RESEARCH PAPER

Smart Water Management: Integrating AI and IoT for Real-Time Monitoring of Water Quality and Supply Networks

Dr. Qamar Sultana[1*]

[1*]Associate Professor, Civil Engineering, Muffakham Jah College of Engineering and Technology, Osmania University, Hyderabad

Email ID: qamarsultana4@gmail.com , Orcid ID:0000-0001-8133-9871

Abstract

Smart water distribution networks need smart continuous monitoring to minimize the impact of transient disturbances which conventional sampling strategies cannot easily detect. The current study proposes the AI-IoT solution to the Anomaly Detection in real-time in Multi-node drinking water distribution systems using high frequency sensor readings. A homogenous preprocessing pipeline standardizes the heterogeneous streams of nodes which in turn models temporal features which capture short-term trends and inconsistencies. In statistic strength with weak supervision strategy, no known contamination labels, proxy anomaly ground truth with multi sense node specific thresholds is obtained. The accuracy of the RF classifier trained on chronological validation was 92.19, recall was 89.18 and ROC-AUC was 0.951 which is a high anomaly detector. The variability of pH and ORP contributed to the prediction of turbidity-related characteristics which became the most important. Prioritization- being based on streaming simulation and severity was used to convert the detections into actionable alerts. The framework illustrates how intelligent monitoring (scaled and deployment oriented) may be applied for improving smart water distribution systems in terms of their safety and resilience.

Keywords: Smart Water Management, Internet of Things (IoT), Anomaly Detection, Machine Learning, Water Distribution Networks

How to cite this article: Sultana Q. Smart Water Management: Integrating AI and IoT for Real-Time Monitoring of Water Quality and Supply Networks. Int J Drug Deliv Technol. 2026;16(12s): 265-274. DOI: 10.25258/ijddt.16.12s.28

1. Introduction

The provision of safe and reliable drinking water supply is one of the most urgent infrastructure issues in the world, especially in the environment of the active urbanization process, population increase, climate change, and the aging distribution infrastructure. The drinking water distribution systems (DWDS) are complicated networks that are spatially dispersed and that water that is treated is delivered to consumers via pipes, valves, storage tanks, and pumping stations. Water quality in such networks may vary dynamically because of hydraulic transients, corrosion of pipes, biofilm processes, intrusion processes or operational interventions like flushing, and adjusting the valves. Such upheavals can occur suddenly and spread through parts of the network, endangering the health of the people and regulatory standards. Recent innovations underscore the fact that the integration of machine learning and IoT comes at a time when the concepts are transforming the concept of water quality monitoring by facilitating continuous monitoring, automated analytics, and intelligent decision support at temporal scales commensurate with network characteristics.[1] Nevertheless, even with the advancement in technology, most utilities still use periodic sampling and laboratory-based analysis method, which is sparse in space and time and thus unable to capture transient anomalies.

The conventional monitoring models are usually reliant on manual grab sampling and stationary threshold-based alarms that are built into supervisory control systems. Although these methods can be used to aid compliance reporting, they are ill designed to identify multi-parameter deviations that can indicate an early contamination or hydraulic disturbance. Extensive surveys of IoT-based water quality measurement systems reiterate that fixed threshold detectors frequently produce false positives because of the sensor drift or noise, and at the same time, they cannot generate patterns of correlated values of several parameters.[2] Moreover, traditional systems seldom have built-in predictive intelligence which restricts their capability to foresee the changing risks or prioritize the responses in relation to the severity of anomaly. With the dawn of artificial intelligence in water quality control, researchers emphasize that learnable non-linear relationships between variables and extracting actionable information in streams of high-frequency sensors are required.[3] Complex analyses of embedded IoT platforms also highlight the need to have harmonized data preprocessing, timing compatibility, and computational performance in the implementation of machine learning models in resource-training assets.[4] Real-time IoT implementations in water treatment plants demonstrate the operational value of continuous parameter tracking, particularly for turbidity, pH, residual disinfectant, and oxidation–reduction potential.[5] However, these implementations are often limited to facility boundaries and not made thoroughly into distribution networks, where local disturbances can occur where the treatment is not under control. Machine learning in river and watershed monitoring have been found to have high predictive power on water quality indices, but it is difficult to apply this directly to a distribution system because the sampling frequencies are high, nonstationarity is more considerable, and there is spatial heterogeneity among nodes.[6] Existing smart water monitoring prototypes often emphasize hardware integration and connectivity but devote limited attention

*Author for Correspondence: qamarsultana4@gmail.com

to time-aware model validation or realistic streaming deployment scenarios.[7] Integrated IoT and machine learning frameworks have been proposed for water quality assessment and device monitoring; however, many studies focus on single-node environments or small-scale testbeds, restricting their applicability to large, distributed supply networks.[8] Recent advances in AI-agent coordination within IoT ecosystems further illustrate the potential of distributed intelligent sensing, where multiple nodes collaborate to enhance environmental monitoring fidelity.[9] In spite of such advances, there are still pronounced gaps in solving the problem of the multi-node sensor fusion, labeling of anomalies of a high quality with the limited ground truth, and evaluation in deployment scenarios that reflects the reality of operating a system in real-time.

Out of these constraints, this paper creates an AI-based smart water management model that incorporates the use of multi-node IoT sensors along with machine learning to detect anomalies in real time in drinking water distribution systems. This work is guided by three major objectives. First, to create a multi-node anomaly detection architecture that has the capability of learning coordinated deviations between temperature, turbidity, pH, ORP and conductivity measurements taken at five-minute intervals. Second, to apply a statistically powerful weak-supervision labeling procedure that powers off proxy anomaly ground truth in the absence of known contamination incidences whilst reducing the effects of individual sensor noise. Third, to emulate real-time streaming detection and include prioritization mechanism based on severity which converts model outputs into operational decision support alert levels.

To meet these goals, the proposed framework creates a common preprocessing pipeline to standardize heterogeneous flows of data across various monitoring nodes, align their time-temporal patterns and construct rolling statistical and temporal dependency properties to reflect short-term tendencies, volatility and sudden changes of parameters. Since few verified labels of contamination exist, the node-specific statistical bounds based on the interquartile range and extreme quantile thresholds are used to generate binary labels of anomaly in accordance with a multi-sensor agreement rule. The method gives a possibility of supervised learning and retains methodological transparency in terms of label provenance. The reason to use a Random Forest classifier is its strength with regard to multicollinearity and nonlinear interaction among correlated parameters of water quality. To avoid temporal leakage, model evaluation takes a strictly chronological division to reflect real-world conditions of forecasting the water quality, which is consistent with best practices in predictive water quality modeling.[10]

In addition to the retrospective accuracy measure, the framework has an inference layer that is sequential and, as such, simulates streaming deployment by processing observations in time order and issuing alerts with timestamps. The detection rate and false alert rate are calculated to gauge the possibility of the infrastructure to operate, and it is known that the high recall is a condition of infrastructure safety whereas a large number of false alarms may clog utility operations. To make the alerts more actionable, a severity scoring system that is based on deviation is used to classify the alerts as Low, Medium, and High based on the severity of multi-parameter deviation to node-specific baselines. Such a prioritization allows utilities to spend their inspection resources efficiently and respond quickly to high-risk events and track lower-severity variations to determine persistence.

The research of this work has triple contributions. First, it shows how to provide a scalable approach to the integration of heterogeneous multi-node water quality datasets into a single analysis framework applicable to AI-based monitoring. Second, it presents a statistically based proxy labeling process capable of detecting abnormalities with considerably scarce ground truth. Third, it introduces deployment-based assessment paradigm which mediates between offline classification performance and real-time alert model and severity-based prioritization. Together, these works contribute to the practical implementation of the use of AI and IoT in smart water management, which will help to shift the mode of reactive monitoring to the mode of proactive and data-driven protection of drinking water distribution systems.

2. Literature review

The evolution of smart water management has been strongly influenced by advances in multivariate analytics and high-frequency sensing within distribution networks. Recent work on multivariate functional data analysis combined with machine learning demonstrates that anomaly detection performance improves when temporal dependencies and cross-parameter correlations are explicitly modeled, particularly in dense sensor environments where traditional univariate thresholds are insufficient.[11] As intelligent monitoring systems expand, the demand for transparency and trust has increased, leading to systematic investigations into explainable AI frameworks for water quality monitoring, emphasizing interpretability, accountability, and model validation in safety-critical applications.[12] In parallel, anomaly detection within smart water metering networks has been comprehensively reviewed, highlighting the growing use of supervised and semi-supervised learning for identifying irregular consumption or hydraulic behavior patterns in distributed sensing infrastructures.[13] IoT-based water monitoring systems typically deploy physicochemical sensors measuring parameters such as pH, ORP, turbidity, conductivity, and temperature across distribution nodes. These systems increasingly leverage cloud and edge architectures to enable near real-time processing and scalable storage. However, leakage detection remains a dominant focus in distribution network analytics. Reviews of leakage detection technologies reveal that acoustic, pressure, and flow-based monitoring combined with artificial intelligence can enhance early detection capability, yet many frameworks prioritize hydraulic anomalies over water quality disturbances.[14] Comparative surveys of AIdriven leak detection approaches further demonstrate the diversity of classification algorithms employed, including neural networks, support vector machines, and ensemble methods, but also note variability in validation rigor and scalability.[15] Broader assessments of machine learning applications in smart water distribution systems reinforce the importance of integrating heterogeneous data streams while addressing challenges such as nonstationarity, sensor drift, and computational constraints.[16] High-frequency pressure and acoustic sensing has enabled smart network diagnostics in real-world systems, yet integration with multivariate quality analytics remains limited.[17]

Machine learning for water quality prediction has advanced significantly, particularly in the context of leakage localization and detection frameworks that combine feature extraction with classification strategies for rapid infrastructure diagnostics.[18] Fast detection strategies leveraging optimized learning pipelines demonstrate improved localization accuracy; however, many studies focus on structural health indicators rather than multivariate quality deviations.[19] In broader environmental contexts, self-optimizing machine learning approaches using multi-source remote sensing data have shown promise for water quality monitoring in river systems, underscoring the potential of adaptive modeling strategies in dynamic aquatic environments.[20] Hybrid machine learning and data mining algorithms have also been applied to water quality indexing, demonstrating improved predictive performance compared to single-model approaches, yet these frameworks often rely on labeled datasets and static evaluation paradigms.[21]

Even though there has been significant progress, there are still areas of significant gaps. Numerous studies focus on the deployment of single nodes or monitor facilities, which is not generalizable to multi-node distribution networks. The simulation of streams and time-aware validation are often not under focus, which limits the understanding of feasibility of deployment in real-time. Moreover, although the problem of anomaly detection is broadly examined, there is a lack of literature that implements statistically supported weak supervision to overcome the lack of validated labeling of contamination. Mechanisms of severity-based prioritization that convert the detection of anomalies into risk stratification that can be acted upon are also understudied. These weaknesses point to the necessity of scalable multi-node architectures incorporating robust proxy labeling, time and operational based prioritization into intelligent water quality monitoring parameters.

3. Methodology

3.1 Dataset Description

3.1.1 Data Source

The suggested framework was created based on a publicly available multi-node water quality dataset, which was retrieved from the Mendeley Data repository.[22] The dataset was collected using two distinct monitoring systems deployed across four different locations and connection types within a drinking water distribution setting. Specifically, the YSI EXO 3 monitoring system was installed at three locations and connected to fire hydrants, while the ATI MetriNet Q52 system was installed on a service line within an academic building. The study area corresponds to a monitored drinking water distribution setting comprising four deployment contexts, including three fire-hydrant-connected locations and one service-line location within an academic building. The observation periods ranged from several weeks to several months, depending on the dataset.

The dataset is comprised of measurements from five independent monitoring nodes that are installed in a drinking water distribution network. In the present study, the five separate datasets were treated as five monitoring nodes within a unified analytical framework. A node is a unique sensing point, and this allows the supply infrastructure to be monitored spatially. The complete repository is divided into five separate datasets in .csv format, labeled as dataset_#, with each dataset containing labeled columns and parameter values in the appropriate measurement units. Additional supplementary files, including maintenance logs for the YSI and ATI sensors, are also provided in the repository. The sensor measurements were sampled at constant 5 minutes intervals, which gave high-resolution time series data suitable in time series modeling and real-time monitoring simulation. With the combination of all the five node datasets into a single analytical framework, the study simulates the centralized smart water management platform that can process distributed streams of IoT sensors. The multi-node design can be used to both achieve the spatial heterogeneity in water behavior and to enable the model to scale to larger network applications.

3.1.2 Water Quality Parameters

The original repository includes multiple water quality parameters, including temperature, conductivity, specific conductivity, turbidity, fDOM, pH, ORP, and combined chlorine. However, in order to make nodes consistent, only parameters shared by all datasets were kept for analysis. The last variables were Temperature (°C), Oxidation Reduction Potential (ORP, mV), pH, Turbidity, and Electrical Conductivity. These parameters indicate essential physicochemical parameters that are commonly used in the water quality monitoring. Temperature affects reaction dynamics, ORP is the measure of oxidative stability, pH is the measure of chemical equilibrium, turbidity measures the disturbances of particulate matter, and conductivity is the measure of dissolved ionic content. Together, they provide a full picture of the conditions of distribution networks.

3.2 Data Preprocessing

3.2.1 Column Standardization

Small differences found in column names were addressed to achieve a common schema between all the datasets. Before merging all parameter names were standardized. Each dataset was given a unique Node_ID in order to maintain a traceability and spatial identity.

This enabled processing of nodes and centralized model development.

3.2.2 Missing Value Handling

The missing values have been addressed by node-wise linear interpolation to ensure that imputation was only applied on a per sensing location basis without cross-node contamination. Following interpolation, the stream of nodes was resampled to a frequency of 5 minutes in order to achieve consistency in time. In order to prevent over-smoothing, no post-resampling interpolation was done on gaps longer than 30 minutes (six consecutive intervals). Records with greater discontinuities were discarded in order to maintain reliability.

3.2.3 Time-Series Alignment

Before resampling, duplicate timestamps using nodes were removed. The purged node-wise data sets were then combined and sorted in time sequence by time and node number. This provided both good temporal sequencing and supported the later chronological modelling.

3.3 Feature Engineering

3.3.1 Rolling Statistical Features

To capture short-term water quality dynamics, rolling statistical features were derived with the help of a 30-minute sliding window, equivalent to 6 observations. For each of the monitored parameters, rolling mean and rolling standard deviation values were computed. The rolling mean was indicative of trend behavior for the local area and the rolling standard deviation indicated short-term variability. These features helped the model to identify departures from normal operating patterns.

3.3.2 Temporal Dependency Features

The lag-1 features were developed in order to add instantaneous history in the feature space. Along with this, first-order derivatives were computed in order to measure how the successive observations change. These characteristics enhanced the capacity of the model to identify sudden spikes, sudden changes and rapidly changing conditions relating to anomalous events.

3.3.3 Multi-Sensor Fusion Strategy

All the engineered variables were compiled into one feature matrix. This enables the model to be trained on cross-parameter classification and identify grouped deviations instead of detecting single threshold violations. Since significant anomalies in drinking water systems can be observed to have multiple physicochemical parameters simultaneously, this method enhanced resistance to single-sensor noise and detection of patterns of complex anomalies.

3.4 Statistical Anomaly Labeling (Weak Supervision)

3.4.1 Node-Specific Robust Thresholding

were used as interquartile range (IQR) fences of each parameter and node. Extreme quantile bounds were also calculated to minimize the trends of skewed distributions. The more stringent of the two limits was chosen as the final limit so that anomaly detection was conservative. Q1 − 1.5×IQR and Q3 + 1.5×IQR were used as interquartile range (IQR) fences of each parameter and node. Extreme quantile bounds were also calculated to minimize the trends of skewed distributions. The more stringent of the two limits was chosen as the final limit so that anomaly detection was conservative.

3.4.2 Multi-Sensor Agreement Rule

An anomalous label was given only when two or more parameters crossed their specific thresholds for each node simultaneously. This guideline minimized the false alarms due to peaks in sensor rectifications and guaranteed that the abnormalities were as a result of cross-dimensional abnormalities. The binary labels obtained were subsequently supervised learned.

3.4.3 Justification of Proxy Ground Truth

These labels are not the verified cases of contamination but they are principled approximations to the statistically significant deviations in limited data realities. This proxy labeling model made it possible to train under supervision but make it transparent and reproducible.

3.5 Machine Learning Model Development

A Random Forest classifier was chosen because it can be used to capture nonlinear relationships, it can handle multicollinearity among correlated sensor features and does not seem to be sensitive to noise and outliers. In order to prevent temporal leakage, the collected dataset was chronologically split, with the first 80% of observations for training the model and the remaining 20% for model testing. This design is more realistic of actual forecasting conditions. The model was implemented using 200 decision trees, and a fixed random state to make it easy to reproduce. All engineered features were kept so that the ensemble could learn the importance of features automatically.

3.6 Model Evaluation

Model performance was evaluated with Accuracy, Precision, Recall and F1-score. to get an unbiased evaluation on model classification performance. Due to the moderate degree of class imbalance in the dataset, both ROC-AUC and Precision-Recall (PR) curves are also analyzed, where the accuracy of anomaly detection can be understood more informatively by the analysis of PR. Besides that, the confusion matrix was applied to measure the true positives, false positives, false negatives, and true negatives hence to indicate trade-offs in operations between missed detections and false alerts.

3.7 Real-Time Monitoring Simulation and Severity-Based Prioritization

To study the applicability in the context of operation, a sequential inference system was employed to simulate IoT deployment in time within the distribution network. Observations were processed in chronological order, and whenever an anomaly was detected by the classifier, a node specific and time stamped alert was generated. Detection rate and false alert rate were used as the operational feasibility. In order to make the outputs of the anomaly more actionable, IQR-normalized deviations of node-parameter medians were used to obtain severity scores. Alerts were then ranked into Low, Medium or High, so that response efforts can be prioritized more effectively.

4. Results

4.1 Dataset Characteristics

The final dataset was the consolidated dataset that consisted of five uniformly spaced, 5-minute monitoring nodes, and was the result of the preprocessing and weakly supervised labeling. The total aberration percentage was just around 6.48, which confirmed that there was moderate imbalance of classes to be used in the model of anomaly detection. Temporal aggregation found out that the anomalies were not equally distributed throughout the monitoring horizon but instead they were concentrated in small clusters. This intermittent action is seen to be in line with local hydraulic variations, operational ups and downs or temporary contamination type features as opposed to global system-wide degradation.

Post-spatial analysis of the distribution showed that there was heterogeneity among nodes. The anomaly density of Nodes 4 and 5 was greater than that of Nodes 1-3 indicating the variability of operational conditions between the segments of supply. Figure 1 shows the daily count of proxy anomalies and expected alerts where the model-based alerts track the time trend of identified anomalies.

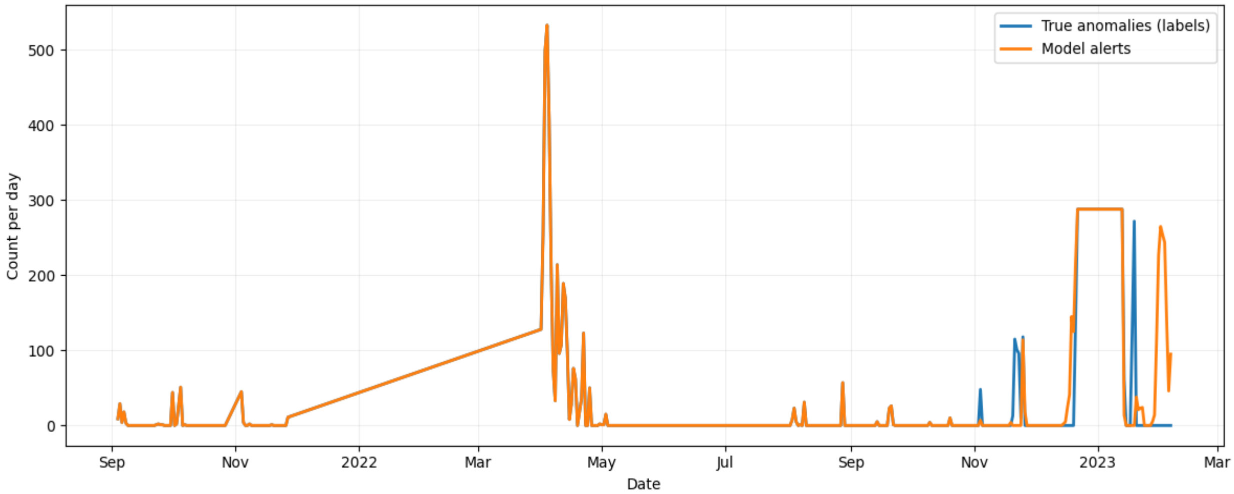

Figure 1. Network-Level Real-Time Monitoring: Daily True Anomalies vs Model Alerts.

Figure 1 indicates that there were sharp peaks around the beginning of April 2022 and the beginning of December 2022-January 2023 when the number of anomalies was growing rapidly. These peaks are in high agreement with the predicted alerts, which shows that there is temporal compatibility between the model output and ground-truth proxy labels. Excessive caution on very large spikes is an indication of a conservative sensitivity that is beneficial in water safety systems where false alarms that are moderate may not be as serious as the miss of an event.

4.2 Classification Performance

A chronological test split consisting of the last 20% of observations was used to assess model performance so as to avoid time leakage. The performance measures were summarized in Table 1.

Table 1. Classification Performance on Time-Aware Test Set

Metric	Value
Accuracy	0.9219
Precision	0.7789
Recall	0.8918
F1-score	0.8315
ROC-AUC	0.9510

The classifier obtained a high accuracy of 92.19, which shows that the classifier has got good overall discrimination ability between normal and anomalous states. The accuracy of 77.89% indicates proper control of false alarms, whereas the recall of 89.18% indicates that it is sensitive to anomalous events. The value of 0.8315 in the F1-score verifies the balanced predictive behavior. The ROC-AUC value 0.9510 has a high level of separability which is a confirmation that the model is able to differentiate between abnormal patterns and normal operational fluctuations with high degree of reliability. The confusion matrix in Figure 2 gives more information about the classification behavior, and it measures detection and misclassification results.

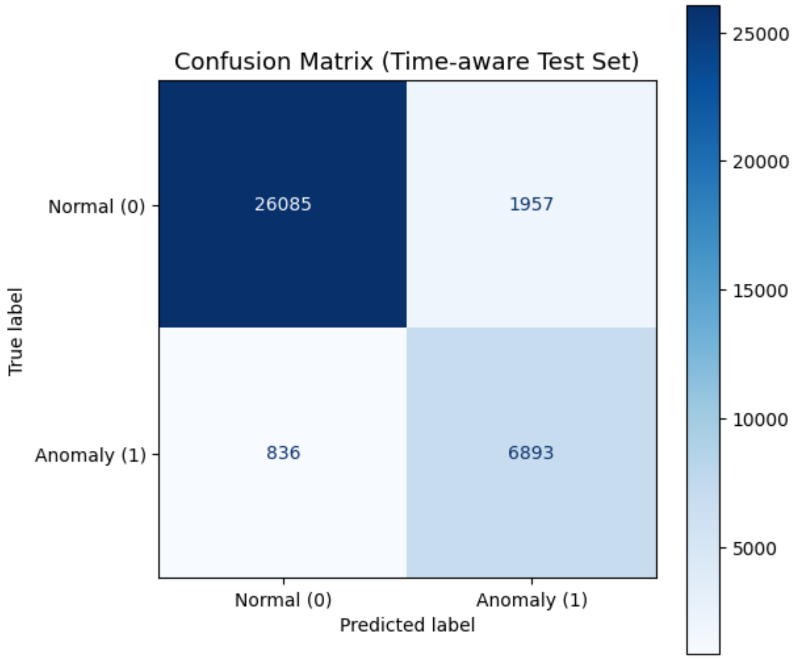

Figure 2. Confusion Matrix (Time-Aware Test Set).

Figure 2 indicates that there were 6,893 correctly identified anomalies (true positives) and 836 anomalies that were missed (false negatives). False positives are 1,957, which is against 26,085 true negatives. The fact that the number of false negatives is relatively low proves the high level of anomaly detection, which is an essential factor in protecting infrastructure. There are false positives, but their percentage is operationally manageable as compared to overall observations, and it facilitates the practical implementation of deploying it.

4.3 Feature Importance Analysis

The analysis of feature importance based on the Random Forest model showed that the features based on turbidity were the most prominent predictors. Rolling mean, rolling standard deviation and lag based turbidity features had the greatest contributions to anomaly discrimination. This observation is in line with the expectations of the domain because turbidity is very sensitive to intrusion of particles, mobilization of sediments, and hydraulic disruptions.

The pH and ORP features including rolling variability, first-order derivative components were found to contribute to secondary contributions. The parameters indicate the chemical instability and oxidative variances that tend to come with the disturbance of quality. Conductivity and temperature characteristics did not have a relatively high level of significance, which means that the extreme shifts in ionic or thermal conductivity were not as common anomaly drivers in the analyzed data. The real impact of turbidity processes is demonstrated in Figure 3, where the changes in turbidity are shown with the identified alerts.

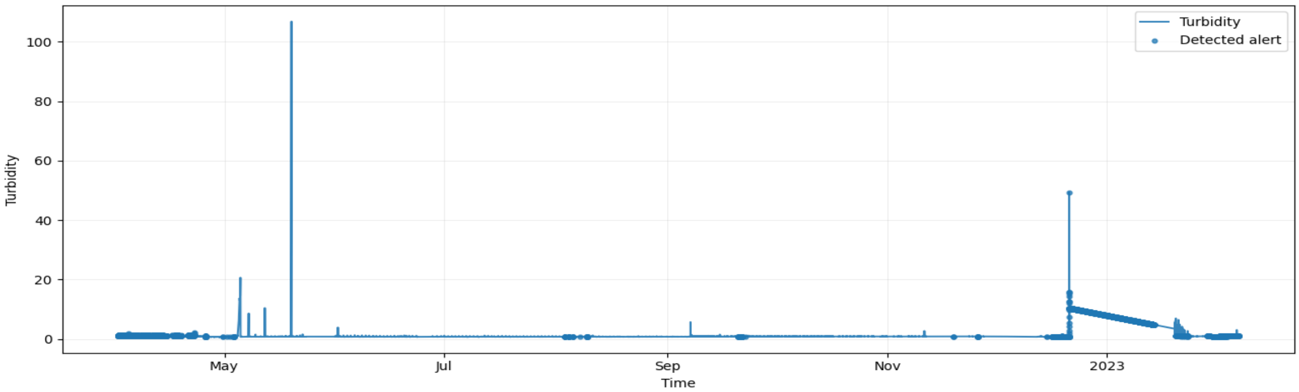

Figure 3. Real-Time Timeline: Turbidity Monitoring with Detected Alerts (Node_4).

Figure 3 demonstrates that significant turbidity deviations are in direct proportion to groups of forecasted alerts. A sharp spike of more than 100 NTU corresponds to extreme alert production, which proves that the model is responsive to extreme disturbances of particulate matter. Moderate levels of turbidity in the

*Author for Correspondence: qamarsultana4@gmail.com

beginning of 2023 also generate corresponding proportional alerts. The temporal consistency confirms the rank of importance and proves that anomaly detection is a physically consistent difference and not a random noise.

4.4 Real-Time Simulation Outcomes

The operational feasibility was assessed by a sequential streaming simulation which took into account the observations in order of time. Some of the key monitoring indicators based on this simulation are summarized in Table 2.

Table 2. Real-Time Monitoring Performance Indicators

Indicator	**Value**
Total Alerts Raised	8,850
True Anomalies	7,729
Detection Rate	0.8918
False Alert Rate	0.0630

An 89.18% detection rate establishes consistent anomaly recognition with 89.18% under simulated deployment, which is similar to test-set recall. The false alert of 6.3 percent in comparison to the total observations is also acceptable alert burden to continuous monitoring settings. The intensity of spatial alerts at the nodes is depicted in Figure 4 and indicates the variability across nodes and a local disturbance pattern.

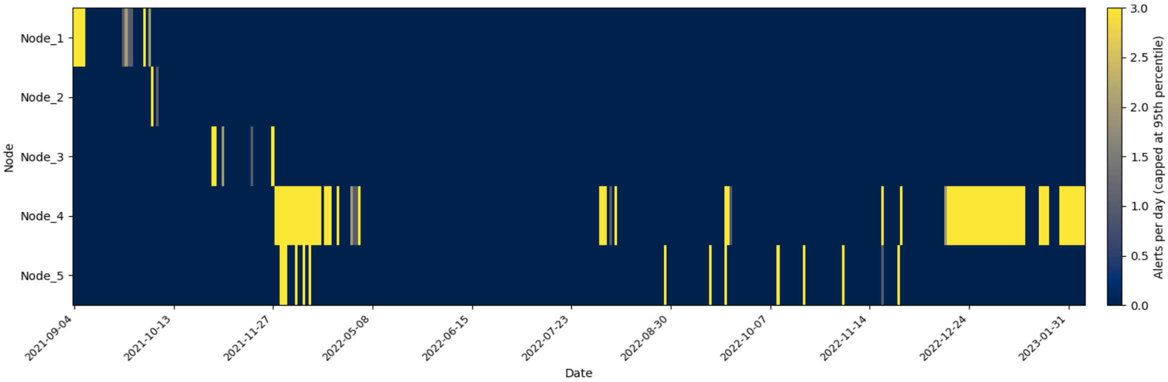

Figure 4. Supply Network Heatmap: Daily Alert Intensity (Robust-Scaled).

Figure 4 indicates that there were high-intensity alert clusters sustained at Node_4 in April 2022 and late 2022, whilst Nodes 1, 2 and 3 have sparse and intermittent activity. The node 5 has temporary spikes and no long-term clustering. This is because the lack of coordinated high-intensity notifications on all nodes confirms that the system does not generate artificial network-wide triggering and instead records localized anomaly behaviour.

4.5 Node-Wise Monitoring Performance

Comparisons between cross-nodes indicate that cross-node comparison is consistent across anomaly density variations. The nodes that are high in activity produce proportionately more alerts whilst retaining the same sensitivity, which means that the model has generalization properties across the heterogeneous spatial conditions. The temporal responsiveness at peak periods of disturbance is also explained in Figure 5.

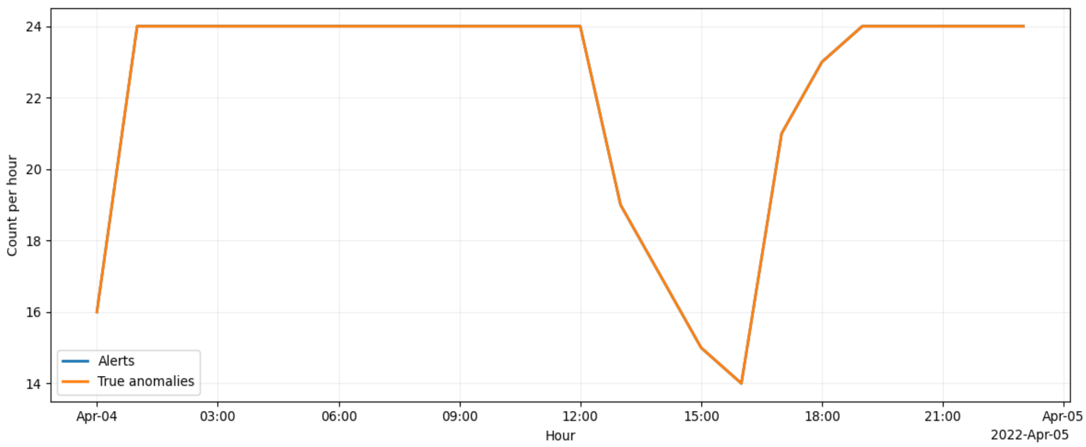

Figure 5. High-Impact Day Drilldown: Hourly Alerts vs Anomalies (2022-04-04).

Figure 5 demonstrates that hourly proxy anomalies and predicted alerts are close to each other on a high-impact day. The number of alerts monitors the variation of anomalies during the 24 hours with minimum lags. Minor distortions in the periods of transition are indicators of threshold sensitivity and not systematic misclassification. The almost parallel behavior proves to be appropriate in real-time implementation in dynamic operational conditions.

4.6 Severity Distribution of Alerts

IQR-normalized deviation severity classification shows that about 52, 33 and 15 percent of the alerts were classified as Low, Medium, and High severity, respectively. The multiple-parameter excursions with the turbidity, pH, and ORP deviations were mostly linked to high-severity alerts. Such stratification allows using high-severity events as a priority to examine incidents immediately and follow low-severity notifications to identify persistence. The distribution shows that the framework separates the marginal threshold exceedances and the significant physicochemical disturbances, which contributes to the decision support of smart water monitoring systems.

5. Discussion

The results prove that multi-sensor, time-sensitive anomaly detection offers a powerful model to track drinking water distribution systems with the realistic constrained data. The obtained recall of 0.8918 means that most of the statistically defined anomalous states have been detected and the precision of 0.7789 means that the rate of false alarm with safety-critical infrastructure is manageable. The temporal correlation of predicted alerts and clustered anomaly regions is an indication that the classifier identified coordinated multi-parameter anomalies, not individual anomalies. This justifies the fact that distribution network disturbances tend to be multivariate, meaning they comprise coupled modifications in turbidity, pH, ORP, and other indicators as opposed to unit sensor threshold violations.[3]

One of the core reasons that make performance possible is the introduction of temporal feature engineering. The model had rolling mean and rolling standard deviation features with lag and first-order derivative components, which helped the model to identify gradual drift, as well as abrupt spikes. These artificial descriptors introduce historical context (in the short term) into every prediction and they are less sensitive to short-term noises. The importance of time-series features in nonlinear dependence and enhancing generalization across time-varying operational conditions is highlighted in prior studies in predictive water quality modeling.[10,23] The prevalence of turbidity-related characteristics, which is facilitated by the variation of pH and ORP, is compatible with the physical interpretation of network perturbations, which supports the idea that the obtained decision-boundary is based on significant physicochemical behavior as opposed to statistical features.

The proposed framework has obvious benefits compared to the traditional monitoring based on threshold. The systems are non-dynamic, and each parameter is treated separately and based on pre-determined limits that might not be able to adjust to node-specific baselines or seasonal changes. Multi-Sensor Agreement rule and learned ensemble model, on the contrary, incorporate correlated signals which renders the model less sensitive to sensor drift or single spikes. This strategy is indicative of new trends in the multivariate detection of anomalies, as computational and machine learning techniques are showing better performance in complex environmental systems than univariate thresholding.[11] Moreover, chronological train-test split makes sure that estimates of performance are indicative of actual deployment environments with no optimistic bias induced by random partitioning. It has been found that time-aware validation is a key to reliable anomaly detecting with smart infrastructure networks.[13] This practical relevance is especially important in water supply distribution systems, where infrastructure behavior and monitoring requirements are closely linked to network management and service reliability.[24]

The combination of severity based prioritization mechanism enhances the operational significance of the framework. Through triage on the statistical detection, the system converts the statistical level of alerts into triage: Low, Medium, and High, defined by the magnitude of deviation. This is in line with the demands of explainable and interpretable AI systems which offer clear explanatory routes and risk levels instead of binary choices.[12] Practically, this prioritization facilitates the effective resource allocation of the inspection process, especially in the large urban supply grids where the alerts can be received in different nodes at the same time. The visualization of the node-level alert intensity in the form of a heatmap also allows operators to differentiate local disturbances and network-wide anomalies, which increases the situational awareness in smart city water infrastructures.[16]

Although the results are promising, there are a number of limitations that should be considered. To begin with, anomaly labels were created using statistically robust proxy, as opposed to confirmed cases of contamination. As a result, performance measures are used to measure detection of statistically determined deviations, as opposed to confirmed health events. Second, it is possible that the imbalance in the density of node data is biased against nodes with larger historical records, and thus there is a limitation on sensitivity to rare behavior in nodes with sparse sampling. Third, the streaming assessment is an offline model and does not consider real-world limitations like communication delay, sensor failure, cyber-physical interference, or response operator feedback. Lastly, the classification relied on a fixed probability threshold; this was effective in this study, but there are other cost-sensitive thresholding methods that can be more effective based on the operational priorities.

This framework can be expanded in a number of ways in future research. The edge computing architectures integration could also minimize the latency and increase resilience through the local inference of sensor nodes. Optimization of threshold based on cost-sensitive or risk-sensitive criteria would match the sensitivity of detection with utility-sensitive response capabilities. Deep learning models based on sequences, e.g. LSTM or Transformer architecture, can potentially capture longer-term temporal dependencies than the existing window-based features. Lastly, causal interpretation with datasets of proven contamination or controlled disturbance events would enhance causal interpretation and provide the connection between anomaly detection by statistics and a proven operational event of contamination, which would make AI-based water quality management systems more reliable and scalable.

6. Conclusion

This paper introduces a unified AI-IoT system to detect anomalies in the drinking water network of the city in real-time, which proves the possibility of integrating multi-node sensor fusion, time-conscious machine learning, and deployment-focused evaluation as a part of a single monitoring structure. The proposed system successfully predicted coordinated physicochemical deviations in turbidity, pH, ORP, conductivity, and temperature by synchronizing heterogeneous data sets of five distributed nodes and developing temporal-informed features. The loosely guided labeling approach allowed the supervised learning with a restricted ground truth, and chronological validation allowed the realistic performance evaluation. Findings have shown that overall classification accuracy and strong recall of anomalous events is high, which proves the ability of the system to reduce the number of missed disturbances in safety-critical settings. The analysis of the importance of features revealed that turbidity was the main indicator of an anomaly, with the support of chemical variability defined by the dynamics of pH and ORP. Streaming alert simulation and severity-based prioritization also improve the operational applicability of the model by converting model results into decision support. Together, they can further the cause of intelligent water monitoring toward enhancing the ability to deliver early warnings, reinforcing sustainable intelligent infrastructure management, and offering a scalable framework of AI-based analytics to distribution systems in the real world.

References:

1. Essamlali I, Nhaila H, El Khaili M. Advances in machine learning and IoT for water quality monitoring: A comprehensive review. Heliyon. 2024 Mar 30;10(6).
2. Dharmarathne G, Abekoon AM, Bogahawaththa M, Alawatugoda J, Meddage DP. A review of machine learning and internet-of-things on the water quality assessment: Methods, applications and future trends. Results in Engineering. 2025 Jun 1;26:105182.
3. Zou S, Ju H, Zhang J. Water quality management in the age of AI: Applications, challenges, and prospects. Water. 2025 May 28;17(11):1641.
4. Vicente EC, Silva LA, da Rocha Fernandes AM, Parreira WD. A structured review of IoT-based embedded systems and machine learning for water quality monitoring. Applied Sciences. 2025 Nov 3;15(21):11719.
5. Forhad HM, Uddin MR, Chakrovorty RS, Ruhul AM, Faruk HM, Kamruzzaman S, Sharmin N, Jamal AS, Haque MM, Morshed AM. IoT based real-time water quality monitoring system in water treatment plants (WTPs). Heliyon. 2024 Dec 15;10(23).
6. Cojbasic S, Dmitrasinovic S, Kostic M, Turk Sekulic M, Radonic J, Dodig A, Stojkovic M. Application of machine learning in river water quality management: a review. Water Science & Technology. 2023 Nov 1;88(9):2297-308.
7. Hemdan EE, Essa YM, Shouman M, El-Sayed A, Moustafa AN. An efficient IoT based smart water quality monitoring system. Multimedia tools and applications. 2023 Aug;82(19):28827-51.
8. Bhardwaj A, Dagar V, Khan MO, Aggarwal A, Alvarado R, Kumar M, Irfan M, Proshad R. Smart IoT and machine learning-based framework for water quality assessment and device component monitoring. Environmental Science and Pollution Research. 2022 Jun;29(30):46018-36.
9. Miller T, Durlik I, Kostecka E, Kozlovska P, Łobodzińska A, Sokołowska S, Nowy A. Integrating artificial intelligence agents with the internet of things for enhanced environmental monitoring: applications in water quality and climate data. Electronics. 2025 Feb 11;14(4):696.
10. Ahmed AN, Othman FB, Afan HA, Ibrahim RK, Fai CM, Hossain MS, Ehteram M, Elshafie A. Machine learning methods for better water quality prediction. Journal of Hydrology. 2019 Nov 1;578:124084.
11. Rigueira X, Olivieri D, Araujo M, Saavedra A, Pazo M. Multivariate functional data analysis and machine learning methods for anomaly detection in water quality sensor data. Environmental Modelling & Software. 2025 May 30;190:106443.
12. Aderemi IA, Kehinde TO, Okwor UD, Ahmad KH, Adjei KY, Ekechi CC. Explainable AI for Water Quality Monitoring: A Systematic Review of Transparency Interpretability and Trust. IEEE Sensors Reviews. 2025;2:419-43.
13. Kanyama MN, Shava FB, Gamundani AM, Hartmann A. Machine learning applications for anomaly detection in Smart Water Metering Networks: A systematic review. Physics and Chemistry of the Earth, Parts A/B/C. 2024 Jun 1;134:103558.
14. Islam MR, Azam S, Shanmugam B, Mathur D. A review on current technologies and future direction of water leakage detection in water distribution network. IEEE access. 2022 Oct 6;10:107177-201.
15. Kammoun M, Kammoun A, Abid M. Leak detection methods in water distribution networks: a comparative survey on artificial intelligence applications. Journal of Pipeline Systems Engineering and Practice. 2022 Aug 1;13(3):04022024.

16. Taloma RJ, Cuomo F, Comminiello D, Pisani P. Machine learning for smart water distribution systems: exploring applications, challenges and future perspectives. Artificial Intelligence Review. 2025 Jan 31;58(4):120.
17. Rousso BZ, Lambert M, Gong J. Smart water networks: A systematic review of applications using high-frequency pressure and acoustic sensors in real water distribution systems. Journal of Cleaner Production. 2023 Jul 15;410:137193.
18. Fan X, Yu X. An innovative machine learning based framework for water distribution network leakage detection and localization. Structural Health Monitoring. 2022 Jul;21(4):1626-44.
19. Fan X, Zhang X, Yu X. Machine learning model and strategy for fast and accurate detection of leaks in water supply network. Journal of Infrastructure Preservation and Resilience. 2021 Apr 15;2(1):10.
20. Chen P, Wang B, Wu Y, Wang Q, Huang Z, Wang C. Urban river water quality monitoring based on self-optimizing machine learning method using multi-source remote sensing data. Ecological Indicators. 2023 Feb 1;146:109750.
21. Aslam B, Maqsoom A, Cheema AH, Ullah F, Alharbi A, Imran M. Water quality management using hybrid machine learning and data mining algorithms: An indexing approach. IEEE Access. 2022 Nov 10;10:1
22. Sela L, Trimble T, Vosburgh E, Ling J, Katz L, Kinney K, Zigler C, Werth C. Dataset of multiple water quality parameters measured in a drinking water distribution system. Mendeley Data. 2023;V1. doi:10.17632/wwy39n8p7z.1.
23. Sultana Q. Prediction of ground water quality index using artificial neural networks. published in Science and Engineering Journal. 2020 Aug;24(8):283-95.
24. Sultana A, Sultana Q. Design of water supply distribution system: a case study. International Journal of Scientific Research and Review. 2019 Jun;7(6):435-53.