a proportional number of alignment and embedding operations as compared to the frames with 2-3 students. Secondary factors are operating system scheduling jitter, and memory bandwidth contention due to running system processes.

E. Multi-Face Live Recognition.

These are representative frames of the live system, running with a full pipeline, as seen in Figures 8 and 9. Green bounding boxes indicate detected face areas; overlay labels have matching identity and cosine similarity rating. Figure 8 identifies six students at once, the range of scores being between 0.38 (Hemchand who is partly out of camera view) and 0.80 (Guru Kishore who is almost facing the camera at a medium distance). Figure 9 shows that there are 5 students in a close-range frame scoring between 0.41 and 0.71.

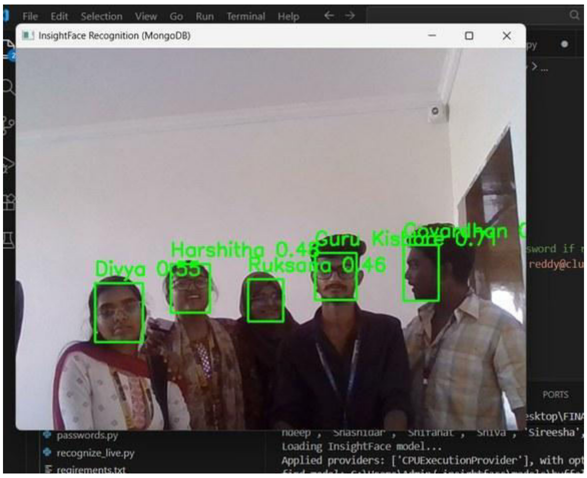

Fig.8. Live Recognition — 5 Students with Bounding Boxes and Cosine Scores

International Journal of Engineering & Extended Technologies Research (IJEETR)

|ISSN: 2322-0163| www.ijeetr.com | A Bimonthly, Peer Reviewed, Scholarly Indexed Journal |

| Volume 8, Issue 2, March - April 2026 |

DOI:10.15662/IJEETR.2026.0802408

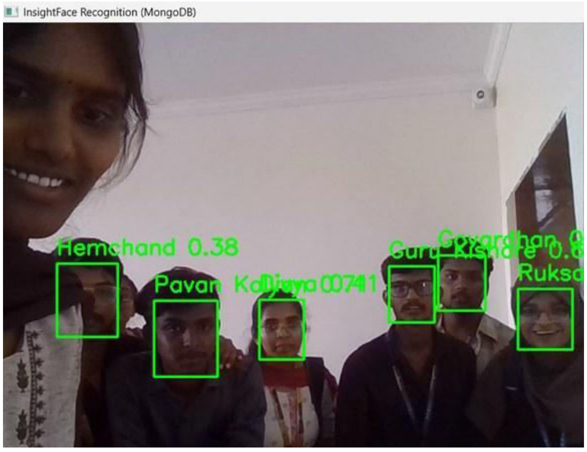

Fig.9. = Live Recognition — 6 Students, Varied Poses and Distances

The findings prove that the system works with real group set-ups that contain different poses, corrective lenses and diverse distance settings of the camera. The individual scoring 0.38 same but slightly above the threshold can explain the need to tune τ: an even higher threshold would have discarded this occurrence, whereas a slightly lower one would open up the system to marginal false acceptances by other people.

TABLE I. AttendAI Performance Summary

Metric	Value	Notes
Peak Accuracy	96%	τ = 0.35
Mean Frame Time	~152 ms	CPU only
Min / Max Frame	120 / 178 ms	20-frame test
Max Simultaneous	6 faces	Group frame
Blur Threshold	60 (Lap. var.)	Auto-discard
Duplicate Guard	Session set	Per-day/subj.
Timetable Lock	HTTP 403	Out-of-slot
ArcFace vs. LBP	+20 pp	Accuracy gain
ArcFace vs. PCA	+13 pp	Accuracy gain

IX. CONCLUSION

AttendAI, a full production-state automated facial recognition student attendance management system (combining facial recognition with deep learning and Flask REST backends and MongoDB Atlas cloud storage) was presented in this paper. The system addresses the three major vulnerabilities of the manual attendance system, time based cost, limit on manual-transcription and vulnerability to proxy frauds by utilization of web-based teacher workflow that does not need any additional equipment but a web camera which is already available.

On a real-world dataset of classroom data at a cosine similarity cosine threshold of τ = 0.35, the InsightFace ArcFace recognition pipeline achieves a 96% recognition accuracy, in comparison to LBP (75) and PCA Eigenface (82) baseline scores 20 and 13 percentage points respectively. Per-frame CPU hardware processing varies between 120 ms and 178 ms, well within the 2-second sampling period, and proves there is no need to deploy with a GPU. Practically the real classroom group size can be scaled by recognizing up to six students at once, in a frame.

The access control is an integrity based system using time-based access in the server level, which enforces the integrity of a schedule, prohibiting out-of-slot attendance cheating, even in the case of current timetable engineering. Within-session pose and expression variance is empowered by the use of mean-embedding enrollment plan. A set of these design choices will give a precise and secure system, which can be implemented with few or no special infrastructure.

X. FUTURE WORK

Devices There are several research and engineering directions, where AttendAI can be developed. Firstly, it would take longer than 150-ms per-frame inferences to just go to under 15 ms with the addition of the ability to do inference on a GPU (even a low-end one that's consumer-friendly) and would permit streaming sampling rather than periodic sampling, and larger, higher-end inferences at a no-extra latency cost. Second, the passive liveness detector component would be added, which would prevent the attacks of photograph-based spoofing, which are the main malicious risk to face-recognition attendance systems. Third, with dynamically determined percentages of attendance we would obtain by replacing the constant fixed TOTAL_CLASSES with the size of classes calculated on a given subject basis by the Timetable collection. Fourth, a mobile application that presents real-time records of attendance to students will enhance transparency and assist students to rectify problematic trends of attendance in time. Fifth, testing on full cohort groups of 60 or more students who represent out of gallery should be performed, so as to achieve statistically rigorous open-set recognition results, to have formal deployment certification requirements.

REFERENCES

1. P. Viola and M. J. Jones, "Robust real-time face detection," Int. J. Comput. Vis., vol. 57, no. 2, pp. 137–154, May 2004.
2. J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR), Long Beach, CA, USA, Jun. 2019, pp. 4690–4699.
3. J. Guo, J. Deng, A. Lattas, and S. Zafeiriou, "Sample and computation redistribution for efficient face detection," arXiv:2105.04714, May 2021.
4. P. S. Neto, G. Holanda, R. Trinta, and J. Souza, "FRAS: A face recognition attendance system," in Proc. ACM Symp. Appl. Comput. (SAC), Sierre, Switzerland, Mar. 2010, pp. 1528–1535.
5. 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.
6. C.Nagarajan and M.Madheswaran - 'Stability Analysis of Series Parallel Resonant Converter with Fuzzy Logic Controller Using State Space Techniques'- Taylor &Francis, Electric Power Components and Systems, Vol.39 (8), pp.780-793, May 2011. DOI: 10.1080/15325008.2010.541746
7. C.Nagarajan and M.Madheswaran - 'Experimental verification and stability state space analysis of CLL-T Series Parallel Resonant Converter' - Journal of Electrical Engineering, Vol.63 (6), pp.365-372, Dec.2012. DOI: 10.2478/v10187-012-0054-2
8. C.Nagarajan and M.Madheswaran - 'Performance Analysis of LCL-T Resonant Converter with Fuzzy/PID Using State Space Analysis'- Springer, Electrical Engineering, Vol.93 (3), pp.167-178, September 2011. DOI 10.1007/s00202-011-0203-9
9. S.Tamilselvi, R.Prakash, C.Nagarajan,"Solar System Integrated Smart Grid Utilizing Hybrid Coot-Genetic Algorithm Optimized ANN Controller" Iranian Journal Of Science And Technology-Transactions Of Electrical Engineering, DOI10.1007/s40998-025-00917-z,2025
10. S.Tamilselvi, R.Prakash, C.Nagarajan," Adaptive sliding mode control of multilevel grid-connected inverters using reinforcement learning for enhanced LVRT performance" Electric Power Systems Research 253 (2026) 112428, doi.org/10.1016/j.epsr.2025.112428
11. S.Thirunavukkarasu, C. Nagarajan, 2024, "Performance Investigation on OCF and SCF study in BLDC machine using FTANN Controller," Journal of Electrical Engineering And Technology, Volume 20, pages 2675–2688, (2025), doi.org/10.1007/s42835-024-02126-w

12. C. Nagarajan, M.Madheswaran and D.Ramasubramanian- ‘Development of DSP based Robust Control Method for General Resonant Converter Topologies using Transfer Function Model’- Acta Electrotechnica et Informatica Journal , Vol.13 (2), pp.18-31,April-June.2013, DOI: 10.2478/aeei-2013-0025.
13. C.Nagarajan and M.Madheswaran - ‘DSP Based Fuzzy Controller for Series Parallel Resonant converter’- Springer, Frontiers of Electrical and Electronic Engineering, Vol. 7(4), pp. 438-446, Dec.12. DOI 10.1007/s11460-012-0212-0.
14. C.Nagarajan and M.Madheswaran - ‘Experimental Study and steady state stability analysis of CLL-T Series Parallel Resonant Converter with Fuzzy controller using State Space Analysis’- Iranian Journal of Electrical & Electronic Engineering, Vol.8 (3), pp.259-267, September 2012.
15. C.Nagarajan and M.Madheswaran, “Analysis and Simulation of LCL Series Resonant Full Bridge Converter Using PWM Technique with Load Independent Operation” has been presented in ICTES’08, a IEEE / IET International Conference organized by M.G.R.University, Chennai.Vol.no.1, pp.190-195, Dec.2007
16. Suganthi Mullainathan, Ramesh Natarajan, “An SPSS and CNN modelling based quality assessment using ceramic materials and membrane filtration techniques”, Revista Materia (Rio J.) Vol. 30, 2025, DOI: https://doi.org/10.1590/1517-7076-RMAT-2024-0721
17. M Suganthi, N Ramesh, “Treatment of water using natural zeolite as membrane filter”, Journal of Environmental Protection and Ecology, Volume 23, Issue 2, pp: 520-530,2022
18. Z. Zhang, X. Song, and H. Cui, "Automated classroom attendance system with deep learning and edge computing," IEEE Access, vol. 9, pp. 34512–34521, 2021.
19. Y. Taigman, M. Yang, M. Ranzato, and L. Wolf, "DeepFace: Closing the gap to human-level performance in face verification," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Columbus, OH, USA, Jun. 2014, pp. 1701–1708.
20. W. Liu, Y. Wen, Z. Yu, M. Li, B. Raj, and L. Song, "SphereFace: Deep hypersphere embedding for face recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Honolulu, HI, USA, Jul. 2017, pp. 212–220.
21. C. Wang et al., "The devil of face recognition is in the noise," in Proc. Eur. Conf. Comput. Vis. (ECCV), Munich, Germany, Sep. 2018, pp. 765–780.
22. K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Las Vegas, NV, USA, Jun. 2016, pp. 770–778.
23. Anand, L., Maurya, M., Seetha, J., Nagaraju, D., Ravuri, A., & Vidhya, R. G. (2023, July). An intelligent approach to segment the liver cancer using Machine Learning Method. In 2023 4th international conference on electronics and sustainable communication systems (ICESC) (pp. 1488-1493). IEEE.
24. Rajendran, S., Sundarapandi, A. M. S., Krishnamurthy, A., & Thanarajan, T. (2022). An intelligent face recognition technology for iot-based smart city application using condition-cnn with foraging learning pso model. International Journal of Pattern Recognition and Artificial Intelligence, 36(14), 2256018.
25. Murugeshwari, B., & Sujatha, R. (2014). Preservation of Privacy for Multiparty Computation System with Homomorphic Encryption. International Journal of Emerging Technology and Advanced Engineering, 4(3), 530-535.
26. Sugumar, R. (2025). Unified AI Framework for Predictive Data Engineering and Real Time Prescription and Billing Systems. International Journal of Advanced Engineering Science and Information Technology (IJAESIT), 8(5), 17261.
27. Samrat, B., Thomas, P. K., Kumar, S., Benila, A., Bhardwaj, R., & Vigenesh, M. (2024, December). Industrial informatics in optimizing software-defined vehicles for logistics. In 2024 IEEE 2nd International Conference on Innovations in High Speed Communication and Signal Processing (IHCSP) (pp. 1-9). IEEE.
28. Soundappan, S. J. (2024). AI-driven customer intelligence in enterprise lakehouse systems Sentiment Mining Governance-Aware Analytics and Real-Time Data Synchronization. International Journal of Advanced Engineering Science and Information Technology.
29. Rajasekar, M. (2024). AI-Powered Cyber-Secure Federated Learning on AWS for Next-Generation Digital Banking Analytics. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(3).
30. Deivendran, P., Babu, P. S., Malathi, G., Anbazhagan, K., & Kumar, R. S. (2023). Emotion Recognition for Challenged People Facial Appearance in Social using Neural Network. arXiv preprint arXiv:2305.06842.

31. Sugumar, R., & Murugeshwari, B. (2016). An Efficient MChord based Authentication for Vehicular Ad-Hoc Networks.
32. Pandey, V. K., Mishra, S., Rengarajan, A., Savita, & Roomi, M. M. (2024, March). Enhancing Weather Forecasting with Machine Learning Techniques. In International Conference on Renewable Power (pp. 147-156). Singapore: Springer Nature Singapore.
33. Mathew, A., & Alex, H. (2025). Federated Learning for Secure Genomic Research: Privacy-Preserving AI Solutions for Precision Medicine. Science and Technology: Developments and Applications Vol. 9, 36-43.
34. Selvi, G. V., Anbarasan, A. B., Murthy, B. A., & Prabavathy, S. (2023). An Application Oriented Integrated Unequal Clustering Algorithm for Wireless Sensor Network. In Underwater Vehicle Control and Communication Systems Based on Machine Learning Techniques (pp. 140-154). CRC Press.
35. Soundappan, S. J. (2025). Next Generation AI Enabled Holistic Cognitive Platform for Secure Cloud Network Intelligence Enterprise Systems and Digital Trust Optimization. International Journal of Computer Technology and Electronics Communication, 8(5), 11534-11542.
36. Rajasekar, M. (2024). Real-Time Predictive DevOps Intelligence for Risk-Aware Digital Business Processes in Cloud and SAP Ecosystems. International Journal of Advanced Research in Computer Science & Technology (IJARCST), 7(4), 10713-10718.
37. Jagadeesh, S., & Sugumar, R. (2017). A comparative study on artificial bee colony with modified ABC algorithm. European Journal of Applied Sciences, 9(5), 243–248.
38. Murugeshwari, B., Sarukesi, K., & Jayakumar, C. (2010, March). An efficient method for knowledge hiding through database extension. In 2010 International Conference on Recent Trends in Information, Telecommunication and Computing (pp. 342-344). IEEE.
39. Reddy, K. V. V. K., & Vimal, V. R. (2024, July). A novel approach on improved segmentation and classification of remote sensing images using AlexNet compared over linear discriminant analysis with improved accuracy. In 2024 Second International Conference on Advances in Information Technology (ICAIT) (Vol. 1, pp. 1-6). IEEE.
40. Gowthami, D., & Vigenesh, M. (2024). Distributed and Lightweight Intrusion Detection for IoT: A Lightweight Pyramidal U-Net With Tri-Level Dual Inception-Based Framework. In The Convergence of Self-Sustaining Systems With AI and IoT (pp. 154-173). IGI Global Scientific Publishing.
41. Anand, P. V., & Anand, L. (2023, December). An Enhanced Breast Cancer Diagnosis using RESNET50. In 2023 International Conference on Innovative Computing, Intelligent Communication and Smart Electrical Systems (ICSES) (pp. 1-5). IEEE.
42. Mathew, A. (2022). Leveraging Big Data Analytics to Power AI and ML (Machine Learning) Automation. Educational Research (IJMCER), 4(5), 131-134.
43. Dhinakaran, D. (2022). Joe Prathap P. M, Selvaraj D, Arul Kumar D and Murugeshwari B," Mining Privacy-Preserving Association Rules based on Parallel Processing in Cloud Computing,". International Journal of Engineering Trends and Technology, 70(3), 284-294.
44. Poornima, G., & Anand, L. (2024, April). Effective Machine Learning Methods for the Detection of Pulmonary Carcinoma. In 2024 Ninth International Conference on Science Technology Engineering and Mathematics (ICONSTEM) (pp. 1-7). IEEE.
45. Rengarajan, A., Jayakumar, C., & Sugumar, R. (2012). Optimization Of Recent Attacks Using Internet Protocol. National Journal of System and Information Technology, 5(1), 8.
46. Mathew, A., & Romasco, L. (2024). Forensic Investigation of Artificial Intelligence Systems. Research Updates in Mathematics and Computer Science Vol. 4, 154-164.
47. Vekariya, V., Kumar, S., & Rengarajan, A. (2024). A distinctive and smart agricultural knowledge-based framework using ontology. In Sustainability in Digital Transformation Era: Driving Innovative & Growth (pp. 207-213). CRC Press.
48. Soundappan, S. J. (2020). Big data analytics in healthcare: Applications for pandemic forecasting. International Journal of Advanced Research in Computer Science & Technology, 3.
49. Sugumar, R. (2024). AI-Augmented Quality Engineering for Performance Optimization and Test Orchestration in Distributed Systems. International Journal of Science, Research and Technology, 7(5), 12835-12846.
50. Soundappan, S. J., & Sugumar, R. (2016). Optimal knowledge extraction technique based on hybridisation of improved artificial bee colony algorithm and cuckoo search algorithm. International Journal of Business Intelligence and Data Mining, 11(4), 338–356.

51. Mathew, A. (2025). Ahead of the breach: Predictive threat intelligence in aviation inspired by Scattered Spider attacks. Multidisciplinary International Journal of Research and Development (MIJRD), 4(6), 54–58.
52. Soundappan, S. J. (2021). DataOps: Orchestrating Reliable ML Data Pipelines. International Journal of Research and Applied Innovations, 4(4), 5533-5537.
53. Garg, V. K., Soundappan, S. J., & Kaur, E. M. (2020). Enhancement in intrusion detection system for WLAN using genetic algorithms. South Asian Research Journal of Engineering and Technology, 2(6), 62–64.
54. Anand, L., Tyagi, R., & Mehta, V. (2024, January). Food recognition using deep learning for recipe and restaurant recommendation. In Proceedings of Eighth International Conference on Information System Design and Intelligent Applications (pp. 269-279). Singapore: Springer Nature Singapore.
55. Kumar, A., & Anand, L. (2025). A Novel EEG-Based Deep Learning Framework for Enhancing Communication in Locked-In Syndrome Using P300 Speller and Attention Mechanisms. KSII Transactions on Internet and Information Systems (TIIS), 19(11), 3841-3855.
56. Soundappan, S. J. (2022). AI-Based Fault Detection and Isolation for Reliability in Modern Power Systems. International Journal of Research Publications in Engineering, Technology and Management (IJRPETM), 5(4), 7106-7110.
57. Chandra, S., Rengarajan, A., Sahoo, G. S., & Sharma[4], S. (2024, October). Identifying Neuronal Damage and Plasticity by Analyzing Changes in Diffusion Tensor. In Proceedings of the 5th International Conference on Data Science, Machine Learning and Applications; Volume 2: ICDSMLA 2023, 15–16 December, Hyderabad, India (Vol. 2, p. 433). Springer Nature.