



iJRASET

International Journal For Research in
Applied Science and Engineering Technology



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** X **Month of publication:** October 2026

DOI: <https://doi.org/10.22214/ijraset.2026.85107>

www.ijraset.com

Call: ☎ 08813907089

E-mail ID: ijraset@gmail.com

Aqua Sense AI: AI-Powered Water Quality Intelligence Platform

K. Swarupa¹, M. Lalith², Muskan³, L. Raneeth⁴

^{1,3,4,5} Undergraduate Students, Hyderabad Institute of Technology and Management, Hyderabad, Telangana, India

² Associate Professor, Department of CSE, Hyderabad Institute of Technology and Management, Hyderabad, Telangana, India

Abstract: Water quality assessment is essential for determining whether water is suitable for human consumption and enabling informed decision-making regarding water resources. Aqua Sense AI is an AI-powered water quality intelligence platform developed using Python and Streamlit. The system provides a preliminary classification of water samples as Safe (Potable) or Unsafe (Non-Potable) based on nine physicochemical parameters: pH, Hardness, Total Dissolved Solids (TDS), Chloramines, Sulfate, Conductivity, Organic Carbon, Trihalomethanes, and Turbidity.

A Random Forest Classifier is employed for the classification task, delivering real-time predictions along with prediction confidence scores, a composite Water Quality Score (0–100), and risk level assessments (Low, Medium, or High). Parameter values are entered through an intuitive interactive web form.

The platform features an Analytics Dashboard powered by interactive Plotly visualizations, Explainable Artificial Intelligence (XAI) providing feature importance rankings and natural language reasoning, dynamic parameter-specific water treatment recommendations, and session-based prediction history tracking. Designed in alignment with UN Sustainable Development Goal 6 (Clean Water and Sanitation), the system serves as a preliminary decision-support tool and does not replace certified laboratory testing. Overall, the work demonstrates an integrated approach that combines machine learning, explainability, interactive data analytics, and actionable guidance within a unified water quality assessment platform.

Keywords: Water Quality, Machine Learning, Random Forest, Water Potability, Streamlit, Explainable Artificial Intelligence (XAI), Data Analytics, Decision Support System, Sustainable Development Goal 6 (SDG 6).

I. INTRODUCTION

Water is a fundamental natural resource vital for human survival, ecosystem stability, agricultural productivity, and global socio-economic progress. Despite its critical importance, rapid urbanization, industrial discharge, agricultural runoff, and climate-induced stressors have severely exacerbated water contamination and depleted potable freshwater reserves worldwide. According to global environmental bodies such as the World Health Organization (WHO) and the United States Environmental Protection Agency (EPA), ensuring water safety requires rigorous monitoring across multiple physicochemical and microbiological parameters. However, conventional water quality assessment protocols typically depend on centralized laboratory testing, which—while precise—frequently entails substantial turnaround times, high operating costs, and raw statistical readouts that require specialized biochemical expertise to interpret. The challenge of modern water management lies in the multidimensional and non-linear interactions among water quality metrics. An evaluation of water potability cannot rely on isolated threshold comparisons alone; parameters such as pH, Hardness, Total Dissolved Solids (TDS), Chloramines, Sulfate, Electrical Conductivity, Total Organic Carbon (TOC), Trihalomethanes (THMs), and Turbidity exhibit complex interdependencies. For instance, elevated turbidity may shield harmful microorganisms from disinfection, while excessive organic carbon in the presence of chlorine fosters carcinogenic disinfection byproducts such as trihalomethanes. For local water administrators, environmental officers, and end consumers lacking specialized training in aquatic chemistry, translating raw quantitative measurements into an immediate, actionable safety assessment presents a major operational bottleneck.

To bridge this operational and interpretive gap, machine learning (ML) has emerged as an indispensable paradigm in computational environmental analytics. By leveraging historical monitoring benchmarks, supervised classification algorithms can identify non-linear decision boundaries and capture subtle multi-parameter dependencies that standard linear indices often overlook. Among various machine learning architectures, the Random Forest Classifier is particularly suited for tabular environmental data. As an ensemble learning method constructed from numerous decorrelated decision trees, Random Forest delivers superior classification accuracy, exhibits resilience against noise and variance, effectively suppresses overfitting, and provides inherent probabilistic confidence scores along with quantitative feature importance metrics [1].

Motivated by these operational challenges and technological opportunities, Aqua Sense AI was designed and developed as an integrated, AI-driven water intelligence platform built using Python and the Streamlit framework. The platform provides an end-to-end analytical pipeline that evaluates water samples across nine critical physicochemical indicators. Through a user-friendly, responsive interface, the system accepts empirical testing parameters and applies a trained Random Forest model to yield real-time potability classifications (Safe vs. Unsafe). Beyond binary classification, the architecture integrates a mathematical evaluation engine that calculates a composite Water Quality Score (0–100) and stratifies potential risks into categorical levels (Low, Medium, and High). A critical barrier to deploying artificial intelligence in public health and environmental decision-making is the "black-box" nature of complex models. To ensure transparency, accountability, and user trust, Aqua Safe integrates Explainable Artificial Intelligence (XAI) principles. By computing parameter importance rankings and providing natural-language contextual explanations, the system elucidates the primary chemical factors driving each prediction. Furthermore, the application features an interactive Analytics Dashboard powered by Plotly for trend and distribution exploration, dynamic water treatment and remediation recommendations tailored to detected anomalies, and structured session-based reporting mechanisms.

The conceptual and functional framework of Aqua Sense AI directly aligns with United Nations Sustainable Development Goal 6 (SDG 6: Clean Water and Sanitation), specifically supporting Target 6.1 (universal and equitable access to safe drinking water) and Target 6.3 (improving water quality through pollution reduction and safe management) [2]. The system is formulated not as a replacement for certified laboratory instrumentation, but as an accessible, rapid, and explainable decision-support tool capable of democratizing environmental intelligence.

Key Contributions:

The primary contributions of this work are summarized as follows:

- 1) **Development of an End-to-End AI-Powered Assessment Platform:** Designed and implemented Aqua Sense AI, a unified, multi-page web platform utilizing Python and Streamlit that provides real-time, preliminary water potability assessment across nine critical physicochemical indicators (pH, Hardness, TDS, Chloramines, Sulfate, Conductivity, Organic Carbon, Trihalomethanes, and Turbidity).
- 2) **Accurate and Robust Machine Learning Classification:** Employed an ensemble Random Forest Classifier trained to model complex, non-linear relationships among chemical water parameters. The model delivers rapid potability classification (<100 ms inference time) with high predictive reliability (~92% estimated accuracy) and class-probability-based confidence scoring.
- 3) **Multi-Dimensional Quality Scoring and Risk Stratification:** Extended standard binary classification by integrating a custom evaluation algorithm in `utils.py`. The system computes a continuous, normalized Water Quality Score (0–100) and stratifies samples into categorical Risk Levels (Low, Medium, High) based on international drinking water standards (WHO and EPA).
- 4) **Incorporation of Explainable Artificial Intelligence (XAI):** Addressed the "black-box" nature of ensemble learning by incorporating feature-importance rankings and automated natural-language explanations. This enables both technical and non-technical stakeholders to identify the specific chemical drivers and contaminants responsible for an "Unsafe" verdict.
- 5) **Rule-Based Remediation and Decision Support Engine:** Engineered an automated recommendation module that translates detected parameter anomalies into actionable, domain-specific mitigation steps—including immediate emergency precautions, long-term filtration/treatment protocols, and health impact advisories.
- 6) **Interactive Visual Analytics and Session-Based Reporting:** Constructed an interactive Analytics Dashboard using Plotly to deliver dynamic visual insights (Safe vs. Unsafe distributions, parameter histograms, and quality gauge metrics) alongside session-state-driven prediction history management and report generation.
- 7) **Empirical Alignment with Sustainable Development Goal 6 (SDG 6):** Demonstrated a practical, cost-effective computational framework supporting UN SDG 6 (Targets 6.1 and 6.3), fostering accessible and decentralized water quality monitoring for community water managers, researchers, and public health practitioners.

II. LITERATURE REVIEW

Computational water quality assessment intersects environmental chemistry, supervised machine learning, explainable artificial intelligence (XAI), and interactive environmental decision-support systems.

A. Physicochemical Water Quality Indices and Deterministic Modeling

Traditional water potability evaluation relies heavily on deterministic Water Quality Index (WQI) formulations, originally introduced by Horton [3] and subsequently refined by Brown et al. [4] and the Canadian Council of Ministers of the Environment (CCME) [5]. Conventional WQI methodologies assign empirical weights to individual parameters—such as pH, Total Dissolved Solids (TDS), Turbidity, and Hardness—aggregating them through sub-index arithmetic or geometric equations to generate a single composite score. However, environmental researchers have highlighted substantial structural limitations in traditional WQI frameworks. Tyagi et al. [6] observed that rigid aggregation functions frequently suffer from "eclipsing" (where a critically toxic concentration of a single parameter is masked by optimal values in others) and "ambiguity" (where boundary thresholds produce inconsistent safety categorizations). Furthermore, deterministic formulations struggle to account for non-linear chemical interactions—such as the relationship between organic carbon, free chlorine, and the formation of carcinogenic trihalomethanes (THMs). These limitations emphasize the necessity of modern computational approaches capable of modeling high-dimensional, non-linear dependencies across multiple water characteristics simultaneously.

B. Machine Learning Architectures for Water Potability Classification

To overcome the constraints of deterministic indexing, researchers have increasingly deployed supervised machine learning paradigms for environmental monitoring. Early comparative studies evaluated statistical classifiers including Logistic Regression, Naive Bayes, and Support Vector Machines (SVMs) [7]. While SVMs map chemical observations into high-dimensional hyperspaces to establish optimal separating hyperplanes, they remain computationally intensive on large tabular benchmarks and exhibit sensitivity to hyperparameter tuning and noise. Recent literature demonstrates the superior empirical performance of tree-based ensemble methods in environmental and hydrological classification tasks. Ahmed et al. [8] evaluated diverse classification architectures across benchmark water quality datasets and established that the **Random Forest** algorithm consistently outclassed individual decision trees, k -Nearest Neighbors (k -NN), and Multilayer Perceptrons (MLPs), achieving classification accuracies in excess of 90%. By constructing an ensemble of decorrelated decision trees through bootstrap aggregation (bagging) and random feature subspace selection, Random Forest effectively suppresses model variance, resists overfitting on complex chemical data, handles collinearity among water parameters, and inherently produces probabilistic confidence estimates [1]. Consequently, Random Forest represents an optimal algorithmic choice for robust, tabular water potability screening.

C. Explainable Artificial Intelligence (XAI) and Environmental Decision Support

Despite high classification accuracy, the pervasive deployment of complex machine learning models in public health and environmental engineering is frequently constrained by their "black-box" nature. When a model classifies a water source as "Unsafe," public health officials, environmental technicians, and community stakeholders require transparent justifications to determine appropriate remediation strategies. To bridge this trust gap, contemporary research integrates Explainable Artificial Intelligence (XAI) frameworks into predictive environmental systems. Lundberg and Lee [9] introduced SHapley Additive exPlanations (SHAP), a game-theoretic approach that computes the additive marginal contribution of each input variable to an individual prediction. Studies by Wang et al. [10] confirm that providing explicit feature importance rankings—identifying whether elevated sulfate, extreme pH deviation, or high turbidity governed a model's classification—significantly enhances user trust, regulatory compliance, and diagnostic speed. Furthermore, embedding these explainable models within lightweight, interactive web dashboards (e.g., Streamlit) bridges the divide between theoretical computational models and practical, real-time decision support aligned with United Nations Sustainable Development Goal 6 (SDG 6) [2].

III. RESEARCH GAP AND RESEARCH QUESTIONS

A. Research Gap

The reviewed literature indicates a significant gap between theoretical machine-learning models and practical, deployable decision-support systems for water quality assessment. While numerous studies have demonstrated high classification accuracy using supervised algorithms, existing solutions frequently suffer from three primary limitations:

- 1) **Black-Box Opacity:** High-performing predictive models are rarely paired with Explainable Artificial Intelligence (XAI) mechanisms, leaving end users and water managers unable to identify which specific chemical contaminants triggered an "Unsafe" classification.
- 2) **Rigid Binary Classification:** Most research frames water potability as a strict binary outcome (Safe vs. Unsafe), failing to provide a continuous, multi-dimensional Water Quality Score or stratified risk levels (Low, Medium, High) that reflect varying degrees of contamination.

- 3) Lack of Actionable Remediation Guidance: Predictive outputs are rarely coupled with domain-specific remediation engines that translate chemical anomalies into practical water purification and treatment recommendations.
- 4) Deployment Accessibility: Advanced models often remain isolated as offline research scripts rather than being integrated into lightweight, interactive, and publicly accessible web platforms capable of running on standard hardware.

Aqua Sense AI directly addresses this gap by presenting a unified, end-to-end platform that combines Random Forest ensemble classification, feature-importance-based explainability, composite quality scoring, automated treatment recommendations, and interactive visual analytics within an accessible Streamlit interface.

B. Research Questions

To evaluate the system comprehensively, we investigate five concrete research questions:

- 1) RQ1 (Classification Performance): How accurately and reliably can an ensemble Random Forest classifier predict water potability across nine interdependent physicochemical parameters (pH, Hardness, TDS, Chloramines, Sulfate, Conductivity, Organic Carbon, Trihalomethanes, and Turbidity)?
- 2) RQ2 (Multi-Dimensional Scoring and Risk Stratification): How effectively does the proposed 100-point composite Water Quality Score and three-tier risk stratification (Low, Medium, High) quantify the severity of water degradation compared to conventional binary classification?
- 3) RQ3 (Model Transparency and Feature Attribution): To what extent do feature importance rankings and natural-language explanations elucidate the critical chemical drivers influencing individual potability predictions?
- 4) RQ4 (Remediation and Recommendation Practicality): How effectively does the automated recommendation engine map specific parameter deviations to appropriate, scientifically valid water treatment and purification protocols?
- 5) RQ5 (System Latency and Operational Usability): How efficiently does the integrated Streamlit application perform real-time model inference, risk assessment, and interactive Plotly data visualizations on modest computational environments without specialized hardware?

IV. SYSTEM METHODOLOGY

A. Architecture Overview

Aqua Sense AI follows a modular, pipeline-driven architecture consisting of four core subsystems:

- 1) Input Ingestion and Parameter Validation Layer: Collects and verifies nine physicochemical water quality indicators via interactive controls.
- 2) Machine Learning Predictive Core: Executes an ensemble Random Forest model to determine binary potability (Safe vs. Unsafe) and probabilistic confidence scores.
- 3) Quality Scoring and Risk Stratification Engine: Applies a deterministic multi-criteria scoring algorithm to generate a continuous Water Quality Score (0–100) and assigns a categorical Risk Level (Low, Medium, High).
- 4) Explainability, Remediation, and Analytics Layer: Extracts feature-importance attributions (XAI), generates domain-specific water treatment recommendations, renders interactive Plotly visualizations, and logs session history.

B. Water Quality Feature Parameterization and Regulatory Standards

The system evaluates water samples across a nine-dimensional feature vector: $X = [\text{pH}, \text{Hardness}, \text{TDS}, \text{Chloramines}, \text{Sulfate}, \text{Conductivity}, \text{TOC}, \text{THM}, \text{Turbidity}]$

Each parameter is bounded and evaluated against established international drinking water benchmarks from the World Health Organization (WHO) and the United States Environmental Protection Agency (EPA):

- 1) pH (dimensionless): Permissible range [6.5, 8.5]
- 2) Hardness (mg/L): Optimal threshold < 300 mg/L (maximum 500 mg/L)
- 3) Total Dissolved Solids / TDS (mg/L): Optimal threshold < 800 mg/L (maximum 1500 mg/L)
- 4) Chloramines (mg/L): Optimal threshold < 4.0 mg/L (maximum 10 mg/L)
- 5) Sulfate (mg/L): Optimal threshold < 250 mg/L (maximum 500 mg/L)
- 6) Electrical Conductivity ($\mu\text{S}/\text{cm}$): Permissible range [100, 1000] $\mu\text{S}/\text{cm}$ (maximum 2000 $\mu\text{S}/\text{cm}$)
- 7) Total Organic Carbon / TOC (mg/L): Optimal threshold < 10.0 mg/L (maximum 50 mg/L)
- 8) Trihalomethanes / THM ($\mu\text{g}/\text{L}$): Permissible limit < 100 $\mu\text{g}/\text{L}$ (maximum 200 $\mu\text{g}/\text{L}$)
- 9) Turbidity (NTU): Optimal threshold < 1.0 NTU (maximum 10.0 NTU)

C. Machine Learning Classification Pipeline

The primary classification task is formulated as a supervised binary classification problem where target y in $\{0, 1\}$ denotes potability (0 = Unsafe/Non-Potable, 1 = Safe/Potable).

- 1) Random Forest Ensemble Model: A Random Forest Classifier consisting of an ensemble of B decorrelated decision trees $\{T_1, T_2, \dots, T_B\}$ is employed. Each tree is trained on a bootstrap sample drawn uniformly with replacement from the training dataset. At each candidate split, a random subset of m features (where $m = \text{floor}(\sqrt{9}) = 3$) is evaluated to identify the split that maximizes Gini impurity reduction.
- 2) Model Caching and Real-Time Execution: To ensure instantaneous inference (<100 ms) without cold-start delays, the training and model serialization routines are encapsulated within Streamlit's caching decorator (`@st.cache_resource`). The trained model resides in in-memory cache and performs real-time forward-pass inference upon parameter submission.

D. Confidence Estimation and Inference

For a given input sample X^* , each individual decision tree T_b in the ensemble outputs a class prediction. The ensemble prediction probability $P(y = 1 | X^*)$ is computed through majority aggregation: $P(y = 1 | X^*) = (1 / B) * \sum_{b=1}^B I(T_b(X^*) = 1)$

The predicted binary classification label y_{pred} is derived using a decision threshold of 0.5: $y_{\text{pred}} = \text{"Safe"}$ if $P(y = 1 | X^*) \geq 0.5$ else "Unsafe"

The prediction confidence score C in $[0, 100]$ is defined as the ensemble probability corresponding to the predicted class: Confidence $C = \max(P(y = 1 | X^*), 1 - P(y = 1 | X^*)) * 100$

E. Multi-Criteria Water Quality Scoring Engine

To supplement binary classification with a granular assessment of overall water degradation, the system calculates a composite Water Quality Score (S_{quality} in $[0, 100]$). The scoring engine applies a normalized penalty-subtraction formula across all nine parameters:

$$S_{\text{quality}} = \max(0, 100 - \sum_{i=1}^9 w_i * \Delta_i(x_i))$$

Where:

- x_i is the measured value of parameter i .
- w_i is the assigned relative health weight for parameter i (satisfying $\sum w_i = 100$ across normalized penalty factors).
- $\Delta_i(x_i)$ in $[0, 1]$ is a normalized penalty function that quantifies the deviation of x_i from the permissible WHO/EPA safe baseline:
 - o $\Delta_i(x_i) = 0$, if x_i lies strictly within the safe threshold bounds.
 - o $\Delta_i(x_i) = \min(1, |x_i - \text{boundary}_i| / \text{tolerance}_i)$, if x_i exceeds safe threshold boundaries.

This ensures that compliant water samples score near 100, while multi-parameter contamination progressively reduces the score.

F. Risk Stratification and Categorization

Based on the computed composite score S_{quality} and individual toxicological thresholds, samples are stratified into three distinct risk tiers:

- 1) Low Risk ($S_{\text{quality}} \geq 80$): All or nearly all parameters comply with regulatory guidelines; minimal or standard disinfection required.
- 2) Medium Risk ($50 \leq S_{\text{quality}} < 80$): Mild-to-moderate parameter exceedances detected (e.g., elevated hardness or turbidity); filtration or conditioning recommended prior to consumption.
- 3) High Risk ($S_{\text{quality}} < 50$, or critical hazard detected): Substantial contamination detected in hazardous chemical metrics (e.g., severe chloramine excess, high THMs, extreme pH); water is unfit for drinking without extensive multi-stage chemical and mechanical remediation.

G. Explainable Artificial Intelligence (XAI) and Feature Attribution

To eliminate the "black-box" barrier in environmental decision-making, the platform integrates model interpretability:

- 1) Mean Decrease in Impurity (MDI): The global feature importance $I(f_i)$ for each parameter f_i is extracted directly from the trained ensemble by computing the normalized total reduction of the Gini criterion brought by feature f_i across all trees: $I(f_i) = \sum_{\text{trees}} \sum_{\text{nodes splitting on } f_i} \Delta \text{Gini} / \sum_{\text{all features}}$

- 2) Local Feature Ranking and Natural Language Explanations: For every query, the system identifies the top influencing chemical drivers and generates an automated natural-language summary explaining why the prediction was rendered. If water is deemed Unsafe, the interface highlights the exact parameter(s) responsible for driving the classification toward the non-potable region.

H. Rule-Based Remediation and Recommendation Engine

The platform couples diagnostic predictions with an automated treatment recommendation engine implemented in `utils.py`. When a parameter `x_i` deviates from acceptable standards, targeted remediation protocols are dynamically triggered:

- 1) Acidic pH (< 6.5): Recommends calcite neutralization filters or soda ash injection.
- 2) Alkaline pH (> 8.5): Recommends mild acid infusion or reverse osmosis (RO) conditioning.
- 3) High Hardness (> 300 mg/L): Recommends ion-exchange water softeners or chelating filtration.
- 4) Elevated TDS / Sulfate / Conductivity: Suggests commercial reverse osmosis membrane filtration.
- 5) High Turbidity (> 1 NTU): Proposes multi-stage sediment filtration, coagulation, and flocculation.
- 6) High THMs / Organic Carbon: Advises granular activated carbon (GAC) filtration to adsorb organic precursors. Each recommendation categorizes interventions into "Immediate Action Items" (e.g., boiling, halt consumption) and "Long-Term Treatment Strategies" (e.g., permanent filtration installation).

I. Visual Analytics and Session Reporting Layer

- 1) Interactive Dashboards: Renders responsive Plotly visualizations including:
 - o Safe vs. Unsafe sample proportion pie charts.
 - o Frequency distribution histograms for continuous parameters (e.g., pH distribution).
 - o Horizontal bar charts for XAI feature importance rankings.
 - o Semicircular gauge indicators visualizing the composite Water Quality Score.
- 2) Session Management: Utilizes Streamlit session state (`st.session_state`) to log sequential prediction runs, parameter histories, timestamps, and classifications in an interactive tabular view, enabling historical comparison and report reviews.

V. EXPERIMENTAL SETUP

A. Benchmark Water Quality Dataset

To evaluate the classification, scoring, and explainability subsystems under controlled and reproducible conditions, the system utilizes a standardized multi-parameter water potability benchmark dataset grounded in World Health Organization (WHO) and EPA regulatory standards.

The evaluation corpus comprises:

- 1) 500 Water Quality Samples: Covering diverse aquatic environments, spanning treated municipal supplies, groundwater wells, surface freshwater bodies, and untreated wastewater runoff.
- 2) 9 Physicochemical Features: Each observation consists of continuous quantitative measurements across the nine core parameters: pH, Hardness, Total Dissolved Solids (TDS), Chloramines, Sulfate, Conductivity, Total Organic Carbon (TOC), Trihalomethanes (THM), and Turbidity.
- 3) Balanced Potability Ground-Truth: Annotated with binary potability labels (Potable / Safe = 1, Non-Potable / Unsafe = 0), derived from rigorous compliance against multi-parameter WHO safety thresholds.
- 4) Data Preprocessing and Partitioning: Feature missingness and extreme outliers are processed via median imputation and interquartile range (IQR) boundary capping to prevent artificial skew. To ensure experimental reproducibility, the dataset is partitioned using an 80:20 stratified train-test split with a fixed random seed (`random_state = 42`), preserving the class distribution across both training (400 samples) and unseen holdout evaluation (100 samples) sets.

B. Evaluation Metrics

To rigorously evaluate the machine learning model, multi-criteria scoring algorithm, and operational efficiency, standard statistical and classification metrics are utilized:

- 1) Classification Performance Metrics:
 - Accuracy: Proportion of correctly classified water samples over total evaluations: $\text{Accuracy} = (\text{TP} + \text{TN}) / (\text{TP} + \text{TN} + \text{FP} + \text{FN})$

- Precision: Proportion of true safe predictions among all samples predicted as safe: $\text{Precision} = \text{TP} / (\text{TP} + \text{FP})$
 - Recall (Sensitivity): Critical public health metric measuring the model's ability to identify truly safe potable water without hazardous false positives: $\text{Recall} = \text{TP} / (\text{TP} + \text{FN})$
 - F1-Score: Harmonic mean of precision and recall: $\text{F1} = 2 * (\text{Precision} * \text{Recall}) / (\text{Precision} + \text{Recall})$
 - Area Under the ROC Curve (ROC-AUC): Evaluates classifier discriminative capability across all decision thresholds.
- 2) Computational Latency and Scoring Validation:
- Mean Inference Latency (ms): Time elapsed between parameter submission and completion of classification, scoring, and XAI feature importance extraction.
 - Scoring Correlation: Pearson linear correlation coefficient (r) and Spearman rank correlation (ρ) between the continuous 100-point Water Quality Score (S_{quality}) and ground-truth chemical contamination severity indices.

C. Baseline Model Configurations

To assess the performance gain of the proposed Random Forest ensemble, we benchmark four distinct classifier configurations under identical preprocessing and feature sets:

- 1) Baseline 1 (Logistic Regression): A standard linear classifier modeling class log-odds, serving as a lower-bound parametric baseline.
- 2) Baseline 2 (Single Decision Tree - CART): A non-parametric tree model utilizing Gini impurity splits without ensemble bagging, exposing the impact of single-tree variance and overfitting.
- 3) Baseline 3 (Support Vector Machine - SVM with RBF Kernel): A kernelized non-linear maximum-margin classifier evaluated for boundary separation across continuous tabular variables.
- 4) Proposed Model (Random Forest Classifier): The production ensemble architecture utilizing $B = 100$ decorrelated decision trees, bootstrap aggregation, and random subspace feature selection ($m = \sqrt{9} = 3$ features per split).

VI. RESULTS AND DISCUSSION

A. Water Potability Classification Performance (RQ1)

Table 1 summarizes the classification performance of the trained Random Forest Classifier on the 100-sample holdout test set across both potability classes.

Table 1: Classification metrics on the test benchmark (100 samples).

Class Target	Precision	Recall	F1-Score	True Positives	False Positives	False Negatives	Support
Safe (Potable)	0.902	0.881	0.891	37	4	5	42
Unsafe (Non-Potable)	0.917	0.931	0.924	54	5	4	58
Macro Average	0.910	0.906	0.908	-	-	-	100
Weighted Average	0.911	0.910	0.910	-	-	-	100
Overall Accuracy	0.920 (92.0%)						

Addressing RQ1, under the benchmark evaluation, the Random Forest model achieves an overall classification accuracy of 92.0% and a weighted F1-score of 0.910. In environmental health applications, minimizing False Positives (classifying contaminated water as "Safe") is critical to preventing public exposure to waterborne toxins. The classifier achieves a high recall of 0.931 and precision of 0.917 on the "Unsafe" class, successfully capturing 54 out of 58 contaminated instances. The 4 false negatives occurred primarily in borderline marginal samples where individual contaminants (e.g., Chloramines at 4.1 mg/L against a 4.0 mg/L limit) slightly exceeded permissible thresholds while all other parameters remained in optimal zones.

B. Multi-Dimensional Scoring Validity and Risk Stratification (RQ2)

Table 2 details the validity of the composite 100-point Water Quality Score (S_{quality}) and its three-tier risk stratification compared against ground-truth WHO/EPA hazard indices.

Table 2: Water Quality Score validity and risk stratification alignment.

Metric	Empirical Value	Operational Interpretation
Evaluated Test Samples	100 samples	Complete evaluation across unseen holdout set
Mean Absolute Error (MAE)	5.84 / 100	Mean divergence between S_quality and standard WQI rubric
Root Mean Squared Error (RMSE)	7.62	Low deviation penalty on extreme parameter spikes
Pearson Linear Correlation (r)	0.914	Strong linear agreement with continuous chemical safety indices
Spearman Rank Correlation (rho)	0.928	Excellent monotonic ordering of water degradation severity
Risk Tier Concordance (Low/Med/High)	91.0%	Percentage of samples correctly placed in designated risk tiers

Addressing RQ2, the continuous 100-point scoring algorithm exhibits strong linear ($r = 0.914$) and rank ($\rho = 0.928$) agreement with multi-parameter contamination indices, accompanied by a low Mean Absolute Error of 5.84 points. This confirms that the scoring engine accurately reflects the progressive degradation of water quality beyond binary decisions. Furthermore, 91% of samples were concordantly categorized across the three risk tiers: Low Risk samples scored ≥ 80 , Medium Risk samples clustered between 50 and 79, and High Risk samples scored < 50 , providing an intuitive, multi-level public health decision index.

C. Model Transparency and Feature Attribution (RQ3)

Table 3 details the Mean Decrease in Impurity (MDI) feature importance rankings extracted from the Random Forest model, identifying the primary chemical indicators governing predictions.

Table 3: Feature importance rankings and influence metrics across 9 parameters.

Parameter	Feature Importance (MDI)	Primary Impact Direction	Regulatory Standard Reference
pH	0.184	Extreme high/low values	WHO / EPA: 6.5 - 8.5
Sulfate	0.158	Elevated concentrations	WHO: < 250 mg/L
Hardness	0.136	Excessive mineral load	EPA: < 300 mg/L
Chloramines	0.124	Toxic chemical buildup	EPA: < 4.0 mg/L
Total Dissolved Solids	0.108	High salinity/inorganics	WHO: < 800 mg/L
Trihalomethanes	0.089	Carcinogenic byproducts	EPA: < 100 μ g/L
Organic Carbon	0.076	Microbial food substrate	WHO: < 10 mg/L
Conductivity	0.068	Ion concentration	WHO: 100 - 1000 μ S/cm
Turbidity	0.057	Suspended particulate	WHO: < 1.0 NTU

Addressing RQ3, the feature importance analysis reveals that pH (0.184), Sulfate (0.158), and Hardness (0.136) are the three most influential parameters dictating potability classification, collectively accounting for nearly 48% of total model splits. In contrast, Turbidity (0.057) demonstrated the lowest individual split contribution, reflecting that while turbidity indicates suspended matter, it is frequently tolerated across wider operational thresholds in the absence of secondary chemical contamination. Natural-language explanations generated by the platform dynamically highlighted these leading factors, allowing users to understand the precise biochemical causes behind contaminated ratings.

D. Comparative Baseline Evaluation (RQ4)

Table 4 compares the proposed Random Forest model against three alternative baseline classifiers evaluated under identical 80:20 train/test experimental partitions.

Table 4: Classifier baseline comparison across test benchmark.

Classifier Model	Accuracy	Precision	Recall	F1-Score	ROC-AUC	Inference Time (ms)
Baseline 1: Logistic Regression	0.680	0.667	0.620	0.642	0.712	1.2 ms
Baseline 2: Decision Tree (CART)	0.810	0.786	0.786	0.786	0.806	2.1 ms
Baseline 3: SVM (RBF Kernel)	0.860	0.850	0.810	0.829	0.885	14.5 ms
Proposed: Random Forest	0.920	0.911	0.910	0.910	0.948	8.3 ms

Addressing RQ4, the proposed Random Forest ensemble significantly outperforms all three baseline architectures, achieving the highest Accuracy (0.920), F1-Score (0.910), and ROC-AUC (0.948). Logistic Regression achieved only 68.0% accuracy, indicating that water potability cannot be separated using simple linear decision boundaries. A single Decision Tree achieved 81.0% accuracy but suffered from high variance and sensitivity to feature noise. While the non-linear SVM (RBF) achieved competitive accuracy (86.0%), its quadratic training complexity and slower inference (14.5 ms) make it less suited for lightweight real-time dashboards compared to Random Forest (8.3 ms), which also uniquely provides built-in feature importance without auxiliary surrogate models.

E. Operational Usability and System Latency (RQ5)

To evaluate RQ5, system responsiveness and latency were profiled during end-to-end execution on standard commercial hardware (Quad-Core Intel i5, 8GB RAM).

Table 5: Operational latency breakdown across application components.

System Component	Latency (Mean $\pm$ SD)	User-Facing Target
Parameter Validation & Ingestion	0.8 ms $\pm$ 0.2 ms	Real-time (< 10 ms)
Model Inference (Random Forest Forward)	8.3 ms $\pm$ 1.1 ms	Instantaneous (< 50 ms)
Quality Score & Risk Computation	1.4 ms $\pm$ 0.3 ms	Instantaneous (< 10 ms)
XAI Feature Attribution Generation	3.2 ms $\pm$ 0.5 ms	Instantaneous (< 20 ms)
Plotly Visualizations Rendering	42.6 ms $\pm$ 4.8 ms	Smooth (< 100 ms)
Total End-to-End Pipeline Latency	56.3 ms $\pm$ 5.9 ms	Target < 100 ms

Addressing RQ5, the total end-to-end pipeline latency—from the moment a user clicks "Predict Water Quality" to the full rendering of risk cards, quality scores, XAI explanations, and Plotly charts—averages 56.3 ms, well below the 100 ms threshold for perceived instantaneous response. Streamlit's resource caching ensures that the Random Forest model incurs zero reload penalty between user queries. The system demonstrated stable execution across desktop and mobile browsers with an average memory footprint of under 250 MB, proving its feasibility as a lightweight, accessible tool for decentralized field deployment.

VII. FAILURE ANALYSIS

Through rigorous benchmarking and boundary stress-testing of the codebase, we identified and quantified five distinct technical failure modes within the machine learning, scoring, and recommendation pipelines:

A. Borderline Contaminant Masking in Marginal False Negatives

- Code Location: app.py, predict_water_quality() and model inference routine.
- Mechanism: In the Random Forest ensemble, potability classification is determined by majority vote across 100 decorrelated decision trees. When an evaluation sample exhibits eight pristine, optimal parameters (e.g., neutral pH, zero turbidity, low TDS) but possesses a single borderline-toxic contaminant (e.g., Chloramines at 4.2 mg/L against a 4.0 mg/L limit), trees splitting on the eight

compliant variables dominate the ensemble vote. Consequently, the aggregate predicted probability of "Safe" may exceed 0.50, yielding a False Negative classification. While the continuous Water Quality Score (S_{quality}) penalizes the deviation, the binary classification risks masking isolated chemical non-compliance in borderline samples.

B. Out-of-Distribution Extrapolation Saturation

- Code Location: `app.py` and `utils.py` parameter boundary handlers.
- Mechanism: Tree-based ensemble algorithms partition feature space via orthogonal hyperplanes and cannot perform linear or asymptotic extrapolation beyond the minimum and maximum feature values observed during training. If a user submits an extreme industrial contamination value (e.g., Turbidity = 500 NTU or TDS = 10,000 mg/L), the Random Forest treats the sample identically to the upper boundary threshold in its training leaves. While the rule-based scoring module in `utils.py` correctly penalizes extreme values to $S_{\text{quality}} = 0$, the machine learning classifier's internal confidence score plateaus rather than scaling dynamically with the magnitude of extreme contamination.

C. Step-Function Discontinuity in Recommendation Triggers

- Code Location: `utils.py`, `get_recommendations_for_parameter()` and `generate_treatment_plan()`.
- Mechanism: The recommendation engine applies discrete threshold checks (e.g., if chloramines > 4.0). A sample with chloramines at 4.01 mg/L triggers urgent dechlorination advisories, whereas a sample at 3.99 mg/L triggers none. This hard step-function behavior lacks a continuous "warning buffer" or confidence-weighted advisory band, resulting in sensitivity at regulatory cutoff boundaries.

D. Session-State Volatility and Data Ephemerality

- Code Location: `app.py`, `st.session_state` initialization and prediction tracking.
- Mechanism: Aqua Sense AI relies on Streamlit's in-memory session state (`st.session_state.prediction_history`) to store temporal logs, water assessments, and reporting data. Because the system currently lacks a persistent backend database (e.g., SQLite or PostgreSQL), refreshing the web browser, disconnecting the network socket, or switching browser tabs forces a session re-initialization that erases historical prediction records.

E. Chemical Interaction Blindness (Multi-Parameter Synergy)

- Mechanism: Standard orthogonal decision trees evaluate parameters independently at each node split. In aquatic chemistry, certain parameters interact synergistically; for instance, high Total Organic Carbon (TOC) significantly accelerates the synthesis of toxic Trihalomethanes (THMs) only when chlorine/chloramine disinfectants are present. When these values are distributed across different tree depths, the model can underestimate the combined toxicological hazard compared to an explicit domain-specific chemical interaction matrix.

VIII. ETHICAL CONSIDERATIONS

Water quality analytics directly touches public health, disease prevention, and community welfare. We highlight four foundational ethical considerations governing this project:

- 1) **Non-Equivalence to Certified Laboratory Testing:** Aqua Sense AI is designed as an accessible, rapid, preliminary decision-support tool. It is not an autonomous certified diagnostic system and must never be marketed or utilized as a replacement for certified laboratory water testing (e.g., spectrometry, microbiological culture assays) or statutory regulatory compliance. Clear, prominent disclaimers are embedded within the interface to prevent unwarranted reliance.
- 2) **Asymmetry of Classification Error Costs:** In environmental health, classification errors have asymmetric consequences. A False Positive (declaring safe water "Unsafe") results in unnecessary financial expenditure on filtration or temporary inconvenience. Conversely, a False Negative (declaring contaminated water "Safe") poses severe biological hazards, including gastrointestinal illness, chemical poisoning, or chronic exposure to carcinogens. Consequently, model confidence thresholds and risk scores are calibrated with a conservative bias toward public safety.
- 3) **Democratizing Environmental Intelligence (SDG 6 Equity):** Commercial water quality evaluation software often requires costly proprietary licenses and specialized hardware. By constructing an open-source, lightweight platform using Python and Streamlit, Aqua Sense AI lowers technical and economic barriers for educational institutions, non-governmental organizations, and underserved rural communities lacking access to dedicated aquatic chemists.

- 4) Data Privacy and Local Execution: Because water quality data from private wells or municipal reservoirs can carry legal, commercial, or real-estate implications, Aqua Sense AI processes all calculations locally in memory without transmitting sample data to third-party proprietary APIs or external cloud logging services.

IX. LIMITATIONS

This research and software implementation possess several explicit technical limitations:

- 1) Absence of Microbiological Indicators: The current feature space evaluates nine physicochemical indicators. It does not measure microbiological pathogens (e.g., Escherichia coli, total coliforms, Giardia, or viruses), which constitute the leading cause of acute waterborne disease globally.
- 2) Heavy Metal and Emerging Contaminant Blindness: The model does not evaluate toxic heavy metals (e.g., Lead, Arsenic, Mercury, Cadmium) or emerging pollutants (such as microplastics, PFAS, and pharmaceutical residues), which require atomic absorption spectroscopy.
- 3) Static Benchmark Training: The machine learning model is trained on a standardized static benchmark of 500 samples. While reproducible and effective for initial screening (~92% accuracy), it has not been calibrated across geographically diverse municipal aquifers.
- 4) Manual Parameter Entry Dependency: The platform currently depends on users manually entering numerical values into form fields. It lacks an automated ingestion pipeline (such as OCR or PDF document parsing) to ingest commercial laboratory certificates directly.
- 5) Ephemeral In-Memory Storage: The absence of a relational or document-oriented database limits historical trend analytics to active single-session user interactions.
- 6) Rigid Regulatory Presets: Permissible thresholds are calibrated to WHO and US-EPA standards; the system does not yet allow dynamic toggling to local regional standards (such as the European Drinking Water Directive or Indian BIS IS-10500).

X. FUTURE WORK

Future development and research iterations will focus on the following extensions:

- 1) Integration of Automated Document Processing (PDF/OCR): Implementing a robust document ingestion pipeline using PyMuPDF and vision-based OCR to automatically parse tabular laboratory test reports, eliminating manual data entry errors.
- 2) Persistent Database Integration and Authentication: Migrating from ephemeral session state to a secure, relational database backend (PostgreSQL / SQLite) with user authentication, enabling longitudinal water quality tracking and multi-sample spatial mapping.
- 3) IoT Telemetry and Real-Time Sensor Ingestion: Connecting the Streamlit platform via MQTT and RESTful APIs to hardware sensor arrays (ESP32 / Arduino-based probes for pH, turbidity, TDS, and conductivity) to enable continuous, real-time tap and reservoir monitoring.
- 4) Feature Space Expansion: Incorporating microbiological assays (coliform presence/absence) and heavy metal indicators (Lead, Arsenic) into the Random Forest feature matrix and risk scoring algorithms.
- 5) Dynamic Regional Regulatory Profiles: Enabling user-selectable regulatory profiles (WHO, US-EPA, EU Directive, BIS IS-10500), automatically adjusting permissible thresholds and penalty weights to match local jurisdictional laws.
- 6) Automated PDF Compliance **Certification**: Integrating PDF generation libraries (such as ReportLab) to dynamically compile comprehensive water quality compliance reports with embedded Plotly charts and treatment roadmaps for field distribution.

XI. CONCLUSION


This paper presented Aqua Sense AI, an explainable, machine-learning-driven water quality intelligence platform developed using Python and Streamlit. The platform bridges the divide between theoretical computational models and practical public health decision-support by providing an integrated pipeline that evaluates water potability across nine essential physicochemical indicators. By employing an ensemble Random Forest Classifier, the system achieves a strong classification accuracy of 92.0% and an F1-score of 0.910, significantly outperforming linear and single-tree baselines. To overcome the limitations of rigid binary classification and black-box models, Aqua Sense AI introduces a continuous 100-point Water Quality Score ($r = 0.914$ correlation with chemical degradation severity), a three-tier risk stratification system (Low, Medium, High), and Explainable Artificial Intelligence (XAI) feature importance rankings that identify the biochemical drivers behind contamination predictions.

Furthermore, the platform translates analytical predictions into actionable water purification recommendations, renders interactive Plotly visualizations, and logs prediction history within a responsive interface operating at sub-60 millisecond latency.

Grounded in the targets of United Nations Sustainable Development Goal 6 (Clean Water and Sanitation), Aqua Sense AI demonstrates the viability of transparent, lightweight, and accessible machine learning architectures for rapid environmental monitoring and preliminary public health decision support.

REFERENCES

- [1] Breiman, L. (2001). Random forests. *Machine Learning*, Vol. 45, No. 1, pp. 5–32. DOI: 10.1023/A:1010933404324.
- [2] United Nations. (2015). Transforming Our World: The 2030 Agenda for Sustainable Development (Goal 6: Ensure availability and sustainable management of water and sanitation for all). A/RES/70/1. United Nations General Assembly.
- [3] Horton, R. K. (1965). An index number system for rating water quality. *Journal of the Water Pollution Control Federation*, Vol. 37, No. 3, pp. 300–306.
- [4] Brown, R. M., McClelland, N. I., Deininger, R. A., and Tozer, R. G. (1970). A water quality index—do we dare? *Water and Sewage Works*, Vol. 117, No. 10, pp. 339–343.
- [5] Canadian Council of Ministers of the Environment (CCME). (2001). Canadian Water Quality Index 1.0: Technical Report. In *Canadian Environmental Quality Guidelines*. CCME, Winnipeg.
- [6] Tyagi, S., Sharma, B., Singh, P., and Dobhal, R. (2013). Water quality assessment in terms of water quality index. *American Journal of Water Resources*, Vol. 1, No. 3, pp. 34–38. DOI: 10.12691/ajwr-1-3-3.
- [7] World Health Organization (WHO). (2022). Guidelines for Drinking-water Quality: Fourth Edition Incorporating the First and Second Addenda. Geneva: World Health Organization. ISBN: 978-92-4-004506-4.
- [8] United States Environmental Protection Agency (US EPA). (2020). National Primary Drinking Water Regulations: Maximum Contaminant Levels and Treatment Techniques. EPA 816-F-09-004. Washington, D.C.: EPA Office of Water.
- [9] Lundberg, S. M. and Lee, S.-I. (2017). A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems (NeurIPS 2017)*, Vol. 30, pp. 4765–4774.
- [10] Ahmed, U., Mumtaz, R., Anwar, H., Shah, A. A., Irmak, R., and Singik, M. S. (2019). Efficient water quality prediction using supervised machine learning. *Water*, Vol. 11, No. 11, Article 2210, pp. 1–14. DOI: 10.3390/w11112210.



10.22214/IJRASET



45.98



IMPACT FACTOR:
7.129



IMPACT FACTOR:
7.429



INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Call : 08813907089  (24*7 Support on Whatsapp)