DOI.ORG/10.67505/IJIAMR.v6.i1.61010
E-ISSN: 2583-4452

SmartFace: An Offline-First AI-Based Attendance Management System for Rural Schools Using MobileFaceNet

Muskan Behl[1], Parishi Mishra[2*], Deepali Bajaj[3] and Pushkar Gole[4]
[1, 2, 3, 4] Department of Computer Science, Shaheed Rajguru College of Applied Sciences for Women, University of Delhi, Delhi
[*]Corresponding Author ✉ parishimishra19@gmail.com

ABSTRACT

Manual recording of student attendance in rural schools is a time-consuming and error-prone task. Moreover, attendance records are necessary for keeping a check on educational activities and promoting social welfare schemes. In literature, most existing attendance automation systems rely on continuous internet connectivity, cloud-based processing, or costly hardware infrastructure, which can limit their scalability, affordability, and ease of deployment. Hence, a lightweight attendance system named SmartFace has been designed and developed in this research work which uses facial recognition technology for marking students' attendance digitally. SmartFace makes use of MobileFaceNet, a Convolutional Neural Network model that allows facial recognition on the device itself, conserving processing resources and eliminating the need for server-side processing. The proposed system is developed as a Progressive Web Application (PWA) and it can function without an internet connection and synchronize data with Firebase Firestore once connectivity is restored. The MobileFaceNet model of SmartFace application was trained and tested on the Indian Faces Image Classification Dataset obtained from Kaggle data repository. On evaluation, the fine-tuned model achieved a training accuracy of 99.43%, with a validation accuracy of 93.75% and a test accuracy of 96.88%, demonstrating the viability of the proposed approach on the dataset used. These findings suggest the potential of lightweight facial recognition technology as a viable approach for efficient and economical attendance management in rural schools, warranting further evaluation on larger and more diverse datasets.

Keywords: *Face Recognition, Rural Schools, Automated Attendance System, MobileFaceNet, Progressive Web Application (PWA), Offline-First Architecture*

1 Introduction

Attendance management in Indian schools has historically relied on manual roll calls, a practice that is both time-consuming and susceptible to inaccuracies. Manual attendance recording remains a burdensome administrative task, often reducing the instructional time available to teachers during the school day. Apart from this, manually marking attendance may allow proxy attendance, resulting in false attendance records. This adversely affects the actual beneficiaries from availing the benefit of government funded programs such as Mid-Day Meal schemes and other scholarships. These inaccuracies hinder allocation of government funding to the schools and further lead to a drop in student retention rates.

The increase in availability of affordable smartphones and tablets in rural India has also paved the way to digitize school administration. The existing digital attendance solutions have been designed according to available urban infrastructure requiring dedicated hardware, consistent high-speed internet connectivity, and trained technical personnel. However, these solutions are not feasible for about 1.23 million rural schools of India due to unreliable network access and power outages.

A unified deep embedding framework that maps facial images to a compact Euclidean space and achieves high level of accuracy on standard benchmarks was first introduced by Schroff et al. (2015). MobileFaceNets, a collection of highly efficient Convolutional Neural Network (CNN) models created especially for accurate, real-time face verification on mobile devices, was later introduced by Chen et al. (2018). Recent research including ArcFace (J. Deng et al., 2022) and improved versions of MobileFaceNet (Hassanpour & Kowsari, 2023), has demonstrated that high quality recognition is still achievable with reduced computational costs. Although several attendance management systems have been developed, none have been designed for infrastructure limited settings; server-side processing, dedicated hardware, and a stable internet connection remain the primary barriers.

Received on: 13 Jun 2026 | Accepted on: 01 Jul 2026 | Published Online on: 25 July 2026

To solve the above problems, the paper provides SmartFace, an AI-driven attendance management system that has been built to eliminate infrastructure deficiencies. SmartFace is designed as a Progressive Web Application that can run on standard web browsers without requiring installation. Using a lightweight deep learning model, facial recognition is performed on-device (Chen et al., 2018) which has been selected for its compact architecture and suitability for browser-based inference (Hassanpour & Kowsari, 2023). Attendance records are synchronized immediately when connectivity is available, and in case of poor network availability attendance records are stored locally, then synchronized with cloud storage when connectivity permits, ensuring no data loss. The objectives of this research are: (1) to develop a lightweight AI-based attendance management system using facial recognition, (2) to enable on-device face recognition without server-side computation, (3) to support offline attendance functionality using an offline-first architecture, and (4) to evaluate the effectiveness of the proposed system using a publicly available face dataset. The remainder of this paper is structured as follows: Section 2 reviews relevant literature and the research gap addressed by this work; Section 3 describes the methodology; Section 4 presents the experimental setup; Section 5 reports the results; and Section 6 concludes the paper with directions for future research.

2 Literature Review

Face recognition is a subfield of machine learning that has received increasing attention over the last decades. A number of models with different requirements have been developed. In Schroff et al. (2015), the FaceNet architecture, which maps a face image into a compact Euclidean space where distances are indicative of face similarity, is introduced. The FaceNet is trained on millions of faces using a triplet loss function, achieving 99.63% accuracy on the Labelled Faces in the Wild (LFW) dataset. This model, however, is computationally intensive and cannot be directly deployed on mobile devices or embedded systems without severe optimization.

To meet the demands of mobile inference, MobileFaceNets were developed in Chen et al. (2018). Mobile-FaceNets is a family of lightweight CNN architectures optimized to perform real-time face verification on mobile platforms. These models adopt depthwise separable convolutions that are also part of the MobileNet architecture (Sandler et al., 2018), and a global depthwise convolution layer replaces the fully connected layers, thus significantly reducing both the number of parameters and computation. Lightweight model research has produced other competitive architectures, including ShuffleFaceNet models by Mart´ýnez-D´ýaz et al. (2019) under 20 MB and averaging 37ms of CPU inference time. A more recent contribution is the 3.5MB few-shot face recognition model presented in Z.-Y. Deng et al. (2023), which is about 30 times smaller than the FaceNet model, while achieving comparable recognition accuracy and real-time performance. Further in this line, Hassanpour and Kowsari (2023) improved on the MobileFaceNet model to achieve top ranked performance on EFaR-2023 face recognition competition under efficiency constraints. Another very light 3.5MB FewFaceNet by Sufian et al. (2023) has shown effective few-shot face authentication for edge cameras with just one enrolment image. State-of-the-art of unconstrained face recognition has been improved to 99.83% on LFW by J. Deng et al. (2022) using the Arc-Face method. In Sun et al. (2024), a landmark based self-supervised model LAFS is introduced that outperforms the state of the art by a significant margin on few-shot recognition benchmarks without requiring large, labelled datasets.

2.1 Automated Attendance Systems

Deep learning has been applied to automate student attendance. The most common pipeline described in the literature comprises face detection using MTCNN (Zhang et al., 2016) followed by embedding based matching of student images. The system has been shown to work efficiently on urban or well-funded institutional environments, but evaluation has been restricted to those environments. Chitrapu et al. (2025) even showed that an ensemble of deep learning models could achieve reliable face recognition performance, though such architectures also entail a higher computational overhead.

2.2 Research Gap

Existing attendance automation systems found in the literature were almost exclusively designed and evaluated in contexts very different to the realities faced in rural Indian schools. The most common architecture design pattern is cloud based: face detection and recognition are offloaded to remote servers and a consistent network connection is needed, which translates in latency, higher cost and lack of functionality when the network connection is intermittent. The hardware aspect is a second issue: a few researchers have used specialized enrolment terminals, IP cameras with high resolution or powerful edge computers like the Raspberry Pi with GPU accelerators. Such hardware might be justifiable in a wealthy context, but the cost is far too high for rural schools. Lightweight models (Z.-Y. Deng et al., 2023; Hassanpour & Kowsari, 2023; Mart´ýnez-D´ýaz et al., 2019) achieve a much better performance-resource ratio, but this line of research has not yet been applied to the deployment in rural educational contexts. SmartFace tackles these problems, and leverages MobileFaceNet (Chen et al., 2018) for browser-based inference so no server-side computation per recognition event occurs. Offline-first IndexedDB storage allows the system to remain functional even when internet connectivity is unavailable, which matches the infrastructure-less design paradigm promoted by Sufian et al. (2023) who developed FewFaceNet particularly for edge cameras.

Table 1 summarizes representative face recognition systems alongside SmartFace for comparative reference.

Table 1: *Summary of Related Face Recognition Systems*

Reference	Model / Approach	Dataset	Key Results
Chen et al. (2018)	MobileFaceNets	LFW, MegaFace	99.55% (LFW)
Schroff et al. (2015)	FaceNet (triplet loss)	LFW, YTF	99.63% (LFW)
J. Deng et al. (2022)	ArcFace	MS1MV2	99.83% (LFW)
Z.-Y. Deng et al. (2023)	FN8 (3.5 MB few-shot model)	CASIA-WebFace	98.4% (LFW)
Hassanpour & Kowsari (2023)	Improved Mobile-FaceNet	EFaR-2023 benchmark	Ranked 1st on EFaR-2023 benchmark under efficiency constraints
Mart´ýnezD´ýaz et al. (2019)	ShuffleFaceNet	LFW, MegaFace	Outperformed comparable lightweight models
SmartFace (This work)	MobileFaceNet (on-device)	Indian Faces Image Classification Dataset (Garg et al., 2023)	99.43% train accuracy (Val: 93.75%, Test: 96.88%)

3 Methodology

This section is organized into 5 subsections. Section 3.1 describes the dataset used in the current study. Section 3.2 describes the data pre-processing pipeline used for the raw images. Section 3.3 describes the model used in this study for on-device face recognition. Section 3.4 describes the training methodology used to fine-tune the model. Section 3.5 describes the overall end-to-end attendance recognition system used in the SmartFace application.

3.1 Dataset

The Indian Faces Image Classification Dataset (Garg et al., 2023) hosted on Kaggle has been used for the training and testing of the model. Students created this dataset where every student contributed 20 facial images each. The images were collected under various circumstances; e.g., with differing illumination, varying face angles, with and without specs, with diverse backgrounds. All in all, there were 320 images consisting of 16 students. After splitting the data into train and test, 14 images per student were selected for training (70%) and the rest 6 for testing (30%). Dataset statistics are summarized in Table 2.

Table 2: *Dataset Statistics*

Property	Value
Total Students	16
Total Images	320
Images per Student	20
Training Images per Student	14
Testing Images per Student	6
Train Split	70% (224 images; 179 train + 45 val)
Validation Split	20% of train pool (45 images)
Test Split	30% (96 images)

3.2 Data Preprocessing

All images were preprocessed using a fixed pipeline before training. Faces were first detected and cropped from images to generate standard inputs. Cropped regions were resized to the MobileFaceNet input resolution, and their pixels were normalized to prevent numerical instability during training. The detected faces were then aligned using landmark-based transformation (Zhang et al., 2016). This process centers facial keypoints (pupils, nose tip, lip corners) onto predefined positions, which is critical to achieve discriminative representations and avoid different embeddings from the same face but with varying orientation. To prevent over-fitting data augmentation including horizontal flipping and subtle changes in brightness were applied.

3.3 Model Architecture

MobileFaceNet (Chen et al., 2018) was used as the backbone due to its lightweight nature, which is beneficial for deployment on embedded devices with limited computational resources. The architecture extends MobileNetV2 (Sandler et al., 2018) to adapt it for face recognition. Instead of fully connected layers at the end, it uses a global depthwise convolution and a bottleneck design to minimize memory usage at inference time. These properties make the model ideal for on-device execution in a web browser. Unlike heavily parameterized models such as ArcFace (J. Deng et al., 2022) or FaceNet (Schroff et al., 2015), it demonstrates effective recognition performance with acceptable parameter count for JavaScript-based on-device inference, which is aligned with the design goals of ShuffleFaceNet (Mart´ýnez-D´ýaz et al., 2019) and its successor (Hassanpour & Kowsari, 2023).

This model is trained to output a 128-dimensional facial embedding vector for each input image. During enrolment, the student's face is imaged and its average embedding representation is saved locally. In the recognition phase, the live embedding of an input face image is then computed and cosine distance against stored embedding is calculated. A match is identified when the cosine distance falls below a predefined threshold value.

3.4 Training Procedure

The model was fine-tuned on the collected dataset using the training split described in Section 3.1. A standard cross-entropy classification loss was used for the 16-class identity problem. Images were loaded from the structured folder layout (16 student folders), preprocessed, and passed through the network in mini batches of size 16. Parameters were updated using stochastic gradient descent with momentum. Training was performed for a maximum of 100 epochs with early stopping using a patience value of 20 to reduce overfitting. A learning rate scheduler (StepLR) was used to improve training stability. After training, the average enrolment embedding per student was computed as the stored reference representation for live recognition.

Table 3: *Training Configuration*

Hyperparameter	Value
Input Resolution	112×112
Embedding Dimension	128
Loss Function	Cross-Entropy
Optimizer	SGD with Momentum (0.9)
Learning Rate	0.01
Weight Decay	5×10^{-4}
Batch Size	16
Learning Rate Scheduler	StepLR
Maximum Epochs	100
Early Stopping Patience	20

3.5 Attendance Recognition Pipeline

The end-to-end attendance workflow follows a structured sequence:

1. Teacher logs in via Google Firebase Authentication (Bharti et al., 2020);
2. School and class are selected, loading the corresponding student enrolment data;
3. The device camera is activated and face detection begins scanning the video feed in real time;
4. MobileFaceNet (Chen et al., 2018) generates an embedding vector for each detected face;
5. Cosine distance matching against stored embeddings marks the student as present;

6. Unmatched faces cause a fallback to QR code or manual marking;
7. Attendance details are stored on the local IndexedDB, and they get synced with Firebase when a network connection becomes available.
8. A CSV report is generated for record-keeping and compliance purposes.

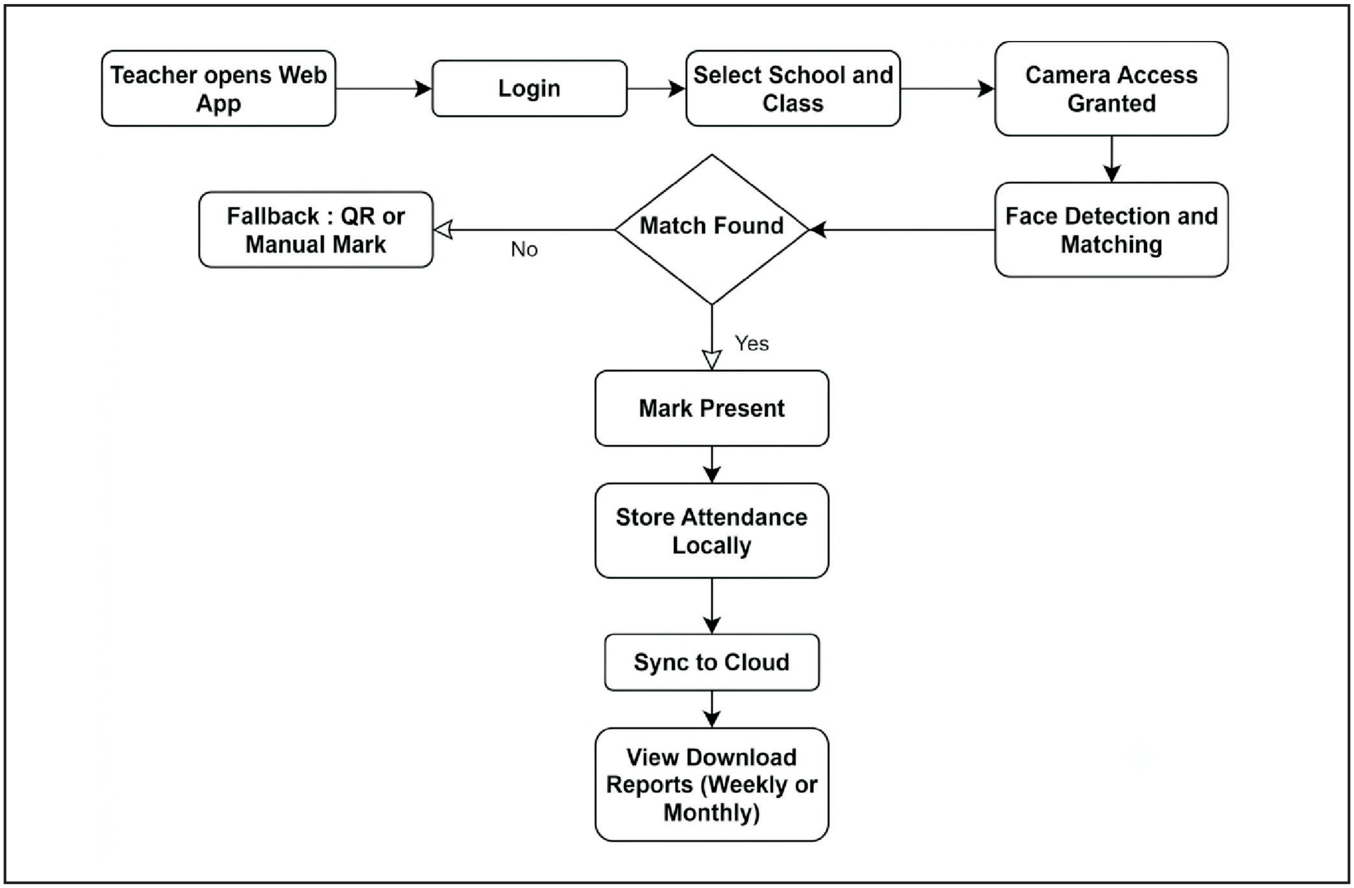

Fig. 1: *SmartFace Technical Workflow and Attendance Recognition Pipeline*

4 Experimental Setup

This section summarizes the details of experimental procedures. There are two sub-sections in this section. Section 4.1 describes the features of the dataset and the metric used in experiment. Section 4.2 lists the detailed hardware configuration and provides details about the software environment used in the experiment.

4.1 Dataset and Evaluation Metric

All experiments were conducted using the Indian Faces Image Classification Dataset (Garg et al., 2023), containing 320 images of 16 individuals, split into a 70% training set (224 images) and a 30% test set (96 images). From the training set a further 20% split was drawn to provide a 45 image validation set, training was reduced to 179 images and 96 images remained in the test set. The validation set was used to monitor model generalisation during training and support early stopping. Evaluation metrics and calculations were made using the standard scikit-learn python library whereas inference latency was calculated from average times of runs on images in the test set.

MobileFaceNet (Chen et al., 2018) was selected as the recognition model because of its lightweight design and high performance on mobile devices. Recognition accuracy was used as the primary evaluation metric in this study. Furthermore, the precision, recall, F1-score, AUC-ROC, and inference latency were also included in the evaluation criteria.

4.2 Hardware and Software Specifications

The SmartFace prototype was created and experimented on consumer devices to model the resource limited environment where the SmartFace system will be placed. The web application was implemented as a Progressive Web Application (PWA) using HTML, CSS, and JavaScript with cloud storage through Firebase Firestore. Face recognition was implemented using a lightweight MobileFaceNet model and integrated with browser-based inference.

The application was tested on both laptop and Android smartphone devices to check its cross platform compatibility, offline operability and real-time attendance tracking with IndexedDB used as the storage for offline entries and firebase firestore to automatically sync the offline entries when connection becomes available.

5 Results and Discussion

This section is divided into four subsections. Section 5.1 presents and discusses the recognition performance of the proposed system. Section 5.2 compares SmartFace against existing attendance methods. Section 5.3 describes the features and user interfaces of the SmartFace application. Section 5.4 discusses the limitations and threats to validity of the current study.

5.1 Recognition Performance

The fine-tuned MobileFaceNet model achieved a training accuracy of 99.43%, a validation accuracy of 93.75%, and a test accuracy of 96.88% on the held-out 96-image split across 16 student classes. This result is consistent with the high accuracy reported for lightweight models on small, controlled datasets: Z.-Y. Deng et al. (2023) similarly observed high performance for their 3.5 MB model on in-domain test splits, while Hassanpour and Kowsari (2023) demonstrated that MobileFaceNet variants retain competitive recognition under significant parameter constraints. However, the reported accuracy should be interpreted in the context of the dataset used, which consists of 320 images from 16 subjects collected under controlled conditions. Though the experimental results indicate the feasibility of the proposed method, the method could be verified on a larger and more heterogeneous dataset to thoroughly assess its robustness on different realistic scenarios as future research.

Table 4: *Performance Metrics of SmartFace*

Metric	Value
Train Accuracy	99.43%
Validation Accuracy	93.75%
Test Accuracy	96.88%
Macro Precision	97.32%
Macro Recall	96.88%
Macro F1-score	96.85%
Weighted F1-score	96.85%
Mean AUC (OvR)	0.9999
Mean Inference Latency	382.4 ms (P95: 1025.0 ms)

The training curves in Figure 2 show steady convergence with no signs of severe overfitting; the small gap between training and validation loss indicates that the model generalises well within this dataset. A practical strength observed

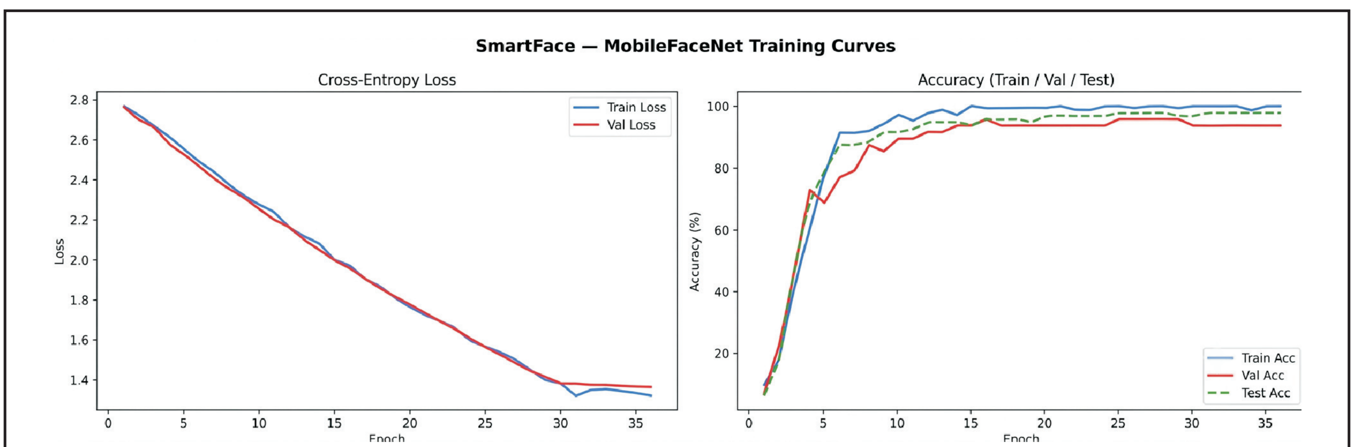

Fig. 2: *MobileFaceNet Training Curves: Cross-Entropy Loss and Train/Validation/Test Accuracy over 36 Epochs*

during testing was the model's resilience to spectacles and moderate lighting changes, which were deliberately included in the image collection protocol, consistent with the multi-condition robustness reported by Sufian et al. (2023) for FewFaceNet under similarly varied enrolment conditions.

To further validate the choice of recognition strategy, the 128-dimensional embeddings produced by the trained model were evaluated under four different classification approaches on the same test set. The results are summarised in Table 5.

Table 5: *Classifier Comparison on Frozen MobileFaceNet Embeddings (Test Set, n = 96)*

Classification Method	Test Accuracy	Classification Method	Test Accuracy
Cosine Nearest-Centroid	**96.88%**	Softmax Classifier Head	95.83%
KNN (*k*=1, cosine distance)	95.83%	SVM (RBF kernel)	95.83%

The cosine nearest-centroid approach used in SmartFace outperformed all alternative classifiers, confirming the suitability of the paper's recognition strategy. This method is also the most efficient at inference time, as it requires only a single cosine distance computation per enrolled student rather than a learned classifier head. Figure 3 illustrates the accuracy comparison across all four methods.

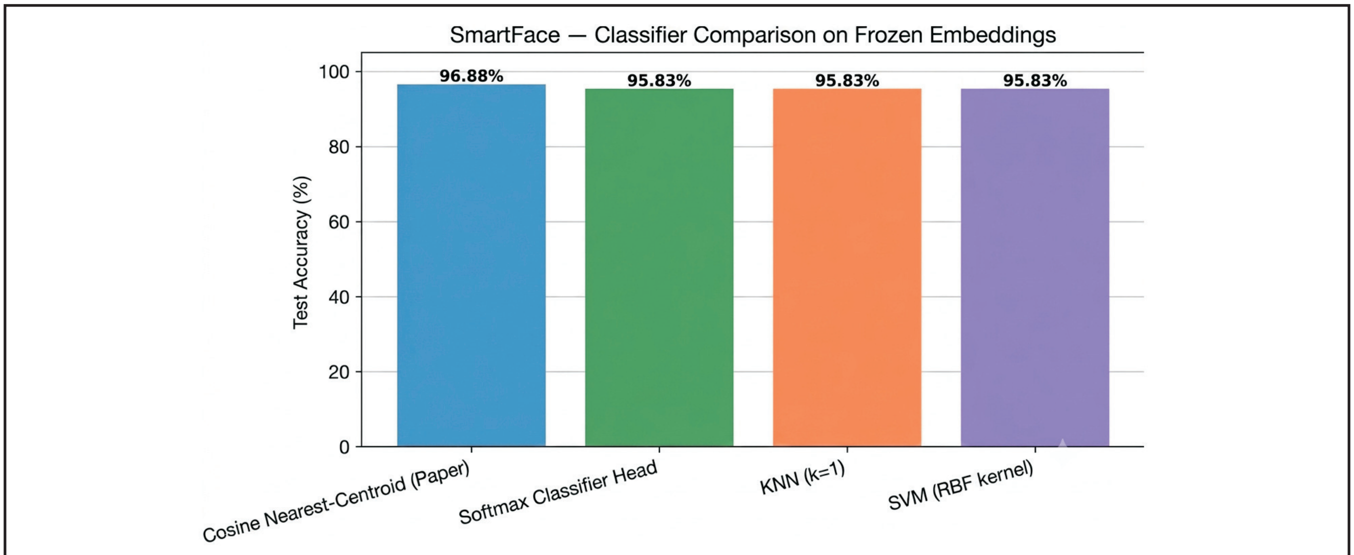

Fig. 3: *Classifier Comparison on Frozen MobileFaceNet Embeddings: Test Accuracy of Cosine Nearest-Centroid, Softmax, KNN, and SVM Methods*

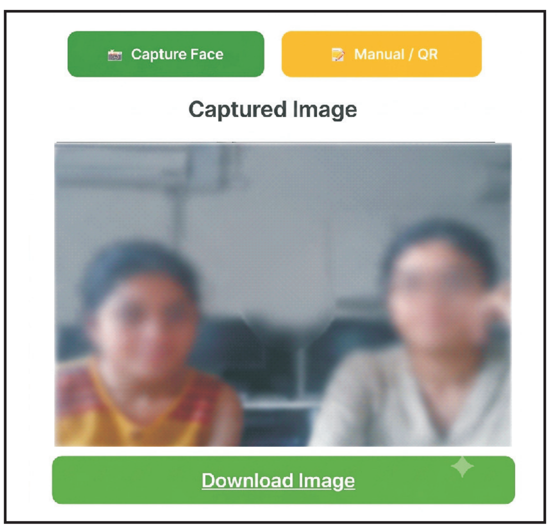

Fig. 4: *Attendance Capture Interface*

5.2 Comparison with Existing Approaches

Table 6 compares SmartFace against traditional manual attendance method and typical cloud-based face recognition systems across five parameters significant for rural school deployment.

As Table 6 demonstrates, SmartFace combines the offline capability of traditional methods with the minimal manual effort characteristic of cloud-based systems, while reducing the hardware and infrastructure costs associated with server-dependent deployments. Such a combination of factors would be especially useful in the setting of rural schools, wherein both dedicated hardware investment and ubiquitous Internet access are neither practical nor economically sustainable.

5.3 Features and User Interfaces of the SmartFace Application

SmartFace aims to support rural schools wherein the usage of digital administrative tools has been low so far. Teachers will be able to use SmartFace on existing hardware capable of taking photographs, so there will be no additional costs and the use of the application will be unaffected by unreliable rural network. Figures 5–7 illustrate the main user interfaces of the SmartFace application.

In addition to rural schools, the architecture is applicable for colleges and vocational institutes where, in order to meet

Table 6: *Comparison with Existing Attendance Methods*

Feature	Traditional Attendance	Cloud-Based Recognition	SmartFace
Manual Effort	High	Low	Low
Internet Required	No	Yes (continuous)	No (sync when available)
Hardware Cost	None	Moderate to High	Near Zero
Offline Capability	Full	None	Full (offline-first)
Deployment Simplicity	High (no setup)	Low (server config)	High (browser link)

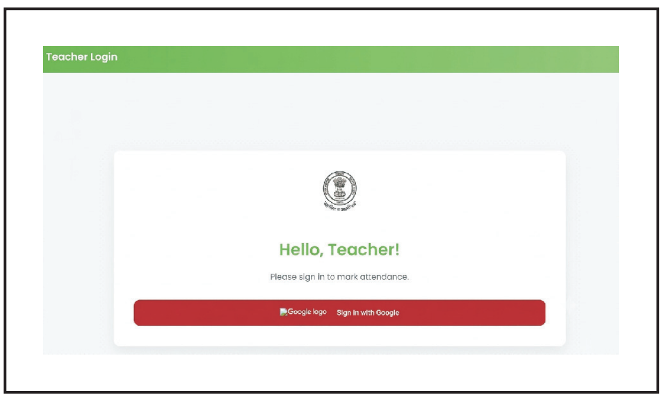

Fig. 5: *Teacher Login Dashboard*

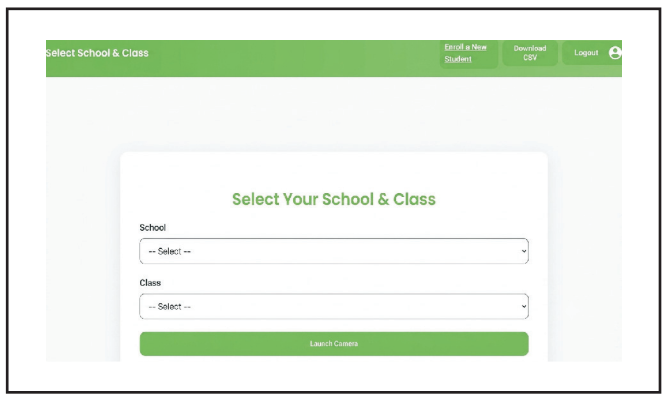

Fig. 6: *School and Class Selection Interface*

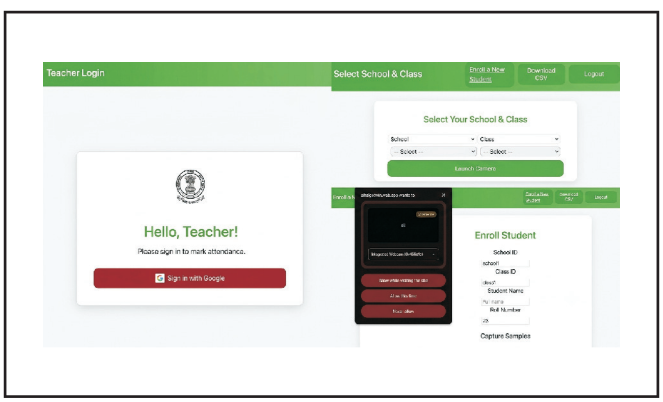

Fig. 7: *Student Enrolment Interface*

scholarship/certification criteria, verifiable records of attendance need to be maintained. Also, at exam centers, time stamped photo identity authentication can be performed to supplement manual checking. Government schemes that need to track participant attendance-e.g., skill development and community health, can benefit from the SmartFace system.

5.4 Limitations and Threats to Validity

Despite achieving a training accuracy of 99.43% and a test accuracy of 96.88%, several limitations warrant acknowledgement. First, the evaluation was conducted on a relatively small dataset (Garg et al., 2023) comprising 320 images from 16 subjects, which may not fully represent real-world classroom diversity. Second, the experiments were performed under semi-controlled conditions. Challenging scenarios such as severe illumination changes, motion blur, crowded classrooms, or significant facial occlusions were not evaluated. Prior research has established that broader validation on larger and more diverse datasets is necessary to assess recognition robustness conclusively (Chitrapu et al., 2025). Finally, the system has not yet undergone long-term deployment in rural schools, and practical factors such as user adoption, operational reliability, and biometric privacy considerations require further investigation prior to large-scale implementation (Sufian et al., 2023).

6 Conclusion and Future Work

Conclusion: This paper presents the design, development and preliminary evaluation of SmartFace, an AI-based on-device face recognition system to manage school attendance in rural Indian schools. The system integrates face recognition using MobileFaceNet (Chen et al., 2018) with a locally stored offline-first data architecture to facilitate recording attendance data without requiring a constant internet connection or dedicated server hardware. The fine-tuned model achieved a training accuracy of 99.43%, a validation accuracy of 93.75%, and a test accuracy of 96.88% on a dataset of 320 images of 16 students.

Future Work: Several directions are planned for future development:

1. A government-facing reporting dashboard for compliance and resource allocation;
2. A teacher analytics dashboard providing attendance trends to identify at-risk students;
3. Anti-spoofing detection using liveness detection techniques such as blink detection or depth estimation;
4. Significant dataset expansion through partnerships with rural schools across multiple states, enabling rigorous real-world evaluation;
5. Multi-classroom deployment coordination; and
6. A parent notification system for attendance alerts.

Acknowledgements

The authors would like to thank the students who contributed facial image data for the Indian Faces Image Classification Dataset (Garg et al., 2023), and the Department of Computer Science at Shaheed Rajguru College of Applied Sciences for Women, Delhi University, for their institutional support.

Conflict of Interest Statement

The authors declare no conflict of interest.

Ethical Declarations

The authors declare that this research does not involve human participants, animals, or any procedures requiring ethical approval. All data used in this study were obtained and processed in accordance with applicable ethical and legal guidelines.

References

1. Bharti, U., Bajaj, D., Tulika, Budhiraja, P., Juyal, M., & Baral, S. (2020). Android based e-voting mobile app using Google Firebase as BaaS. *Lecture Notes on Data Engineering and Communications Technologies*, *39*, 231–241. https://doi.org/10.1007/978-3-030-34515-0_24
2. Chen, S., Liu, Y., Gao, X., & Han, Z. (2018). MobileFaceNets: Efficient CNNs for accurate real-time face verification on mobile devices. *Biometric Recognition: 13th Chinese Conference, CCBR 2018*, *10996*, 428–438. https://doi.org/10.1007/978-3-319-97909-0_46
3. Chitrapu, P., Kumar Morampudi, M., & Kumar Kalluri, H. (2025). Robust face recognition using deep learning and ensemble classification. *IEEE Access*, *13*, 99957–99969. https://doi.org/10.1109/ACCESS.2025. 3575192
4. Deng, J., Guo, J., Yang, J., Xue, N., Kotsia, I., & Zafeiriou, S. (2022). ArcFace: Additive angular margin loss for deep face recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *44*(10), 5962–5979. https://doi.org/10.1109/TPAMI.2021.3087709

5. Deng, Z.-Y., Chiang, H.-H., Kang, L.-W., & Li, H.-C. (2023). A lightweight deep learning model for real-time face recognition. *IET Image Processing*, *17*(13), 3869–3883. https://doi.org/10.1049/ipr2.12903
6. Garg, S., et al. (2023). Indian faces image classification dataset [[Data set]]. https://www.kaggle.com/datasets/shubh48/indian-faces-image-classification
7. Hassanpour, A., & Kowsari, Y. (2023). Lightweight face recognition: An improved MobileFaceNet model. https://arxiv.org/abs/2311.15326
8. Mart´ ýnez-D´ ýaz, Y., Me´ndez-Va´zquez, H., Nicola´s-D´ ýaz, M., Luevano, L. S., Chang, L., & Gonzalez-Mendoza, M. (2019). ShuffleFaceNet: A lightweight face architecture for efficient and highly-accurate face recognition. *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*, 2721–2728. https://doi.org/10.1109/ICCVW.2019.00333
9. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., & Chen, L.-C. (2018). MobileNetV2: Inverted residuals and linear bottlenecks. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 4510–4520. https://doi.org/10.1109/CVPR.2018.00474
10. Schroff, F., Kalenichenko, D., & Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clus-tering. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 815–823. https://doi.org/10.1109/CVPR.2015.7298682
11. Sufian, A., Ghosh, A., Barman, D., Leo, M., Distante, C., & Li, B. (2023). FewFaceNet: A lightweight few-shot learning-based incremental face authentication for edge cameras. *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops (ICCVW)*. https://openaccess.thecvf.com/content/ICCV2023W/ACVR/papers/Sufian_FewFaceNet_A_Lightweight_Few-Shot_Learning-Based_Incremental_Face_Authentication_for_Edge_ICCVW_2023_paper.pdf
12. Sun, Z., Feng, C., Patras, I., & Tzimiropoulos, G. (2024). LAFS: Landmark-based facial self-supervised learn-ing for face recognition. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 1639–1649. https://openaccess.thecvf.com/content/CVPR2024/html/Sun_LAFS_Landmark-based_Facial_Self-supervised_Learning_for_Face_Recognition_CVPR_2024_paper.html
13. Zhang, K., Zhang, Z., Li, Z., & Qiao, Y. (2016). Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Processing Letters*, *23*(10), 1499–1503. https://doi.org/10.1109/ LSP.2016.2603342

❖ ❖ ❖