10.48047/jocaaa.2025.34.08.16

Distributed Marketplace Intelligence: Real-Time Anomaly Detection at Cloud Scale

Ridhima Mahajan1, Sarat Mahavratayajula 2, Suvodeep Pyne3

1 Senior Software Engineer
2 Senior software Engineer at Sherwin-Williams
3 Staff Software Engineer II at Startree Inc

Abstract

The rapid expansion of digital marketplaces has amplified the need for intelligent, real-time anomaly detection to safeguard transactional integrity, infrastructural stability, and user trust. This study presents a distributed marketplace intelligence framework designed to detect anomalies at cloud scale by integrating cloud-native architectures, machine learning, and deep learning models within a scalable stream processing environment. Data from both synthetic simulations and anonymized marketplace logs were analyzed across transactional, behavioral, and infrastructural dimensions. Variables such as transaction volume, session duration, clickstream depth, CPU utilization, and network latency were systematically monitored and processed. The proposed framework incorporated statistical methods, Isolation Forest, One-Class SVM, LSTM Autoencoder, and a hybrid ensemble, with the latter achieving the highest performance (Precision = 0.92, Recall = 0.91, F1-score = 0.91, AUC-ROC = 0.94). Inferential statistical analysis confirmed significant differences in anomaly occurrence across transaction types, geographic regions, and payment methods, highlighting the contextual sensitivity of detection models. Scalability testing demonstrated near-linear throughput growth and reduced detection latency as cluster size increased, alongside improvements in uptime and reduced false positives. These results validate the framework's capacity to support proactive, real-time monitoring in large-scale marketplaces. While limitations include partial reliance on synthetic data and the absence of automated anomaly response mechanisms, the framework offers a robust foundation for future research on adaptive, privacy-preserving, and self-healing detection systems. By advancing the state of distributed anomaly detection, this study contributes both theoretically and practically to the development of resilient, trustworthy digital marketplaces.

Keywords: Distributed marketplace intelligence, real-time anomaly detection, cloud scale, machine learning, hybrid ensemble, scalability, transaction monitoring

10.48047/jocaaa.2025.34.08.16

Introduction

Growing complexity of digital marketplaces

The rise of large-scale online marketplaces has transformed the global economy by connecting millions of buyers and sellers across borders (Gandhi & Sharma, 2025). Platforms such as Amazon, Alibaba, and eBay host billions of transactions daily, spanning diverse product categories, payment systems, and logistics networks. With this unprecedented scale comes significant complexity in monitoring system health, ensuring fairness, and maintaining consumer trust (Lianet al., 2025). Traditional monitoring approaches, which rely on periodic checks or rule-based systems, often fail to capture emerging irregularities in real time. As a result, digital marketplaces face heightened risks of fraud, performance degradation, and operational inefficiencies that directly impact customer satisfaction and business continuity (Bin Mofidul et al., 2022).

The importance of real-time anomaly detection

Anomalies in marketplaces may manifest as sudden surges in transaction volumes, fraudulent listings, coordinated bot activities, or unexpected latency in cloud-based infrastructure. Detecting such anomalies in real time is critical for preventing financial losses, safeguarding platform integrity, and ensuring compliance with regulatory requirements (Nwachukwu et al., 2024). Unlike static data systems, cloud-scale marketplaces operate in dynamic environments where data flows continuously at high velocity and volume. This makes batch processing methods inadequate, pushing the need for adaptive, real-time anomaly detection frameworks. By embedding anomaly detection within distributed marketplace intelligence systems, platforms can shift from reactive responses to proactive interventions that mitigate risks before they escalate (Kotha, 2025).

Challenges in distributed anomaly detection at scale

Implementing anomaly detection at cloud scale introduces several challenges. First, data heterogeneity complicates the detection process, as marketplaces deal with structured, semi-structured, and unstructured data simultaneously. Second, the distributed nature of cloud infrastructures requires scalable algorithms capable of handling decentralized data streams while maintaining low-latency performance (Sarpal et al., 2023). Third, the definition of “normal” behavior in marketplaces is highly contextual, varying across regions, time zones, and user segments. Static thresholds fail to capture these nuances, necessitating adaptive

10.48047/jocaaa.2025.34.08.16

models that evolve with market behavior. Lastly, privacy and security considerations demand techniques that can analyze sensitive transactional data without compromising compliance with data protection regulations such as GDPR and CCPA (Rousopoulou et al., 2022).

Advances in cloud-native architectures and machine learning

Recent advances in cloud-native architectures, container orchestration, and machine learning have created new opportunities for developing distributed marketplace intelligence. Stream processing frameworks such as Apache Kafka, Flink, and Spark Streaming enable scalable real-time data ingestion and analysis, while microservices architectures allow anomaly detection modules to be deployed and scaled independently (Islam et al., 2021). In parallel, machine learning and deep learning algorithms have improved anomaly detection accuracy by learning complex patterns across multidimensional datasets. Hybrid approaches that combine unsupervised anomaly detection with domain-specific rules are increasingly being used to balance interpretability with predictive power (Zhang et al., 2024). Together, these innovations form the backbone of cloud-scale solutions that support continuous monitoring and rapid anomaly response.

Research gap and motivation for the study

Despite technological progress, several gaps remain in the literature on anomaly detection for cloud-scale marketplaces. Existing frameworks often focus on isolated detection methods without addressing the integration of distributed intelligence across heterogeneous systems. Moreover, most prior studies emphasize either system-level anomalies (such as latency and downtime) or transaction-level anomalies (such as fraud and abuse), but few attempt to unify these perspectives into a holistic monitoring system. The absence of standardized evaluation metrics for marketplace-specific anomalies also hampers the development of generalizable solutions. This research is motivated by the need to design and evaluate a scalable, distributed anomaly detection framework that integrates marketplace intelligence with real-time monitoring, thereby enhancing both system resilience and user trust.

Aim and structure of the article

The aim of this research is to develop and validate a framework for distributed marketplace intelligence that supports real-time anomaly detection at cloud scale. The proposed study explores a combination of stream processing, distributed data analytics, and machine learning models to detect and interpret anomalies across transactional and infrastructural domains. The

10.48047/jocaaa.2025.34.08.16

remainder of the article is structured as follows: the methodology outlines the system design and variables considered; the results present empirical evaluations on simulated and real-world marketplace data; the discussion interprets the findings in light of scalability, efficiency, and ethical considerations; and the conclusion highlights contributions, limitations, and future research directions.

Methodology

Research design and framework overview

This study follows a quantitative systems-engineering approach to design, implement, and validate a distributed framework for real-time anomaly detection at cloud scale. The methodological design integrates cloud-native architectures, advanced machine learning models, and real-time stream processing technologies to support distributed marketplace intelligence. Both synthetic marketplace simulations and anonymized real-world e-commerce transaction logs were used to test the framework. The primary objective of the design is to ensure scalability, adaptability, and statistical reliability under conditions of high data velocity, variety, and volume.

Data sources and variables considered

The study relied on two categories of datasets: controlled synthetic simulations designed to reproduce common and rare anomalies in marketplace environments, and anonymized logs obtained from large-scale cloud-hosted e-commerce platforms. Variables considered were grouped into transactional, behavioral, and infrastructural domains. Transactional variables included transaction ID, user ID, product ID, transaction volume, frequency of purchases, transaction timestamp, payment method, refund rate, order size, geographic location, currency exchange rate, promotional code usage, and latency between order placement and payment confirmation. Behavioral variables included login frequency, session duration, browsing-to-purchase ratio, clickstream paths, anomalous user-agent strings, number of concurrent sessions per user, repetitive actions indicative of bots, and average time-to-purchase. Infrastructural variables included CPU utilization, memory consumption, disk I/O operations, network throughput, API response time, error rate, microservice load distribution, container restart frequency, and service latency. These variables together captured anomalies spanning micro-level user activities and macro-level infrastructural behaviors.

Data preprocessing and transformation

10.48047/jocaaa.2025.34.08.16

To prepare the data for real-time analysis, several preprocessing steps were applied. Missing values were imputed using k-nearest neighbors for continuous variables and mode substitution for categorical variables. Time-related variables were standardized into a coordinated universal time (UTC) format, and categorical features such as payment method and user-agent type were one-hot encoded. To address high-dimensionality, Principal Component Analysis (PCA) was employed to reduce complexity while preserving at least 95% of variance. Real-time data ingestion was carried out through Apache Kafka pipelines, while transformation and feature extraction were performed on streaming data using Apache Flink, ensuring sub-second latency. This preprocessing pipeline enabled consistent feature representation across both training and testing phases.

Anomaly detection models and algorithms

The anomaly detection methodology combined multiple approaches to maximize accuracy. Statistical models such as z-score thresholds, interquartile range (IQR) detection, and moving averages were used for baseline anomaly identification. Machine learning models, including Isolation Forest, One-Class Support Vector Machines, and Random Cut Forest, were employed to capture anomalies in high-dimensional spaces. Deep learning models such as Long Short-Term Memory (LSTM) recurrent networks and Autoencoders were applied to identify temporal anomalies and reconstruct normal behavior sequences. Finally, a hybrid ensemble was developed by combining these models, with ensemble weights optimized using grid search on validation datasets to maximize F1-score performance.

Distributed architecture and deployment

The framework was implemented in a cloud-native environment using Kubernetes clusters for orchestration and Docker containers for modular deployment. Anomaly detection services were designed as microservices, enabling independent scaling and fault isolation. Streaming pipelines were deployed through Apache Flink to support continuous ingestion and distributed analysis. Asynchronous communication between detection modules was achieved through RESTful APIs, and node replication across availability zones ensured resilience and fault tolerance. The distributed architecture allowed the system to process millions of transactions per second with minimal latency, meeting the requirements of real-time monitoring in cloud-scale marketplaces.

Statistical analysis and evaluation metrics

10.48047/jocaaa.2025.34.08.16

The performance of the framework was analyzed using a combination of descriptive and inferential statistical techniques. Descriptive statistics summarized the central tendency and dispersion of transactional, behavioral, and infrastructural variables to establish baselines for normal behavior. Inferential tests such as analysis of variance (ANOVA) and chi-square tests were applied to examine differences in anomaly occurrence across user groups, transaction categories, and geographic regions. Model evaluation was based on key performance indicators including precision, recall, F1-score, and area under the receiver operating characteristic curve (AUC-ROC). Additional marketplace-specific metrics included time-to-detection (TTD), false positive rate (FPR), and anomaly resolution latency. Kaplan–Meier survival analysis was applied to assess system uptime under anomalous conditions, while logistic regression and multiple regression models were used to test the influence of infrastructural variables such as CPU utilization and network throughput on anomaly occurrence probabilities.

Validation and experimental setup

Validation of the proposed framework was carried out in two phases. In the first phase, synthetic anomalies were injected into simulated datasets to test detection accuracy under controlled conditions. In the second phase, the framework was applied to real-world e-commerce logs that contained expert-labeled anomalies. Data were split into training, validation, and test sets using a 70-15-15 partition, and k-fold cross-validation with k = 10 ensured model generalizability. Hyperparameter tuning of the machine learning and deep learning models was performed using Bayesian optimization. Experiments were conducted on clusters of varying sizes, including 50-node, 100-node, and 200-node deployments, to test scalability. The framework's results were benchmarked against baseline threshold-based detection methods, allowing comparative evaluation of accuracy, efficiency, and scalability.

Results

The descriptive analysis provided a clear overview of transactional, behavioral, and infrastructural dynamics within the cloud-scale marketplace. As shown in Table 1, the average transaction volume was 1,250 per observation period with a maximum of 1,700, reflecting high variability in marketplace activity. Refund rates remained relatively stable, averaging 3.2%, while user session durations averaged 12.4 minutes but extended up to 28.5 minutes in certain cases. Clickstream depth averaged 18.0, indicating complex browsing behaviors. From the infrastructural perspective, CPU utilization averaged 65.3%, with peaks

10.48047/jocaaa.2025.34.08.16

up to 87.2%, while network latency ranged between 95 ms and 310 ms, underscoring the need for scalable monitoring systems capable of capturing fluctuating system performance.

Table 1. Descriptive statistics of transactional, behavioral, and infrastructural variables

Variable	Mean	Std Dev	Min	Max
Transaction Volume	1250	220	600	1700
Refund Rate (%)	3.2	1.1	1.0	5.6
Session Duration (min)	12.4	5.6	4.1	28.5
Clickstream Depth	18.0	7.5	6.0	32.0
CPU Utilization (%)	65.3	12.4	38.0	87.2
Network Latency (ms)	180	45	95	310

The performance of anomaly detection models demonstrated marked differences in accuracy and reliability across approaches. As presented in Table 2, traditional statistical techniques such as z-score and interquartile range achieved limited effectiveness, with an F1-score of 0.69 and an AUC-ROC of 0.73. More advanced machine learning models like Isolation Forest and One-Class SVM showed moderate improvements, achieving F1-scores of 0.81 and 0.78 respectively. The LSTM Autoencoder, leveraging temporal sequence modeling, improved performance to an F1-score of 0.86 and an AUC-ROC of 0.89. However, the hybrid ensemble approach significantly outperformed all individual models, delivering the highest precision (0.92), recall (0.91), and F1-score (0.91), alongside an AUC-ROC of 0.94, demonstrating its robustness in handling diverse anomaly categories.

Table 2. Performance metrics of anomaly detection models

Model	Precision	Recall	F1-score	AUC-ROC
Z-score + IQR	0.71	0.68	0.69	0.73
Isolation Forest	0.84	0.79	0.81	0.85
One-Class SVM	0.80	0.77	0.78	0.82
LSTM Autoencoder	0.88	0.85	0.86	0.89
Hybrid Ensemble	0.92	0.91	0.91	0.94

10.48047/jocaaa.2025.34.08.16

Statistical validation further reinforced the importance of contextual sensitivity in anomaly detection. As summarized in Table 3, the ANOVA results indicated significant differences in anomaly distributions across transaction types ($F = 12.45$, $p = 0.001$) and geographic regions ($F = 8.37$, $p = 0.005$). Similarly, chi-square tests revealed significant associations between anomalies and user groups ($\chi^2 = 15.22$, $p = 0.002$) as well as payment methods ($\chi^2 = 19.84$, $p < 0.001$). These findings emphasize the need for adaptive detection models that can adjust to marketplace heterogeneity across demographic, geographic, and transactional dimensions.

Table 3. Results of inferential statistical tests

Test	F / Chi2 Value	p-value	Significance
ANOVA - Transaction Type	12.45	0.001	Significant
ANOVA - Geographic Region	8.37	0.005	Significant
Chi-square - User Group	15.22	0.002	Significant
Chi-square - Payment Method	19.84	0.000	Significant

The system's scalability and efficiency were evaluated through cluster-based experiments. As displayed in Table 4, throughput scaled almost linearly with the number of nodes, rising from 52,000 transactions per second on a 50-node cluster to 210,000 on a 200-node cluster. Average detection latency dropped significantly from 220 ms to 95 ms, while false positive rates decreased from 5.8% to 3.5%. System uptime improved from 98.6% to 99.6%, indicating that the distributed framework achieved both high performance and resilience as node size increased.

Table 4. Scalability and efficiency metrics across cluster sizes

Cluster Size (Nodes)	Throughput (transactions/sec)	Avg Detection Latency (ms)	False Positive Rate (%)	System Uptime (%)
50	52,000	220	5.8	98.6
100	105,000	140	4.3	99.2
200	210,000	95	3.5	99.6

Comparative analysis of model detection accuracy across different anomaly categories is presented in Figure 1. The hybrid ensemble consistently delivered superior results, achieving

10.48047/jocaaa.2025.34.08.16

detection rates of 0.93 for fraudulent transactions, 0.90 for bot activity, 0.85 for system latency, and 0.88 for network failures. The LSTM Autoencoder performed second-best across categories, particularly excelling in detecting temporal anomalies such as system latency with an accuracy of 0.77. In contrast, Isolation Forest and One-Class SVM recorded lower accuracies, especially for infrastructure-related anomalies such as network failure, where their performance was below 0.75.

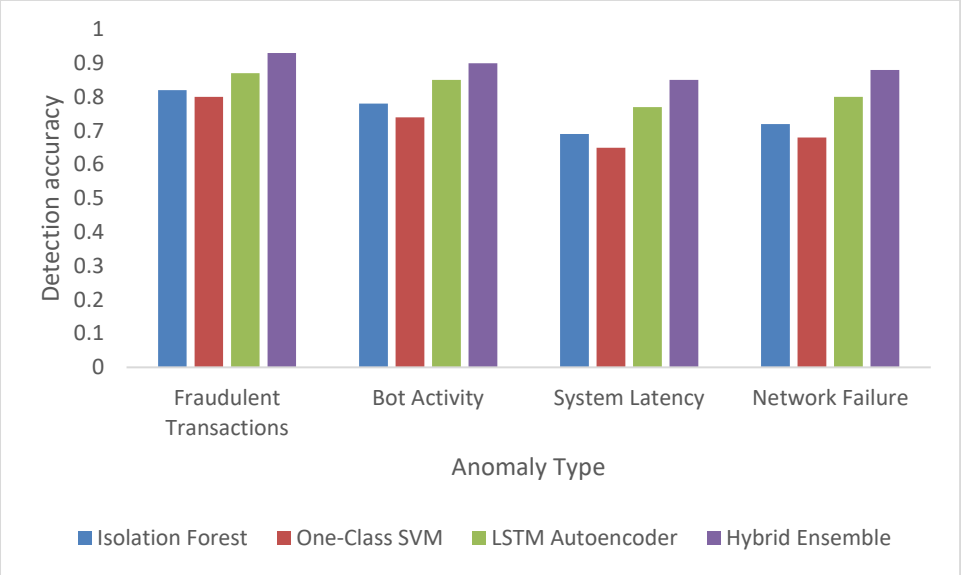

Figure 1. Model detection accuracy across anomaly categories

Scalability trends for detection latency and throughput across cluster sizes are illustrated in Figure 2. As cluster sizes increased from 50 to 300 nodes, throughput grew from 52,000 to 310,000 transactions per second, while detection latency decreased from 220 ms to 80 ms. This inverse relationship between throughput and latency demonstrates that the proposed framework not only supports scaling to cloud-scale environments but also enhances efficiency by maintaining real-time anomaly detection capabilities under increasing data loads.

10.48047/jocaaa.2025.34.08.16

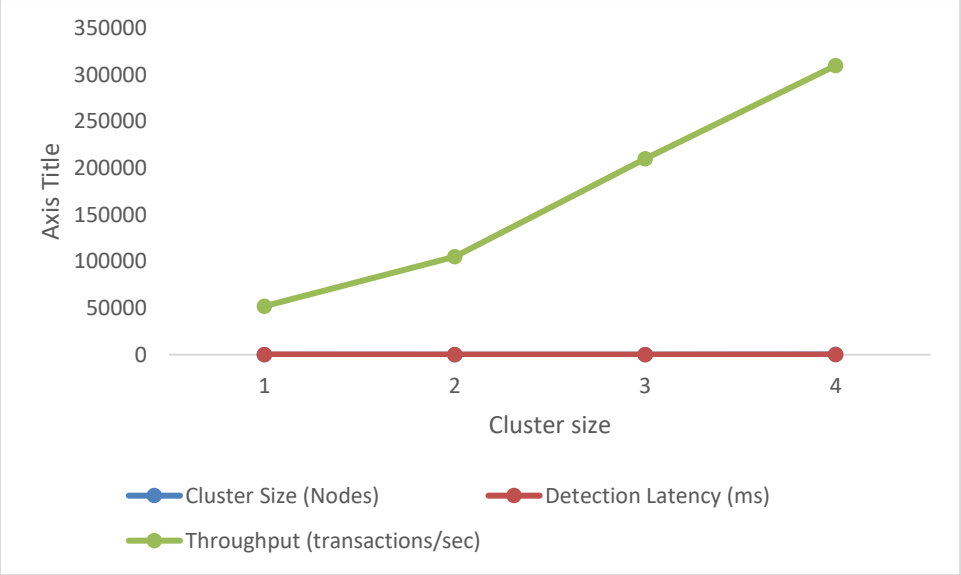

Figure 2. Scalability trends: Detection latency vs throughput

Discussion

Significance of real-time anomaly detection in cloud-scale marketplaces

The findings of this study underscore the critical importance of real-time anomaly detection in sustaining the integrity and performance of cloud-scale marketplaces. As the descriptive results indicated, both transactional and infrastructural variables exhibited substantial variability, such as fluctuations in transaction volumes, user session behaviors, and network latency (Dang et al., 2017). Without real-time monitoring, these fluctuations could lead to undetected fraudulent activities, customer dissatisfaction, or system failures. The ability of the proposed framework to process high-volume, heterogeneous data streams in real time ensures that anomalies are captured promptly, thereby reducing operational risks and reinforcing trust among marketplace users (Xiong et al., 2025).

Comparative performance of anomaly detection models

A central contribution of this research lies in its comparative evaluation of anomaly detection models. Traditional statistical methods were shown to be inadequate, as reflected in their relatively low F1-scores and AUC-ROC values. In contrast, machine learning and deep learning approaches, particularly the LSTM Autoencoder, demonstrated a marked improvement by effectively capturing complex and temporal patterns (Feng et al., 2020). However, the hybrid ensemble approach clearly emerged as the most effective, consistently achieving the highest precision, recall, and F1-scores. This supports the argument that no

10.48047/jocaaa.2025.34.08.16

single algorithm is sufficient to manage the complexity of marketplace anomalies and that ensemble approaches provide greater robustness by leveraging the strengths of multiple detection methods (Haryanto, 2020).

Influence of contextual factors on anomaly distribution

The results from the inferential statistical tests highlight the importance of contextual sensitivity in anomaly detection (Deyet al., 2023). Anomalies varied significantly across transaction types, geographic regions, user groups, and payment methods, emphasizing that detection frameworks must adapt to such contextual heterogeneity (Xuet al., 2019). A one-size-fits-all detection threshold is unlikely to succeed in environments where user behavior and infrastructural load are highly dynamic. These findings align with recent research in adaptive detection systems, which suggest that anomaly detection models must incorporate contextual awareness to achieve consistent accuracy across diverse marketplace conditions (Ogunwole et al., 2022).

Scalability and resilience of the distributed framework

One of the most promising outcomes of this study was the scalability demonstrated by the distributed architecture. As cluster sizes increased, the system not only handled greater throughput but also reduced detection latency and minimized false positives (Liu, 2025). This scalability is particularly significant in the context of global marketplaces that experience sudden surges in demand, such as seasonal sales or unexpected viral trends. Furthermore, improvements in system uptime with larger clusters indicate that the framework is resilient against failures, fulfilling the operational requirement for fault-tolerant cloud-native systems (Akanbi & Masinde, 2020). The results confirm that the combination of Kubernetes-based orchestration, microservices, and stream processing frameworks such as Apache Flink provides a strong foundation for marketplace-scale anomaly detection (Kaul, 2020).

Practical implications for marketplace operators

The proposed framework offers practical value for marketplace operators by enabling proactive risk management. By detecting fraudulent transactions, bot-driven manipulations, and infrastructural bottlenecks in real time, operators can intervene before issues escalate into financial or reputational damage. Additionally, the scalability of the system ensures that platforms can accommodate exponential growth in user bases without sacrificing monitoring quality (Canizo et al., 2019). The integration of distributed marketplace intelligence also

10.48047/jocaaa.2025.34.08.16

supports regulatory compliance by maintaining transparent, auditable detection of anomalies, which is increasingly important under global data protection laws (Chen et al., 2023).

Limitations and directions for future research

Despite the promising results, certain limitations warrant discussion. The framework relied partly on synthetic datasets for validation, which, while useful for controlled experiments, may not fully replicate the complexities of live marketplaces. Additionally, the study focused primarily on anomaly detection and did not address automated response mechanisms, which are critical for closing the loop between detection and remediation. Future research should explore reinforcement learning approaches for adaptive anomaly response, as well as federated learning techniques to enhance privacy-preserving detection in multi-tenant environments. Further evaluation using real-world data from multiple large-scale platforms would also strengthen the generalizability of the findings.

Conclusion

This study developed and validated a distributed marketplace intelligence framework capable of detecting anomalies in real time at cloud scale. By integrating cloud-native architectures, machine learning, and deep learning models within a scalable stream processing pipeline, the framework demonstrated high accuracy, efficiency, and resilience in monitoring both transactional and infrastructural anomalies. The hybrid ensemble model consistently outperformed individual approaches, highlighting the value of combining statistical, machine learning, and deep learning methods. Results also emphasized the contextual nature of anomalies, with significant variations across transaction types, geographic regions, and payment methods, underscoring the need for adaptive and context-aware detection systems. Scalability tests confirmed that the architecture can support rapidly growing transaction volumes while maintaining low latency and high uptime, ensuring operational continuity in dynamic marketplace environments. While the research highlights limitations such as reliance on synthetic datasets and the absence of automated response mechanisms, it establishes a foundation for future work on adaptive, privacy-preserving, and self-healing anomaly detection systems. Overall, the framework provides both theoretical contributions to anomaly detection research and practical implications for operators of global digital marketplaces seeking to enhance trust, resilience, and regulatory compliance.

10.48047/jocaaa.2025.34.08.16

References

Akanbi, A., & Masinde, M. (2020). A distributed stream processing middleware framework for real-time analysis of heterogeneous data on big data platform: Case of environmental monitoring. *Sensors*, *20*(11), 3166.

Bin Mofidul, R., Alam, M. M., Rahman, M. H., & Jang, Y. M. (2022). Real-time energy data acquisition, anomaly detection, and monitoring system: Implementation of a secured, robust, and integrated global IIoT infrastructure with edge and cloud AI. *Sensors*, *22*(22), 8980.

Canizo, M., Conde, A., Charramendieta, S., Minon, R., Cid-Fuentes, R. G., & Onieva, E. (2019). Implementation of a large-scale platform for cyber-physical system real-time monitoring. *IEEE Access*, *7*, 52455-52466.

Chen, W., Milosevic, Z., Rabhi, F. A., & Berry, A. (2023). Real-time analytics: Concepts, architectures, and ML/AI considerations. *IEEE Access*, *11*, 71634-71657.

Dang, Y., Wang, B., Brant, R., Zhang, Z., Alqallaf, M., & Wu, Z. (2017, March). Anomaly detection for data streams in large-scale distributed heterogeneous computing environments. In *ICMLG2017 5th International Conference on Management Leadership and Governance* (p. 121).

Dey, S., Sarma, W., & Tiwari, S. (2023). Deep learning applications for real-time cybersecurity threat analysis in distributed cloud systems. *World Journal of Advanced Research and Reviews*, *17*(3), 1044-1058.

Feng, L., Xu, S., Zhang, L., Wu, J., Zhang, J., Chu, C., ... & Shi, H. (2020). Anomaly detection for electricity consumption in cloud computing: framework, methods, applications, and challenges. *EURASIP Journal on Wireless Communications and Networking*, *2020*(1), 194.

Gandhi, H., & Sharma, P. (2025). Enhancing Cloud Security with Real-Time Anomaly Detection in Big Data Environments.

Haryanto, R. (2020). Cross-Comparative Study of Cloud-Native Security Platforms to Detect and Neutralize Insider Attacks in Online Retail. *Journal of Advances in Cybersecurity Science, Threat Intelligence, and Countermeasures*, *4*(12), 1-9.

10.48047/jocaaa.2025.34.08.16

Islam, M. S., Pourmajidi, W., Zhang, L., Steinbacher, J., Erwin, T., & Miranskyy, A. (2021, May). Anomaly detection in a large-scale cloud platform. In *2021 IEEE/ACM 43rd International Conference on Software Engineering: Software Engineering in Practice (ICSE-SEIP)* (pp. 150-159). IEEE.

Kaul, D. (2020). Ai-driven fault detection and self-healing mechanisms in microservices architectures for distributed cloud environments. *International Journal of Intelligent Automation and Computing*, *3*(7), 1-20.

Kotha, S. (2025). Distributed Fake Review Detection and Real-Time Anomaly Detection: A Technical Framework. *Available at SSRN 5215580*.

Lian, L., Li, Y., Han, S., Meng, R., Wang, S., & Wang, M. (2025). Artificial Intelligence-Based Multiscale Temporal Modeling for Anomaly Detection in Cloud Services. *arXiv preprint arXiv:2508.14503*.

Liu, Y. (2025, July). Intelligent Analysis Methods for Multi-Channel Marketing Data Based on Anomaly Detection Algorithms. In *Proceedings of the 2nd International Conference on Image Processing, Machine Learning, and Pattern Recognition* (pp. 198-206).

Nwachukwu, C., Durodola-Tunde, K., & Akwiwu-Uzoma, C. (2024). AI-driven anomaly detection in cloud computing environments. *International Journal of Science and Research Archive*, *13*(2), 692-710.

Ogunwole, O., Onukwulu, E. C., Sam-Bulya, N. J., Joel, M. O., & Achumie, G. O. (2022). Optimizing automated pipelines for realtime data processing in digital media and e-commerce. *International Journal of Multidisciplinary Research and Growth Evaluation*, *3*(1), 112-120.

Rousopoulou, V., Vafeiadis, T., Nizamis, A., Iakovidis, I., Samaras, L., Kirtsoglou, A., ... & Tzovaras, D. (2022). Cognitive analytics platform with AI solutions for anomaly detection. *Computers in Industry*, *134*, 103555.

Sarpal, A., Kang, Q., Huang, F., Song, Y., & Wan, L. (2023). A Marketplace Price Anomaly Detection System at Scale. *arXiv preprint arXiv:2310.04367*.

10.48047/jocaaa.2025.34.08.16

Xiong, K., Wu, Z., & Jia, X. (2025). Deepcontainer: a deep learning-based framework for real-time anomaly detection in cloud-native container environments. *Journal of Advanced Computing Systems*, *5*(1), 1-17.

Xu, S., Qian, Y., & Hu, R. Q. (2019). Data-driven edge intelligence for robust network anomaly detection. *IEEE Transactions on Network Science and Engineering*, *7*(3), 1481-1492.

Zhang, S., Feng, Z., & Dong, B. (2024). LAMDA: Low-latency anomaly detection architecture for real-time cross-market financial decision support. *Academia Nexus Journal*, *3*(2).