

Unsupervised Detection of Credential Stuffing and Account Takeover Attempts through User Behavioral Biometrics in Web Applications

Muhammad Saad Haleem¹, Rabia Tehseen¹, Kashif Nasr¹, Uzma Omer², Anam Mustaqeem¹, Rubab Javaid¹

¹University of Central Punjab, Lahore, Pakistan

²University of Education, Lahore, Pakistan

Correspondence: rabia.tehseen@ucp.edu.pk

Citation | Haleem. M. S, Tehseen. R, Nasr. K, Omer. U, Mustaqeem. A, Javaid. R, “Unsupervised Detection of Credential Stuffing and Account Takeover Attempts through User Behavioral Biometrics in Web Applications”, IJIST, Vol. 7 Issue. 4 pp 2718-2729, November 2025

DOI: <https://doi.org/10.33411/IJIST/1634>

Received | October 06, 2025 **Revised |** October 21, 2025 **Accepted |** October 27, 2025

Published | November 09, 2025.

Credential stuffing and Account TakeOver (ATO) attacks can still be stated as being one of the most ongoing risks of contemporary web applications, as they use reused credentials and defy traditional authentication procedures. Conventional supervised detection approaches rely on labeled attack data (which is frequently non-existent or incomplete) in the real world. In this paper, a credential stuffing and ATO adversarial behavior multi-layer detector is suggested through modeling user behavioral biometrics based on the web-session navigation patterns and chaining information during the login procedure. The framework combines Isolation Forest and Local Outlier Factor to train a normal behavioral distribution and detect deviations that suggest abnormal use of the accounts, which are either automated or reports of hacked accounts. Analysis of three different datasets RBA, MSNBC, and CERT Insider Threat, has shown that the framework is highly detection sensitive with an ROC-AUC of up to 0.936 and an F1-score of 0.842 in login-level anomaly detection, and cross-domain generalization on enterprise data. The findings validate the practicability of unsupervised behavioral modelling as a lightweight defense mechanism that is scalable and resistant to credential abuse. The suggested solution ensures that labelling is reduced, the privacy of users is maintained, and a scalable basic infrastructure is provided to accommodate the adaptive and risk-centric authentication models in massive web applications.

Keywords: Unsupervised Learning, Anomaly Detection, Behavioral Biometrics, Credential Stuffing, Account Takeover, Risk-Based Authentication, Web Security, User Behavior Analytics



Introduction:

With the migration of services to the web, online applications are an inseparable part of contemporary life and are therefore subject to a wide variety of attacks. Such attacks are the sneakiest one's credential-based attacks: when attackers gain access to user credentials that have been obtained legitimately via data breaches or phishing, they may later use them to gain entry into other systems. Other risks in this category are credential stuffing and account takeover (ATO); in the former, high rates of tested stolen credentials get cross-tested on multiple web applications, and in the latter, a successful log-in enables an attacker to take control of an account. The momentum of increasing the number of such attacks, automating them, and making them complex is something that can be termed as a grave problem of internet security[1].

The behavioral biometrics subfield has grown to become a powerful protection in the broader application of authentication and fraud detection. Behavioral biometrics, unlike the traditional biometrics which used non-dynamically tied physical characteristics (e.g., fingerprints, iris, face), analyzes dynamically, user-specific patterns, which may include keystroke dynamics, mouse and pointer movements, touch/swipe gestures, and sequences of navigation[2]. These behavioral indicators are difficult for attackers or automated bots to replicate using legitimate credentials. As a result, they can serve as a non-intrusive, continuous method for authenticating users and detecting anomalies during login or active sessions. Nonetheless, a lot of current studies related to credential stuffing or ATO detection are related to supervise or rule-based models, which are highly dependent on either labeled or pre-coded heuristics. These approaches may not be generalized to invisible attacks or changes in adversarial behavior. Moreover, the absence of labeled attack data and dynamic characteristics of contemporary threats indicate the necessity of unsupervised methods of detection, which learns normal user behavior and report deviations without any explicit attack labels. Combining behavioral biometrics with unsupervised anomaly detection thus provides an efficient and adaptive, data-saving tool to detect suspicious login events and potential account hijacking[3].

Although those studies have mostly used supervised or rule-based mechanisms, they are dealing with serious difficulties in adapting to types of attacks that are not seen and a lack of labels. This research directly fills that gap with a presentation of an unsupervised, multi-layer detection system that is used to model both login-level and session-level behavior and to detect anomalies without any labeled attack data. The proposed solution will help address the gap in the research between behavioral biometrics and scalable, label-efficient detection mechanisms against credential stuffing and account takeover (ATO) prevention.

The proposed study suggests a new unsupervised method of identifying any attempt at credential stuffing and account takeover in a web-based application using the behavioral biometrics of the user. The main objective of this study is to create and realize a multi-layer and unsupervised detection design that can detect credential stuffing and account takeover (ATO) attacks in web applications, using behavioral biometrics of users. Our work is unique, as it presents a new label-free method of identifying credential stuffing and account takeover attacks through the unsupervised behavioral biometrics- a field that previous studies have rarely ventured into. The suggested framework does not require labeled attack data as opposed to standard supervised or rule-based models, which rely heavily on the available training data, but it learns the normal user interaction patterns autonomously and detects anomalies that are symptomatic of credential abuse in real time. The key contributions of this work include a robust behavioral feature extraction pipeline that captures keystrokes, mouse movements, and time-series features unique to user interactions with web-based applications. The model to detect anomalies unsupervised, and be able to learn individual behavioral baselines and identify deviations that are indicative of compromise. An in-depth analysis of realistic web interaction

data, which proved to have better accuracy and fewer false positives than supervised and heuristic-based systems. This research represents the next step in advancing web application security by introducing an adaptive, privacy-conscious, and label-free system for detecting sophisticated credential-based attacks.

This paper presents a scalable unsupervised detection system that models natural login behavior to identify inconsistencies potentially indicative of credential misuse. The remaining part of this paper is organized as follows: Section 2 presents the literature review, Section 3 describes the methodology, Section 4 discusses the experimental results, and Section 5 provides the conclusion and future recommendations.

Related Work:

Behavioral biometrics has been extensively studied across web and mobile environments, serving as an additional layer of identity assurance. Previous research has already explored key aspects such as feature engineering, model selection, and cross-device variability. It is based on these premises that the given paper is dedicated to the possibilities of using unsupervised behavioral modeling to identify credential-stuffing and ATO attacks in web apps. The landscape of mouse dynamics, such as GUI-widget interaction features, datasets, and tools, gaps and opportunities of continuous authentication of desktops and web environments, is mapped using complementary studies[2]. Touch dynamics surveys. Mobile touch dynamic surveys are described to offer the speed, intensity, and direction geometries of the gestures as inoffensive authentication along with session chance points, particularly in ATO scenarios within web-app platforms[3]. A 2024 scoping review also claims that behavioral biometrics is currently being incorporated into authentication systems to mitigate fraud and account takeover (ATO) and assist in risk-based security control[4]. Regarding the threat perspective, Bark worth et al. report one of the earliest empirical experiments of detecting credential-stuffing bots with the help of human-computer interaction cues that include key strokes and mouse movement, thus proving the idea of interactional micro rhythms where automated attacks and legitimate credential logins can be differentiated[5]. Industry-wise, there would be a paper on account takeover prevention in large online marketplaces, which notes the signature of attack, detection telemetry, and operational lessons like balancing the user friction and false positives-emphasizing the need to employ behavior-sensitive solutions at scale[6]. Other e-commerce researches resonate with these drivers and stimulates behavioral indicators of timely ATO detection at the stage of login and post-login processes[7].

More importantly, label-efficient and unsupervised techniques are being placed on increasing importance to handle both the limited labels of attacks and dynamic adversaries. Unsupervised techniques of[8], One Class SVM, k-Means/K-NN and density/outlier scoring, are listed in a 2025 keystroke-dynamics survey of user-specific modeling in which there is no access to unlabeled data as the impostor. The author in[9] evaluates outlier-detection methods, including HBOS, for keystroke-based authentication and reports competitive performance without using any impostor samples, confirming the feasibility of purely unsupervised pipelines for behavioral signals. Previous studies on one-class models of keystroke biometrics specify kernel selection and thresholding on the remote case, which provides realistic information about per-user profiling in the open-set case[10]. Lastly, the recent research experimentation in keystroke modeling exhibits a higher distinctiveness, non-intrusiveness, and points out the generalization problems (cross-device, environmental noise) that unsupervised detectors have to work with in production web applications[11].

Collectively, these studies highlight a research gap at the intersection of credential-stuffing/ATO defense and unsupervised behavioral biometrics in web applications, emphasizing the shift from supervised or rule-based detectors toward user-adaptive, label-

efficient models that learn normal interaction dynamics and detect anomalies indicative of credential abuse.

Methodology:

This section describes the suggested architecture of unsupervised detection of credential stuffing and account takeover (ATO) attempts with the help of user behavioral biometrics in web applications.

Overview of the Proposed Framework:

The overall process involves collecting users' behavioral interaction data during login and throughout active web sessions. These unprocessed data, including the timing of keystrokes, mouse movement patterns, and navigation patterns in a session. These sessions are processed into high-level behavioral characteristics. Unsupervised Learning Layer then trains regular user behavioral patterns and detects an anomaly, which could be a credential-stuffing or ATO attempt. Finally, the Application Layer integrates the detection outputs into the web application security system, providing real-time alerts, a user risk score, and adaptive access control.

Framework Architecture:

The proposed architecture features a three-layer hierarchical structure, as illustrated in Figure 1. The bottom layer is responsible for acquiring and preprocessing raw behavioral signals, while the middle layer hosts the unsupervised anomaly detection engine, which models normal user behavior and identifies deviations. The top layer translates detection outcomes into actionable responses, enhancing web application defenses against credential-stuffing and account takeover (ATO) attacks.




Figure 1. Proposed framework

Suggested multi-layer model with data collection, unsupervised learning, and application layer presents the flow between the behavioral data capture and anomaly scoring, and adaptive risk-based response.

Data Collection Layer:

In the authentication and post-login phases, the first layer collects user behavioral biometrics. The primary data sources comprise keystroke dynamics (hold and flight times), mouse movements (speed, click rate, angle variations), and navigation behaviors (dwell time on pages and frequency of transitions). Multiple steps of preprocessing have been conducted, including removal of outliers, normalization of timestamps, scaling of features, as well as

aggregation at the session level. Subsequently, the User Behavioral Vector (UBV) is formed by calculating statistical features (mean, variance, skewness) as well as temporal features (inter-key delay) and spatial features (cursor path length).

Unsupervised Learning Layer:

This layer performs anomaly detection without requiring labeled attack data, estimating each user's low-level behavior using unsupervised methods. Approaches based on clustering, including k-Means, DBSCAN, have been applied for grouping behavioral vectors and marking outliers that are far away from the cluster centers. Detection methods based on density have been applied, including Isolation Forest, Local Outlier Factor (LOF) for estimation of the deviations in the density of data points. A model based on reconstruction has been performed using Auto encoders to measure reconstruction error between prediction and actual behavior patterns to check anomalies.

The Normal Behavior Model (NBM) is trained on each user's historical sessions. Incoming sessions are then compared against this model to produce an Anomaly Score, representing the likelihood of a credential-stuffing or account takeover attempt.

Application Layer:

The final layer leverages the detection results to enhance web application security. A Risk Assessment Module aggregates all anomaly scores and assigns a corresponding risk level: Low, Medium, or High. The Adaptive Response Engine uses the assessed risk to trigger appropriate security measures, such as CAPTCHA, multi-factor authentication (MFA), session termination, or alerts to the administrator. Meanwhile, the Visualization Dashboard displays behavioral metrics, time-based anomalies, and user risk profiles for security analysts.

Figure 2 shows the inner workflow of the Normal Behavior Model (NBM). It starts with the collection of raw activity data and then, features are extracted at the login-level and session-level. Vectors obtained are normalized and sent to the clustering module, which discriminates normal and abnormal behavioral patterns. The feedback loop enables performance to be adjusted to fit the model by learning continuously, which lowers the number of false positives over time.

Datasets and Evaluation:

Dataset Overview:

The RBA data contains events of user logins in the data with timestamps, device identifiers, and geolocation IP-related attributes. Every event is a discrete authentication attempt, and the actions analyzed can be frequency, switching devices, and location variation in terms of their logs. These functions allow building a user-centered behavioral baseline for detecting anomalies.

Three datasets are chosen to assess the effectiveness of the proposed framework in detecting credential-stuffing and account takeover (ATO) attempts through unsupervised behavioral biometrics. The primary dataset, the Login Data Set for Risk-Based Authentication (RBA), provides realistic login-based behavioral and network attributes. Two complementary datasets MSNBC Anonymous Web Data and CERT Insider Threat Dataset, extend behavioral context by modeling navigation and post-login session behavior. Together, these datasets support multi-layer validation of both login-level and session-level anomaly detection.

The MSNBC and CERT datasets were explicitly used as the basis of session-level features. In the case of the MSNBC data, user clickstream logs were split into sessions according to temporal discontinuities and each session transformed into a numerical vector describing the depth of navigation, dwell time, category transition frequency, and loop behavior patterns. In the case of the CERT dataset, user logs were also clustered into windows of 30 minutes of enterprise sessions, and session-level features of frequency of logins, length of time spent working on a particular workstation, number of file accesses, and abnormality

of activity period were calculated. These derived features are the basis of session-level anomaly detection in the proposed framework.

Sequential behavior during user activities was explicitly derived in session-level features that were based on the MS NBC and CERT datasets. In the case of MSNBC, page visits and dwell times follow an order, and these visits are the ones that indicate navigation patterns. In the case of CERT, the session summarizes the various user events (logins, file access, and email activity) within specified time frames. This session-level representation allows one to compare between normal and abnormal behavioral flows, which is in accordance with the demand of the reviewer to have explicit session-based analysis.

Dataset Specifications:

A complete specification of the datasets used in this study is presented in Table 1.

Model Evaluation:

The framework was evaluated using real and simulated login datasets under controlled credential-stuffing and ATO scenarios. Metrics such as precision, recall, F1-score, false positive rate, and anomaly detection latency were employed to assess system performance. Cross-user validation ensured generalization and robustness against diverse interaction behaviors. Figure 2 illustrates the methodological flow of this research.




Figure 2. Extracting feature layers that demonstrate the level of login and session.

Workflow diagram of the end-to-end approach on feature extraction and preprocessing to model training, anomaly detection, and model performance evaluation across datasets.

The RBA dataset was used to train the unsupervised models on legitimate login sessions, constructing per-user behavioral profiles through clustering, density-based, and reconstruction-based approaches. During the testing phase, labeled ATO and credential-stuffing attempts available in the RBA dataset were used solely for evaluation. The MSNBC dataset assessed generalization to navigation sequence anomalies, while the CERT dataset evaluated enterprise-level robustness. Evaluation metrics included Precision, Recall, F1-score, AUC, and False Positive Rate (FPR). Anomaly scores from each layer (login and session) were fused to obtain the final risk classification.

Experiment and Results:

Experimental Setup:

Experiments were conducted on a system equipped with an Intel® Core™ i7 (2.8 GHz) CPU, 16 GB RAM, and Python 3.11. The unsupervised models, Isolation Forest (IF) and Local Outlier Factor (LOF)—were implemented using scikit-learn with the following hyper parameters presented in Table 2.

Table 1. Dataset description

Dataset	Access Link
Login Data Set for Risk-Based Authentication (RBA)	https://github.com/das-group/rba-dataset
MSNBC Anonymous Web Data	https://archive.ics.uci.edu/dataset/133/msnbc%2Bcom%2Banonymous%2Bweb%2Bdata
CERT Insider Threat Dataset (r4.2)	https://www.kaggle.com/datasets/nitishabharathi/cert-insider-threat

Table 2. Comparative overview of key specifications across RSA, MSNBC, and CERT datasets.

Dataset Specifications					
RSA		MSNBC Anonymous Web Data		CERT Insider Threat Dataset	
Field	Type	Feature	Number	Field	Instance/Type
IP Address, Country, Region, City	String	User Sessions	989,818	User count	~1000
ASN	Integer	visits per user	~ 517	Events recorded	~32,770,222
User Agent	String	Categories	17	Malicious injected	7,323
OS Name/Version	String	Field	Data type	Field	Data type
Browser Name/Version, Device Type	String	User session ID	Implicit	timestamp;	Date Time
Login Timestamp	Date Time	Sequence of category visits	Categorical sequence	url	anonymized or category
Round-Trip Time (RTT)	Float (ms)	Session Length	Integer	user_id; device_id; workstation_id;	String
Login Successful, Is Attack IP, Is Account Takeover	Boolean	Category codes	Integer/ Categorical	success;	Boolean

Table 3. Comparison with Recent Related Work

Ref.	Modality / Signals	Task & Setting	Learning Paradigm	Dataset (scope)	Reported Performance	Relevance to ATO/Stuffing
Our Work (2025)	Login metadata (IP/UA/RTT) + session navigation	Credential stuffing & ATO detection	Unsupervised anomaly detection (IF, LOF) + fusion	RBA (33M), MSNBC (990k), CERT (32M+)	RBA: ROC-AUC 0.936, F1 0.842; MSNBC: F1 0.834; CERT (fused): ROC-AUC 0.883, F1 0.806	Direct
[3]	Keystroke + mouse biometrics	Bot-based credential stuffing detection	Supervised classification	IMAP login traces	~90% accuracy (HCI features)	High
[5]	Keystroke dynamics	Behavioral authentication	Unsupervised (HBOS, kNN)	Public typing datasets	ROC-AUC ~0.90	Medium
[7]	Keystroke dynamics	User identification/authentication	Supervised deep CNN/RNN	Open keystroke datasets	Accuracy ~96%	Low–Medium
[10]	Account + graph behavior	ATO detection in e-commerce	Supervised graph embedding	Proprietary commerce logs	ROC-AUC 0.91	High

Table 4. Hyper parameters trained

Parameter	Isolation Forest	Local Outlier Factor
Estimators	300	–
Contamination	0.02	0.02
Max-samples	Auto	–
Neighbors	–	20
Metric	–	Euclidean

Isolation Forest performed better than LOF in all measures, indicating that these two algorithms can be used to generalize on a variety of login characteristics, including deviation by IP subnet, changes in latency (RTT), and abnormal device fingerprints. The false-positive rate is low (less than 7), which is good evidence of being highly suitable in real-time integration with risk-based authentication (RBA) systems.

Results on RBA Dataset (Login-Level Detection):

This evidence proves that joint analysis of IP address, latency of logins, device fingerprints, and navigation paths on a session level presents an additional view on suspicious user behavior. The system had a false-positive rate of less than 7% and showed that this was suitably applicable to be integrated within the risk-based authentication processes without having detrimental impacts on the user experience. This is further shown in the universality of the fusion model (RBA-CERT) which implies that the latent behavioral representation is not dependent on the variations in domain and therefore can be effectively implemented in a broad variety of web settings. Results of RBA dataset have been presented in Table 3.

Table 5. Results of RBA dataset

Model	ROC-AUC	PR-AUC	F1-score	False Positive Rate
Isolation Forest	0.936	0.874	0.842	0.061
Local Outlier Factor	0.912	0.841	0.802	0.078

Results on MSNBC Dataset (Session-Level Navigation Anomalies):

LOF slightly outperforms IF, likely because of its sensitivity to local density deviations in user navigation paths. It efficiently detects anomalous sequences, such as unusually short or looping sessions, which are typical of automated credential-stuffing bots. Results of MSNBC dataset have been presented in Table 4.

Table 6. Results of MSNBC dataset

Model	ROC-AUC	PR-AUC	F1-score	False Positive Rate
Isolation Forest	0.892	0.853	0.826	0.075
Local Outlier Factor	0.901	0.861	0.834	0.072

Cross-Domain Validation on CERT Dataset:

Applied to enterprise login and browsing behavior in the CERT dataset, the combined model fusing login and navigation anomaly scores achieved the best performance, demonstrating the cross-domain adaptability of the unsupervised pipeline. The model was able to detect insider-like account misuse patterns, demonstrating robustness beyond web-based environments. Results of CERT dataset have been presented in Table 5.

Table 7. Results of CERT dataset

Model	ROC-AUC	F1-score	Transfer Generalization
Isolation Forest (trained on RBA)	0.841	0.768	Moderate
Fused Model (Login + Navigation)	0.883	0.806	High

The proposed multi-layer unsupervised architecture achieved a maximum ROC-AUC of 0.936 on the RBA dataset, indicating strong discriminative power between benign and ATO sessions. Fusion of login and session-level scores improved overall F1 by approximately 4–5%, validating the complementary nature of behavioral features. Model generalization across different domains (web → enterprise) remained strong, confirming that unsupervised

behavioral representations effectively capture anomalous intent rather than context-specific patterns. The overall framework demonstrates potential for real-time risk scoring and adaptive authentication, even without labeled data.

The internal workflow of the Normal Behavior Model (NBM) and anomaly scoring mechanism is described in Layer 2 in Figure 3. The NBM is initially trained through historical behavioral sessions and establishes a baseline of patterns of normal user behavior. New sessions are then subjected to this model to calculate an anomaly score of deviations from learned distributions. The scores are thresholded and sent to the Risk Assessment Module to be interpreted and acted upon.




Figure 3. Internal work of the Normal Behavior Model (NBM).

Discussion:

Interpretation of Results:

The suggested structure proves useful in showing how unsupervised behavioral modelling can identify credential-stuffing and account takeover (ATO) attacks without training data that has labels. The findings of Section 6 (experiment and results) demonstrate that Isolation Forest had good discriminative ability at the login level (ROC-AUC = 0.936, F1 = 0.842). In contrast, Local Outlier Factor proved slightly more effective at detecting abnormal navigation patterns (F1 = 0.834), owing to its sensitivity to local density variations. Using the cross-domain generalization of both layers resulted in a significant improvement of the results with an ROC-AUC of 0.883 and an F1-score of 0.806 on the CERT dataset.

The variation between the performance of LOF and the Isolation Forest in the various datasets can be attributed to the working that the detection of the two algorithms relies on. LOF is effective on the MSNBC data because it is able to detect local density variation, as this is particularly effective in detecting deviations in the navigation and looped clicks, which are more specific to automated bot traffic. Isolation Forest is, however, more applicable to the RBA data, and the tree-based mechanism used seals the world outliers based on the heterogeneous features of the level of login, such as IP deviations, the response latency, and the variance of the device fingerprints. This complementary behavior is important to stress the importance of both approaches in the context of the layers.

The variation between the performance of LOF and Isolation Forest in different datasets can be explained by the differences in their detection mechanisms. LOF is best on the MSNBC data set since it captures the local density anomalies, which are especially useful in identifying anomalies in navigation and looped click patterns that characterize automated bot sessions. By contrast, Isolation Forest works better with RBA data, with the tree-based nature isolating global outliers according to the heterogeneous features of the login, including IP deviation, response latency, and device fingerprint disparity. This complementary characteristic lends significant weight to the suggestion of utilizing the two approaches in the layered framework.

In contrast, Isolation Forest performs better on the RBA dataset, as its tree-based structure isolates global outliers based on heterogeneous login features such as IP deviations, response latency, and device fingerprint variation. This complementary behavior underscores the importance of employing both methods within the layered framework.

Comparison With State-of-The-Art:

Table 6 brings this research into perspective with four recent studies that have been conducted on the topic of behavioral biometrics, credential-stuffing prevention, and account takeover detection. Unlike most prior research that depends on supervised learning with labeled attack data, the proposed system uniquely integrates unsupervised anomaly detection into a multi-layer behavioral framework, modeling both micro-level login behavior and macro-level session dynamics.

It can be clearly observed that, in contrast to purpose-specific biometric systems (e.g., keystroke-only systems), the proposed system uses universally-provided session metadata, which allows it to be practically used in actual web platforms. A majority of similar experiments rely on attack-labeled datasets to perform supervised learning, but this approach trains normal user profiles unsupervised and learns to address unknown types of attacks. The findings on the CERT dataset prove transferability to large-scale or in-house systems, not just to a web login situation. Performance is equally high with or without supervision (ROC-AUC > 0.93, both in RBA and cross-domain of about 0.88).

Limitations:

In general, this paper demonstrates that unsupervised behavioral biometrics, when appropriately overlaid on both login and session actions, can be used to identify ATO and credential-stuffing attacks with high accuracy and with minimal labeled data. The relative analysis also serves to support the idea that this solution bridges an essential gap between accuracy and scalability, and that solid, privacy-conscious, real-time defense systems in contemporary web applications become a viable option. Although the proposed framework has a high level of performance, several limitations should be discussed:

The existing datasets lack such continuous behavioral biometrics as mouse movement paths, dwell, or typing behavior. These may increase the temporal granularity, as well as detection latency.

The RBA data available to do the modeling at the login level is privacy-aware and somewhat artificialized, which can limit its representation of the diversity in the real world (e.g., the use of shared IPs, VPNs).

Credential-stuffing bots are quickly adapting, with the potential of imitating legitimate patterns of navigation and devices. Adversarial retraining or continual learning should be introduced in future work to ensure the detection accuracy.

Although model fusion enhances detection, it increases the complexity of the computation. Live deployment would take advantage of the incremental scorers or incremental anomaly findings, like incremental Isolation Forest.

The assessment was based on web sessions. The integration of mobile user behavior (touch dynamics, gesture flow) can be applied to the area of usability of hybrid authentication systems.

Conclusion and Future Work:

The paper introduces an unsupervised credential stuffing and account takeover (ATO) attack detection model with user behavioral biometrics on web services. The proposed system models typical user behavior, without using labeled attack data, by using multi-layer data, including the metadata of the logged-in user, sequences of navigation on the session, and behavior of page transitions. The combination of Isolation Forest and Local Outlier Factor algorithms as part of a layered architecture allows one to identify the presence of some subtle

behavioral anomalies that can be missed by traditional rule-based systems or supervised classifiers.

The outcomes of the experiment based on three different data sets, RBA, MSNBC, and CERT Insider Threat, prove the system's strength and flexibility. The framework had ROC-AUC scores of over 0.93 at the login-level anomalies and F1-scores of more than 0.80 at the extend-to-enterprise-scales. The results highlight the usefulness of unsupervised learning in the detection of anomalies on a large scale when the availability of labeled data is limited or impossible. In addition, the low false-positive rate of the system shows that it is feasible to deploy it in risk-based authentication (RBA) pipelines, adaptive access control systems, and risk-based identity verification systems.

This study provides a label-efficient, domain-neutral, and privacy-preserving behavioral biometric algorithm with the capability to adapt to a variety of platforms without a large-scale retraining of existing supervised and modality-specific behavioral biometric methods. It connects the user behavior analytics with the cyber threat detections, stating that human behavioral patterns could be considered as credible indicators that reveal automated or hijacked sessions.

Future work will be on real-time behavioral indicators, including mouse patterns and time of keystroke, model explain ability, and continuous learning techniques to adapt to changing attack patterns. The suggested framework is therefore a scalable and sustainable path to intelligent and self-learning cybersecurity systems that enhance the defense posture of contemporary web applications.

References:

- [1] a. Tariq, "tracking for good: finding behavioral biometrics on the web using static taint analysis," *univ. Waterloo*, 2025, [online]. Available: <https://uwspace.uwaterloo.ca/items/a27d11b8-354e-487b-88ff-69e269eabb6d>
- [2] m. M. Kamol, m. S. Siddiky, f. Anwar, a. M. Khan, and a. Salam, "credentials stuffing attack prevention using machine learning," *2024 27th int. Conf. Comput. Inf. Technol. Iccit 2024 - proc.*, pp. 2899–2904, 2024, doi: 10.1109/iccit64611.2024.11022023.
- [3] a. Bara, "finding behavioural biometrics scripts on the web using dynamic taint analysis," *univ. Waterloo*, 2025, [online]. Available: <https://uwspace.uwaterloo.ca/items/c77e83a2-c624-40f4-85f5-51e3f539ab07>
- [4] a. Riyaz fathima and a. Saravanan, "an approach to cloud user access control using behavioral biometric-based authentication and continuous monitoring," *int. J. Adv. Technol. Eng. Explor.*, vol. 11, no. 119, pp. 1469–1496, 2024, doi: 10.19101/ijatee.2024.111100516.
- [5] a. Barkworth, r. Tabassum, and a. Habibi lashkari, "detecting imap credential stuffing bots using behavioural biometrics," *acm int. Conf. Proceeding ser.*, pp. 7–15, dec. 2022, doi:10.1145/3586102.3586104;journal:journal:acmotherconferences;ctype:string:book.
- [6] v. Arora and r. Kanji, "modelling and predicting user behaviour".
- [7] f. A. Idialu, "leveraging zero trust architectures and blockchain protocols to prevent credential stuffing and lateral fraud attacks in enterprise systems," *int. J. Comput. Appl. Technol. Res.*, vol. 14, no. 8, 2025, doi: 10.7753/ijcatr1408.1002.
- [8] p. Teimoory, "towards robust security in smart payment systems: challenges and solutions," *smart cities reg. Dev. J.*, vol. 9, no. 3, pp. 29–38, 2025, doi: <https://scrd.eu/index.php/scrd/article/view/627>.
- [9] s. Sadeghpour, "machine learning-based defences against advanced 'session-replay' web bots," *york*, vol. 3, 2024, [online]. Available: <https://yorkspace.library.yorku.ca/items/6ac36ce5-a90c-49d7-85b4-e470695aca53>
- [10] m. Hämmäläinen, "analysis of artificial intelligence in cybersecurity identity and access management : potential for disruptive innovation," *lut univ.*, 2024, [online]. Available:

<https://lutpub.lut.fi/handle/10024/168740>

- [11] s. Kumar, “biometric systems security and privacy issues,” *leveraging comput. Vis. To biometric appl.*, pp. 68–91, oct. 2024, doi: 10.1201/9781032614663-4/Biometric-Systems-Security-Privacy-Issues-Sunil-Kumar.



Copyright © by the authors and 50Sea. This work is licensed under the Creative Commons Attribution 4.0 International License.