www.ijarp.com ISSN 2456-9992 Page: 01- 14

International Journal Advanced Research Publications

Research Article

Volume: 02
Issue: 03

AQUA-SENSE: AN IOT-DRIVEN HEALTH SURVEILLANCE ARCHITECTURE FOR REAL-TIME WATER QUALITY MONITORING AND PREDICTIVE EPIDEMIOLOGICAL ALERTING

***[1]Antony Alex S, [2]Dr. Rajakumar G,[3] Subhashini Kiruthika S,[4] Snofa D,[5] Arishvanth Sriganesh M N,[6] and Suguna A**

[1,3,4,5,6,]Department of Electronics and Communication Engineering Francis Xavier Engineering College Tirunelveli - 627 003, Tamil Nadu, India.

[2]Professor & COE - Department of Electronics and Communication Francis Xavier Engineering College Tirunelveli - 627 003, Tamil Nadu, India.

Article Received: 27 January 2026, Article Revised: 15 February 2026, Published on: 07 March 2026

***Corresponding Author: Antony Alex S**

Department of Electronics and Communication Engineering Francis Xavier Engineering College Tirunelveli - 627 003, Tamil Nadu, India.

DOI: https://doi-doi.org/101555/ijarp.4837

ABSTRACT

KEYWORDS: Internet of Things (IoT), Water Quality Monitoring, Machine Learning, Disease Prediction, Public Health Surveillance, Random Forest Classifier, Real-Time Alerting Systems.

INTRODUCTION

Water is an indispensable resource for human survival, yet the consumption of contaminated water remains a primary catalyst for severe public health crises worldwide. Traditional water quality management has historically relied on manual sampling followed by rigorous, laboratory-based chemical and biological analyses. While these conventional methodologies are highly accurate, they are fundamentally flawed when applied to dynamic environments due to their inherent latency. By the time laboratory results confirm the presence of harmful pathogens, the affected population has often already consumed the hazardous water, transforming a preventable risk into an active epidemiological outbreak. In recent years, the proliferation of the Internet of Things (IoT) has initiated a paradigm shift in environmental monitoring, enabling the continuous acquisition of physicochemical parameters such as Total Dissolved Solids (TDS), The proposed architecture is further fortified by a high-performance

Node.js middleware layer utilizing WebSocket integration for sub-second, bidirectional data transmission, ensuring that local health authorities are viewing genuinely real-time analytics. To combat the pervasive issue of spam generation, the system features a highly optimized active alerting mechanism equipped with strict cooldown logic, delivering prioritized and actionable intelligence directly to stakeholders. Ultimately, Aqua-Sense bridges the gap between raw environmental engineering and predictive healthcare, empowering administrators to preemptively intercept waterborne disease outbreaks before mass consumption occurs.

Ph, and Turbidity. However, a critical review of Contemporary IoT architectures reveals that the majority remain functionally passive. They operate primarily as sophisticated data loggers, transmitting raw metrics to static dashboards and placing the cognitive burden entirely on human operators to deduce the underlying health implications. Furthermore, when automated alerts are integrated into these systems, they frequently rely on rigid, static thresholds. This primitive approach triggers an overwhelming volume of redundant notifications during minor, natural sensor fluctuations, inevitably leading to administrative "alert fatigue" and delayed emergency responses.

To address these critical operational gaps, this paper introduces "Aqua-Sense," an intelligent, end-to-end health surveillance system designed to transition water management from a state of passive observation to proactive disease prevention. By establishing a robust hardware perception layer utilizing microcontrollers and specialized sensor arrays, the system continuously tracks key water metrics at the source. Rather than settling for a generic water quality index or a binary "potable versus non-potable" classification, Aqua-Sense incorporates a dynamic Random Forest Machine Learning classifier. This computational intelligence layer is uniquely trained to map complex intersections of TDS, pH, and Turbidity to calculate the real-time probability of five specific pathogenic threats: Cholera, Typhoid, Hepatitis A, Dysentery, and Diarrheal infections.

Related Work

The pursuit of safeguarding public health through environmental engineering has increasingly relied on the integration of digital technologies. Recent literature highlights significant advancements across sensor networks, predictive analytics, and automated notification systems, though critical operational gaps remain in translating raw data into actionable medical intelligence.

A. The Evolution of IoT in Environmental Monitoring

The deployment of Internet of Things (IoT) architectures has precipitated a major shift in ecological data acquisition, moving the industry away from delayed, manual laboratory sampling toward continuous, automated edge surveillance. Contemporary research frequently showcases architectures built upon microcontrollers, such as Arduino and ESP platforms, which interface with physicochemical sensors to capture real-time metrics. These fundamental setups typically rely on standard protocols like HTTP or MQTT to push analog signals to centralized, third-party cloud platforms for visualization. While these hardware ecosystems successfully digitize the physical environment, a pervasive limitation across current paradigms is their functional passivity. Existing IoT networks treat the successful transmission and graphical representation of data as the ultimate goal, forcing human administrators to manually interpret the epidemiological severity of fluctuating sensor metrics. There is a pronounced need for architectures that utilize advanced, low-latency streams (such as WebSocket) not just to log data, but to feed an autonomous computational layer.

B. Machine Learning Applications and Predictive Bottlenecks

To decipher the complex patterns hidden within environmental datasets, researchers have increasingly adopted Machine Learning (ML), demonstrating a strong preference for supervised algorithms like Support Vector Machines (SVM), Regression models, and Random Forest (RF) classifiers. Recent studies underscore the efficacy of these models in handling volatile sensor data, including missing value imputation and basic parameter evaluation. However, a critical analysis reveals that the predictive scope of current ML implementations remains severely restricted. The prevailing methodology in recent literature is confined either to calculating a generalized Water Quality Index (WQI) or executing rigid binary classifications, simply labeling water as "Potable" or "Non-Potable". Furthermore, these models are overwhelmingly static; they are trained strictly offline using historical datasets and lack the infrastructure for dynamic, on-the-fly retraining at the deployment edge. The academic landscape currently lacks robust, multi-class predictive frameworks capable of translating raw physicochemical anomalies (such as specific intersections of pH, Turbidity, and TDS) into precise, localized probabilities for distinct pathogens, such as identifying the specific environmental signature of a Cholera or Typhoid outbreak. dispatches an immediate notification. Because aquatic environments are inherently dynamic and sensors naturally experience minor noise, this rigid approach routinely generates an overwhelming volume of redundant warnings, inducing severe "alert fatigue" among administrators. The literature demonstrates a distinct absence of intelligent throttling mechanisms or mathematically driven "cooldown" logic within notification pipelines. Additionally, alerts typically transmit only raw data values, placing the burden of medical deduction on the recipient. Modern platforms require prioritized, multi-tiered alerting engines capable of distinguishing between routine hardware maintenance and critical epidemiological emergencies, ensuring actionable intelligence is delivered without desensitizing the end user.

D. Research Gap and Proposed Contribution

Synthesizing the current literature establishes a clear technological void that the Aqua-Sense architecture addresses. By replacing monolithic third-party dashboards with a sophisticated hybrid infrastructure that pairs NoSQL databases with localized flat-file logging, the proposed system enables seamless, on-demand ML retraining. It transcends binary potability assessments by employing a dynamic RF classifier to generate real-time, multi-class disease probabilities. Finally, it entirely mitigates the pervasive issue of alert fatigue through the engineering of an intelligent notification engine governed by strict, parameter-specific cooldown algorithms, transforming a conventional data logger into a resilient, proactive public health instrument.

C. Vulnerabilities in Smart Public Health Alerting

The ultimate efficacy of any health surveillance framework hinges on its ability to rapidly and reliably notify stakeholders before an environmental anomaly escalates into a public health crisis. A review of modern automated notification systems exposes significant vulnerabilities in alert management. Most contemporary industrial models rely on primitive, rule-based threshold triggers; if a raw parameter marginally exceeds a static limit, the system.

Proposed System Architecture

The architectural foundation of Aqua-Sense is engineered to overcome the latency and functional passivity typical of conventional environmental monitoring setups. To guarantee rapid data acquisition, resilient storage, and scalable predictive capabilities, the proposed framework abandons the traditional linear edge-to-cloud pipeline. Instead, it employs a highly modular, multi-tiered hybrid architecture divided into four deeply integrated tiers: the Perception Layer, the Middleware Layer, the Computational Intelligence Layer, and the Application Layer.

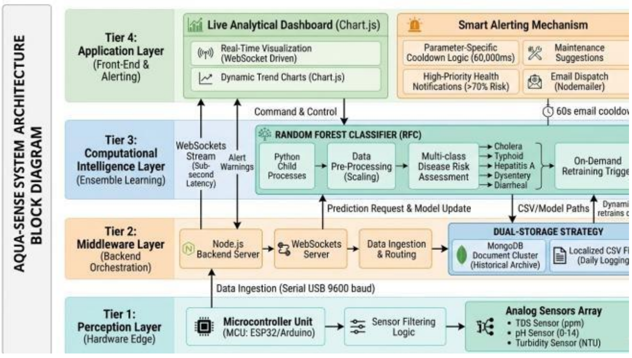

Fig 1: System Architecture Diagram.

D. The Hardware Perception Layer

Operating at the physical edge, the perception layer is anchored by a robust microcontroller unit (utilizing the Arduino or ESP32 architecture) that serves as the primary hub for data acquisition. This microcontroller interfaces directly with a specialized array of physicochemical sensors designed to capture vital environmental metrics. A Total Dissolved Solids (TDS) sensor measures the concentration of dissolved inorganic and organic matter in parts per million (ppm), serving as a primary indicator of general water purity. Simultaneously, a pH sensor evaluates the hydrogen-ion activity on a 0–14 scale to detect acidic or alkaline anomalies that could foster pathogen survival. Finally, a Turbidity sensor assesses optical clarity in Nephelometric Turbidity Units (NTU), tracking suspended particulates that frequently indicate microbial contamination. The microcontroller performs localized signal filtering before formatting these readings into unified data strings and transmitting them to the central server via a stable USB serial connection at 9600 baud.

E. Middleware Orchestration and Dual- Storage Strategy

The Middleware Layer acts as the central orchestrator of the system, powered by an asynchronous, event-driven Node.js backend. To eliminate the inherent network overhead and latency associated with traditional HTTP polling, this layer natively integrates the WebSocket protocol. This establishes a persistent, full-duplex communication channel that instantly broadcasts serialized payload data to the frontend Application Layer, achieving true sub-second latency for live analytical visualizations.

To ensure both robust data persistence and optimized algorithmic performance, the framework implements a unique dual-database storage strategy. For scalable historical archiving and complex query aggregations, the system utilizes MongoDB, a flexible NoSQL document database. Concurrently, the backend executes a highly localized flat-file logging protocol, appending real-time records into daily Comma- Separated Values (CSV) files. This structural decoupling allows the machine learning scripts to natively and rapidly ingest the CSV data without relying on external network requests to the database cluster.

F. Computational Intelligence and Machine Learning

The Computational Intelligence Layer isolates intensive analytical workloads from the main backend event loop by utilizing independent Python subprocesses. Because the raw hardware data operates on vastly different numerical scales, an automated pre-processing pipeline employs a standard scaler to normalize the features to unit variance, ensuring equitable algorithmic weighting.

The core of this layer is a Random Forest (RF) classifier, chosen for its resilience against overfitting and its capacity to map complex, non- linear environmental datasets. Escaping the limitations of binary potability assessments, the RF model evaluates the normalized TDS, pH, and Turbidity arrays to compute granular, real-time probability percentages for five distinct public health threats: Cholera, Typhoid, Hepatitis A, Dysentery, and Diarrheal infections. Furthermore, the architecture supports dynamic, on-demand intelligence; the backend can trigger the Python scripts to ingest the most recent localized CSV logs, allowing the ensemble model to continuously refine its decision trees based on evolving environmental conditions.

G. Active Alerting and Cooldown Mechanism

The true operational value of the Application

Layer lies in its ability to translate these predictive metrics into actionable interventions without inducing administrative alert fatigue. The system features an active notification engine that logically separates anomalies into routine maintenance suggestions (e.g., severe pH fluctuations) and critical health alerts (triggered exclusively when the ML model calculates a specific disease risk exceeding a severe 70% threshold).

To prevent the inbox flooding frequently caused by minor sensor noise, the architecture executes a strict mathematical throttling algorithm. The backend maintains an active memory registry that records the exact timestamp an alert is dispatched for a specific parameter or disease. If a subsequent anomaly is detected within a predefined 60,000- millisecond (one-minute) window, the system suppresses the redundant notification. This intelligent cooldown logic guarantees that health authorities receive prioritized, immediate warnings during genuine epidemiological emergencies while maintaining the integrity and usability of their communication channels.

II. Experimental Results and Performance Evaluation

To rigorously evaluate the operational efficacy of the Aqua-Sense architecture, the system was deployed in a controlled testing environment. The evaluation was designed to simulate a wide spectrum of environmental conditions, ranging from baseline potable water metrics to severe, multi-parameter contamination events. The performance assessment focused on three critical dimensions: the latency of real-time data visualization, the accuracy of the multi-class predictive model, and the functional resilience of the smart alerting mechanism.

A. Real-Time Data Acquisition and Dashboard Latency:

The primary objective of the middleware layer was to eliminate the inherent delays associated with traditional HTTP polling. Empirical observation during the testing phase confirmed that the integration of the WebSocket protocol successfully facilitated persistent, full-duplex communication between the edge hardware and the application layer.

Fig 2: Live data Dashboard and Graphs.

As illustrated in **Figure 2**, the glassmorphism-styled dashboard renders incoming physicochemical metrics with true sub-second latency. Under baseline testing, the interface accurately reflected stable parameters, such as a safe Total Dissolved Solids (TDS) level of 480.0 ppm and a neutral pH of 7.8, while dynamically flagging a slight turbidity elevation (6.1 NTU) as a early warning. Furthermore, the system demonstrated robust hardware fault tolerance; when the physical serial connection was deliberately interrupted, the UI instantly propagated a critical "Sensors Disconnected" alert, ensuring administrators were immediately aware of the surveillance blind spot.

B. Machine Learning Multi-Class Prediction:

A core contribution of the Aqua-Sense framework is its transition from binary potability classification to specific epidemiological risk assessment. To validate the Random Forest (RF) classifier, the system was subjected to simulated extreme contamination data (e.g., highly acidic pH combined with extreme turbidity).

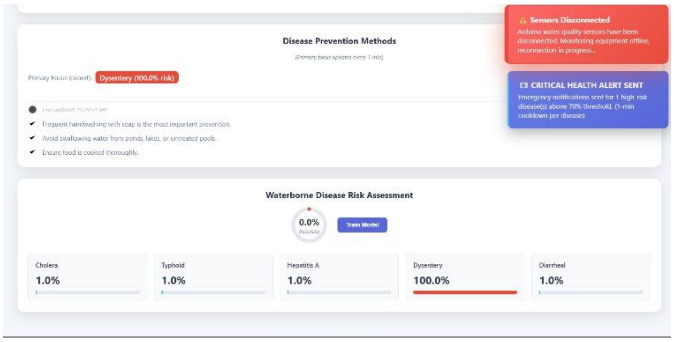

Fig 3: Image of the Prediction Section.

Unlike conventional models that would simply output a generic "Non-Portable" flag, the RF algorithm successfully mapped these severe nonlinear anomalies to distinct pathogen profiles. As demonstrated in **Figure 3**, during an induced high- turbidity and acidic event, the computational intelligence layer accurately isolated the environmental signature associated with Dysentery, computing a definitive 100.0% risk probability. Concurrently, the model correctly suppressed the probability metrics for unrelated pathogens like Cholera and Typhoid. This granular, multi-class output proves the system's capability to deliver precise, targeted intelligence rather than generalized panic.

C. Efficacy of the Smart Alerting and Cooldown Engine:

The most pervasive issue in modern environmental IoT networks is "alert fatigue" caused by continuous, redundant notifications during prolonged contamination events. Aqua-Sense’s active alerting mechanism was tested to ensure it could deliver critical warnings without overwhelming administrative channels.

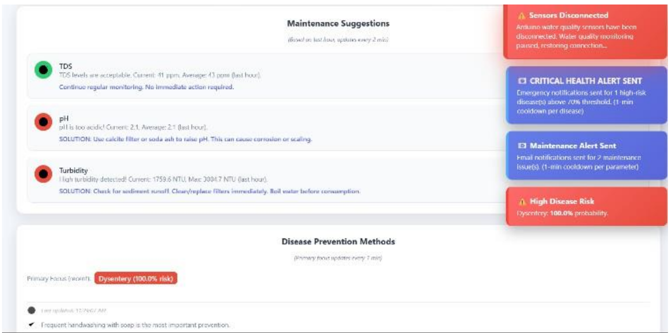

Fig 4: Image of the Suggestion Section.

The local UI successfully generates context-aware maintenance suggestions, dynamically calculating hourly averages and max values to recommend specific corrective actions—such as utilizing calcite filters for severe acidic drops (**Figure 4**).

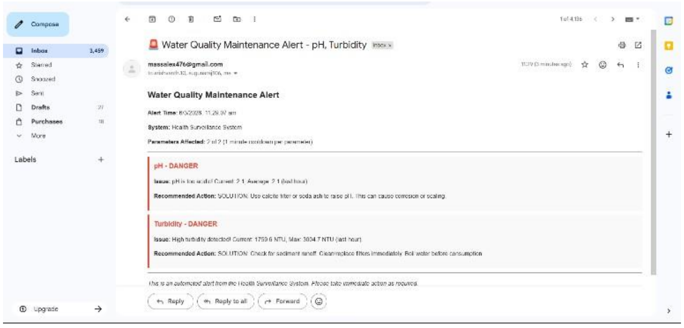

Fig 5: Alert Notification through Email Crucially, the external email dispatch system.

Operated exactly as engineered. When the ML model flagged the critical Dysentery risk alongside the dangerous pH (2.1) and Turbidity (1759.6 NTU) deviations, the Node.js backend instantly formatted and dispatched an HTML-rich warning to predefined stakeholders (**Figure 5**). During continuous simulated anomaly streaming, the backend's active memory registry successfully executed the strict 60,000-millisecond cooldown logic. Redundant triggers were silently suppressed internally, proving that the system can sustain an ongoing public health emergency by delivering prioritized, actionable intelligence while entirely mitigating inbox spam.

III.Discussion and Practical Implications

The empirical findings of this study underscore a fundamental transition in how environmental monitoring can—and should—be approached in the modern era. While traditional IoT architectures have successfully digitized the collection of water quality metrics, they have largely remained passive observers, leaving the critical task of epidemiological deduction to human administrators. The successful deployment of the Aqua-Sense framework demonstrates that integrating localized computational intelligence directly into the data pipeline can effectively close this gap.

A. Interpreting System Capabilities:

The implementation of the Random Forest classifier proved highly effective in processing multi- dimensional sensor data. Instead of relying on rigid, one-dimensional thresholds to declare water simply "safe" or "unsafe", the ensemble model successfully learned the complex, non-linear environmental signatures associated with specific pathogens. The system's demonstrated ability to accurately isolate a Dysentery threat from a convergence of high turbidity and extreme acidity proves that predictive modeling can provide granular, actionable intelligence. Furthermore, the WebSocket-driven middleware layer successfully eliminated the latency bottlenecks inherent in standard HTTP polling, ensuring that the analytical dashboard reflects the true, real-time state of the water source.

Equally significant is the validation of the system's smart alerting mechanism. "Alert fatigue" is a well- documented vulnerability in industrial and healthcare monitoring, where operators become desensitized to constant, redundant alarms triggered by minor sensor noise. By enforcing a strict 60,000- millisecond cooldown logic per specific parameter and disease threat, Aqua-Sense mathematically guarantees that stakeholders receive prioritized warnings without suffering from inbox flooding. This transforms the notification engine from a potential nuisance into a reliable emergency broadcast tool.

B. Practical Implications for Public Health:

From a practical standpoint, the Aqua-Sense architecture offers a highly scalable blueprint for municipal water management and rural health administration. In developing regions or disaster- stricken areas where laboratory infrastructure is severely lacking, the deployment of this low-cost, edge-computed system can serve as a first line of defense. Because the hardware relies on accessible ESP32/Arduino microcontrollers and standard physicochemical sensors, the financial barrier to implementation is significantly lower than establishing traditional testing facilities.

The most profound implication of this research, however, lies in its potential impact on public health policy. Historically, the management of waterborne illnesses—such as Cholera, Typhoid, and Hepatitis A—has been reactive; medical interventions are typically mobilized only after patients begin presenting symptoms at local clinics. Aqua-Sense introduces a preemptive paradigm. By instantly notifying authorities when the environmental preconditions for an outbreak reach a critical 70% probability , local governments can execute immediate, targeted interventions. Authorities can issue localized boil-water advisories, deploy emergency chemical treatments, or halt specific distribution nodes hours or days before the contaminated water is widely consumed. Ultimately, transitioning from passive measurement to active, predictive surveillance possesses the tangible potential to save lives and significantly reduce the economic burden of preventable epidemics.

CONCLUSION AND FUTURE WORK

C. Conclusion

The development of the Aqua-Sense framework marks a critical departure from traditional, reactive water quality management methodologies. By successfully bridging the gap between raw physicochemical sensor data and actionable epidemiological intelligence, this system translates environmental monitoring into a proactive, life- saving public health tool. The integration of a Random Forest machine learning classifier allows the architecture to move far beyond simplistic binary potability assessments, accurately computing granular, real-time risk probabilities for severe pathogenic threats, specifically Cholera, Typhoid, Hepatitis A, Dysentery, and Diarrheal infections.

Furthermore, the implementation of a highly resilient, event-driven middleware layer utilizing WebSocket ensures sub-second data visualization, completely eliminating the network latency inherent in legacy polling methods. Crucially, the engineering of a smart alerting mechanism—governed by strict, parameter-specific cooldown logic—successfully mitigates the pervasive industry issue of administrative alert fatigue. By seamlessly combining edge-level hardware, hybrid database management, and autonomous predictive analytics, Aqua-Sense evolves the standard IoT data logger into a fully automated, highly scalable surveillance platform capable of preempting waterborne outbreaks before mass consumption occurs.

D. Future Research Directions

While the current iteration of the Aqua-Sense architecture demonstrates robust predictive and operational capabilities, several viable avenues for future research and system enhancement remain.

First, the physical perception layer can be significantly expanded. Integrating more advanced, specialized sensors—such as those designed to detect Dissolved Oxygen (DO), residual chlorine, heavy metals, or direct biological fluorescence— would supply the computational layer with a richer feature set, thereby increasing the diagnostic accuracy and scope of the machine learning model.

Second, the system's deployment architecture can be optimized for distributed, off-grid applications.

Transitioning the edge communication framework from a tethered serial interface to Low-Power Wide- Area Networks (LPWAN), such as LoRaWAN or NB-IoT, would facilitate the deployment of a vast, interconnected mesh of sensor nodes across entire municipal water grids or remote rural rivers. This hardware evolution would naturally necessitate the integration of renewable energy harvesting, such as compact solar panels and advanced battery management protocols, to ensure continuous, autonomous operation.

Finally, future computational frameworks should explore the integration of advanced neural networks or Federated Learning techniques. By allowing a decentralized network of Aqua-Sense nodes to collaboratively train a shared predictive model at the edge, the system could continuously adapt to geographically diverse environmental anomalies without the need to centralize massive volumes of raw data, further solidifying its role as a next- generation public health infrastructure.

ACKNOWLEDGEMENT

The authors would like to express their profound gratitude to the management and faculty of Francis Xavier Engineering College for providing the technical infrastructure, laboratory resources, and an encouraging academic environment that made the development of the Aqua-Sense architecture possible.

We also wish to acknowledge the collaborative efforts of the entire project team. The successful integration of the hybrid Node.js middleware, the Random Forest predictive models, and the real-time alerting mechanisms was a deeply collective endeavor, relying on countless hours of testing and calibration.

Finally, we extend our sincere appreciation to the supporting platforms and initiatives—such as the Smart India Hackathon framework—that provided the competitive motivation and the critical problem- statement focus necessary to innovate within the public health and environmental engineering sectors.

REFERENCES

1. Wiryasaputra, Rita & Huang, Chin-Yin & Lin, Yu- Ju & Yang, Chao-Tung. (2024). An IoT Real-Time Potable Water Quality Monitoring and Prediction Model Based on Cloud Computing Architecture. Sensors. 24. 1180. 10.3390/s24041180.
2. Hilary, Anabi & Alsharif, Mohammed & Anya, Anya & Bonet, Mathias & Uyi, Aiyudubie & Uthansakul, Peerapong & Jamel, Nebhen & Aly, Ayman. (2021). Design and Implementation of a Low-Cost Portable Water Quality Monitoring System. Computers, Materials & Continua. 69. 2405-2424. 10.32604/cmc.2021.018686.
3. Singh, Yurav & Walingo, Tom. (2024). Smart Water Quality Monitoring with IoT Wireless Sensor Networks. Sensors. 24. 2871. 10.3390/s24092871.
4. Rahman, Md.Mahbubur & Bepery, Chinmay & Hossain, Mohammad & Hassan, Zahid & Hossain, G M & Islam, Md. (2020). Internet of Things (IoT) Based Water Quality Monitoring System. 2. 168- 180.
5. Ren, Jianxun & Zhu, Qiliang & Wang, Changsheng. (2022). Edge Computing for Water

Quality Monitoring Systems. Mobile Information Systems. 2022. 1-7. 10.1155/2022/5056606.

6. Dragulinescu, Ana & Balaceanu, Cristina & Osiac, Filip & Roscaneanu, Roxana & Chedea, Veronica & Suciu, George & Paun, Mirel & Bucuci, Stefania. (2021). IoT-based Smart Water Management Systems. 51-56. 10.1109/SIITME53254.2021.9663611.
7. Dallaf, Amina. (2025). Edge Computing in IoT Networks: Enhancing Efficiency, Reducing Latency, and Improving Scalability. International Journal of Advanced Network, Monitoring and Controls. 10. 2025. 10.2478/ijanmc-2025-0009.
8. Hong, Wong & Shamsuddin, Norazanita & Abas, Pg Emeroylariffion & Apong, Rosyzie & Masri, Zarifi & Suhaimi, Hazwani & Gödeke, Stefan & Noh, Muhammad. (2021). Water Quality Monitoring with Arduino Based Sensors. Environments. 8. 6. 10.3390/environments8010006.
9. Purohit, Manish & Vyas, Archana & Bora, Avnish. (2025). Energy-Efficient Water Quality Monitoring Using Wireless Sensor Networks with Public Health Integration. Journal of Applied Bioanalysis. 299-320. 10.53555/jab.v11si7.2179.
10. Sharma, Astha & Sharma, Richa & Rana, Rishi & Kalia, Anshul. (2024). Water quality prediction using Machine Learning Models. E3S Web of Conferences. 596. 10.1051/e3sconf/202459601025.
11. Ahmad, Abu Fuad & Alshammari, Khaznah & Ahmed, Istiaque & Sayed, MD. (2024). Machine Learning for Missing Value Imputation. 10.48550/arXiv.2410.08308.
12. Ahmed, Umair & Mumtaz, Rafia & Anwar, Hirra & Shah, Asad & Irfan, Rabia & García-Nieto, José. (2019). Efficient Water Quality Prediction Using Supervised Machine Learning. Water. 11. 2210. 10.3390/w11112210.
13. Mesiar, Radko & Sheikhi, Ayyub. (2021). Nonlinear Random Forest Classification, a Copula- Based Approach. Applied Sciences. 11. 7140. 10.3390/app11157140.
14. Aljarah, Fatin & Çetin, Aydın. (2023). Prediction of Water Quality with Ensemble Learning Algorithms. Advances in Artificial Intelligence Research. 3. 36-44. 10.54569/aair.1200695.
15. Thakker, Zalak & Harshadkant, Buch. (2024). Effect of Feature Scaling Pre-processing Techniques on Machine Learning Algorithms to Predict Particulate Matter Concentration for Gandhinagar, Gujarat, India. International Journal of Scientific Research in Science and Technology. 410-419. 10.32628/IJSRST52411150.
16. B., Priya & Sharma, Geetanjali & Bongale, Anupkumar & Dharrao, Deepak & Seitshiro, Modisane. (2025). Exploratory Data Analysis and Water Potability Classification using

Supervised Machine Learning Algorithms. Engineering, Technology & Applied Science Research. 15. 20898-20903. 10.48084/etasr.8904.

17. Halgamuge, Malka & Niyato, Dusit. (2024). Adaptive edge security framework for dynamic IoTsecurity policies in diverse environments. Computers & Security. 148. 104128. 10.1016/j.cose.2024.104128.
18. Dogo, Eustace & Nwulu, Nnamdi & Twala, Bhekisipho & Aigbavboa, Clinton. (2021). Accessing Imbalance Learning Using Dynamic Selection Approach in Water Quality Anomaly Detection. Symmetry. 13. 818.10.3390/sym13050818.
19. Pandya, Suhagbhai. (2021). Predictive Analytics in Smart Grids: Leveraging Machine Learning for Renewable Energy Sources. International Journal of Current Engineering and Technology. 11. 677-683. 10.14741/ijcet/v.11.6.12.
20. Nwabor, Ozioma Forstinus & Nnamonu, Emmanuel & Martins, Paul & Christiana, Ani. (2016). Water and Waterborne Diseases: A Review. International Journal of TROPICAL DISEASE & Health. 12. 1-14. 10.9734/IJTDH/2016/21895.
21. Diallo, Abdoul Rahime & Homri, Lazhar & Boeuf, Thomas & Dantan, Jean-Yves & Bonnet, Frédéric. (2025). Quantifying and mitigating alarm fatigue caused by fault detection systems. Reliability Engineering & System Safety. 267. 111890. 10.1016/j.ress.2025.111890.
22. Ezeh, Chinedu & Anioke, Sandra & Oyewole, Sola & David, Munakur. (2024). The role of predictive analytics in enhancing public health surveillance: Proactive and data-driven interventions. World Journal of Advanced Research and Reviews. 24. 3059-3077. 10.30574/wjarr.2024.24.3.3909.
23. Nandy, Manish & Ramesh, Hajare. (2024). Design of Iot-Enabled Early Warning System For Health Disease Diagnosis. South Eastern European Journal of Public Health. 338-343. 10.70135/seejph.vi.857.
24. Nguyen, Thuy & Kim, Chaelin & Goucher, Gerard & Kim, Jong-Hoon. (2023). Associations of Water Quality with Cholera in Case-Control Studies: A Systematic Review and Meta-Analysis. 10.1101/2023.09.06.23295113.