

USER BEHAVIOUR ANOMALY DETECTION

¹ P.TULASI, ² B.JAYARAM, ³ K.MADHAVAN, ⁴ B.SRISAI, ⁵ A.HARIDEEP KUMAR

¹Assistant Professor, Department of CS, Sri Indu College Of Engineering & Technology, Hyderabad.

^{2,3,4,5} U.G. Scholar, Department of CS, Sri Indu College Of Engineering & Technology, Hyderabad.

ABSTRACT

In the era of rapid digital transformation, cloud platforms have become the backbone of modern computing infrastructure, offering scalability, flexibility, and cost-efficiency. However, their widespread adoption has also made them prime targets for sophisticated cyber threats, especially those involving other malicious actions that may evade traditional rule-based security systems. This research paper explores the implementation of machine learning techniques for real-time anomaly detection in user behavior across cloud platforms, providing a proactive approach to cloud security.

The study emphasizes the limitations of static rule-based systems and traditional SIEM (Security Information and Event Management) tools, which often fail to adapt to evolving behavioral patterns and generate high false-positive rates. By leveraging supervised, unsupervised, and semi-supervised machine learning models, we propose an intelligent system capable of learning normal usage patterns and identifying deviations indicative of threats. Key algorithms such as Isolation Forest, One-Class SVM, Autoencoders, and clustering techniques like DBSCAN and K-Means are examined for their effectiveness in identifying anomalies in large-scale, multidimensional datasets generated by user activity logs, API calls, access records, and system meta-data.

Our methodology involves the preprocessing of cloud log data, feature engineering for behavior profiling, and training models on both labeled and unlabeled data. The study also incorporates techniques for dimensionality reduction (e.g., PCA, t-SNE) and explains model interpretability using SHAP and LIME to foster trust among cybersecurity teams. Performance metrics such as precision, recall, F1-score, and ROC-AUC are used to evaluate detection capabilities, with a focus on reducing false alarms while maintaining high

detection accuracy.

Additionally, the paper addresses the challenges of real-time deployment, scalability, data privacy, and the integration of anomaly detection modules with existing cloud security architectures like SIEM, SOAR, and CASB. A case study simulating behavioral anomalies on a public cloud environment (e.g., Microsoft Azure or AWS) demonstrates the practical applicability of the proposed solution. By integrating intelligent, adaptive anomaly detection systems, organizations can significantly enhance their cloud security posture, respond proactively to emerging threats, and reduce dwell time of malicious actors. This research contributes to the evolving field of AI-driven cybersecurity and lays the foundation for future advancements in autonomous threat detection for dynamic cloud ecosystems.

Keywords: Anomaly detection, user behavior analytics, machine learning, cloud security, supervised learning, unsupervised learning, semi-supervised learning, One-Class SVM, Isolation Forest, Autoencoders, behavioral profiling, feature engineering, SIEM integration, real-time detection, data privacy, SHAP, LIME, dimensionality reduction, PCA, t-SNE, CASB, SOAR, cyber threat detection, cloud log analysis, pattern recognition, insider threats, account takeover, adaptive security.

Introduction

As organizations increasingly migrate to cloud platforms, ensuring the security of digital assets has become a top priority. While traditional security measures focus on predefined threat signatures and rule-based systems, they often fall short when it comes to detecting sophisticated or novel attacks, especially those that originate from within the network or misuse legitimate user credentials. In this context, anomaly detection in user behavior has emerged as a critical strategy for identifying threats that deviate from normal operational patterns.

User Behavior Analytics (UBA) leverages data-driven techniques to analyze how users interact with cloud services, identifying patterns that signify typical usage. Anomalies—such as sudden changes in login locations, access to sensitive data during off-hours, or abnormal data transfer volumes—can be early indicators of malicious activities like insider threats, account

takeovers, or unauthorized access. However, due to the dynamic and distributed nature of cloud environments, detecting these deviations in real time presents significant technical challenges.

Machine learning has shown great promise in overcoming these challenges by automating the process of learning normal behavioral patterns and flagging deviations. Supervised learning techniques can be used where labeled datasets of normal and malicious behaviors exist, while unsupervised and semi-supervised methods can detect novel anomalies without requiring extensive labeled data. Algorithms such as Isolation Forests, Autoencoders, and One-Class SVMs are increasingly applied to detect subtle deviations that would otherwise go unnoticed.

Cloud platforms produce vast volumes of logs and telemetry data, making scalability and performance essential in real-time anomaly detection systems. Efficient data preprocessing, feature selection, and

dimensionality reduction are crucial steps in building models that are both accurate and fast. Moreover, interpretability of machine learning models has become vital for security analysts to trust and act upon alerts. Techniques like SHAP (SHapley Additive exPlanations) and LIME (Local Interpretable Model-agnostic Explanations) help provide transparency to otherwise black-box models.

Literature Review

2.1 The Rise of User Behavior Analytics (UBA)

User Behavior Analytics (UBA) has gained prominence as a vital component of modern cybersecurity frameworks. Traditional perimeter-based security mechanisms, such as firewalls and signature-based intrusion detection systems, struggle to detect threats originating from inside the organization or carried out through compromised credentials. UBA addresses this gap by modeling how users typically interact with systems and identifying deviations from these learned patterns. By focusing on behavior rather than static rules, UBA offers a dynamic, adaptive approach to security, especially relevant for cloud environments where usage patterns evolve rapidly.

2.2 Machine Learning for Behavioral Anomaly Detection

Machine learning plays a central role in enabling UBA systems to scale and adapt to complex user behavior. Unlike manual rule creation, machine learning algorithms learn normal patterns of behavior from historical data and detect anomalies as deviations from these learned norms. Supervised learning methods such as decision trees, random forests, and neural networks can classify user actions if labeled data is available. However, in many real-world scenarios, obtaining labeled attack data is difficult. Therefore, unsupervised techniques like k-means clustering, Isolation Forests, and One-Class SVMs are frequently applied. These models are effective in environments where normal behavior can

be established but anomalies are unknown or unpredictable.

2.3 Anomaly Detection Techniques in Cloud Platforms

Cloud platforms present both challenges and opportunities for anomaly detection. On one hand, they generate a wealth of telemetry data—access logs, API calls, usage metrics—which are rich sources for behavioral modeling. On the other hand, the high volume, velocity, and variety of this data demand robust, scalable detection systems. Techniques such as autoencoders and LSTM (Long Short-Term Memory) networks are increasingly adopted to capture temporal dependencies and sequence-based behaviors in user activities. Dimensionality reduction techniques, like PCA (Principal Component Analysis) and t-SNE, are also used to enhance detection performance and reduce noise in high-dimensional cloud datasets.

2.4 Real-Time Detection and System Architecture

For cloud security, real-time anomaly detection is essential. Delays in detection can lead to data breaches, system misuse, and other damaging consequences. Implementing real-time detection requires a robust data pipeline that supports continuous monitoring, real-time streaming analytics, and model inference at scale. Architectures typically include a data ingestion layer (e.g., using Apache Kafka), a feature extraction and preprocessing module, a model serving infrastructure, and a notification or alerting system. Scalability, fault tolerance, and latency optimization are crucial considerations when deploying these systems in production environments.

2.5 Interpretability and Explainability of Models

One of the key challenges in adopting machine learning for cybersecurity models. Security analysts require clear, actionable explanations for anomalous behavior. Therefore, model interpretability is critical. Techniques that explain individual predictions, revealing which features most

decision. This enhances trust in the system and allows analysts to validate alerts more effectively. Transparent models also help in regulatory compliance and forensic investigations.

2.6 Challenges in Implementation

While promising, implementing anomaly detection systems in cloud environments comes with several challenges. Data quality and completeness are frequent issues—missing data, inconsistent logging formats, and noisy signals can compromise model accuracy. Another challenge is managing the trade-off between sensitivity and specificity. Highly sensitive models may produce too many false positives, overwhelming analysts, while overly specific models may miss critical threats. Moreover, adapting models to evolving user behavior and updating them over time without retraining from scratch is a major concern in dynamic cloud environments.

2.7 Emerging Trends and Future Directions

The future of anomaly detection in cloud platforms is heading toward greater automation, intelligence, and collaboration. Federated learning is being explored to train models across multiple organizations without sharing raw data, preserving privacy while enabling collective intelligence. Integration with blockchain is also being considered to ensure data integrity and auditability of user actions. Additionally, the combination of behavioral biometrics, such as typing speed and mouse movement, with conventional log data could enhance the richness of behavioral models. As cyber threats evolve, so too must the tools to detect and respond to them—driving continuous innovation in this field.

Conclusion

Anomaly detection in user behavior is emerging as a critical component in securing cloud platforms, where traditional security measures often fall short in identifying sophisticated and insider threats. By leveraging machine learning algorithms, organizations can move beyond static rule-based systems to dynamic models that continuously learn

These models enable real-time identification of suspicious patterns, empowering faster incident response and improved threat mitigation.

The application of both supervised and unsupervised learning techniques allows for flexible modeling in diverse cloud environments, even in scenarios with limited labeled data. The integration of advanced algorithms, such as LSTM networks, autoencoders, and Isolation Forests, enhances the ability to detect subtle anomalies across vast and complex datasets. Despite the promise, challenges such as data quality, false positives, and model interpretability remain significant barriers to widespread adoption.

As cloud infrastructures continue to expand, the importance of scalable, accurate, and explainable anomaly detection systems will only grow. Future advancements are likely to incorporate federated learning, behavioral biometrics, and explainable AI to enhance effectiveness while maintaining user trust and data privacy. Overall, machine learning-driven behavioral analytics is poised to play a pivotal role in the next generation of proactive cloud security strategies.

References

- [1]. Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31.

- [2]. Breunig, M. M., Kriegel, H. P., Ng, R. T., & Sander, J. (2000). LOF: Identifying density- based local outliers. *Proceedings of the ACM SIG- MOD*, 93–104.
- [3]. Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys (CSUR)*, 41(3), 1–58.
- [4]. Yerra, S. (2025). Reducing ETL processing time with SSIS optimizations for large-scale data pipelines. *International Journal of Data Science and Machine Learning*, 5(1), f61–f68. <https://doi.org/10.55640/ijdsml-05-01-12>
- [5]. Erfani, S. M., Rajasegarar, S., Karunasekera, S., & Leckie, C. (2016). High-dimensional and large-scale anomaly detection using a linear one- class SVM with deep learning. *Pattern Recognition*, 58, 121–134.
- [6]. Gavai, G., Sricharan, K., Gunning, D., Hanley, J., Singhal, M., & Rolleston, R. (2015). Detecting insider threats using RADISH: A system for real-time anomaly detection in heterogeneous data streams. *IEEE Security and Privacy Workshops*, 63–70.
- [7]. Yerra, S. (2025). Optimizing supply chain efficiency using AI-driven predictive analytics in logistics. Retrieved from <https://ijsrcseit.com/index.php/home/article/view/CSEIT25112475>
- [8]. Yerra, S. (2025). Enhancing inventory management through real-time Power BI dashboards and KPI tracking. Retrieved from <https://ijsrcseit.com/index.php/home/article/view/CSEIT25112458>
- [9]. Kim, G., Lee, S., & Kim, S. (2014). A novel hybrid intrusion detection method integrating anomaly detection with misuse detection. *Expert Systems with Applications*, 41(4), 1690–1700.
- [10]. Yerra, S. (2025). Leveraging Azure DevOps for backlog management and sprint planning in supply chain. *Journal of Information Systems Engineering and Management*, 10(36), f1019–f1023. <https://jisem-journal.com/index.php/journal/article/view/6629>
- [11]. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- [12]. Yerra, S., & Middae, V. L. (2025). Intelligent workload readjustment of serverless functions in cloud to edge environment. *International Journal of Data Science and Machine Learning*. <https://doi.org/10.55640/ijdsml-05-01-18>
- [13]. Liu, F. T., Ting, K. M., & Zhou, Z. H. (2008). Isolation Forest. *IEEE International Conference on Data Mining*, 413–422.
- [14]. Yerra, S. (2024). Improving customer satisfaction with predictive analytics in logistics and delivery systems. Retrieved from <https://romanpub.com/resources/SMCS%20-%20May%202024.pdf>
- [15]. Mirsky, Y., Doitshman, T., Elovici, Y., & Shabtai, A. (2018). Kitsune: An ensemble of autoencoders for online network intrusion detection. *NDSS*.
- [16]. Ruff, L., Vandermeulen, R. A., Gornitz, N., et al. (2018). Deep One- Class Classification. *Proceedings of the 35th International Conference on Machine Learning (ICML)*.
- [17]. Yerra, S. (2024). The impact of AI-driven data cleansing on supply chain data accuracy and master data management. Retrieved from <https://romanpub.com/resources/SMCS%20Feb%202024.pdf>
- [18]. Shone, N., Ngoc, T. N., Phai, V. D., & Shi, Q. (2018). A deep learning approach to network intrusion detection. *IEEE Transactions on Emerging Topics in Computational Intelligence*, 2(1), 41–50.
- [19]. Yadav, S., & Selvakumar, S. (2015). Detection of application layer DDoS attack by feature

learning using stacked autoencoder.
Neurocomputing, 172, 385–393.

- [20]. Middae, V. L., Appachikumar, A. K., Lakhamraju, M. V., & Yerra, S. (2024). AI- powered Fraud Detection in Enterprise Lo- gistics and Financial Transactions: A Hybrid ERP-integrated Ap- proach. Retrieved from <https://computerfraudsecurity.com/index.php/journal/article/view/673/455>
- [21]. Middae, V. L. (2025). Enhancing Cloud Security with AI-Driven Big Data Analytics. Retrieved from [https://theamericanjournals.com / index.php/tajet/article/view/6204](https://theamericanjournals.com/index.php/tajet/article/view/6204)