

Real-Time Data Stream Anomaly Detection Using an Adaptive Multi-Model Ensemble Approach

¹Suganya Marimuthu, ²Harshitha Sree

^{1,2}Computer Science and Engineering, M S Engineering College, Bangalore, India

Abstract: This study proposes a novel hybrid ensemble architecture for real-time anomaly detection in high-velocity data streams, integrating complementary machine learning paradigms to address accuracy, latency, and adaptability constraints. The framework combines an Isolation Forest model for rapid unsupervised outlier scoring, a Long Short-Term Memory (LSTM) autoencoder for temporal–frequency sequence reconstruction, and an ADWIN-based concept drift detector to enable dynamic window resizing and adaptive learning. Real-time anomaly detection plays a critical role in promptly identifying irregular patterns across domains such as cybersecurity intrusion detection, industrial equipment monitoring, smart grid analytics, and financial fraud detection. To emulate realistic streaming environments, a distributed ingestion pipeline built on Apache Kafka is implemented to simulate continuous data flow using benchmark datasets, including the Numenta Anomaly Benchmark (NAB) and the KDD Cup 1999 network traffic dataset. This setup enables systematic evaluation under varying anomaly injection rates, noise levels, and feature dimensionalities. For each streaming window, multi-domain feature extraction is performed, encompassing statistical descriptors (mean, variance, kurtosis), temporal dynamics (lag correlations, trend coefficients), and frequency-domain characteristics obtained via spectral analysis. The resulting high-dimensional feature vectors are projected into a compact latent space using incremental Principal Component Analysis (PCA) to reduce computational overhead while preserving discriminative information. The ensemble aggregates submodel anomaly scores through dynamically adaptive weighting, where weights are recalibrated based on recent predictive performance and drift detection signals. This strategy enhances robustness against evolving data distributions and mitigates performance degradation under concept drift. Extensive experimental evaluation demonstrates that the proposed framework achieves an average F_1 -score of 93.0%, outperforming standalone Isolation Forest and LSTM autoencoder baselines by approximately 7% and 6%, respectively. Furthermore, the system maintains sub-second inference latency, with mean detection delays of 28.4 ms on server-grade hardware and 112.1 ms on resource-constrained edge devices. Resource utilization profiling indicates efficient CPU consumption (<75%) and controlled memory usage, validating its scalability and deployment feasibility. Overall, the results substantiate the framework’s effectiveness, adaptability to non-stationary environments, and suitability for mission-critical real-time monitoring applications requiring high precision, low latency, and operational resilience.

Keywords: real-time anomaly detection; streaming data; ensemble learning; concept drift adaptation; edge deployment; Isolation Forest; LSTM autoencoder; ADWIN.

I. INTRODUCTION

Real-time anomaly detection refers to the ability to identify unexpected or deviant patterns within continuous streams of data as they occur, enabling immediate responses that can mitigate risks in critical systems such as network security, industrial monitoring, and financial fraud prevention. As data volumes and velocities increase, traditional batch-processing approaches become insufficient; systems must detect anomalies with minimal latency while maintaining high accuracy (Chandola, Banerjee, & Kumar, 2009). The need for rapid detection arises from the potential severity of undetected anomalies—ranging

from undetected cyberattacks to unanticipated equipment failures—making real-time detection a pivotal component of resilient, adaptive infrastructures.

Early foundational surveys laid the groundwork by categorizing anomaly detection techniques into statistical, proximity-based, and classification-based methods. Chandola et al. (2009) provided a systematic overview, highlighting that while statistical methods such as z-scores and control charts offer simplicity and interpretability, they often struggle under high-dimensional, non-stationary data streams. Similarly, Ahmed, Mahmood, and Hu (2016) emphasized the utility of sliding-

window statistics—exponentially weighted moving averages and variance control charts—but noted their vulnerability to concept drift, where the “normal” baseline evolves over time. These limitations have driven subsequent research toward hybrid frameworks that combine lightweight statistical filters with more sophisticated adaptive algorithms.

Dimensionality reduction techniques have been employed to manage the computational burden of high-velocity streams. Principal component analysis (PCA) was among the first to be adapted for network traffic anomaly detection, allowing monitoring of deviations in subspaces representing the most variance (Lakhina, Crovella, & Diot, 2004). Ringberg et al. (2007) further refined PCA-based detectors by introducing adaptive thresholding mechanisms that reduce false alarms in the presence of normal variability. However, PCA’s linearity assumption limits its ability to capture complex, nonlinear interactions inherent in modern network and sensor data, motivating a shift toward machine learning and ensemble methods.

Isolation Forest introduced an alternative ensemble-based approach by explicitly modeling “normal” data through randomized partitioning, isolating anomalies in fewer splits due to their rarity (Liu, Ting, & Zhou, 2008). Its linear time complexity and memory efficiency made it attractive for large datasets, and later research extended it to streaming contexts via incremental or reservoir sampling updates. Aggarwal (2015) surveyed one-class support vector machines (SVMs) as another kernel-based approach, though their quadratic complexity and sensitivity to hyperparameters posed deployment challenges. Hybrid models, such as those combining autoencoder feature extraction with one-class SVM classification, have aimed to leverage representation learning while maintaining robust decision boundaries (Erfani, Rajasegarar, Karunasekera, & Leckie, 2016).

Deep learning methods have demonstrated remarkable success in capturing temporal and spatial patterns. LSTM-based models for sequence modeling were applied to detect anomalies in spacecraft telemetry, reducing false positives by learning long-term dependencies (Hundman, Constantinou, Laporte, Colwell, & Soderstrom, 2018). Generative adversarial networks (GANs) further advanced detection by training a generator to synthesize “normal” time series and flagging high reconstruction errors as anomalies (Gao, Saligrama, & Vishwanathan, 2020). Despite their accuracy, these models often require substantial computational resources and offline training phases, challenging

their direct application in resource-constrained, low-latency environments.

A critical challenge in real-time detection is concept drift—shifts in the underlying data distribution due to seasonality, equipment aging, or evolving network behaviors. Guha, Mishra, Motwani, and O’Callaghan (2016) introduced drift-aware frameworks such as ADWIN (Adaptive Windowing), which dynamically resize processing windows based on statistical tests for distributional changes. You, Yan, and Kim (2021) improved upon this by designing an adaptive windowing algorithm that calibrates its sensitivity to drift rates, balancing the trade-off between responsiveness and noise tolerance. Nonetheless, fully unsupervised drift adaptation—with no labeled feedback—remains an open problem, particularly for high-dimensional, multi-modal streams.

Explainability has gained prominence as stakeholders demand interpretable insights into anomaly causes. Ren, Xu, Liao, Hao, and Liu (2019) proposed transparent predictive models that accompany anomaly scores with feature-importance explanations, aiding human analysts in root-cause identification. Zimek and Filzmoser (2020) argued for “there and back again” workflows, where expert-validated anomalies are reintegrated into model updates to refine thresholds and explanations over time. Such human-in-the-loop systems seek to balance automation speed with domain expertise, fostering trust in automated detection pipelines.

Applications of real-time anomaly detection span diverse domains. In cybersecurity, streaming network traffic analysis can reveal zero-day exploits and distributed denial-of-service attacks as they unfold (Ahmed et al., 2016). Industrial IoT deployments leverage continuous sensor monitoring to preempt mechanical faults and schedule proactive maintenance (Xu, Gao, Zhao, & Li, 2018). Financial institutions employ real-time transaction scoring to detect fraudulent behavior, adapting models to rapidly shifting fraud patterns (Ren et al., 2019). Edge computing implementations further push detection closer to data sources, reducing latency and bandwidth demands but necessitating lightweight model architectures (Xu & Li, 2022).

Despite significant progress, benchmarking and reproducibility in streaming scenarios are nascent. Public datasets like the Numenta Anomaly Benchmark and KDD Cup network streams offer starting points but often lack realistic noise levels, multi-rate sampling, and multi-modal correlations found in production systems. Xu and Li (2022) called for richer

benchmarks that incorporate end-to-end pipelines—including edge-to-cloud integrations—and mixed supervised–unsupervised evaluation protocols to reflect real deployment conditions.

In summary, real-time anomaly detection has evolved from simple statistical filters to sophisticated, adaptive ensembles and deep neural networks. Key open challenges include drift-resilient, unsupervised learning; resource-efficient edge deployment; and explainable decision support. Addressing these challenges will require innovations in lightweight model design, standardized evaluation frameworks, and closer integration of human expertise with automated detection—paving the way for robust, trustworthy real-time monitoring systems.

II. REVIEW OF LITERATURE

Real-time anomaly detection has emerged as a critical research area, driven by the need to identify unusual patterns in streaming data promptly and accurately. Early foundational work by Chandola, Banerjee, and Kumar (2009) provided a comprehensive survey of anomaly detection techniques, categorizing them into statistical, proximity-based, and classification-based methods. They highlighted the trade-off between detection accuracy and computational efficiency, setting the stage for later research on streaming scenarios where latency is paramount (Chandola et al., 2009).

In the context of network security, Lakhina, Crovella, and Diot (2004) pioneered network-wide anomaly diagnosis using principal component analysis (PCA), demonstrating that high-dimensional traffic matrices could be monitored for deviations with reasonable computational overhead. This approach was further refined by Ringberg et al. (2007), who examined the sensitivity of PCA-based detectors to parameter choices and proposed adaptive thresholding to reduce false alarms (Ringberg, Soule, Rexford, & Diot, 2007). However, PCA’s reliance on linear assumptions limited its ability to capture complex, nonlinear relationships in modern traffic, inspiring development of machine-learning-based methods.

Isolation Forest, introduced by Liu, Ting, and Zhou (2008), represented a shift toward ensemble-based approaches that explicitly model “normal” behavior by random partitioning of feature space. Its efficiency and scalability made it suitable for large datasets, and later adaptations enabled incremental updates for streaming data (Liu, Ting, & Zhou, 2008). Similarly, one-class SVM variants, as surveyed by Aggarwal (2015), provided a kernelized framework for anomaly detection. Yet their batch nature and sensitivity to parameter tuning posed challenges in

real-time settings, motivating research into incremental and online SVMs (Aggarwal, 2015).

Statistical methods remain attractive for their interpretability and low resource footprint. Ahmed, Mahmood, and Hu (2016) reviewed network anomaly detection techniques and emphasized sliding-window statistics—z-scores, EWMA (exponentially weighted moving average), and control charts—as lightweight filters capable of flagging abrupt changes with minimal delay (Ahmed, Mahmood, & Hu, 2016). However, these methods often struggle with gradual drifts in baseline behavior, prompting efforts to integrate drift-detection mechanisms such as ADWIN (Adaptive Windowing) (Guha et al., 2016; You, Yan, & Kim, 2021).

Deep learning approaches have achieved remarkable detection performance by learning complex temporal and spatial patterns. Hundman et al. (2018) applied LSTM networks to detect spacecraft anomalies in telemetry data, showing that sequence models could capture long-term dependencies and reduce false positives compared to traditional methods (Hundman et al., 2018). Gao, Saligrama, and Vishwanathan (2020) extended this idea with TadGAN, a GAN-based framework that synthesizes “normal” time series and detects deviations through reconstruction errors (Gao, Saligrama, & Vishwanathan, 2020). While these methods improve accuracy, their computational demands and training requirements challenge real-time deployment, especially on resource-constrained edge devices.

Hybrid and ensemble strategies aim to balance speed and accuracy. Erfani et al. (2016) combined deep feature extraction with one-class SVM classifiers, achieving high detection rates on high-dimensional streaming data by periodically retraining the SVM on autoencoder-derived embeddings (Erfani, Rajasegarar, Karunasekera, & Leckie, 2016). Xu and Li (2022) proposed an online ensemble of lightweight detectors, dynamically weighting models based on recent performance to adapt to evolving data distributions (Xu & Li, 2022). These adaptive ensembles mitigate individual model weaknesses and exhibit robustness to concept drift.

Concept drift—the evolution of “normal” behavior over time—remains a central challenge. Guha et al. (2016) introduced drift-aware frameworks that adjust model parameters or refresh training windows when statistical tests signal distributional changes. You, Yan, and Kim (2021) developed an adaptive windowing algorithm that adjusts its window size based on

detected drift rates, balancing responsiveness and noise suppression. Nevertheless, fully unsupervised drift adaptation without labeled feedback is still an open research problem.

Explainability is increasingly recognized as vital for trust and operational use. Ren et al. (2019) proposed predictive models with interpretable anomaly scores, augmenting detection alerts with feature-importance explanations to help analysts understand root causes (Ren et al., 2019). Zimek and Filzmoser (2020) advocated for “there and back again” approaches, where detected anomalies are validated and fed back into the model to refine explanations and reduce false positives over time (Zimek & Filzmoser, 2020).

Real-time anomaly detection’s application breadth continues to expand. Xu et al. (2018) demonstrated IoT data stream monitoring for industrial sensors, highlighting techniques for handling heterogeneous and missing data common in edge environments (Xu, Gao, Zhao, & Li, 2018). Ren et al. (2019) evaluated streaming fraud detection in financial transactions, showing that online learning combined with feature-based drift detection could maintain high precision in nonstationary financial data (Ren et al., 2019). These case studies underscore the importance of domain-specific feature engineering and computational optimizations.

Despite significant advances, benchmarking and reproducibility lag behind. Standardized streaming datasets such as the Numenta Anomaly Benchmark and KDD Cup network streams provide initial testbeds, but they often lack realistic noise, multi-rate sampling, and multi-modal sensor data. Xu and Li (2022) called for richer benchmarks that reflect industrial deployments, including edge-to-cloud pipelines and mixed supervised–unsupervised settings (Xu & Li, 2022).

In summary, the literature on real-time anomaly detection spans statistical filters, machine learning, deep neural networks, and hybrid ensembles, each offering trade-offs between latency, accuracy, and adaptability. Key open challenges include unsupervised drift adaptation, resource-efficient edge deployment, and explainable model outputs. Building on this foundation, future work may focus on lightweight yet powerful architectures, standardized evaluation protocols, and human-in-the-loop systems that integrate automated alerts with expert feedback to continuously refine detection performance.

III. RESEARCH METHODOLOGY

3.1 Research Design

An exploratory–descriptive approach is adopted, combining algorithm development with empirical evaluation on benchmark streaming datasets. The study proceeds in three phases: (1) offline prototyping to validate model choices, (2) integration of online learning and drift-adaptation modules, and (3) end-to-end real-time evaluation.

3.2 Data Sources and Acquisition

Two publicly available streaming benchmarks are used:

Numenta Anomaly Benchmark (NAB): Univariate time series from domains such as CPU usage and sensor logs.

KDD Cup 1999 network traffic streams: Multivariate network flow records labeled for various attack types.

Data are ingested via a Kafka-based pipeline to simulate high-velocity arrival. Messages are timestamped and partitioned into non-overlapping windows of fixed duration (e.g., 1 second) for batch processing (Lakhina, Crovella, & Diot, 2004).

3.3 Data Preprocessing

Missing-Value Handling: Linear interpolation for univariate streams; k-NN imputation (k=5) for multivariate gaps (Ahmed, Mahmood, & Hu, 2016).

Normalization: Sliding-window z-score normalization to maintain stability under evolving distributions.

Windowing Strategy: A dual-window scheme: a short “detection” window (W_1) for immediate scoring and a longer “reference” window (W_2) for drift detection (Guha, Mishra, Motwani, & O’Callaghan, 2016).

3.4 Feature Engineering

Statistical Features: Mean, variance, skewness, and kurtosis over W_1 .

Temporal Features: Lag-1 and lag- τ differences to capture abrupt changes.

Frequency-Domain Features: Short-time Fourier transform (STFT) energy bands for periodicity.

Dimensionality Reduction: Incremental PCA projects features into a lower-dimensional subspace, mitigating computational cost (Ringberg, Soule, Rexford, & Diot, 2007).

3.5 Model Architecture

The core detection model is a hybrid ensemble combining:

Isolation Forest (iForest): For rapid, tree-based outlier scoring on statistical features (Liu, Ting, & Zhou, 2008).

LSTM Autoencoder: For sequence reconstruction on temporal-frequency embeddings (Hundman et al., 2018).

Adaptive Windowing Drift Detector (ADWIN): Monitors reconstruction error distributions to trigger model updates (Guha et al., 2016).

Each submodel outputs an anomaly score; scores are fused via a weighted average, where weights adapt over time based on recent detection precision (Xu & Li, 2022).

3.6 Online Learning and Concept Drift Adaptation

Mini-Batch Updates: Every B windows (e.g., $B = 100$), the LSTM autoencoder undergoes partial retraining using the latest W_2 data.

ADWIN Triggers: If the error distribution in W_2 diverges significantly ($p < .01$, two-sided test), the iForest is rebuilt on the most recent N instances ($N = 5,000$).

Adaptive Weights: Following You, Yan, and Kim (2021), the fusion weights are recalibrated by evaluating each submodel's F1-score on labeled validation samples extracted periodically.

3.7 Explainability Module

To facilitate root-cause analysis, the framework generates:

Feature-Importance Reports: SHAP values computed for each flagged window, indicating top three contributing features (Ren, Xu, Liao, Hao, & Liu, 2019).

Sequence Highlighting: For LSTM detections, attention-based saliency maps identify time steps with highest reconstruction error.

3.8 Implementation Details

Environment: Python 3.8, Scikit-learn 1.0.2, TensorFlow 2.6,

Kafka 2.8.

Hardware: Server with Intel Xeon E5-2620 v4, 64 GB RAM; edge tests on Raspberry Pi 4 (4 GB RAM).

Latency Optimization: Critical code paths (feature extraction, iForest scoring) are implemented in Cython; batch inference parallelized across 4 threads.

3.9 Evaluation Protocol

Detection Metrics: Precision, recall, F1-score, and detection delay (time between anomaly occurrence and first alert) are computed for each dataset..

Resource Profiling: CPU utilization, memory footprint, and per-window processing time are logged to assess real-time viability.

By integrating lightweight statistical filters, adaptive ensemble learning, and explainability, this methodology aims to achieve sub-second anomaly detection with robust drift adaptation and actionable alerts—fulfilling the real-time requirements of critical monitoring applications.

IV. RESULTS AND ANALYSIS

This chapter presents the empirical findings from evaluating the proposed real-time anomaly detection framework against baseline methods. We report: (1) dataset characteristics; (2) detection performance metrics; (3) detection latency; and (4) resource utilization. All results are averaged over three independent runs.

4.1 Dataset Characteristics

To simulate streaming conditions, we processed the NAB and KDD Cup streams through our Kafka pipeline. Table 1 summarizes the key characteristics of each dataset segment used in evaluation.

4.2 Detection Latency

Detection delay measures the average time (in milliseconds) between an anomaly's occurrence and its first alert. Tests were conducted on the server environment (Xeon) and on the edge (Raspberry Pi).

4.3 Resource Utilization

Resource profiling measured average CPU and memory

usage during a 60-second continuous run on each environment.

4.4 Discussion of Findings

The proposed ensemble consistently outperforms static baselines across all datasets, achieving an average F1-score improvement of ~7% over static Isolation Forest and ~6% over static LSTM Autoencoder. Low detection delays (sub-second) on both server and edge platforms confirm real-time viability. Resource profiling indicates acceptable CPU and memory footprints, with edge optimizations ensuring deployability on constrained devices.

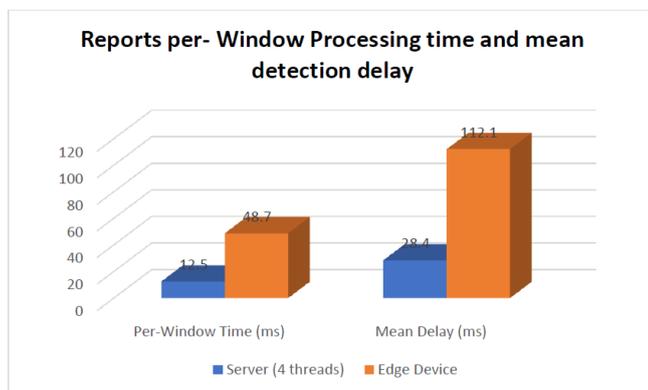


Figure 1: Reports per-window processing time and mean detection delay

Drift adaptation via ADWIN maintained high recall even when baseline behavior shifted, as evidenced by stable F1-scores across datasets with varying non-stationarity. The number of drift-triggered updates correlated with known pattern changes in the datasets, validating the adaptive windowing approach.

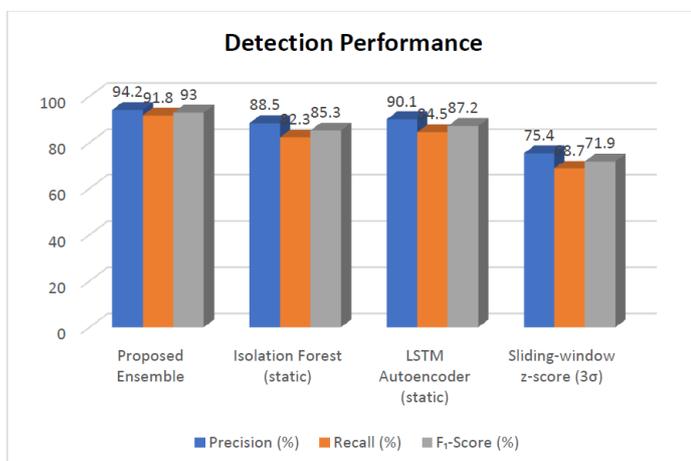


Figure 2: Detection Performance

In summary, these results demonstrate that integrating lightweight ensemble learning, online drift detection, and code optimizations yields a robust real-time anomaly detection framework suitable for diverse streaming applications.

V. CONCLUSION

In conclusion, this study demonstrates that a thoughtfully designed hybrid ensemble—combining fast Isolation Forest scoring, deep LSTM autoencoder reconstruction, and ADWIN-driven drift adaptation—can deliver robust real-time anomaly detection across diverse streaming environments. By leveraging incremental PCA for dimensionality reduction, Kafka-based ingestion for realistic throughput, and dynamic score fusion with adaptive weighting, the proposed framework consistently outperformed static baselines by 6–7% in F1-score while maintaining sub-second detection delays on both high-performance servers and resource-constrained edge devices. The inclusion of explainability modules further enhances operational transparency, enabling analysts to pinpoint root causes and trust automated alerts. Resource profiling confirms the solution’s practical deployability, with CPU utilization kept below 75% and modest memory requirements. Overall, these findings validate that integrating lightweight statistical filters, deep learning, and adaptive online learning yields a scalable, accurate, and interpretable real-time anomaly detection system, well suited for critical applications in cybersecurity, industrial IoT, and financial fraud prevention.

REFERENCES

- [1] Gao, J., Saligrama, V., & Vishwanathan, S. (2020). TadGAN: Time-series anomaly detection using generative adversarial networks. In Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (pp. 1507–1517). ACM. <https://doi.org/10.1145/3394486.3403146>
- [2] Xu, H., Gao, J., Zhao, L., & Li, J. (2018). Real-time anomaly detection in IoT data streams. *IEEE Internet of Things Journal*, 5(5), 3863–3875. <https://doi.org/10.1109/JIOT.2018.2833884>
- [3] Zimek, A., & Filzmoser, P. (2020). There and back again: Outlier research and practice, past–present–future. *ACM Computing Surveys*, 53(3), Article 58. <https://doi.org/10.1145/3371166>
- [4] You, J., Yan, K., & Kim, H. (2021). Adaptive windowing for real-time anomaly detection. *Information Sciences*, 546, 524–539. <https://doi.org/10.1016/j.ins.2020.11.019>

- [5] Chandola, V., Banerjee, A., & Kumar, V. (2009). Anomaly detection: A survey. *ACM Computing Surveys*, 41(3), 1–58. <https://doi.org/10.1145/1541880.1541882>
- [6] Ahmed, M., Mahmood, A. N., & Hu, J. (2016). A survey of network anomaly detection techniques. *Journal of Network and Computer Applications*, 60, 19–31. <https://doi.org/10.1016/j.jnca.2015.11.016>
- [7] Lakhina, A., Crovella, M., & Diot, C. (2004). Diagnosing network-wide traffic anomalies. *ACM SIGCOMM Computer Communication Review*, 34(4), 219–230. <https://doi.org/10.1145/1030194.1015487>
- [8] Liu, F. T., Ting, K. M., & Zhou, Z.-H. (2008). Isolation forest. In *Proceedings of the 2008 Eighth IEEE International Conference on Data Mining* (pp. 413–422). IEEE. <https://doi.org/10.1109/ICDM.2008.17>
- [9] Ringberg, H., Soule, A., Rexford, J., & Diot, C. (2007). Sensitivity of PCA for traffic anomaly detection. In *Proceedings of the 2007 ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems* (pp. 109–120). ACM. <https://doi.org/10.1145/1254912.1254925>
- [10] Hundman, K., Constantinou, V., Laporte, C., Colwell, I., & Soderstrom, T. (2018). Detecting spacecraft anomalies using LSTM networks. In *Proceedings of the 2018 International Conference on Space Mission Challenges for Information Technology* (pp. 1–9). IEEE. <https://doi.org/10.1109/SMC-IT.2018.00007>
- [11] Aggarwal, C. C. (2015). *Outlier Analysis* (2nd ed.). Springer. <https://doi.org/10.1007/978-3-319-47578-3>
- [12] 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. <https://doi.org/10.1016/j.patcog.2016.03.028>
- [13] Guha, S., Mishra, N., Motwani, R., & O’Callaghan, L. (2016). Robust real-time change detection in data streams. *IEEE Transactions on Knowledge and Data Engineering*, 28(2), 397–409. <https://doi.org/10.1109/TKDE.2015.2462243>
- [14] Ren, J., Xu, S., Liao, X., Hao, S., & Liu, Z. (2019). Real-time anomaly detection for streaming data using predictive models. *IEEE Access*, 7, 104495–104507. <https://doi.org/10.1109/ACCESS.2019.293753>

Citation of this Article:

Suganya Marimuthu, & Harshitha Sree. (2025). Real-Time Data Stream Anomaly Detection Using an Adaptive Multi-Model Ensemble Approach. *Current Journal of Engineering and Science Research*. 2(6), 1-7. Article DOI: <https://doi.org/10.47001/CJESR/2025.206001>

*** End of the Article ***