International Journal on Advanced Computer Engineering and Communication Technology

ISSN: 2278-5140

Volume 15 Issue 2 (2026) | Pages: 264-273

BehaviorGuard-AI: A Context Aware UEBA Framework for Insider Threat Detection Using Unsupervised Behavioral Analytics

Vinjamuri Bhuvanesh[1], Gennepally Sreehitha[2], Mohammed Abrar Hamed Khan[3], Namita Parati[4]

[1,2,3,4]Department of Computer Science and Engineering (Data Science),
MVSR Engineering College, Hyderabad, Telangana, India

Peer Review Information

Type: *Article*
Received: *12 June 2026*
Revised: *15 July 2026*
Accepted: *20 Aug 2026*
Published: *23 Sept 2026*

Abstract

Insider threats are especially difficult to detect because they come from legitimate users with authorized access, and traditional defences which are mainly tuned to detect external attacks, often miss the subtle behavioral deviations exhibited by insiders. This paper presents BehaviorGuard-AI, a context aware UEBA framework that detects anomalous user behavior using behavioral analytics and unsupervised learning. The system builds meaningful baselines through attributes like contextual user profiling, engineered behavioral features, shift based segmentation, and role aware peer grouping. A hybrid detection pipeline combines HDBSCAN clustering with Isolation Forest validation to surface rare or unusual patterns in large scale enterprise authentication logs.

BehaviorGuard-AI provides investigative depth beyond statistical anomaly scores with an Email Compliance Auditing Pipeline. This pipeline detects high risk users by anomalies and correlates their email activity to anomalous windows in time and applies heuristic threat filters and Retrieval Augmented Generation (RAG) workflow to evaluate potentially sensitive messages against organizational policy. The system narrows down the workload and cuts down on the computational cost by confining the email analysis to the windows flagged as anomalous, while providing policy relevant context for analysts.

The framework was evaluated on a synthetic enterprise data set containing more than 3.3 million user activity records per hour and ten engineered behavioral features. The system identified 2,764 behavioral clusters and labeled approximately 6.34% of observations as anomalous without any labeled threat data. Results suggest that combining behavioral anomaly detection with targeted and policy aware email auditing provides analysts with a combination of clear anomaly signals and actionable context that improves the effectiveness of insider threat investigation.

Keywords: User and Entity Behavior Analytics; Insider Threat Detection; Behavioral Analytics; Unsupervised Learning; Isolation Forest; HDBSCAN; Anomaly Detection; Email Compliance Auditing; Cyber Security

How to Cite This Article

Bhuvanesh, V., Sreehitha, G., Khan, M. A. H., & Parati, N. (2026). BehaviorGuard-AI: A context aware UEBA framework for insider threat detection using unsupervised behavioral analytics. *International Journal on Advanced Computer Engineering and Communication Technology,* 15(2), 264–273.

Introduction

One of the biggest security concerns for modern organizations is insider threats. They are difficult to detect because they come from people who already have access to company systems. These threats can come from malicious employees, compromised accounts, or basic mistakes. They can result in loss of money, disruption of operations, and damage to the reputation of an organization. Spotting bad behavior is both important and hard. Organizations create huge amounts of data about user activity every day.

Firewalls, IDS and rule based monitoring work well against many external attacks but have difficulty with the subtle deviations that indicate insider misuse. Fixed rules and predefined thresholds do not take into account the differences between user roles, work schedules and normal behavior that generate many false alarms. This has caused researchers to begin using smarter analytical techniques which grasp user activity in its proper context.

User and Entity Behavior Analytics (UEBA) relies on data analytics and machine learning to understand what users normally do and then looks for actions that deviate from that behavior. It looks for unusual events by checking login patterns, device usage, access sequences and other user actions. Even with these methods, it is still hard to distinguish between a real insider threat and normal changes in day-to-day work, especially in large organizations where the employees have different roles and work in different ways.

We developed BehaviorGuard-AI, a context-aware UEBA framework, to tackle these challenges by employing behavioral analytics and unsupervised learning to detect insider threats. The framework establishes user baselines by leveraging a combination of contextual user profiles, engineered behavioral features, shift-based segmentation and role-based peer grouping. Then it uses HDBSCAN to find natural groups of behavior and Isolation Forest to look for statistically unusual activities. This combination allows to detect rare behavior patterns without the need of labeled attack data .

The email auditing layer does not analyze all of the emails sent in an organization, but only those messages that are tied to time periods where unusual behavior was detected. It first picks high risk users and associates their emails with the related behavioral windows and runs some simple heuristic checks before sending the selected emails to a Retrieval Augmented Generation (RAG) workflow. This targeting reduces extraneous processing and gives analysts policy-related information to help them understand why a user's activity might warrant further investigation.

We tested BehaviorGuard-AI on a large synthetic enterprise dataset to see how well it could detect anomalous user activity. The framework performed context-based user grouping before anomaly detection, which helped to interpret the results over different roles and work shifts. It also linked behavioral anomalies to targeted email compliance checks, providing analysts with more policy-related context rather than just an anomaly score. The results show that the framework can aid in insider threat investigations by limiting the analysis to users that require further attention.

Contributions

The main contributions of this paper are:

- A context aware User and Entity Behavior Analytics (UEBA) framework that combines contextual user profiling, behavioral feature engineering, shift based segmentation and role aware peer grouping to build meaningful behavioral baselines for insider threat detection.
- A hybrid anomaly detection approach, combining HDBSCAN clustering and Isolation Forest validation, to identify anomalous user behavior without the need for labeled threat data.
- An Email Compliance Auditing Pipeline that extends anomaly detection to the investigation phase, via anomaly driven user selection, temporal email correlation, heuristic threat filtering, and targeted communication analysis.
- A Retrieval Augmented Generation (RAG) based policy evaluation framework that supports explanations and policy references, for analyzing suspicious email communications against organizational security and privacy policies to generate compliance verdicts.
- A unified and explainable insider threat detection architecture that combines behavioral analytics with policy aware email auditing, allowing security analysts to perform more effective and context rich security investigations.

Related Works

Insider threat detection is a hot topic in security research since it is getting more and more difficult to detect malicious activities of legitimate users. Traditional security solutions are mainly focused on external attacks and signature-based detection mechanisms that are not very effective against insider threats that usually present subtle behavioral deviations. To address this problem, researchers have increasingly resorted to User and Entity Behavior Analytics (UEBA) frameworks that employ machine learning and behavioral modeling techniques to develop user baselines and detect anomalous activities.

Some studies have examined anomaly detection techniques for insider threat detection using authentication logs, access logs, and user activity data. Unsupervised learning approaches have grown in popularity as real world insider threat datasets have rarely enough labeled attacks [3]. Among these methods, Isolation Forest is popular because of its scalability, computational efficiency and ability to identify rare observations in high dimensional datasets [5]. However, the Isolation Forest can be challenged when the behavioral distributions involve multiple overlapping user populations or heterogeneous activity patterns.

Recent work has also looked at density based clustering techniques such as HDBSCAN for detecting behavioral anomalies [6]. HDBSCAN, unlike tree based methods that attempt to isolate individual observations, identifies dense behavioral regions and labels observations of low density as potential anomalies. This capability makes it amenable to the discovery of rare behavioral patterns that may not be captured using traditional anomaly detection techniques. However, clustering based methods alone might not perform well in complex and dynamic enterprise environments.

To overcome these limitations, recent UEBA systems integrate contextual information, behavioral features and multiple anomaly detection methods . Based on this idea, we propose the BehaviorGuard-AI framework which leverages contextual user profiling, role aware peer grouping, HDBSCAN clustering, and Isolation Forest to detect anomalous user behavior. Unlike many existing UEBA systems that stop at detecting anomalies, BehaviorGuard-AI also has an Email Compliance Auditing Pipeline. So, this pipeline takes high risk users and connects their emails to suspicious activity. Then it does some heuristic filtering and uses Retrieval Augmented Generation (RAG) to check emails against organization policies. This allows security analysts to better understand whether the detected anomalies are true insider threats.

Proposed System

BehaviorGuard-AI is a context aware User and Entity Behavior Analytics (UEBA) system which detects insider threats through behavioral analytics, unsupervised learning and policy aware communication auditing. Instead of applying fixed rules or pre-set thresholds, the framework establishes behavioral baselines for users and detects anomalies by spotting deviations from normal activity. It provides useful security insights through analyzing large scale authentication logs with a sequence of analytical steps.

The framework starts from the integration of user profiles and organizational context to build situation awareness. Authentication events are aggregated into hourly behavioral windows and a behavioral feature engineering process extracts indicators such as login deviations, device usage anomalies, burst activity, session gaps and temporal behavior patterns. To decrease the heterogeneity in behavior across the enterprise, user activities are grouped according to work shifts and organizational role groups, so that anomaly detection is performed on behaviorally similar peer populations.

The contextualized behavioral data is then analyzed with HDBSCAN, a density based clustering algorithm that finds natural behavioral groups and sparse observations as potential anomalies. Isolation forest is adopted as an additional anomaly detection mechanism to validate statistically unusual activities and improve detection robustness. The risk scores are generated by combining the outputs of the two models and the high risk users are identified for further investigation of their behavior.

BehaviorGuard-AI goes beyond statistical anomaly detection with an Email Compliance Auditing Pipeline that performs contextual investigation of high risk users. Rather than evaluate all communication within an organization, the pipeline evaluates emails relevant to anomalous behavioral windows. To reduce unnecessary analysis overhead, a multi stage filtering process is used to identify potentially sensitive communication. This process consists of anomaly driven user selection, temporal correlation and heuristic threat scoring.

Emails that pass filtering criteria are processed through a Retrieval Augmented Generation (RAG) architecture for policy aware compliance assessment. In the first stage, the system finds relevant parts of the organizational security and privacy policy based on the content of the email observed. In the second stage, a Large Language Model (LLM) analyzes together the selected policy sections and email communications to generate compliance verdicts, explanations and supporting policy references. This allows the framework to go beyond simple anomaly detection and provide explainable contextual evidence for insider threat investigations.

The framework produces anomaly indicators, behavioral risk scores, compliance verdicts, policy references, and actionable security alerts as the final output, which help the security analysts in their investigation and response activities. BehaviorGuard-AI merges behavioral anomaly detection and policy-aware email auditing into one workflow so security analysts don’t have to cross-reference both separately to connect unusual user activity with relevant policy evidence.

Figure 1 shows the overall workflow of the proposed BehaviorGuard-AI framework.

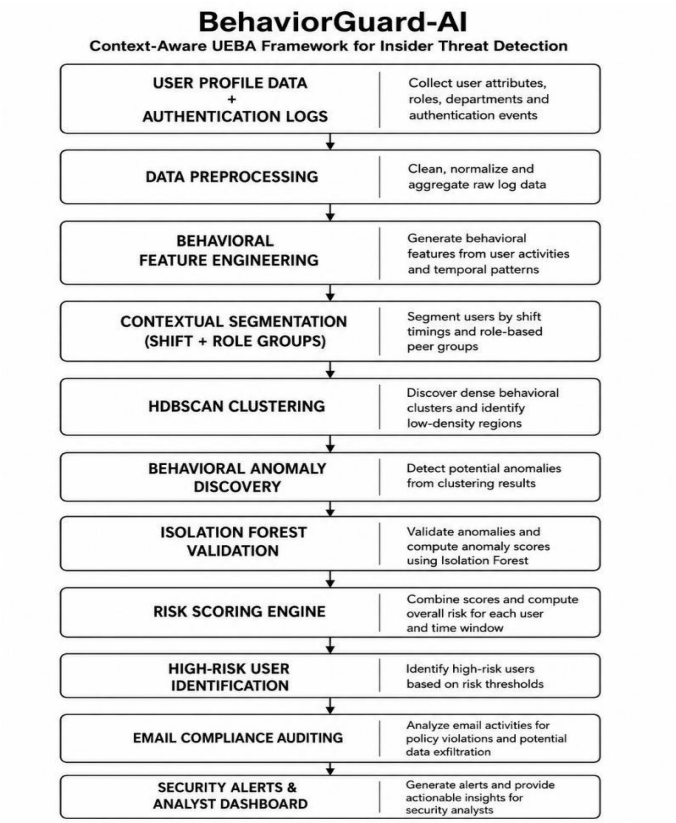

***Fig. 1**. Architecture of the Proposed BehaviorGuard-AI Framework*

The proposed framework is depicted in Figure 1. It combines contextual user profiling, behavioral feature engineering, contextual segmentation, anomaly detection, risk assessment and email compliance auditing in a single UEBA pipeline.

Methodology

Dataset Description

We conducted experimental evaluation of BehaviorGuard-AI on a large synthetic enterprise authentication dataset modeled on a CERT style insider threat scenario. The data set includes approximately 3,366,264 hourly user activity records from January 2010 to June 2011 and captures the behavior of almost 3,000 enterprise users. Each record captures authentication related events (user identity, timestamp, device attributes and logon/logoff actions) providing a rich timeline of access activity. To create analysis ready inputs, raw logs were aggregated into hourly windows of user activity, transforming discrete events into structured behavioral observations, amenable to anomaly detection and temporal analysis.

Besides the raw authentication events, the dataset was enriched by adding organizational context information such as user roles, departments, shift schedules and peer group memberships. This context enables BehaviorGuard-AI to establish more accurate baseline comparisons by comparing users to behaviorally similar peers rather than using global averages that can obscure local variation. The combination of high resolution temporal windows and role aware contextual features make the dataset well suited for evaluating how the framework addresses realistic enterprise challenges, such as large scale, diverse user behaviors, shift driven patterns, and limited labeled threats. To summarize, the dataset provides a realistic, large scale environment for rigorous evaluation of insider threat detection techniques, model robustness and the framework's ability to surface meaningful anomalies with explainable context for analysts.

Data Preprocessing

The raw authentication logs were subject to several preprocessing steps before behavioral analysis. First, logon and logoff records were ingested into a structured processing environment, and checked for completeness and consistency. Timestamps were normalized to a common datetime format to facilitate temporal analysis, and duplicate or malformed entries were scrubbed to improve data quality. Events were then ordered chronologically and associated to the correct users and devices.

Events were aggregated into fixed one-hour windows of activity, to allow for scalable behavioral modeling. We then calculated activity statistics per user-hour like login count, logout count and device usage. This transformed millions of low level events to compact, structured behavioral observations and reduced noise and computational load. These aggregated records were combined with additional metadata such

as user roles, departments, and peer group memberships, which provided each observation with the relevant contextual dimensions. The resulting preprocessed dataset was organized in terms of temporal and contextual dimensions and fed to the feature engineering and anomaly detection modules.

Behavioral Feature Engineering

Prior to running the anomaly detection models, we engineered 10 behavioral features from the authentication logs to better capture unusual user activity. Besides basic activity measures like login count, logout count, and device usage frequency, the proposed framework also generates a set of behavioral deviation features, intended to capture abnormal activity patterns based on historical user behavior.

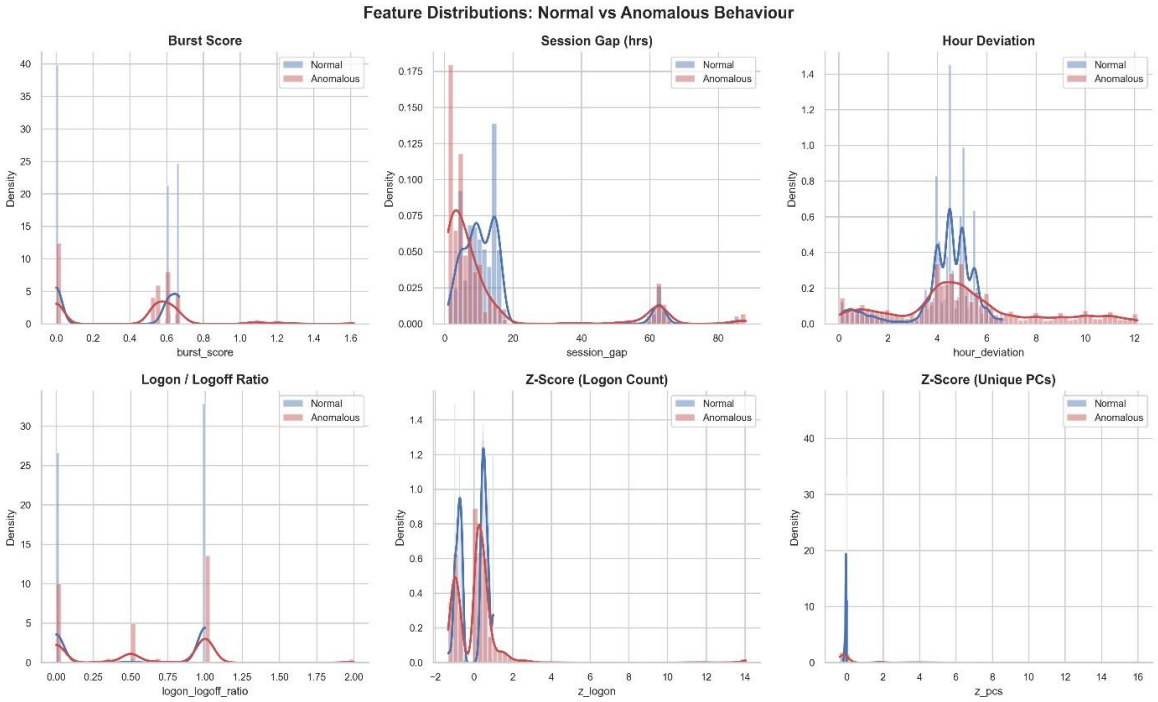

Fig. 2*. Distribution of Engineered Behavioral Features for Normal and Anomalous Users.*

The distribution of engineered behavioral features for normal and anomalous observations are shown in Figure 2. The engineered feature set consists of standardized login and device usage deviations (z_logon, z_pcs), logon deviation, device deviation, device ratio, burst score, hour deviation, session gap, logon-logoff ratio, and nite activity indicator. These indicators characterize temporal behavior, activity intensity, device utilization patterns, and deviations from established user baselines.

Of these features, the standardized deviation metrics are especially important because they are relative to the user's historical behavior rather than in absolute numbers. The approach enhances the robustness of anomaly detection by considering the differences in the activity profiles of the users. The output feature space is a comprehensive representation of user behavior and serves as the input for models used for contextual clustering and anomaly detection.

Contextual Segmentation

One of the main challenges in behavioral anomaly detection is the variability among users. When a single anomaly detection model is deployed across an entire organization, the anomaly rate can be inflated by comparing users with fundamentally different responsibilities, schedules, and activity patterns. To tackle this challenge, the proposed framework adopts a contextual segmentation before the anomaly detection.

The first level of segmentation is work shift based. The user activities are divided into Day (09:00-16:00), Evening (17:00-21:00) and Nite (22:00-08:00). This separation of users ensures that they are evaluated in similar time contexts, thus minimizing the effect of natural variations in activity patterns at different times of the day. Modeling with shifts taken into account improves behavioral consistency and enables more meaningful anomaly detection.

The second level of segmentation is organizational role groups. Rather than modeling all users at once, employees are grouped by functionally similar roles. This approach allows us to compare users to behaviorally similar peers to reduce the structural heterogeneity in the dataset. The framework combines shift based and role based segmentation to form contextual groups which provide a more stable environment for behavioral modeling and anomaly detection.

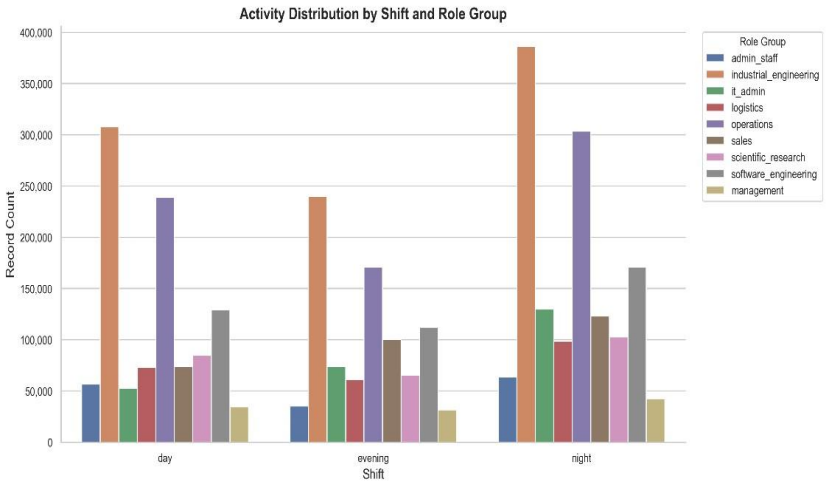

***Fig. 3**. Activity Distribution Across Shift and Role Groups.*

This resulting Shift × Role Group structure is the basis for the later stages of clustering and anomaly detection. Such a contextual approach reduces false positives, improves anomaly specificity and allows the framework to differentiate genuine behavioral anomalies from legitimate changes in user activity.

HDBSCAN-Based Behavioral Clustering

The proposed framework uses Hierarchical Density Based Spatial Clustering of Applications with Noise (HDBSCAN) [6] to identify natural behavioral patterns in the contextual user groups. Unlike traditional clustering algorithms where the number of clusters has to be defined beforehand, HDBSCAN automatically finds behavioral groups of varying density and simultaneously detects low density observations that could represent anomalous behavior.

Behavioral feature values were scaled using RobustScaler prior to clustering to reduce the impact of outliers and to enhance clustering stability. Then the HDBSCAN was applied separately within each group of contexts with the minimum cluster size of 50 observations and minimum sample size of 10. These parameters were chosen to enable the formation of meaningful behavioral groups while retaining sensitivity to abnormal activity patterns.

The clustering process identified 2,764 behavioral clusters across the data set. Observations assigned to dense clusters were assumed to represent normal behavior patterns and observations assigned to noise (cluster = -1) were assumed to represent potential anomalies. HDBSCAN directly finds observations that do not belong to any dense behavioral region, which is an effective method for finding rare and unusual activity patterns without needing labeled threat data.

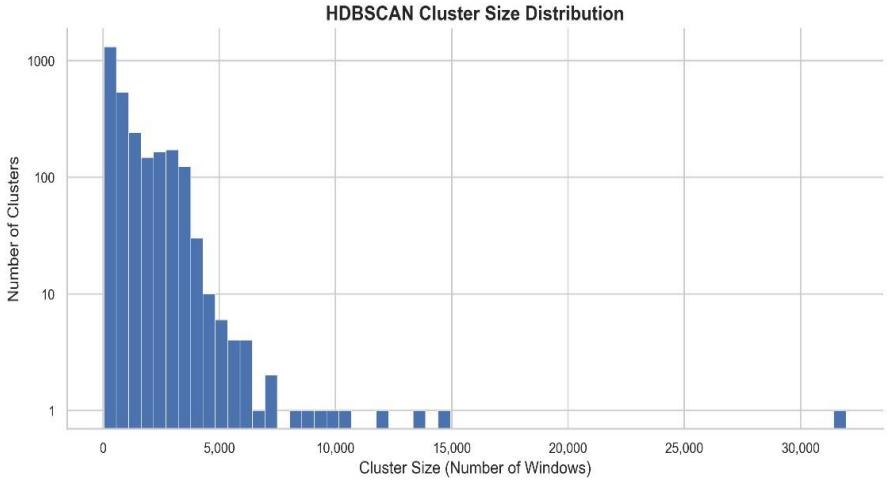

***Fig. 4**. Distribution of HDBSCAN Cluster Sizes Across Contextual User Groups.*

Figure 4 shows clusters of varying sizes produced by HDBSCAN, indicative of heterogeneous user behavior. The experimental analysis indicated that the clustering framework captured different behavioral modes in the organization and identified about 213,628 observations

as anomalous, i.e., an anomaly rate of about 6.34%. The results show the usefulness of density based behavioral modeling for detection of large-scale insider threats.

Isolation Forest Validation

Although HDBSCAN provides an effective mechanism for discovering density based behavioral anomalies, additional validation is required to increase confidence in anomaly detection results. To achieve this, the proposed framework adds Isolation Forest as a second anomaly detection model in the same behavioral feature space [5].

The Isolation Forest finds anomalies by recursively splitting observations on randomly chosen features and split values. Anomalies are observations which can be separated with less partitions , therefore they are quite different from most of the data . Compared to density based clustering, Isolation Forest directly estimates anomaly scores and is computationally efficient for large scale datasets.

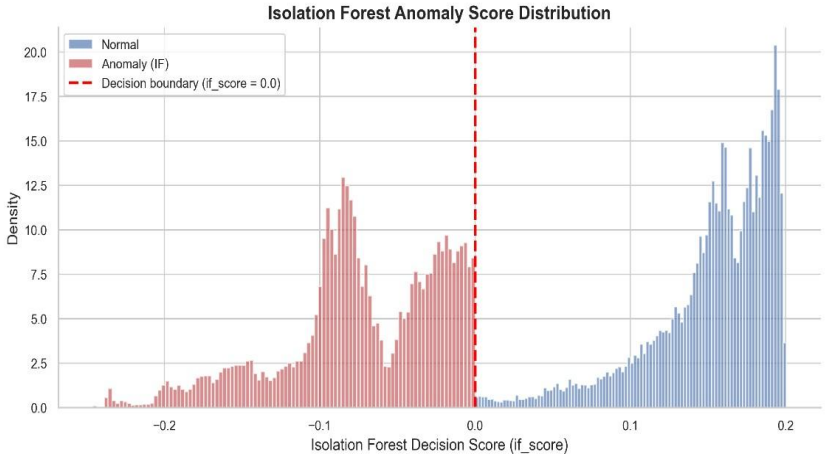

Fig. 5. *Isolation Forest Decision Score Distribution Showing Normal and Anomalous Samples*.

The distribution of anomaly score obtained by Isolation Forest is shown in Figure 5.

In the proposed framework, Isolation Forest is employed as a secondary validation mechanism to confirm anomalies spotted by HDBSCAN. The system combines density based behavioral discovery with tree based anomaly scoring, thus decreasing dependency on a single modeling approach and increasing detection robustness. This hybrid approach allows the framework to identify low density behavioral observations as well as statistically isolated activity patterns that may indicate potential insider threats.

Risk Scoring and Alert Generation

The final stage in the proposed framework involves translating the results of anomaly detection into actionable security intelligence. Results from HDBSCAN clustering and Isolation Forest validation are combined to produce behavioral risk indicators for each user activity window. Each model (or both) marks observations as anomalous and provide a higher risk score. Normal observations are scored with lower risk.

Risk scores are aggregated over windows of user activity to detect persistent abnormal behavior. High risk users are identified and prioritized for further investigation, for repeated anomalous patterns shown by the users. This allows security analysts to focus their efforts on users whose activity falls outside the behavioral baselines that have been set over time.

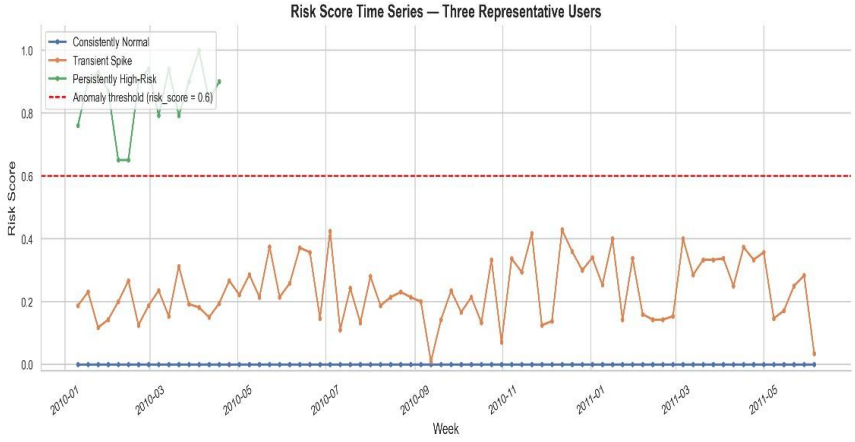

Fig. 6. *Temporal Risk Score Evolution for Representative Users.*

Figure 6 shows representative risk score trajectories over time for different user behaviors.

The framework supports an optional layer of email compliance auditing to provide an additional contextual assessment. Email activity within anomalous behavioral windows can be examined for policy violations, suspicious communication patterns, or indicators of potential data exfiltration. The final system output contains anomaly labels, risk scores, high risk user indicators and security alerts that can be displayed through an analyst dashboard for investigation and decision support.

Email Compliance Auditing Pipeline

Behavioral anomaly detection can identify suspicious user activities, but often provides limited contextual evidence of the intent and impact of those activities. To overcome this limitation, BehaviorGuard-AI incorporates an Email Compliance Auditing Pipeline that performs policy aware investigation on users identified to be high risk by the anomaly detection framework.

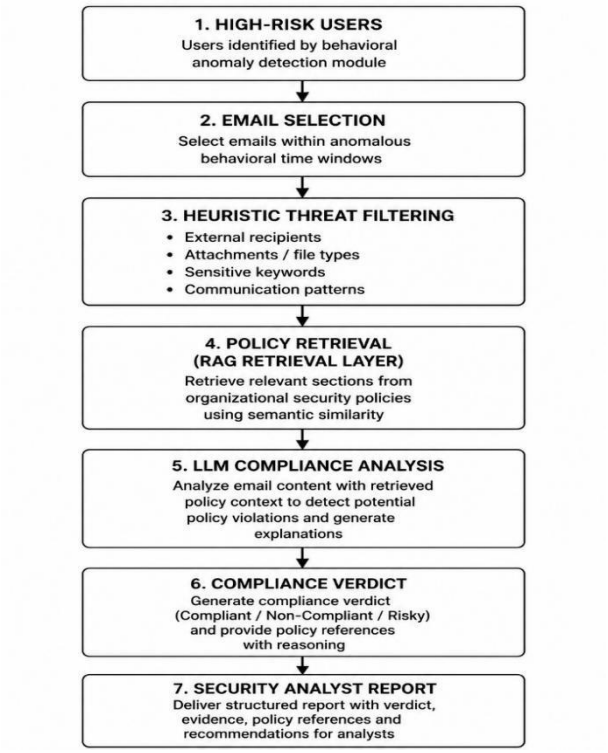

Fig. 7. *Workflow of the Proposed Email Compliance Auditing Pipeline*

The pipeline is only triggered for users with behavioral risk scores above certain thresholds. The anomaly-driven approach reduces computational overhead by focusing email analysis on a small subset of potentially suspicious users. Temporal correlation of email records is first performed to anomalous behavioral windows such that only communications pertaining to suspect activity periods are taken up for investigation.

After temporal correlation, a heuristic threat filtering step analyzes emails based on indicators like external recipients, file attachments, sensitive keywords, and communication patterns. Emails that pass the filtering criteria are sent to a Retrieval Augmented Generation (RAG) framework for policy aware compliance assessment.

The RAG pipeline is performed in two stages. It begins by using semantic similarity search to extract the most relevant sections of the organization's security and privacy policy. Second, the chosen policy content and email communications are jointly analyzed by a Large Language Model (LLM) to detect potential policy violations, generate compliance verdicts, and furnish supporting explanations. The resulting compliance reports provide contextual evidence that supports behavioral anomaly scores and supports security analysts in insider threat investigations.

Results and Discussion

We evaluated BehaviorGuard-AI on a large synthetic enterprise authentication dataset with approximately 3,366,264 hourly activity records from nearly 3,000 users. The goal was to evaluate the framework's effectiveness at identifying anomalous behavioral patterns without the use of labeled insider threat examples.

The system performed shift aware segmentation and role based peer grouping to divide users into behaviorally comparable cohorts. This contextual segmentation reduced the effect of behavioral heterogeneity and made the results of anomaly easier to interpret. HDBSCAN found 2,764 unique behavioral clusters and identified 213,628 observations as anomalous which is an overall anomaly rate of around 6.34%. This demonstrates that the framework is capable of identifying rare patterns while the majority of activity is classified as normal.

Fig. 8. *Anomaly Rate Across Shift and Role Groups.*

Figure 8 summarizes the anomaly rates for each of the shift and role combinations.

We added Isolation Forest as a further validation step. Confidence in detections was increased by combining density based cluster discovery with tree based anomaly scoring, decreasing reliance on any single method. The resulting risk scores helped to rank and prioritize high risk users for analyst review.

The Email Compliance Auditing module delivered targeted contextual analysis for users at elevated risk. Instead of scanning all communications, the auditing pipeline only focused on emails temporally linked to anomalous windows, thus cutting unnecessary processing and improving investigation efficiency. Emails were scored against organizational policy and correlated with anomalous activity through a Retrieval Augmented Generation (RAG) workflow that provided policy relevant evidence to supplement raw anomaly indicators.

The experiments showed that the proposed BehaviorGuard-AI was able to detect abnormal behavior in a large enterprise dataset without using labeled attack data. HDBSCAN reduced millions of activity records into meaningful behavioral groups, and only about 6.34% of the observations were flagged as anomalous. That meant a lot fewer events for analysts to pore over. Instead of having to audit each user activity record, the framework helped to narrow the investigation down to a manageable set of high risk users prior to the email auditing stage.

The key experimental outcomes obtained from the proposed framework are summarized in Table 1.

Table 1. *Summary of Experimental Results*

Metric	**Value**
Total Activity Records	3,366,264
Enterprise Users	~3,000
Behavioral Features	10
Clusters Discovered	2,764
Anomalous Observations	213,628
Anomaly Rate	6.34%
Detection Approach	HDBSCAN+Isolation Forest

As shown in Table 1, the proposed framework was able to identify 2,764 behavioral clusters and detect anomalous observations, which correspond to approximately 6.34% of the activity records that were analyzed.

The Email Compliance Auditing Pipeline increases contextual investigation for high risk users. However, the primary goal of the proposed framework is behavioral anomaly detection. The pipeline does not analyze all organizational communications, but rather selectively analyzes

emails related to anomalous behavioral windows identified by the UEBA framework. The net effect is a more focused approach, eliminating unnecessary analysis but putting investigative resources on activities that appear suspicious.

The framework can provide contextual compliance assessments by combining anomaly driven user selection, temporal email correlation, heuristic threat filtering and RAG based policy evaluation. The system connects the communication content with the organizational security policies, which provides supporting evidence for security analysts in addition to behavioral anomaly scores. This capability improves the explainability of insider threat investigations and supports analysts in making a determination as to whether anomalous behavior may be related to policy violations, exposure of sensitive information, or other security issues.

Conclusion and Future Work

In this paper, we propose BehaviorGuard-AI, a context aware UEBA framework that leverages behavioral analytics and unsupervised learning to detect insider threats. We detect anomalous user activity in large enterprise environments by using contextual user profiling, engineered behavioral features, shift based segmentation, role aware peer grouping combined with HDBSCAN clustering and Isolation Forest validation. Experiments on a synthetic enterprise dataset showed that the framework was able to find meaningful behavioral clusters and surface anomalous observations without the need for labeled threat data.

BehaviorGuard-AI goes beyond detecting statistical anomalies and adds an Email Compliance Auditing Pipeline to support investigations. The pipeline consists of an anomaly driven selection, temporal correlation of emails with anomalous windows, heuristic filtering and a Retrieval Augmented Generation (RAG) workflow for policy evaluation. This targeted, policy aware analysis provides analysts with contextual evidence to determine if anomalous behavior is related to policy violations, exposure of sensitive data or other risks.

The framework was successfully applied to the synthetic enterprise dataset used in this study but there are still some limitations. The experiments were not performed on real enterprise logs with confirmed insider threat labels, thus it is difficult to measure the detection accuracy with metrics, such as precision and recall . The HDBSCAN parameters were also manually chosen, and may need to be adapted for organizations with different patterns of user behavior. We intend to test the framework in real enterprise settings, study adaptive risk thresholds for different groups of roles, and improve the email auditing stage with more complete organizational policy data.

Reference

1. M. Bishop and D. Gates, “Defining the Insider Threat,” Proceedings of the Cyber Security and Information Intelligence Research Workshop, pp. 1–3, 2008.
2. R. Mitchell and I. R. Chen, “Behavior Rule Specification-Based Intrusion Detection for Safety Critical
3. Medical Cyber Physical Systems,” IEEE Transactions on Dependable and Secure Computing, vol. 12, no. 1, pp. 16–30, 2015.
4. A. Tuor, S. Kaplan, B. Hutchinson, N. Nichols, and S. Robinson, “Deep Learning for Unsupervised Insider Threat Detection in Structured Cybersecurity Data Streams,” AAAI Workshops, 2017.
5. M. Goldstein and S. Uchida, “A Comparative Evaluation of Unsupervised Anomaly Detection Algorithms for Multivariate Data,” PLoS ONE, vol. 11, no. 4, 2016.
6. F. T. Liu, K. M. Ting, and Z. H. Zhou, “Isolation Forest,” Proceedings of the IEEE International Conference on Data Mining (ICDM), pp. 413–422, 2008.
7. R. J. G. B. Campello, D. Moulavi, and J. Sander, “Density-Based Clustering Based on Hierarchical Density Estimates,” Pacific-Asia Conference on Knowledge Discovery and Data Mining, pp. 160–172, 2013.
8. M. A. Maloof and R. S. Stephens, “Elicit: A System for Detecting Insider Threats Using Behavioral Analysis,” Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications, vol. 4, no. 4, pp. 73–95, 2013.
9. S. Axelsson, “The Base-Rate Fallacy and the Difficulty of Intrusion Detection,” ACM Transactions on Information and System Security, vol. 3, no. 3, pp. 186–205, 2000.