And., Intl.,J., Appl Sci., Eng., Tech, Vol 06, No 03, November 2026
e-ISSN: 2797-0442

Real-Time Face Recognition for Smart Door Lock Systems: A Comparative Study of FaceNet, FaceNet512, and InsightFace on Raspberry Pi

Muhammad Iqbal Halid, Ivan Tanra, Marvin Yonathan Hadiyanto*
Department of Electrical Engineering, Faculty of Smart Technology, Krida Wacana Christian University, Jakarta, Indonesia
*Corresponding author, e-mail: marvin.yonathan@ukrida.ac.id

Abstract—Face recognition has emerged as an effective biometric authentication technology for intelligent access control systems because of its contactless operation, convenience, and security advantages. Although deep learning has substantially improved recognition performance, the suitability of different models for real-time deployment on embedded platforms remains insufficiently explored. This study presents a real-time, face recognition-based smart door lock system implemented on a Raspberry Pi 5 and compares three face recognition models: FaceNet, FaceNet512, and InsightFace. The system integrates a webcam, a relay module, and an electromagnetic door lock to provide automated access control for authorized users. Facial images were collected from 20 individuals, yielding 400 training images and 200 authorized-user testing images. Model performance was evaluated using accuracy, precision, sensitivity, specificity, F1-score, latency, frames per second (FPS), and recognition robustness at different distances. InsightFace achieved the highest performance, with 100% accuracy, precision, sensitivity, specificity, and F1-score on the testing dataset. During real-time testing, InsightFace maintained authorized-user recognition rates above 97% at all tested distances and achieved a 100% unauthorized-user rejection rate. It also recorded the lowest average latency (0.243 s) and the highest average processing speed (4.153 FPS), demonstrating its suitability for embedded smart door lock applications.

Keywords: Face Recognition, InsightFace, Smart Door Lock, Raspberry Pi, Biometric Security, Real-Time Recognition.

This article *is licensed under the CC–BY-SA license.*

1. Introduction

Advances in artificial intelligence and computer vision, particularly in face recognition, have significantly transformed biometric security systems [1]. Traditional authentication methods, including physical keys, passwords, and RFID cards, remain vulnerable to key duplication, password disclosure, card loss, and unauthorized access [2]. Face recognition provides several advantages as a biometric authentication method, including contactless operation, ease of use, and automatic identification in real time. By analyzing distinctive facial characteristics that are difficult to replicate or forge, face recognition offers a promising approach to improving security and user convenience [3]. Moreover, the proliferation of smart systems and Internet of Things (IoT) technologies has accelerated the adoption of face recognition in attendance systems, surveillance applications, and smart home automation [4]. Among these applications, smart door lock systems have attracted considerable attention because they enable automated access management while improving security and user experience [5].

Recent advances in deep learning, particularly convolutional neural network (CNN)-based approaches, have substantially improved the accuracy, robustness, and efficiency of face recognition systems under diverse environmental conditions [6], [7], [8]. Several established face recognition frameworks, including FaceNet, FaceNet512, and InsightFace, have demonstrated strong recognition

Received: October 29, 2026, Revised: November 13, 2026, Accepted: November 17, 2026
https://doi.org/10.25077/aijaset.v6i3.384

performance and have been adopted in research and industrial applications [9]. FaceNet employs a triplet-loss-based embedding strategy to learn discriminative facial representations by reducing the distance between embeddings of the same identity while increasing the distance between different identities [10], [11]. FaceNet512 is a variant of FaceNet that generates 512-dimensional facial embeddings, providing a higher-dimensional feature representation for face recognition tasks [9]. In contrast, the ArcFace-based recognition approach supported by InsightFace employs additive angular margin loss to improve inter-class separability and intra-class compactness, thereby producing highly discriminative facial embeddings [12]. Nevertheless, recognition performance in practical deployments remains sensitive to changes in illumination, facial pose, occlusion, image quality, camera distance, and the computational constraints of real-time processing [13], [14]. Evaluating modern face recognition models under realistic operating conditions is therefore essential for assessing their reliability in practical security applications.

Previous studies have demonstrated the feasibility of integrating face recognition into smart door lock and access control systems. Yedulapuram et al. developed an automatic door lock system using a Haar classifier on a Raspberry Pi platform. Their system recognized authorized users and recorded unauthorized access attempts [15]. Similarly, Saputra and Surantha implemented a Raspberry Pi-based home security system that used Haar classifier-based face recognition to identify elderly users, achieving 98% accuracy with a processing time of 1.472 s [16]. Syafeeza et al. proposed an IoT-enabled facial recognition access control system based on AlexNet, incorporating real-time monitoring and remote control through the Blynk platform [17]. Irjanto et al. also developed a Raspberry Pi-based smart door security system using AlexNet and reported approximately 97.5% accuracy in recognizing authorized homeowners [18]. Although these studies reported promising results, they primarily evaluated systems using a single face recognition model and provided limited analysis of the trade-offs between recognition performance and computational efficiency in embedded deployments.

Modern face recognition frameworks differ in their computational requirements, which can significantly affect their suitability for resource-constrained embedded platforms such as the Raspberry Pi. Previous studies have demonstrated that deploying deep learning-based face recognition systems on embedded devices presents challenges related to limited processing capacity, memory availability, and inference latency [19], [20], [21], [22]. Consequently, a model that achieves high accuracy during offline evaluation may not necessarily provide optimal performance in real-time applications when computational efficiency and hardware limitations are considered. Despite the widespread use of FaceNet, FaceNet512, and InsightFace, comparative evaluations of these models under identical deployment conditions remain limited. Accordingly, this study compares the performance of FaceNet, FaceNet512, and InsightFace in a real-time smart door lock application. The proposed system integrates a Raspberry Pi, webcam, relay module, and electromagnetic door lock. The evaluation considers classification metrics, including accuracy, precision, sensitivity, specificity, and F1-score, alongside deployment-oriented metrics such as latency, FPS, and recognition robustness at different camera distances. This study makes three main contributions. First, it systematically compares FaceNet, FaceNet512, and InsightFace using a facial image dataset collected from Indonesian students. Second, it evaluates their recognition effectiveness and computational efficiency when deployed on a Raspberry Pi-based smart door lock system. Third, it examines the effect of camera distance on recognition robustness, providing practical insights for the development of reliable face recognition-based access control systems.

2. Method

2.1 Smart Door Lock Systems

This section describes the hardware configuration of the proposed smart door lock system. A Raspberry Pi 5 with 8 GB of RAM, running the 64-bit Raspberry Pi OS, serves as the main controller. A PC02 2K webcam captures facial images for recognition. The locking mechanism consists of a 12 V electromagnetic lock rated at 600 lbs and installed above the laboratory entrance. GPIO18 on the Raspberry Pi is connected to a relay driver circuit comprising a 220 Ω resistor, a 1N4007 diode, and an

NPN transistor. The transistor drives an SRD-05VDC-SL-C 5 V relay module, which acts as an electronic switch for the electromagnetic door lock. An S-240-12 power supply provides 12 V DC for the lock, while an LM2596 step-down DC-DC converter reduces the 12 V supply to a stable 5 V for the relay driver circuit. The Raspberry Pi is powered separately by a 5 V, 3 A USB Type-C adapter. During operation, the webcam captures the face of a person standing in front of the door, and the Raspberry Pi processes the captured frame using the selected face recognition model. When an authorized user is recognized, the Raspberry Pi sends a HIGH signal through GPIO18 to activate the NPN transistor and energize the relay. The relay contacts then allow the electromagnetic lock to be released. When the detected person is classified as unauthorized, the relay remains inactive, and the door stays locked. Figure 1 illustrates the system's hardware configuration. The system was deployed at the entrance to the Electrical Engineering Laboratory at UKRIDA to control access by registered users. Figure 2 shows the installation location.

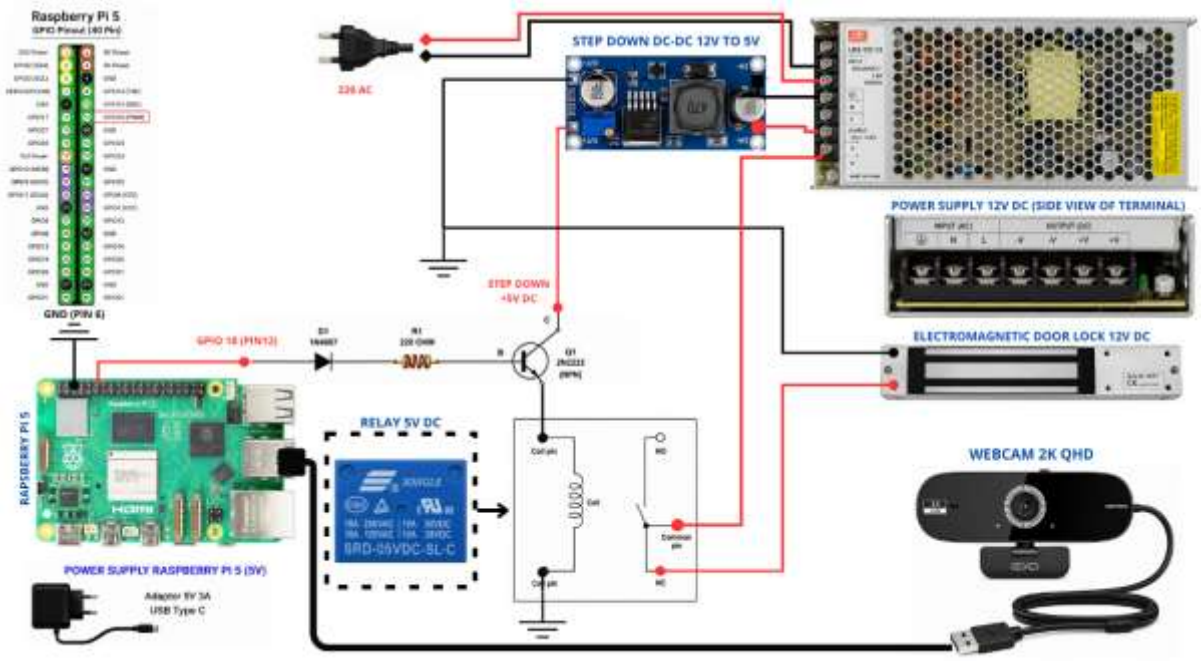


Figure 1. Hardware configuration of the proposed smart door lock system

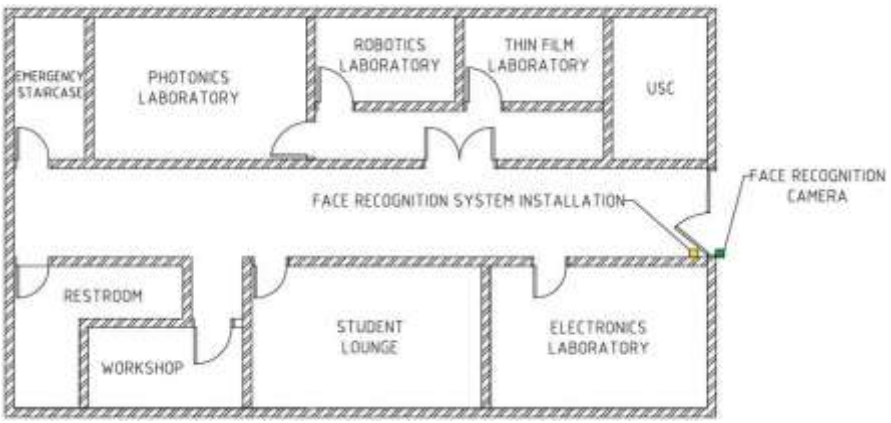


Figure 2. Installation location of the smart door lock system

AIJASET – Vol. 06, No. 03, November 2026. 387-396
https://doi.org/10.25077/aijaset.v6i3.384

2.2 Dataset Preparation and Face Recognition

Facial images were collected manually using the webcam. The training dataset included 20 students, with 20 images captured per individual, yielding 400 training images. For the same 20 students, an additional 10 images per individual were collected for testing, resulting in 200 authorized-user testing images. All participants were Indonesian citizens. Images were acquired under varying lighting conditions and with different facial poses and expressions to improve recognition robustness. The images were organized into separate folders according to participant identity, as illustrated in Figure 3. In addition, 50 facial images of unregistered individuals were collected to evaluate unauthorized-user rejection. The complete testing dataset therefore contained 250 images: 200 from authorized users and 50 from unauthorized users.



(a) (b)

Figure 3. (a) Dataset folder organization and (b) variations in facial images

To ensure a fair and consistent comparison, FaceNet, FaceNet512, and InsightFace were evaluated using the same facial image datasets and experimental procedures. All three models were implemented on the same Raspberry Pi 5 platform and evaluated using the same webcam and facial image acquisition conditions. The collected images underwent face detection and alignment before feature extraction, while model-specific input processing requirements were maintained according to each model's architecture. The same evaluation procedure and performance metrics were applied to all three models. Model-specific recognition thresholds were determined using the validation dataset and subsequently applied to the independent testing dataset. This standardized evaluation framework was intended to minimize the influence of differences in experimental conditions on the comparative results.

Before model training and recognition, the collected facial images were preprocessed to improve image quality and support reliable feature extraction. Each image underwent face detection and alignment to normalize facial position and orientation and reduce variations caused by rotation and head tilt. Three face recognition models, namely InsightFace, FaceNet, and FaceNet512, were evaluated. These models extract distinctive facial features and represent them as numerical embedding vectors, allowing facial identities to be compared in a feature space [10], [12], [21]. During recognition, the embedding vector generated from the input face was compared with the reference embeddings stored in the database using cosine similarity, which measures the similarity between two vectors based on their angular relationship [21], [22]. If the calculated similarity score exceeded a predefined threshold, the input face was assigned to the corresponding registered identity. Otherwise, the individual was classified as unknown.

2.3 System Implementation

The real-time face recognition-based smart door lock system was installed at the entrance to the Electrical Engineering Laboratory to control access by authorized users. As shown in Figure 4, system operation comprises two main stages: registration and real-time recognition. During registration, facial images of authorized users are collected and preprocessed, and the selected face recognition model generates facial embeddings. These embeddings are stored in the database as reference templates for subsequent identification. Only individuals who have been registered are eligible to access the laboratory. During real-time recognition, the camera continuously captures video frames and checks

for faces. When a face is detected, the system extracts its embedding and compares it with the reference embeddings stored in the database. If the individual is identified as an authorized user, the Raspberry Pi sends a control signal to the relay module, releasing the electromagnetic lock and granting access. The system also records the user's identity and access time in a cloud-based spreadsheet for attendance tracking and access logging.

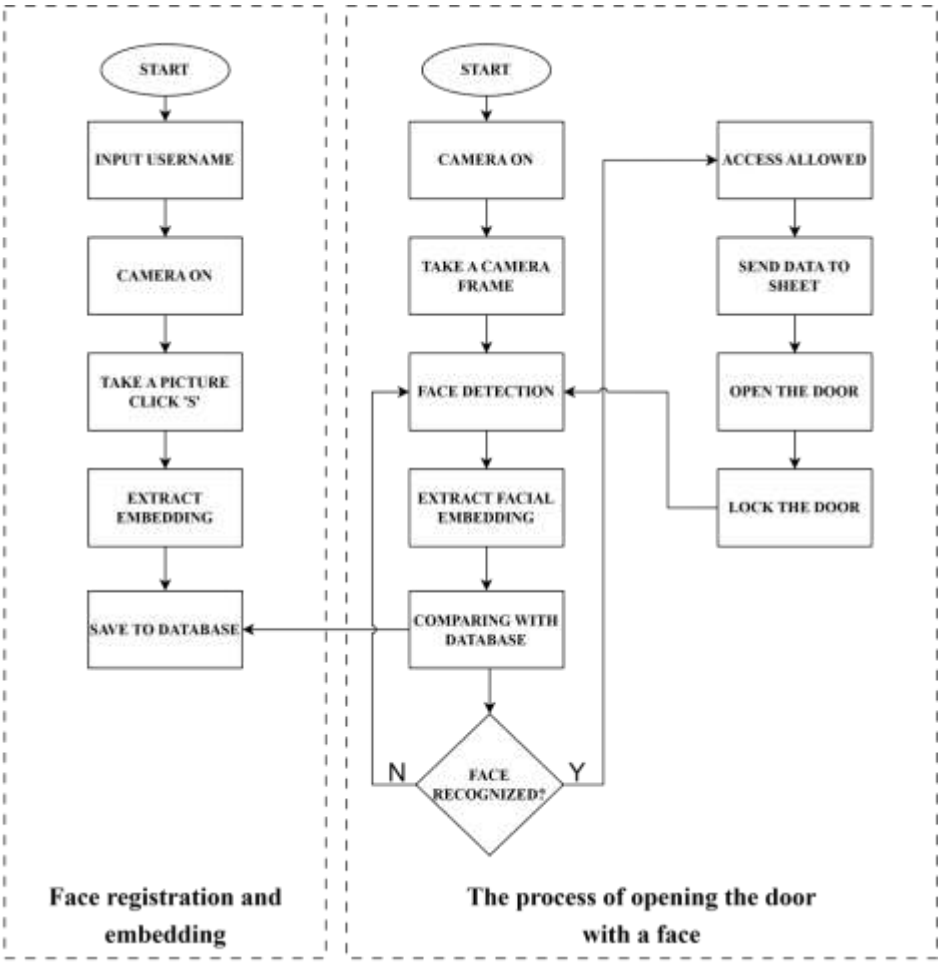


Figure 4. Workflow of the real-time face recognition-based smart door lock system

3. Results and Discussion

Threshold selection was performed using the validation dataset, which contained different facial images from those used in the final testing dataset, despite involving the same individuals and maintaining the same data composition. The selection process aimed to determine the optimal cosine similarity thresholds by balancing false-positive (FP) and false-negative (FN) counts. As shown in Figures 5(a) and 5(b), FaceNet and FaceNet512 exhibited several candidate thresholds that provided reasonable trade-offs between FP and FN counts. However, because granting access to an unauthorized individual presents a greater security risk than denying access to an authorized user, minimizing false positives was prioritized. Consequently, a threshold of 0.75 was selected for both FaceNet and FaceNet512. For InsightFace, several thresholds between 0.35 and 0.65 satisfied the selection criteria, as illustrated in Figure 5(c). Based on the observed FP-FN trade-off, a threshold of 0.60 was selected. These thresholds were subsequently fixed and applied to the separate testing dataset to generate the confusion matrices and calculate the final recognition performance metrics without further optimization.

AIJASET – Vol. 06, No. 03, November 2026. 387-396
https://doi.org/10.25077/aijaset.v6i3.384

As shown in Figure 6, FaceNet, FaceNet512, and InsightFace correctly rejected 43, 46, and 50 unauthorized samples, respectively. InsightFace achieved the best results, with no false positives. Table 1 summarizes the recognition performance in terms of accuracy, precision, sensitivity, specificity, and F1-score. FaceNet produced 167 true positives, 43 true negatives, 7 false positives, and 33 false negatives. FaceNet512 produced 190 true positives, 46 true negatives, 4 false positives, and 10 false negatives. InsightFace correctly classified all 200 authorized and 50 unauthorized samples, with no false positives or false negatives. The superior performance of InsightFace may be attributed to its face recognition architecture and optimization strategy, particularly the ArcFace-based angular margin loss. This approach encourages discriminative facial embeddings by reducing intra-class variation and increasing inter-class separation, potentially producing clearer decision boundaries between identities [12]. These characteristics are consistent with InsightFace's observed performance, which reached 100% on all five classification metrics in the testing dataset.

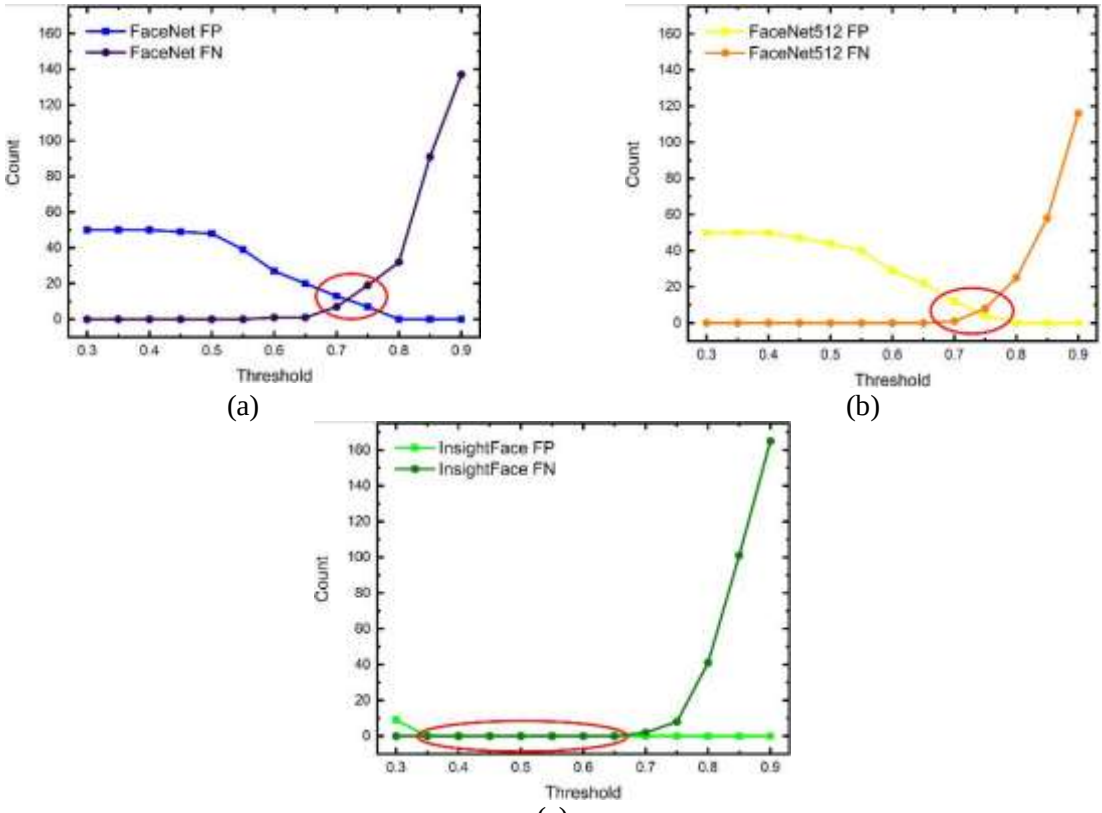


Figure 5. FP and FN counts at different thresholds using the validation dataset for (a) FaceNet, (b) FaceNet512, and (c) InsightFace

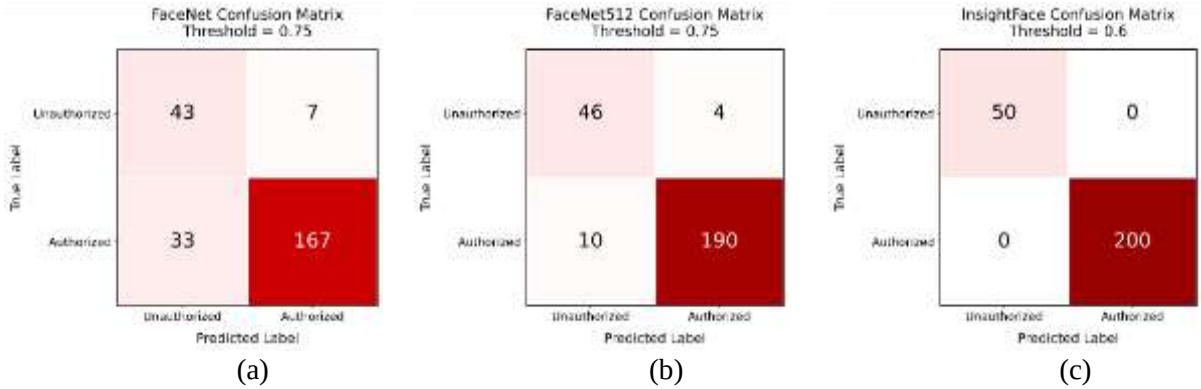


Figure 6. Confusion matrices obtained from the independent testing dataset for (a) FaceNet, (b) FaceNet512, and (c) InsightFace

Table 1. Face recognition performance metrics

Metric	FaceNet	FaceNet512	InsightFace
Accuracy	84.00%	94.40%	100.00%
Precision	95.98%	97.94%	100.00%
Sensitivity	83.50%	95.00%	100.00%
F1-score	89.30%	96.45%	100.00%
Specificity	86.00%	92.00%	100.00%

Real-time recognition performance was evaluated at distances of 0.5, 1.0, and 1.5 m between the user and the webcam. For each testing scenario, five recognition trials were conducted per individual at each camera distance for all three evaluated models. The distance-based evaluation involved 20 authorized users and 5 unauthorized users, resulting in 100 authorized-user trials and 25 unauthorized-user trials per distance for each model. Two scenarios were examined: authorized-user recognition and unauthorized-user rejection. A successful result in the authorized-user scenario indicated that a registered user was correctly recognized, whereas a successful result in the unauthorized-user scenario indicated that an unregistered person was correctly rejected. For authorized users, FaceNet achieved recognition rates of 92.00%, 84.00%, and 76.00% at 0.5, 1.0, and 1.5 m, respectively. The corresponding rates for FaceNet512 were 96.00%, 98.76%, and 89.33%, while InsightFace achieved 100.00%, 100.00%, and 97.33%. For unauthorized users, FaceNet achieved rejection rates of 100.00%, 88.00%, and 100.00% at the three distances, respectively. FaceNet512 and InsightFace both achieved 100.00% unauthorized-user rejection at every tested distance. As summarized in Table 2, InsightFace was the most robust to changes in camera distance and maintained high recognition performance even as faces appeared smaller in the captured frames.

Table 2. Authorized-user recognition and unauthorized-user rejection rates for FaceNet, FaceNet512, and InsightFace at different camera distances

Model	**FaceNet**			**FaceNet512**			**InsightFace**		
Distance (m)	**0.5**	**1.0**	**1.5**	**0.5**	**1.0**	**1.5**	**0.5**	**1.0**	**1.5**
Authorized (%)	92.00	84.00	76.00	96.00	98.76	89.33	100.00	100.00	97.33
Unauthorized (%)	100.00	88.00	100.00	100.00	100.00	100.00	100.00	100.00	100.00

The real-time processing performance of the three models was evaluated using latency and FPS, as summarized in Table 3. Latency was defined as the time required to process a captured camera frame and produce an authentication decision. Measurement began after the webcam successfully captured a frame and ended after face detection, embedding extraction, similarity matching, and classification as authorized or unauthorized. Model loading, camera initialization, display rendering, CSV writing, relay activation, and electromagnetic lock response time were excluded. Latency and FPS were measured for five users, with 30 testing iterations per user for each model. All models were evaluated on the same Raspberry Pi 5 hardware using an equivalent experimental procedure. These consistent measurement boundaries enabled a comparative assessment of the models' computational performance under the tested deployment conditions. FaceNet achieved an average latency of 0.614 s and an average processing speed of 1.978 FPS. FaceNet512 achieved an average latency of 0.628 s and a processing speed of 1.957 FPS. InsightFace recorded the lowest average latency (0.243 s) and the highest average processing speed (4.153 FPS). These results indicate that InsightFace provides the fastest authentication decisions among the evaluated models and is well suited to this Raspberry Pi-based smart door lock implementation.

Table 3. Comparison of real-time processing latency and FPS across the evaluated models

Statistic	FaceNet		FaceNet512		InsightFace	
	Latency (s)	FPS	Latency (s)	FPS	Latency (s)	FPS
Average	0.614	1.978	0.628	1.957	0.243	4.153
Median	0.492	2.032	0.497	2.014	0.244	4.151
Min	0.480	0.241	0.488	0.228	0.232	3.875
Max	4.175	2.082	4.441	2.049	0.261	4.316
Std	0.672	0.329	0.720	0.327	0.007	0.100

4. Conclusion

This study designed and implemented a real-time face recognition-based smart door lock system on a Raspberry Pi and compared three face recognition models: FaceNet, FaceNet512, and InsightFace. The system integrated a webcam, relay module, electromagnetic door lock, and access logbook that was updated in real time to automate access control for authorized laboratory users. The experiments showed that InsightFace outperformed FaceNet and FaceNet512 in both recognition accuracy and real-time deployment performance. On the testing dataset, InsightFace achieved 100% accuracy, precision, sensitivity, specificity, and F1-score, compared with accuracies of 94.40% for FaceNet512 and 84.00% for FaceNet. InsightFace also maintained high recognition rates across camera distances of 0.5, 1.0, and 1.5 m, achieving 100.00%, 100.00%, and 97.33%, respectively, together with a 100.00% unauthorized-user rejection rate at all tested distances. It achieved the lowest average latency (0.243 s) and the highest average processing speed (4.153 FPS). Overall, these findings highlight the importance of considering recognition accuracy, robustness, and computational efficiency when selecting face recognition models for embedded access control systems. Among the models evaluated, InsightFace offered the most favorable combination of these characteristics.

Future work should evaluate the proposed system using larger and more diverse datasets covering a wider range of demographic characteristics, lighting conditions, facial expressions, occlusions, and operating environments to assess model generalization. Additional experiments should also examine resistance to presentation attacks, including printed photographs, replayed images on digital screens, and facial masks, and investigate the integration of face anti-spoofing techniques to strengthen security. Finally, future studies could assess larger-scale deployments involving multiple cameras, centralized user management, cloud-based monitoring, and integration with broader IoT security infrastructures.

Acknowledgment

The authors thank the Lembaga Penelitian dan Pengabdian kepada Masyarakat of Universitas Kristen Krida Wacana (UKRIDA) for financially supporting the publication of this study.

References

[1] Y. Kortli, M. Jridi, A. Al Falou, and M. Atri, “Face Recognition Systems: A Survey,” Sensors, vol. 20, no. 2, p. 342, Jan. 2020, doi: 10.3390/s20020342.

[2] A. K. Jain, A. Ross, and S. Prabhakar, “An Introduction to Biometric Recognition,” IEEE Transactions on Circuits and Systems for Video Technology, vol. 14, no. 1, pp. 4–20, Jan. 2004, doi: 10.1109/TCSVT.2003.818349.

[3] T. Nguyen, B. Lakshmanan, and W. Sheng, “A Smart Security System with Face Recognition,” arXiv preprint arXiv:1812.09127, Dec. 2018, doi: https://doi.org/10.48550/arXiv.1812.09127.

[4] M. R., R. Y., R. R., and S. A., “Smart Home Security System using Iot, Face Recognition and Raspberry Pi,” Int. J. Comput. Appl., vol. 176, no. 13, pp. 45–47, Apr. 2020, doi: 10.5120/ijca2020920105.

[5] Mehrunnisa, J. S. Priyanka, and B. Sharath Chandra, "Automatic Door Un-locking and Security System Using Raspberry Pi," in 2022 IEEE International Conference on Distributed Computing and Electrical Circuits and Electronics (ICDCECE), IEEE, Apr. 2022, pp. 1–5. doi: 10.1109/ICDCECE53908.2022.9793143.

[6] S. I. Serengil and A. Ozpinar, "Boosted LightFace: A Hybrid DNN and GBM Model for Boosted Facial Recognition," Gazi University Journal of Science, vol. 39, no. 1, pp. 452–466, Mar. 2026, doi: 10.35378/gujs.1794891.

[7] Wirianto and T. Mauritsius, "The development of face recognition model in indonesia pandemic context based on dcnn and arcface loss function," International Journal of Innovative Computing, Information and Control, vol. 17, no. 5, pp. 1513–1530, Oct. 2021, doi: 10.24507/ijicic.17.05.1513.

[8] M. Wang and W. Deng, "Deep face recognition: A survey," Neurocomputing, vol. 429, pp. 215–244, Mar. 2021, doi: 10.1016/j.neucom.2020.10.081.

[9] A. Firmansyah, T. F. Kusumasari, and E. N. Alam, "Comparison of Face Recognition Accuracy of ArcFace, Facenet and Facenet512 Models on Deepface Framework," in 2023 International Conference on Computer Science, Information Technology and Engineering (ICCoSITE), IEEE, Feb. 2023, pp. 535–539. doi: 10.1109/ICCoSITE57641.2023.10127799.

[10] F. Schroff, D. Kalenichenko, and J. Philbin, "FaceNet: A unified embedding for face recognition and clustering," in 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Jun. 2015, pp. 815–823. doi: 10.1109/CVPR.2015.7298682.

[11] I. William, D. R. Ignatius Moses Setiadi, E. H. Rachmawanto, H. A. Santoso, and C. A. Sari, "Face Recognition using FaceNet (Survey, Performance Test, and Comparison)," in 2019 Fourth International Conference on Informatics and Computing (ICIC), IEEE, Oct. 2019, pp. 1–6. doi: 10.1109/ICIC47613.2019.8985786.

[12] J. Deng, J. Guo, J. Yang, N. Xue, I. Kotsia, and S. Zafeiriou, "ArcFace: Additive Angular Margin Loss for Deep Face Recognition," IEEE Trans. Pattern Anal. Mach. Intell., vol. 44, no. 10, pp. 5962–5979, Oct. 2022, doi: 10.1109/TPAMI.2021.3087709.

[13] D. Zeng, R. Veldhuis, and L. Spreeuwers, "A survey of face recognition techniques under occlusion," IET Biom., vol. 10, no. 6, pp. 581–606, Nov. 2021, doi: 10.1049/bme2.12029.

[14] S. Karahan, M. K. Yildirim, K. Kirtac, F. S. Rende, G. Butun, and H. K. Ekenel, "How Image Degradations Affect Deep CNN-Based Face Recognition?," in Proc. 15th Int. Conf. Biometrics Special Interest Group (BIOSIG), 2016, pp. 313–320, doi: 10.1109/BIOSIG.2016.7736924.

[15] S. Yedulapuram, R. Arabelli, K. Mahender, and C. Sidhardha, "Automatic Door Lock System by Face Recognition," IOP Conf. Ser. Mater. Sci. Eng., vol. 981, no. 3, p. 032036, Dec. 2020, doi: 10.1088/1757-899X/981/3/032036.

[16] R. Saputra and N. Surantha, "Smart and real-time door lock system for an elderly user based on face recognition," Bulletin of Electrical Engineering and Informatics, vol. 10, no. 3, pp. 1345–1355, Jun. 2021, doi: 10.11591/eei.v10i3.2955.

[17] S. A. Radzi, M. K. M. F. Alif, Y. N. Athirah, A. S. Jaafar, A. H. Norihan, and M. S. Saleha, "IoT based facial recognition door access control home security system using raspberry pi," International Journal of Power Electronics and Drive Systems (IJPEDS), vol. 11, no. 1, p. 417, Mar. 2020, doi: 10.11591/ijpeds.v11.i1.pp417-424.

[18] N. S. Irjanto and N. Surantha, "Home Security System with Face Recognition Based on Convolutional Neural Network," Int. J. Adv. Comput. Sci. Appl., vol. 11, no. 11, pp. 408–412, 2020, doi: 10.14569/IJACSA.2020.0111152.

[19] R. Cordova-Cardenas, D. Amor, and Á. Gutiérrez, "Edge AI in Practice: A Survey and Deployment Framework for Neural Networks on Embedded Systems," Electronics (Basel)., vol. 14, no. 24, p. 4877, Dec. 2025, doi: 10.3390/electronics14244877.

[20] A. A. A. Aboluhom and I. Kandilli, "Face recognition using deep learning on Raspberry Pi," The Computer Journal, vol. 67, no. 10, pp. 3020–3030, Oct. 2024, doi: 10.1093/comjnl/bxae066.

[21] S. I. Serengil and A. Ozpinar, "LightFace: A Hybrid Deep Face Recognition Framework," in 2020 Innovations in Intelligent Systems and Applications Conference (ASYU), 2020, pp. 23–27, doi: 10.1109/ASYU50717.2020.9259802.

[22] P. Hofer, M. Roland, P. Schwarz, and R. Mayrhofer, “Face to Face with Efficiency: Real-Time Face Recognition Pipelines on Embedded Devices,” in Advances in Mobile Computing and Multimedia Intelligence (MoMM 2023), 2023, doi: 10.1007/978-3-031-48348-6_11.