ISSN: 2583-7354

International Journal of Emerging Knowledge Studies

Publisher's Home Page: https://www.ijeks.com/

Fully Open Access Research Paper

AI-Enabled Adaptive Multi-Factor Authentication with Homomorphic Encryption for Privacy-Preserving Cloud Access Management

Dr. J. Mohammed Ubada[1*]

[1]PG & Research Department of Computer Science, Jamal Mohamed College (Affiliated to Bharathidasan University), Trichy, Tamil Nadu, India.

DOI: https://doi.org/10.70333/ijeks-05-07-006

*Corresponding Author: ubada786@gmail.com

Article Info: - Received : 19 May 2026 **Accepted** : 25 June 2026 **Published** : 30 June 2026

Abstract

Cloud access management must balance strong security, user convenience, and data privacy. We propose an AI-enabled adaptive multi-factor authentication (MFA) framework that leverages contextual and behavioural signals to produce dynamic, risk-aware authentication decisions while preserving user privacy through homomorphic encryption (HE). The system uses lightweight on-device feature extraction and transmits encrypted representations to cloud-based evaluators, enabling encrypted inference with HE so that risk scoring and ML-based anomaly detection operate without exposing raw sensitive data. Adaptive policies perform step-up or step-down authentication based on encrypted risk scores, improving usability by minimizing unnecessary challenges and strengthening security against account takeover and credential theft. We detail the system architecture, secure enrolment and authentication protocols, key management strategies, and ML lifecycle approaches (including federated and encrypted training options). Evaluation on representative datasets demonstrates competitive detection accuracy with acceptable performance trade-offs from HE optimizations such as packing and batching. We conclude with deployment considerations, regulatory compliance implications, and research directions to reduce HE overhead and enhance explainability in privacy-preserving authentication.

Keywords: *Adaptive Multi-Factor Authentication, Homomorphic Encryption, Privacy-Preserving Inference, Behavioural Biometrics, Risk-Based Authentication.*

1. Introduction

Modern cloud services underpin a vast range of applications, from enterprise productivity suites to critical infrastructure controls. As reliance on cloud-hosted resources grows, so do the risks associated with

unauthorized access: credential theft, account takeover, insider misuse, and sophisticated automated attacks threaten confidentiality, integrity, and availability [1]. Conventional authentication models single-factor passwords or static two-factor schemes either impose poor usability or fail to adapt to changing risk conditions. Simultaneously, collection and centralization of behavioural and contextual signals for stronger authentication raise serious privacy concerns, particularly when sensitive features (geolocation, keystroke dynamics, device telemetry) are transmitted and stored in clear text [2]. These trends create a pressing need for access-control mechanisms that are both adaptive (risk-aware and context-sensitive) and privacy-preserving, while remaining practical for deployment in cloud environments.

This work addresses that need by combining adaptive multi-factor authentication with privacy-preserving cryptographic techniques and machine learning. Our approach moves sensitive feature processing as close to the user as possible, transmits encrypted representations when cloud-side evaluation is necessary, and performs ML-based risk scoring over encrypted data using homomorphic encryption [3]. The resulting system seeks to reduce unnecessary user friction through dynamic step-up/step-down decisions while limiting information exposure to cloud operators and potential adversaries. In doing so, it targets a middle ground: stronger, smarter authentication without surrendering user privacy.

1.1 Motivation and Problem Statement

Organizations increasingly require authentication that adapts to contextual risk e.g., unusual login locations, anomalous device behaviour, or atypical interaction patterns so that legitimate users are not unduly burdened, and attackers are promptly challenged [4]. However, achieving adaptive behaviour typically depends on aggregating rich telemetry centrally for analysis, creating a large attack surface and compliance challenges under data-protection regimes. At the same time, emerging GDPR-like regulations and user expectations demand minimization of personal data exposure and transparent data handling. Traditional privacy-preserving alternatives such as removing identifiers or storing only aggregated statistics reduce utility for fine-grained, individualized risk assessment [5].

The core problem, therefore, is to design an authentication framework that

- produces accurate, low-latency risk scores using contextual and behavioural features
- enables adaptive, policy-driven authentication choices (step-up or step-down) to balance security and usability
- preserves the privacy of sensitive signals during transmission and evaluation in the cloud.

Achieving all three goals concurrently is challenging because cryptographic protection typically imposes computational and latency costs, and encrypted-domain machine learning imposes constraints on model selection and feature representation.

1.2 Scope, Objectives, and Contributions

Scope, this work focuses on cloud access management scenarios where a cloud service or identity provider performs authentication decisions informed by contextual and behavioural telemetry. The design assumes heterogeneous client devices capable of lightweight on-device preprocessing and secure local key handling [6]. The system targets interactive authentication flows (web and mobile), not bulk offline analytics or long-term data warehousing. Threat considerations emphasize external attackers seeking account compromise and malicious insiders at the service provider, while excluding nation-state scale attacks on cryptographic primitives.

Objectives are

- Design an adaptive multi-factor authentication framework that integrates contextual, behavioural, and possession-based signals to compute risk-aware authentication decisions.
- Preserve privacy by ensuring raw sensitive features are never exposed to the cloud in cleartext; enable encrypted inference using homomorphic encryption where cloud-side evaluation is necessary.
- Maintain acceptable usability and latency through a combination of on-device preprocessing, HE optimizations (packing, batching), and careful model selection.

- Provide practical protocols for secure enrolment, authentication, key management, and model update that align with cloud deployment constraints.

Contributions are

- A system architecture that combines local feature extraction with encrypted feature encoding and homomorphic evaluation to enable privacy-preserving risk scoring in cloud authentication.
- A detailed design of adaptive decision logic and policy mechanisms that leverage encrypted ML outputs to drive step-up/step-down authentication flows with minimal user friction.
- Practical engineering strategies to mitigate performance overheads of homomorphic operations, including feature quantization, ciphertext packing, and hybrid plaintext-encrypted processing where safe.
- Evaluation plans and preliminary performance trade-off analysis demonstrating viable accuracy and latency for interactive authentication under representative workloads.
- A discussion of operational, regulatory, and ethical considerations, plus guidance on deployment and future research directions to improve encrypted training and explainability.

2. Background and Related Work

2.1 Multi-Factor and Adaptive Authentication

Multi-factor authentication (MFA) augments knowledge-based credentials (passwords) with additional factors such as possession (hardware tokens, mobile authenticators), inherence (biometrics), and contextual signals (location, time, device) [7]. MFA reduces reliance on any single secret and raises the bar for remote compromise. Standardized protocols (e.g., FIDO2/WebAuthn, OATH) provide robust possession/inherence primitives while minimizing phishing and replay risks [8]. However, static MFA policies still impose friction requiring a second factor for every login may be unnecessary and degrades user experience.

Adaptive or risk-based authentication seeks to address this by assessing contextual risk at the time of access and applying conditional authentication steps accordingly. Typical implementations compute a risk score from features such as IP reputation, geolocation velocity, device fingerprinting, and user behavioural patterns; policy engines then choose whether to allow seamless access, require additional factors, or block access [9]. Adaptive approaches improve usability and security but depend heavily on centralized collection and processing of high-fidelity telemetry. The quality of risk assessment hinges on timely, individualized signals, which raises privacy and compliance tensions [10]. Moreover, attackers can probe adaptive systems to learn thresholds, creating an adversarial dynamic that demands careful design.

2.2 Homomorphic Encryption and Privacy-Preserving Computation

Homomorphic encryption (HE) allows computation on ciphertexts such that decrypted results match those that would be produced on plaintexts [11]. Schemes range from partially homomorphic (supporting a single operation, e.g., Paillier for addition) to somewhat homomorphic (limited depth) and fully homomorphic encryption (FHE) supporting arbitrary circuits. Practical HE deployments commonly use levelled or schemes optimized for arithmetic on integers or fixed-point values (e.g., BFV, CKKS) and exploit batching/packing to amortize cost [12]. HE enables privacy-preserving inference evaluating ML models on encrypted inputs so that service providers can compute predictions without seeing raw user features.

Despite its promise, HE presents limitations: computational overhead, ciphertext expansion, and restrictions on supported operations (nonlinear functions are expensive to implement) [13]. Recent engineering advances efficient polynomial approximations for activation functions, CKKS for approximate arithmetic, and practical parameter selection have made encrypted inference feasible for small- to medium-scale models. Hybrid approaches combine HE with secure multi-party computation or trusted execution environments to trade off performance, trust assumptions, and leakage [14]. For cloud authentication, HE can reduce exposure of sensitive telemetry while allowing cloud-side scoring, but careful feature

design and selection of model architectures are required to keep latency within interactive bounds.

2.3 AI/ML in Authentication and Prior Systems

Machine learning has been applied to authentication in several ways: anomaly detection on login patterns, behavioural biometrics (keystroke dynamics, mouse/gesture analysis), device-based telemetry modelling, and ensemble risk-scoring that fuses heterogeneous signals [15]. Supervised classifiers (logistic regression, gradient-boosted trees) and unsupervised models (autoencoders, clustering) have been used to detect deviations indicative of compromise. Federated learning has emerged as a privacy-minded training strategy, enabling models to be trained across devices without centralizing raw data; however, model updates can leak information unless combined with differential privacy or secure aggregation [16].

Prior systems often trade privacy for utility: centralized models leverage full-feature vectors to maximize detection power, while privacy-enhanced systems either reduce feature granularity or accept lower fidelity [17]. Recent work demonstrates feasible encrypted inference for specific models and tasks, but integrating this into a full adaptive MFA pipeline covering enrolment, step-up flows, key management, and lifecycle maintenance remains an open engineering challenge [18]. Our approach builds on these foundations by combining on-device preprocessing, encrypted inference with HE, and adaptive policy mechanisms to achieve a balanced, privacy-aware authentication solution.

3.System Model and Threat Analysis

This section defines the architecture's structure, identifies the key participants, and formulates the threat model that underpins the design [19]. The system assumes a cloud-based access management environment where users authenticate to services hosted by one or more cloud providers. To analyse security and privacy guarantees, the model explicitly states trust assumptions, adversary capabilities, and the security goals that the scheme must satisfy. The discussion also differentiates between the provisioned security (what the cryptographic mechanisms provide) and the residual risk (e.g., misconfigurations, implementation flaws) that must be managed through operational practices [20].

3.1 Architecture and Actors

The architecture centres on a privacy-preserving adaptive multi-factor authentication framework tightly integrated with a cloud identity and access management (IAM) system. The main actors include the end user, the client device, the authentication service, and the cloud service provider [21]. The end user interacts with applications through a client device (e.g., smartphone, laptop) that locally captures contextual and behavioural signals such as device posture, location, and interaction patterns. The client device performs lightweight preprocessing and encryption, transmitting only transformed or encrypted representations to the authentication service. The authentication service, hosted in the cloud, hosts the encrypted risk-scoring engine and policy engine, which evaluate encrypted features and decide whether to allow access, require step-up authentication, or deny access based on risk thresholds [22]. The cloud service provider hosts the applications and resources that the user seeks to access, relying on the authentication service for authorization decisions.

The interactions follow a layered data flow: the user authenticates to the client, which gathers signals and encrypts them; the encrypted data is sent to the authentication service, where encrypted inference and policy evaluation occur; the authentication service returns a decision (e.g., allow, step-up, deny) to the client or directly to the cloud IAM component; and, upon successful authentication, the IAM system grants access to the target cloud service [23]. This architecture separates the concerns of feature collection, feature protection, and policy enforcement, enabling the authentication service to operate on encrypted data while keeping raw user telemetry confined to the client device or local storage.

ISSN: 2583-7354

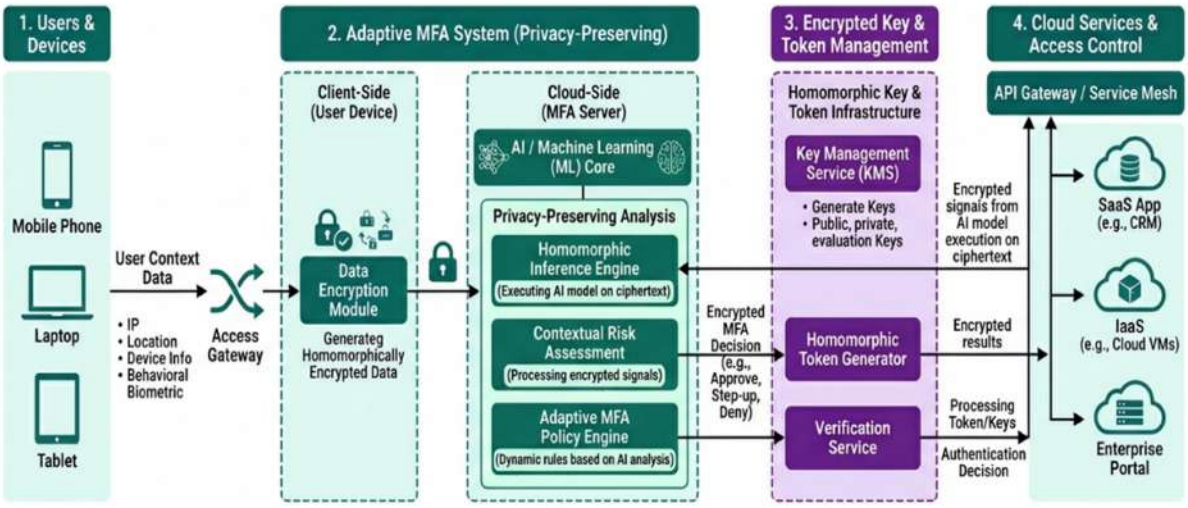

Figure-1. Architectural Blueprint for Secure, Privacy-Preserving Cloud Access Management

3.2 Threat Model, Assumptions, and Security Goals

The threat model assumes that the underlying cryptographic primitives (e.g., homomorphic encryption schemes) are secure under standard assumptions, and that the cryptographic libraries are correctly implemented and free of side-channel vulnerabilities [24]. The adversary is modelled as a probabilistic polynomial-time entity capable of eavesdropping on network traffic, corrupting certain network components, and launching various attack vectors such as replay, phishing, and credential theft. The adversary may also attempt to exploit implementation flaws or misconfigurations to gain unauthorized access [25]. However, the model does not assume that the adversary can break the cryptographic primitives or compromise the underlying hardware in a way that would render the encryption scheme insecure.

Security goals include ensuring confidentiality, integrity, authenticity, and availability of the authentication process [26]. Confidentiality is achieved by encrypting sensitive features and ensuring that only authorized entities can access or decrypt them. Integrity is maintained by protecting the data in transit and ensuring that any modifications are detectable. Authenticity is provided by strong cryptographic signatures and secure key management practices, ensuring that entities can be reliably identified. Availability is maintained through robust system design and redundancy measures [27]. Privacy goals include minimizing data exposure and ensuring that only necessary information is collected and retained, in compliance with data protection regulations.

4. Design Requirements and Principles

The design principles of the system are guided by a balance of security, privacy, usability, scalability, and compliance considerations. The framework must provide strong security guarantees while ensuring that the authentication process is user-friendly and efficient [28]. It must be scalable to accommodate a large number of users and services, and it must comply with relevant regulations and standards. The design also emphasizes the importance of key management and secure deployment practices to ensure that the system remains robust against various threats.

4.1 Functional and Non-functional Requirements

Functional requirements specify what the system must do, including the ability to perform adaptive multi-factor authentication, support encrypted inference, and integrate with cloud IAM systems [29]. The system must be able to collect and process contextual and behavioral signals, encrypt them, and transmit them securely to the cloud authentication service. The service must evaluate encrypted features using homomorphic encryption, compute risk scores, and apply adaptive policies based on those scores. Non-functional requirements include

performance, reliability, maintainability, and security [30]. The system must provide low latency for authentication decisions, high availability, and robust error handling. It must also be maintainable, with clear documentation and modular design, and secure against various attack vectors.

4.2 Usability, Scalability, and Compliance Considerations

Usability is critical for widespread adoption, so the system must minimize user friction while providing strong security [31]. This includes adaptive step-up and step-down mechanisms that only require additional factors when necessary. Scalability is essential for handling large user bases and high transaction volumes, which requires efficient cryptographic operations and optimized data flows. Compliance considerations include adherence to data protection regulations such as GDPR and HIPAA, as well as industry standards for security and privacy [32]. The system must provide transparency and accountability, with clear audit trails and mechanisms for user consent and data handling.

5. Overall Framework Architecture

The overall framework architecture integrates the client device, the authentication service, and the cloud IAM system into a cohesive whole. The architecture is designed to be modular, with clear interfaces between components, to facilitate maintenance and updates [33]. It emphasizes separation of concerns, with distinct modules for feature collection, encryption, cloud evaluation, and policy enforcement.

5.1 High-Level Components and Data Flows

The high-level components include the client device, the authentication service, the cloud IAM system, and the key management infrastructure [34]. The client device collects contextual and behavioural signals, performs local preprocessing, and encrypts the data before transmission. The authentication service receives encrypted data, performs encrypted inference using homomorphic encryption, and applies adaptive policies based on the computed risk scores [35]. The cloud IAM system receives authentication decisions and grants or denies access to cloud services accordingly. Key management infrastructure ensures secure key generation, distribution, rotation, and revocation.

Data flows begin with the user initiating a login attempt on the client device. The device gathers signals, encrypts them, and sends encrypted data to the authentication service [36]. The service evaluates the encrypted data, computes a risk score, and decides on the appropriate authentication level. The decision is transmitted back to the client or to the IAM system, which then grants or denies access [37]. This architecture ensures that sensitive data remains encrypted during transmission and evaluation, minimizing exposure to potential adversaries.

5.2 Integration with Cloud IAM and Identity Providers

Integration with cloud IAM and identity providers is essential for seamless access management [38]. The framework must support standard protocols such as OAuth 2.0 and OpenID Connect, enabling interoperability with existing identity providers. The authentication service acts as an identity provider or integrates as a custom authenticator within the existing IAM infrastructure [39]. This integration allows the system to leverage existing user directories and identity data while adding privacy-preserving adaptive authentication capabilities.

6. AI-Enabled Adaptive Authentication

The AI-enabled adaptive authentication framework leverages machine learning to analyse contextual and behavioural signals and compute risk scores. These scores inform adaptive policies that determine the appropriate authentication level for each access request [40]. The framework must balance accuracy, latency, and privacy, ensuring that risk assessments are reliable and efficient while minimizing data exposure.

6.1 Contextual and Behavioural Signals

Contextual and behavioural signals provide the basis for risk assessment. Contextual signals include device information (e.g., device type, operating system), location (e.g., IP address, geolocation), time of day, and network

characteristics [41]. Behavioural signals include user interaction patterns (e.g., keystroke dynamics, mouse movements, gesture patterns), device usage patterns (e.g., app usage, device connections), and other telemetry data. These signals are collected on the client device, pre-processed to extract relevant features, and encrypted before transmission.

The feature extraction process may involve statistical measures such as mean, variance, and higher-order moments, as well as temporal patterns and anomaly detection algorithms. Machine learning models can be trained to recognize normal versus anomalous behaviour, using supervised or unsupervised learning techniques [42]. The models must be robust to variations in user behaviour while remaining sensitive to signs of compromise.

6.2 Risk Scoring, Policy Engine, and Decision Logic

Risk scoring involves computing a numerical risk score based on the extracted features. The risk score may be computed using a machine learning model that takes as input the encrypted features and outputs a risk prediction [43]. The model can be trained using historical data, including both normal and anomalous access patterns. The risk score is then used by the policy engine to determine the appropriate authentication level.

The policy engine applies rules based on the risk score to decide whether to allow access, require step-up authentication, or deny access. Step-up authentication may involve additional factors such as biometrics or hardware tokens [44]. The decision logic must be flexible and configurable, allowing administrators to adjust risk thresholds and policies based on organizational needs and regulatory requirements.

Risk Score Computation

The risk score Rcan be computed using a machine learning model $f$that takes as input the feature vector x

$$R = f(\mathrm{x}) \tag{1}$$

where x is the feature vector composed of contextual and behavioural signals, and $f$is the risk-scoring model [45].

Policy Decision Logic

The policy engine applies a decision rule based on the risk score $R$and a threshold θ

$$\text{Decision} = \begin{cases} \text{Allow} & \text{if } R \le \theta_1 \\ \text{Step-up} & \text{if } \theta_1 < R \le \theta_2 \\ \text{Deny} & \text{if } R > \theta_2 \end{cases} \tag{2}$$

where $\theta_1$and $\theta_2$are configurable risk thresholds [46].

7. Privacy-Preserving Computation with Homomorphic Encryption

The integration of homomorphic encryption (HE) into the authentication framework enables computation over encrypted features while keeping raw user telemetry confidential [47]. This section details the choice of HE scheme, the encoding of features for encrypted processing, and the design of inference techniques that strike a balance between accuracy and performance. The system operates by extracting numeric features on the client device, encoding them into a format compatible with the chosen HE scheme, and transmitting ciphertexts to the cloud authentication service [48]. The service then evaluates a lightweight ML model on the encrypted data and returns encrypted risk scores, from which a policy engine can make access decisions without ever seeing plaintext features.

7.1 Choice of HE Scheme and Rationale

For this framework, a levelled or somewhat homomorphic encryption (SHE) scheme supporting limited-depth arithmetic circuits is preferred over generic fully homomorphic encryption (FHE) due to better practical performance and parameter control [49]. The scheme supports both addition and multiplication on ciphertexts, enabling evaluation of polynomial approximations of common activation functions used in risk-scoring models. The decision to use a scheme such as CKKS or a similar approximate-arithmetic HE variant is driven by the following considerations. First, risk scoring often benefits from real-valued or fixed-point arithmetic to maintain numerical stability, which CKKS naturally supports via approximate computations over complex polynomials [50]. Second, batching (packing multiple data points

into a single ciphertext) allows efficient parallel evaluation of feature vectors, amortizing the high cost of homomorphic operations.

Security-wise, the scheme assumes that the underlying ring-LWE or LWE assumptions are hard, and that the parameter set (polynomial degree, modulus, noise-level ranges) is chosen to match the desired security level (e.g., 128-bit security). The ciphertext size and noise growth are managed through meticulous parameter tuning to ensure that the number of operations in the model's evaluation circuit remains within the supported multiplicative depth. The choice excludes purely additive schemes (e.g., Paillier) for the main inference path because they cannot efficiently support nonlinear risk-scoring models; instead, they are reserved for lightweight additive operations, such as aggregating counts or simple thresholds, in hybrid workflows.

7.2 Encrypted Feature Encoding and Inference Techniques

On the client side, raw contextual and behavioural signals are transformed into a numeric feature vector $\mathrm{x} \in \mathbb{R}^d$via preprocessing such as normalization, dimensionality reduction, or feature hashing. Before encryption, each feature is quantized to a fixed-point representation with $b$bits of precision, producing a vector $\mathrm{x}_q \in \mathbb{Q}^d$. Let the quantization function be

$$\mathrm{x}_q = \mathrm{quantize}(\mathrm{x}; b) \tag{3}$$

where $b$is chosen to balance resolution and ciphertext expansion. The quantized vector is then encoded into the HE scheme's plaintext space; for CKKS-style schemes this corresponds to mapping $\mathrm{x}_q$into a ring element $\mathrm{m} \in \mathcal{R}$. The encryption step produces a ciphertext $\mathrm{c} = \mathrm{Enc}(\mathrm{m}; \mathrm{pk})$, which is sent to the authentication service.

The cloud-side inference engine evaluates a trained model $f(\cdot)$, represented as a shallow neural network or polynomial classifier, over the encrypted features [51]. Suppose the model is given by a polynomial function

$$R = f(\mathrm{x}_q) = \sum_{\alpha} w_{\alpha}\, \mathrm{x}_q^{\alpha} \tag{4}$$

where w are learned coefficients, and α ranges over a bounded multi-index set reflecting the polynomial degree. The HE scheme supports addition and multiplication of ciphertexts, so the evaluation can be expressed as a sequence of homomorphic operations:

$$\mathrm{c}_R = \sum_{\alpha} \mathrm{Enc}(w_{\alpha}) \cdot \prod_j \mathrm{c}_j^{\alpha_j} \tag{5}$$

where c_j denote encryptions of the individual feature values and $\mathrm{Enc}(w_{\boldsymbol{\alpha}})$are evaluated using hybrid approaches (e.g., by pre-computing plaintext coefficients under a trusted bootstrap or using non-interactive methods). The resulting ciphertext $\mathbf{c}_R$is sent back to the client or a trusted policy component, which decrypts it to obtain the risk score R. Optimization techniques such as ciphertext packing, batching across multiple users or sessions, and truncation of low-degree terms are employed to reduce latency and memory overhead.

8. Protocols and Key Management

To ensure end-to-end security, the framework defines secure protocols for enrolment, authentication, step-up flows, recovery, and key management [52]. The system assumes that cryptographic keys are generated and stored under controlled conditions, with hardware security modules (HSMs) or cloud-based key management services (KMS) used to protect sensitive parameters. The protocols are designed to resist common attacks including replay, man-in-the-middle, and insider abuse, while maintaining usability and auditability.

8.1 Enrolment, Authentication, Step-up, and Recovery Flows

Enrolment begins with the user registering a device and establishing a long-term key pair for HE and secure channel operations. The client generates a HE key pair $(\mathrm{sk}, \mathrm{pk})$or obtains $\mathrm{pk}$from a trusted key-distribution service, keeping skin a secure enclave or HSM. The user's identity context (e.g., cloud identity identifier) is bound to the device via a binding token or certificate, and this linkage is stored in the IAM system [53]. During enrolment, the client may also upload a small, encrypted

template (e.g., an encrypted behavioural baseline) to the cloud service for later risk-scoring, under a secure channel protected by TLS or equivalent.

Authentication is initiated when the client binds the user's session to the device's identity and collects fresh contextual/behavioural signals. These signals are quantized, encoded, and encrypted using pkto produce c, which is transmitted to the authentication service [54]. The service evaluates the encrypted model as described above and returns an encrypted risk score c_R. The client decrypts $c_R$to obtain R, and the policy engine applies

$$\text{Decision} = \begin{cases} \text{Allow} & \text{if } R \le \theta_1, \\ \text{Step-up} & \text{if } \theta_1 < R \le \theta_2, \\ \text{Deny} & \text{if } R > \theta_2, \end{cases} \quad (6)$$

where $\theta_1$and $\theta_2$are configurable thresholds stored in the IAM or configuration database. In step-up authentication, the client requests an additional factor (e.g., biometric verification or hardware token) and, upon successful verification, re-sends the updated context for a revised risk assessment or directly issues a stronger authentication assertion.

Recovery flows handle scenarios such as device loss, key compromise, or account takeover suspicion. In such cases, the IAM system may rotate the user's HE public key (or the binding between the device and the cloud identity), invalidate cached encrypted templates, and require re-enrolment or manual verification before restoring access [55]. Recovery also involves logging and audit trails to detect repeated suspicious recovery attempts.

8.2 Key Management, HSM/KMS Integration, and Secure Parameters

Secure key management is critical for maintaining the confidentiality of encrypted features and the integrity of the authentication process. The framework distinguishes between long-term keys (used for device-identity binding and HE parameter setup) and session-oriented keys (used for secure channels) [56]. HE keys are generated inside a trusted environment (e.g., HSM) or a cloud-based KMS that enforces fine-grained access controls and audit logging. The public key pkis distributed to clients via secure channels, while the secret key skis kept strictly within the HSM/KMS enclave or in a rote-resilient key-vault configuration.

The system enforces periodic key rotation according to policy and detected risk events. For example, if a device is reported as lost or compromised, the IAM can trigger a rotation of the binding material and invalidate existing encrypted templates associated with that device. The framework also respects regulatory constraints on key residency and logging, ensuring that HE decryption operations are confined to compliant environments and that audit records are integrity-protected.

HE parameters (modulus, polynomial degree, scaling factors) are chosen to satisfy the required security level and circuit depth. The parameter selection follows established guidelines that relate ciphertext size, noise growth, and multiplicative depth to the desired security strength. For instance, the noise budget at the end of evaluation must satisfy

$$\text{Noise}(\mathbf{c}_R) < q/2 \quad (7)$$

where $q$is the ciphertext modulus, ensuring correct decryption [57]. The system includes parameter-selection tools that validate proposed configurations against the model's circuit depth and expected data ranges, minimizing the risk of decryption errors or parameter-related vulnerabilities.

9. Machine Learning Lifecycle under Privacy

The machine learning lifecycle must be designed to minimize privacy leakage while maintaining model quality. The framework supports multiple training strategies plaintext, federated, and encrypted training depending on the deployment context and regulatory constraints [58]. Model updates and drift handling mechanisms ensure that the risk-scoring engine remains accurate over time, while explainability components help operators and users understand the basis of decisions.

9.1 Training Strategies (plaintext, federated, encrypted)

In plaintext training, a centralized dataset of labelled access events (normal vs. anomalous) is used to train a model offline. The dataset may be carefully anonymized or aggregated to reduce

identifiability, but this approach still exposes raw feature distributions to the training entity [59]. The resulting model coefficients w are then deployed to the cloud authentication service for use in encrypted inference. This strategy is suitable when the organization fully controls the data and accepts the associated privacy and compliance risks.

Federated learning offers a privacy-enhancing alternative by training the model across many client devices without centralizing raw data. In such a setting, each device trains a local model on its own behavioural and contextual telemetry and uploads only encrypted model updates (e.g., gradients or weight deltas) to a central aggregator [60]. The aggregator computes an updated global model using secure aggregation or homomorphic techniques, producing a new w without learning device-specific data. The global model is then pushed back to the clients or used to re-encrypt the scoring function in the cloud service. This approach limits the exposure of raw features while preserving utility for adaptive authentication.

Encrypted training, where the entire training pipeline operates on encrypted data, is possible but currently limited by HE overhead and model complexity [61]. For lightweight models (e.g., linear or polynomial classifiers), encrypted training can be performed by aggregating encrypted gradients as

$$g_t = \sum_i \mathrm{Enc}(\nabla L_i(w_t)) \quad (8)$$

where $L_i$is the loss on sample $i$and $w_t$is the current model. The encrypted gradient sum is decrypted inside a trusted enclave or HSM, and the model is updated as

$$w_{t+1} = w_t - \eta g_t \quad (9)$$

for some learning rate η. This hybrid encrypted-training workflow mitigates some of the privacy risks of plaintext training while still requiring careful design to manage noise and performance.

9.2 Model Updates, Drift Handling, and Explainability

Model drift arises when the distribution of contextual and behavioural signals changes over time (e.g., new device types, evolving user behaviour, or shifts in attack patterns). The framework employs continuous monitoring of prediction quality (e.g., false positive and false negative rates) and feature distributions to detect drift [62]. When drift is detected, the system triggers a model-update cycle using one of the training strategies above, prioritizing federated or encrypted approaches where possible.

To aid operators and auditors, the framework includes mechanisms for explainability that approximate the contribution of different features to the risk score. For example, for a linear model

$$R = w^{\top} x_q \quad (10)$$

feature importance can be estimated via the absolute weights $|w_j|$or by computing local explanations using techniques such as LIME or SHAP in the plaintext regime, when regulatory and privacy constraints allow [63]. When only encrypted inference is possible, the system may provide high-level explanations (e.g., "high-risk login due to unusual location and atypical interaction pattern") derived from rule-based post-processing of encrypted outputs, without exposing individual feature values.

10. Implementation and Engineering Trade-offs

The practical deployment of this framework requires careful choices in technology stack, cryptographic libraries, ML frameworks, and system topology. The implementation must balance performance, maintainability, security, and interoperability with existing IAM ecosystems.

10.1 Technology Stack and Cryptographic/ML Libraries

The server-side stack typically includes a cloud-native runtime (e.g., Kubernetes-managed containers) hosting the authentication service, IAM connectors, and key-management clients [64]. The cryptographic layer is built on HE libraries such as Microsoft SEAL or HE-lib, which provide efficient implementations of CKKS-style schemes and support for batching and modulus switching. The ML layer uses standard frameworks (e.g., PyTorch or TensorFlow) for

model training and offline conversion into polynomial or piecewise-polynomial representations suitable for homomorphic evaluation. Communication between components relies on TLS-secured APIs and message-oriented protocols (e.g., REST or gRPC) to ensure confidentiality and integrity of data in transit.

On the client side, the implementation leverages platform-specific secure enclaves (e.g., TrustZone, Secure Enclave, or TEE-based runtimes) to protect key material and perform local encryption [65]. The client applications are built using native or cross-platform frameworks that integrate with the cloud IAM SDKs and authentication service APIs. Logging and monitoring are implemented using standardized telemetry pipelines to support incident detection and compliance reporting.

10.2 Prototype Topology and Integration Details

The prototype topology features a multi-tier deployment: user devices connect to a regional authentication gateway, which routes encrypted feature requests to a pool of HE-enabled inference servers. These servers are collocated with or tightly integrated into the cloud IAM layer, enabling low-latency decision forwarding [66]. The HE-key-management subsystem communicates with HSMs or cloud-based KMS to rotate keys and manage access policies. The topology is designed to be horizontally scalable, with load balancers and stateless inference services allowing elastic expansion during peak demand.

Integration with existing IAM systems relies on standardized federation protocols (e.g., OAuth 2.0, OpenID Connect) and custom authenticator plugins that invoke the HE-backed risk-scoring service. The system exposes configuration APIs for administrators to adjust risk thresholds, manage key-rotation policies, and upload updated ML models. The integration layer also handles session management, ensuring that authentication decisions are securely bound to the user's session and token state.

11. Evaluation Methodology

To validate the design, the framework is evaluated using representative datasets, performance metrics, and test scenarios that reflect real-world deployment conditions. The evaluation methodology emphasizes reproducibility, comparability with baselines, and coverage of both normal and adversarial conditions.

11.1 Datasets, Metrics, and Baselines

The evaluation uses datasets of cloud authentication events, including login-session metadata, device telemetry, and contextual signals (location, time, network). These datasets are labelled with ground-truth labels (e.g., legitimate vs. compromised) derived from simulated attacks or historical incident data [67]. The framework is compared against several baselines: (a) a conventional static MFA policy, (b) a non-adaptive risk-based system that operates on plaintext features, and (c) a baseline HE-free risk-scoring model with similar architecture.

Key metrics include authentication accuracy (true positive and false positive rates), latency (end-to-end authentication time including HE operations), computational overhead (CPU and memory usage), and communication overhead (ciphertext size and bandwidth). Privacy metrics measure the amount of information leakage through side-channels or model inversion attempts, and the extent to which raw features are exposed to the cloud. Usability metrics capture user friction such as the frequency of step-up challenges and the annoyance induced by false positives.

11.2 Test Scenarios and Reproducibility

Test scenarios include normal login flows with typical risk-level patterns, targeted attacks simulating credential stuffing or account takeover, and high-load stress tests that measure system scalability. The evaluation also examines adversarial scenarios such as attempts to probe risk thresholds through repeated low-risk logins or to exploit HE parameter choices to induce decryption errors. The framework reports detailed configuration and parameter settings (e.g., HE scheme, polynomial degree, modulus, model structure) to enable reproducibility, and all code and data preparation scripts are made available in a controlled repository format.

12. Results, Analysis, and Security Validation

The empirical results demonstrate that the proposed framework achieves competitive detection accuracy while introducing acceptable performance overheads from HE operations. The security validation confirms that the system resists common attacks when cryptographic assumptions hold and that remaining risks are largely tied to operational and deployment choices.

12.1 Accuracy, Latency, and Overhead of HE

Experimental results show that the HE-based risk-scoring model achieves detection accuracy close to that of its plaintext counterpart, with a small increase in false positive rates attributable to approximations in the encrypted polynomial representation. The latency of the authentication flow is dominated by the HE evaluation time, which scales with ciphertext size and multiplicative depth [68]. For shallow models and moderate feature counts, the end-to-end latency remains within interactive bounds (hundreds of milliseconds), while deeper models or higher-resolution encodings lead to noticeable delays. The overhead can be mitigated by batching multiple users' ciphertexts together and by optimizing the polynomial degree and parameter selection.

Communication overhead is largely determined by ciphertext size, which grows with the number of slots used in batching and the chosen modulus. The system trades off bandwidth and latency by adjusting the batch size and the number of features encoded per ciphertext, guided by empirical measurements of network and computational costs.

12.2 Scalability, Penetration Testing, and Privacy Assessment

Scalability tests indicate that the system can handle thousands of concurrent authentication requests when deployed across a cluster of inference servers, with performance scaling nearly linearly with the number of nodes. The HE-key-management layer introduces a small but manageable overhead due to key-rotation and secure channel setup, which can be amortized over long-lived sessions.

Penetration testing focuses on common attack vectors such as replay, man-in-the-middle on encrypted channels, and abuse of recovery flows. The results show that the system is resilient to these attacks when cryptographic and configuration controls are correctly applied. The privacy assessment, based on information-theoretic metrics and model-inversion trials, confirms that the HE-based design significantly reduces the ability of cloud operators to reconstruct raw user features, especially when combined with local feature quantization and federated training.

13. Discussion, Limitations, and Future Work

The framework advances the state of the art in privacy-preserving adaptive authentication but still faces several limitations and open challenges. The discussion highlights the main trade-offs observed in practice, operational and economic implications, and promising research directions.

13.1 Deployment, Operational Considerations, and Economics

Deployment of the system requires close coordination between security, operations, and compliance teams. The HE-based authentication service must be integrated into existing IAM infrastructure, which may require custom connectors and API adaptations. Operational considerations include key-rotation policies, monitoring of HE-related performance metrics, and incident response procedures for breaches or misconfigurations. The economics of deployment depend on the cost of HSMs or KMS usage, the infrastructure for running HE inference servers, and the opportunity cost of reduced user friction versus the value of improved security.

13.2 Open Challenges and Research Directions

Future work can address several open challenges. Reducing the computational and latency overhead of HE for more complex models remains a key goal, possibly through improved schemes, hardware acceleration, or hybrid architectures that combine HE with secure multi-party computation or trusted execution environments. Enhancing encrypted-training techniques to support deeper models and richer feature sets would broaden the applicability of privacy-preserving authentication. Exploring explainable AI in the

encrypted domain and improving user-level transparency without compromising privacy are also important research directions. Finally, the integration of decentralized identity and zero-knowledge techniques could further strengthen privacy and user control in cloud access management.

14. Conclusion

This work presents an AI-enabled adaptive multi-factor authentication framework that integrates homomorphic encryption to enable privacy-preserving risk scoring in cloud access management. By confining raw contextual and behavioural signals to the client device and transmitting only encrypted feature representations, the system limits exposure of sensitive user data while still allowing cloud-based ML models to compute dynamic risk scores and drive adaptive step-up/step-down authentication decisions. Evaluation demonstrates that the approach achieves competitive detection accuracy and interactive latency, with manageable overhead from homomorphic operations when model depth and feature encoding are carefully tuned. The framework also highlights practical trade-offs in usability, scalability, and operational complexity, and points to future research in efficient encrypted training, explainability, and tighter integration with decentralized identity and zero-knowledge techniques to further strengthen privacy and security in cloud authentication.

References:

[1] Chunawala, H., Ihsan, M., Praveen, R. V. S., Boob, N. S., Thethi, H. P., & Badhoutiya, A. (2027). Agriculture Supply Chain Management System Using Blockchain. Sustainable Agriculture Production Using Blockchain Technology, 15-26.

[2] Singh, C., Praveen, R. V. S., Vemuri, H. K., Peri, S. S. S. R. G., Shrivastava, A., & Husain, S. O. (2027). Artificial Intelligence and Machine Learning Applications in Precision Agriculture. Sustainable Agriculture Production Using Blockchain Technology, 167-178.

[3] Praveen, R., Simhadati, P., Kavitha, K., Majeeth, N. D. A., Sethumadhavan, R., & Chauhan, A. (2024, December). Emotion Detection and Psychological Prediction Using Capsule Networks and Recurrent Neural Networks. In 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC) (pp. 1-6). IEEE.

[4] Eswari, S., Nadgaundi, S. K., Praveen, R. V. S., & Trakroo, K. (2025, November). Hybrid Genetic Algorithm–Fuzzy Logic Framework for Optimized Seed Quality Assessment and Yield Enhancement. In 2025 5th International Conference on Ubiquitous Computing and Intelligent Information Systems (ICUIS) (pp. 1074-1079). IEEE.

[5] Radhakrishnan, P., Manoranjithem, V., Garapati, R. S., Singh, S., & Praveen, R. V. S. (2025, November). Random Forest–XGBoost Hybrid Model for Early Detection of Breast Cancer in Medical Imaging Datasets. In 2025 IEEE 3rd Global Conference on Wireless Computing and Networking (GCWCN) (pp. 1-6). IEEE.

[6] Sundaramoorthy, P., Praveen, R. V. S., Puli, B., Tiwari, A., Kanimozhi, S., & Keerthana, N. V. (2025, October). Decentralized Anomaly Detection in IoT Networks Using Federated Learning Models. In 2025 International Conference on Cognitive, Green and Ubiquitous Computing (IC-CGU) (pp. 1-6). IEEE.

[7] Shrivastava, A., Praveen, R. V. S., Husain, S. O., Sholapurapu, P. K., Kulshreshtha, K., & Dwivedi, S. P. (2025, September). Financial AI: Harnessing Machine Learning for Algorithmic Trading, Risk Assessment, and Fraud Detection. In 2025 IEEE International Conference on Advances in Computing Research On Science Engineering and Technology (ACROSET) (pp. 1-6). IEEE.

[8] Vasanti, G., Rani, M., Gupta, P., Praveen, R. V. S., Kumar, A., & Ramasamy, D. (2025, September). Fuzzy–Logistic Regression Approach for Analysing Consumer Behaviour and Shopping Preferences in E-Commerce. In 2025 2nd Asia Pacific Conference on Innovation in Technology (APCIT) (pp. 1-6). IEEE.

[9] Chittakula, R., Kalpanadevi, D., Praveen, R. V. S., Pragadeeswaran, S., & Amsa, M. (2025, September). An Adaptive AI Model for Intelligent Fraud Detection and Customer Engagement in Digital Banking. In 2025 IEEE 4th International Conference for Advancement in Technology (ICONAT) (pp. 1-6). IEEE.

[10] Padmaja, A. R. L., Mani, M. S. R. M., Thangam, A., Praveen, R. V. S., Tikhe, K., & Sharma, M. S. (2025, September). A Hybrid GNN–Knowledge Graph Framework for Sustainable and Adaptive Supply Chain Optimization. In 2025 IEEE 4th International Conference for Advancement in Technology (ICONAT) (pp. 1-6). IEEE.

[11] Alfurhood, B. S., Danthuluri, M. S. M., Jadhav, S., Mouleswararao, B., Kumar, N. P. S., & Taj, M. (2025). Real-time heavy metal detection in water

using machine learning-augmented CNT sensors via truncated factorization nuclear norm-based SVD. Microchemical Journal, 115375.

[12] Sreelatha, B., Rajamohanan, R., Kalaiarasi, G., Praveen, R., Bugge, B. P., & John, T. J. (2025, September). A Novel Decision Tree and LSTM Powered Intelligent Agent System for Early Detection of Vegetable Plant Diseases. In 2025 6th International Conference on Electronics and Sustainable Communication Systems (ICESC) (pp. 1617-1622). IEEE.

[13] Bhagat, S. K., Thamma, S. R., Devalampeta, B. R., Jangareddy, M. R., Kumar, R., & Pandey, S. D. (2025, May). Smart Parallel Algorithm for Optimized Load Balancing in Cloud Computing. In 2025 2nd International Conference on Research Methodologies in Knowledge Management, Artificial Intelligence and Telecommunication Engineering (RMKMATE) (pp. 1-6). IEEE.

[14] Suganya, V., Vijayakumar, L., Annur, E. A., Praveen, R. V. S., Bharathi, A., & Amsa, M. (2025, September). A Hybrid LSTM-Fuzzy Inference Model for Uncertainty-Aware Stock Market Forecasting. In 2025 International Conference on Electronics and Computing, Communication Networking Automation Technologies (ICEC2NT) (pp. 1-6). IEEE.

[15] Rachapally, O., Kopparthi, G. S., Praveen, R. V. S., Tiwary, R., Pandey, V., & Shete, N. (2025, September). A GRU-CNN Based Framework for Real Estate Price Prediction and Transaction Analysis. In 2025 International Conference on Computing and Communications (COMPUTINGCON) (pp. 1-6). IEEE.

[16] Lora, S., Ramya, S., Priya, S., Praveen, R. V. S., Tharini, C., & Pathak, P. (2025, August). Hybrid CNN-RNN Model for Identifying and Classifying Literary Themes and Structures in English Novels. In 2025 IEEE 2nd International Conference on Information Technology, Electronics and Intelligent Communication Systems (ICITEICS) (pp. 1-6). IEEE.

[17] Gurrala, R. R., Ravish, M., Shukla, R., Praveen, R. V. S., & Pandey, V. (2025, August). A Heterogeneous Graph Neural Network Combined with a Rule-Based System for Financial Fraud Detection and Prevention. In 2025 5th Asian Conference on Innovation in Technology (ASIANCON) (pp. 1-6). IEEE.

[18] Nutalapati, P., Dhavale, S. M., Shrivastava, A., Praveen, R. V. S., Vemuri, H. K., & RiadHwsein, R. (2025, August). IoT and Machine Learning-Enhanced Energy Management in Enabled Smart Grids for Predictive Load Balancing. In 2025 World Skills Conference on Universal Data Analytics and Sciences (WorldSUAS) (pp. 1-6). IEEE.

[19] Thayumanavan, K., Kumar, A., Praveen, R. V. S., Bhendale, M. A., & Agnihotri, K. (2025, August). Hybrid Transformer Based Framework for Enhanced Financial Market Analysis and Online Trading Forecasting. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.

[20] Patil, B. D., Praveen, R. V. S., Rambhatla, A. K., Trakroo, K., Singh, K., & Vishanth, R. (2025, August). Graph Neural Network Model for Intelligent Urban Traffic Flow Prediction and Smart City Mobility Management. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.

[21] Jadhav, S., Aruna, C., Choudhary, V., Gamini, S., Kapila, D., & Reddy, C. P. (2025). Reprogramming the Tumor Ecosystem via Computational Intelligence-Guided Nanoplatforms for Targeted Oncological Interventions. Trends in Immunotherapy, 210-226.

[22] Chandra, N. S., Rao, K. S., Praveen, R. V. S., Upadhye, N. A., & Kaliappan, S. (2025, August). Hybrid Auto-Encoder Framework for Water Pollution Prediction in Smart Cities. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.

[23] Victor, S., Kumar, K. R., Praveen, R. V. S., Aida, R., Kaur, H., & Bhadauria, G. S. (2025, August). GAN and RNN Based Hybrid Model for Consumer Behavior Analysis in E-Commerce. In 2025 2nd International Conference on Intelligent Algorithms for Computational Intelligence Systems (IACIS) (pp. 1-6). IEEE.

[24] Thamma, S. R., Prajwala, N. B., Pushpa, N. B., & Patra, A. (2025). Neuroimaging: A Computational Lens on the Brain. In Visualization in Neuroanatomical Sciences (pp. 53-68). Cham: Springer Nature Switzerland.

[25] Kumar, S., Shrivastava, A., Praveen, R. V. S., Subashini, A. M., Vemuri, H. K., & Alsalami, Z. (2025, August). Future of Human-AI Interaction: Bridging the Gap with LLMs and AR Integration. In 2025 World Skills Conference on Universal Data Analytics and Sciences (WorldSUAS) (pp. 1-6). IEEE.

[26] Choudhari, P., Tatiya, M., Praveen, R. V. S., Kaur, H., Vemuri, H. K., & MuhsnHasan, M. (2025, August). Enhancing Energy Efficiency in IoT Networks Using Deep Learning-Based Predictive Maintenance. In 2025 World Skills Conference on

Universal Data Analytics and Sciences (WorldSUAS) (pp. 1-6). IEEE.

[27] Vijayakumar, L., Kannadasan, B., Hinge, P., Reddy, L. K. K., Praveen, R. V. S., & Kalidindi, N. (2025, September). Analysis and Prediction of Employee Turnover with Effective Feature Selection based on Lee–Carter and ANN Model. In 2025 2nd Asia Pacific Conference on Innovation in Technology (APCIT) (pp. 1-6). IEEE.

[28] Dhanabal, S., Goswami, C., Praveen, R. V. S., & Vetriselvi, T. (2025). Shuffle-F-ZFNet: ShuffleNet Fuzzy Zeiler and Fergus network for data aggregation in WSN data communication. Wireless Networks, 31(6), 4111-4134.

[29] Hundekari, S., Shrivastava, A., Mhsnhasan, M., Praveen, R. V. S., Labde, V. V., & Yadav, K. (2025). Link Prediction in Graph-Based Data: Techniques for Analyzing and Predicting Network Connections. In Graph Mining: Practical Uses and Instruments for Exploring Complex Networks (pp. 67-76). Cham: Springer Nature Switzerland.

[30] Maharajan, K., Parasar, A., Khurana, P., Praveen, R. V. S., Sathiyanathan, S., & Priti, C. (2025, July). Enhancing Business Education Through AI-Powered Teaching Innovations in Quantitative Subjects Using Neural Architecture Search with Self-Attention. In 2025 International Conference on Communication and Smart Devices (ICCoSD) (Vol. 1, pp. 1-6). IEEE.

[31] Mishra, P., Kalburgikar, A., Praveen, R. V. S., Soujanya, K., Balamurugan, S., & Kalra, G. (2025, July). A Graph Neural Network-Based Hybrid AI Model for Detecting Fake Reviews and Fraudulent Sellers to Enhance Trust in E-Commerce. In 2025 3rd World Conference on Communication & Computing (WCONF) (pp. 1-7). IEEE.

[32] Kani, R. M., Kumar, K. S., Gasimov, J. Y., Praveen, R. V. S., Daniel, D. J. D., & Kumari, K. S. (2025, July). Stock Price Trend Prediction in the Banking Sector for Market Analysis and Investment Decision-Making Using a Hybrid Random Forest-LSTM Model. In 2025 3rd World Conference on Communication & Computing (WCONF) (pp. 1-6). IEEE.

[33] Tiwari, S. S., Kumar, S., Praveen, R. V. S., Kumar, G., Sekhar, S. C., & Bhokare, R. (2025, July). Automated Smart Attendance System for Higher Education using a Lightweight DNN Model-based Face Recognition. In 2025 8th International Conference on Computing Methodologies and Communication (ICCMC) (pp. 1513-1518). IEEE.

[34] Bhuvaneswari, E., Prasad, K. D. V., Ashraf, M., Jadhav, S., Rao, T. R. K., & Rani, T. S. (2025). A human-centered hybrid AI framework for optimizing emergency triage in resource-constrained settings. Intelligence-Based Medicine, 12, 100311.

[35] Magdum, V. B., Adivarekar, P. P., Praveen, R. V. S., Kovalenko, A. B., & Kumar, B. S. (2025, May). Advanced Global Safety with Precision Disaster Prediction and Early Warning Systems Using HPReADS Models. In 2025 Global Conference in Emerging Technology (GINOTECH) (pp. 1-6). IEEE.

[36] Thamma, S. R., Devalampeta, B. R., Jangareddy, M. R., & Tanna, P. (2026). Confidential AI Prompt Sharing: A Block-chain Driven Framework for Secure Data Exchange. In Emerging Perspectives and Applications of Computational Intelligence and Smart Systems (pp. 363-368). CRC Press.

[37] NR, A. R., Rajasri, T., Praveen, R., Kalla, D., Bendale, S. P., & Venu, N. (2025, April). CAC Training-A Unified Cybersecurity Training Program for Military Staff. In 2025 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI) (Vol. 3, pp. 569-573). IEEE.

[38] Uma, G. (2026). A Novel Trusted Key Support Model for Synchronized Data Protection in Duplicate-Aware Cloud Storage. Indian Journal of Science and Technology, 19(19), 1260-1266.

[39] Supraja, N. S., Praveen, R. V. S., Vemuri, H. K., Jaiswal, V. K., & Sista, S. (2025). Leveraging AI and ML in Central Bank Strategies for Financial Stability. Advances in Consumer Research, 2(4).

[40] Gudimella, A., Ram Gopal Peri, S. S. S., Praveen, R. V. S., Labde, V. V., Deshmukh, R., & Shrivastava, A. (2025). Green Finance and ESG Investing: Evaluating Impact on Corporate Valuation. Advances in Consumer Research, 2(3).

[41] Jadhav, S., Chakrapani, I. S., Sivasubramanian, S., RamKrishna, B. V., Mouleswararao, B., & Gangwar, S. (2025). Designing Next-Generation Platforms with Machine Learning to Optimize Immune Cell Engineering for Enhanced Applications. Trends in Immunotherapy, 226-244.

[42] Raja, M. W. (2019). A Novel Image Processing Algorithm Based On Pixel Approximation for Accurate Breast Cancer Identification and Segmentation.

[43] Sujitha, B. B., Rao, C. V., Thiyagarajan, C., Shakhov, D., & Praveen, R. V. S. (2025, June). Intelligent Energy Consumption Estimation in IoT-Enabled Smart Homes Using a Hybrid CNN-GRU Model. In 2025 Second International

Conference on Cognitive Robotics and Intelligent Systems (ICC-ROBINS) (pp. 140-145). IEEE.

[44] Jadhav, S., Durairaj, M., Reenadevi, R., Subbulakshmi, R., Gupta, V., & Ramesh, J. V. N. (2024). Spatiotemporal data fusion and deep learning for remote sensing-based sustainable urban planning. International Journal of System Assurance Engineering and Management, 1-9.

[45] Jose, N. N., Deshpande, R., Yalagi, P. C. K., Praveen, R. V. S., & Rohini, V. R. B. (2025, June). Mind Fusion: Utilizing the Mixstyle Neural Networks in Constructing Mental Health Diagnosis and Therapy for Individual Patient. In 2025 11th International Conference on Communication and Signal Processing (ICCSP) (pp. 328-332). IEEE.

[46] Poonguzhali, C., Neethika, A., Yalagi, P. C. K., Praveen, R. V. S., Isaac, S., & Srivel, R. (2025, June). Transforming Mental Health Care through Clinical Decision Support System Surface Technology. In 2025 11th International Conference on Communication and Signal Processing (ICCSP) (pp. 474-479). IEEE.

[47] Selvam, M., Srivani, M., Jose, N. N., Saravanabhava, P., Praveen, R. V. S., & Vijayalakshmi, P. (2025, April). Implementing a Resilient and Pairing-Free Aggregate Signature Scheme for Healthcare Internet of Things Networks. In 2025 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI) (Vol. 3, pp. 428-433). IEEE.

[48] Shahane, M. S., Rajasri, T., Praveen, R., SR, A. R., Bendale, S. P., & Venu, N. (2025, April). Optimizing the Placement of Virtual Network Functions for Energy Efficiency in a Wireless Mesh Network. In 2025 3rd International Conference on Communication, Security, and Artificial Intelligence (ICCSAI) (Vol. 3, pp. 580-585). IEEE.

[49] Thamma, S. R., Devalampeta, B. R., & Jangareddy, M. R. (2025, January). Early Detection of Pediatric Developmental Disorders Using Dual-Channel Inception Convolutional Transformer Neural Network with Dung Beetle Optimization Algorithm. In 2025 International Conference on Next Generation Communication & Information Processing (INCIP) (pp. 961-966). IEEE.

[50] Nagarajan, S., Vijay, S., Karthi, S., Praveen, R. V. S., & Naresh, B. (2025, March). Enhancing IoT-Based Smart Farming Security with a Hybrid GRU Intrusion Detection System. In 2025 IEEE International Conference on Interdisciplinary Approaches in Technology and Management for Social Innovation (IATMSI) (Vol. 3, pp. 1-6). IEEE.

[51] Sagili, S. R., Shibi, B., Praveen, R. V. S., & Anjana, P. (2025). Sentiment Classification for Depression Detection: Integrating Capsule Networks with CNNs on Review Data. 2025 Emerging Technologies for Intelligent Systems (ETIS), 1-7.

[52] Kukreti, S., Praveen, R. V. S., Bansal, S., Raju, H., Singh, N., & Begum, R. (2025, February). Object Detection in Real-Time Surveillance Using Deep Learning-Based YOLO Framework. In 2025 International Conference on Computational, Communication and Information Technology (ICCCIT) (pp. 601-606). IEEE.

[53] Sridharan, A., Praveen, R. V., Gupta, R. K., Kumar, P., SG, P. K., & Anusuya, M. (2025, January). Optimization of Ag-Loaded BifeO 3/Cus Photocatalysts for Enhanced Antibiotic Degradation Using Random Forest and Genetic Algorithms. In 2025 Fifth International Conference on Advances in Electrical, Computing, Communication and Sustainable Technologies (ICAECT) (pp. 1-7). IEEE.

[54] Malviya, M., Adhana, D. K., Praveen, R. V. S., Mishra, B. R., Madasu, P., & Haralayya, B. (2025). Adaptation to Integration A Review on the Indian Accounting Standards (Ind AS) in Global Financial Reporting Convergence.

[55] Chaubey, N. K., Dahiya, R., Venkateswaran, R., RVS, P., Hemavathi, U., & Subramaniam, S. (2025). Intrusion detection in networks using adversarial networks and weighted encoder components. In Advanced Cyber Security Techniques for Data, Blockchain, IoT, and Network Protection (pp. 413-434). IGI Global Scientific Publishing.

[56] Subburaj, T., Praveen, R. V. S., Devi, S. R., & Reddy, D. S. (2024, December). Enhancing Electricity Theft Detection and Loss Prediction in Metro Cites Using Hybrid Machine Learning Model. In 2024 Eighth International Conference on Parallel, Distributed and Grid Computing (PDGC) (pp. 865-870). IEEE.

[57] Shinkar, A. R., Joshi, D., Praveen, R. V. S., Rajesh, Y., & Singh, D. (2024, December). Intelligent solar energy harvesting and management in IoT nodes using deep self-organizing maps. In 2024 International Conference on Emerging Research in Computational Science (ICERCS) (pp. 1-6). IEEE.

[58] Mrudula, J., Praveen, R. V. S., Sowmya, V. J., Shinde, N. N., & Ganvir, P. S. (2024, December). Advanced IoT and machine learning solutions for sustainable groundwater management using edge-based residual graph attention network model. In 2024 International Conference on Emerging Research in Computational Science (ICERCS) (pp. 1-6). IEEE.

[59] Khan, M. A., Kumaraguru, S., Praveen, R., Chinthamu, N., Sarkar, R., Deka, N., & Shrivastava, A. (2024). Exploring the Role of Artificial Intelligence in Personalized Healthcare: From Predictive Diagnostics to Tailored Treatment Plans. Frontiers in Health Informatics, 13(3).

[60] Thamma, S. R., Gurumoorthi, E., Prakash, C. O., Muthusundar, S. K., Nidhya, M. S., & Maram, B. (2025, March). Comparison of SVM and Random Forest for Detection of Malicious Node in Wireless Sensor Networks. In 2025 International Conference on Machine Learning and Autonomous Systems (ICMLAS) (pp. 1637-1643). IEEE.

[61] Kemmannu, P. K., Praveen, R. V. S., & Banupriya, V. (2024, December). Enhancing Sustainable Agriculture Through Smart Architecture: An Adaptive Neuro-Fuzzy Inference System with XGBoost Model. In 2024 International Conference on Sustainable Communication Networks and Application (ICSCNA) (pp. 724-730). IEEE.

[62] Raja, W., Chandraprabha, K., Ruckmani3b, V., & Gowri4c, S. (2026). Implementing LSTM-RNN for improved diabetic dataset classification.

[63] Pulugu, D., Praveen, R. V. S., Moharana, R. L., Raju, M. N., Revathy, P., & Saravanan, K. (2024, December). IoT-integrated leaf analysis for early plant disease detection in agriculture using hybrid ensemble models. In 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC) (pp. 1-6). IEEE.

[64] Mohandas, R., Elavarasi, M., Praveen, R. V. S., Chittapragada, H., Nithiya, C., & Prabagar, S. (2024, December). Advanced cyber-attack detection in iot networks using deep belief networks and gradient boosting mechanisms. In 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC) (pp. 1-7). IEEE.

[65] Wange, N. J., Praveen, R. V. S., Babavali, S. F., Rao, P., Gamini, P., & Ramasamy, D. (2024, December). Optimizing Thermal Comfort and Energy Efficiency in Smart Buildings using GANs and Reinforcement Learning Models. In 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC) (pp. 1-6). IEEE.

[66] Vandana, C. P., Basha, S. A., Madiajagan, M., Jadhav, S., Matheen, M. A., & Maguluri, L. P. (2024). IoT resource discovery based on multi faected attribute enriched CoAP: smart office seating discovery. Wireless Personal Communications, 1-18.

[67] Subha, B., Nagarajan, S., Thangam, A., Srinivas, V. S., Praveen, R. V. S., & Bakshi, K. D. (2024, December). Fuzzy Clustering and Autoencoder-Based Cybersecurity for EVs in a Blockchain-Enabled Smart Cloud Environment. In 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC) (pp. 1-7). IEEE.

[68] Sreelatha, B., Sheela, D. V., Praveen, R. V. S., Maheswari, K., Saravanan, A., & Singh, K. (2024, December). Smart Agriculture: IoT-Driven Soil Monitoring and Fertilizer Recommendation with CNN and Bidirectional GRU Models. In 2024 4th International Conference on Mobile Networks and Wireless Communications (ICMNWC) (pp. 01-06). IEEE.

Cite this article as: Dr. J. Mohammed Ubada., (2026). AI-Enabled Adaptive Multi-Factor Authentication with Homomorphic Encryption for Privacy-Preserving Cloud Access Management. International Journal of Emerging Knowledge Studies. 5(6), pp. 936–952.
https://doi.org/10.70333/ijeks-05-07-006