

Runtime and Security Evaluation of Lightweight Biometric Models for Edge Based Access Control

Paul Ndudiri Ohia ¹, Evaristus Chibuzor Nwoke ²

¹ Dept of Information Technology

² Dept of Computer Science

¹ University of Port Harcourt River State, Nigeria

² University of Agriculture and Environmental Sciences Umuagwo Imo state Nigeria

Abstract- This study presents a runtime and security evaluation of lightweight biometric models for edge-based access control. The work evaluates MobileNetV2 for face recognition and EfficientNetB0 for fingerprint recognition within an edge-aware access-control environment. The models were assessed using locally captured biometric datasets and public benchmark datasets, with emphasis on verification accuracy, false acceptance rate (FAR), false rejection rate (FRR), equal error rate (EER), ROC-AUC, model size, inference latency, throughput, and real-time decision performance. The results show that the lightweight biometric models can provide reliable identity verification while maintaining practical runtime performance for edge deployment. The face model achieved strong verification performance with low error rates, while the fingerprint model provided complementary biometric evidence for improved access reliability. The findings demonstrate that lightweight deep learning models can support secure, fast, and scalable biometric access control in edge computing environments where low latency, privacy protection, and real-time authentication are required.

Keywords- edge computing, bimodal biometrics, access control, fingerprint recognition, face recognition, contextual risk analysis, Chinese Wall policy, IPFS

I. INTRODUCTION

Access control remains a fundamental security requirement in distributed systems because it determines who can access protected resources, under which conditions, and at what privilege level. In edge and IoT environments, this requirement becomes more demanding because users, devices, and services operate across heterogeneous and latency-sensitive infrastructures [1], [2], [3]. Modern Zero Trust thinking further emphasizes that access should not be granted merely because a requester appears to be in a familiar location or on a trusted network; identity, device posture, and request context should all be evaluated explicitly before access is approved [4].

Although context-aware access control improves risk sensitivity, context alone does not always prove that the requester is the legitimate enrolled user. A trusted device, expected login time, or familiar location may still be exploited by an impostor. This study therefore extends a context-driven framework by adding bimodal biometric verification through face and fingerprint recognition, while retaining contextual reasoning and policy enforcement. The proposed system is intended to support low-latency decision making at the edge while improving resistance to impersonation, credential misuse, and unauthorized access [5], [6], [7].

II. RELATED WORK

Related works on access control have proposed policy-based and attribute-based methods for controlling access in distributed systems. Attribute-Based Access Control (ABAC) improves flexibility by evaluating user, resource, action, and environmental attributes before granting access [8]. More recently, adaptive and risk-aware models have been introduced for access control at the network edge. Muhammad et al. proposed an adaptive risk-based access control system for edge environments, using contextual factors such as resource sensitivity, action severity, and risk history to determine whether access should be granted or denied [9]. However, such context-based systems may still be limited because they do not always verify whether the requester is the real authorized user. This creates a gap for biometric-enhanced access control, where face and fingerprint verification can strengthen identity assurance before edge-based access decisions are finalized [6]. Within this setting, the key research gap is not the absence of contextual risk analysis, but the absence of strong identity assurance inside context-aware access workflows. The proposed study addresses this gap by combining biometric identity evidence with contextual and policy-based reasoning in a unified edge architecture.

A. Access Control at the Network Edge

Access control at the network edge has become important because edge computing moves computation, storage, and security decisions closer to users, devices, and

data sources. In this architecture, the Device Layer generates access requests through sensors, smartphones, biometric scanners, and IoT devices; the Edge Layer performs local authentication, risk evaluation, and policy enforcement; while the Cloud Layer supports long-term storage, model updates, and centralized management [10]. Prior studies show that edge computing reduces latency and supports real-time decision-making, but its distributed structure also increases exposure to threats such as unauthorized access, compromised nodes, spoofing, and data leakage [2], [3]. Recent work on edge computing highlights the need to move security decisions closer to users and devices in order to reduce latency and support real-time services [2], [3]. At the same time, the distribution of computation across exposed edge nodes introduces new concerns for authentication, privacy, and policy enforcement. These concerns make adaptive and identity-sensitive access control more important than purely static credential checks [1], [4].

B. Edge Computing Architecture

Edge computing brings computation and data storage closer to end-user devices and data sources instead of relying only on distant cloud servers [11], [12]. In an access-control environment, data from cameras, sensors, smartphones, and biometric scanners can be processed near the source, reducing latency, bandwidth use, and response time [11]. The architecture can be viewed in three layers: Device Layer, where biometric scanners, cameras, sensors, and user devices generate access data; Edge Layer, where gateways or edge servers process biometric and contextual information for real-time access decisions; and Cloud Layer, where large-scale storage, model updates, backup records, and long-term analytics are maintained. This layered approach supports fast, secure, and responsive biometric access control in edge environments [10].




Fig. 1. Edge computing architecture adapted from XenonStack [13].

C. Biometric System Theory

Biometric systems verify identity by comparing captured physiological traits with enrolled reference templates. In this study, face and fingerprint evidence provided direct identity assurance before the contextual and policy layers made the final access decision.

D. Biometric Technology

Biometric technology converts human traits into digital templates for identity matching. The proposed system used facial and fingerprint inputs as complementary modalities because face recognition is convenient, while fingerprint recognition provides stable ridge-pattern evidence.




Fig. 2. Fingerprint recognition process adapted from Maltoni et al. [17].

Bimodal biometric authentication combines complementary physiological evidence in order to reduce the weaknesses of unimodal verification. In this study, face recognition supported convenient and contact-light identity capture, while fingerprint recognition supplied a more stable and distinctive physiological factor. This design aligns with biometric theory and with recent multimodal authentication studies that seek stronger resilience against spoofing and identity misuse [6], [7], [15], [16].

E. Edge Computing and Biometric Authentication Systems

Edge computing has become important in biometric authentication systems because identity verification often requires fast, secure, and reliable processing. Traditional cloud-based biometric systems may experience high latency, network dependence, and privacy concerns when raw biometric data must be transmitted to remote servers. By processing biometric data closer to the user, edge computing reduces response time, limits unnecessary data transmission, and supports real-time access-control decisions [11], [12]. In bimodal biometric systems, edge devices can process facial and fingerprint data locally using deep learning models. Convolutional Neural Networks (CNNs) are useful for this purpose because they can extract discriminative features from

biometric images. Lightweight CNN architectures such as MobileNetV2 and EfficientNetB0 are suitable for edge environments because they are computationally efficient and can support faster recognition on resource-constrained devices [18], [19].

F. Machine Learning Models

Machine learning models improve access-control systems by enabling them to learn from historical data, detect patterns, assess contextual risk, and support adaptive decision-making. Instead of relying only on static rules, ML-based access control can evaluate factors such as user behavior, device trust, access history, and abnormal activity before granting or denying access [20], [21]. In this study, CNN-based models were applied to biometric recognition, while a 1D CNN/Conv1D model was used for contextual access-risk analysis.

G. Convolutional Neural Network

A Convolutional Neural Network (CNN) was used as the feature-learning mechanism for image-based biometric recognition. MobileNetV2 supported face recognition, while EfficientNetB0 supported fingerprint recognition; their confidence scores were passed into the wider access-control decision workflow.




Fig. 3. Structure of the convolutional neural network (CNN).

H. Edge Computing Deployment Models

Edge computing can be deployed in different ways depending on where computation and storage resources are placed. In a device-edge model, processing is performed directly on smart devices, sensors, cameras, or biometric scanners with limited dependence on external servers. In an edge-server model, nearby gateways, routers, or local edge servers handle real-time processing, authentication, and access-control decisions close to users. In a cloud-edge collaborative model, urgent tasks are processed at the edge, while the cloud supports large-scale storage, model training, backup, and long-term analytics [23], [11]. This layered

deployment approach is suitable for biometric access control because identity verification and risk analysis can be performed quickly at the edge, while records and system updates can still be managed through the cloud.

I. Benefits of Edge Computing

Edge computing provides several benefits for real-time access-control systems. It reduces latency by processing data close to the source instead of sending every request to a distant cloud server. It reduces bandwidth consumption because only selected or processed data is transmitted to the cloud. Edge computing improves privacy by allowing sensitive biometric and contextual data to be processed locally. It supports faster decision-making, which is important for security systems that require immediate grant, step-up, or deny responses [11], [12]. In the proposed system, these benefits allow biometric verification, contextual risk analysis, and policy enforcement to occur close to the access point.

III. MATERIALS AND METHODS

The study adopted an object-oriented analysis and design methodology for developing the proposed access-control system. This methodology was suitable because the system consists of interacting entities such as users, biometric inputs, access requests, edge nodes, risk-analysis modules, policy enforcement components, and administrative records. The approach supported clear modelling of system actors, workflows, data objects, and decision processes before implementation. The proposed prototype was implemented as a web-based access-control platform using Python and Flask, with MySQL for structured records and IPFS for content-addressed storage of biometric-related artifacts [24].

A. Proposed System Architecture

The proposed system architecture was implemented as an edge-aware web application using Python and Flask. MySQL stored structured entities such as users, biometric labels, access logs, contextual records, and historical access decisions, while IPFS stored biometric-related artifacts using content-addressed identifiers. The intelligence layer of the system integrated MobileNetV2 CNN for facial recognition, EfficientNetB0 CNN for fingerprint recognition, and a 1D CNN/Conv1D model for contextual risk analysis. These models worked together to verify user identity, analyze access conditions, and support real-time decision-making at the edge. During an access request, the face and fingerprint models generated biometric confidence scores, while the contextual model evaluated attributes such as device trust, login time, known location, resource sensitivity, and access history. The

system then combined biometric evidence, contextual risk output, and policy restrictions to produce a grant, step-up, or deny decision. This combination allowed the system to provide transactional control, biometric identity assurance, contextual risk evaluation, policy-aware authorization, and verifiable storage of sensitive supporting files [24]. Fig. 4 shows the main architecture of the proposed system.




Fig. 4. Proposed system architecture for the edge-based bimodal biometric access-control framework.

B. Edge Deployment and Runtime Measurement

The prototype was developed and evaluated using a local edge aware configuration consisting of a computer system with at least 8 GB RAM and 500 GB storage, a USB fingerprint reader, an HD webcam, network connectivity through a router, and a Raspberry Pi or local edgesserver option. The software environment included Python 3.10, Flask, MySQL 8.x, TensorFlow, OpenCV, NumPy, scikit-learn, IPFS, XAMPP, and a standard browser interface.

Inference timing was measured on the tested Windows CPU environment using the saved biometric models. The MobileNetV2 face model size was 12.95 MB, with batch inference time of 7,887.58 ms for 105 test images, throughput of 13.31 samples/s, mean single-image latency of 227.23 ms, median latency of 204.15 ms, and p95 latency of 351.25 ms. The EfficientNetB0 fingerprint model size was 20.04 MB, with batch inference time of 18,512.65 ms for 240 test images, throughput of 12.96 samples/s, mean single-image latency of 237.63 ms, median latency of 235.42 ms, and p95 latency of 251.91 ms.

These values support prototype level edge-aware runtime performance on the tested CPU system.

TABLE 1. Runtime Model Metrics Showing Model Size, Latency, P95 Latency, And Throughput.

No.	Metric	Face (Mobile-NetV2)	Fingerprint (EfficientNetB0)
1	Batch time (ms)	7,887.58	18,512.65
2	Throughput (samples/s)	13.31	12.96
3	Mean latency (ms)	227.23	237.63
4	Median latency (ms)	204.15	235.42
5	P95 latency (ms)	351.25	251.91
6	Model size (MB)	12.95	20.04
7	Latency samples	12	12

```

Windows PowerShell
ical results due to floating-point round-off errors from different computation orders. To turn them off, set the environment variable 'TF_ENABLE_ONEDNN_OPTS=0'.
WARNING: All log messages before absl::InitializeLog() is called are written to STDERR
[2024-09-06 00:00:17.90093677 772485 17836 port.cc:153] onednn custom operations are on. You may see slightly different numerical results due to floating-point round-off errors from different computation orders. To turn them off, set the environment variable 'TF_ENABLE_ONEDNN_OPTS=0'.
[2024-09-06 00:00:17.90093683 976409 17836 cpu_feature_guard.cc:227] This TensorFlow binary is optimized to use available CPU instructions in performance-critical operations.
To enable the following instructions: SSE3 SSE4.1 SSE4.2 AVX, in other operations, rebuild TensorFlow with the appropriate compiler flags.
WARNING:tensorflow:GPU support is not available on native Windows for TensorFlow >= 2.11. Even if CUDA/cuDNN are installed, GPU will not be used. Please use MSIL or the TensorFlow-DirectML plugin.
Saved reviewer metric outputs:
- C:\Users\Upauli\Documents\Myproject\training_reports\reviewer_metrics\authentication_metrics.csv
- C:\Users\Upauli\Documents\Myproject\training_reports\reviewer_metrics\runtime_model_metrics.csv
- C:\Users\Upauli\Documents\Myproject\training_reports\reviewer_metrics\single_image_audit.csv
- C:\Users\Upauli\Documents\Myproject\training_reports\reviewer_metrics\reviewer_metrics_full.json
Authentication metrics:
system threshold genuine_attempts impostor_attempts FAR FRR EER EER_threshold ROC-AUC verification_accuracy model_size_mb batch_total_ms batch_throughput_samples_per_sec single_mean_ms single_std_ms
face 0.7 305 630 0.000000 0.000000 0.000000 0.72185 1.000000 195.5866 240.918695 12.949924 11.538092 203.899450
fingerprint 0.7 240 2160 0.000463 0.195833 0.013194 0.230437 0.998679 235.7855 265.502265 12.9 12.265594 238.959595
fused_mean_face_fingerprint 0.7 105 630 0.000000 0.047619 0.000000 0.533593 1.000000 0.993197 NaN NaN NaN NaN NaN
Split audit:
system v_disjoint identity_overlap_count file_overlap dataset_dir train_split test_split identity_interpretation
samples_per_class
face C:\Users\Upauli\Documents\Myproject\datasets\face_identity_balanced test_samples_per_class 0.7 0.3
false 7 0 image-disjoint but not identity-disjoint

```

Fig.5: Terminal output showing runtime and authentication metrics.

This output shows the biometric verification performance of the system. The face model performed very strongly, with FAR = 0.0000, FRR = 0.0000, EER = 0.0000, ROC-AUC = 1.0000, and verification accuracy = 1.0000. This means the face model correctly separated genuine users from impostors in the test evaluation. The fingerprint model also performed well, with a very low FAR of 0.000463 and ROC-AUC of 0.998679. However, its FRR is higher at 0.195833, meaning some genuine fingerprint attempts were rejected at the selected threshold of 0.70. The fused face-fingerprint decision improved the overall balance, achieving FAR = 0.0000, FRR = 0.047619, and verification accuracy = 0.993197. This shows that combining both biometric

modalities improves reliability compared with relying on fingerprint alone.




Fig.6: Latency comparison chart for face versus fingerprint models.




Fig.7: Model size comparison chart for face versus fingerprint models

C. Dataset Analysis

The study used both public and locally captured datasets to develop the proposed bimodal biometric access-control system. VGGFace2 and SOCOFing supported early experimentation with face and fingerprint recognition, while locally captured enrolled-user samples were used for the final operational training to reflect the real deployment environment [25], [26]. Face and fingerprint data were organized into identity-based folders for supervised learning, while contextual access records were prepared as tabular feature vectors for risk analysis. Preprocessing involved image resizing, normalization, labeling, contextual data cleaning, encoding, scaling, and class balancing. The datasets were split into 70% training and 30% testing. Evaluation was based on accuracy, precision, recall, F1-score, confusion matrices, and functional system testing to confirm support for identity verification, risk evaluation, and real-time edge-based access decisions.

D. Dataset Composition and Split Protocol

The locally captured operational dataset consisted of enrolled identity folders/users. The complete face dataset in the main work contained 1,090 images: myself-face = 190 images and user2-face to user10-face = 100 images each. The fingerprint dataset table contained 1,013 samples: myself-finger = 112, user2-finger = 101, user3-finger = 101, user4-finger to user9-finger = 100 each, and user10-finger = 101.

Public datasets were used for preliminary experimentation and comparison only. VGGFace2 supported face-recognition experimentation, SOCOFing supported fingerprint-recognition experimentation, and CMU/KiltHub supported contextual access-record analysis. The final operational evaluation emphasized locally captured enrolled-user samples because they matched the intended access-control environment.

For the verification-metric evaluation, the 70/30 split was image-disjoint but not identity-disjoint. For the face dataset, the split used seven identities with 50 images each: 35 images per identity for training and 15 images per identity for testing. For the fingerprint dataset, the split used 10 identities with 80 images each: 56 images per identity for training and 24 images per identity for testing. No image file appeared in both partitions, but the same identity classes appeared in both training and testing. Therefore, the experiment evaluated recognition of enrolled identities using unseen image samples rather than generalization to completely unseen users.

Table 2. Face Dataset Split Audit.

No.	Audit item	Result
1	Dataset	face_identity_balanced
2	Training images	70%
3	Test images	30%
4	Identities in both splits	7
5	Files shared between splits	0
6	Split type	Separate images of the same enrolled identities

shows that the face dataset used a 70/30 training and test split that was image-disjoint but not identity-disjoint. No image file appeared in both sets, but the same seven enrolled identities appeared in each. The results therefore measure

recognition of new images of enrolled users, not recognition of previously unseen identities.



Fig. 8. Public face image dataset (VGGFace2 from Kaggle).




Fig. 9. Public fingerprint dataset (SOCOFing from Kaggle).




Fig. 10. Locally captured face image dataset.




Fig. 11. Locally captured fingerprint dataset.

E. Bias Mitigation

Bias mitigation was considered important because biometric and contextual access-control models can produce unfair or unreliable decisions when trained on imbalanced or

poorly representative data. In the proposed system, bias was reduced by combining public datasets with locally captured enrolled-user samples, ensuring that the final model reflected the actual deployment environment. The face and fingerprint datasets were organized by user identity, while pre-processing steps such as resizing, normalization, labeling, and controlled class balancing were applied before training.

For contextual access analysis, records were cleaned, encoded, scaled, and balanced to reduce the effect of dominant classes on the 1DCNN/Conv1D model. Model performance was also evaluated using accuracy, precision, recall, F1-score, and confusion matrices instead of relying on accuracy alone. This helped identify whether the models performed consistently across classes. In addition, the system combined biometric confidence scores with contextual risk analysis and Chinese Wall policy enforcement, reducing overdependence on a single model output during final access decisions.

F. Technology Stack and Data Organization

Table 3 summarizes the main technologies and analytical components integrated into the proposed prototype. The system was implemented using Python and Flask for backend logic, routing, and web-based access-control operations, while MySQL was used to store structured records such as user profiles, biometric labels, access logs, contextual attributes, and decision history. IPFS supported decentralized and content-addressed storage of biometric-related artifacts, improving traceability and reducing dependence on a single storage point.

The analytical layer combined MobileNetV2 for face recognition, EfficientNetB0 for fingerprint recognition, and a 1D CNN/Conv1D model for contextual risk analysis. These models processed locally captured and public training evidence together with structured access-control records.

TABLE 3. Core Technologies, Analytical Modules, And Functions In The Proposed System.

Component	Technology and function
Backend service	Python for core application logic and system integration
Web framework	Flask for routing, forms, dashboards, and API handling
Relational database	MySQL for structured storage of users, logs, and decisions
Decentralized storage	IPFS for content-addressed storage of biometric-related artifacts
Face module	MobileNetV2 CNN for face feature learning and identity

	verification
Fingerprint module	EfficientNetB0 CNN for primary fingerprint verification
Context module	Conv1D CNN for contextual risk analysis and decision support

Note: the table summarizes the primary implementation and decision-support components used in the prototype.

CXID	UID	USER ID	ROLE	SERVICE FILE	ACTIVITY	ACCESS TIME	HOUR	PURPOSE	DEVICE	STATUS	LOCATION
CTX-001	USR-1001	Admin	Officer	\\fs1\civc	USB Connect	2008-05-01 09:07:11	9	PC-117	Trusted	Known	
CTX-002	USR-1002	Client Officer	\\fs1tp.civc	HTTP Visit	2008-05-01 12:24:12	12	PC-114	Trusted	Known		
CTX-003	USR-1003	Research Staff	\\fs1\2010-01-01	Database Query	2008-05-01 18:12:11	18	PC-120	Trusted	Known		
CTX-004	USR-1004	Client Officer	\\fs1tp.civc	Login	2008-05-01 12:44:18	12	PC-116	Trusted	Known		
CTX-005	USR-1005	Operations User	\\fs1tp.civc	USB Connect	2008-05-01 21:35:21	21	PC-118	Trusted	Known		
CTX-006	USR-1006	Security Analyst	\\fs1tp.civc	HTTP Visit	2008-05-01 06:42:08	6	PC-1202	Trusted	Known		
CTX-007	USR-1007	Research User	\\fs1\2010-01-01	Database Query	2008-05-01 03:47:17	3	PC-1219	Trusted	Known		
CTX-008	USR-1008	Finance Analyst	\\fs1tp.civc	Login	2008-05-01 06:56:28	6	PC-120	Trusted	Known		
CTX-009	USR-1009	Admin Officer	\\fs1tp.civc	USB Connect	2008-05-01 09:03:39	9	PC-1203	Trusted	Known		
CTX-010	USR-1010	Client Officer	\\fs1tp.civc	HTTP Visit	2008-05-01 12:10:50	12	PC-125	Trusted	Known		
CTX-011	USR-1011	Research Staff	\\fs1\2010-01-01	Database Query	2008-05-01 15:17:15	15	PC-1287	Untrusted	Known		
CTX-012	USR-1012	IT Admin	\\fs1tp.civc	Login	2008-05-01 18:21:12	18	PC-1304	Trusted	Known		
CTX-013	USR-1013	Operations User	\\fs1tp.civc	USB Connect	2008-05-01 21:34:21	21	PC-121	Trusted	Unknown		
CTX-014	USR-1014	Research Analyst	\\fs1tp.civc	Database Query	2008-05-01 03:47:34	3	PC-120	Trusted	Known		
CTX-015	USR-1015	Research User	\\fs1\2010-01-01	Database Query	2008-05-01 03:26:35	3	PC-1355	Trusted	Known		
CTX-016	USR-1016	Finance Analyst	\\fs1tp.civc	Login	2008-05-01 06:52:56	6	PC-1372	Trusted	Known		
CTX-017	USR-1017	Admin Officer	\\fs1tp.civc	USB Connect	2008-05-01 09:05:39	9	PC-1389	Untrusted	Known		
CTX-018	USR-1018	Client Officer	\\fs1tp.civc	HTTP Visit	2008-05-01 12:36:18	12	PC-1406	Trusted	Known		
CTX-019	USR-1019	Research Staff	\\fs1\2010-01-01	Database Query	2008-05-01 15:12:15	15	PC-1423	Trusted	Unknown		

Fig. 12. Contextual dataset used for access-risk analysis.

The contextual dataset contains access-related attributes used to evaluate the risk level of each access request. It includes factors such as device trust, known location, login time, resource sensitivity, access history, and user behavior. Unlike biometric datasets, which are image-based, the contextual dataset is structured as tabular data and is used by the 1D CNN/Conv1D model to learn access patterns and support risk-aware decision-making. In the proposed system, this dataset helped determine whether a verified user should be granted access, required to perform step-up verification, or denied access based on the surrounding access conditions.

G. CNN Model Training Process

```

Select Command Prompt

Training face model from C:\Users\Paul\Documents\MyProject\datasets\face_identity ...
Training class distribution: ['0', '1', '2', '3', '4', '5', '6', '7', '8', '9', '10', '11', '12', '13', '14', '15', '16', '17', '18', '19', '20', '21', '22', '23', '24', '25', '26', '27', '28', '29', '30', '31', '32', '33', '34', '35', '36', '37', '38', '39', '40', '41', '42', '43', '44', '45', '46', '47', '48', '49', '50', '51', '52', '53', '54', '55', '56', '57', '58', '59', '60', '61', '62', '63', '64', '65', '66', '67', '68', '69', '70', '71', '72', '73', '74', '75', '76', '77', '78', '79', '80', '81', '82', '83', '84', '85', '86', '87', '88', '89', '90', '91', '92', '93', '94', '95', '96', '97', '98', '99']
Priority labels: ['0', '1', '2', '3', '4', '5', '6', '7', '8', '9', '10', '11', '12', '13', '14', '15', '16', '17', '18', '19', '20', '21', '22', '23', '24', '25', '26', '27', '28', '29', '30', '31', '32', '33', '34', '35', '36', '37', '38', '39', '40', '41', '42', '43', '44', '45', '46', '47', '48', '49', '50', '51', '52', '53', '54', '55', '56', '57', '58', '59', '60', '61', '62', '63', '64', '65', '66', '67', '68', '69', '70', '71', '72', '73', '74', '75', '76', '77', '78', '79', '80', '81', '82', '83', '84', '85', '86', '87', '88', '89', '90', '91', '92', '93', '94', '95', '96', '97', '98', '99']
Instructions in performance-critical operations: S54 S56.1 S56.2 AUX, in other operations, rebuild TensorFlow with the appropriate compiler flags.
WARNING: TensorFlow GPU support is not available on native machines for TensorFlow > 2.11. Even if CUDA/cuDNN is installed, GPU will not work. Please use MKL or the TensorFlow-DirectML plugin.
Epoch 1/25
[1m30.3b] [m-12s] [m-37m] [m-148ms] [m-90ms/step - accuracy: 0.1574 - loss: 2.9153
Epoch 2/25
[1m30.3b] [m-12s] [m-37m] [m-148ms] [m-90ms/step - accuracy: 0.1776 - loss: 2.8182 - val_accuracy: 0.1356 - val_loss: 2.8814 - learning_rate: 5.0000e-04
Epoch 3/25
[1m30.3b] [m-12s] [m-37m] [m-148ms] [m-91ms/step - accuracy: 0.1774 - loss: 2.5386
Epoch 4/25
[1m30.3b] [m-12s] [m-37m] [m-148ms] [m-90ms/step - accuracy: 0.2068 - loss: 2.5197 - val_accuracy: 0.1516 - val_loss: 2.2064 - learning_rate: 0.00012500000000000000
Epoch 5/25
[1m30.3b] [m-12s] [m-37m] [m-148ms] [m-87ms/step - accuracy: 0.2132 - loss: 2.4512
Epoch 6/25
[1m30.3b] [m-12s] [m-37m] [m-148ms] [m-94ms/step - accuracy: 0.2431 - loss: 2.4266 - val_accuracy: 0.1559 - val_loss: 2.3559

```

Fig. 13. Model training process of facial biometric data.

[illegible]

Fig. 14. Fingerprint model training process.

Figures 13 and 14 show the training process for the face and fingerprint recognition models used in the proposed system. Figure 13 presents the MobileNetV2 face-recognition model training, where facial images were resized, normalized, labeled, and passed through the CNN model to learn identity-based facial features. The model was trained to distinguish enrolled users and generate confidence scores during authentication.

Figure 14 presents the EfficientNetB0 fingerprint-recognition model training, where fingerprint samples were preprocessed and organized into identity classes before training. The model learned distinctive fingerprint patterns and produced identity confidence scores for access-control decisions. Together, the face and fingerprint models formed the bimodal biometric layer of the system.

H. Training Configuration and Reproducibility

MobileNetV2 and EfficientNetB0 used fixed image inputs, Adam optimizer, categorical cross-entropy loss, 20 epochs, batch size 16, learning rate 0.0003, ImageNet weights, and a 70/30 train/test split. The face model used MobileNetV2 preprocessing, while the fingerprint model used EfficientNetB0 preprocessing. Training used controlled augmentation, including rotation, shifting, zooming, brightness variation, and horizontal flipping where appropriate for face images. Fingerprint augmentation used lighter geometric variation to avoid damaging ridge structure.

The contextual Conv1D model used tabular access-control records with seven core features: login hour, trusted-device status, known-location status, face score, fingerprint score, resource sensitivity, and Chinese Wall conflict status. The Conv1D implementation used 45 epochs, batch size 8, convolution layers with batch normalization, max pooling, dense layers, dropout of 0.30, and contextual labels mapped to verified, uncertain, and failed classes. The biometric training and evaluation scripts used a fixed random seed of 42 to support reproducible partitioning and reporting.

The Conv1D model was selected because the contextual inputs were ordered numerical access-control features rather than images, allowing the model to learn local relationships among time, device, location, biometric confidence, resource sensitivity, and policy-conflict indicators.

I. Model Performance Evaluation




Fig. 15. MobileNetV2 real face confusion matrix.




Fig. 16. EfficientNetB0 real fingerprint confusion matrix.




Fig. 17. Contextual Conv1D confusion matrix.

Figures 15, 16, and 17 show the confusion matrices for the three intelligent models used in the proposed system.

Figure 15 presents the MobileNetV2 face-recognition confusion matrix, Figure 16 shows the EfficientNetB0 fingerprint-recognition confusion matrix, and Figure 14 presents the Contextual Conv1D confusion matrix for access-condition classification.

J. System Deployment and Testing

System deployment and testing covered registration, biometric capture, login, contextual scoring, Chinese Wall enforcement, and administrative reporting. Functional testing confirmed that enrolled face and fingerprint samples could be captured, stored, and evaluated within the unified application workflow.

Fig. 18. Enrollment form page.

Fig. 19. Biometric face capture page.




Fig. 20. Biometric fingerprint capture page.

K. Proposed System Workflow

The proposed workflow begins when a user submits credentials and biometric evidence at an access point. The edge node captures face and fingerprint inputs, verifies them using MobileNetV2 CNN and EfficientNetB0 CNN, and evaluates contextual attributes such as device trust, location, login time, resource sensitivity, and access history using a 1D CNN Conv1D model. The system then applies Chinese Wall policy restrictions before producing a final grant, step-up, or deny decision.

L. Multimodal Fusion and Final Access Decision

Face verification was accepted when the predicted identity matched the claimed identity and the face confidence score was at least 0.70, or when reference similarity reached the configured strong-match threshold of 0.90. Fingerprint verification was accepted when the predicted fingerprint identity matched the claimed identity and the confidence score was at least 0.70.

The bimodal biometric decision used a decision-level fusion rule. If both `face_match` and `fingerprint_match` were true, `biometric_status` was classified as verified. If only one biometric modality matched, `biometric_status` was classified as uncertain and the system required step-up verification. If neither modality matched, `biometric_status` was classified as failed. The contextual Conv1D model then classified the request as verified, uncertain, or failed using device, location, time, resource sensitivity, biometric scores, and Chinese Wall conflict information.

The final access decision combined biometric status, contextual risk, and Chinese Wall policy. A verified biometric

result with low contextual risk and no Chinese Wall conflict produced GRANT. A verified biometric result with uncertain contextual risk produced STEP-UP. An uncertain biometric result under low or medium contextual risk also produced STEP-UP. A failed biometric result, high contextual risk, or Chinese Wall violation produced DENY.

M. Context-Aware Authorization and Policy Enforcement

Context-aware authorization remains essential because authentication alone does not determine what a verified user should be allowed to access under a specific situation. In the proposed system, a user first submits biometric and access-request information through the edge-aware application. The system then verifies identity using facial recognition and fingerprint recognition, while contextual attributes such as device trust, known location, login hour, resource sensitivity, access history, and conflict-of-interest status are evaluated before the final decision is made. The 1D CNN/Conv1D contextual model supports risk analysis from these access conditions, while the Chinese Wall policy prevents access where organizational conflicts exist. This workflow aligns with attribute-based and risk-adaptive access control principles by combining identity assurance, contextual risk evaluation, and policy enforcement within a real-time edge computing environment [8], [21], [27].

IV. RESULTS

The biometric and contextual modules produced strong results. The MobileNetV2 face model achieved about 0.96 overall accuracy, while the enrolled-user face class recorded 0.92 precision, 0.94 recall, and 0.93 F1-score. The EfficientNetB0 fingerprint model achieved about 93.36% overall classification accuracy; the enrolled-user fingerprint class recorded 95.83% precision, 95.83% recall, and 95.83% F1-score, and the highest validation accuracy reached 93.78% during training. The contextual Conv1D model achieved 84.08% test accuracy on 157 records, with the verified class producing 96.77% recall and 90.91% F1-score. Integrated system testing showed that biometric evidence, contextual analysis, and Chinese Wall enforcement worked together to support real-time grant, step-up, and deny decisions.

A. Authentication Verification Metrics

In addition to accuracy, precision, recall, F1-score, and confusion matrices, the biometric modules were evaluated using authentication-oriented metrics. The evaluation reported threshold, FAR, FRR, EER, ROC-AUC, genuine attempts, impostor attempts, and verification accuracy separately for face, fingerprint, and the fused face-fingerprint decision.

The configured decision thresholds were face threshold = 0.70, face reference threshold = 0.68, strong face-match threshold = 0.90, fingerprint threshold = 0.70, and fused biometric threshold = 0.70.

At threshold 0.70, the face model achieved FAR = 0.000000, FRR = 0.000000, EER = 0.000000, ROC-AUC = 1.000000, and verification accuracy = 1.000000. The fingerprint model achieved FAR = 0.000463, FRR = 0.195833, EER = 0.013194, ROC-AUC = 0.998679, and verification accuracy = 0.980000. The fused mean face-fingerprint decision achieved FAR = 0.000000, FRR = 0.047619, EER = 0.000000, ROC-AUC = 1.000000, and verification accuracy = 0.993197.

Table 4. Authentication Metrics For Face, Fingerprint, And Fused Face-Fingerprint Decisions.

No .	Metric	Face	Fingerpri nt	Fused face + fingerprint
1	Decision threshold	0.70	0.70	0.70
2	Genuine attempts	105	240	105
3	Impostor attempts	630	2,160	630
4	False acceptance rate (FAR)	0%	0.046%	0%
5	False rejection rate (FRR)	0%	19.58%	4.76%
6	Equal error rate (EER)	0%	1.32%	0%
7	EER threshold	0.9721	0.2304	0.5336
8	ROC AUC	1.0000	0.9987	1.0000
9	Reported verification accuracy	100%	98%	99.32%
10	Model size (MB)	12.95	20.04	—
11	Batch time (ms)	5,531.62	14,362.86	—
12	Throughput (samples/s)	18.98	16.71	—

13	Mean latency (ms)	134.86	167.93	—
14	Median latency (ms)	126.28	167.14	—
15	P95 latency (ms)	162.96	176.94	—
16	Latency samples	12	12	—




Fig. 21. ROC curves for face, fingerprint, and fused biometric verification.

The ROC curves show how well the biometric models distinguish genuine users from impostors. The face model and fused biometric decision achieved very strong separation, with ROC-AUC values close to 1.0, indicating excellent verification performance. The fingerprint model also performed strongly, with a high ROC-AUC, showing that it reliably supports identity verification. Overall, the fused biometric approach provides the most dependable authentication decision because it combines evidence from both face and fingerprint recognition.

The face model shows a strong verification performance, with FAR and FRR staying very low across the selected threshold range. This means the model can correctly accept genuine users while rejecting impostors with minimal error.




Fig. 22. Fingerprint EER threshold curve.

The fingerprint curve shows the balance point where FAR and FRR are closest. The low EER value indicates that the fingerprint model is reliable, although its FRR is higher than the face model, meaning some genuine fingerprint attempts may be rejected at stricter thresholds.




Fig. 23. Fused biometric FAR/FRR threshold curve.

The fused face-fingerprint curve gives the strongest overall decision balance. By combining both modalities, the system reduces false acceptance and improves confidence in access decisions, making the final biometric verification more suitable for security-sensitive access control.

B. Enhanced Performance Evaluation




Fig. 24. CNN identity confidence score distribution.




Fig. 25. Grant decision by identity confidence score.

Figures 22, 23, and 24, 25 show the FAR versus FRR curve, identity-confidence distribution, and grant-by-identity-confidence output of the biometric-enhanced access-control system. The FAR/FRR curve examines the balance between false acceptance and false rejection, while the confidence-score outputs show how biometric evidence influenced final access decisions.

V. PERFORMANCE IMPROVEMENT

The results indicate that the proposed system improves on context-only access control by adding direct biometric identity assurance to edge-based authorization. The study also showed that locally captured biometric samples improved operational dependability because they reduced the domain gap between training data and live deployment conditions. The proposed system shows clear improvements over context-only access control based on the following findings

A. Ablation and System-Level Comparison

The system-level comparison is organized as an ablation study that separates the effect of each major component. The comparison should report context only, face only, fingerprint only, bimodal biometrics, biometrics plus context, and the complete system with Chinese Wall policy enforcement.

The ablation table should compare the input evidence used, decision rule, verification accuracy or access-decision accuracy, FAR, FRR, step-up rate, deny rate, and security observation for each configuration. This structure shows whether the claimed improvement comes from identity evidence, contextual risk analysis, policy enforcement, or their combined use.

The contextual-only and biometric-enhanced decision outputs already provide the baseline comparison, while the biometric metrics provide face-only, fingerprint-only, and fused biometric evidence for the expanded ablation table.

Table 5. Ablation Comparison Of Context Only, Face Only, Fingerprint Only, Bimodal Biometrics, Biometrics Plus Context, And Complete System With Chinese Wall Policy.

Configuration	Main Result
Context only	83.83% accuracy
Face only	100.00% verification accuracy
Fingerprint only	98.00% verification accuracy
Bimodal biometrics	99.32% verification accuracy
Biometrics + context	Combines fused biometric verification with contextual risk analysis to reduce overdependence on either identity evidence or context alone.
Complete system with Chinese Wall policy	Adds policy enforcement to biometric and contextual decisions, ensuring that access is denied when conflict-of-interest rules are violated.

The ablation comparison shows that the contextual-only model achieved 83.83% accuracy, while unimodal biometric verification achieved 100.00% for face and 98.00% for fingerprint. The fused bimodal biometric decision improved overall reliability, achieving 99.32% verification accuracy with FAR = 0.000000 and FRR = 0.047619. This confirms that combining face and fingerprint evidence improves biometric decision reliability before contextual and Chinese Wall policy enforcement are applied.




Fig. 26. Ablation system-level comparison chart.

The ablation chart compares the contribution of each system component. The contextual-only model achieved 83.83% accuracy, while face-only verification achieved 100.00%, fingerprint-only verification achieved 98.00%, and fused bimodal biometrics achieved 99.32%. The result shows that biometric evidence improves the identity-verification layer, while the context and Chinese Wall stages add risk-aware and policy-aware access control.

1. Added direct identity verification to context-based access control.
2. Reduced the risk of granting access based only on trusted device or location.
3. Improved detection of uncertain or weak identity-confidence cases.
4. Supported better grant, step-up, and deny decision separation.
5. Strengthened edge-based authorization by combining biometric evidence, contextual risk, and policy enforcement.




Fig. 27. Contextual-only decision output.




Fig. 28. Biometric-enhanced decision output.




Fig. 29. Decision transition heatmap.

VI. DISCUSSION

The Contextual Only Decision Output represents the baseline system, where access decisions were based on factors such as device trust, location, login time, access history, resource sensitivity, and policy restrictions. Although useful, this approach could still grant access when context appeared normal without confirming the requester's real identity. The Biometric Enhanced Decision Output shows the improved model after face and fingerprint verification were added. This made the system more identity-aware and reduced overdependence on contextual trust. The Decision Transition Heatmap shows that some requests previously granted by the contextual-only model were changed to step-up or deny decisions when biometric confidence was weak, confirming stronger and more reliable access control.

VII. CONCLUSION

The study demonstrates that combining bimodal biometrics, contextual CNN analysis, and Chinese Wall policy enforcement can produce a more secure and adaptive access-control framework for edge computing environments. The

final system delivered stronger identity verification, improved decision reliability, and practical real-time operation through edge-aware processing. Future work should expand the locally captured biometric dataset, refine multimodal score fusion, and validate the framework with a larger user population and more deployment sites.

VIII. RECOMMENDATION AND FUTURE WORK

Future work should expand the real-user face and fingerprint dataset to improve verification accuracy and reduce false acceptance and false rejection rates. Locally captured biometric data should remain the main training source, while public datasets should be used for benchmarking and comparison. Further studies should also improve threshold calibration, multimodal fusion, and test the system with more users in real edge-based environments such as universities, healthcare, banking, and enterprise offices.

IX. CONTRIBUTION TO KNOWLEDGE

This study contributes to knowledge by developing an enhanced access control framework that combines bimodal biometric verification, 1D CNN-based contextual risk analysis, and Chinese Wall policy enforcement in an edge computing environment. It shows that adding face and fingerprint verification strengthens contextual access control by providing direct identity evidence for grant, step-up, or deny decisions.

The study also demonstrates that focused enrolled-user biometric training can improve verification performance when datasets are limited or imbalanced. In addition, it shows that a 1D CNN can be applied to structured contextual access data, allowing features such as trusted device, known location, login time, and access condition to support adaptive risk-based decision-making.

X. DATA AVAILABILITY

Public benchmark datasets referenced in this study include VGGFace2 for face-recognition experimentation, SOCOFing for fingerprint-recognition experimentation, and the KiltHub Insider Threat Test Dataset for contextual access-record analysis. VGGFace2 is available at https://github.com/ox-vgg/vgg_face2, SOCOFing is available at <https://www.kaggle.com/ruizgara/socofing>, and the KiltHub Insider Threat Test Dataset is available at https://kilthub.cmu.edu/articles/dataset/Insider_Threat_Test_Dataset/12841247. Locally captured biometric samples and implementation artifacts are available from the corresponding author upon reasonable request.

XIII. ETHICS AND BIOMETRIC DATA PRIVACY

Locally captured facial and fingerprint samples were collected from enrolled project participants/users for research and prototype evaluation after permission was obtained. The samples were treated as sensitive biometric data, organized with identity labels rather than public personal identifiers, and were not released with the manuscript. Access to biometric folders, templates, database records, and IPFS references was restricted to the researcher and authorized supervisors. Locally captured samples were retained only for research verification and prototype evaluation and should be deleted or archived according to the institution's approved data-retention requirements after the study.

REFERENCES

- [1] Hameed, S., et al. Securing distributed systems: A survey on access control techniques for cloud, blockchain, IoT and SDN. *Cyber Security and Applications*. 2023, 1, 100015.
- [2] Shakarami, A., et al. Contemporary advances in multi-access edge computing: A survey of fundamentals, architecture, technologies, deployment cases, security, challenges, and directions. *Journal of Network and Computer Applications*. 2023, 219, 103726.
- [3] Kaur, N. A systematic review on security aspects of fog computing environment: Challenges, solutions and future directions. *Computer Science Review*. 2024, 54, 100688.
- [4] Chandramouli, R., Butcher, Z. A Zero Trust Architecture Model for Access Control in Cloud-Native Applications in Multi-Cloud Environments. NIST Special Publication 800-207. National Institute of Standards and Technology. 2023.
- [5] Alam, T., et al. Differential privacy in edge computing-based smart city applications: Security issues, solutions and future directions. *Array*. 2023, 19, 100293.
- [6] Xu, G., Ren, J., Zhang, Y., Lin, C. A multimodal biometric authentication scheme for mobile edge computing. *Information Sciences*. 2021, 553, 1-15. <https://doi.org/10.1016/j.ins.2020.11.029>
- [7] Yao, H., Zhang, J., Liu, X., Wang, Y. Multimodal biometric authentication by federated learning. *Biomedical Signal Processing and Control*. 2023, 85, 104932. <https://doi.org/10.1016/j.bspc.2023.104932>
- [8] Hu, V. C., Ferraiolo, D. F., Kuhn, R., Schnitzer, A., Sandlin, K., Miller, R., Scarfone, K. Guide to Attribute Based Access Control (ABAC) Definition and Considerations. NIST SP 800-162. National Institute of Standards and Technology. 2015. <https://doi.org/10.6028/NIST.SP.800-162>
- [9] Muhammad, B. A., Garba, M., Gabi, D., Suru, H. U., Argungu, M. S. Dynamic access control at the network edge using an adaptive risk-based access control system (ad-RACs). 2024, 1(2), 12-19.
- [10] Goudarzi, H., Liao, X., Zhang, Q. A review on edge computing architecture and its applications in IoT and AI. *Future Generation Computer Systems*. 2023, 142, 312-328.
- [11] Shi, W., Cao, J., Zhang, Q., Li, Y., Xu, L. Edge Computing: Vision and Challenges. *IEEE Internet of Things Journal*. 2016, 3(5), 637-646. <https://doi.org/10.1109/JIOT.2016.2579198>
- [12] Satyanarayanan, M. The Emergence of Edge Computing. *Computer*. 2017, 50(1), 30-39. <https://doi.org/10.1109/MC.2017.9>
- [13] XenonStack. Edge Computing Solutions and Capabilities. Available from: <https://www.xenonstack.com/enterprise/edge-computing/> (accessed 15 September 2026).
- [14] Jain, A. K., Ross, A., Prabhakar, S. An Introduction to Biometric Recognition. *IEEE Transactions on Circuits and Systems for Video Technology*. 2004, 14(1), 4-20. <https://doi.org/10.1109/TCSVT.2003.818349>
- [15] Ratha, N. K., Connell, J. H., Bolle, R. M. Enhancing security and privacy in biometrics-based authentication systems. *IBM Systems Journal*. 2001, 40(3), 614-634. <https://doi.org/10.1147/sj.403.0614>
- [16] Jain, A. K., Ross, A. Introduction to Biometrics. In *Handbook of Biometrics*. Springer. 2008, pp. 1-22.
- [17] Maltoni, D., Maio, D., Jain, A. K., Prabhakar, S. *Handbook of Fingerprint Recognition*. 2nd ed. Springer. 2009.
- [18] Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L. C. MobileNetV2: Inverted Residuals and Linear Bottlenecks. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2018, pp. 4510-4520. <https://doi.org/10.1109/CVPR.2018.00474>
- [19] Tan, M., Le, Q. V. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. *Proceedings of the 36th International Conference on Machine Learning*. 2019, 97, 6105-6114.
- [20] Nobi, M. Machine learning in access control systems: A survey. *ACM Computing Surveys*. 2022, 55(3), 1-36.
- [21] Rahman, M. S., Li, X., Chen, J. Risk-adaptive access control using reinforcement learning. *IEEE Transactions on Information Forensics and Security*. 2023, 18, 408-421. <https://doi.org/10.1109/TIFS.2022.3218593>
- [22] LeCun, Y., Bengio, Y., Hinton, G. Deep learning. *Nature*. 2015, 521, 436-444. <https://doi.org/10.1038/nature14539>
- [23] Abbas, N., Zhang, Y., Taherkordi, A., Skeie, T. Mobile edge computing: A survey. *IEEE Internet of Things*

- Journal. 2018, 5(1), 450-465.
<https://doi.org/10.1109/JIOT.2017.2770187>
- [24] Jahan, N., Reno, S., Ahmed, M. Securing E-Passport Management Using Private-Permissioned Blockchain and IPFS. In 3rd International Conference on Electrical, Computer and Communication Engineering. 2023.
<https://doi.org/10.1109/ECCE57851.2023.10101496>
- [25] Cao, Q., Shen, L., Xie, W., Parkhi, O. M., Zisserman, A. VGGFace2: A dataset for recognising faces across pose and age. 13th IEEE International Conference on Automatic Face and Gesture Recognition. 2018, pp. 67-74. <https://doi.org/10.1109/FG.2018.00020>
- [26] Ruiz Garcia, S. SOCOFing Dataset. Kaggle. Available from: <https://www.kaggle.com/datasets/ruizgara/socofing> (accessed 15 September 2026).
- [27] Brewer, D. F. C., Nash, M. J. The Chinese Wall Security Policy. Proceedings of the IEEE Symposium on Security and Privacy. 1989, pp. 206-214.
<https://doi.org/10.1109/SECPRI.1989.36295>