engineering proceedings

MDPI

Proceeding Paper

Facial Identification and Attendance Infrastructure for Adaptive Smart Classrooms: An Edge Prototype and Open-Set Benchmark †

Ramazan Ertuğrul Aydoğan [1,2], Fatima Sapundzhi [1,*], Berat Çiftçi [1,3], Öykü Boğazkesenli [1,4] and Slavi Georgiev [5,6,*]

[1] Department of Communication and Computer Engineering, Faculty of Engineering, South-West University "Neofit Rilski", 66 Ivan Mihailov Str., 2700 Blagoevgrad, Bulgaria; ra011054@mail2.gantep.edu.tr (R.E.A.); beratciftciofficial@gmail.com (B.Ç.); oyku.bogazkesenli@std.yildiz.edu.tr (Ö.B.)
[2] Department of Electrical and Electronics Engineering, Faculty of Engineering, Gaziantep University, 27310 Gaziantep, Türkiye
[3] Department of Statistics, Eskisehir Technical University, 26555 Eskisehir, Türkiye
[4] Department of Computer Education and Instruction Technologies, Yıldız Technical University, 34220 Istanbul, Türkiye
[5] Department of Applied Mathematics and Statistics, Faculty of Natural Sciences and Education, University of Ruse, 8 Studentska Str.,7004 Ruse, Bulgaria
[6] Institute of Mathematics and Informatics, Bulgarian Academy of Sciences, 8 Acad. Georgi Bonchev Str., 1113 Sofia, Bulgaria
* Correspondence: sapundzhi@swu.bg (F.S.); sggeorgiev@math.bas.bg (S.G.)
† Presented at the 2nd International Conference on Engineering, Technology, and Applied Science Innovations (ICETASI 2026), Zakynthos, Greece, 24–26 September 2026.

Abstract

This study presents a two-level evaluation of a smart-classroom attendance infrastructure that combines an embedded edge prototype with a separate open-set face-recognition benchmark. A Raspberry Pi 4 prototype was developed to integrate facial recognition, attendance registration, LCD feedback, and servo-based access simulation. Device-level testing comprised 20 access attempts under each of three controlled conditions. The prototype achieved end-to-end correct identification rates of 95.0% under standard indoor lighting and 85.0% under reduced lighting, with mean response times of 0.8 and 1.2 s, respectively. All 20 unknown-face attempts were correctly rejected, with a mean response time of 0.6 s. The dlib and SCRFD-ArcFace 1-NN pipelines were compared offline on a separate workstation using a custom open-set recognition protocol on the LFW-deepfunneled dataset. At the fixed, non-matched operating thresholds used in this study, SCRFD-ArcFace achieved 98.3% overall accuracy, a 96.7% true positive identification rate (TPIR), and a 0.4% false positive identification rate (FPIR), compared with 83.8%, 93.1%, and 23.3% for dlib, respectively. The main contribution is a clear distinction between the complete Raspberry Pi prototype tests and the algorithmic benchmarks conducted on a workstation. The prototype provides a basis for further development and validation of secure and automated educational environments, including systems intended to support learners with special educational needs.

Keywords: facial recognition; attendance registration; Raspberry Pi; edge computing; open-set identification; data protection; smart classroom; Internet of Things

Check for updates

Academic Editor: Konstantinos Poirazidis

Published: 6 November 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

In recent years, there has been a rapid development of artificial intelligence (AI), embedded systems, and sensor technologies, which improve the possibilities for building intelligent educational environments. Intelligent classrooms are seen as interactive, data-driven spaces in which sensors and intelligent software support both learning processes and classroom management [1,2]. As a key direction of innovation in education, AI emphasizes personalization, learning analytics, intelligent feedback, and context-aware support [3,4].

Attendance registration and classroom access control should not be viewed as mere administrative functions. In smart educational environments, these processes can provide verified attendance data, support reliable interaction workflows, and contribute to the reliable operation of digital learning systems. IoT-based attendance architectures have already demonstrated that automated attendance registration can be integrated with low-cost hardware, local or networked processing, and real-time education management scenarios [5,6].

Automated attendance has been investigated through IoT-based frameworks [5], classroom face detection [6], face recognition under challenging conditions [7], automatic attendance systems [8], embedded Raspberry Pi-based identification [9], biometric technology acceptance models [10], and BLE-based indoor positioning [11]. Face recognition with masks has also been explored [12]. At the algorithmic level, ArcFace [13], SCRFD [14], LFW [15,16], and dlib (19.24.6) [17] provide directly relevant foundations. Device-level and end-to-end validation and large-scale open-set recognition benchmarking are often reported without a clear separation between their hardware, datasets, and performance claims.

Facial recognition and biometrics in education can, however, raise important questions. The use of visual identification systems can improve contactless and efficient attendance management, but such systems must also be considered in relation to privacy, governance, secure data processing, and applicable regulatory requirements [18,19]. Therefore, engineering prototypes should be interpreted cautiously: not as fully deployable institutional solutions but as research demonstrations that explore limited technical functionality under controlled conditions.

The main contributions of the present study are: (1) the development of a Raspberry Pi 4 prototype integrating facial identification, attendance registration, LCD feedback, and servo-based access simulation; (2) an end-to-end device-level protocol with registered-user and unknown-face outcomes; (3) an offline comparison of dlib 1-NN and SCRFD-ArcFace 1-NN using a filtered open-set protocol on LFW; and (4) an analysis of technical, privacy, security, and deployment limitations of the prototype. The results from the workstation and the Raspberry Pi are reported separately.

2. Materials and Methods

The study comprised two evaluation environments. The first was a proof-of-concept Raspberry Pi 4 (Raspberry Pi Ltd., Cambridge, UK) system used to evaluate end-to-end integration of face identification, attendance registration, LCD feedback, and servo-based access-control simulation under controlled conditions. The second was a separate workstation used exclusively for offline benchmarking of dlib and SCRFD-ArcFace 1-NN on the LFW-deepfunneled dataset. The Raspberry Pi tests evaluate the end-to-end prototype, while the workstation tests focus on recognition accuracy and processing performance.

2.1. Embedded Raspberry Pi Architecture

The prototype was designed to perform four main functions: contactless identification of registered users, automated registration of attendance in a local database, immediate status feedback on an LCD display, and servo-based simulation of an access mechanism.

https://doi.org/10.3390/engproc2026167010

The system was designed as a controlled laboratory prototype, not as production-ready institutional access control.

Based on contemporary engineering directions in smart education, embedded computing devices are used in real-time classroom analytics and educational interaction [1,5,9]. This also reflects the practical need to balance accessibility, functionality, and data handling when developing experimental educational prototypes [1,3].

The hardware configuration was based on a Raspberry Pi 4 Model B with 4 GB RAM, which acts as the main processor for image capture, recognition, attendance registration, and peripheral control. Raspberry Pi 4 handles real-time video processing and controls hardware peripherals. The system utilizes the Pi's GPIO (General Purpose Input/Output) pins to interface with external components. Specifically, GPIO Pin 17 is configured for Pulse Width Modulation (PWM) to control a servo motor, simulating the door locking mechanism.

The hardware organization of the proposed prototype is presented in Figure 1. The diagram summarizes the main embedded components and their interconnections, including the Raspberry Pi 4 processing unit, Camera Module 3, I2C LCD display, servo motor, local face-encoding store, attendance log, and power supply. The camera module captured video frames at a resolution of 640 × 480 pixels. A 16 × 2 LCD module connected to the I2C interface at address 0x27 displayed status information during recognition.

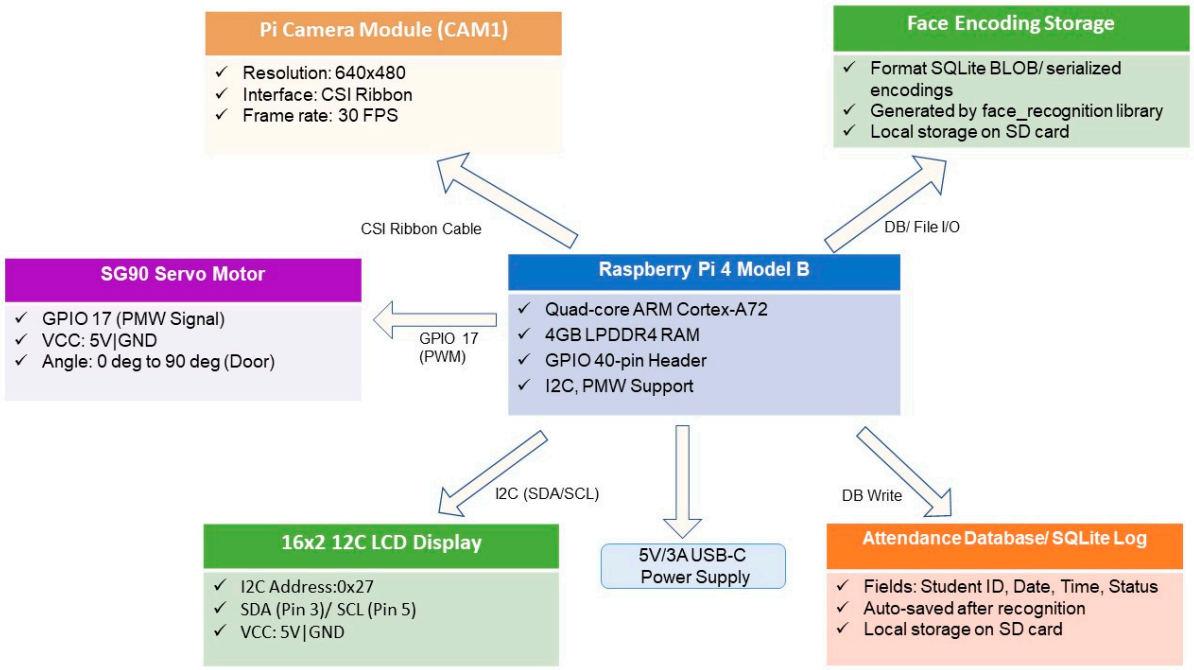


Figure 1. Hardware diagram of the smart-classroom entry and attendance prototype.

As shown in Figure 1, the Raspberry Pi 4 serves as the central processing and control unit of the prototype. The camera module provides visual input for face detection and recognition, while the LCD gives immediate user feedback during the identification and attendance process. The servo motor, controlled through GPIO 17 using PWM, simulates the physical access mechanism. Face encodings and attendance records are stored locally, which makes the prototype suitable for controlled laboratory validation of classroom identification workflows. Camera Module 3 provides 640 × 480-pixel frames for real-time identification, and the 16 × 2 LCD at I2C address 0 × 27 displays scanning, access-granted, access-denied, and attendance-status messages. After successful authentication, the servo motor is actuated only to simulate physical access.

2.2. Software Environment and Recognition Pipelines

The embedded software was implemented in Python 3.12 (Python Software Foundation, Beaverton, OR, USA) on Raspberry Pi OS 12 (Bookworm) (Raspberry Pi Ltd., Cambridge, UK). The workstation SCRFD–ArcFace pipeline was implemented using InsightFace (0.7.3), maintained by Jia Guo and Jiankang Deng, with ONNX Runtime 1.22.0 (Microsoft Corporation, Redmond, WA, USA) and the DirectML execution provider.. The software managed the local database and controlled the LCD display and servo motor through the Raspberry Pi GPIO interface.

The embedded pipeline used dlib's histogram-of-oriented-gradients (HOG) face detector and generated a 128-dimensional descriptor for each detected face [17]. The device-level gallery contained eight registered adult participants with ten facial images acquired per person at approximately 1 m from the Raspberry Pi camera. The resulting descriptors were stored separately, giving ten reference templates per identity. For each probe, Euclidean distances to all stored templates were calculated, and the identity associated with the minimum distance was selected. The match was accepted when the distance $d_{min} \leq \tau_{emb}$, with $\tau_{emb} = 0.55$, otherwise the detected face was classified as unknown. The same procedure was used in all three device-level conditions.

A separate workstation pipeline combined SCRFD-based face detection and five-point facial landmark localization with an ArcFace recognition model using a ResNet-50 backbone [13,14]. ArcFace generated 512-dimensional L2-normalized embeddings, which were compared with gallery templates using cosine similarity. For the workstation comparison, both pipelines used a 1-NN decision rule.In the SCRFD-ArcFace pipeline, the gallery identity with the highest cosine similarity was accepted when the maximum similarity was at least 0.42; otherwise, the probe was classified as unknown. This is a pipeline-level comparison: the detector, alignment procedure, embedding model, and embedding dimensionality change simultaneously, so the observed difference should not be attributed to the embedding alone.

The offline benchmark was conducted entirely on a Windows 11 (Microsoft Corporation, Redmond, WA, USA) workstation equipped with an AMD Ryzen AI 7 350 processor (Advanced Micro Devices, Inc., Santa Clara, CA, USA), 32 GB of RAM, and an NVIDIA (NVIDIA Corporation, Santa Clara, CA, USA) GeForce RTX 5060 Laptop GPU with 8 GB of video memory. The dlib baseline ran on the workstation CPU, whereas SCRFD-ArcFace used ONNX Runtime 1.22.0 with the DirectML execution provider for GPU acceleration. SCRFD-ArcFace was not executed on the Raspberry Pi.

2.3. Local Database and Operational Workflow

The local database was implemented using SQLite 3.40.1 (SQLite Development Team, Hipp, Wyrick & Company, Inc., Charlotte, NC, USA) and contained a registration table and an attendance log. The registration table stored the user identifier, name, class/group, and reference face descriptors, while the attendance log stored the identifier, date, time, and attendance status generated during operation. The software also controlled LCD feedback and servo-motor actuation through the GPIO interface. Table 1 presents the structure of the registration database used for learner identification and attendance handling.

The graphical interface used in this study was run on Windows 11, but it is also compatible with Raspberry Pi OS. It supports user management, live recognition, and access-log review, while the physical prototype and its LCD and servo controls were tested on the Raspberry Pi.

The application initializes the components, loads the registered descriptors, and compares incoming probe descriptors with the stored gallery. When a match is successful, the system checks attendance status, records attendance when applicable, updates the LCD,

and activates the servo motor. When the match fails, the system displays an "Access denied" message and returns to scanning.

Table 1. Structure of the student registration database used in the system.

Field Name	Data Type	Description
Student_ID	Integer	Unique identifier of the learner
First_Name	Text	First name of the registered user
Last_Name	Text	Last name of the registered user
Class_Group	Text	Assigned classroom/group
Face_Encoding	BLOB	Stored facial feature representation

The prototype stores facial descriptors and attendance records locally. Facial descriptors processed for the purpose of uniquely identifying a person constitute biometric data and are special-category personal data under Article 9 GDPR [18]. Local storage alone is not a complete data-protection measure; institutional deployment would require an appropriate lawful basis, purpose limitation, data minimization, role-based access control, encryption at rest and in transit, retention and deletion policies, audit logging, and incident procedures. Depending on the intended purpose and deployment context, an institutional system of this type may also fall within high-risk categories under Regulation (EU) 2024/1689 (the AI Act), including certain remote biometric identification systems and certain AI systems used to determine access or admission to educational institutions [19].

For the present laboratory study, participation was voluntary, and oral informed consent was obtained before data collection. In the GDPR, consent is a legal basis for processing personal and biometric data under Article 6(1)(a) and explicit consent under Article 9(2)(a), respectively. The consent was obtained orally and was not documented using a written consent form.

The graphical interface includes a "Delete person (GDPR erase)" control that removes the selected user's profile, associated metadata, and stored facial embeddings from the local database. Raw facial images are not stored because they are processed in memory during enrollment and then discarded. The deletion function does not remove previously recorded entries from the access log. The present implementation should be interpreted as a laboratory prototype.

For an unknown-person attempt, the system was expected to deny access, display an "Access denied" message, leave the servo motor inactive, and create no attendance record. An attempt was counted as successful only when all of these conditions were met. The prototype was tested under standard indoor lighting, reduced indoor lighting, and unknown-person conditions.

Three measures were used to evaluate the embedded prototype: successful access attempts by registered users, correct rejections of unknown users, and end-to-end response time.

The device-level evaluation involved 20 adults: eight registered participants and 12 unregistered participants. The registered group comprised three adult members of the research team and five adult volunteers; one research-team member participated as an intentionally unregistered user. Twenty access attempts were performed under each of three conditions, giving 60 attempts in total. The 20 trials of a stranger were drawn from the 12 unregistered adults. Some participants made more than one attempt, and the attempts were not independent at the participant level.

A registered-user attempt was successful if the system (1) correctly identified the user, (2) created the corresponding attendance record, (3) displayed the access-granted message on the LCD screen, and (4) activated the servo-based access simulation. An

unknown-person attempt was considered successful only if the system rejected the identity, created no attendance record, displayed an access-denied message, and left the servo motor inactive. The three conditions were standard indoor lighting, reduced indoor lighting, and unknown-person attempts. In all three conditions, the dlib threshold was fixed at 0.55. Figure 2 summarizes the operational workflow for recognition, attendance registration, LCD feedback, and servo-based access simulation.

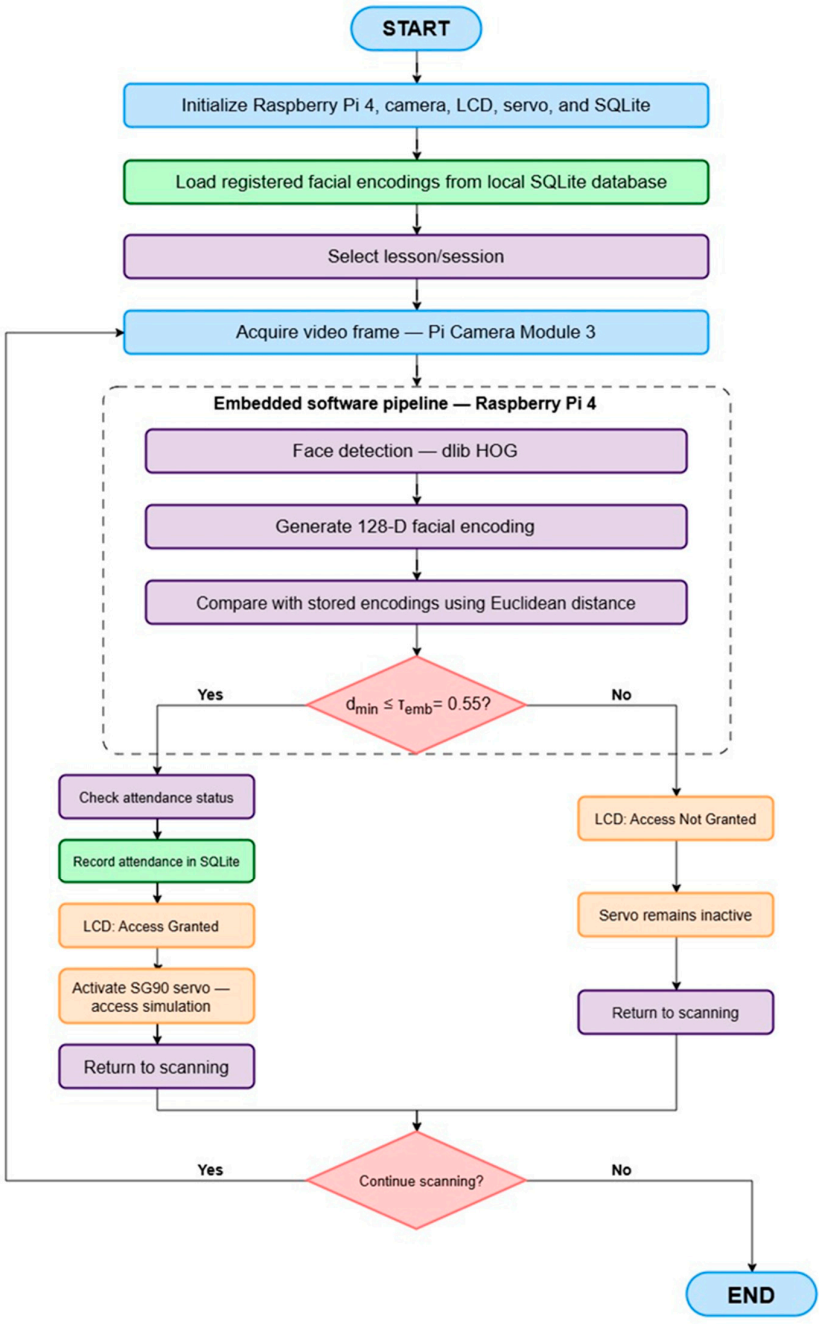


Figure 2. Operational workflow of face detection, identity verification, attendance registration, LCD feedback, and servo-based access simulation.

2.4. Offline LFW Benchmark Protocol

The offline comparison used the LFW-deepfunneled dataset, which contains 13,233 images of 5749 individuals [15,16]. Only identities represented by at least five admissible images were eligible for enrollment. Each selected image had to contain exactly one detected face,

a detected face width of at least 60 pixels, and a variance-of-Laplacian sharpness score of at least 60. Identities for which the available images did not meet the detection or quality criteria were excluded.

For each enrolled identity, the first image in the directory order that produced a valid face encoding was selected for the gallery. The embedding generated from this image was used as the identity template, while all remaining valid images of the same person were assigned to the known-probe set.

The resulting gallery contained 168 templates—one per enrolled identity—and the known-probe set contained 3032 images. Embeddings from multiple images were not averaged. The remaining 5581 identities were treated as unenrolled, resulting in 9223 unknown probe images. The same gallery/probe partition was used for dlib and SCRFD-ArcFace. The benchmark partition retained 12,423 images of the 13,233 source images, while 810 images were excluded during partition construction, face detection, or application of the quality criteria. Unlike the LFW gallery, which used one template per identity, the device-level gallery retained ten templates per identity. Because the galleries were constructed differently, the results of the two evaluations should not be compared directly.

For dlib, a face was accepted when the minimum Euclidean distance was at most 0.55, while for SCRFD-ArcFace a probe was accepted when the maximum cosine similarity was at least 0.42. The identity was assigned using the nearest match (1-NN); if the relevant threshold was not met, the probe was classified as unknown. These thresholds were fixed a priori and were not tuned on LFW to obtain a common FPIR. Consequently, the two pipelines are evaluated at different operating points. A stricter dlib threshold would be expected to reduce FPIR, generally at the cost of a lower TPIR; the reported accuracy gap should therefore not be interpreted as a pure measure of representation quality.

Runtime performance was evaluated using the same set of 150 frames at 640 × 480 resolution for both pipelines. Performance was reported as mean FPS, median latency, and 95th-percentile latency, and the nominal frame time was calculated as 1000/FPS.

For the open-set LFW evaluation, the true positive identification rate (TPIR) was defined as the proportion of enrolled probes that were correctly identified. The false positive identification rate (FPIR) was defined as the proportion of unenrolled probes that were incorrectly assigned to an enrolled identity. At the selected operating point, the false negative identification rate (FNIR) was reported as 1—TPIR, i.e., the proportion of enrolled probes that were not correctly identified.

Overall open-set accuracy was defined as the number of correctly identified known probes plus correctly rejected unknown probes divided by the total number of known and unknown probes. Because the known and unknown probe sets differ in size, TPIR, FNIR, and FPIR are also reported separately.

3. Results and Discussion

3.1. Embedded Prototype Evaluation

The proposed prototype integrated face-based learner identification, attendance registration, LCD feedback, and servo-based access simulation in the Raspberry Pi environment. The physical hardware prototype (Raspberry Pi 4 Model B, CSI camera, I2C LCD display, and SG90 servo motor) is shown in Figure 3. A 16 × 2 I2C LCD module and an SG90 microservomotor, supplied by SunFounder (Shenzhen, China), were used for status feedback and access-control simulation, respectively. The user interface was upgraded from a legacy console workflow to a graphical user interface (GUI) supporting user registration, authorization management, live recognition, access-log review, registry refresh, and deletion of registered users (Figure 4).

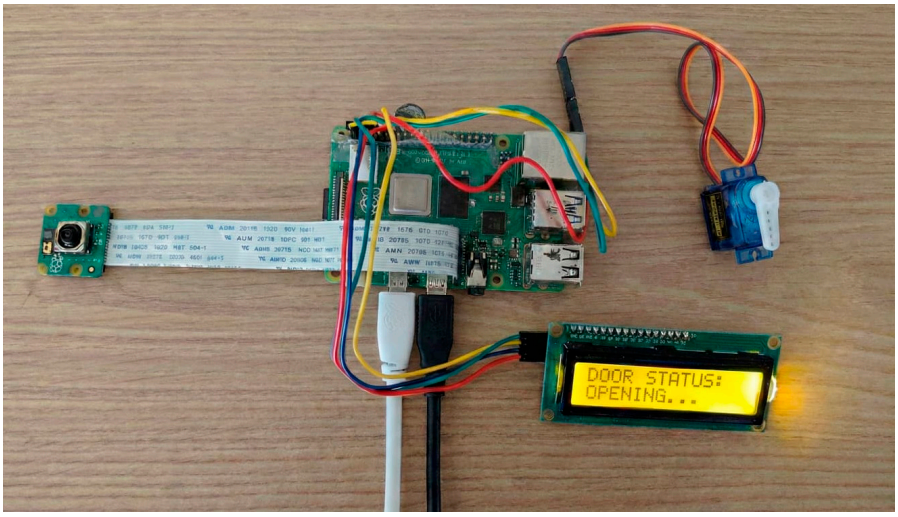


Figure 3. Implemented prototype with Raspberry Pi 4, camera, I2C LCD display, GPIO wiring, and servo motor for access simulation.

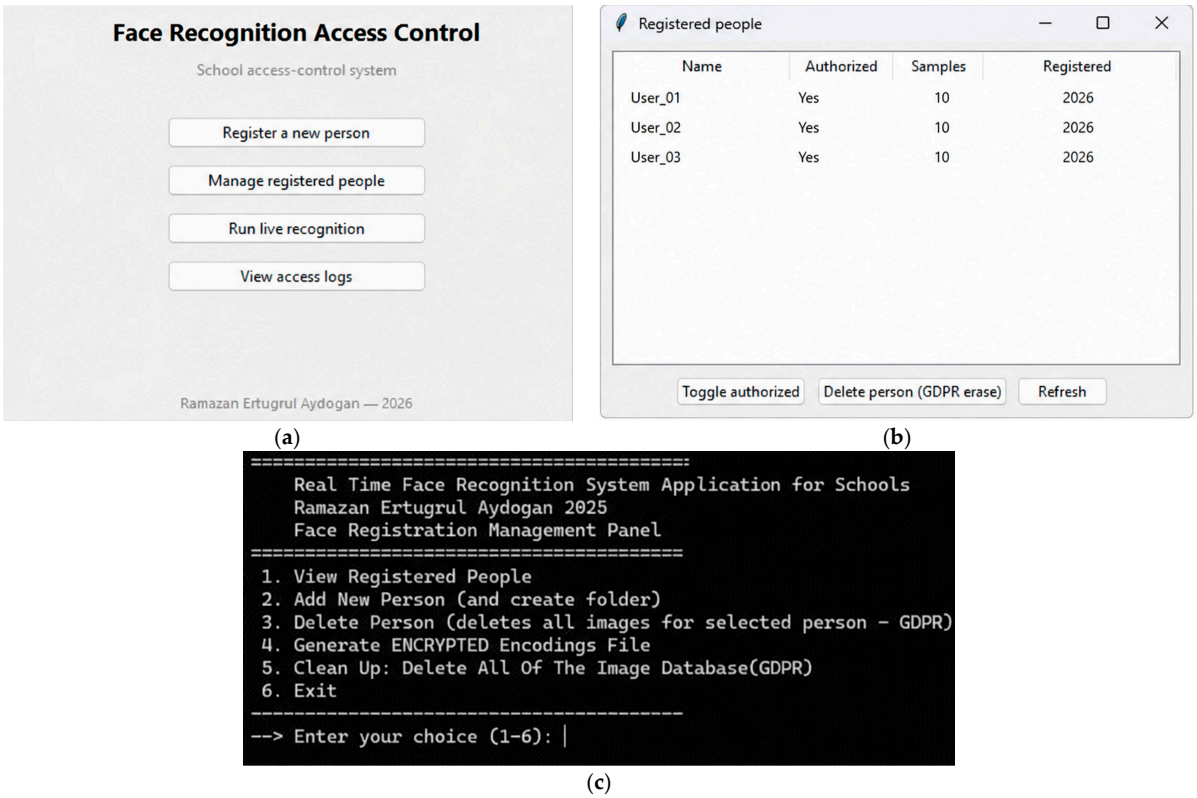


Figure 4. Graphical user interface for managing control (**a**), anonymized list of registered users and local data deletion controls (**b**), and legacy console interface (**c**).

Research on smart classrooms and smart attendance systems demonstrates that embedded and IoT infrastructures can automate classroom workflows when tasks are clearly defined and matched to the available hardware [1,4].

Under standard indoor lighting, 19 of 20 registered-user attempts met all success criteria (95.0%), with a mean response time of 0.8 s. Under reduced indoor lighting, 17 of 20 attempts met all success criteria (85.0%), with a mean response time of 1.2 s. The results are shown in Table 2.

Table 2. End-to-end device outcomes under the three laboratory conditions.

Test Condition	Attempts	Outcome	Rate (%)	Mean Time (s)
Standard lighting	20	19 registered successes	95.0	0.8
Reduced lighting	20	17 registered successes	85.0	1.2
Unknown face	20	20 correct rejections	100.0	0.6

The device-level results are consistent with the practical emphasis of prior embedded and attendance-system studies, but direct numerical comparison requires caution because the evaluation protocols differ. Pecolt et al. [9] reported an average recognition time of 10.5 s for a Raspberry Pi-based recognition system tested across different lighting conditions, whereas the present prototype produced mean end-to-end times of 0.8–1.2 s for registered-user trials; the algorithms, participant groups, thresholds, and success criteria are not equivalent, so this difference should not be interpreted as a controlled speedup. Nguyen et al. [5] and Lin and Huang [8] use distributed or server-assisted attendance architectures, while the present device-level evaluation reports the local workflow through recognition, database update, user feedback, and simulated actuation.

Reduced indoor lighting was associated with a 10-percentage-point decrease in successful registered-user attempts and 0.4 s increase in mean response time relative to standard indoor lighting (Figure 5).

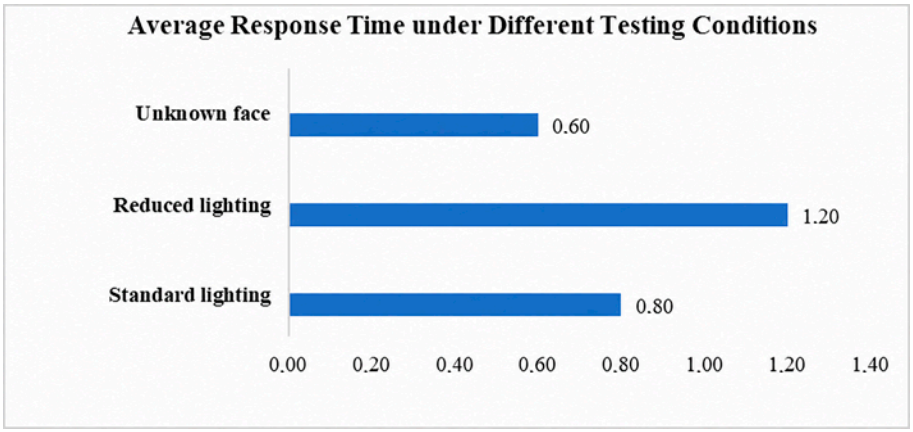


Figure 5. Average response time of the proposed prototype under different testing conditions.

3.2. Offline Open-Set Benchmark on the Workstation

The LFW-deepfunneled benchmark used 168 enrolled identities, 168 gallery templates, 3032 known probes, 5581 unenrolled identities, and 9223 unknown probes. The dlib and SCRFD-ArcFace 1-NN pipelines were evaluated on exactly the same partition. Runtime measurements used the same set of 150 frames on the workstation; dlib ran on the CPU, whereas SCRFD-ArcFace used DirectML GPU acceleration.

As shown in Table 3, at the fixed thresholds used in this study, SCRFD-ArcFace produced a higher TPIR (96.71% vs. 93.11%) and a substantially lower FPIR (0.42% vs. 23.34%) than dlib. Because these thresholds were not calibrated to a matched FPIR, this is a comparison of the two complete pipelines at different operating points rather than an embedding-only comparison. Mean throughput was 3.3 FPS for dlib on the CPU and 64.5 FPS for SCRFD-ArcFace on the GPU, corresponding to nominal frame times of 303.0 ms/frame and 15.5 ms/frame, respectively, a 19.5-fold difference. Median processing times were 343.3 ms/frame for dlib and 13.6 ms/frame for SCRFD-ArcFace, while the 95th-percentile

times were 376.3 and 24.0 ms/frame, respectively. All measurements were obtained on the workstation.

Table 3. Offline 1-NN recognition results on the LFW-deepfunneled dataset at the fixed operating thresholds and workstation runtime results obtained from 150 frames.

Metric	Dlib 1-NN	SCRFD-ArcFace 1-NN
Overall accuracy	83.82%	98.34%
TPIR (correct identification)	93.11%	96.71%
FPIR (false positive identification)	23.34%	0.42%
FNIR (1—TPIR)	6.89%	3.29%
Mean FPS	3.3	64.5

3.3. Limitations and Deployment Considerations

The prototype was tested with only eight registered adults and 60 attempts in a controlled laboratory; children and school users were not included. The lighting conditions were described qualitatively, but illuminance was not measured in lux, limiting reproducibility of the reduced-light condition; future assessments will use a calibrated lux meter. No formal institutional ethics review was obtained, which is a limitation of the study's governance reporting even though the evaluation was non-interventional and involved adults. The LFW comparison also uses fixed, non-matched decision thresholds and changes both the detector and embedding pipeline, while the LFW and device-level galleries use different numbers of templates per identity. These factors limit direct causal attribution and direct comparison across the two evaluation levels. Broader participant samples, liveness and spoofing tests, calibrated illumination, threshold calibration, and security and data-protection validation are required before institutional deployment. Related integrated system architectures are discussed in [20,21].

4. Conclusions

This study presented a Raspberry Pi 4-based prototype for facial identification and local attendance registration. The system provides feedback on an LCD and uses a servo motor to simulate access control. In the laboratory tests, registered users were successfully recognized in 95.0% of attempts under standard indoor lighting and 85.0% under reduced lighting. All 20 unknown-face attempts were correctly rejected.

A separate workstation benchmark compared the complete dlib and SCRFD-ArcFace 1-NN pipelines using the same custom open-set LFW partition. At the fixed, non-matched operating thresholds, SCRFD-ArcFace showed a higher TPIR, lower FPIR, and faster processing than dlib. These results characterize the workstation pipelines at the selected operating points and should not be interpreted as an embedding-only ablation.

The prototype demonstrates how facial identification, attendance recording, and simulated access control can be integrated in an educational laboratory environment. The results provide a basis for further development and validation of adaptive and automated learning-environment infrastructure. Institutional implementation would require broader validation, liveness and security testing, calibrated operating thresholds, formal governance review, and appropriate data-protection controls.

Author Contributions: Conceptualization, F.S. and S.G.; methodology, F.S.; software, B.Ç., Ö.B. and R.E.A.; validation, F.S., B.Ç., Ö.B., R.E.A. and S.G.; formal analysis, F.S.; investigation, B.Ç., Ö.B. and R.E.A.; resources, S.G.; data curation, F.S., B.Ç., Ö.B. and R.E.A.; writing—original draft preparation, F.S. and S.G.; writing—review and editing, F.S., B.Ç., Ö.B., R.E.A. and S.G.; visualization, F.S.; supervision, F.S.; project administration, F.S. and S.G.; funding acquisition, S.G. All authors have read and agreed to the published version of the manuscript.

Funding: This study was supported by Project 2026-FNSE—02, financed by the "Scientific Research" Fund of Ruse University and Project КП-06-H100//7/10.12.2025 "Intelligent Adaptive Approaches to Supporting Primary Education for SSEN through Machine Learning and Robotic Technologies", Bulgarian National Science Fund.

Institutional Review Board Statement: No formal institutional ethics approval was obtained, as the study was conducted as a non-interventional engineering evaluation involving only adult members of the research team and adult volunteers. No minors, patients, or vulnerable participants were involved.

Informed Consent Statement: Oral informed consent was obtained from all participants prior to data collection. Participants were informed about the purpose of the study, the processing and local storage of their personal data, the voluntary nature of participation, and their right to refuse. The publication does not include identifiable images of the participants' faces.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The LFW dataset is publicly available from its source.

Acknowledgments: This study was carried out within the framework of Project КП-06-H100//7/10.12. 2025 "Intelligent Adaptive Approaches to Supporting Primary Education for SSEN through Machine Learning and Robotic Technologies", Bulgarian National Science Fund.

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. Zhang, X.; Ding, Y.; Huang, X.; Li, W.; Long, L.; Ding, S. Smart Classrooms: How Sensors and AI Are Shaping Educational Paradigms. *Sensors* **2024**, *24*, 5487. [CrossRef]
2. Dimitriadou, E.; Lanitis, A. A Critical Evaluation, Challenges, and Future Perspectives of Using Artificial Intelligence and Emerging Technologies in Smart Classrooms. *Smart Learn. Environ.* **2023**, *10*, 12. [CrossRef]
3. Simon, P.; Zeng, L. Behind the Scenes of Adaptive Learning: A Scoping Review of Teachers' Perspectives on the Use of Adaptive Learning Technologies. *Educ. Sci.* **2024**, *14*, 1413. [CrossRef]
4. Mustafa, M.Y.; Tlili, A.; Lampropoulos, G.; Huang, R.; Jandrić, P.; Zhao, J.; Salha, S.; Xu, L.; Panda, S.; Kinshuk; et al. A Systematic Review of Literature Reviews on Artificial Intelligence in Education (AIED): A Roadmap to a Future Research Agenda. *Smart Learn. Environ.* **2024**, *11*, 59. [CrossRef]
5. Nguyen, V.D.; Khoa, H.V.; Kieu, T.N.; Huh, E.-N. Internet of Things-Based Intelligent Attendance System: Framework, Practice Implementation, and Application. *Electronics* **2022**, *11*, 3151. [CrossRef]
6. Sun, C.; Wen, P.; Zhang, S.; Wu, X.; Zhang, J.; Gong, H. A Face Detector with Adaptive Feature Fusion in Classroom Environment. *Electronics* **2023**, *12*, 1738. [CrossRef]
7. Zhalgas, A.; Amirgaliyev, B.; Sovet, A. Robust Face Recognition Under Challenging Conditions: A Comprehensive Review of Deep Learning Methods and Challenges. *Appl. Sci.* **2025**, *15*, 9390. [CrossRef]
8. Lin, C.-L.; Huang, Y.-H. The Application of Adaptive Tolerance and Serialized Facial Feature Extraction to Automatic Attendance Systems. *Electronics* **2022**, *11*, 2278. [CrossRef]
9. Pecolt, S.; Błażejewski, A.; Królikowski, T.; Maciejewski, I.; Gierula, K.; Głowiński, S. Personal Identification Using Embedded Raspberry Pi-Based Face Recognition Systems. *Appl. Sci.* **2025**, *15*, 887. [CrossRef]
10. Rukhiran, M.; Wong-In, S.; Netinant, P. User Acceptance Factors Related to Biometric Recognition Technologies of Examination Attendance in Higher Education: TAM Model. *Sustainability* **2023**, *15*, 3092. [CrossRef]
11. Puckdeevongs, A.; Tripathi, N.K.; Witayangkurn, A.; Saengudomlert, P. Classroom Attendance Systems Based on Bluetooth Low Energy Indoor Positioning Technology for Smart Campus. *Information* **2020**, *11*, 329. [CrossRef]
12. Tsai, T.-H.; Lu, J.-X.; Chou, X.-Y.; Wang, C.-Y. Joint Masked Face Recognition and Temperature Measurement System Using Convolutional Neural Networks. *Sensors* **2023**, *23*, 2901. [CrossRef]
13. Deng, J.; Guo, J.; Xue, N.; Zafeiriou, S. ArcFace: Additive Angular Margin Loss for Deep Face Recognition. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*; IEEE: Long Beach, CA, USA, 2019; pp. 4685–4694. [CrossRef]
14. Guo, J.; Deng, J.; Lattas, A.; Zafeiriou, S. Sample and Computation Redistribution for Efficient Face Detection. *arXiv* **2021**. [CrossRef]
15. Huang, G.B.; Ramesh, M.; Berg, T.; Learned-Miller, E. *Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments*; Technical Report 07–49; University of Massachusetts: Amherst, MA, USA, 2007.
16. Lee, H.; Learned-Miller, E. Learning to Align from Scratch. In *Advances in Neural Information Processing Systems 25*; Curran Associates, Inc.: Red Hook, NY, USA, 2012.

17. King, D.E. Dlib-ml: A Machine Learning Toolkit. *J. Mach. Learn. Res.* **2009**, *10*, 1755–1758.
18. European Parliament; Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation). 2016. Available online: https://eur-lex.europa.eu/legal-content/EN/ALL/?uri=celex:32016R0679 (accessed on 3 November 2026).
19. European Parliament; Council of the European Union. Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 Laying down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act). 2024. Available online: https://eur-lex.europa.eu/eli/reg/2024/1689/oj/eng (accessed on 3 November 2026).
20. Sapundzhi, F.; Aydoğan, R.E.; Georgiev, S.; Nikitov, N. Architecture and Performance Evaluation of Real-Time Facial Recognition for Access Control. *Eng. Proc.* **2026**, *150*, 11. [CrossRef]
21. Ganchev, N.; Sapundzhi, F.; Georgiev, S. Architecture-Driven Logic Design Strategies for Distributed Educational IoT Environments with Robotic Support. *WSEAS Trans. Comput. Res.* **2026**, *14*, 495–504. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.