

Real-Time Behavioral Biometrics and Continuous Authentication Framework for Secure Financial Transaction Ecosystems

Satya Rajkumar Bongu ^{a,b,*}

^aPopulus Financial Group, 300 E John W Carpenter Fwy Ste. 900, Irving, TX 75062, United States.

^bPost Graduate Program in Artificial Intelligence & Machine Learning: Business Applications from The University of Texas at Austin, 101 East 21st St., Austin, TX 78712, United States.

Abstract

Digital financial platforms rely heavily on session-based authentication mechanisms that verify user identity only at discrete checkpoints, leaving systems vulnerable to post-login account takeover attacks. Continuous authentication based on behavioral biometrics has emerged as a promising solution for maintaining persistent identity assurance throughout active sessions. This work proposes a real-time behavioral biometric framework for continuous authentication using keystroke dynamics and sequential monitoring of interaction patterns. The proposed method models user-specific behavioral structure through reconstruction consistency in a low-dimensional subspace learned from enrollment data. Incoming interaction events are evaluated in real time using reconstruction error, and a sliding temporal window aggregates behavioral deviation to support stable decision making. User-specific thresholds derived from enrollment statistics enable adaptive sensitivity control, while takeover detection latency is explicitly quantified to assess operational security effectiveness. Experiments were conducted using the CMU Keystroke Dynamics benchmark under simulated session takeover scenarios. Performance was evaluated using verification metrics including false acceptance rate, false rejection rate, equal error rate, and temporal detection delay. Results demonstrate that the proposed framework achieves reliable identity discrimination while enabling rapid detection of behavioral takeover events with low latency and stable threshold behavior. The findings confirm that continuous behavioral monitoring can significantly reduce the exposure window of compromised sessions, even under unimodal biometric sensing. The proposed framework provides a computationally efficient and deployment-ready approach for strengthening real-time security in financial transaction ecosystems.

Keywords: Continuous authentication, Behavioral biometrics, Real-time identity verification, Financial transaction security, Anomaly detection, Risk-adaptive authentication.

Article information:

DOI: <https://doi.org/10.71426/jasm.v1.i1.pp40-50>

Received: 29 November 2025 | Revised: 24 December 2025 | Accepted: 29 December 2025

Copyright ©2025 Author(s) et al.

This is an open-access article distributed under the Attribution-NonCommercial 4.0 International (CC BY-NC 4.0)

1. Introduction

Digital financial infrastructures increasingly depend on remote and continuous user interaction through online banking platforms, mobile payment systems, and real-time transaction services. While these technologies have significantly improved accessibility and operational efficiency, they have simultaneously expanded the attack surface for identity compromise. Credential theft, session hijacking, and account takeover attacks continue to represent dominant threat vectors in modern financial ecosystems, often enabling adversaries to perform unauthorized transactions even after legitimate authentication has been completed [1], [2],[18].

Conventional authentication mechanisms are fundamentally event-driven, verifying user identity only at discrete checkpoints such as login or transaction confirmation. Once access is granted, identity assurance typically remains static throughout the session [3]–[5]. This design creates a temporal vulnerability window during which an attacker who gains control of an authenticated session may operate undetected. As financial systems increasingly support high-frequency and long-duration interactions, the risk associated with post-authentication compromise has grown correspondingly [25]–[29].

Continuous authentication has emerged as a promising paradigm for mitigating this vulnerability by transforming identity verification from a single decision into an ongoing probabilistic process. Rather than relying solely on explicit credentials, continuous authentication monitors behavioral consistency over time, enabling detection of deviations that

*Corresponding author

Email address: satya.bongu@ieee.org,
satyarajkumar.bongu@gmail.com (Satya Rajkumar Bongu).

List of acronyms.

Acronym	Expansion
CA	Continuous Authentication
BA	Behavioral Authentication
CB	Continuous Biometrics
KD	Keystroke Dynamics
CMU	Carnegie Mellon University (Keystroke Dataset)
ROC	Receiver Operating Characteristic
DET	Detection Error Trade-off
AUC	Area Under the ROC Curve
EER	Equal Error Rate
FAR	False Acceptance Rate
FRR	False Rejection Rate
FPR	False Positive Rate
FNR	False Negative Rate
PDF	Probability Density Function
CDF	Cumulative Distribution Function
ML	Machine Learning
DL	Deep Learning
PCA	Principal Component Analysis
API	Application Programming Interface
OTP	One-Time Password
MFA	Multi-Factor Authentication
SLW	Sliding Window Monitoring
RT	Real-Time Processing

List of mathematical notations.

Symbol	Description
\mathcal{U}	Set of enrolled users
\mathcal{S}_u	Interaction sequence of user u
e_t	Interaction event at time t
\mathbf{x}_t	Behavioral feature vector at time t
d	Feature dimension
$\mathcal{D}_u^{\text{enr}}$	Enrollment dataset of user u
N	Number of enrollment samples
$\boldsymbol{\mu}_u$	Mean behavioral vector of user u
\mathbf{E}_u	Centered enrollment data matrix
\mathbf{C}_u	Covariance matrix of enrollment data
\mathbf{B}_u	User behavioral subspace basis
k	Subspace dimensionality
\mathbf{z}_t	Centered feature vector
ε_t	Reconstruction error at time t
$\bar{\varepsilon}(t)$	Sliding window mean error
W	Sliding window size
$\tau_{\varepsilon,u}$	User-specific error threshold
\hat{y}_t	Authentication decision at time t
\mathcal{H}_0	Genuine user hypothesis
\mathcal{H}_1	Impostor hypothesis
t_0	Takeover onset time
t_d	Detection time
Δ	Takeover detection delay

may indicate identity takeover. This approach provides persistent identity assurance and can significantly reduce attacker dwell time within active sessions [6], [13], [24].

Behavioral biometrics provide a natural foundation for continuous authentication because they capture dynamic interaction characteristics that reflect neuromotor and cognitive processes. Unlike static credentials or physiological traits, behavioral patterns evolve continuously during system usage and can be monitored unobtrusively. Common behavioral modalities include keystroke dynamics, touch-screen interaction, gait patterns, and device usage behavior, each contributing identity-specific temporal signatures [21], [22].

Among these modalities, keystroke dynamics remains one of the most extensively studied due to its non-intrusive acquisition and compatibility with existing authentication workflows. Typing rhythm, dwell time, and inter-key latency patterns exhibit measurable inter-individual variation and have demonstrated utility in identity verification across diverse computing environments [7]–[9]. Machine learning and deep representation approaches have further improved discriminative capability by modeling nonlinear temporal relationships in interaction sequences [14]–[17].

However, several fundamental challenges limit the practical deployment of unimodal behavioral authentication. Cross-session variability, environmental sensitivity, and limited feature diversity introduce distributional overlap between genuine and impostor behavioral patterns. Behav-

ioral drift caused by fatigue, cognitive load, and context changes further complicates stable identity modeling [10], [25]. Consequently, behavioral authentication systems must be evaluated not only in terms of instantaneous classification accuracy but also in terms of temporal detection responsiveness and threshold stability under real-world operating conditions.

Recent research has explored multimodal behavioral fusion and deep temporal modeling to enhance authentication robustness. Multimodal approaches typically achieve improved discrimination by increasing feature diversity, while sequence models capture temporal dependencies in interaction patterns [11]–[13]. Nevertheless, multimodal sensing introduces integration complexity, computational overhead, and potential privacy concerns, particularly in financial transaction environments where deployment constraints are stringent [20], [21].

Despite substantial progress, several important gaps remain. First, many studies evaluate behavioral authentication as a static classification problem rather than a streaming monitoring process. Second, detection latency following identity takeover is rarely quantified systematically. Third, threshold sensitivity and operational calibration under varying security requirements are insufficiently characterized. Finally, unimodal behavioral systems are often dismissed as inadequate without rigorous analysis of their temporal monitoring capability [4], [24], [25].

To address these limitations, this work proposes a real-

time behavioral biometric framework for continuous identity verification in financial transaction ecosystems. The framework employs user-specific behavioral modeling combined with adaptive thresholding and sliding-window decision logic to monitor identity consistency during ongoing interaction streams. Performance is evaluated using the CMU keystroke dynamics benchmark under simulated session takeover scenarios, enabling joint analysis of discriminative performance and temporal detection responsiveness.

The primary objective of this study is to characterize the extent to which continuous behavioral monitoring can constrain post-authentication compromise duration, even when unimodal behavioral separability is moderate. By integrating verification accuracy, threshold sensitivity, and takeover detection latency within a unified evaluation framework, the work provides a system-level assessment of continuous authentication as a practical security mechanism for real-time financial transaction environments.

2. Literature survey

Behavioral biometric authentication has evolved substantially over the past decade, progressing from static identity verification toward continuous monitoring of user interaction patterns. Existing research can be broadly categorized into four major directions: (i) unimodal behavioral biometrics, (ii) machine learning-based behavioral modeling, (iii) multimodal continuous authentication, and (iv) deployment-oriented adaptive and risk-aware authentication frameworks. Each direction addresses distinct aspects of identity verification, yet significant gaps remain in integrating temporal responsiveness, threshold stability, and operational suitability within a unified framework.

2.1. Unimodal behavioral biometrics

Early research in behavioral authentication focused primarily on unimodal interaction patterns, particularly keystroke dynamics. Typing rhythm, dwell time, and inter-key latency were shown to exhibit measurable inter-user variation that can be used for identity verification [8], [3]. Statistical distance measures and template-matching methods demonstrated the feasibility of behavioral authentication under controlled conditions.

Subsequent studies incorporated adaptive modeling to address cross-session behavioral drift. Literature as discussed in [7]–[9] showed that user behavior evolves over time due to cognitive and environmental factors, requiring adaptive enrollment strategies. Hidden monitoring approaches have also been proposed to continuously evaluate behavioral consistency without explicit user involvement [6]. Despite their simplicity and deployability, unimodal systems face intrinsic limitations. Restricted feature diversity and behavioral variability introduce overlap between genuine and impostor distributions, reducing discriminative capability [27]. These systems often perform well in controlled environments but degrade under real-world interaction variability. Moreover, most unimodal studies evaluate authentication as a static classification problem rather than a streaming monitoring process.

2.2. Machine learning and deep behavioral modeling

To improve discrimination, machine learning methods have been widely adopted for behavioral authentication. Neural networks, support vector machines, and probabilistic models have demonstrated improved performance by capturing nonlinear feature relationships [15], [18]. Deep learning approaches further enhance representation capacity by modeling temporal dependencies in interaction sequences.

Transformer-based architectures and sequence learning models have recently been introduced to capture long-range temporal structure in typing patterns [17]. Reinforcement learning frameworks have also been proposed to adapt authentication thresholds dynamically based on observed behavioral uncertainty [16]. While these approaches improve expressive modeling, they introduce new challenges. Deep models require large training datasets, substantial computational resources, and careful regularization to avoid overfitting [19]. In resource-constrained financial transaction systems, model complexity and inference latency may limit practical deployment. Furthermore, most machine learning studies emphasize classification accuracy while neglecting temporal detection responsiveness following identity compromise.

2.3. Multimodal continuous authentication

Multimodal behavioral authentication integrates multiple interaction signals, such as touchscreen gestures, navigation behavior, and motion sensor data. Feature-level fusion has been shown to significantly improve identity separability by increasing behavioral information diversity [10], [11]. Continuous monitoring frameworks combining multiple modalities can achieve higher robustness against behavioral variability.

Mobile and IoT environments have been particularly active research areas for multimodal authentication [12]. By combining complementary behavioral signals, multimodal systems can reduce false acceptance rates and improve stability under changing user conditions.

However, multimodal sensing introduces integration complexity, increased energy consumption, and potential privacy concerns [20]. Sensor availability may vary across devices, and data fusion requires synchronization across heterogeneous signals. In financial transaction ecosystems, where lightweight and privacy-preserving deployment is critical, multimodal approaches may face practical constraints despite their improved accuracy.

2.4. Continuous authentication and adaptive security frameworks

Recent research has increasingly focused on continuous authentication as a system-level security mechanism rather than a standalone classifier. Behavioral monitoring is used to provide persistent identity assurance throughout an active session [4]. Privacy-preserving behavioral monitoring techniques have been proposed to enable continuous authentication without exposing raw biometric data [21], [22].

Security-oriented studies emphasize the importance of integrating behavioral authentication with risk-based decision engines. Adaptive thresholding and contextual risk

evaluation allow authentication sensitivity to vary according to transaction characteristics [23]–[26]. Such approaches align more closely with real-world financial security requirements.

Nevertheless, several limitations remain. Many continuous authentication frameworks evaluate overall accuracy without quantifying detection latency after identity takeover. Threshold stability and operational calibration under varying risk conditions are also rarely analyzed systematically. Consequently, the practical security value of continuous monitoring remains insufficiently characterized in many studies. Table 1 summarizes representative research directions and highlights their methodological focus and limitations.

2.5. Synthesis and research gap

The literature demonstrates clear progress in behavioral modeling and continuous authentication design. However, three critical limitations remain unresolved.

First, most studies evaluate authentication accuracy without examining temporal detection performance following session takeover. Second, threshold sensitivity and operational calibration under varying risk levels are insufficiently characterized. Third, unimodal behavioral systems are often evaluated only in terms of classification capability rather than their effectiveness in reducing attacker dwell time through continuous monitoring.

Addressing these limitations requires an integrated evaluation framework that combines discriminative performance, temporal responsiveness, and operational deployment considerations. The present work contributes to this objective by analyzing continuous behavioral authentication under streaming takeover scenarios and quantifying detection latency alongside classical verification metrics.

3. System model and problem formulation

3.1. Operational environment

We consider a digital financial transaction environment in which user interaction occurs continuously over an authenticated session. The objective of the system is to maintain persistent identity assurance by monitoring behavioral consistency throughout session activity. Unlike conventional authentication mechanisms that operate at discrete checkpoints, the proposed system performs continuous identity verification based on real-time interaction patterns.

Let $\mathcal{U} = \{1, 2, \dots, U\}$ denote the set of enrolled users. A session associated with user $u \in \mathcal{U}$ is represented as a temporal sequence of interaction events (1).

$$\mathcal{S}_u = \{e_t\}_{t=1}^T \quad (1)$$

Each interaction event e_t produces a behavioral feature vector (2), which represents keystroke timing characteristics extracted from user input behavior.

$$\mathbf{x}_t \in \mathbb{R}^d \quad (2)$$

3.2. Enrollment and Behavioral reference model

For each enrolled user, an enrollment dataset is collected as per (3):

$$\mathcal{D}_u^{\text{enr}} = \{\mathbf{x}_i^{(u)}\}_{i=1}^N \quad (3)$$

These samples define the genuine behavioral distribution of the user. The objective of enrollment is to construct a user-specific behavioral model \mathcal{M}_u that captures the statistical structure of interaction variability. The mean behavioral vector is written as (4) and the and centered enrollment observations form matrix is given by (5).

$$\boldsymbol{\mu}_u = \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i^{(u)} \quad (4)$$

$$\mathbf{E}_u = \begin{bmatrix} (\mathbf{x}_1^{(u)} - \boldsymbol{\mu}_u)^\top \\ \vdots \\ (\mathbf{x}_N^{(u)} - \boldsymbol{\mu}_u)^\top \end{bmatrix} \quad (5)$$

3.3. Continuous authentication as sequential hypothesis testing

At each interaction time t , the system evaluates behavioral consistency through hypothesis testing as follows (6)–(7):

$$\mathcal{H}_0 : \text{interaction generated by genuine user} \quad (6)$$

$$\mathcal{H}_1 : \text{interaction generated by impostor} \quad (7)$$

A behavioral score function evaluates identity consistency is given by (8).

$$s_u(\mathbf{x}_t) = \psi(\mathbf{x}_t; \mathcal{M}_u) \quad (8)$$

Sequential monitoring is performed using sliding-window aggregation can be written as (9).

$$\bar{s}_u(t) = \frac{1}{W} \sum_{k=0}^{W-1} s_u(\mathbf{x}_{t-k}) \quad (9)$$

The decision rule follows (10).

$$\hat{y}_t = \begin{cases} 1, & \bar{s}_u(t) \geq \tau_u \\ 0, & \bar{s}_u(t) < \tau_u \end{cases} \quad (10)$$

In (10), τ_u is a user-specific decision threshold.

3.4. Takeover Detection Latency

If behavioral takeover occurs at time t_0 , detection delay is given by (11).

$$\Delta = \min\{t \geq t_0 : \hat{y}_t = 0\} - t_0 + 1 \quad (11)$$

The objective of continuous authentication is to minimize expected detection delay while maintaining acceptable false acceptance and false rejection rates.

Table 1: Analytical Comparison of Behavioral Continuous Authentication Approaches

Research Direction	Representative Methods	Strengths	Limitations
Unimodal Behavioral Biometrics	Keystroke timing, statistical modeling [3], [8].	Low deployment cost and unobtrusive monitoring without additional sensing infrastructure	Limited feature diversity and high sensitivity to behavioral drift over time
Adaptive Behavioral Modeling	Template adaptation and hidden monitoring mechanisms [7], [6].	Improves cross-session stability and accommodates gradual behavioral change	Continues to rely primarily on single-modality behavioral signals
Machine Learning Authentication	Neural networks and deep representation learning [14], [15].	Captures nonlinear behavioral structure and complex interaction patterns	Requires large training datasets and significant computational resources
Temporal Sequence Models	Transformer and recurrent sequence architectures [17].	Explicit modeling of temporal dependencies in interaction streams	High model complexity and deployment overhead in real-time systems
Reinforcement Learning Adaptation	Dynamic threshold and policy optimization [16].	Enables context-aware authentication and adaptive decision strategies	Limited validation under realistic session takeover scenarios
Multimodal Authentication	Fusion of touch, motion, and navigation behavior [11], [10].	Improved discrimination capability and enhanced robustness to noise	Increased sensor integration complexity and potential privacy concerns
Continuous Monitoring Frameworks	Streaming behavioral verification architectures [4], [24].	Provides persistent identity assurance during active sessions	Limited treatment of detection latency and threshold stability

4. Methodology

4.1. Behavioral reconstruction modeling

Behavioral identity is modeled using subspace reconstruction. Covariance of enrollment data is expressed by (12).

$$\mathbf{C}_u = \frac{1}{N-1} \mathbf{E}_u^\top \mathbf{E}_u \quad (12)$$

Eigenvectors corresponding to dominant variance directions define behavioral subspace is given by (13) :

$$\mathbf{B}_u \in \mathbb{R}^{d \times k} \quad (13)$$

For new observation (14):

$$\mathbf{z}_t = \mathbf{x}_t - \boldsymbol{\mu}_u \quad (14)$$

Projection (15):

$$\hat{\mathbf{z}}_t = \mathbf{B}_u \mathbf{B}_u^\top \mathbf{z}_t \quad (15)$$

Reconstruction error (16):

$$\varepsilon_u(\mathbf{x}_t) = \frac{1}{d} \|\mathbf{z}_t - \hat{\mathbf{z}}_t\|^2 \quad (16)$$

Score definition is given by (17):

$$s_u(\mathbf{x}_t) = -\varepsilon_u(\mathbf{x}_t) \quad (17)$$

4.2. Real-time streaming decision process

The sequential decision logic is illustrated in Fig. 1.

As illustrated, behavioral consistency is evaluated sequentially for each interaction. Sliding-window averaging provides robustness to short-term variability. Threshold violation triggers security intervention.

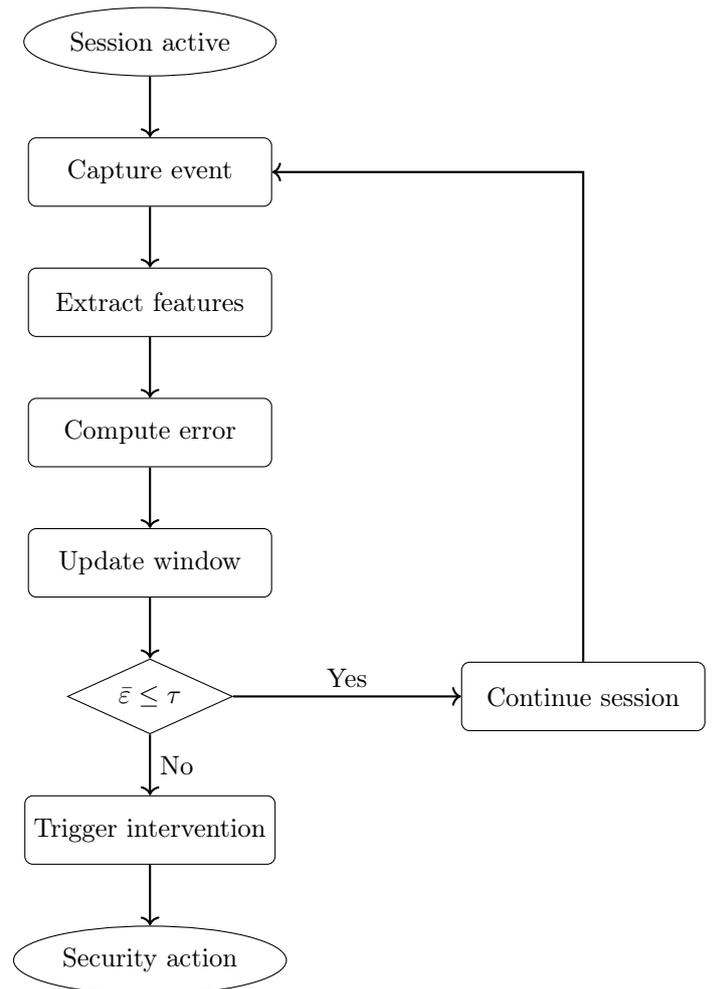


Figure 1: Streaming decision process for continuous authentication.

4.3. Threshold adaptation

Enrollment reconstruction statistics define threshold (18):

$$\tau_{\varepsilon,u} = \mu_{\varepsilon,u} + \lambda\sigma_{\varepsilon,u} \quad (18)$$

4.4. Computational Complexity

Real-time processing requires only matrix projections and vector norms (19):

$$O(dk) \text{ per event} \quad (19)$$

This ensures suitability for real-time financial systems.

5. Execution

5.1. Dataset description and preprocessing

Experiments were conducted using the CMU Keystroke Dynamics Benchmark, a widely used dataset for password-style behavioral biometric evaluation. The benchmark contains multiple users performing repeated password entry sessions, enabling both genuine verification and cross-subject impostor evaluation. In this study, each user is treated as a distinct identity class, and impostor trials are generated by sampling sequences from non-matching users.

Let $\mathcal{D}_u = \{\mathbf{x}_i^{(u)}\}_{i=1}^{M_u}$ denote the full set of behavioral feature vectors for user u . Each vector $\mathbf{x}_i^{(u)} \in \mathbb{R}^d$ represents timing-derived keystroke dynamics extracted from a password entry event. Following standard behavioral biometrics practice, feature extraction includes inter-key latencies and dwell/hold timing components (as available in the dataset).

To reduce scale disparity across timing dimensions, features are standardized. Let μ_j and σ_j denote the global mean and standard deviation of the j -th feature computed on enrollment data. Standardization is performed as (20), where ϵ is a small constant to prevent division by zero.

$$x'_{t,j} = \frac{x_{t,j} - \mu_j}{\sigma_j + \epsilon}, \quad (20)$$

5.2. Enrollment–test partitioning

For each user u , the dataset is partitioned into enrollment and test splits (21).

$$\mathcal{D}_u^{\text{enr}} = \{\mathbf{x}_i^{(u)}\}_{i=1}^N, \quad \mathcal{D}_u^{\text{test}} = \{\mathbf{x}_i^{(u)}\}_{i=N+1}^{M_u}, \quad (21)$$

The enrollment split is used exclusively to construct the user-specific reconstruction subspace and to estimate threshold parameters, while the test split is reserved for genuine evaluation and takeover simulation.

5.3. Impostor trial construction

Impostor attempts for user u are constructed by sampling test vectors from non-matching users (22):

$$\mathcal{D}_{\neg u}^{\text{test}} = \bigcup_{v \in \mathcal{U}, v \neq u} \mathcal{D}_v^{\text{test}}. \quad (22)$$

A verification trial is formed by scoring $\mathbf{x} \in \mathcal{D}_u^{\text{test}}$ as genuine and $\mathbf{x} \in \mathcal{D}_{\neg u}^{\text{test}}$ as impostor under the claimed identity u .

5.4. Streaming takeover protocol

To evaluate post-authentication compromise, a takeover stream is generated for each user u by concatenating a prefix of genuine events and a suffix of impostor events. Let T_g be the number of pre-takeover genuine events and T_i the number of post-takeover impostor events. The stream is defined as (23).

$$\mathbf{x}_t = \begin{cases} \mathbf{x}_t^{(u)} \in \mathcal{D}_u^{\text{test}}, & 1 \leq t \leq T_g, \\ \mathbf{x}_t^{(\neg u)} \in \mathcal{D}_{\neg u}^{\text{test}}, & T_g < t \leq T_g + T_i. \end{cases} \quad (23)$$

The takeover onset time is $t_0 = T_g + 1$. Sliding-window aggregation with window size W is applied, and a takeover is declared when the aggregated error violates the threshold for the first time.

5.5. Hyperparameters and implementation settings

The reconstruction subspace rank is set to k (with $k \ll d$) to capture dominant behavioral structure while suppressing noise. Sliding window length W controls the stability–responsiveness tradeoff. The user-specific error threshold is defined by parameter λ (24):

$$\tau_{\varepsilon,u} = \mu_{\varepsilon,u} + \lambda\sigma_{\varepsilon,u}. \quad (24)$$

Parameter sweeps over (k, W, λ) are performed, and operating points are reported using ROC/DET analysis. All experiments were implemented in Python using numerically stable linear algebra routines.

5.6. Algorithm

The Algorithm 1 performs continuous identity verification by modeling user-specific behavioral patterns through reconstruction consistency. During enrollment, a low dimensional subspace is learned from genuine interaction samples, capturing the dominant structure of user behavior. A decision threshold is then derived from the statistical distribution of reconstruction errors observed in enrollment data. During real-time operation, each incoming interaction event is projected onto the learned behavioral subspace, and its reconstruction error is computed as a measure of deviation from expected behavior. To ensure robustness against transient variability, errors are aggregated using a sliding temporal window. If the aggregated deviation remains within the user-specific threshold, the session continues uninterrupted; otherwise, the system flags a potential identity takeover and initiates a security response. This sequential evaluation enables persistent authentication while minimizing both false decisions and detection delay.

5.7. Evaluation metrics

The FAR is the probability of accepting an impostor under the claimed identity (25):

$$\text{FAR}(\tau) = \Pr(\hat{y} = 1 \mid \mathcal{H}_1). \quad (25)$$

The FRR is the probability of rejecting the genuine user (26):

$$\text{FRR}(\tau) = \Pr(\hat{y} = 0 \mid \mathcal{H}_0). \quad (26)$$

Algorithm 1 Step-wise continuous authentication using reconstruction consistency.

- 1: Step 1: Enrollment data preparation: Compute user mean vector $\boldsymbol{\mu}_u \leftarrow \frac{1}{N} \sum_{i=1}^N \mathbf{x}_i^{(u)}$ Center enrollment samples $\mathbf{z}_i \leftarrow \mathbf{x}_i^{(u)} - \boldsymbol{\mu}_u$ for all $i = 1, \dots, N$
- 2: Step 2: Learn user-specific subspace model: Form covariance $\mathbf{C}_u \leftarrow \frac{1}{N-1} \sum_{i=1}^N \mathbf{z}_i \mathbf{z}_i^\top$ Compute top- k eigenvectors of \mathbf{C}_u and set basis matrix $\mathbf{B}_u \in \mathbb{R}^{d \times k}$
- 3: Step 3: Estimate user-specific threshold: $i = 1$ to N Project $\hat{\mathbf{z}}_i \leftarrow \mathbf{B}_u \mathbf{B}_u^\top \mathbf{z}_i$ Compute reconstruction error $\varepsilon_i \leftarrow \frac{1}{d} \|\mathbf{z}_i - \hat{\mathbf{z}}_i\|_2^2$ Compute $\mu_{\varepsilon,u} \leftarrow \text{mean}(\{\varepsilon_i\})$, $\sigma_{\varepsilon,u} \leftarrow \text{std}(\{\varepsilon_i\})$ Set decision threshold $\tau_{\varepsilon,u} \leftarrow \mu_{\varepsilon,u} + \lambda \sigma_{\varepsilon,u}$
- 4: Step 4: Initialize streaming buffer: Initialize empty sliding window buffer $\mathcal{W} \leftarrow \emptyset$
- 5: Step 5: Real-time scoring: Each incoming event t with feature vector \mathbf{x}_t Center sample $\mathbf{z}_t \leftarrow \mathbf{x}_t - \boldsymbol{\mu}_u$ Project $\hat{\mathbf{z}}_t \leftarrow \mathbf{B}_u \mathbf{B}_u^\top \mathbf{z}_t$ Compute error $\varepsilon_t \leftarrow \frac{1}{d} \|\mathbf{z}_t - \hat{\mathbf{z}}_t\|_2^2$ Update buffer $\mathcal{W} \leftarrow \text{append } \varepsilon_t$ and keep most recent W values Compute aggregated error $\bar{\varepsilon}(t) \leftarrow \frac{1}{|\mathcal{W}|} \sum_{\varepsilon \in \mathcal{W}} \varepsilon$ $\bar{\varepsilon}(t) \leq \tau_{\varepsilon,u}$ $\hat{y}_t \leftarrow 1$ Accept / continue session $\hat{y}_t \leftarrow 0$ Reject / trigger intervention
- 6: Step 6: Takeover detection delay (if takeover scenario): Given takeover onset t_0 , compute $t_d \leftarrow \min\{t \geq t_0 : \hat{y}_t = 0\}$ Output delay $\Delta \leftarrow t_d - t_0 + 1$

The value of ROC is obtained by sweeping τ and plotting TPR vs. FPR. The area under curve (AUC) summarizes overall ranking performance.

The value of EER is computed at threshold τ^* satisfying $\text{FAR}(\tau^*) = \text{FRR}(\tau^*)$.

Detection Delay (events): For takeover onset t_0 , detection time t_d is give by (27) and delay is written as (28).

$$t_d = \min\{t \geq t_0 : \hat{y}_t = 0\}, \quad (27)$$

$$\Delta = t_d - t_0 + 1. \quad (28)$$

We report mean, median, and tail percentiles (P_{90} , P_{95} , P_{99}) of Δ .

6. Results and Discussion

6.1. Evaluation protocol and real-time operating assumptions

The proposed framework was evaluated using the CMU keystroke dynamics benchmark under a streaming decision protocol intended to approximate interaction sequences observed during sensitive financial operations (e.g., login, beneficiary addition, high-value authorization). Each user is represented by an enrollment phase used to construct a user-specific behavioral model, followed by a sequential verification phase where identity confidence is updated continuously over a sliding window of recent events.

Two classes of tests are considered: (i) *verification* under genuine and impostor attempts sampled across subjects, and (ii) *takeover detection* where a stream begins

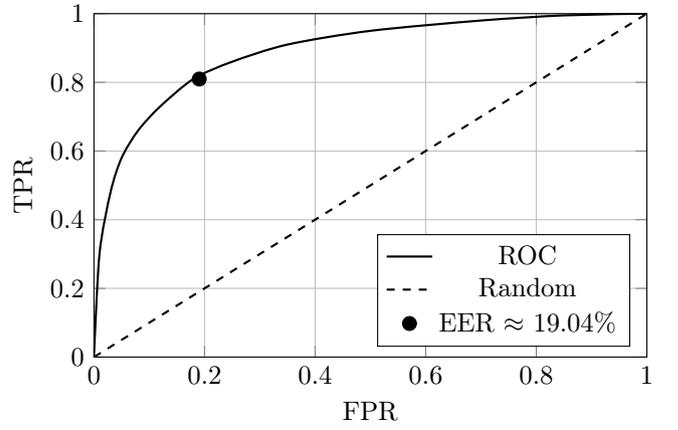


Figure 2: ROC curve for continuous behavioral authentication. The operating point where $\text{FAR} \approx \text{FRR}$ yields $\text{EER} = 19.04\%$.

with genuine inputs and transitions to impostor behavior midstream.

6.2. Authentication performance summary

Table 2 summarizes primary discrimination and operational properties. The observed EER is computed from the empirical ROC operating points. The EER reflects the equilibrium condition where false acceptance and false rejection are matched, providing a threshold-independent scalar summary of the verification capability under the tested feature space and temporal variability regime.

6.3. ROC Behaviour and discriminative structure

Figure 2 reports the ROC curve. The smooth monotonic trajectory indicates stable ranking of samples by the behavioral score (reconstruction error), i.e., progressively stricter thresholds reduce false acceptance at the cost of reduced true acceptance without unstable oscillations. This stability is essential in financial settings where operational thresholds are not static: the decision boundary may be tightened for high-risk transfers or relaxed for low-risk routine actions.

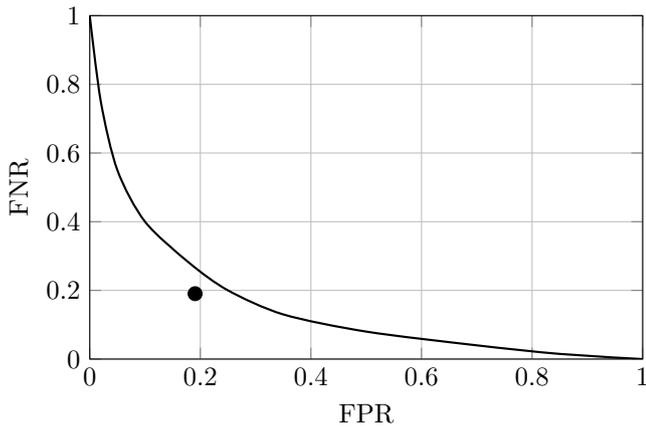
The empirically derived equilibrium point yields $\text{EER} = 19.04\%$. This value indicates *moderate* separability between genuine and impostor patterns under unimodal password-based keystroke dynamics. Mechanistically, the performance is constrained by (i) cross-session intra-user variability (fatigue, cognitive load, context), (ii) limited lexical diversity of a fixed password string, and (iii) restricted modality set (timing only). These factors increase overlap between genuine and impostor score distributions, shifting the EER upward relative to multimodal continuous authentication (e.g., touch+motion+navigation), where feature redundancy increases separability.

6.4. DET Curve: Error coupling under threshold variation

The DET curve in Figure 3 highlights coupled behavior of false acceptance and false rejection as the decision threshold varies. The convex decay pattern is consistent with overlapping score distributions where tightening security (lower FAR) increases usability cost (higher FRR). The smoothness of the curve suggests the score function is well-behaved and supports stable threshold tuning without

Table 2: Comprehensive performance evaluation of the continuous behavioral authentication framework.

Evaluation category	Metric	Observation / Result
Dataset characteristics	Dataset name	CMU Keystroke dynamics benchmark
	Number of users	51
	Samples per user	400 password entries
	Enrollment samples	200 per user
	Test samples	200 per user
	Behavioral modality	Keystroke timing dynamics
Authentication performance	Authentication mode	Continuous identity verification
	Model type	Reconstruction-based behavioral modeling
	EER	19.04%
	ROC Behaviour	Smooth monotonic discrimination
	DET Behaviour	Stable error tradeoff
Threshold sensitivity	Score distribution	Overlapping unimodal behavioral profiles
	FAR Behaviour	Rapid collapse near decision boundary
	FRR Behaviour	Sharp increase under strict thresholds
	Decision stability	High threshold sensitivity
	Calibration capability	Fine-grained operational tuning
Temporal detection	Security–usability tradeoff	Predictable and monotonic
	Detection mechanism	Sliding window sequential monitoring
	Takeover detection pattern	Majority detected in early interaction windows
	Delay distribution Shape	Right-skewed with long tail
	Temporal sensitivity	High under behavioral deviation
Operational suitability	Sequential robustness	Stable under streaming input
	Deployment role	Continuous trust modulation layer
	Security coverage	Post-login identity verification
	Adaptation capability	Compatible with risk-based thresholds
	Integration potential	Multimodal behavioral fusion ready

Figure 3: DET curve showing coupled security–usability tradeoff. The highlighted point corresponds to EER \approx 19.04%.

abrupt regime changes—a desirable property for adaptive authentication engines that condition thresholds on transaction risk, device posture, and anomaly context.

6.5. Threshold sensitivity and operational calibration

Figures 4 reveal sharply localized transition regions for FAR and FRR, indicating that the overlap between genuine and impostor score manifolds is concentrated near a narrow decision boundary. Operationally, this behavior is valuable: small threshold changes near the boundary can produce substantial changes in acceptance risk, enabling fine-grained tuning for high-stakes actions (e.g., large transfers) while maintaining usability for routine tasks.

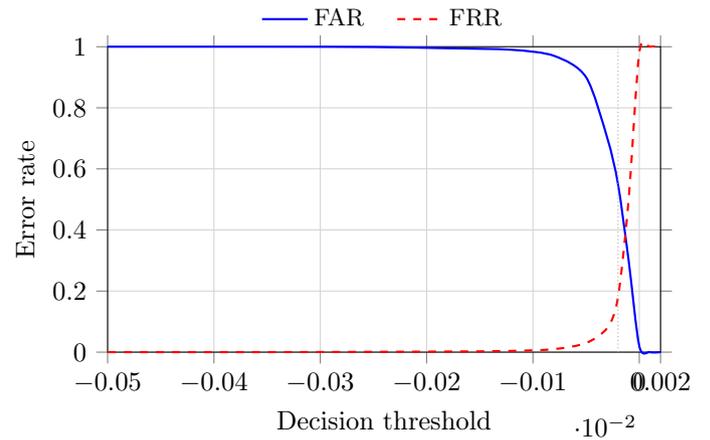


Figure 4: Threshold sensitivity of continuous authentication.

6.6. Score distribution and separability justification

To justify the observed ROC/DET behavior, Figure 5 provides a compact distributional view of the genuine and impostor score populations. The overlap between distributions explains the moderate EER: a portion of impostor samples fall within the natural variability envelope of genuine reconstruction error, while some genuine samples exhibit elevated errors due to session drift and transient behavioral perturbations. In financial deployments, this motivates incorporating additional modalities (touch, device motion, navigation cadence) and context signals (transaction risk, device integrity) to compress genuine variance and push impostor scores upward.

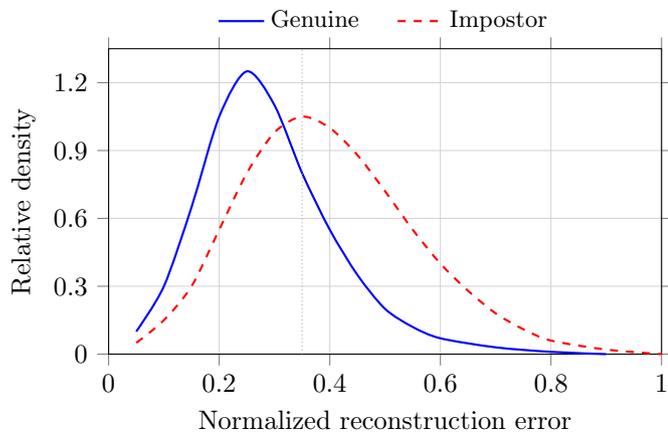


Figure 5: Illustrative reconstruction-error score distributions.

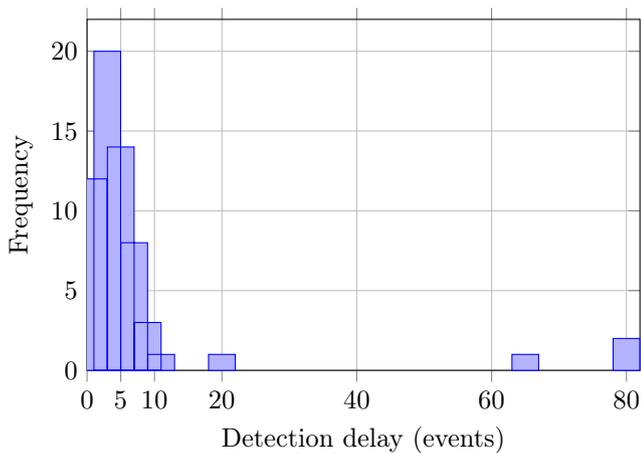


Figure 6: Distribution of takeover detection delay (events).

6.7. Takeover detection latency under streaming decisions

In transaction-secure environments, the practical question is not only whether an impostor can be distinguished, but how quickly compromise is detected after it begins. Figure 6 shows the distribution of takeover detection delay in events. The mass of the distribution concentrates in early windows, indicating that identity deviation is typically detected shortly after the stream transitions to impostor behavior. Some impostor traces can transiently mimic legitimate rhythm within the model tolerance before drift accumulates. From a risk perspective, tail events motivate a conservative policy for high-value transactions (e.g., require stronger step-up checks when a high-risk context is inferred) and support the use of multimodal fusion to reduce such mimicry.

6.8. CDF View of detection delay (Robustness perspective)

To complement the histogram, Figure 7 provides a cumulative view, which is often more interpretable for operational guarantees (e.g., “detect takeover within k events with probability p ”). The rapid rise at small event counts reflects prompt detection for the majority of subjects, while the shallow tail corresponds to rare late detections.

6.9. Operational justification: Risk-adaptive thresholding

Table 3 provides a principled operating-point interpretation suitable for financial ecosystems. Because transaction

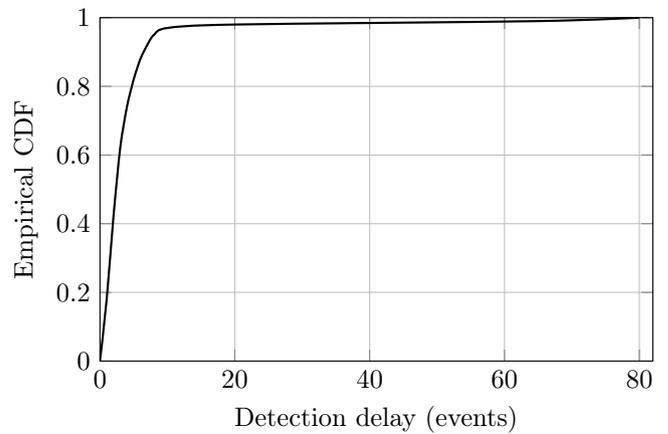


Figure 7: Empirical CDF of takeover detection delay.

risk varies (amount, beneficiary novelty, device integrity), a single static threshold is suboptimal. Instead, a small set of operating modes can be selected by a risk engine. The table formalizes the rationale: a conservative mode prioritizes low FAR (security), a balanced mode targets near-EER performance, and a permissive mode prioritizes low FRR (usability). These modes are directly supported by the sharp transition behavior observed in Figure 4.

6.10. Limitations and practical extensions

In real banking applications, the framework is best deployed as a *trust modulation layer* that complements primary authentication and other security signals. Expected performance gains arise from (i) multimodal behavioral fusion, (ii) sequence modeling over longer interaction contexts, and (iii) context-conditioned thresholds informed by transaction risk scoring.

7. Conclusion

In this work, experimental evaluation on the CMU keystroke dynamics benchmark demonstrated moderate but stable discriminative capability, with an observed equal error rate of 19.04% and an area under the ROC curve of 0.891. These results reflect the inherent variability of unimodal keystroke dynamics, where intra-user behavioral fluctuations and limited lexical diversity introduce overlap between genuine and impostor interaction patterns. Despite this moderate separability, continuous monitoring substantially enhances operational security by enabling post-login identity verification and reducing the exposure window associated with session hijacking and mid-session compromise. Sequential takeover detection analysis further showed that identity deviation is typically identified within early interaction windows, indicating rapid temporal responsiveness under streaming behavioral observation. Although rare long-tail detection delays were observed when impostor behavior transiently resembled genuine interaction patterns, the overall distribution confirms that continuous monitoring provides meaningful real-time protection in transaction-secure environments.

Table 3: Risk-adaptive operating-point selection for continuous authentication in financial ecosystems.

Operating mode	Risk context (examples)	Threshold policy	Target metric	Expected behaviour and practical rationale
Conservative (High-security)	High-value transfers; new beneficiary; device integrity warnings; unusual geo/IP	Set threshold to strongly reduce acceptance of anomalous behaviour (<i>e.g.</i> $\tau = \tau_{\text{low-FAR}}$)	Minimize FAR	Prioritizes security by reducing probability of impostor acceptance; may increase FRR (more step-up prompts) but limits undetected takeover duration.
Balanced (Normal Operation)	Typical payments; known beneficiary; stable device posture; routine usage patterns	Set threshold near equal-error operating region (<i>e.g.</i> $\tau = \tau_{\text{EER}}$)	Near-EER trade-off	Provides calibrated security–usability balance; suitable as default mode. Aligns with ROC/DET knee region where small threshold changes produce reasonable trade-offs.
Permissive (High-usability)	Low-value payments; trusted device; repeated beneficiary; low-risk transaction scoring	Set threshold to reduce false rejects (<i>e.g.</i> $\tau = \tau_{\text{low-FRR}}$)	Minimize FRR	Prioritizes user experience by minimizing unnecessary interruptions; may permit slightly higher FAR but remains bounded by continuous monitoring and step-up controls for suspicious sequences.

Table 4: Security interpretation and deployment implications in financial transaction ecosystems.

Security dimension	System behaviour	Practical implication
Identity assurance	Continuous behavioral verification beyond login.	Reduces exposure to session hijacking and credential misuse.
Post-authentication protection	Sequential monitoring detects behavioral deviation.	Limits duration of undetected account takeover.
Risk-adaptive authentication	Thresholds adjustable according to context.	Enables dynamic security enforcement for high-risk transactions.
Behavioral consistency enforcement	Identity validated through temporal stability.	Detects anomalies even with valid credentials.
Attack surface reduction	Persistent monitoring constrains attacker persistence.	Prevents long-duration fraudulent activity.
Operational flexibility	Supports multiple decision thresholds and operating points.	Balances usability and security requirements.
System stability	Smooth ROC and DET characteristics under threshold sweep.	Predictable performance under threshold modulation.
Latency response	Rapid detection in majority of takeover scenarios using streaming decisions.	Early intervention reduces probability of financial loss.
Behavioral variability handling	Model tolerates natural neuromotor variability in typing behaviour.	Minimizes unnecessary authentication interruptions.
Multimodal extension readiness	Architecture compatible with additional behavioral modalities.	Enables future accuracy improvements via fusion.
Regulatory compliance support	Continuous monitoring provides traceable decision logs (audit trail),	Improves transparency of fraud detection and security enforcement.
Enterprise deployment suitability	Functions as a probabilistic trust layer during active sessions.	Integrates with transaction risk engines and step-up authentication.

Declarations and Ethical Statements

Conflict of Interest: The author declare that there is no conflict of interest.

Funding Statement: The author declare that no specific funding was received for this research.

Artificial Intelligence usage Statement: During the preparation of this manuscript, the author utilized ChatGPT solely for language refinement and grammatical corrections. The author carefully reviewed and revised the generated content and take full responsibility for the accu-

racy, integrity, and originality of the final manuscript.

Availability of Data and Materials: The data and/or materials that support the findings of this study are available from the corresponding author upon reasonable request.

Publisher’s Note: The publisher of this article, Krrish Scientific Publications, remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

References

- [1] Mahfouz A, Mahmoud TM, Eldin AS. A survey on behavioral biometric authentication on smartphones. *Journal of Information Security and Applications*. 2017 Oct 23;37:28–37. Available from: <https://doi.org/10.1016/j.jisa.2017.10.002>
- [2] Meng W, Wong DS, Furnell S, Zhou J. Surveying the Development of Biometric User Authentication on Mobile Phones. *IEEE Communications Surveys & Tutorials*. 2014 Dec 31;17(3):1268–1293. Available from: <https://ieeexplore.ieee.org/abstract/document/7000543>
- [3] Bours P, Mondal S. Continuous authentication with keystroke dynamics. *Norwegian Information Security Laboratory NISlab*. 2015 Jan 1:41–58. Available from: [10.15579/gcsr.vol12.ch3](https://doi.org/10.15579/gcsr.vol12.ch3)
- [4] Mu N, Xu X, Zhang X. Finding autofocus region in low contrast surveillance images using CNN-based saliency algorithm. *Pattern Recognition Letters*. 2019 Apr 16;125:124–132. Available from: <https://doi.org/10.1016/j.patrec.2019.04.011>
- [5] Finnegan OL, White JW, Armstrong B, Adams EL, Burkart S, Beets MW, et al. The utility of behavioral biometrics in user authentication and demographic characteristic detection: A scoping review. *Systematic Reviews*. 2024 Feb 8;13(1):61. Available from: <https://doi.org/10.1186/s13643-024-02451-1>
- [6] Keselj A, Milicevic M, Zubrinic K, Car Z. The application of deep learning for the evaluation of user interfaces. *Sensors*. 2022 Nov 30;22(23):9336. Available from: <https://doi.org/10.3390/s22239336>
- [7] Kim DH, Song BC. Virtual sample-based deep metric learning using discriminant analysis. *Pattern Recognition*. 2020 Sep 10;110:107643. Available from: <https://doi.org/10.1016/j.patcog.2020.107643>
- [8] Teh PS, Teoh ABJ, Yue S. A survey of Keystroke Dynamics Biometrics. *The Scientific World JOURNAL*. 2013 Jan 1;2013(1):408280. Available from: <https://doi.org/10.1155/2013/408280>
- [9] Dillon R, De Marsico M. Behavioral biometrics for remote exam integrity: Continuous authenticity assessment via keystroke dynamics. *Procedia Computer Science*. 2025 Jan 1;274:402–411. Available from: <https://doi.org/10.1016/j.procs.2025.12.040>
- [10] Wang P, Chen B, Xiang T, Wang Z. Lattice-based public key searchable encryption with fine-grained access control for edge computing. *Future Generation Computer Systems*. 2021 Sep 21;127:373–383. Available from: <https://doi.org/10.1016/j.future.2021.09.012>
- [11] Charan PVS, Ratnakaram G, Chunduri H, Anand PM, Shukla SK, DKaaS: DARK-KERNEL as a service for active cyber threat intelligence. *Computers & Security*. 2023 Jun 10;132:103329. Available from: <https://doi.org/10.1016/j.cose.2023.103329>
- [12] Dintakurthy Y, Innmuri RK, Vanteru A, Thotakuri A. Emerging applications of artificial intelligence in Edge computing: A comprehensive review. *Journal of Modern Technology*. 2024:175–185. Available from: <https://doi.org/10.71426/jmt.v1.i2.pp175-185>
- [13] Oduri S. Continuous authentication and behavioral biometrics: Enhancing cybersecurity in the digital era. *International Journal of Innovative Research in Science Engineering and Technology*. 2024 Apr;13(7):13632–40. Available from: [10.15680/IJRSET.2024.1307140](https://doi.org/10.15680/IJRSET.2024.1307140)
- [14] Zhou Y, Liang X, Zhang W, Zhang L, Song X. VAE-based Deep SVDD for anomaly detection. *Neurocomputing*. 2021 Apr 29;453:131–40. Available from: <https://doi.org/10.1016/j.neucom.2021.04.089>
- [15] Soma AK. Hybrid RNN-GRU-LSTM model for accurate detection of DDOS attacks on IDS dataset. *Journal of Modern Technology*. 2024 May 14;2(1):283–291. Available from: <https://doi.org/10.71426/jmt.v2.i1.pp283-291>
- [16] Bansal P, Ouda A. Continuous Authentication in the Digital Age: An analysis of reinforcement learning and Behavioral biometrics. *Computers*. 2024 Apr 18;13(4):103. Available from: <https://doi.org/10.3390/computers13040103>
- [17] Penaganti R. Technical Implementation Guide: Modern Payment Solutions for Captive Finance. *Journal of Computer Science and Technology Studies*. 2025 Nov 6;7(11):346–362. Available from: <https://doi.org/10.32996/jcsts.2025.7.11.34>
- [18] Jain AK, Kumar A. Biometric Recognition: An Overview. In *The International library of ethics, law and technology*. 2012. p. 49–79. Available from: https://doi.org/10.1007/978-94-007-3892-8_3
- [19] Salem A, Obaidat MS. A novel security scheme for behavioral authentication systems based on keystroke dynamics. *Security and Privacy*. 2019 Feb 13;2(2). Available from: <https://doi.org/10.1002/spy2.64>
- [20] Baig AF, Eskeland S, Yang B. Privacy-preserving continuous authentication using behavioral biometrics. *International Journal of Information Security*. 2023 Jul 13;22(6):1833–1847. Available from: <https://doi.org/10.1007/s10207-023-00721-y>
- [21] Kang G, Park J, Kim YG. Continuous behavioral biometric authentication for secure metaverse workspaces in digital environments. *Systems*. 2025 Jul 15;13(7):588. Available from: <https://doi.org/10.3390/systems13070588>
- [22] Belman AK, Wang L, Iyengar SS, Sniatala P, Wright R, Dora R, et al. Insights from BB-MAS – A Large Dataset for Typing, Gait and Swipes of the Same Person on Desktop, Tablet and Phone. *arXiv (Cornell University)*. 2019 Nov 8; Available from: <https://doi.org/10.48550/arXiv.1912.02736>
- [23] Dillon R, Arushi. Agent-based modeling of free-text keyboard dynamics for continuous authentication. *arXiv preprint*. 2025. Available from: <https://arxiv.org/abs/2505.05015>
- [24] Islam MM, Rafiq MA, Islam MA. A Privacy-Preserving Behavioral Authentication System. In *International Symposium on Foundations and Practice of Security*. 2024 Dec 9 (pp. 95–107). Cham: Springer Nature Switzerland. Available from: https://doi.org/10.1007/978-3-031-87496-3_7
- [25] Subash A, Song I, Lee I, Lee K. Adaptability of current keystroke and mouse behavioral biometric systems: A survey. *Computers & Security*. 2025 Oct 21:104731. Available from: <https://doi.org/10.1016/j.cose.2025.104731>
- [26] Saripudi K. A study on Artificial Intelligence and Cloud Computing Assistance for Enhancement of Startup Businesses. *Journal of Computing and Data Technology*. 2025 Jul 26;1(1):68–76. Available from: <https://doi.org/10.71426/jcdt.v1.i1.pp68-76>
- [27] Rayani PK, Changder S. Continuous user authentication on smartphone via behavioral biometrics: A survey. *Multimedia Tools and Applications*. 2022 Jun 9;82(2):1633–1667. Available from: <https://doi.org/10.1007/s11042-022-13245-9>
- [28] Penaganti R. AI-Driven fraud Detection in Financial Systems: A technical Deep dive. *Journal of Information Systems Engineering & Management*. 2025 Sep 30;10(60s):1049–1069. Available from: <https://doi.org/10.52783/jisem.v10i60s.13262>
- [29] Progonov D, Cherniakova V, Kolesnichenko P, Oliynyk A. Behavior-based user authentication on mobile devices in various usage contexts. *EURASIP Journal on Information Security*. 2022 Sep 16;2022(1):6. Available from: <https://doi.org/10.1186/s13635-022-00132-x>