AB Journals

British Journal of Computer, Networking and Information Technology

ISSN 2689-5315 • Volume 9, Issue 3, 2026 • pp. 19–39

RESEARCH ARTICLE Received: 1 Sep 2026, Accepted: 26 Oct 2026, Published: 11 Nov 2026

Comparative Analysis of CatBoost and One-Dimensional Convolutional Neural Network Models for Contextual Risk Prediction in Edge-Based Access Control Systems

Paul Ndudiri Ohia[1*]; Friday Eleonu Onuodu, Prof.[1]; Ugochi Adaku Okengwu[2]

[1]Department of Information Technology, University of Port Harcourt, River State, Nigeria.

[2]Department of Computer Science, University of Port Harcourt, River State, Nigeria.

[*]*Corresponding Author's Email: paul.ohia@uniport.edu.ng*

Abstract

Edge-based access-control systems must evaluate identity, context, and policy constraints quickly enough to protect distributed resources without imposing unnecessary authentication. This study compares CatBoost and a one-dimensional convolutional neural network (1D CNN) as complementary models for contextual risk prediction in an enhanced access-control framework. Structured access records were encoded using login time, device trust, known location, biometric confidence, resource sensitivity, and Chinese Wall conflict status. CatBoost supplied interpretable risk estimates, whereas the 1D CNN classified access states as failed, uncertain, or verified. The 1D CNN achieved 84.08% test accuracy on 157 records; the verified class produced 96.77% recall and a 90.91% F1-score. System-level results further showed that adding biometric evidence and policy checks moved some context-only grants to step-up or deny decisions. The findings indicate that neither model is sufficient as a standalone security mechanism. A layered edge pipeline that combines structured risk scoring, contextual classification, identity confidence, and conflict-of-interest enforcement provides a more defensible basis for real-time grant, step-up, and deny decisions.

Keywords: Adaptive Access Control, CatBoost, Contextual Risk Prediction, Edge Computing, One-Dimensional Convolutional Neural Network

How to Cite: Ohia, P. N., Onuodu, F. E., & Okengwu, U. A. (2026). Comparative Analysis of CatBoost and One-Dimensional Convolutional Neural Network Models for Contextual Risk Prediction in Edge-Based Access Control Systems. British Journal of Computer, Networking and Information Technology, 9(3), 19–39. https://doi.org/10.52589/BJCNIT-NSMYAF

© 2026 by the authors. Published by AB Journals. This is an open access article distributed under the terms of the Creative Commons Attribution (CC BY 4.0) license.

INTRODUCTION

Edge computing places computation near users, devices, and protected resources, reducing communication delay and supporting time-sensitive services (Mahbub & Shubair, 2023; Wang et al., 2025). The same distribution that improves responsiveness also complicates access control. Requests arrive from heterogeneous devices, locations, and network conditions, while protected resources differ in sensitivity. A decision mechanism must therefore consider more than a static user role. It must evaluate the circumstances of each request and respond before an unsafe transaction reaches the resource.

Attribute-based access control (ABAC) provides a policy foundation for using subject, object, action, and environmental attributes in authorization decisions, while recent adaptive ABAC research shows how real-time context can extend fixed policy evaluation (Baccouri & Abdellatif, 2025; Kalaria et al., 2024). Contemporary distributed systems consequently combine policy enforcement with contextual assessment, continuous verification, or risk-aware authentication (Chen et al., 2024; Golightly et al., 2023). Recent zero-trust Internet of Things research likewise treats access as an ongoing decision informed by interactions among users, devices, and protected objects rather than as a one-time perimeter check (Ameer et al., 2024).

Machine learning can support this adaptive process by detecting patterns across historical access records. Quan et al. (2023) demonstrate that CatBoost can learn attribute-based access-control policies from access logs while retaining rule extraction and policy-quality analysis. Neural approaches provide a complementary capacity to learn nonlinear interactions among contextual variables; recent graph-based access-control research also shows that learned representations can improve decisions when organizational and user-resource relationships are encoded explicitly (Yin et al., 2024).

A one-dimensional convolutional neural network (1D CNN) can therefore be evaluated as a compact contextual classifier, provided that feature order, validation, and class-level errors are reported transparently.

The empirical gap is not simply a shortage of access-control models. Current research increasingly supports adaptive and context-aware controls, but fewer studies show how an interpretable tabular learner, a neural contextual classifier, biometric confidence, and a conflict-of-interest policy can be coordinated at the point of access. The systematic mapping by Matter, Martins, and Barbosa (2026) identifies access control as a major application of context-aware security and machine learning, while Xue et al. (2025) emphasise the continuing need for low-latency, resource-conscious, and scalable controls at the edge. These findings justify a comparison that evaluates model roles as well as predictive performance.

This study therefore examines how CatBoost and a 1D CNN contribute to contextual risk prediction in an edge-based access-control system.

It asks three questions: (1) how effectively does the 1D CNN classify failed, uncertain, and verified contextual states; (2) how do CatBoost and the 1D CNN differ in input, output, interpretability, and operational role; and (3) how does combining contextual prediction with biometric confidence and Chinese Wall enforcement alter final access decisions? The contribution is a layered interpretation in which CatBoost estimates structured risk, the 1D CNN classifies contextual states, biometric scores strengthen identity evidence, and the Chinese Wall policy prevents conflicts of interest.

The remainder of the paper is organised as follows. Section 2 synthesises prior work and explains the theoretical and conceptual framework. Section 3 describes the data, preprocessing, models, and evaluation procedure. Section 4 presents the model-level and system-level findings. Section 5 relates those findings to recent literature and states the contribution. Sections 6 and 7 present implications, limitations, future research, and the conclusion.

LITERATURE REVIEW AND THEORETICAL FRAMEWORK

This section reviews previous studies relevant to edge-based contextual access control, adaptive authentication, machine-learning-supported risk assessment, and conflict-of-interest enforcement. It also establishes the theoretical and conceptual foundations of the proposed system. The review focuses on how contextual information, biometric evidence, machine-learning models, and deterministic security policies can be combined to support secure and responsive access-control decisions. Existing studies provide important foundations for this work but also reveal gaps concerning the integration of structured risk estimation, neural contextual classification, biometric verification, and explicit policy enforcement within a unified edge-based framework.

Edge-Based and Context-Aware Access Control

Edge environments require authorization mechanisms that are both responsive and sensitive to changing conditions. Mahbub and Shubair (2023) show that multi-access edge computing supports low-latency applications by moving services towards end devices, while Golightly et al. (2023) identify access-control integration as a continuing challenge across cloud, blockchain, Internet of Things, and software-defined networking environments. More recent edge-IAM research moves authorization and identity data closer to the point of use: Wang et al. (2025) propose local authentication and authorization with context-driven synchronization for dynamic, multi-authority IoT services. These studies indicate that edge security must be evaluated through latency, policy enforcement, auditability, and resource use as well as predictive accuracy.

Recent studies extend this argument from static authorization to adaptive control. Kalaria et al. (2024) distribute context collection and XACML-based policy decisions across fog and cloud nodes, reporting low processing overhead while adapting to changing IoT behaviour. Chen et al. (2024) organize risk-aware authentication into multiple stages that progressively adjust controls according to context, resource sensitivity, and user authenticity. Ameer et al. (2024) similarly propose a zero-trust architecture that recognizes multiple interaction types among users, devices, and objects, while Ali et al. (2024) combine zero-trust authentication, biometrics, and trust-aware task offloading in multi-access edge computing. Collectively, these studies support continuous, evidence-based decisions, but they differ in how risk is represented, explained, and enforced.

Attribute-Based and Risk-Adaptive Access Control

The first theoretical foundation is ABAC. In ABAC, authorization follows from evaluating subject, object, action, and environmental attributes against policy rules. Recent work has shifted from treating those attributes as static inputs to updating policies as context changes. Baccouri and Abdellatif (2025) use real-time policy evaluation and decision-tree processing to support adaptive, scalable ABAC, while Quan et al. (2023) show how CatBoost can derive and refine ABAC rules from historical access logs. This perspective explains why login time,

device trust, location, resource sensitivity, and conflict status are legitimate decision inputs, while also distinguishing contextual evidence from direct proof of identity.

The second foundation is risk-adaptive access control. Rather than mapping every valid identity to the same response, a risk-adaptive mechanism changes the strength of control according to estimated risk. Low-risk evidence can support a grant, ambiguous evidence can trigger step-up verification, and high-risk evidence can cause denial. Chen et al. (2024) provide a multi-stage implementation of this graduated logic.

Halgamuge and Niyato (2025) broaden the same principle to dynamic IoT edge policies by combining risk assessment, conflict resolution, compliance analysis, and AI-driven adaptation. Xue et al. (2025) further show why flexible, resource-conscious, low-latency controls remain essential in IoT edge environments.

The Chinese Wall policy adds a deterministic conflict-of-interest control layer. It organizes sensitive information into protected datasets within conflict-of-interest classes and prevents a user who has accessed one dataset from accessing a competing dataset in the same class (Tu et al., 2025). Unlike the learned models, this policy is not a probability estimate and can therefore override a favourable contextual or biometric result. In the proposed access-control pipeline, the Chinese Wall rule prevents an otherwise valid request from being granted whenever the request violates a recorded conflict-of-interest relationship. This separation is consistent with adaptive-policy architectures that preserve explicit policy and compliance checks alongside learned risk signals (Baccouri & Abdellatif, 2025; Halgamuge & Niyato, 2025).

Machine Learning for Contextual Risk

CatBoost and a 1D CNN offer different inductive biases. CatBoost builds an ensemble of decision trees and is well suited to mixed tabular attributes. Its relevance to access control is supported by Quan et al. (2023), who use it to learn ABAC policies from access logs and extract rules for review. A 1D CNN applies learnable filters across an ordered feature vector and can represent local interactions among encoded contextual variables, but the assigned order becomes a modelling assumption. Recent work on heterogeneous graph learning demonstrates a broader trend towards learned access-control representations that incorporate relationships beyond isolated attributes (Yin et al., 2024). These findings support comparative evaluation, but not the assumption that one model family is automatically superior.

The recent literature reveals three unresolved gaps. First, deployment-oriented frameworks emphasize adaptive policy enforcement and edge locality, whereas learning-oriented studies often emphasize classification or representation quality (Kalaria et al., 2024; Wang et al., 2025; Yin et al., 2024). Second, Matter, Martins, and Barbosa (2026) said that machine learning remains under used in context-aware access control and that datasets, taxonomies, and reproducible comparisons are still limited. Third, the Anubis model shows that contextual machine learning can revise unsafe static permissions, but its synthetic evaluation also illustrates the need for deployment-level validation (Matter, Schardosim, et al., 2026). The present study addresses these gaps by comparing the operational roles of two model families and examining how their outputs interact with biometric confidence and an explicit conflict rule.

Conceptual Framework

Figure 1 presents the author-generated conceptual framework. An access request supplies subject, device, resource, and environmental evidence. CatBoost converts structured evidence into a risk estimate, while the 1D CNN classifies the contextual state. Biometric scores add identity confidence, and attribute-based plus conflict-of-interest rules impose policy constraints. The edge decision engine fuses these signals and releases one of three outcomes: grant, step-up, or deny. This layered structure reflects current research that combines edge-local authorization, adaptive context evaluation, staged authentication, and explicit policy enforcement rather than relying on a single model output (Chen et al., 2024; Kalaria et al., 2024; Wang et al., 2025).

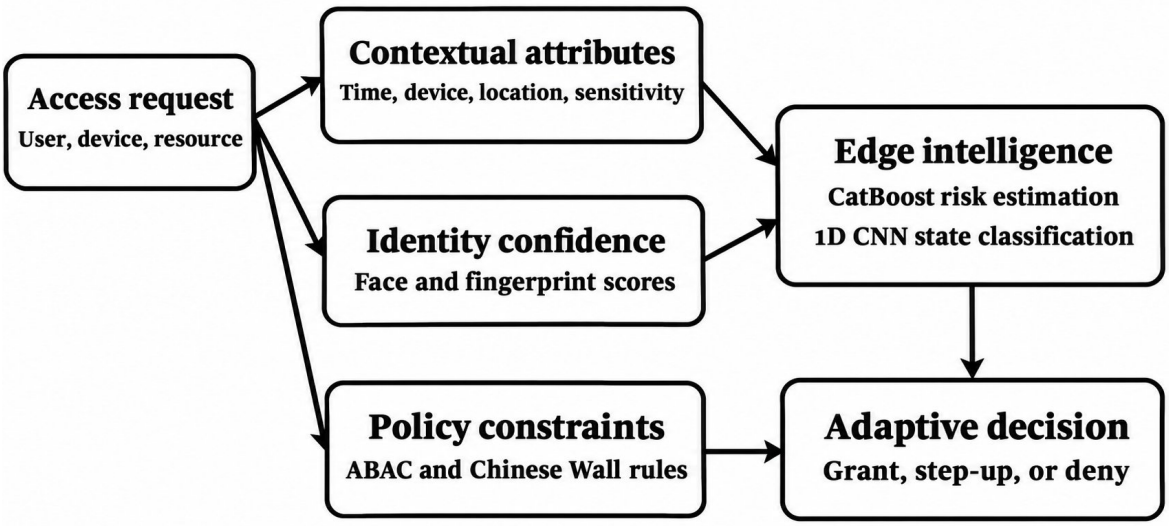


Figure 1: Conceptual framework for the layered edge-based access-control decision process.
Source: Authors' conceptualisation (2026).

Figure 1 conceptual framework shows that authorization is a layered process in which contextual predictions are checked against biometric identity evidence and mandatory policy constraints before a final decision is released.

METHODOLOGY

This section describes the research design, data sources, data integration procedures, missing-value handling, preprocessing techniques, feature-engineering process, proposed system architecture, model configuration, and evaluation measures used in the study. The methodology was designed to compare the complementary roles of CatBoost and a one-dimensional convolutional neural network within an edge-based contextual access-control framework. It also explains how contextual attributes, biometric confidence scores, resource information, and Chinese Wall conflict indicators were combined to support grant, step-up verification, or denial decisions.

Research Design and System Setting

The study used a comparative system-evaluation design. It examined existing experimental outputs from an enhanced access-control application implemented with Python and Flask, with

My Structured Query Language (MySQL) used for structured records and the InterPlanetary File System (IPFS) used for selected decentralised artefacts. No survey respondents or key informants were recruited for this manuscript. The units of analysis were access-request records and the resulting model or policy decisions. This clarification addresses the sampling issue by distinguishing machine-learning records from human respondents.

The analysis operated at two levels. At model level, the 1D CNN was assessed with accuracy, precision, recall, F1-score, support, learning curves, and a confusion matrix. CatBoost was assessed from its structured risk output and its role in the existing decision path because a complete, directly comparable CatBoost classification report was not retained in the supplied experiment artefacts. At system level, contextual-only decisions were compared with decisions that also used biometric confidence and Chinese Wall enforcement.

Data Sources and Integration

The contextual dataset integrated three forms of evidence generated within the access-control environment. First, application access logs supplied request-level information, including access time, device status, location status, the requested resource, and the recorded decision. Second, adapted security-event records provided contextual activity fields relevant to user and device behaviour. Third, outputs from the biometric verification components supplied face and fingerprint confidence scores, while the policy component supplied resource-sensitivity and Chinese Wall conflict indicators. The unit of integration was an access-request record, so that contextual, biometric, resource, and policy evidence described the same decision event.

After integration, each observation represented one access attempt and contained a common set of fields for model development. The retained one-dimensional convolutional neural network (1D CNN) test output contained 157 records distributed across failed, uncertain, and verified classes. Records were allocated to independent training, validation, and test partitions during the original experiment. The training partition supported parameter learning, the validation partition supported model monitoring, and the test partition was reserved for the final class-level evaluation reported in this paper.

Table 1: Contextual features used in the one-dimensional convolutional neural network

Login hour	Represents the time of the access attempt and supports detection of unusual periods.	Environment
Trusted-device status	Indicates whether the request originated from a recognised device.	Subject/device
Known-location status	Shows whether the access location matches expected behaviour.	Environment
Face score	Provides confidence from facial verification.	Identity
Fingerprint score	Provides confidence from fingerprint verification.	Identity
Resource sensitivity	Represents the security importance of the requested resource.	Object
Chinese Wall conflict status	Indicates whether prior access creates a prohibited conflict of interest.	Policy

Note: Features were derived from the supplied system records and experimental artefacts. Source: Authors' study (2026).

Table 1 above summarizes the retained contextual predictors, their operational meanings, and the evidence category represented by each field. The table therefore demonstrates that the classifier combines environmental, identity, resource, and policy evidence rather than relying on a single data type. The feature set covers environmental, identity, resource, and policy evidence, enabling the model to assess an access request from several security perspectives.

Handling Missing Values

Missing values were assessed before encoding, scaling, or partition-specific transformation. Each record was checked across the seven required predictors: login hour, trusted-device status, known-location status, face score, fingerprint score, resource sensitivity, and Chinese Wall conflict status. Missingness was treated as a data-quality and security issue because an absent biometric score, device indicator, resource classification, or policy value can change the meaning of an access request and may conceal a high-risk condition.

Records missing a required identity, context, resource, policy, or target field were excluded from model training when the value could not be recovered reliably from the corresponding access log. No class label was imputed. Where a value could be reconstructed deterministically from the same event for example, login hour from a valid timestamp or conflict status from the associated policy record the reconstructed value was used and the source record was retained. This conservative procedure avoided introducing synthetic security evidence and ensured that only complete, internally consistent observations entered the modelling pipeline.

Data Preprocessing

Preprocessing was performed at the access-record level before model training. Records were checked for missing, invalid, duplicated, or internally inconsistent values. Time fields were converted to a consistent representation, categorical status fields were encoded numerically, and continuous confidence values were converted to a common numeric format. Records that could not be reconciled with the required feature schema were excluded rather than imputed because an incorrect device, location, biometric, or policy value could change the security meaning of an access event.

The target labels were standardised as failed, uncertain, or verified. Numerical predictors were scaled using parameters estimated from the training data, and the same transformation was applied to the validation and test partitions to prevent information leakage. The processed feature matrix was then reshaped into the three-dimensional input required by the Conv1D layer: observations, ordered features, and one channel. Partition boundaries were maintained throughout preprocessing so that test observations did not influence model fitting or validation decisions.

Feature Engineering

Feature engineering translated raw access events into seven security-relevant predictors. Login hour was extracted from the request timestamp to represent temporal context. Trusted-device and known-location fields were encoded as status indicators. Face and fingerprint outputs were retained as continuous confidence scores so that the model could distinguish strong, weak, and conflicting identity evidence. Resource sensitivity was encoded to represent the protection level of the requested object, while Chinese Wall conflict status was represented as a policy indicator that identified prohibited relationships created by prior access.

The engineered variables were ordered consistently for every observation before they were passed to the 1D CNN. This ordering allowed the convolutional filters to learn local relationships among temporal, device, location, biometric, resource, and policy evidence. The same variables were also available to the structured risk pathway, although CatBoost treated them as tabular attributes rather than as positions in a convolutional sequence. No outcome-derived variable was included among the predictors; the final access state remained the target used for supervised learning.

Proposed System Architecture

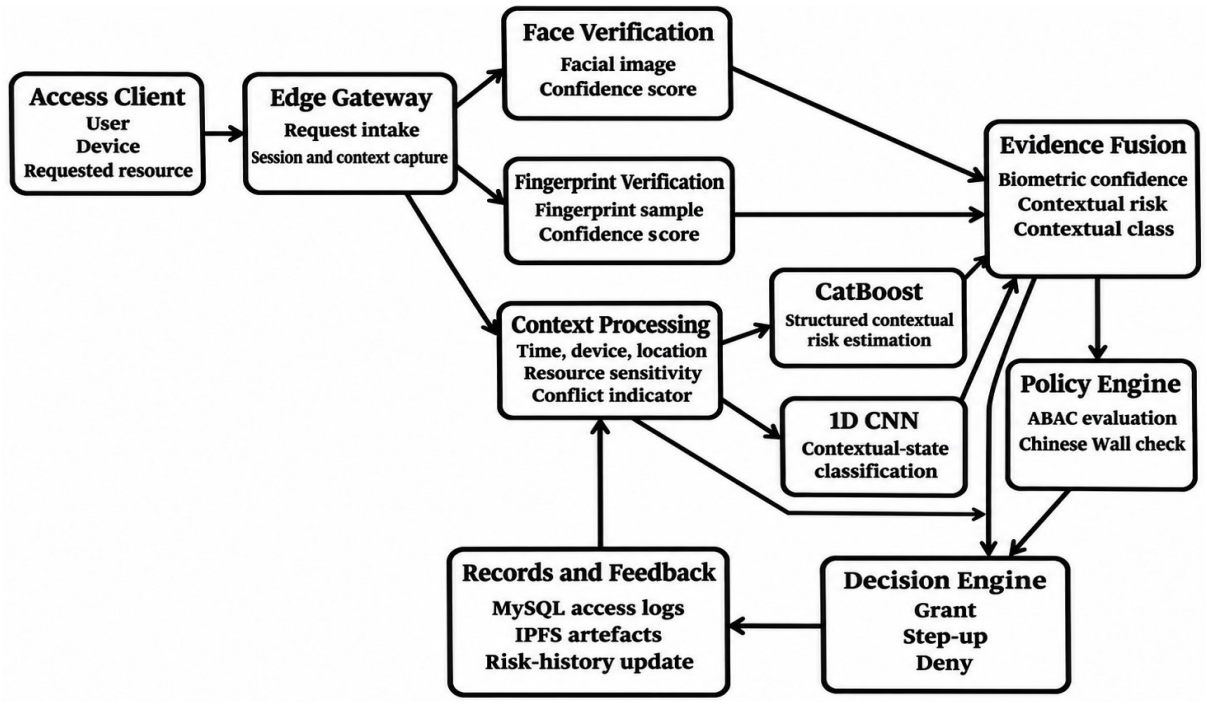


Figure 2: Proposed architecture of the edge-based adaptive access-control system.
Source: Authors' design (2026).

The proposed architecture integrates identity verification, contextual risk analysis, machine-learning classification, and deterministic policy enforcement within an edge-processing pipeline. As shown in Figure 2 above, an access client submits the user, device, and requested-resource information to the edge gateway. The gateway captures the session context and directs facial and fingerprint samples to their respective verification components while contextual attributes are prepared for the CatBoost and 1D CNN paths. CatBoost produces a structured risk estimate, whereas the 1D CNN produces a failed, uncertain, or verified contextual state.

The evidence-fusion component combines biometric confidence, the contextual risk estimate, and the contextual class. The policy engine evaluates attribute-based access-control requirements and applies the Chinese Wall conflict rule before the decision engine releases a grant, step-up, or deny outcome. Each outcome is written to the access log, while selected artefacts and risk-history references are retained through MySQL and IPFS. The recorded history is returned to the contextual path so that subsequent requests can be evaluated against prior behavior and policy relationships. Figure 2 therefore shows that learned risk estimates support the decision, while deterministic identity and conflict checks retain authority over the final outcome. The architecture separates predictive scoring from deterministic enforcement,

so a favourable model output cannot override failed identity verification or a Chinese Wall conflict.

Model Configuration

Incomplete or inconsistent contextual records were removed during the original preparation stage. Categorical attributes were converted to numerical values, continuous features were scaled, and the feature matrix was reshaped into the three-dimensional form required by a Conv1D layer: observations, ordered features, and one channel. Labels represented failed, uncertain, or verified contextual states. The architecture comprised one-dimensional convolution, batch normalisation, pooling, flattening, dropout, and dense classification layers.

CatBoost processed structured contextual attributes and produced risk information for the adaptive decision path. The 1D CNN processed ordered feature vectors and produced a three-class contextual state.

Table 2: Comparative model design and evaluation roles

Model family	Gradient-boosted decision trees	One-dimensional convolutional neural network
Primary input	Structured and categorical context	Scaled ordered feature vector
Primary output	Contextual risk estimate	Failed, uncertain, or verified state
Evaluation evidence	Risk distribution and decision association	Accuracy, class metrics, curves, and confusion matrix
Operational strength	Efficient tabular learning and auditability	Learns local interactions across encoded features
Main limitation	No retained like-for-like classification report	Lower interpretability and order-sensitive encoding

Note: The comparison describes the roles supported by the supplied experimental outputs. Source: Authors' study (2026).

Table 2 above summarizes why the two outputs are complementary rather than strictly interchangeable. CatBoost offers stronger tabular interpretability, whereas the 1D CNN provides a learned representation of interactions across the encoded vector. The comparison explains why the two outputs are interpreted as complementary evidence rather than competing versions of the same result. CatBoost and the 1D CNN contribute different forms of evidence; their value lies in complementary risk estimation and contextual classification rather than direct interchangeability.

Evaluation Measures and Decision Logic

Accuracy measured the proportion of correct predictions. Precision measured how often a predicted class was correct, recall measured how much of the true class was recovered, and the F1-score provided their harmonic mean. Macro averages weighted each class equally, while weighted averages reflected class support. The confusion matrix was used to identify which contextual states were most frequently confused.

The system converted model and policy evidence into three decisions. A grant required acceptable contextual risk, sufficient identity confidence, and no Chinese Wall conflict (Tu et

al., 2025). Step-up verification was used when evidence was incomplete or uncertain but not clearly unsafe. Denial followed failed identity evidence, high contextual risk, or a policy conflict. This rule structure preserves the separation between probabilistic prediction and deterministic policy enforcement. The complete decision mapping is presented in Appendix A (Table A1).

RESULTS AND FINDINGS

This section presents the experimental results and findings from the comparative evaluation of CatBoost and 1D Convolutional Neural Network (1D-CNN) models for contextual risk prediction in edge-based access control systems. The evaluation focuses on the predictive performance, classification accuracy, precision, recall, F1-score, computational efficiency, and suitability of both models for deployment in resource-constrained edge computing environments. The findings provide empirical evidence of the strengths, limitations, and relative effectiveness of the two models in predicting contextual security risks and supporting adaptive access control decisions. These results form the basis for the subsequent discussion and interpretation.

Structured Risk and Identity Confidence Outputs

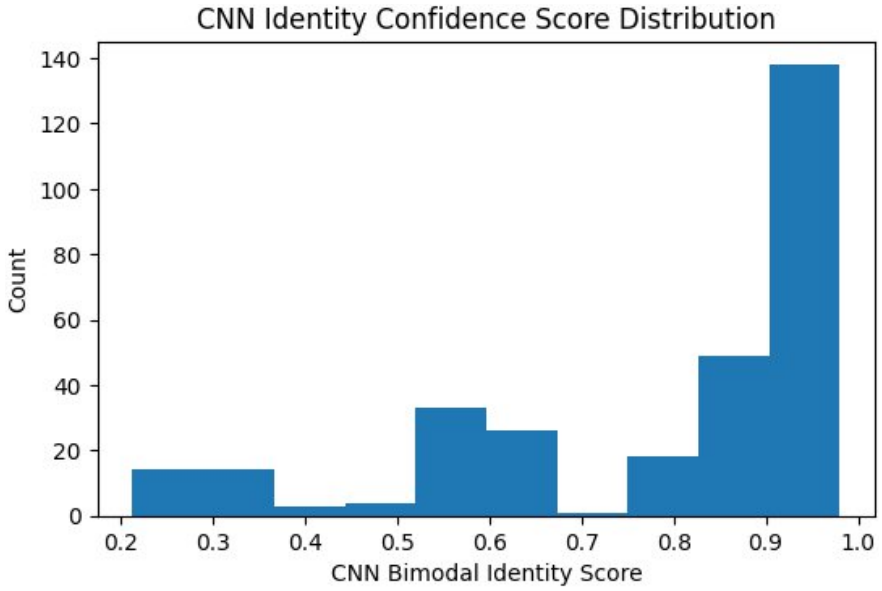


Figure 3: Distribution of combined convolutional neural network identity-confidence scores.
Source: Authors' experimental results (2026).

The retained score distribution in Figure 3 is concentrated towards high confidence, although smaller groups appear at medium and low values. The distribution indicates that most observed requests produced strong combined identity evidence, but a non-trivial minority required additional scrutiny. Most scores cluster near the high-confidence end, but the lower-score groups justify retaining step-up verification for requests with weaker identity evidence

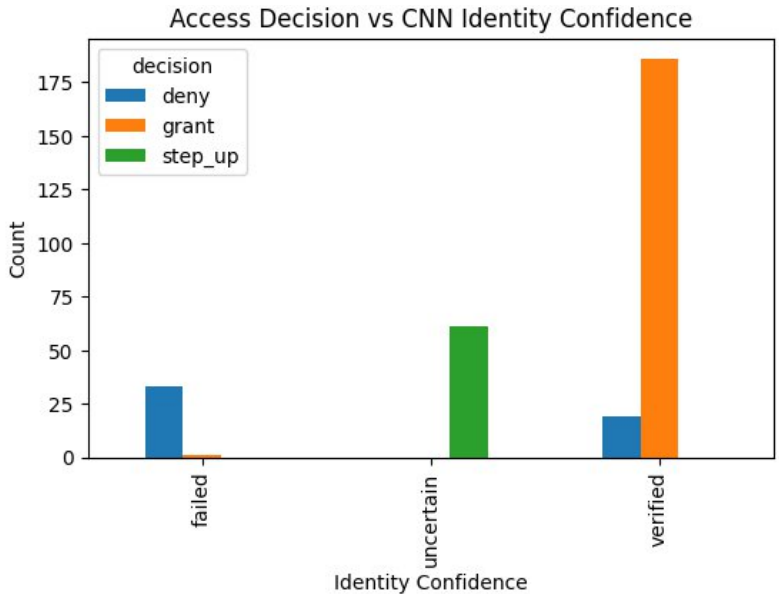


Figure 4: Access decision by combined identity-confidence class.
Source: Authors' experimental results (2026).

Figure 4 confirms the operational effect: verified cases were predominantly granted, uncertain cases were directed to step-up verification, and failed cases were primarily denied. These figures support graduated control rather than a binary accept-or-reject rule.

The decision pattern confirms graduated control, verified identities are mainly granted, uncertain identities are challenged, and failed identities are mainly denied.

One-Dimensional Convolutional Neural Network Performance

The 1D CNN achieved 84.08% accuracy on 157 test records.

Table 3: One-dimensional convolutional neural network classification results

Failed	86.11%	75.61%	80.52%	41
Uncertain	80.39%	75.93%	78.10%	54
Verified	85.71%	96.77%	90.91%	62
Accuracy	-	-	84.08%	157
Macro average	84.07%	82.77%	83.17%	157
Weighted average	83.99%	84.08%	83.79%	157

Note: Percentages were calculated from the retained test classification report. Source: Authors' experimental results (2026).

Table 3 shows that the verified class performed best, with 85.71% precision, 96.77% recall, and a 90.91% F1-score. The model therefore recovered nearly all verified records. The failed class achieved 86.11% precision but lower recall of 75.61%, which means that some failed records were assigned to another class. The uncertain class produced the lowest F1-score, 78.10%, confirming that ambiguity was the most difficult condition to separate.

The verified class is the most reliably identified, while uncertainty remains the most difficult condition and the principal target for model improvement.

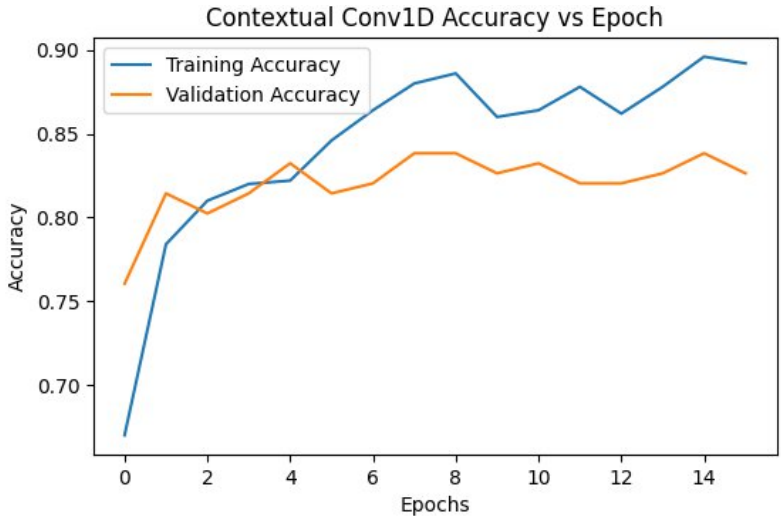


Figure 5: Training and validation accuracy of the contextual 1D CNN across epochs.
Source: Authors' experimental results (2026).

Training accuracy continues to rise after validation accuracy begins to level off, indicating that additional training yields limited generalisation benefit.

Figure 5 shows the training and validation accuracy curves, while Figure 6 shows the corresponding loss curves. Training and validation accuracy rose quickly during the early epochs, after which training accuracy continued to improve while validation accuracy fluctuated around a lower plateau. The corresponding loss curves declined initially, but validation loss later varied and ended above its minimum. This pattern suggests modest overfitting and supports the use of dropout, early stopping, and threshold calibration in future experiments.

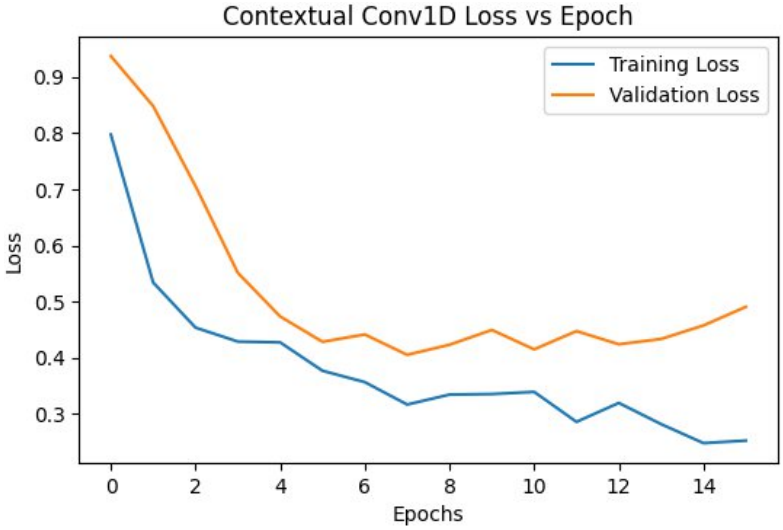


Figure 6: Training and validation loss of the contextual 1D CNN across epochs.
Source: Authors' experimental results (2026).

The late separation between training and validation loss provides further evidence of modest overfitting and supports early stopping and regularisation.

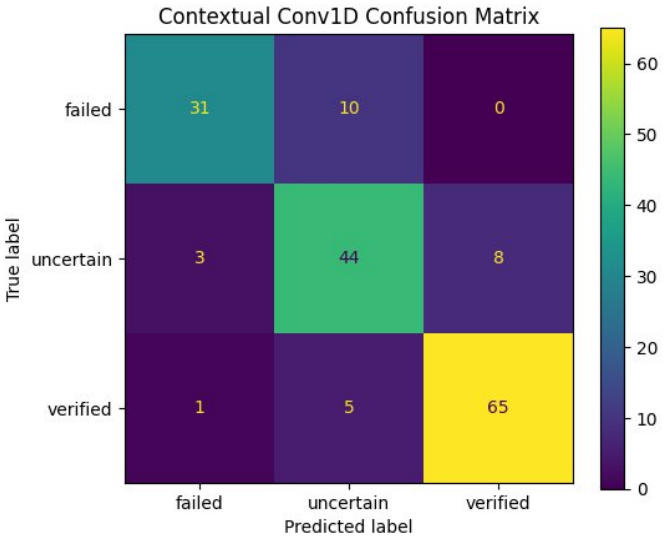


Figure 7: Confusion matrix for contextual 1D CNN classification.
Source: Authors' experimental results (2026).

Figure 7 explains the class metrics. Of the 41 failed records, 31 were classified correctly and 10 were classified as uncertain. Of the 55 records represented in the uncertain row of the matrix, 44 were classified correctly, 3 were assigned to failed, and 8 were assigned to verified. The verified class was the most stable: 65 records were correctly classified, with only six errors. The matrix therefore locates the principal weakness at the boundaries around uncertainty rather than in the recognition of clearly verified cases.

Most errors involve the uncertain class, whereas verified requests are recognised more consistently, locating the main weakness at class boundaries.

System-Level Decision Effects

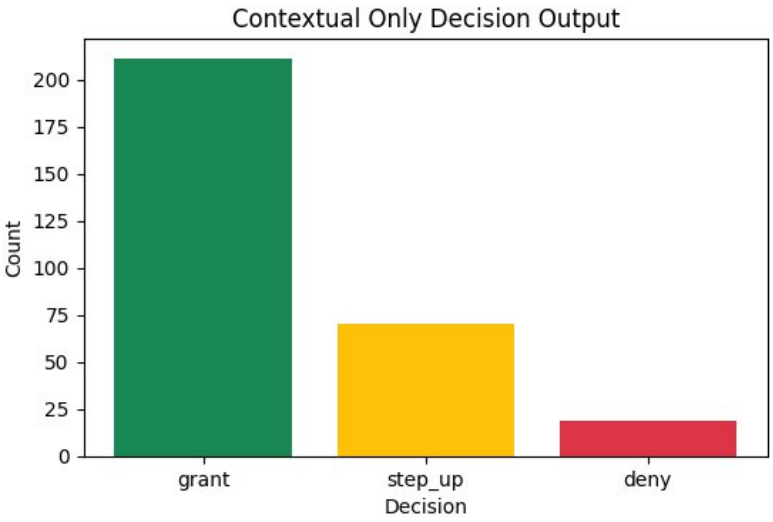


Figure 8: Distribution of contextual-only access decisions.
Source: Authors' experimental results (2026).

Figure 8 shows that the contextual-only stage produced a large majority of grants, approximately 70 step-up decisions, and fewer than 20 denials. After biometric evidence was

introduced, Context alone produces a permissive distribution dominated by grants, showing why contextual plausibility should not be treated as final authorization.

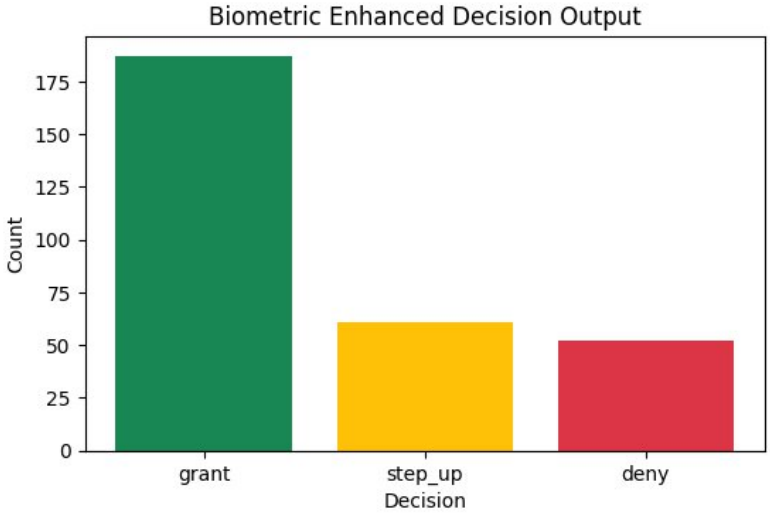


Figure 9: Distribution of access decisions after biometric enhancement.
Source: Authors' experimental results (2026).

Figure 9 shows fewer grants and more denials. The step-up count also declined slightly. The change is security-relevant because it demonstrates that apparently acceptable contextual conditions did not automatically override weak identity evidence. Adding biometric evidence reduces grants and increases denials, demonstrating the security effect of direct identity verification.

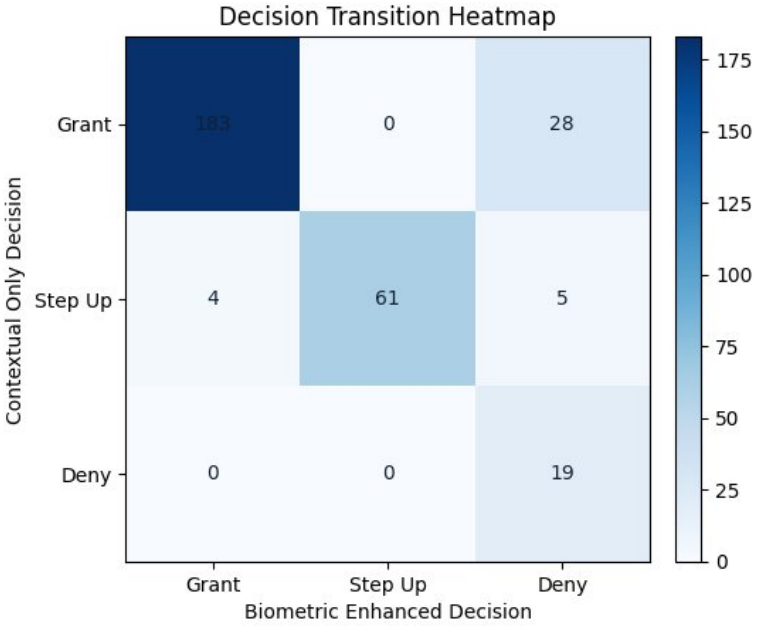


Figure 10: Transition from contextual-only decisions to biometric-enhanced decisions.
Source: Authors' experimental results (2026).

The transition matrix in Figure 10 provides the exact movement between decision categories. Of the requests initially granted, 183 remained grants and 28 moved to denial. No initial grant

moved to step-up in the retained output. Among contextual step-up cases, 61 remained step-up, 4 moved to grant, and 5 moved to denial. All 19 contextual denials remained denials. The most important effect is therefore the 28 grants that became denials when stronger identity and policy evidence was considered. The dominant transition is from contextual grant to biometric-enhanced denial, confirming that stronger identity and policy evidence corrects some initially permissive outcomes.

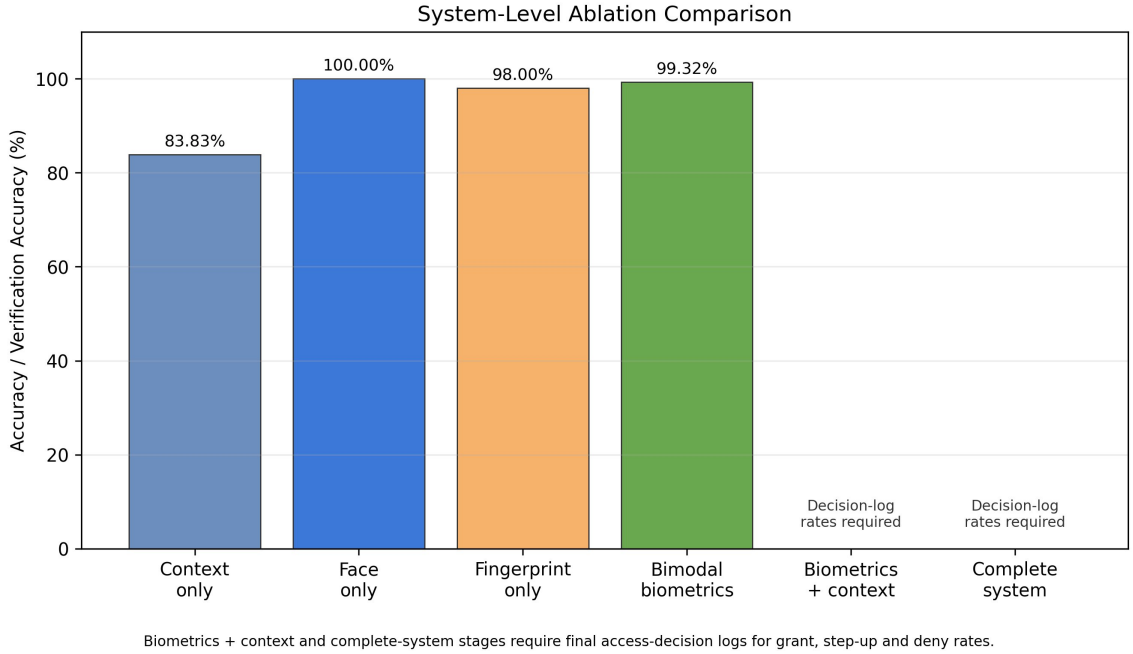


Figure 11: System-level ablation summary for available component evaluations.
Source: Authors' experimental results (2026).

Figure 11 provides the available ablation summary. Context alone achieved 83.83%, while the face-only and fingerprint-only components produced 100.00% and 98.00% verification accuracy in their respective retained evaluation outputs. Bimodal biometrics reached 99.32%. Because complete decision-log rates were not retained for the combined biometrics-plus-context and complete-system stages, those two bars are intentionally not reported. The figure should therefore be interpreted as a component summary rather than a like-for-like benchmark across every system configuration.

The available component results are strong, but the missing combined-system values prevent a complete like-for-like ablation comparison.

Comparative Interpretation

Table 4 connects the model-level results to the system design by comparing the distinct operational contribution of each model.

Table 4: Interpretation of complementary model contributions

Risk representation	Produces a structured risk estimate from tabular context.	Produces a failed, uncertain, or verified state.
Evidence retained	Risk-score distribution and decision association.	84.08% accuracy plus class-level metrics.

Interpretability	Tree-based structure is more amenable to feature attribution.	Requires post-hoc explanation of filter and feature effects.
Edge role	Fast preliminary scoring and auditable risk ranking.	Contextual state classification from the encoded vector.
Best combined use	Supplies risk evidence to the decision engine.	Supplies a learned contextual confidence state.

Note: The table distinguishes supported operational roles; it does not imply identical training targets. Source: Authors' study (2026).

Table 4 clarifies the central finding. CatBoost and the 1D CNN are not fully interchangeable because the retained experiments assign them different targets and report different evidence. CatBoost supports structured risk estimation and auditability. The 1D CNN supplies a contextual state with measurable class performance. Their defensible comparison is therefore functional: each contributes a different form of evidence to the same adaptive decision process.

The comparison supports a layered design in which CatBoost improves structured-risk transparency and the 1D CNN supplies measurable contextual classification.

DISCUSSION

The 84.08% 1D CNN accuracy demonstrates that a compact neural model can learn useful contextual patterns from structured access records. Its strongest result was verified-class recall, whereas uncertain and failed states were more frequently confused. This asymmetry matters operationally. A high verified recall reduces unnecessary challenges for legitimate requests, but lower failed recall can allow risky cases to enter the uncertainty path. The final decision engine should therefore treat the neural class as one piece of evidence rather than as unconditional authorization.

This result aligns with recent adaptive-access-control literature. Chen et al. (2024) argue for multiple stages that respond to context, resource sensitivity, and authenticity. The observed transition from contextual grants to biometric-enhanced denials provides empirical support for that staged logic. It also agrees with Ameer et al. (2024), whose zero-trust architecture recognises that different user-to-object and device-to-object interactions require explicit control. In the present system, device and location context can raise confidence, but biometric evidence and conflict rules can still reverse the initial result.

The findings also refine the role of machine learning at the edge. Xue et al. (2025) identify latency, scalability, flexibility, and resource consumption as persistent challenges for edge access control. CatBoost is attractive in that setting because tree ensembles handle heterogeneous tabular attributes efficiently and are easier to audit. The 1D CNN offers richer interaction learning, but its lower transparency and modest overfitting require monitoring. The models therefore align rather than compete: a tree-based risk estimate can provide a fast and interpretable signal, while the neural classifier can capture interactions that are difficult to specify manually.

The study diverges from approaches that present a single predictive score as the final authority. The conflict-of-interest rule remains deterministic and can deny an otherwise favourable request when a prohibited relationship exists. That separation is consistent with current adaptive frameworks that retain policy, conflict-resolution, and compliance components

alongside AI-supported risk assessment (Baccouri & Abdellatif, 2025; Halgamuge & Niyato, 2025). Similarly, ABAC defines the policy vocabulary that determines which attributes are relevant, while machine learning estimates patterns within those attributes. The contribution is therefore an integration of policy and prediction, not a substitution of one for the other.

A second contribution is methodological restraint. The retained CatBoost evidence does not contain a classification report that is directly comparable with the 1D CNN results. Reporting a numerical winner would therefore be misleading. Instead, the revised analysis distinguishes the models' targets, outputs, and decision roles. This interpretation responds to the broader reproducibility concern in context-aware security identified by Matter, Martins, and Barbosa (2026): comparisons are meaningful only when the data split, target definition, metrics, and deployment constraints are aligned.

Implications for Research and Practice

For research, the results support evaluation frameworks that separate predictive performance from policy effectiveness. Future comparisons should train CatBoost, Random Forest, XGBoost, support vector machines, and the 1D CNN on the same records, target labels, and partitions. They should also report calibration, inference latency, memory use, feature attribution, and robustness to adversarial or missing context. Such a design would show whether improved accuracy remains valuable under the computational constraints of an edge node.

For practice, the findings support a layered decision service. An organisation can use CatBoost for rapid structured-risk scoring, a compact 1D CNN for contextual state classification, biometric verification for direct identity evidence, and explicit policy checks for conflicts that models must not override. Step-up authentication is particularly important because it creates a safe middle response when contextual and identity signals disagree. This reduces the pressure to choose between permissive granting and blanket denial.

Administrators should log each component output and the rule that produced the final decision. Logging enables audit, supports appeal or incident investigation, and supplies data for monitoring drift. Human oversight remains necessary when thresholds, feature encodings, or resource classifications change. Recent risk-aware and zero-trust studies similarly show that adaptation must be governed by coherent policies rather than by unconstrained model output (Chen et al., 2024; Ameer et al., 2024).

Limitations and Future Research

The study has four limitations. First, the retained artefacts do not provide a complete CatBoost classification report on the same test partition used for the 1D CNN. The analysis can therefore compare model roles but cannot establish statistical superiority. Second, the contextual test set is relatively small and reflects a limited application environment. Third, the ordering of tabular features in a 1D CNN is a modelling choice; different orders may change the learned local relationships. Fourth, the system-level graphs summarise retained experimental outputs and do not include confidence intervals, repeated runs, or an external deployment trial.

CONCLUSION

This study compared CatBoost and a 1D CNN within an edge-based contextual access-control framework. The 1D CNN achieved 84.08% test accuracy and performed most strongly on verified requests, but uncertainty remained the main source of error. CatBoost supplied a complementary, more interpretable structured-risk signal. System-level analysis showed that biometric confidence and Chinese Wall enforcement changed some context-only grants to denials, confirming that contextual plausibility is not equivalent to verified identity or policy compliance.

These findings align with recent research that advocates staged, context-sensitive authentication and edge-local policy enforcement. In particular, the movement of some context-only grants to denial after biometric and policy checks supports the graduated decision logic described by Chen et al. (2024) and the adaptive fog/edge enforcement approaches of Kalaria et al. (2024) and Wang et al. (2025). The results also agree with Tu et al. (2025) that conflict-of-interest constraints require explicit policy enforcement. They diverge, however, from learning-centred approaches that evaluate access control primarily through a single predictive output: the present findings show that contextual classification can inform—but should not independently determine—the final authorization decision. Because CatBoost and the 1D CNN were not evaluated on identical targets and partitions, this divergence is architectural rather than evidence of statistical superiority.

The main conclusion is that secure edge authorization should not delegate the final decision to a single predictive model. A layered pipeline can combine tabular risk estimation, neural contextual classification, biometric identity evidence, and deterministic policy constraints. This design supports timely grant, step-up, and deny responses while retaining an auditable separation between learned uncertainty and mandatory rules.

Future Work

Future research should address the following priorities through a consistent, staged evaluation programme.

1. **Dataset reconstruction.** Reconstruct a single versioned dataset and use stratified, leakage-controlled partitions for all models.
2. **Edge-deployment efficiency.** Measure inference latency, memory footprint, and energy use on the intended edge hardware.
3. **Model interpretability.** Use feature-attribution and ablation tests to explain which contextual variables drive each decision.

Declarations

Author Contributions

Paul Ndudiri Ohia: conceptualisation, software, investigation, data curation, visualisation, and writing - original draft. Friday Eleonu Onuodu: supervision, methodology, validation, and writing - review and editing. Ugochi Adaku Okengwu: supervision, validation, and writing - review and editing.

Funding

This research received no external funding.

Ethics Approval and Informed Consent

This manuscript reports a secondary analysis of access-control experiment outputs supplied by the authors. No new participant recruitment or biometric data collection was conducted for this comparative manuscript. Any underlying locally captured biometric data remain subject to the ethics and consent arrangements of the originating project, which the authors should verify against their institutional record before submission.

Data Availability Statement

The contextual records and locally captured biometric artefacts are not publicly available because they may contain security-sensitive and biometric information. De-identified data and implementation details may be requested from the corresponding author, subject to privacy, ethical, and institutional restrictions.

Competing Interests

The authors declare no competing interests.

Acknowledgements

The authors acknowledge the Department of Information Technology and the Department of Computer Science, University of Port Harcourt, for academic and technical support.

Generative AI and AI-Assisted Technologies

OpenAI was used to assist with, document formatting, and consistency numbering checks during revision. The authors reviewed the revised manuscript and remain responsible for its accuracy, originality, interpretation, citations, and final submission.

REFERENCES

Ali, B., Gregory, M. A., Li, S., & Dib, O. A. (2024). Implementing zero trust security with dual fuzzy methodology for trust-aware authentication and task offloading in multi-access edge computing. *Computer Networks, 241*, 110197. https://doi.org/10.1016/j.comnet.2024.110197

Ameer, S., Praharaj, L., Sandhu, R., Bhatt, S., & Gupta, M. (2024). ZTA-IoT: A novel architecture for zero-trust in IoT systems and an ensuing usage control model. *ACM Transactions on Privacy and Security, 27*(3), Article 22, 1-36. https://doi.org/10.1145/3671147

Baccouri, S., & Abdellatif, T. (2025). BIG-ABAC: Leveraging big data for adaptive, scalable, and context-aware access control. *Computer Modeling in Engineering & Sciences, 143*(1), 1071-1093. https://doi.org/10.32604/cmes.2025.062902

Chen, J., Hengartner, U., & Khan, H. (2024). MRAAC: A multi-stage risk-aware adaptive authentication and access control framework for Android. *ACM Transactions on Privacy and Security, 27*(2), Article 17, 1-30. https://doi.org/10.1145/3648372

Golightly, L., Modesti, P., Garcia, R., & Chang, V. (2023). Securing distributed systems: A survey on access control techniques for cloud, blockchain, IoT and SDN. *Cyber Security and Applications, 1*, 100015. https://doi.org/10.1016/j.csa.2023.100015

Halgamuge, M. N., & Niyato, D. (2025). Adaptive edge security framework for dynamic IoT security policies in diverse environments. *Computers & Security, 148*, 104128. https://doi.org/10.1016/j.cose.2024.104128

Kalaria, R., Kayes, A. S. M., Rahayu, W., Pardede, E., & Shahraki, A. S. (2024). Adaptive context-aware access control for IoT environments leveraging fog computing. *International Journal of Information Security, 23*, 3089-3107. https://doi.org/10.1007/s10207-024-00866-4

Mahbub, M., & Shubair, R. M. (2023). 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, 219*, 103726. https://doi.org/10.1016/j.jnca.2023.103726

Matter, V. K., Martins, M. G., & Barbosa, J. L. V. (2026). Context-aware security and machine learning for access control: A systematic mapping and taxonomies. *Computer Science Review, 60*, 100880. https://doi.org/10.1016/j.cosrev.2025.100880

Matter, V. K., Schardosim, T. C., Martins, M. G., & Barbosa, J. L. V. (2026). Anubis: A smart context-aware security model for access control. *Journal of Information Security and Applications, 100*, 104469. https://doi.org/10.1016/j.jisa.2026.104469

Quan, S., Zhao, Y., & Helil, N. (2023). Attribute-based access control policy generation approach from access logs based on the CatBoost. *Computing and Informatics, 42*(3), 615-650. https://doi.org/10.31577/cai_2023_3_615

Tu, H., Xiang, D., Qi, L., & Liu, G. (2025). Modeling and verification of Chinese Wall policy based on Petri nets with data. *Computing and Informatics, 44*(4), 915-932. https://doi.org/10.31577/cai_2025_4_915

Wang, Y., Martínez-Ortega, J.-F., Castillejo, P., & San Emeterio de la Parte, M. (2025). Edge-enabled IAM for IoTs with edge-based access management and context-driven syncservice. *Journal of Systems Architecture, 165*, 103430. https://doi.org/10.1016/j.sysarc.2025.103430

Xue, T., Zhang, Y., Wang, Y., Wang, W., Li, S., & Zhang, H. (2025). Edge computing for IoT: Novel insights from a comparative analysis of access control models. *Computer Networks, 270*, 111468. https://doi.org/10.1016/j.comnet.2025.111468

Yin, J., Chen, G., Hong, W., Cao, J., Wang, H., & Miao, Y. (2024). A heterogeneous graph-based semi-supervised learning framework for access control decision-making. *World Wide Web, 27*, Article 35. https://doi.org/10.1007/s11280-024-01275-2

APPENDIX

List of Abbreviations and Acronyms

Table 5 lists every abbreviation used in the manuscript together with its full meaning. It provides a single reference point for interpreting the model, policy, storage, networking, and evaluation terminology used across the study.

Table 5: Abbreviations And Acronyms used in the Manuscript

1D CNN	One-dimensional convolutional neural network
ABAC	Attribute-based access control

ACM	Association for Computing Machinery
AI	Artificial intelligence
BIG-ABAC	Big-data-enabled attribute-based access control framework
CNN	Convolutional neural network
Conv1D	One-dimensional convolution operation or layer
F1-score	Harmonic mean of precision and recall
IAM	Identity and access management
IPFS	InterPlanetary File System
IoT	Internet of Things
ML	Machine learning
MRAAC	Multi-stage risk-aware adaptive authentication and access control
MySQL	My Structured Query Language (relational database management system)
ORCID	Open Researcher and Contributor ID
SDN	Software-defined networking
SQL	Structured Query Language
XACML	eXtensible Access Control Markup Language
XGBoost	Extreme Gradient Boosting
ZTA-IoT	Zero-Trust Architecture for the Internet of Things

Note: Definitions reflect the terminology used throughout the manuscript. Source: Authors' compilation (2026).

The table standardises the technical vocabulary used throughout the manuscript and allows each acronym to be interpreted consistently.

Appendix A: Decision Rules for the Layered Access-Control Pipeline

This appendix provides the decision rules used by the proposed layered access-control pipeline. Table A1 summarises how contextual risk, identity evidence, and conflict status map to grant, step-up, or deny outcomes.

Table A1: Illustrative mapping from evidence to final access decision

Low	Verified	No	Grant	All required evidence is favourable.
Low or medium	Uncertain	No	Step-up	Additional identity evidence is required.
Medium	Verified	No	Step-up	Context remains ambiguous despite verified identity.
High	Any	No	Deny	Contextual risk exceeds the permitted threshold.
Any	Failed	No	Deny	Identity evidence is insufficient.
Any	Any	Yes	Deny	Chinese Wall conflict overrides learned scores.

Note: Threshold values must be calibrated and approved for the deployment environment. Source: Authors' framework (2026).

The mapping makes the final decision logic explicit, policy conflict or failed identity leads to denial, while incomplete evidence triggers step-up verification.