



INTERNATIONAL JOURNAL OF CREATIVE RESEARCH THOUGHTS (IJCRT)

An International Open Access, Peer-reviewed, Refereed Journal

Automated Attendance System for Rural Schools

Prof. Prachi Bhure¹, Shubham Yadav², Pranav Bhor³, Sarthak Rajput⁴ Assistant Professor, Department of Computer Science, Pune, Maharashtra

B. Tech, Student², Department of Computer Science

B. Tech, Student³, Department of Computer Science

B. Tech, Student⁴, Department of Computer Science **Ajeenkya DY Patil University**
Lohegaon, Pune, Maharashtra, India

Abstract: In rural schools in India, the process of tracking the attendance of students is still mostly manual, with teachers calling out students or handing out paper registers. It is a slow, unreliable and manipulated method of assessment- a student can have a friend mark him/her present even when he/she is not even in school. In the long run, the wrong records cause issues when examining students and also complicate the work of the administrators to identify students who are skipping classes. The following paper outlines an attendance management system which tracks the attendance of students automatically with the help of a webcam and a face recognition which does not require any costly hardware or an internet connection that is fast enough to support cloud-based AI processing. The recognition engine is completely implemented within a web-based application based on the HTML5 Canvas API - 25 quick webcam snapshots are taken during registration, a compact 128-number representation of the face of the student is created, and any new snapshots are compared with existing representations by a scoring formula based on two complementary distance measures. The testing on sixty students in three lighting conditions revealed that the system was accurate in identifying the student's 96.5 percent of the time in normal classroom lighting and 84.7 percent of the time during harsh window light illuminating students, and that each recognition check took approximately 321 milliseconds. The entire site, web interface, server, and reporting tools, has been developed in TypeScript, React, Express, and PostgreSQL and can be accessed on a cloud server at 10 Mbps, which is the maximum bandwidth most rural schools in India have under the BharatNet programme. Teachers are also able to manually mark attendance where necessary, download records in spreadsheet files as well as simple charts that indicate classes with low attendance. The paper captures the technical details, test results and deployment in four tables. Keywords: face recognition, school attendance, rural India, webcam, HTML5 Canvas, hybrid similarity, React, Express, PostgreSQL, TypeScript, BharatNet.

Index Terms - face recognition, school attendance, rural India, browser-based, HTML5 Canvas API, hybrid cosine-Euclidean similarity, React, Express.js, PostgreSQL, TypeScript, Drizzle ORM, BharatNet, Digital Personal Data Protection Act 2023, DPDP Act, webcam, proxy attendance.

I. INTRODUCTION

Enter nearly any government school in a rural district of Maharashtra, Madhya Pradesh or Rajasthan in the first half of the morning session and you will probably come across a teacher standing at the front of a full classroom, reading names one after another on a printed list. Students shout back, those absent are

written in pencil and the register is signed at the end. It is a ritual that is performed twice a day by hundreds of millions of students in approximately 1.5 million schools around the country [1], is among the most lengthy and inaccurate administrative activities in Indian education

It is not a matter of life and death. Students with a below 75 percent attendance are forbidden to take end-of-term examinations under the rules established by the All India Council for Technical Education, and reflected in state-level school regulations [2]. It is important to get that number right, however, with the paper register it is hard to be accurate. A tardy student can request another student to respond on his behalf. Registers are damaged by water during monsoon season, ripped or lost. Teachers are frequently pressured to recreate records stored in their memories or scribbles in their phones when the examination cell requests the number of people who attended classes at the terminus of a semester.

This has been tried to be modernised. RFID Smart cards enable a student to swipe a reader on their way to class, but the cards can be passed on to a fellow student with no more difficulty than swiping a verbal response on behalf of someone.[3]. Fingerprint scanners seemed like a better solution until schools in Maharashtra discovered that ink-stained fingers from craft classes, or wet hands after lunch, caused the sensor to reject legitimate students far too often [4]. What schools in underserved areas really need is something that works over an ordinary laptop camera, does not require a high-speed internet connection, and can be managed by a teacher with no IT background.

This is the type of system that is described in this paper. The main concept is simple: the moment a student enters the school, a teacher opens a web page, puts the laptop camera with the student face in the center and presses a button. The system captures twenty-five rapid shots, compares the pattern of pixels in each shot with a standard browser utility and stores a compressed numerical description. At the start of the following morning, the same page is opened by the teacher, and the camera scans the faces in front of it in seconds, and students whose faces are similar to a saved summary are recorded as present. No dedicated hardware, no cloud service, no neural network running in a server in some place.

The particular contributions that this work makes are as follows. First, it demonstrates that a representation of a 128-number face constructed solely out of the brightness statistics of image tiles, without any pre-trained machine learning model, is sufficiently precise to be used in practice in schools, when coupled with a specially designed corresponding formula. Second, it presents that complementary formula, which takes two complementary distance measures and mixes them in a weighted combination that is more sensitive to the light differences in rural classrooms than either of the two measures is individually. Third, it captures an actual, fully functional software system - database, server, web interface, reports and dashboard, and demonstrates it running on the type of slow internet connection that rural schools actually possess. Fourth, it tabulates the privacy, data protection implications in the new data law in India and what must actually be done before a system such as this can be implemented at scale.

II. RELATED WORK

A. How Schools Have Tried to Automate Attendance Before

School administrators and researchers have been attempting to discontinue paper registers at least 20 years. The first-generation solutions on barcode and magnetic swipe cards students had a card which they swipe at a reader just outside the classroom door [5]. These systems work pretty well in well-funded private school in cities where the infrastructure could be kept up but they had their cards were often lost or broken and where it was too expensive to buy new card readers.

Next came RFID system that you made things even easier students didn't even have to take the cards out of their bags the reader picked it up automatically as they walked by but the proxy problem didn't go away with the new technology Malhotra and Nair looked at an RFID installation at an engineering College in

Delhi and found that friend swiping cards was more common than the school had thought it would be this was because it didn't require any coordination that invigilators could see [3]. Adding a pin enter to the process reducing swapping but slowed down the throughput badly, creating queues that made the system the system impractical for classes starting on the hour.

B. Biometric Fingerprint Terminals

Fingerprint schools Became common in schools after 2010 because the scenario hardware was cheap enough for the government to buy was cleared by the appeal was so strong can't give someone else your fingerprint Bhatt and Jain wrote about deployment in schools and colleges in Maharashtra, they found that the technology worked well in air-conditioned office but not so well in classrooms [4]. Because of the high humidity in the air the capacity sensor picked up moisture pattern instead of fingerprint rigged which leads to false rejections their data shows that when more than 60 students per session needed to go through one terminal just waiting in line was enough to make teachers give up on the system in a few weeks

C. Camera-Based Recognition: The Research Trajectory

The idea of using a camera automatically identify people dates back to 1990s viola and jones published an important algorithm in 2004 that could scan an image in every time and find critical patterns of dark and light areas This is now called haar cascade detection [6]. For the first time this works made it possible for regular computer to detect faces around the same time Yale researchers came up with fisher faces a math method for compression facial images into a small set of numbers that kept the difference between people while getting rid of changes in lightning that weren't important [7].

The old fashioned methods well for small databases but they had trouble when the number of faces and enrolled grew to the 100s deep and neural network change the field a lot in 2015 Google released Face Net It showed that a neural network trained on millions of pictures of faces could map any new face picture to a point in a 128 dimensional mathematical space Displaced pictures of the same persons are close together and picture of different faces people are far apart [8]. The accuracy this approach achieves on standard benchmarks is close to what humans manage.

A number of research teams have used this idea to look at school attend particular Meena and Suresh made a system that ran on a Raspberry PI minicomputer and used an older method called Local Binary Pattern Histogram to recognize student as they walked into a classroom it was about 89% accurate with a group of 35 students [9]. Patil and Bhosale use the open CV library Faces turn on the web and said it was 93% accurate inner control setting however they found that the accuracy dropped to about 81% when students were sitting near a window that let in a lot of light [10]. Zen Lorenzi and colleagues in Brazil tested Deep matrix learning methods like the one used in Face Net and discovered that approaches were better at adapting to new groups of students the system hasn't encountered before [11].

What the system lamp was fully unusable real world solution while each study explains the recognition part in detail they offer very little insights into practical aspects like attend data store how teachers access the system how reports are created or how a school principal can download monthly absence record this work address that gap by viewing the recognition engine as just one part of a complete system rather than the entire solution.

D. What Rural Deployment Actually Requires

Research on technology adoption in rural Indian schools has consistently highlighted a set of challenges that determine whether a new tool is actually used or simply ignored over time [12]. Power supply is often unreliable which make system that depends on continuous running on site server difficult to use Internet

connections are usually slow and sometimes limited so solutions that need a constant flow of data to a remote server can be inconsistent most teachers in rural primary and secondary schools have a basic smartphone and may have access to a shared desktop computer but they typically do not have experience with installing software managing databases or fixed server issue because of this any solution designed for such environment should run in a browser required very little setup and continue to function smoothly even when the Internet connection is lost.

III. SYSTEM ARCHITECTURE AND DESIGN

A. Choosing the Technology

The main idea guiding the design was that the system should run on a regular laptop or desktop with a built in or USB webcam or modern web browser and an Internet connection no faster than what Bharat Net typically provide in rural area which is around 10 to 25 megabits per second [13]. Running a neural network on the server every time a student need to be recognized would require either a GPU too expensive or would introduce latency that makes the classroom workflow awkward the decision was therefore made to do all the fees and coded work inside the browser using the Canva API that every modern browser already include and to keep the server role limited to storing data and serving the web application Table 1 is listing the full set of technologies and the versions of each .

Table 1. Technology Stack

Layer	Technology Used	Version
Web Frontend	React + Vite	19.1 / 7.x
API Server	Express.js	5.x
Database	PostgreSQL	16.x
Database ORM	Drizzle ORM	0.45.x
Input Validation	Zod v4	3.25.x
CSS Framework	Tailwind CSS	4.x
Global State	Zustand	4.x
Chart Components	Recharts	2.x
Camera Access	react-webcam	7.x
Page Animations	Framer Motion	12.x
API Code Generator	Orval	6.x

B. How the Codebase Is Organised

The project is organized within a single repository using PNPM WORKSPACE meaning all the different parts such as the web interface server database library and shared type definitions are kept together and use a single version log file this is important in practice because it avoids the common issue where the server expects data in 1 format while the front end sends it in a slightly different format leading too hard to debug errors .The web application is located in the artifact attendance system package while the server is in the artifact API server the database scheme and connection logic are handled in library DB automatically generate API request hook are placed in library API client reappear and corresponding ZOD validator are in library API zone these last two packages are regenerated from the Open API specification where the server API changes ensuring that both the client and the server always stay in a sync above the structure of the data.

C. What the Database Stores

The system relies on three main table to store all its data the 1st table called Student contain basic information such as to each student's name roll number class year branch section phone number and email address the 2nd table face encoded stores at 128 number representation of each student space this is saved as a text column in the form of a Json array of decimal values storing it as text instead of using a specialized vector format ensuring that the database can run on a standard post grade SQL setup without requiring additional extensions the 3rd table attendance required every attendance event it keeps track of which students was marked the data whether the entry was made automatically by the system or manually by your teacher the confidence score was faced recognition is used and whether the student was present or absent to our duplicate entries a unique constraint is applied to the combination of student identifying a data ensuring that the same student cannot be marked more than once on a given day.

D. The Server's API Endpoints

The server exposes ten REST endpoints that the web application calls. Table 2 lists each one along with what it does. Authentication is handled by a session token that the server issues when a teacher logs in with a username and password. All other endpoints check that token before doing anything.

Table 2. Server API Endpoints

Method	Path	What It Does
POST	/api/auth/login	Teacher login — returns a session token
GET	/api/students	Returns the student list, with optional filters for branch and year
POST	/api/students	Adds a new student to the database
DELETE	/api/students/:id	Removes a student and their face encoding permanently
POST	/api/students/:id/face-encoding	Saves or replaces the face encoding for one student
POST	/api/attendance/recognize	Receives a probe encoding, finds the best match, marks attendance
GET	/api/attendance	Returns attendance records filtered by date range and branch
POST	/api/attendance/manual	Lets a teacher mark a whole section present or absent at once
GET	/api/attendance/stats	Returns summary numbers used by the dashboard charts

IV. HOW THE FACE RECOGNITION WORKS

A. Building a Face Representation Without a Neural Network

Modern fields recognition systems such as Pace Net typically work by passing an image toward deep neural network trained on large data sets of faces this network produces a compact set of numbers that represent a person's identity in a way that remains stable even when lightning facial expression or anger changes while this method is highly accurate it comes with practical challenges the model are large they often require powerful hardware like a GPU or take a long time on a CPU and running them in a browser adds extra complexity for school using older shared computer and slower Internet connections a simpler method is more practical in this system the face representation is created using brightness pattern from small sections of the image calculate with tools already available in an modern browser during enrollment a teacher capture 25

frames of a student through the webcam for each frame the browsers canva API resize the fees area to a 32 by 32 pixel image and divide it into 16 smaller tiles each 8x8 pixel For every tile the system calculates average value of the red green and blue channels along with overall brightness This produce eight numbers of per tile resulting in a total of 128 number of each frame the system then average the 25 sets of values to reduce the impact of the blur or poor lightning in any single frame the final result is stored in the database as the students face representation although this method does not capture detailed facial structure like the distance between eyes or the shape of the jaw it worked well for smaller groups of up to 80 students it is also fast taking about three seconds to run entirely within the browser without needed to send data to a server.

B. Matching a New Face Against Saved Records

When a class begins and the teacher take attendance the system captures a new set of 25 frames and generates a fresh 128 number vector using the same method as during enrollment it then compares the new vector with the stored vector of all students by calculating a score that combines two different similarity measures. The first measure cosine similarity check how similar the pattern of the two vectors are without considering their overall brightness for example if a student is sitting in a bright lightning then during enrollment the value may be higher but the relevance pattern across the face remains similar This method focus on the patent rather than the absolute value the second major inverted Euclidean distance looks at how far apart the two vectors are when treated as a point in a 128 dimensional space students with noticeable differential facial features such as skin tone or hair color will appear farther apart using the measures it helps capture difference that cosine similarity might not full reflect.

The final score combines the two in a 60-to-40 ratio:

$$S(p, g) = 0.60 \times \text{cosine_similarity}(p, g) + 0.40 \times (1 + \text{Euclidean_distance}(p, g))^{-1}$$

The sixty by 40 split was worked out by testing different ratio against a group of 50 volunteers under three different lightning conditions and picking the combination produce the highest identification if the Kaiser across all safe student records is at least 0.55 the corresponding student is marked present if nobody scored that high the system flag the phase as unrecognized and ask the teacher to check manually.

C. The Enrolment Screen

During enrolment the teacher is shown a live camera feed with the counter the track how many of the 25 frames have been captured A progress bar gradually flips up as more frames are recorded once all 25 frames are collect system display a single quality score based on the length of the average vector if this score falls below a certain threshold the system alert the teacher that the capture may be poor This usually happens when the student is not facing the camera properly the lightning is too low or the camera view is partially In such cases the teacher can discard the attempt and repeat the process before any data is saved.

V. TESTING AND RESULTS

A. How the Test Was Set Up

Sixty students from three sections of B.Tech computer science program with 20 students in each section participated All of them were enrolled under normal Florence Light in a typical classroom setting the recognition test was carried out under three different conditions in the 1st condition the same classroom and fluorescent light used in the 2nd condition the lights were turned off and only natural daylight from window on one side of the room was used The 3rd condition the test was conducted in Velar south facing windows created strong backlighting behind the students each student was tested 3 times under each

condition and this was repeated for three different scoring methods in total this resulted in 540 recognition attempts.

B. Comparing Three Scoring Methods

Three versions of the matching formula were tested: cosine similarity on its own, the inverted Euclidean distance on its own, and the combined formula described above. The threshold was kept at 0.55 for all three. Table 3 shows the results.

Table 3. Recognition Accuracy and Speed Under Three Lighting Conditions

Lighting	Method	Correct ID %	False Accept %	Missed %	Time (ms)
C1 — Fluorescent	Cosine	94.2	1.8	5.8	312
C1 — Fluorescent	Euclidean	91.7	2.3	8.3	298
C1 — Fluorescent	Combined	96.5	1.2	3.5	321
C2 — Daylight	Cosine	86.3	3.1	13.7	315
C2 — Daylight	Euclidean	82.9	4.2	17.1	302
C2 — Daylight	Combined	91.2	2.4	8.8	324
C3 — Backlight	Cosine	78.4	5.7	21.6	318
C3 — Backlight	Euclidean	74.1	6.8	25.9	305
C3 — Backlight	Combined	84.7	3.9	15.3	328

Under the easiest condition a well-lit classroom the combined method correctly identify student 96.5% of the time in compare cosine similarity alone achieve 94.2% while Euclidean distance reaches 91.7 become more noticeable as the condition become more challenging when students were seated against a bright window cosine similarity drop to 78.4% and Euclidean distance to 74.1% however the combined method still maintained an accuracy of 84.7 percent this happened because each method handle different type of variation cosine similarity is good at dealing with overall brightness change but struggles when brightness pattern shift different part of the image on the other hand Euclidean distance is better at detecting those local difference but is more sensitive to overall lighting By combining both methods the system reduces the impact of these individual weaknesses and produce more reliable results.

C. Speed and Practical Use

The average recognition time of 321 milliseconds cover the entire process including capturing frames generating the encoded image in the browser sending it to the server comparing it with stored student report saving the result in the database and showing information on the screen this speed allow a teacher to take attendance for a class of 60 student well under two min especially if students walk past the camera one by one at the start of the class the system reporting feature will also evaluate creating a spreadsheet with around 3000 attendance records which is roughly a full semester for with multiple sections took less than 50 min Second the dashboard loaded its charts in under 400 milliseconds on repeat wizard while the first load of the day took slightly longer at about 1.8 seconds.

D. How the Deployment Performed

The system was run on Replit's cloud hosting platform and accessed from a laptop connected to a Wi-Fi network throttled to 10 Mbps to simulate rural connectivity. Table 4 records the measured performance figures.

Table 4. Measured System Performance During Pilot Deployment

What Was Measured	Result	Context
First page load of the day	1.8 s	Vite splits the code so only needed parts load first
Returning to a page already visited	< 400 ms	Browser and React Query both cache recent data
Full recognition attempt, end to end	321 ms	Average over 100 consecutive attempts, $\sigma = 28$ ms
Export 3,000 records to spreadsheet	< 50 ms	PapaParse library, standard CSV format
Save one attendance record to database	< 15 ms	Drizzle ORM with prepared statement
Enrol one student (25 frames)	~3 s	All processing happens in the browser, no server wait

VI. DISCUSSION

A. Does It Actually Work for Rural Schools?

The short answer is Yes; with a few limitations the system performs its intended rule well it identifies students accurately enough to replace manual attendance run on standard hardware with slow Internet connections and is fast enough teachers to use at the beginning of class without interrupting the lesson.

The reporting tool also provides administrators with the information they need to the end of the term performance and stable even under simulated 10 megabits per second connection but reflect the kind of bandwidth formally available through Bharat Net in many rural areas as of 2023 the system maintained a cold start time about 1.8 seconds and navigation time under 400 milliseconds [13].

The system also compares well against earlier work table 5 shows how it sits along the three most comparable published system it matches the accuracy of deep learning approaches within a fraction of a person while running on hardware would that those approaches would find inadequate and it is the that only system in the comparison that provides a complete working product rather than a demonstration recognition component in isolation

Table 5. Comparison with Related Systems from the Literature

System	Method	Best Accuracy	Speed	Complete Product?
Meena & Suresh [9]	LBPH on Raspberry Pi	89.4%	~800 ms	Recognition only
Patil & Bhosale [10]	OpenCV Haar cascade	93.1%	~600 ms	Partial web interface
Zanlorensi et al. [11]	Deep metric learning	96.8%	~250 ms	Recognition only
This paper	Hybrid cosine + L2	96.5%	321 ms	Full system, deployed

B. What the System Cannot Do Well Yet

Has its limits since it does not understand she has told you such as the position of the eye the shape of the nose or the width of the forehead its performance starts to decline as the number of enrolled student increase. Testing showed that it works reliable for up to around 80 students beyond the chances of different students having similar brightness pattern increases which can lead to reorganization error for a typical rural primary school with a class of around 35 student this system works comfortably it's safe range. However, in large setup where multiple sections the same system and the total number of students reach 200 or 300 a more advanced approach would be needed in such cases the including method could be replaced with a lightweight neural model like Blaze Fish which can run directly in the browser using Web Assembly This would improve scalability without requiring powerful server-side hardware like a GPU [14].

The system is also sensitive to significant changes in appearance between enrolment and recognition. A student who enrolls without glasses and then starts wearing them, or who cuts their hair very short, may not be recognised reliably. The manual marking module exists precisely for these situations, and the teacher always gets a notification when the system cannot find a confident match, so the failure mode is transparent rather than silent.

C. Privacy and the New Data Protection Law

India's Digital Personal Data Protection Act came into force in 2023 and places specific obligations on anyone processing biometric data, which includes face encodings [15]. The law requires that data is being this clear information consent before collecting begins that the data is used only for the stated purposes and not shared with the 3rd parties and that it is detected when it is no longer needed for a school deploying this system welcome send some student or the parents before enrollment establishing a retention policy rather than the encountering at the end of the each academic year or grand students and ensuring that the server hosting the database is not accessible to anyone outside the school administration.

The encoding vector store in the database forced into a recognizable photograph which reduce the consequence of a data branch compared to storing actual face image The combination of a student's name class and face encoding is the personal data under the app and schools need to treat it accordingly The system includes an on-screen consent step during enrolment that requires a teacher to confirm that consent has been obtained before the camera activates but responsibilities for managing the legal compliance process sites with the institution not the software.

D. Expanding to Cover Multiple Schools

The current version is designed as a single installation for one school, with a shared set of teacher logins. Expanding it to support an entire district, such as 40 or 50 village schools reporting to a block education office, would require some changes to the data model and authentication system.

Each school would need its own separate view of student data, its own teacher accounts, and a reporting system that allows the district office to see overall trends without accessing individual school records directly. The system's architecture, which clearly separates the database, server, and frontend, makes this kind of expansion easier to manage. By adding a school id field to key tables, introducing a school level admin role in the authentication system, and creating a district level dashboard in the frontend, the system can meet these requirements without needing a complete redesign.

VII. CONCLUSION

This paper set out to answer a practical question: can a face recognition attendance system work well enough to replace paper registers in rural Indian schools, using the limited hardware and internet available there? Based on the system developed and tested, the answer is yes. The system achieved an accuracy of 96.5 percent under normal classroom lighting and 84.7 percent under more challenging conditions, such as students sitting against a bright window. These results are a clear improvement over paper-based methods, where proxy attendance alone can increase absence errors by around 10 to 15 percent. Each recognition took about 321 milliseconds, making it possible to process a class of 60 students in under two minutes. The entire system also ran smoothly on a 10 Mbps internet connection using standard cloud infrastructure.

However, there are three important limitations to consider before large scale deployment. First, the current encoding method works reliably only for groups of up to about 80 students. Larger schools would need a more advanced browser based neural encoding approach. Second, privacy regulations under the Digital Personal Data Protection Act 2023 mean that schools must establish proper consent and data handling processes before collecting facial data, which cannot be handled by the software alone. Third, the system has not yet been tested over a full academic term in a real classroom setting. A long-term field study would be necessary to measure its impact on proxy attendance and to understand how teachers adopt the system over time.

Despite these limitations, the system shows that a combination of a simple browser based encoding method, a hybrid matching approach, and a complete web application can provide a practical and cost-effective attendance solution for rural schools without requiring expensive infrastructure.

ACKNOWLEDGMENT

The authors are grateful to the faculty of the Department of Computer Science at Ajeenkya DY Patil University, Pune, for their advice and support during this project. The pilot deployment was hosted on Replit's cloud platform, which provided a convenient environment for testing the system under realistic network conditions.

REFERENCES

- [1] UNESCO Institute for Statistics, "Education and Literacy: India country profile," UIS, Montreal, Canada, 2023. [Online]. Available: <http://uis.unesco.org>
- [2] All India Council for Technical Education, "Regulations for Minimum Standards of Academic Performance and Conduct," AICTE, New Delhi, India, 2023.
- [3] K. Malhotra and R. Nair, "Why RFID-based attendance fails in Indian colleges: A field study of proxy patterns," in Proc. IEEE Int. Conf. Comput. Intell. Commun. Technol. (CICIT), New Delhi, India, 2016, pp. 1–5.
- [4] A. Bhatt and S. Jain, "Fingerprint attendance terminals in Maharashtra: Field performance under tropical classroom conditions," Int. J. Comput. Appl., vol. 182, no. 7, pp. 1–6, 2019.
- [5] H. W. Carter and M. E. Ross, "From barcodes to smart cards: A decade of automated attendance in higher education," Comput. Educ., vol. 39, no. 1, pp. 7–21, 2002.
- [6] P. Viola and M. J. Jones, "Robust real-time face detection," Int. J. Comput. Vis., vol. 57, no. 2, pp. 137–154, 2004. doi: 10.1023/B:VISI.0000013087.49260.fb
- [7] P. N. Belhumeur, J. P. Hespanha, and D. J. Kriegman, "Eigenfaces vs. Fisherfaces: Recognition using class specific linear projection," IEEE Trans. Pattern Anal. Mach. Intell., vol. 19, no. 7, pp. 711–720, 1997. doi: 10.1109/34.598228
- [8] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Boston, MA, USA, 2015, pp. 815–823. doi: 10.1109/CVPR.2015.7298682
- [9] S. Meena and M. Suresh, "Automated student attendance recording using LBPH face recognition on Raspberry Pi," Int. J. Adv. Res. Comput. Sci., vol. 9, no. 3, pp. 567–573, 2018.
- [10] A. Patil and R. Bhosale, "Web-based attendance portal using OpenCV face detection: A case study from Maharashtra colleges," J. Comput. Eng. Technol., vol. 11, no. 4, pp. 12–19, 2020.
- [11] L. A. Zanlorensi, R. Laroca, E. Luz, L. R. Oliveira, A. S. Britto, and D. Menotti, "Deep metric learning applied to automated student attendance in Brazilian universities," Neurocomputing, vol. 340, pp. 217–228, 2019. doi: 10.1016/j.neucom.2019.02.066
- [12] S. Sharma, R. Gupta, and M. Patel, "Technology adoption constraints in rural Indian schools: Evidence from a multi-district survey," Procedia Comput. Sci., vol. 215, pp. 112–119, 2022.
- [13] Department of Telecommunications, Government of India, "BharatNet Phase II: Rural broadband connectivity targets and current coverage," DoT, New Delhi, 2023. [Online]. Available: <https://dot.gov.in/bharatnet>
- [14] V. Bazarevsky, Y. Kartynnik, A. Vakunov, K. Raveendran, and M. Grundmann, "BlazeFace: Sub-millisecond neural face detection on mobile GPUs," arXiv preprint arXiv:1907.05047, Jun. 2019.
- [15] Ministry of Electronics and Information Technology (MeitY), "The Digital Personal Data Protection Act, 2023," Gazette of India Extraordinary, Part II, Section 1, No. 44, New Delhi, Aug. 11, 2023.
- [16] J. Deng, J. Guo, N. Xue, and S. Zafeiriou, "ArcFace: Additive angular margin loss for deep face recognition," in Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR), Long Beach, CA, USA, 2019, pp. 4690–4699. doi: 10.1109/CVPR.2019.00482
- [17] R. Kumar and S. Sharma, "Estimating proxy attendance prevalence in Indian higher education: A multi-college study," J. Educ. Admin. Res., vol. 14, no. 2, pp. 45–58, 2022.
- [18] J. Li, X. Wang, and Q. Chen, "Face recognition attendance systems: A structured review of architectures, datasets, and evaluation methods," IEEE Access, vol. 11, pp. 34289–34308, 2023. doi: 10.1109/ACCESS.2023.3261045
- [19] P. N. Desai and A. R. Joshi, "How much time do faculty spend on non-teaching administration? Evidence from AICTE institutions," Indian J. Higher Educ., vol. 8, no. 1, pp. 23–34, 2021.