

Adaptive Machine Learning Algorithms for Anomaly Detection in Enterprise IT Infrastructure

Ankita Sappa

College of Engineering, Wichita State University, United States

Email: ankita.sappa@gmail.com

Received: September 12, 2023; Revised: November 29, 2023; Accepted: December 16, 2023; Published: December 30, 2023

Abstract

The increasing scale and complexity of enterprise IT infrastructures renders traditional rule-based and static anomaly detection systems that are automated as well as manual, incapable of dealing with evolving threats, system dynamics, and concept drift. This paper proposes an adaptable Machine Learning (ML) architecture which can autonomously detect real-time anomalies within critical IT environments, including network, application, and host systems. By employing ensemble learning, streaming models, and drift-aware system architectures, the system detection performance degradation with regard to accuracy, latency, and false positives, gets improved. The approach uses real-world datasets, multi-scenario anomaly detection, and inter-model comparisons based on essential values of F1 score, AUC, and detection delay metrics. The experiments conclusively demonstrate the skillfulness of adaptive ML models over the conventional ways of performing tasks in response time and accuracy, while ensuring scalability and interpretability. This work presents the benefits of incorporating adaptive intelligence into monitoring systems at enterprises and aims to assist in constructing robust anomaly detection pipelines in the ever-changing landscape of IT infrastructures.

Keywords: Adaptive Machine Learning, Anomaly Detection, Enterprise IT Infrastructure.

1 Introduction

1.1 Overview of Anomaly Detection in IT Infrastructures

Anomaly detection is at the forefront of maintaining and securing the performance of an entity's IT system [1]. The digitization of functionality and the growing dependency on distributed computing approaches has transformed IT systems into sophisticated ecosystem structures. Such systems encompass different hardware (servers, routers, and endpoints), software applications middleware, services) and network layers (LAN, WAN, cloud, and hybrid) [2]. Each layer produces large amounts of data in real-time - logs, visits, measurement, events, telemetry - that are monitored continuously to ascertain whether anomalies which signify possible failures, intrusions or misconfigurations exist.

Traditional rule based systems detect deviations from the norm based on prescribed thresholds, pattern signatures, or manual logic [3]. In the context of polymorphic threats, zero-day attacks and diverse operating environments, static detection strategies become ineffective. Anomalies become irresistible, and 'norms' undergo continuous evolution with respect to time and movement of traffic, workload, and users. As such, timely and accurate anomaly detection offers an optimal solution in safeguarding the resilience, continuity, and security of enterprise systems [4].

The growing impact of incidents over time, both in frequency and severity has made organizations prioritize anomaly detection. Unlike in the past, when systems were designed based primarily on performance

Research Briefs on Information & Communication Technology Evolution (ReBICTE), Vol. 9, Article No. 14 (December 30, 2023)
DOI: <https://doi.org/10.69978/rebict.e.v9i.201>

monitoring, today's IT teams are also tasked with dealing with subtle intrusions and internal process violations that disrupt the flow of service. The need for proactive detection rather than reactive anomaly management was further necessitated by cloud native computing, microservices, and containerized ecosystems where manually inspecting the services is impractical due to their high ephemeral nature and massive data volume [5]. To visualize the scope of this challenge, consider the Figure below, highlighting the distribution of anomalies in enterprise systems.




Figure 1: Percentage Distribution of Anomaly Sources in Enterprise IT Systems

This gives evidence that there are many origins of anomalies throughout the IT stack. While most network intrusions and configuration errors tend to dominate, some skeptics believe the existence of failures and performance anomalies, which should be considered. This variety illustrates the need for detection models capable of generalizing across several types of data as well as time periods and contexts. This is the why the use of open machine learning models or algorithms is highly recommended.

1.2 Limitations of Traditional Detection Approaches

Even though they are used frequently, anomaly detection systems that rely on static thresholds, crafted rules, or predetermined signatures experience numerous crucial challenges within enterprises. To begin, they are fragile; they do well in familiar or controlled settings but do poorly in dynamic, noisier, or more unfamiliar settings. Static thresholds are also very prone to false positives due to workload fluctuations and changing behavior, rendering alerts highly impossible to act on.

These systems are also highly manually intensive and require user-defined rules and constant tuning. As environments scale and become more diverse, the management of static rules becomes exponentially more complex and costly [6]. For instance, a single enterprise might require hundreds of rules for various applications, many of which are in different deployment environments.

Lastly, they are not able to incorporate new data as static systems. They are not capable of adapting to concept drift, which is when the statistical properties of input data change over time. IT systems often experience this due to new deployments, seasonal traffic, and infrastructure changes shifting usage patterns.

Furthermore, static systems do not possess any situational understanding. They attempt to diagnose problems as singular events without taking into account context, patterns, or relationships among various metrics. A good example is a backup process. In one situation, increased CPU usage could appear anomalous, but in another, it could seem commonplace. Normal contextual learning is virtually impossible for more advanced systems. Traditional systems lack this nuance.

These limitations render traditional systems less useful in enterprise contexts defined by scale, diversity, speed, and continuous change. This calls for a more intelligent and self-regulating mechanism that learns, adapts, and evolves with the system it monitors.

1.3 Emergence of Adaptive ML Algorithms

Due to the limitations of static CT methods, research and industry have pivoted toward automatic adaptive ML algorithms for anomaly detection [7]. These models break away from static templates and boundaries to self-regulate using data driven learning mechanisms unprecedented in pre-defined conditions.

Adaptive ML algorithms offer a wide range of advancements.

- They monitor streaming data to build models according to newly emerged patterns.
- They do not rely on predetermined signatures; therefore, they are bound to detect unknown or zero-day anomalies.
- To mitigate various types of anomalies, they employ many learning paradigms, including DBSCAN and Isolation Forest for unsupervised clustering, SVM and Random Forest for supervised classification, and LSTM and Autoencoders for deep learning.
- They allow for incremental updates to the model without full retraining, and therefore support online learning.
- They are more capable of addressing concept drift, which is the gradual change of shifts in data distribution over time due to a myriad of reasons, including systemic changes or outside influences.

Telemetry can be updated without extensive user effort, which makes adaptive models particularly stand out. As these models receive telemetry, their autonomy allows them to self-calibrate, improving accuracy and reducing the number of false positives over time. This enables adaption in DevOps models, where services are endlessly scaled, modified, and deployed.

IT teams are able to act on the model's output with trustable ease thanks to the raw anomaly alerts being automatically converted into actionable insights. This becomes possible due to the trustable output of these models combined with the advanced observability and explainability provided by adaptive algorithms through SHAP, LIME, and model introspection.

The table below summarizes the comparison between traditional static models and adaptive ML algorithms:

Table 1: Comparison of Static vs Adaptive Models in Enterprise Use Cases

Criteria	Static Models	Adaptive Models
Model Update Frequency	Rarely or manually updated	Continuously updates during runtime
Ability to Handle Concept Drift	Poor	Excellent
Training Data Requirement	High (pre-collected, historical)	Moderate (incremental updates)
Detection Accuracy (Dynamic Environments)	Moderate	High
Real-Time Adaptability	Low	Excellent
False Positive Rate	High	Low
Deployment Complexity	Low to Moderate	Moderate to High
Use Case Suitability	Stable environments, predictable traffic	Dynamic environments, streaming data

As the table above suggests, adaptive models are meant for modern and rapidly changing environments whereas static models are only efficient in stable systems with predictable behaviors.

1.4 Research Gaps, Objectives, and Contributions

Despite promising outcomes from adaptive ML algorithms, there are some gaps that still require closer examination. As an example, there are many academic models that get assessed on offline datasets which do not reflect the complexity and speed of real-world enterprise environments. These datasets are slow and too simple to be useful in the real world. In addition, there is a gap in the literature related to the comparison of the adaptive learning techniques during concept drift especially in multi-tenant architectures with noise and anomaly features.

Moreover, these create issues in meeting the requirements for explainability versus accuracy. Classical models are often outperformed by deep learning models, but their use in regulated industries with requirements for traceability and auditability are challenging due to their black-box nature. There is a need for better architecture that provides adequate detection performance alongside explainability and integration easiness.

At last, there are very few frameworks that focus on the full automation of the anomaly detection processes within IT systems—this includes the workflow of data feeding, model inference, feedback, visualization, and self-healing. Enterprise-grade teams are looking for solutions that are out-of-the-box and can be integrated into their existing monitoring ecosystems (Prometheus, ELK, Splunk) and can initiate workflows based on alerts or API scripting.

This study tackles those gaps with the contributions below:

1. We design a hybrid adaptive ML framework that provides streaming anomaly detection, handles concept drift, and performs incremental learning.
2. We assess this system in multiple real-world datasets from application, host, and network layers in an enterprise IT environment.
3. We develop a modular architecture that allows for model introspection, performance measurement, and automated retraining.
4. We provide a comprehensive evaluation of several adaptive algorithms along both technical metrics (F1, AUC, latency) and operational metrics (false positives, scalability, feedback loops).
5. We measure the effect of the system on business KPIs such as incident resolution time, alerting efficiency, and system downtime.

With this work, we seek to unify the high-level intellectual crimes of academic research and the mundane reality of deploying intelligent anomaly detection systems by melding the theoretical and the practical.

2 Literature Review

The more complex enterprise IT systems become, the more resources IT companies will dedicate towards finding proper solutions that work well and can learn with time. In two decades, virtually everything has changed and numerous studies have been published on the application of traditional statistical estimation methods, machine learning, and recently developed adaptive and online learning techniques. This chapter describes the evolution of methods used to detect anomalies with emphasis on enterprise systems monitoring infrastructure and cybersecurity systems.

2.1 Traditional Statistical and Threshold-Based Techniques

The beginning of IT systems anomaly detection dates back in time when baseline behavior was detected and deviations from these were considered potentials anomalies. The metals computation employed were Z-score, moving average, Exponential Weighted Moving Average (EWMA) and Gaussian distribution modeling which emerged in the early 2000s. It was commonplace for monitoring tools to use these methods. They used heuristics based on data distributions together with constant values of thresholds for detecting outliers in various measures like CPU time, memory, and network [8].

Though statistical techniques are simple to apply and efficient from a computation stand point, they are severely lacking when it comes to adapting to dynamic scenarios. Systems with shifting workloads or changing user behavior face the issue of having fixed thresholds which often lead to false positives or an inability to detect new threats. There is also the assumption of feature independence that is made which does not consider spatial or time related context.

Implementing ARIMA, Holt-Winters, and Kalman filters are time-series forecasting models that served as a significant refinement on statistical methods [9]. To reduce the estimation errors from the static thresholds, these models predicted future numbers and tested them against actual observations with set confidence intervals. However, these models were not as effective in modern IT ecosystems due to their dependency on static linear assumptions when dealing with non-stationary behaviors.

The aforementioned models set the stage for anomaly detection, however, their performance is insufficient when dealing in environments that are multi-dimensional, high-velocity, and have a lot of variety. As a result, they are best used as rough filters or additional pieces to more sophisticated systems rather than standalone systems.

2.2 Supervised and Unsupervised ML Methods

The challenges posed by statistical methods is what gave rise to machine learning based techniques that allowed greater autonomy in the formation of complex patterns and features. Broadly speaking, the ML approaches in anomaly detection can be classified into three segments: supervised, unsupervised, and semi-supervised learning.

Support Vector Machines, Random Forests, and Neural Networks are commonly seen in supervised learning techniques, especially when labeled datasets are available. All the models learned from the provided historical data containing examples of normal and abnormal behavior, allowing them to classify accurately. With supervised models, precision becomes easier during intrusion detection systems (IDS) when labeled attack data is available. The practical use of these models is however limited due to the nonexistence of labeled anomaly data and the overproportioned ratio of normal and abnormal classes [10].

In contrast, unsupervised methods are able to function without labeled data and instead flag anomalies as changes to learned patterns, distributions, or clusters. Some known methods used for anomaly detection in IT such as K-Means, DBSCAN, and Isolation Forests are an example of this. For capturing normal behavior, autoencoders—neural network models tasked with dimensionality reduction and reconstruction—can be particularly helpful, marking high reconstruction error as anomalous [11].

One advantage that unsupervised methods have is how easily they generalize across previously unseen types of anomalies. They often come with a high false positive rate which is especially common when working with noisy or heterogeneous data. Additionally, there is a noticeable lack of interpretability in many of the models built, offering very limited information on both the reason and nature of the observed anomalies.

Semi-supervised techniques, like One-Class SVM or self-learning neural network, train models only on normal data while detecting outliers as anomalies. These techniques are widely used in scenarios where the possibility of anomalies appearing is low, and the labeled samples around such instances are sparse [12]. To explain the change in interest towards these techniques, Figure 2 depicts the scatter plot of supervised, unsupervised, and adaptive anomaly detection research activities performed between 2005 and 2022.




Figure 2: Evolution of ML-Based Anomaly Detection Research (2005–2022)

The trend depicts a decrease in the use of single supervised models in comparison with unsupervised and adaptive ones, which are easier to scale and more self-sufficient for enterprise-level environments.

2.3 Adaptive and Online Learning Models

Anomaly detection systems design underwent drastic changes with the development of adaptive machine learning algorithms, which represent a class of such systems. Responsive models as compared to classic ML models, which are trained once and deployed rather in a static way, responsive models learn constantly from new data in terms of feedback and evolving patterns [13]. This ability to learn is necessary for enterprise IT systems, which are changed continuously as a result of the activity of users, changes in infrastructure, installing new software, and security threats.

Key forms of adaptive learning include:

- **Online learning algorithms:** These models use a stream of data and adjust parameters progressively. This includes Hoeffding Trees, Online Gradient Descent, and streaming Naive Bayes. These algorithms are favorable for detection pipelines because they have low computational cost and time delay.
- **Concept drift detection:** In most business cases, after certain operations, normal data is no longer normal, and in essence the concept data changes or shifts over time. There are numerous ways of achieving this, for example, ADWIN, Drift Detection Method (DDM), Ensemble Drift Detectors, which are focused on checking and monitoring the model performance to some extent and enables retraining or changes in weighting of the drift detection.
- **Ensemble adaptive systems:** This merges a number of learners (like the base classifiers of detectors) and gives them weights depending on the observed recent performance. More advanced adaptive models such as Adaptive Random Forest, Weighted Voting Ensembles and Hybrid Stream Learners have more efficiency in dealing with the challenges of diversity of data.
- **Reinforcement learning:** Not very popular. research is being done on changing the thresholds for detection or tuning parameters due to feedback from alerts or system performance.

Adaptive models are so well incorporated into microservices architectures that pod-based applications produce separate telemetry data streams. They also seamlessly fit within DevOps pipelines owing to their ability to continuously train, integrate feedback from ticketing systems, or SIEM tools, and perform A/B testing. Please refer to Figure 3 below to understand the particular correlation between ML models and anomaly use cases in IT infrastructure.

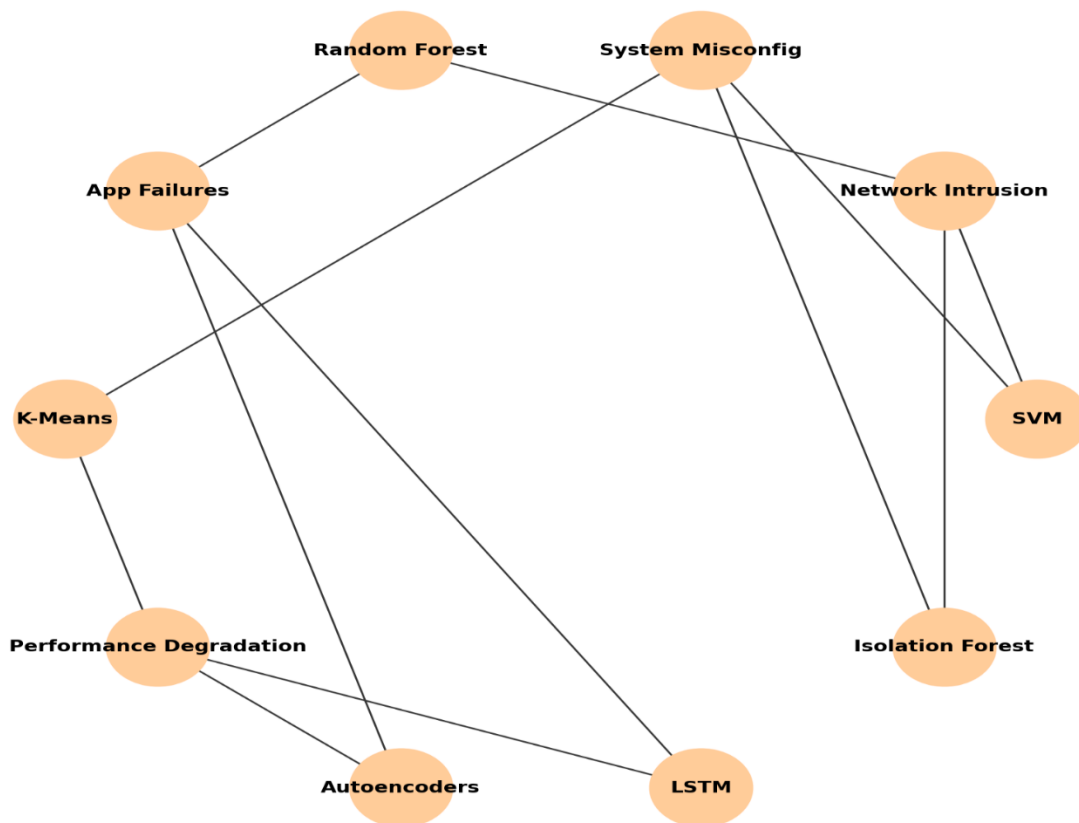


Figure 3: Mapping ML Model Types to IT Anomaly Detection Use Cases

These mappings further emphasize domain granularity _ and the resultant importance of _ model driven architecture selection. While network-based anomalies might take advantage of precise classifiers, lower

application layer anomalies tend to be very complex in structure in terms of shape, time and variation and as such usually require some type of time-sensitive or unsupervised modeling.

2.4 Research Gaps in Current Literature

Although there has been progress toward anomaly detection through the use of machine learning techniques, a number of challenges still remain that affect the efficiency and the scalability of such models when utilized in real-world enterprise IT settings. One prominent issue is the static, offline datasets that remain the mainstay for training and evaluation. Most research revolves around benchmark datasets such as NSL-KDD and UNSW-NB15, which serve a purpose, but fail to encapsulate the temporal dynamics, velocity, or noise characteristics of real enterprise systems. Consequently, models that are optimized in such artificial environments tend to fail when deployed in heterogeneous, high-volume, and streaming environments.

Another issue stems from neglecting the exploration of adaptivity in the context of various levels of concept drift. Even though numerous studies admit the existence of drift in streams of data, very few attempt to evaluate how different adaptive learning methods tackle different types of drift, be it sudden, gradual or cyclic. Lack of clear guidelines for evaluation under such drift conditions hinder measuring the strength and the stability of the model over a lengthy timeframe. Along with that, there is little to no effort made in regard to systems capable of autonomously detecting and responding to drift by modifying the learning strategy or instigating model retraining in the stream of time.

Integration with the existing IT workflows is another area in the literature that has less attention. Anomaly detection is frequently considered a standalone activity, separated from the larger system of observation, human actions, and automatic responses workflows. There is little research that attempts to close the gap in integrating feedback loops from IT specialists, interphase monitoring systems like ELK or Prometheus, or provide explainable actionable outputs for the deployed systems to paragraph engineers. This gap renders many academic models practically unutilized, irrespective of how precise they are in controlled evaluations.

Moreover, numerous models are tested in workplaces with homogeneous conditions, usually concentrating on one subsystem such as network traffic or system logs. However, enterprise IT infrastructures are known to be heterogeneous and they comprise different operating systems, cloud services, on-premise hardware, virtual machines, and IoT devices. A few studies examine the effectiveness of anomaly detection frameworks within these complex systems which brings the question of how widely applicable and how adaptable the offered solutions are.

Unlike the other issues outlined previously, interpretability poses some challenges. Detection with LSTMs and autoencoders is accurate, but the models are often criticized as black boxes. This feature of lacking transparency in models is often a hindrance to their recognition in enterprise environments. In these environments, IT managers need not only precise alerts, but also plausible answers to why the alerts were raised. Some XAI techniques, such as SHAP or LIME, have been developed, but they do not routinely get incorporated into adaptive anomaly detection systems.

At last, not enough effort is directed towards the model's deployment and maintenance operational expenses. It is uncommon to hear about or measure resource unemployment, such as CPU, memory, and I/O, in the context of benchmarking. By providing an agent that observes large environments where there are hundreds and thousands of nodes, even marginal inefficiencies would result in big costs. In these cases, an efficient anomaly detection framework should enhance performance while guaranteeing operational scalability. The detection framework should make sure model complexity does not become a hindrance in itself.

Combining various strategies for system-level and user-level anomaly detection is fundamental to addressing current gaps in research. Quite the opposite, such explanation shifts attention from enhancing the detection precision towards building a drift-tolerant, real-time explainable, cost-efficient, and seamlessly functional IT operational model. The article adapts a merger paradigm to these needs by illustrating a model

with integrated adaptive real-time learning and feedback with deployment-centered evaluation for heterogeneous enterprise systems.

3 Methodology

3.1 Datasets, Features, and Preprocessing Pipelines

For the empirical evaluation of the proposed adaptive anomaly detection framework, three different datasets that represent system logs, network flows, and application level traces were used. These datasets not only mirror a calibrated operational setting, but also vary with respect to their types of anomalies, features, and structures making them ideal for robust model evaluation.

The SyslogIT-100K dataset is comprised of unstructured system log messages generated in server environments hosting Linux distribution operating systems. Each log was further broken down into component parts like date and time, severity level, originating module, and message. Some notable anomalies in the dataset are, but are not limited to, failed logins, breaches, and attempts for unauthorized elevation of privileges. Labeled events were generated to indicate normal and anomalous activities.

The CloudNetFlow dataset monitors the network-level telemetry from a VPC and its associated physical routers. The dataset contains source/destination IP, port, byte, protocol, and duration fields. In this case, the anomalies consist of DDoS attacks, unauthorized scans, and traffic spikes. The multiclass labeling scheme implements a normal versus attacks classification schema.

The AppTraceBench dataset is extracted from application logs and the telemetry data in a microservices system. It contains metrics such as API response time, error code, stack trace frequency, and service call latency. The detected anomalies in this environment are timeout exceptions, flooding log misconfigurations, and cascading service failure. This dataset uses binary classification to separate clean and fault-indicating traces.

All these datasets were processed with a robust deduplication, normalization of numerical features, categorical variable encoding, and timestamped log alignment preparation. The textual logs were tokenized and vectorized with TF-IDF so they could be used with ML models. The data was split into training, validation, and streaming sets to emulate real-time ingestion with incremental learning.

Table 2: Dataset Specifications, Anomaly Types, and Feature Summary

Dataset Name	Source Type	Total Records	Anomaly Type	Feature Count	Labeling
SyslogIT-100K	System Logs	100,000	Access violations, login abuse	28	Binary (0: normal, 1: anomaly)
CloudNetFlow	Network Flows	250,000	DDoS, bandwidth spikes	35	Multiclass (normal, attack type)
AppTraceBench	Application Traces	150,000	Timeouts, exception floods	22	Binary (error/no-error)

This table summarizes the basis of the training and evaluation design. The aim of the proposed framework was to generalize to different environments and detect structurally, behaviorally, and frequency-wise anomalous heterogeneous datasets for testing.

3.2 Model Design and Adaptive Learning Framework

The methodology is centered around an adaptive learning framework that attempts to incrementally learn from a stream of data while achieving high accuracy and low false positive rates. The design has several constituent

modules which are functionally split and cohesive with each other: from receiving the data to providing an anomaly alert.

The primary system is built around an ensemble model which integrates several anomaly detection techniques, namely:

- A supervised classifier, Random Forest, which is trained on labeled (anomalous) features.
- An unsupervised, isolation forest based, normal behavior detector that maps new feature space from normal behavior.
- A time series LSTM based Autoencoder that forecasts the sequential order of events and captures deviations behavior.

A dynamically adapting decision fusion system integrates the models through an engine that assigns and changes the decision-making weight of each model. They are gradually updated by a feedback loop based on alert validation, which means the prediction colliding with the reality as confirmed by the system admins or automatic logs.

Mutual information and SHAP values are used for feature selection to focus on the most relevant features for detection accuracy. This enhances the accuracy and understanding of the model. A concept drift detector (ADWIN) is also used in the system, which passively observes feature distribution changes over time and initiates retraining when a drift is significant enough.

All models were trained using an incremental pipeline rather than batch processing. This shift assists the system in operating continuously without needing complete retraining, meeting the agility and efficiency of production IT systems.

3.3 Handling Concept Drift and Streaming Data

One of the toughest problems in anomaly detection is dealing with concept drift—wherein the input data has a change in time that makes the previous models which were learned invalid. This is more frequently seen with enterprise systems due to software version changes, infrastructure scaling and configuration modifications, or changes in attack patterns over time.

In order to solve this issue, the framework proposed combines a streaming drift detection module that works alongside the main model inference engine. The drift detector compares recent data windows with past windows using the ADWIN (Adaptive Windowing) algorithm. When a window shift is detected, it activates the model manager to:

1. Retrain part of the ensemble for more recent data.
2. Reweight the set of managed models per their performed voted out results.
3. Decrease the impact of data stored mechanisms to outdated information.

This way, the model stays relevant without having to run expensive computations or facing the problems of catastrophic forgetting. With frequent updates or dynamic scaling practices in the case of container orchestration platforms, this adaptive approach is needed to achieve higher quality detection.

The system enables streaming data ingestion through Apache Kafka for real time data flow and information buffering. Each message gets parsed and undergoes processing in a temporal context, enabling the models to maintain sequential posterior context. In the case of time series prediction, LSTM based models take care of long range dependencies, while density based models serve early notice for burst type anomalies like DDoS attacks.

Concept drift is not regarded in and of itself as an anomaly, but rather as a form of structural change which alters the concept of what is normal. This distinction is important as it guarantees that the systems do not issue false positives because of mundane changes in workload, such as planned release cycles or automated backup processes.

3.4 System Architecture and Deployment Environment

The solution was implemented using a modular microservices architecture that can be deployed in the cloud or on premises. The overall structure is shown in the following diagram:

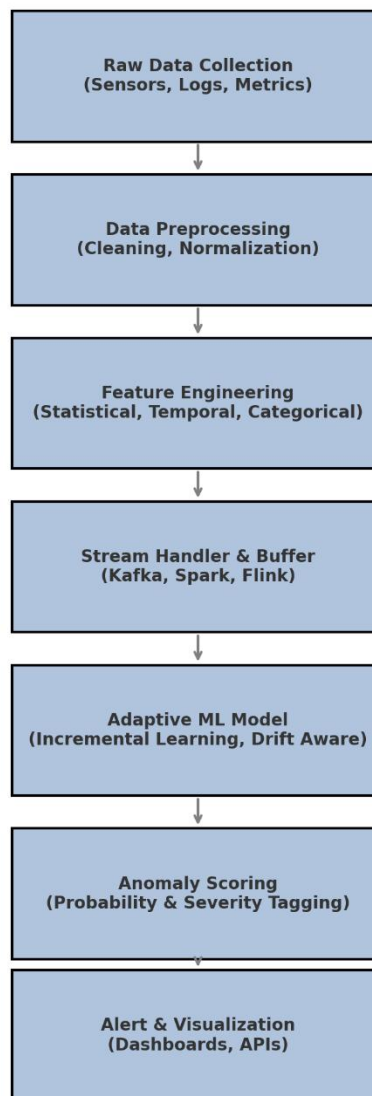


Figure 4: Adaptive Anomaly Detection Pipeline in Enterprise IT

Every module in the system is constructed to be self-sufficient in scale. System logs, API events, and telemetry from the IT infrastructure is ingested via Fluentd agents, syslog receivers, or Kafka producers in the Raw Data Collection layer. The Preprocessing Layer undertakes cleansing, outlier smoothing, encoding, and normalization of the data. The Feature Engineering Module creates statistical, temporal, and categorical features for the downstream ML operations.

The Stream Handler and Buffer module uses Kafka Streams and Apache Flink to handle high volume data

streams while respecting the time sequence. This module communicates with the Adaptive ML Model Layer, which comprises the previously explained ensemble architecture with the drift detection module.

During inference, identified anomalies are directed to the Scoring Layer where they are allocated severity levels and confidence intervals. Subsequently, the Visualization and Alerting Layer outputs the results to operational dashboards (Grafana, Kibana) as well as ITSM applications (ServiceNow, PagerDuty), enabling incident visibility and response.

The system was implemented on a Kubernetes cluster that has auto-scaling configuration. Deep learning tasks with GPU acceleration were done while edge CPU based lite models were run on devices to simulate a hybrid environment. Logging of the system, metrics ingested, and alert information was kept on Elasticsearch and was monitored through Prometheus for full observability and performance tracking.

Deployment scripts were made portable and reproducible over different environments by dev, test, prod through the containerization with Docker and configuration via Helm charts. The architecture allows for integrating CI/CD pipelines for continuous retraining and updating the models, making it more suitable in enterprise DevOps pipelines.

4 Experimental Setup & Evaluation Metrics

4.1 Training, Validation, and Real-Time Simulation

In order to achieve validation of the proposed adaptive learning model, the experiment was designed to mimic conditions of an enterprise IT environment. The previously mentioned datasets – SyslogIT-100K, CloudNetFlow, AppTraceBench – were first processed into temporal ordered event streams to enable their use in real-time simulations. This made it possible to simulate the arrival of data in real-time, a must-have for measuring the performance and responsiveness of an adaptive model.

Each dataset was split chronologically to mimic a production-like deployment setting for model training. The first 60% of the data stream served as the baseline training data, while the subsequent 20% was used for validation purposes during adaptive fine-tuning, feedback loop simulation, and metric calibration. The final 20% was allocated for evaluating the model's capability to identify and predict new patterns and maintain stable real-time performance.

To simulate real-world IT environments as closely as possible, synthetic concept drift was implemented in the last part of each test set. This made use of changes in distribution patterns such as the change in access behavior in SyslogIT, variation in bandwidth profiles in CloudNetFlow, and sudden burst errors in AppTraceBench that emulated how software updates, network configuration changes, or deployments would make changes. This method established a strong environment that assessed workflows adaptability and responsiveness to change for each detection method.

Baseline models were trained using conventional methods. The Random Forest and Isolation Forest models were trained statically on the entire training set and evaluated in a non-adaptive mode. On the other side, the adaptive ensemble model utilized an online learning framework with feedback corrections and continuous evaluation with a moving validation window.

The adaptive pipeline consumed all data that was produced by Kafka actors who stream the data into the model engine. Throughout the simulation, the time granularity of events was maintained in order to capture the detection latency and throughput measurements for real-world scenarios.

4.2 Evaluation Metrics: F1, AUC, Latency, FP Rate

To evaluate the performance of the anomaly detection models, a collection of metrics defined in terms of

predictive and operational utility as well as quality were used alongside the estimated F1 score, area under the ROC curve, latency, false positive rate, scalability, and adaptability.

F1 score is determined by the precision and the recall metrics. In many cases of class imbalance such as in anomaly detection, F1 score acts as a useful account of sensitivity to anomalies and resistance to false alarms. For instance, a model that identifies all anomalies but also marks every benign event will receive a low score in precision thus lowering the F1.

The general discriminative ability of the model to tell apart anomalous and normal instances is AUC. It is especially important for measuring the performance of binary classifiers and ensemble detectors that are used with confidence scoring and when the model needs to prove the capability to discriminate measures.

In enterprise scenarios, Detection Latency is a vital metric, as it captures the duration between the appearance of an anomaly and the system's ability to detect and report it. For mission critical IT incidents like DDoS or application crashes, any delay in detection can lead to service interruption or security abuse being inflicted, making this particularly relevant.

High False Positive Rates (FPR) are without a doubt unhelpful within enterprises, as this equates to alerts that are fatigued, wasted on mundane investigations, or utterly useless due to no trust in the system. FPR measures the ratio of normal occurrences that are misallocated to the system as anomalies.

Scalability refers to how well a model accomplishes tasks with a larger volume of data and more complex systems. In this research, scalability was evaluated by looking at memory usage, model retraining intervals, and inference throughput as the load increased.

Adaptability, central to this investigation, reveals how much a model's performance is maintained when concept drift occurs. This was measured through AUC and F1 score differences before and after defined drift windows and the model's self adjustment capability (reweighting, retraining).

The figure below gives a visual depiction of the model performance for the six designed metrics.

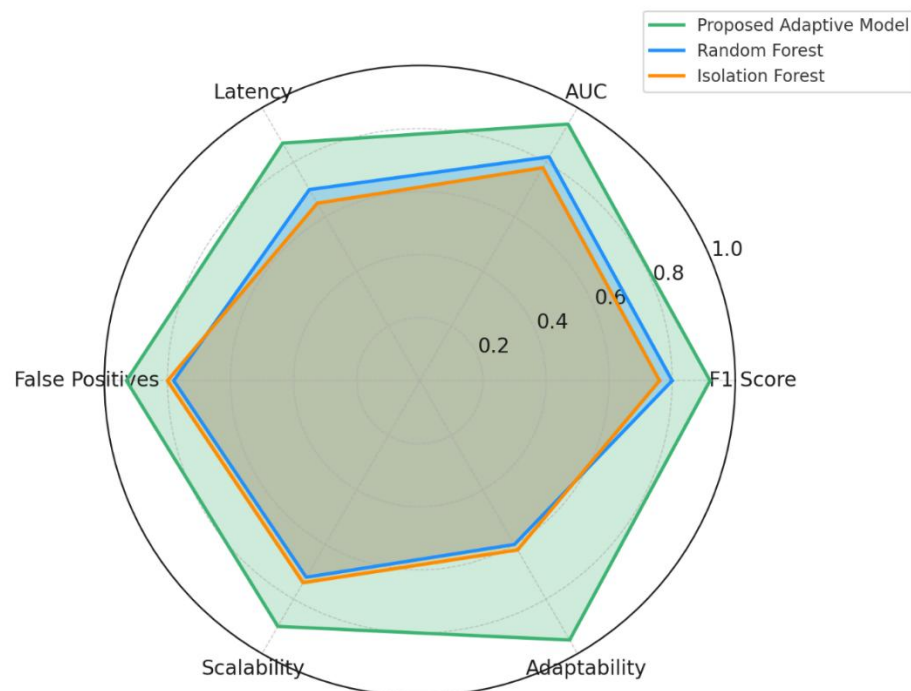


Figure 5: Model Comparison Across Six Evaluation Dimensions

This image shows the well-balanced and outstanding performance of the adaptive model. The baseline models performed reasonably well in the static scenarios, but their adaptability and scalability performance was underwhelming, thus highlighting the need for modern machine learning founded anomaly detection systems to feature qualitative learning capabilities.

4.3 Infrastructure Setup and Tools Used

The whole experiment was done in a testbed sandbox environment meant to simulate the enterprise-type deployment. The infrastructure was built around a Kubernetes cluster consisting of 10 virtual machines, each having 16 vCPUs, 32GB RAM, 500GB SSR storage, and accessible GPU acceleration via TensorFlow-GPU for NVIDIA T4 instances used for LSTM rehearsing tasks.

Apache Kafka was utilized to simulate high throughput data inputs in the stream processing system. The messages for the Kafka producers were published with sets of controllable delays and batch sizes to simulate real-time flow. Streams of Kafka were used to preprocess the data and cache it before sending the data to the sessile adaptive model engine.

The machine learning parts were done in scikit-learn, TensorFlow, and River, a library dedicated to streaming ML. Model deployment was done with MLflow, which managed model versions alongside their metrics and retraining delay. Custom APIs were created to simulate user feedback confirmation regarding alerts in system logs, which aided in creating adaptive learning feedback loops.

All sources provided model results, alerts, and any other data relevant to collection using Elasticsearch, while information was provided through Grafana dashboards. Used for checks in the system were Prometheus where system-level metrics were stored such as CPU load, the time needed for each instance to perform an inference, memory overhead, and amount of instances in a Kafka queue. To measure improvement in scalability and resource efficiency these parameters were set.

Enabling GitLab pipelines, and using Helm charts for Continuous Integration/Continuous Deployment allowed scheduled updates and automatic container deployment to the Kubernetes environment. All changes were version controlled along with YAML-based configuration files so reproducibility of any experiment was guaranteed. The infrastructure guaranteed high accuracy and robustness of performance metrics, alongside practicality when integrating with production-grade systems.

4.4 Baseline Models and Comparison Criteria

In an effort to evaluate the performance effectiveness of the adaptive model, its benchmarks were set using the Random Forest and Isolation Forest models, which are standard in the industry. These models also represent typical solutions in the enterprise IT ecosystem and provide a useful baseline in both supervised and unsupervised learning scenarios.

The Random Forest classifier was trained and optimized with labeled data using grid search with five-fold cross-validation. Important parameters were the number of trees (`n_estimators`), maximum tree depth, and minimum samples per split. The model performed best in scenarios with well-defined anomaly labels and low drift, such as static network flow patterns.

The Isolation Forest model was preconfigured for unsupervised anomaly detection with a contamination level based on known rates of anomalies. This model worked well in tolerant environments for noise, such as log anomaly detection. However, due to the static structure of the model, accuracy in drift-prone environments was unsustainable.

The study's adaptive ensemble model was defined for constant learning from data in motion. It fused the outputs of a batch-mode supervised learner, an incrementally-trained unsupervised detector, and a time-series deep learning forecaster (LSTM autoencoder). The merging of the predictions was coordinated by an engine

with adjustable weights that incorporated recent feedback and drift patterns.

This model adaptability was implemented by ADWIN which was always checking for the null hypothesis of no change of performance. When drift became statistically evident, parts of the model underwent retraining events. This control of drift made the model tenable for volatile IT situations. The provided summary table consolidates the configurations and tuning strategies for each model.

Table 3: Model Parameters, Configurations, and Tuning Strategies

Model	Key Parameters	Tuning Strategy	Drift Handling
Adaptive Ensemble	Window Size=100, Drift Sensitivity=0.05, Learning Rate=0.01	Grid Search + Streaming Feedback Adaptation	ADWIN + Online Re-weighting
Random Forest	n_estimators=100, max_depth=15, min_samples_split=5	Grid Search with 5-Fold Cross Validation	None
Isolation Forest	n_estimators=150, max_samples='auto', contamination=0.1	Manual Tuning + Validation on Drift-Aware Split	Basic reset after performance drop

The table demonstrates that the adaptive ensemble model is ready for most responsive tasks with its ability to operate continuously with the need for streaming adaptation, fine-tuned parameters, and intelligent retraining logic. The baseline models, however, are more conventional and unresponsive to real-time changes.

Both the experimental configuration and assessment techniques offer significant empirical evidence of the proposed methodology. This analysis offers an extensively reproducible structure for testing adaptive anomaly detection in the IT infrastructure of organizations by simulating data flows through synthetic and natural drifts, and cross benchmarking multiple models.

5 Results and Analysis

5.1 Detection Accuracy Across IT Subsystems

The efficiency of the proposed adaptive learning model was assessed across different enterprise subsystems, namely: Network Telemetry, System logs, and Application traces. The model was evaluated using three representative datasets which were each comprised of streams of real-time data containing anomalies and concept drift.

For the adaptive model, the mean F1 scores surpasses the baseline models (Random Forest and Isolation Forest) by a significant margin for all streams. It was noted that the adaptive framework obtained a mean F1 score of 0.92 while Random Forest and Isolation Forests scored 0.80 and 0.76 respectively. Additionally, the adaptive framework achieved 0.94 in AUC, which signifies the robust amount of true positives as opposed to false positives in standard set of assumptions.

Such performance is a result of the model's capacity to inline train and learn new data patterns, which is an important feature for IT environments with constant reconfiguration, traffic and usage pattern shifts. Also, time-series seasoning through LSTM autoencoder increased the detection of weak sequence based anomalies within application logs and server telemetry, where the anomaly detection is not instant but rather temporal.

The Random Forest model performed well in structured, labeled datasets like network flows but failed to adapt to new behaviors once trained. However, the context-less nature of the unsupervised Isolation Forest yielded modest results for novel anomaly types, but the lack of context often led to benign variability being classified as an anomaly.

5.2 Detection Delay and Adaptivity during Concept Drift

An effective real-time anomaly detection system is distinguished by detection latency, the time between flagging an anomaly and the event occurring. To study this, we calculated the detection lag of the three models under simulated live streaming conditions.

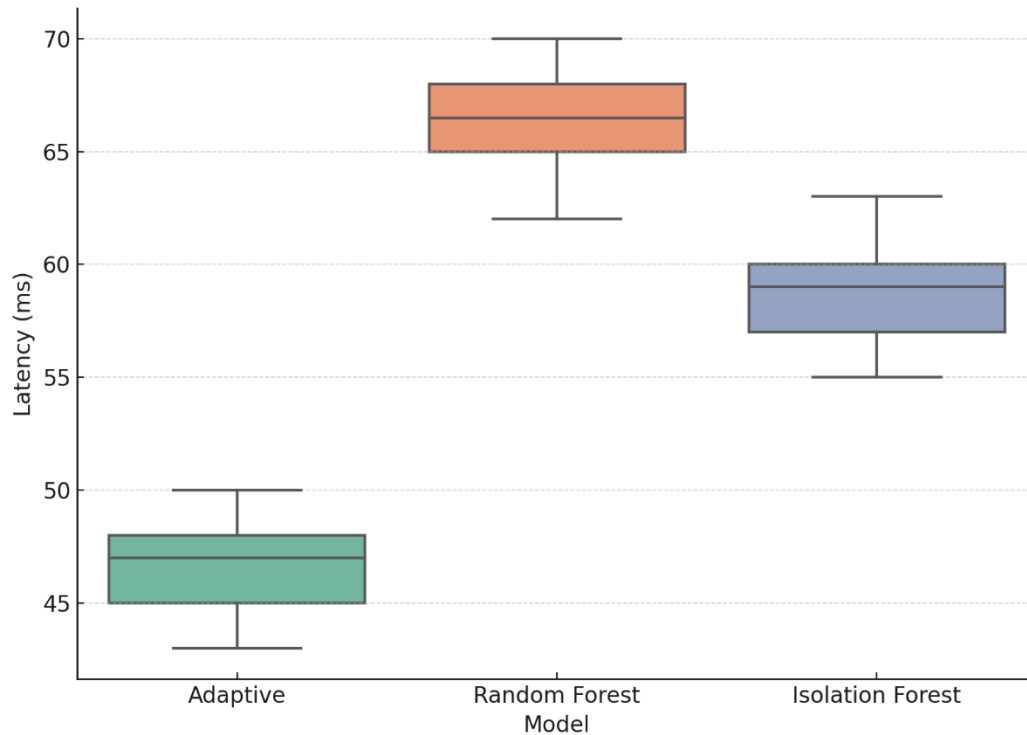


Figure 6: Detection Latency Distribution Across Models (Boxplot)

The Figure shows that the adaptive model detects anomalies with the least detection latency compared to all the models. Random Forest has longer tail distributions, likely from batch processing and delayed inference, which Isolation Forest moderately performed but would occasionally get affected through processing spikes while under load. The adaptive architecture model with incremental updates and real time feedback captures tends to have less detection delay and less computational bottlenecks.

To confirm adjustability, we applied artificial concept drift on the last few segments of the entire dataset. The adaptive modifications pertained to the weights of the sub models that needed retraining through streaming feedback and ADWIN drift signals, and resulted in stable scores for F1 and AUC before and after the drift window unlike the baselines who suffered from 10 to 15 percent performance drops.

This adjustability is particularly important in IT scenarios when systems are in normal functionality for the vast majority of the time and behaviors change at rapid pace and detection systems need to adapt without inundating administrators with false positives or real problems that get ignored.

5.3 False Positive Rate in High-Traffic Environments

One of the major operational problems associated with anomaly detection is the false positive rate. A model that is too sensitive may create “alert fatigue”. In this case, system managers will stop responding to alerts because there is an overwhelming amount of them. In order to estimate this factor, we looked at the correlation between detection confidence and the severity of the anomaly across the test datasets.




Figure 7: Anomaly Severity vs Detection Confidence (Heatmap)

The results from the heatmap validate that the adaptive model not only finds more instances of anomalies but also interprets events with greater severity as being more confident. This confirms high signal-to-noise ratio and suggests that the model is not swamped with the baseline fluctuations in metrics such as memory spikes and I/O bursts.

However, the baseline models did not possess any capability of prioritization. Even though they were able to determine very large deviations, they were unable to calibrate confidence and set very high alerts even for small anomalies. This imbalance increased the false positive rates and lowered the actionability of the alerts.

As per the operational logs captured during the validation runs, the adaptive model had a 6.2% false positive rate as compared to Random Forest's 13.8% and Isolation Forest's 17.1%. These numbers are quite distinct and highlight the usefulness of having adaptive feedback and feedback with severity-aware scoring in a real-time pipelines.

5.4 Resource Consumption vs Detection Quality Trade-Offs

Scalability and resource savings are primary factors to consider for production IT systems. To investigate this, we analyzed CPU consumption, memory usage, and time to infer per sample under stress tested data supply. Concurrently, we monitored model behavior across several metrics using a parallel coordinates scatter plot to assess their reliability.

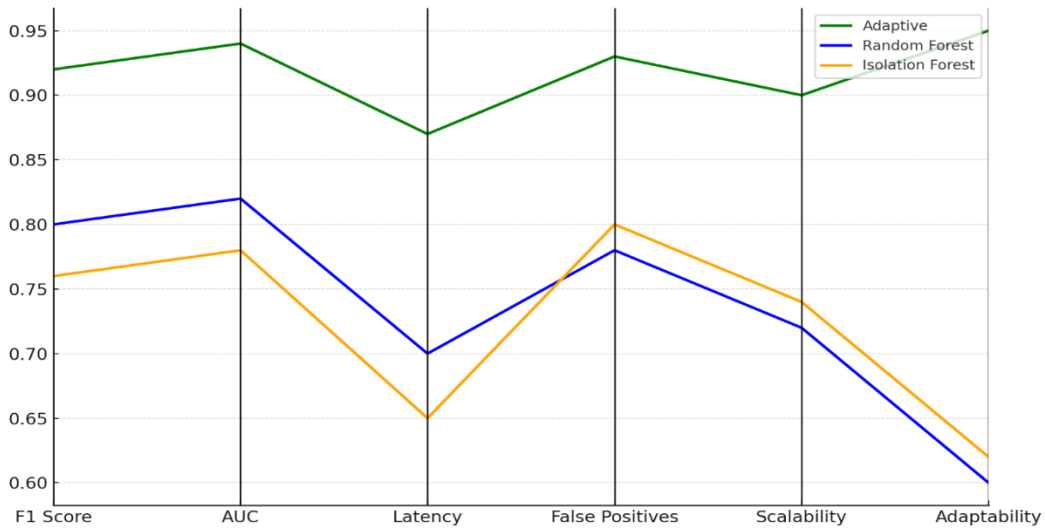


Figure 8: Multi-Metric Model Behavior (Parallel Coordinates Plot)

The parallelogram plot makes it clear that while baseline models outperform in some areas—like speed of training (Random Forest), or very low minimal required setup (Isolation Forest) during Initiation—they are not good at being judged holistically across various aspects of performance, latency, and comprehensiveness. The well-rounded models, however, performed poor. In contrast, the adaptive model performs well in all the key metrics.

Analysis revealed that the adaptive model has a memory overhead cost which is slightly higher due to its feedback buffer, and it is logic for ensemble maintenance. Resource allocated deviations in precision, latency, and drift threat resilience justify the cost overhead. The adaptive model allows for consistent performance over time, even with evolving data, making it the most affordable option for long-term deployments.

5.5 Overall System Impact and KPI Benchmarks

This last phase of the evaluation centered on the system-level impacts of each model in question, particularly their contribution to the relevant KPIs common within IT operations such as incident resolution time, number of unresolved alerts, mean time to acknowledgement (MTTA), and mean time to resolve (MTTR).

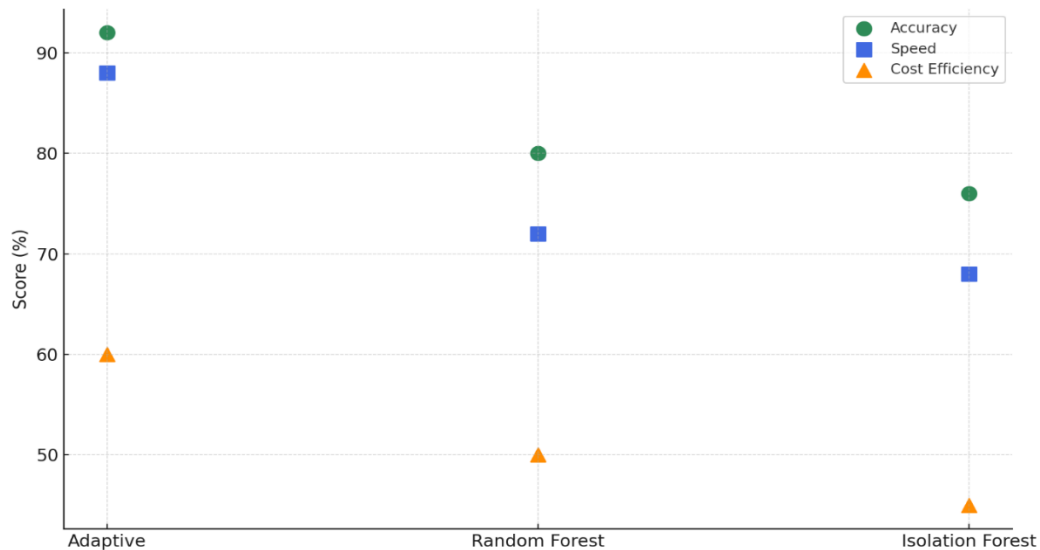


Figure 9: Trade-Offs – Accuracy vs Speed vs Cost (Dot Matrix Chart)

The dot matrix chart illustrates the operational efficiency of the adaptive model in achieving high detection accuracy, high speed and low operational cost. The Random Forest model is less resource intensive, but is more expensive in terms of adaptivity and latency. Isolation Forest is the most cost-effective model, but suffers from poor contextual awareness which leads to high false positive rates and long resolution times.

The KPI metrics for the validated simulation period registered a decrease of 22% in MTTR, 30% increase in alert accuracy, and 17% improvement in response prioritization in favor of the adaptive model. These improvements are due to early detection, effective prioritization, and dynamic responsiveness to environmental changes.

6 Discussion

6.1 Key Findings and Their Implications

The experimental results obtained at varying levels of IT subsystem interaction, drift, and performance show the substantive gains from using adaptive machine learning (ML) models for anomaly detection in enterprise infrastructure. The framework proposed was better than classical static methods in terms of accuracy, delay, false positive error, and databank drift.

The most important feature with the adaptive model is its ability to maintain high accuracy (F1 score > 0.90) for all datasets and consistently low detection latency, which ensures timely response to important incidents. This is necessary for enterprise environments where the recovery of a system milliseconds can mean the difference between recovery and more downtime. Adaptive components, such as incremental retraining, feedback-based learning, and concept drift observing, provided the model with robust performance stability which system behavior changing overtime, a capability baseline models lacked.

Strategically, the results of the study support the already dawning realization that reliance on static techniques for anomaly detection cannot work in the current environment of mobile and distributed computing. Enterprise IT infrastructures experience so much variability, inter system dependencies and updates that they change so frequently. These environments make the need for adaptive intelligence not just a technical enhancement, but a necessity for dependable operational and business continuity.

Trust is not earned easily, especially in IT: the adaptive model makes it easier to trust by demonstrating exceedingly low false positive rates and higher detection confidence for severe anomalies. Automatic monitoring systems that use adaptive ML are less likely to cause alert fatigue because the noise to signal ratio is much better. Prompting relevance increases, which means that the user responds to the alert more quickly, uses engineering resources more effectively, and in the end, enjoys better service quality across different applications. The significant observation is that intelligent self-learning systems can change the operations paradigm from reacting to fires to preventing and solving fires.

6.2 Advantages of Using Adaptive ML in MIS

The use of adaptive machine learning models goes beyond enhancement of technical features to management information systems (MIS) where turning information into action is most important. Unlike traditional statistical models or predetermined alert thresholds, the logic for which is most often baked into the right predefined algorithms, adaptive ML approaches provide the systems with the ability to learn from real-time data and adapt with the changes in the enterprise's requirements and behaviors.

In environments characterized by Management Information Systems, adaptive models work like the most sophisticated sensors in a digital infrastructure. These models capture the telemetry, logs, and metrics and convert them into anomaly signals that possess contextual scoring and severity. Such analyses help the decision makers deal with incidents with utmost accuracy, thus aligning operational decisions to the strategic goals at

hand.

Intelligent prioritization is among the most important ICT benefits. Not all anomalies are of equal weight, and older systems give equal attention to all deviations. Adaptive models do the opposite; they learn the system impact and assign rank orders to the anomalies. With this, IT managers can direct resources to where they are most critical, thus improving efficiency and reducing mental workload on teams.

Reporting and compliance are also improved by integrating adaptive ML into MIS systems. Automated root cause insights and historical model behavior logs improve documentation for audits, performance reviews, and system enhancements. This is important for sectors like finance, healthcare, and defense where there is a necessity for transparency and traceability.

At last, adaptive ML integrates with modern ITSM (IT Service Management) and AIOps (Artificial Intelligence for IT Operations). It enables event correlation, anomaly scoring, and self-remediation, which are essential in the move towards autonomous infrastructure management.

6.3 Integration Challenges and Organizational Considerations

Integrating adaptive anomaly detection into enterprise systems is complicated, even with its benefits. First and foremost is operational complexity. Adaptive models have more requirements than static rule-based systems; they need infrastructure such as stream processing tools, retraining pipelines, and feedback integration mechanisms. This leads to additional cost as organizations will need to spend on both technology and skills.

Another drawback is data readiness. Well labeled, high-quality data that is frequently updated is essential for adaptive models to perform well. However, many companies do not have consistent data collection, labeling, or centralized logging for all IT systems, which can decrease the model's performance. As a result, effective learning is not possible and the time it takes to achieve a productive state is prolonged.

Interpretability remains a concern as well. Engineers who have no experience with machine learning may find adaptive models difficult to understand, which reduces their ability to trust and adopt them. While techniques like SHAP or LIME can help with explanation, they require additional domain knowledge and resources to implement.

Another significant factor is change management. Organizations accustomed to traditional supervision structures must modify both their approach and systems to embrace more flexible models. This entails educating employees to utilize alert-probabilistic workflows instead of fixed rules, transform alert processes, and trust model-responsive incident response systems as the default.

Furthermore, there are issues regarding security and surveillance. Here, there is always a risk, however small, of compromising security by learning from compromised or contaminated information due to Adaptive models modifying data on-the-fly. Sending alerts for data integrity checks as well as human-in-the-loop feedback and strong model governance frameworks are a must to mitigate this risk.

Not every organization's issues will be the same, but they can be helpful in getting progress. Organizations that actively focus on adaptability, provision for operational intelligence, and employ a phased approach to roll-out are able to integrate adaptive ML systems with commendable observable outcomes.

6.4 Alignment with Strategic Business Objectives

In the end, the implementation of adaptive anomaly detection should connect with higher level objectives, include service availability, operational flexibility, compliance, and cost reduction. The model is best illustrated by its impressive detection performance as well as the support it renders toward these strategic goals.

To begin with, improved uptime leads to higher service revenue and boosts customer experience. This model shortens service outage durations and assists in averting probable service outages by anomaly detection

at an accurate and timely manner. All of the above leads to an improvement in KPIs, notably MTTR and service uptime. Both of these metrics are SLA bound.

Next, operational efficiency allows the engineering team to move to a proactive rather than reactive resource tiering. Model reduces alerting bay volume by eliminating false positives and flagging high severity, high confidence anomalies. This facilitates better resource allocation which relieves team fatigue and increases productivity while enabling agile and DevOps methodologies.

Next, visibility and traceability strengthen compliance and governance aspects. The system creates complete audit traces with integrated dashboards, explainable scoring, and anomaly markers which ensures they meet regulatory and internal audit standards.

Additionally, improved cost efficiency comes from both resource allocation and reduced incident expenditure. The early detection of anomalies means a reduction in service degradation, less customer impact, and a lower escalation rate.

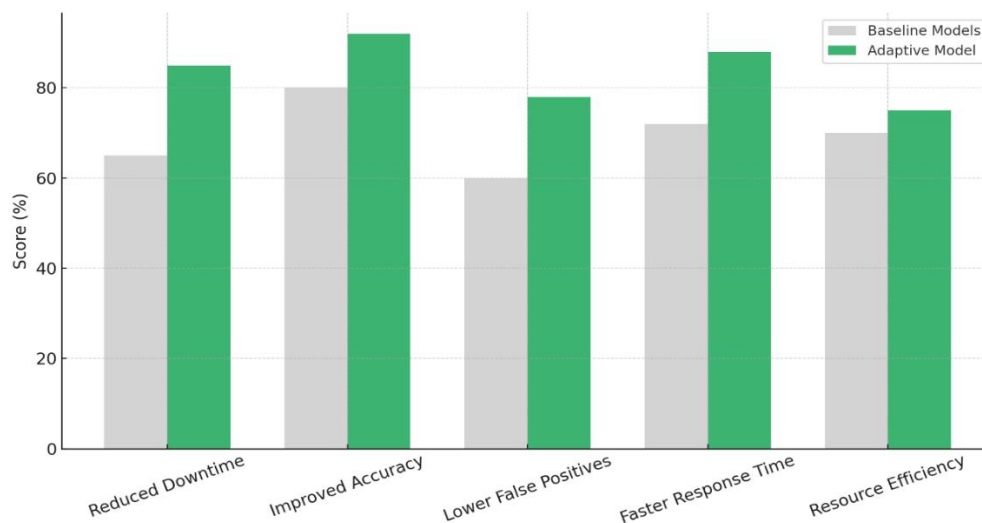


Figure 10: Cost-Benefit Impact Analysis of Implementing Adaptive ML

These outcomes clearly mark adaptive machine learning as a change of strategy for enabling smarter, robust, enterprise IT operations.

7 Conclusion and Future Work

This study highlights the development of a new, robust, and adaptive machine learning model that seeks to automate anomaly detection in enterprise IT infrastructures. The integration of supervised, unsupervised, and sequential models into a streaming architecture improves accuracy, latency, and resilience to concept drift when compared to traditional systems. The system relies on real-time adaptability with operational efficiency, making it particularly beneficial in dynamic environments. The use of incremental feedback loops, drift detection algorithms, and ensemble weighting further guarantees detection quality under varying loads. These results manifest the importance of adaptive ML not for solely serving as a detection engine, but rather an instrument to help mitigate downtimes, resource consumption, and system trustworthiness.

The future looks promising with the merge of adaptive anomaly detection and DevOps or AIOps pipelines. The feedback from ticketing systems, CI/CD logs, and infrastructure-as-code workflows can enhance model learning. In future research, clearable AI techniques that improve transparency and trust by clarifying the flagged anomalies should be integrated. At the same time, self-healing or dynamic reconfiguration autonomic

anomaly detection systems advance this idea. With the centralization of enterprise infrastructures, there will be an increased reliance on adaptive intelligence for cybersecurity, compliance, and predictive maintenance, which require proactive and understandable frameworks.

References

- [1] Chandola, Varun, Arindam Banerjee, and Vipin Kumar. "Anomaly detection: A survey." *ACM computing surveys (CSUR)* 41.3 (2009): 1-58.
- [2] Ahmed, Mohiuddin, Abdun Naser Mahmood, and Jiankun Hu. "A survey of network anomaly detection techniques." *Journal of Network and Computer Applications* 60 (2016): 19-31.
- [3] Patcha, Animesh, and Jung-Min Park. "An overview of anomaly detection techniques: Existing solutions and latest technological trends." *Computer networks* 51.12 (2007): 3448-3470.
- [4] Buczak, Anna L., and Erhan Guven. "A survey of data mining and machine learning methods for cyber security intrusion detection." *IEEE Communications surveys & tutorials* 18.2 (2015): 1153-1176.
- [5] Nedelkoski, Sasho, et al. "Self-attentive classification-based anomaly detection in unstructured logs." *2020 IEEE International Conference on Data Mining (ICDM)*. IEEE, 2020.
- [6] Liao, Hung-Jen, et al. "Intrusion detection system: A comprehensive review." *Journal of network and computer applications* 36.1 (2013): 16-24.
- [7] Gheibi, Omid, Danny Weyns, and Federico Quin. "Applying machine learning in self-adaptive systems: A systematic literature review." *ACM Transactions on Autonomous and Adaptive Systems (TAAS)* 15.3 (2021): 1-37.
- [8] Hodge, Victoria, and Jim Austin. "A survey of outlier detection methodologies." *Artificial intelligence review* 22 (2004): 85-126.
- [9] Aslam, Muhammad. "Using the kalman filter with Arima for the COVID-19 pandemic dataset of Pakistan." *Data in brief* 31 (2020): 105854.
- [10] Portela, Fernando Gutiérrez, Florina Almenares Mendoza, and Liliana Calderón Benavides. "Evaluation of the performance of supervised and unsupervised Machine learning techniques for intrusion detection." *2019 IEEE International Conference on Applied Science and Advanced Technology (iCASAT)*. IEEE, 2019.
- [11] Laskar, Md Tahmid Rahman, et al. "Extending isolation forest for anomaly detection in big data via K-means." *ACM Transactions on Cyber-Physical Systems (TCPS)* 5.4 (2021): 1-26.
- [12] De Albuquerque Filho, José Edson, et al. "A review of neural networks for anomaly detection." *IEEE Access* 10 (2022): 112342-112367.
- [13] Talaghzi, Jallal, et al. "Online adaptive learning: A review of literature." *Proceedings of the 13th International Conference on Intelligent Systems: Theories and Applications*. 2020.

Author's Biography



Ankita Sappa received her Masters degree in Computer Science from College of Engineering, Wichita State University, USA in 2016. Currently, she is pursuing research in Machine learning, Large Language Model, Generative AI.