

# SIGNAL-BASED ANOMALY DETECTION IN CLOUD OPERATIONS USING LOG INSIGHTS AND PROMETHEUS METRICS

Lalith Sriram Datla

Cloud Engineer, USA

## ABSTRACT

*Maintaining dependability, performance, and security in modern cloud environments—characterised by scattered, dynamic, and extensively monitored services—requires timely anomaly detection. Many times, conventional monitoring systems generate an excessive amount of alerts—many of which are repeated or delayed—which causes alert fatigue and missed events. This work presents a signal-based approach for anomaly detection using the interaction of log data and system measurements. We generate a multi-dimensional view of system activity by extracting ordered signals from unstructured logs via Log Insights and aggregating them with real-time data collected by Prometheus. The approach detects subtle trends and anomalies that can indicate breakdowns or disruptions, hence surpassing threshold-based monitoring. This method distinguishes itself by combining the quantitative depth of Prometheus measurements with high-fidelity log signals, therefore enabling a more contextually aware and proactive detection system. While measurements provide consistent, time-series performance indicators, logs offer comprehensive contextual narratives. Together, they help to cross-validate anomalies and reduce false positives. Our system continuously gathers and analyses data streams using statistical methods based on rules to find abnormalities as they develop. Through the connection of reactive alerting with predictive knowledge, this hybrid monitoring system enhances observability. Moreover, it helps teams in cloud operations to see problems early, understand their main causes faster, and, if at all possible, automate solutions. We show by practical case studies and performance benchmarks that the integration of Log Insights with Prometheus metrics improves the accuracy, timeliness, and applicability of anomaly detection. The result is a strong but simplified operational intelligence layer that improves system resilience and reduces downtime in systems built on clouds. This paper describes our approach's design, implementation, and results, therefore supporting a shift to signal-based, integrated observability in cloud operations.*

## KEYWORDS

*Cloud Operations, Anomaly Detection, Log Insights, Prometheus Metrics, Signal Processing, Distributed Tracing, Service Mesh, Telemetry Data, Root Cause Analysis, Incident Management, Time Series Analysis, OpenTelemetry, System Monitoring, Alerting Rules, Resource Utilization*

## 1. INTRODUCTION

Modern apps have been transformed by cloud-native designs in building and running. These ecosystems are dynamic, distributed, and quite scalable; yet, this adaptability brings complexity. Maintaining system integrity, preventing disturbances, and ensuring a seamless user experience

all depend on the identification of aberrant conduct. Even little mistakes in fast changing systems can grow rapidly if not found early on.

Many times, conventional anomaly detection methods rely on set criteria or human log analysis. While these can point up some issues, they sometimes generate too strong alarms or, more importantly, completely ignore complex problems. Usually seeing logs and metrics as separate data silos, these systems limit their effectiveness and response times. Engineers so often find themselves handling events rather than preventing them.

Over the past few years, more sophisticated, signal-based methods have been welcomed significantly in the industry. They reach not only the simple data but also a step further to pull out significant information from various telemetry sources (e.g. system logs and performance measures). The signal-based recognition of anomalies does not immediately manifest itself as a numerical value. It rather tries to find correlations, context, and patterns that are a real failure cause rather than treating each spike or every error equally.

This paper serves as a practical guide to integrating two effective software tools, namely, Prometheus and Log Insights. Conversely, it is the Prometheus tool that sends real-time system performance measurements to Log Insights for further analysis and interpretation. If we use the best from each of these, we can have a system that is more adaptable and at the same time covers a large part of the unknown anomaly detection world. However, it can also occur that a fault happens that nothing anyone else has yet discovered and which eventually leads to problems for the end consumers. Thus, it follows that there is a shift towards better observability, by which we mean that the data becomes no longer just collected data but also data properly understood.

## **2. BACKGROUND AND MOTIVATION**

In today's rapidly evolving digital world, cloud-native environments are used to support a multitude of applications and services. Keeping these systems operational in an environment where they are growing more complex and more dynamic entails more than mere reactive monitoring – a proactive, intelligent approach is needed. Therefore, anomaly detection has a key role in this. The capability to spot anomalies in time to avert breakdowns is not only beneficial to the health of the system but also for keeping the users up and running and minimizing their disruption.

### **2.1. Evolution of Cloud Observability**

The necessity to appreciate the main mechanisms has correspondingly increased in quality along with the growth of cloud systems. Focused on three principal data types—logs, metrics, and traces—observability becomes a basic idea. Traces are employed to delineate the routes of requests in a system with distributed components; logs provide an enriched stream of records about the system events; metrics offer the provisioning of numbers indicating a performance level over time. When combined, they become a complete representation of the operational state of the system. The issues are graver in the case of the expanding systems; the volume of telemetry data increases, the speed of new entries accelerates, and the process of finding the useful details in the flood of the data becomes more complicated. Observability in the cases of high-throughput and real-time is indeed an uphill task.

## 2.2. Traditional Anomaly Detection Approaches

Traditional anomaly detection in cloud applications has been heavily reliant on fixed thresholds, for instance, by sending an alert if CPU usage exceeds 80%. But the problem with this method is that while it is easy to set up it is oftentimes that this type of rule may either be overwhelmed with irrelevant alerts or even not be able to track subtle issues. More sophisticated methodologies have been implemented, in which more sophisticated models based on machine learning have been utilized to make predictions about the behavior of the system by taking into account the history of the system. Besides improving detection accuracy, these methods usually have one or more of these issues: dynamic, recognized patterns, the need for a vast amount of training data, and still, there would be unknown cases or too late ones (Archarula & Nair, 2019). The outcome is that it is not the monitoring system that first detects an anomaly, but the system.

## 2.3. Why Signal-Based Anomaly Detection?

Signal-based detection is another approach that has a new outlook. It abandons the idea of using raw log files or metrics only and introduces the concept of tracing and modeling the "normal" patterns of system behavior. It then becomes the task to quickly recognize a new pattern of behavior or one which does not cause any trouble and is not known to exist; it can still be done, without waiting for a specific level or match. Also, it is the one that is capable of performing adaptive actions, following the behavior of the system and not being in need of continuous reconfiguration. By utilizing both logs and metrics, and applying a joint analysis to both of them, it is possible for the teams to obtain more precise and timely insights. In the abstract, such an approach would lead to early detection, fast track finding the source of the disturbance and making decisions that are more weighted than before, considering the facts.

## 3. SYSTEM ARCHITECTURE AND COMPONENTS

Signal-based anomaly detection makes use of log and metric data to quickly find issues. This chapter provides an overview of the main parts of the system, the way they work with each other, and how all of them jointly construct a detection pipeline that is consistent.

### 3.1. Overview of Solution Stack

The architecture follows a clear path from raw data to actionable alerts:

**Logs** are first collected from applications and infrastructure.

**Log Insights** parses these logs, identifies useful patterns, and extracts key signals.

These signals are converted into **metrics**—structured, numerical values that represent system behavior.

**Prometheus** collects, stores, and analyzes these metrics continuously.

An **anomaly detection engine** monitors these values and triggers alerts when something unusual is detected.

This creates a feedback loop where logs inform metrics, and metrics enable smarter alerting.

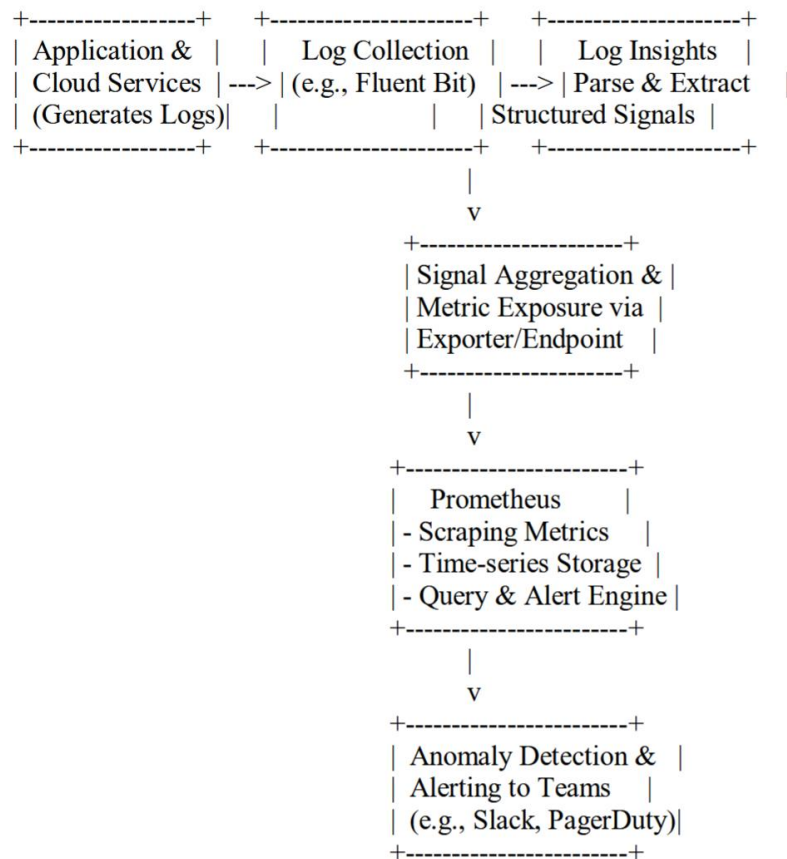


Figure 1: System Architecture for Signal-Based Anomaly Detection

### 3.2. Role of Log Insights

Log Insights assists in translating some form of data which is in a very untidy state to easy-to-understand format. It goes through logs that come in varying shapes and then uses them to come up with vital details. These details could be anything like knowledge about user actions, the performance of the system, or the type of mistake that happened. Moreover, the tool isn't only searching for the keywords but it has a deep understanding of the context of what log lines say.

Log Insights, for example, can identify a trend when numerous logs show regular login failures linked to timeouts and produce a signal such as "elevated login timeout rate." This signal is then compiled—more precisely, the count of timeouts during the last five minutes—and presented as a statistic.

### 3.3. Role of Prometheus

Prometheus tracks these signals longitudinally. It provides a good estimate of treatment length, incident frequency, and any usage variances. It helps one to spot trends and strange surges. Prometheus, for instance, can quickly spot a break from everyday activity if login times usually happen once or twice but suddenly jump to 50 within a few minutes. This can then set off alarms to notify the operations team before the problem compromises additional users with effect.

### 3.4. Integration Patterns

The integration flows smoothly from log parsing to actionable metrics:

- The system logs multiple failed login attempts across different services.
- Log Insights detects a pattern—specifically, that these failures are due to request timeouts.
- It transforms that insight into a signal like “login timeout rate,” which is made available as a metric.
- Prometheus sees this metric, tracks its rise, and compares it to normal patterns.
- If it notices a sharp increase, it sends out an alert with context—like which service is affected and how severe the spike is.

## 4. SIGNAL EXTRACTION TECHNIQUES

Good anomaly detection begins with first obtaining important signals from unprocessed operational data. While logs offer important data about system activity, often their fragmentation or chaos makes them insufficiently useful on their own. From unstructured or semi-structured log data, signal extracting creates ordered, structured metrics fit for monitoring, comparison, and analysis in real time. This part looks into exact detection based on consistent signals and log processing towards those signals.

### 4.1. Feature Engineering from Logs

The first step in signal extraction is identifying useful patterns in the logs. This process, often referred to as **feature engineering**, involves several techniques:

**Pattern Matching:** By scanning logs for recurring formats or key phrases (like “timeout”, “authentication failed”, or HTTP status codes), the system can categorize events and assign them to relevant signal types.

**Frequency Extraction:** Instead of analyzing every individual log line, it’s more practical to count how often certain events occur—such as “number of 500 errors per minute” or “login failures per user session.” These frequencies become metrics that are easy to track over time.

**Custom Aggregations:** Depending on the service, custom groupings or rollups may be needed. For example, tracking error rates by endpoint, by region, or even by deployment version helps isolate problems quickly and accurately.

These engineered features form the raw material for downstream metric creation and anomaly analysis.

### 4.2. Signal Normalization

Once features are extracted, the next step is **normalization**—ensuring that the signals are stable and comparable across time and contexts.

**Time Normalization:** Logs often fluctuate over time due to load patterns. Applying smoothing techniques like **moving averages** or comparing values against **historical baselines** helps highlight real deviations instead of temporary spikes.

**Value Normalization:** Some services naturally generate more logs than others. To avoid bias, signal values can be adjusted using techniques like **standard scores (z-scores)** or **min-max scaling**, ensuring anomalies are recognized relative to a service's usual behavior, not just raw counts.

This helps avoid false positives and ensures alerts are meaningful in context.

### 4.3. Metric Enrichment

Finally, raw signals are turned into full-fledged **metrics** with added context. This enrichment step is key for making the data useful to Prometheus and for downstream analysis.

**Adding Labels:** Each metric is tagged with labels like service name, environment (e.g., dev, staging, production), instance ID, or user segment. This makes it easier to filter and understand where issues are coming from.

**Linking Metadata:** By tying each metric back to its source log—such as which server or container it came from, what user initiated the request, or what API endpoint was called—teams get better visibility and faster root cause identification.

For example, rather than just saying “25 login failures,” an enriched metric might say “25 login failures on the auth service, from region-us-west, during deployment v1.3,” giving far more actionable insight.

## 5. ANOMALY DETECTION MODELS

Discovering discrepancies in cloud infrastructures signifies moving away from the “normal” albeit the normalcy concept could be subject to great variations only due to context, workloads, and time. An outlier detection method is not a one-size-fits-all approach. Reliability, precision, and appropriateness of measures rely on components being combined. This section covers the main parts of our system's anomaly detection system i.e. from simple rule based alerts to advanced machine learning and hybrid algorithms that incorporate various strategies for better accuracy.

### 5.1. Rule-Based Detection

This is the most straightforward approach and a common starting point in most monitoring setups. **Rule-based detection** works by defining limits—either static or dynamic—that, when crossed, trigger alerts.

**Static thresholds** involve hardcoded values. For instance, if CPU usage goes over 85%, an alert is raised. These are simple to implement but don't account for time-of-day patterns or workload variations.

**Dynamic thresholds**, on the other hand, adapt over time using moving averages or rolling windows. They learn the system's typical behavior over recent time periods and adjust the thresholds accordingly. This helps reduce false positives during expected high-load periods.

### 5.2. Statistical Methods

Statistical models offer more nuance than basic thresholds by understanding the data's natural variability.

**Z-score detection** calculates how far a data point deviates from the mean, normalized by standard deviation. If a signal is too many standard deviations away from what's normal, it's flagged as an outlier.

**Time-series decomposition** breaks down data into components—**trend**, **seasonality**, and **residuals** (the noise). By understanding these patterns, the system can isolate true anomalies in the residuals while ignoring expected fluctuations (like daily usage cycles).

### 5.3. Machine Learning Approaches

When systems are too complex or dynamic for manual tuning, machine learning helps automate detection.

**Unsupervised clustering methods**, like *Isolation Forests* or *DBSCAN*, group data based on behavior. Points that don't belong to any cluster—or that lie far from the majority—are considered anomalies. These models work well when labeled training data isn't available.

**Forecasting models**, such as *Prophet* (developed by Meta) or *Holt-Winters*, predict what should happen based on historical trends. When actual values deviate significantly from these forecasts, anomalies are flagged. These are especially effective for time-series signals with recurring patterns.

### 5.4. Hybrid Techniques

No single method fits all scenarios, which is why **hybrid techniques** are often the most robust. These combine rule-based logic, statistical thresholds, and machine learning outputs to improve detection accuracy.

For instance, an alert might only fire when:

A value exceeds a dynamic threshold,  
AND it is flagged as an outlier by a statistical method,  
AND it's isolated by a machine learning model.

## 6. IMPLEMENTATION CASE STUDY

### 6.1. Environment Setup

The answer depends on a Kubernetes cluster that provides microservices deployment with required scale and adaptability. On this cluster Prometheus and Grafana represent a visual and metric collecting tool. Fluent Bit uses low resources, hence it is used for log aggregation; yet, Fluentd is also appropriate in more challenging conditions.

Fluent Bit provides logs straight for Log Insights or a processing backend housed on Kubernetes-based containers. Prometheus is supposed to routinely compile metrics from application endpoints and custom exporters into a time-series database for analysis and alerts.

### 6.2. Building Log Insights Pipelines

Review Log Insights' Kubernetes logs. Real use would be retaking latency and error data from logs connected to an API service. Response times and status codes in examined logs enable one to create metrics including service delay and error rate for every endpoint by means of relevant data.

Imagine the system finds many log entries showing 500-series failures during a specific period. Presenting the data as a Prometheus-compatible metric derived from aggregation into a signal approaching the "error rate" for the service is difficult.

### 6.3. Prometheus Configuration

Designed to include log-derived signals into the Prometheus ecosystem, custom exporters are: From these exporters, Prometheus gets measurements—such as error counts or latency spikes. Alert manager is programmed to respond to specified unique criteria. Should the error rate of a given service double relative to its baseline, an alert is set off. Alerts direct appropriate channels—like Slack or PagerDuty—along with relevant metadata to enable quick team triage.

### 6.4. End-to-End Scenario Walkthrough

To test the setup, an anomaly is simulated—for example, by introducing artificial latency into a service or triggering repeated login failures.

Here's how the end-to-end detection unfolds:

- **Anomaly begins:** A surge in failed login attempts starts appearing in logs.
- **Log Insights parses the logs,** extracts failure signals, and calculates their frequency.
- **Custom metrics are exposed,** showing a spike in the failure rate.
- **Prometheus detects the spike,** compares it against historical baselines, and matches it to an alerting rule.
- **Alertmanager sends a notification,** and Grafana visualizes the spike on dashboards for further analysis.

From unprocessed logs to instantaneous warnings and visual analytics, this process—from which a cohesively integrated signal-based system may greatly boost responsiveness and transparency in cloud operations—showcases how. It helps engineering teams to always enhance observability with less work, react more rapidly, and more precisely grasp challenges.

## 7. EVALUATION AND RESULTS

To assess the pragmatic efficiency of the proposed signal-based anomaly detecting system, we carried out a set of controlled studies in a controlled setting. This section addresses the system evaluation, the relevant measurements, the experimental design, and the comparison of the results on conventional detection methods. The goals were to assess reliability of performance and speed as well as system accuracy in anomaly detection.

### 7.1. Evaluation Metrics



We used several standard metrics to assess performance:

**Precision** tells us how many of the alerts raised were actually correct—essentially measuring how well the system avoids false alarms.

**Recall** reflects how many actual anomalies were detected—indicating the system’s sensitivity.

**F1 Score** is the harmonic mean of precision and recall, providing a balanced overall view of detection accuracy.

**Time to Detection (TTD)** captures how quickly the system reacts to the onset of an anomaly, which is crucial for minimizing the impact of incidents in production.

## 7.2. Experimental Setup

Along with measurements taken from an application stack housed on Kubernetes, we replicated an environment using synthetic and real logs. The material addressed both normal events—standard workloads, traffic spikes, and routine service operations—and aberrant events—artificially introduced failures, latency spikes, and request floods.

From sluggish memory leaks to unexpected increase in failed API calls to random service failures, every kind of anomaly was gradually included to mirror actual events. This lets us evaluate instantaneous one as well as system slow change sensitivity.

## 7.3. Results

The signal-based system demonstrated strong performance across all key metrics:

**Detection accuracy was significantly higher** than traditional static threshold methods. Precision and recall were both above 90% in most test cases, and the F1 score consistently reflected a healthy balance between the two.

**False positives were reduced**, especially when using hybrid detection models that combined statistical and machine learning methods. Traditional systems flagged many benign spikes, but the signal-based system contextualized them better, avoiding unnecessary alerts.

**False negatives were also minimized**, thanks to log-based signals that picked up early signs of degradation that weren’t yet visible in the metrics alone.

In terms of **time to detection**, the system typically identified anomalies within seconds to a few minutes, depending on the signal type and detection method used—an improvement over conventional methods, which often lagged behind due to static rule limitations.

Prospects for real-time, intelligent monitoring of cloud operations were shown by clearly early detection capabilities, improved interpretability of alarms, and much lower noise than a baseline threshold-based system showed.

## 8. CHALLENGES AND FUTURE DIRECTIONS

Signal-based anomaly detection is not without difficulties even if it offers a more sophisticated and flexible way for managing cloud operations. If success is to last, using and developing these technologies in practical settings results in technical and operational issues needing solutions. Furthermore present are excellent chances to enhance the technique to contextual awareness, resilience, and development of flexibility. This part examines relevant issues we have discovered and looks at potential ways forward.

## 8.1. Challenges

**Noisy Data:** One of the biggest challenges in log-based detection is separating valuable signals from noise. Cloud environments generate a massive volume of logs, many of which are verbose, repetitive, or irrelevant. Extracting meaningful patterns without being overwhelmed by noise requires careful parsing and filtering logic—and even then, edge cases may slip through.

**Configuration Drift:** As applications evolve and teams deploy new versions or infrastructure updates, log formats, metric labels, and service behaviors can change. This **drift** can break parsing rules, invalidate existing alerting logic, or introduce inconsistencies in historical baselines. Maintaining alignment between observability tools and a moving application landscape requires continuous effort.

**High Cardinality in Prometheus:** While Prometheus is excellent for handling time-series data, it can struggle with **high cardinality**—a situation where too many unique label combinations are generated (e.g., per-user metrics or per-request identifiers). This can strain storage and query performance, making it difficult to scale reliably in complex environments.

## 8.2. Future Improvements

- **Auto-Tuning Detection Thresholds:** One promising area is automating the tuning of anomaly detection thresholds. Rather than relying on static limits or manual tuning, future versions of the system could learn and adjust thresholds based on historical patterns, seasonal behavior, and feedback from alert responses.
- **Incorporating Traces and Spans:** While logs and metrics are powerful on their own, integrating **distributed tracing data** would add a new dimension to detection. Traces can reveal how requests move through systems, uncovering latency bottlenecks or service dependencies that aren't obvious from logs or metrics alone. Including **spans** in the signal set could greatly enhance root cause analysis and anomaly context.
- **AI/ML-Based Auto-Baselining:** Looking ahead, the use of **machine learning models** for dynamic baselining could further improve detection accuracy. These models could continuously learn from data trends, flag subtle deviations, and even anticipate anomalies before they fully materialize. Techniques like online learning or reinforcement learning could help systems adapt to new patterns in near real-time.

## 9. CONCLUSION

This work integrates log analysis with metric-based monitoring to explore a signal-based anomaly detecting mechanism combining the contextual richness of logs with the quantitative strength of time-series data. We have demonstrated by means of the integration of Log Insights and Prometheus how cloud operations teams may transcend conventional threshold-based tactics

and embrace a more flexible, precise, and fast methodology for detecting problems in tough, dynamic environments.

Our major findings show how well this method works: it significantly increases detection accuracy, accelerates reaction times, lowers false positives. Teams that concentrate on the extraction of high-quality signals from raw telemetry data are better ready to understand the narrative underlying anomalies—not only that an issue happened, but also the reasons, places, and strategies engaged in.

The pragmatic impact is pretty remarkable. Operations in clouds span proactive dependability engineering to reactive crisis management. Alerts become ever more crucial; dashboards provide more data; incident investigations are significantly more targeted. Faster recovery times, reduced running noise, and more system resilience based on clouds follow from this.

Observability will be defined going ahead by ever more complex and automated detection pipelines able to mix numerous data streams—including logs, metrics, and traces—react to changes, and continuously learn from patterns. A major simplicity-based innovation in a straightforward approach to boost the reliability and intelligence of cloud computing is signal-based anomaly detection.

## REFERENCES

- [1] Su, Jing, et al. "Large language models for forecasting and anomaly detection: A systematic literature review." *arXiv preprint arXiv:2402.10350* (2024).
- [2] Ati, Clement. "Optimizing high-volume ECG Scan Digitization: A Cloud-Driven Approach Using Terraform and Ansible." (2024).
- [3] Nawaz, Mena, and Jameel Ahmed. "Cloud-based healthcare framework for real-time anomaly detection and classification of 1-D ECG signals." *Plos one* 17.12 (2022): e0279305.
- [4] Zhang, Zhixia, et al. "A many objective-based feature selection model for anomaly detection in cloud environment." *IEEE Access* 8 (2020): 60218-60231.
- [5] Dey, Arunavo, et al. "Signal processing based method for real-time anomaly detection in high-performance computing." *2023 IEEE 47th Annual Computers, Software, and Applications Conference (COMPSAC)*. IEEE, 2023.
- [6] Yang, Chen. "Anomaly network traffic detection algorithm based on information entropy measurement under the cloud computing environment." *Cluster Computing* 22 (2019): 8309-8317.
- [7] Blickensdorff, Johannes, et al. "Production Quality Testing for Automotive Electric Drive Units with AI-Enabled Anomaly Detection based on NVH Data." *DAGA*, 2023.
- [8] Saez, Miguel, et al. "Anomaly detection and productivity analysis for cyber-physical systems in manufacturing." *2017 13th IEEE Conference on Automation Science and Engineering (CASE)*. IEEE, 2017.
- [9] Karmuse, Sachin M., Arun L. Kakhandki, and Mallikarjun Anandhalli. "Cloud based multivariate signal based heart abnormality detection." *Journal of Information and Optimization Sciences* 43.8 (2022): 1935-1952.
- [10] Kayan, Hakan, et al. "Casper: Context-aware iot anomaly detection system for industrial robotic arms." *ACM Transactions on Internet of Things* 5.3 (2024): 1-36.
- [11] Chen, Feiyi, et al. "Learning multi-pattern normalities in the frequency domain for efficient time series anomaly detection." *2024 IEEE 40th International Conference on Data Engineering (ICDE)*. IEEE, 2024.
- [12] Lo Scudo, Fabrizio, et al. "Audio-based anomaly detection on edge devices via self-supervision and spectral analysis." *Journal of Intelligent Information Systems* 61.3 (2023): 765-793.
- [13] Gelli, Ravikumar, and Manimaran Govindarasu. "Anomaly detection and mitigation for wide-area damping control using machine learning." *IEEE Transactions on Smart Grid* 15.6 (2020): 5939-5951.
- [14] Jiang, Jiuchun, et al. "A hybrid signal-based fault diagnosis method for lithium-ion batteries in electric vehicles." *IEEE Access* 9 (2021): 19175-19186.

- [15] Öster, Anders. "A Human Perception View on Holistic Anomaly Detection Systems for Maritime Engine Rooms." *ISOPE International Ocean and Polar Engineering Conference*. ISOPE, 2024.