**AVE** **Trends in Intelligent Applied Sciences**

# Predictive Water Quality Assessment through Machine Learning and Ensemble Intelligence for Sustainable Water Resource Management

**Raja Brahmendra Chowdary Veerepalli[1], S. Rubin Bose[2*], J. Angelin Jeba[3], R. Regin[4], T. Shynu[5], M. Mohamed Sameer Ali[6]**

[1]Department of Computer Science and Engineering, Texas State University, San Marcos, Texas, United States of America.
[2,4]School of Computer Science and Engineering, SRM Institute of Science and Technology, Ramapuram, Chennai, Tamil Nadu, India.
[3]Department of Electronics and Communication Engineering, S.A. Engineering College, Chennai, Tamil Nadu, India.
[5,6]Department of Research and Development, Dhaanish Ahmed College of Engineering, Chennai, Tamil Nadu, India.
kgm120@txstate.edu[1], rubinbos@srmist.edu.in[2], angelinjeba@saec.ac.in[3], reginr@srmist.edu.in[4], shynu469@gmail.com[5], sameerali7650@gmail.com[6]

*Corresponding author

**Abstract:** In this paper, researchers propose a machine learning (ML) - based predictive analytics platform for water quality assessment without the use of physical sensors. The proposed system uses advanced regression models and ensemble learning techniques to estimate critical water quality indicators from historical datasets, environmental parameters, and contextual factors. The platform provides scalable, efficient, and accurate predictions of water health through a data-driven architecture that replaces costly and maintenance-intensive sensor networks. Experimental studies have shown that it is robust in predicting metrics such as pH, turbidity, and dissolved oxygen with high fidelity, even when the data are noisy and incomplete. Besides its technical performance, the framework provides timely insights for early detection of contamination, resource optimisation, and sustainable management of aquatic ecosystems. The results show the potential of ML-driven platforms to revolutionise water monitoring practices, reduce reliance on infrastructure, and help decision-makers maintain safe, dependable, and environmentally friendly water management. Future work will include real-time environmental feeds, dynamic model adaptation, integration with IoT platforms, if necessary, and the development of explainable AI methods to enhance transparency and trust in water quality predictions.

**Keywords:** Water Quality Prediction; Artificial Intelligence; Machine Learning; Predictive Analytics; Sensor-Free Monitoring; Environmental Data Modelling; Sustainable Water Management; ML-Driven Platforms.

**Cite as:** R. B. C. Veerepalli, S. R. Bose, J. A. Jeba, R. Regin, T. Shynu, and M. M. S. Ali, "Predictive Water Quality Assessment through Machine Learning and Ensemble Intelligence for Sustainable Water Resource Management," *AVE Trends in Intelligent Applied Sciences*, vol. 2, no. 1, pp. 30–44, 2026.

**Journal Homepage:** https://www.avepubs.com/user/journals/details/ATIAS

**Received on:** 11/12/2024, **Revised on:** 02/02/2025, **Accepted on:** 19/04/2025, **Published on:** 03/03/2026

**DOI:** https://doi.org/10.64091/ATIAS.2026.000294

## 1. Introduction

Water is one of the most important natural resources for supporting life, economic development, ecological balance and public health. Water quality directly affects human health, agricultural productivity, industrial activity and environmental integrity.

However, with rapid urbanisation, industrialisation, population growth, agricultural runoff, and climate change, the risk of water pollution has increased significantly worldwide [1]. Fresh water is threatened by heavy metals, pathogens, organic pollutants, chemicals, and excess nutrients, among others, which pose serious environmental and public health problems. Globally, billions of people lack access to safely managed drinking water services, and waterborne diseases remain a major cause of illness and death in many developing countries [2]. Continuous monitoring and assessment of water quality are therefore important requirements for sustainable management of water resources and protection of aquatic ecosystems. Conventional water quality monitoring generally comprises laboratory analysis, field sampling and sensor networks. Such approaches can measure important physicochemical and biological parameters, including pH, dissolved oxygen (DO), turbidity, conductivity, total dissolved solids (TDS), biochemical oxygen demand (BOD), and nutrient concentrations [3]. These methods, while capable of providing accurate results, can entail high operating expenses, the need for specialised equipment and trained personnel, and considerable maintenance. Additionally, deploying large-scale sensor networks across spatially distributed water bodies may be economically challenging, particularly in developing countries and remote areas with limited resources. Sensor failures, calibration problems, communication disruptions, and environmental degradation can also compromise data reliability and monitoring effectiveness [4]. The growing need for intelligent, scalable alternatives that can complement or replace reliance on conventional monitoring architectures drives this. Recent advances in artificial intelligence (AI), machine learning (ML) and data analytics have created new opportunities to address these challenges. The ability of machine learning techniques to extract meaningful patterns from large, complex datasets is remarkable, enabling accurate prediction, classification, and decision-making across many domains [5].

In environmental science, ML-based approaches have been increasingly used to model complex ecological systems, predict environmental changes, and support sustainable resource management. Unlike classical statistical methods, which often rely on strict assumptions about data distribution and linearity, machine learning algorithms can successfully identify nonlinear relationships, interactions and hidden dependencies among multiple variables. [6] This capacity makes them particularly suitable for assessing water quality, where environmental conditions are influenced by many interrelated physical, chemical, meteorological, and anthropogenic factors. The enhanced availability of historical datasets on environmental conditions, meteorological and hydrological records, land-use data, and contextual parameters has further accelerated the adoption of machine learning for water quality prediction. Predictive models can estimate current and future water quality conditions using historical observations and environmental indicators, rather than relying solely on direct sensor observations [7]. These approaches offer a promising alternative to conventional monitoring systems by reducing the infrastructure requirements while still providing reliable predictive performance. Machine learning-based predictive analytics can continuously monitor water health, detect contamination threats early, and perform anomaly detection and smart prediction of key water quality parameters [8]. These capabilities are crucial in supporting proactive interventions and evidence-based decisions in the management of water resources. Several machine learning techniques have been successfully used to predict and assess water quality. Regression-based models such as Linear Regression, Support Vector Regression (SVR), and Decision Tree Regression have been widely used to estimate individual water quality parameters. Ensemble learning methods, such as Random Forests (RF), Gradient Boosting Machines (GBM), Extreme Gradient Boosting (XG-Boost), and AdaBoost, have received significant attention for their ability to improve predictive accuracy by combining multiple weak learners into a stronger ensemble [9].

These algorithms can successfully handle complex datasets, control overfitting, and operate reliably across various environmental conditions. Recently, deep learning architectures such as Artificial Neural Networks (ANNs), Long Short-Term Memory (LSTM) networks, Convolutional Neural Networks (CNNs), and Transformer-based models have shown great potential in modelling temporal dynamics and nonlinear relationships in environmental datasets [10]. The incorporation of these advanced learning techniques has greatly enhanced the accuracy and reliability of water quality forecasting systems. Despite these advances, many existing water quality monitoring frameworks remain heavily dependent on physical sensors, Internet of Things (IoT) infrastructure, or remote sensing platforms. These technologies enable real-time data acquisition, but they are challenging due to deployment costs, maintenance complexity, power consumption, communication requirements, and scalability [11]. Moreover, data loss, sensor drift, hardware failures and environmental interference can all hamper performance monitoring. These limitations point to the need for alternative predictive frameworks that leverage existing environmental and historical data without requiring extensive sensor infrastructure. Such data-driven solutions are particularly useful for areas with limited technical resources, limited monitoring coverage or limited budgets [12]. The advent of sensor-agnostic predictive analytics represents a major paradigm shift in environmental monitoring. Machine learning models based on historical water quality data, climatic variables, hydrological characteristics, seasonal trends, and contextual environmental factors can accurately predict critical water quality indicators. These forecasting systems can provide timely insights even without continuous sensor measurements, thereby reducing operational costs and increasing the accessibility of monitoring. Furthermore, modern machine learning algorithms are better suited to real-world environmental applications because they can handle noisy, incomplete, and heterogeneous data, which are common problems in data quality [13].

The value of such predictive systems is especially apparent in nations with serious water quality problems. For instance, in India, rapid industrial growth, intensification of agriculture, urban expansion, and inadequate wastewater treatment have resulted in widespread water pollution. Unsafe drinking water still affects millions of people and has significant health, social and economic consequences [14]. Many developing countries face similar problems in which a comprehensive monitoring infrastructure is not in place or is difficult to maintain. In such scenarios, machine learning-based predictive analytics can serve as a valuable decision-support tool to identify contamination risks, optimise resource allocation, and implement early intervention strategies before water quality degradation reaches critical thresholds [15]. These challenges and opportunities motivate this study to propose a machine learning-driven predictive analytics platform for water quality assessment using historical datasets, environmental variables and contextual information. The framework uses advanced regression and ensemble learning techniques to estimate important water quality parameters, including pH, turbidity, dissolved oxygen and associated parameters. The proposed platform aims to enable accurate, scalable, and cost-effective water quality predictions by minimising reliance on expensive monitoring infrastructure and leveraging intelligent, data-driven modelling. The system will assist environmental agencies, policymakers, researchers, and water resource managers in making informed decisions on water safety, contamination control, and ecosystem protection [9]. The main contribution of this research is demonstrating the effectiveness of machine learning-based predictive analytics as a sustainable alternative for water quality assessment. The proposed framework is a step towards intelligent environmental monitoring systems through robust modelling techniques, thorough evaluation, and practical applicability. Ultimately, this work supports the broader vision of smart water management, sustainable development, improved public health outcomes and resilient environmental governance in an increasingly data-driven world [11].

## 2. Literature Review

In recent years, the application of Artificial Intelligence (AI), Machine Learning (ML), Deep Learning (DL), the Internet of Things (IoT), Remote Sensing (RS), and virtual sensing technologies has substantially changed water quality monitoring and assessment. Researchers have focused more on developing intelligent, scalable and cost-effective approaches to overcome the limitations of traditional laboratory-based and sensor-dependent monitoring systems. Hasan et al. [1] comprehensively reviewed water quality monitoring systems using machine learning and IoT technologies. They explored how smart sensors, cloud computing, wireless communication, and predictive analytics could be combined to monitor the environment in real time. The authors point out that machine learning algorithms can enhance contamination detection, automate decision-making, and improve the efficiency of water resource management systems. Mohan et al. [2] reviewed the combination of machine learning and remote sensing technologies for water quality monitoring and prediction. Their work proved that satellite images combined with ML algorithms can be employed for large-scale, cost-effective, and continuous monitoring of water bodies. The study highlights the importance of remote sensing in overcoming the spatial limitations of traditional monitoring networks while supporting long-term environmental management. Emerging roles of deep learning in water quality assessment and prediction. Zhi et al. [3] studied. In their review, they covered advanced neural network architectures such as Convolutional Neural Networks (CNNs), Recurrent Neural Networks (RNNs), Long Short-Term Memory (LSTM) networks and Transformer models. The authors concluded that deep learning techniques are better at handling high-dimensional environmental data and learning complex nonlinear relationships among water quality variables. Essamlali et al. [4] provided a detailed review of recent advances in machine learning and IoT-enabled water quality monitoring systems.

Their research showed how to combine smart analytics, sensor technologies, edge computing, and cloud platforms to construct real-time water-quality intelligence systems. The authors also discussed challenges related to sensor reliability, data security, interoperability, and deployment scalability. Achuthankutty et al. [5] proposed a framework for water quality assessment based on deep learning that employs IoT sensor data and Long Short-Term Memory (LSTM) models. Explainable artificial intelligence techniques were used in the study to enhance the model's transparency and interpretability. Results showed that deep learning models can effectively capture temporal dependencies in water quality data and provide accurate, explainable predictions. Paepae et al. [6] reviewed the development of water quality assessment technologies from traditional physical sensing systems to virtual sensing techniques. Their work examined the application of machine learning algorithms to estimate water quality parameters from indirect environmental measurements and historical datasets. The authors considered virtual sensing a promising solution to reduce monitoring costs and extend coverage in a resource-constrained environment. Zhang et al. [7] proposed a deep learning-based framework for retrieving water quality from Landsat-8 satellite images over Dongping Lake. Their study demonstrated the effectiveness of combining remote sensing data with deep neural networks to estimate water quality parameters. The proposed approach successfully captured spatial and temporal variations in water quality and demonstrated the potential of satellite-driven monitoring systems. Zhu et al. [8] reviewed the application of machine learning techniques to water quality evaluation. The study explored several algorithms, including Support Vector Machines (SVMs), Random Forests (RFs), Artificial Neural Networks (ANNs), Decision Trees, and ensemble learning approaches. The authors highlighted the superior predictive accuracy and flexibility of machine learning methods compared to traditional statistical techniques.

Rodríguez-López et al. [9] studied the integration of machine learning and remote sensing in water quality analysis of Lake Ranco, Southern Chile. Their work demonstrated the power of satellite-derived environmental information and predictive analytics to produce accurate spatial assessments of water quality. The study emphasised the significance of geospatial intelligence in regional water resource planning and environmental sustainability. Rejula et al. [10] proposed a decentralised water quality classification framework based on federated learning and a recurrent neural network. The proposed system allows multiple monitoring stations to collaboratively train predictive models without exchanging raw data, thus preserving privacy and reducing communication overhead. Their results demonstrated the efficacy of federated learning in distributed environmental monitoring applications. Peerzade and Kamat [11] investigated machine learning as a decision-support tool for predicting water quality across different aquatic environments, including rivers, lakes, groundwater, coastal waters, and wastewater treatment plants. They found that ML models can provide near-real-time predictions and support proactive management strategies to maintain water quality standards. Kushwaha and Pandey [12] proposed an intelligent IoT-based water quality monitoring and predictive analytics system using machine learning techniques. They integrated low-cost sensing devices and algorithms, such as Random Forest and LSTM, to predict water quality parameters and compute Water Quality Index (WQI) values in their framework. "Not only is it feasible to do both, but this study demonstrated the ability to integrate predictive analytics with real-time environmental surveillance," he said. Basirian et al. [13] studied water quality monitoring in coastal hypoxic areas using a combination of satellite imagery and machine learning models. Their research demonstrated that remote sensing data is an effective tool for detecting oxygen-poor areas in coastal ecosystems and can be used to develop early warning systems and support marine environmental protection efforts.

Zounemat-Kermani and Kheimi [14] studied the application of ensemble boosting methods for surface water quality modelling. Their study explored advanced algorithms such as Gradient Boosting Machines (GBM), Extreme Gradient Boosting (XGBoost), AdaBoost, CatBoost and LightGBM. The authors concluded that boosting-based ensemble methods consistently outperformed traditional machine learning methods due to their ability to improve predictive accuracy, reduce bias, and model nonlinear environmental processes effectively. Das [15] proposed an integrated framework for predicting urban surface water quality scenarios in the Baitarani River Basin, Odisha, based on the Water Quality Index (WQI), multivariate statistical techniques, and machine learning models. The water quality condition assessment was performed by combining Principal Component Analysis (PCA), Pearson correlation analysis, GIS-based spatial mapping, hydrogeochemical interpretation and machine learning classification. The Optimised Forest classifier showed better predictive performance, suggesting the benefits of data-driven approaches for contamination assessment and sustainable water resources management. Despite remarkable progress in machine learning, deep learning, IoT, remote sensing, federated learning, and virtual sensing technologies for water quality monitoring, most current approaches rely on physical sensors, satellite observations, or dedicated monitoring infrastructure. Deployment costs, maintenance requirements, data quality, communication overhead, and scalability remain challenging issues. Hence, there is an increasing demand for sensor-independent predictive analytics frameworks that utilise historical datasets, environmental variables, and contextual information to accurately estimate water quality conditions. These smart systems can reduce reliance on infrastructure, enhance monitoring availability, enable early detection of contamination, and promote sustainable management of water resources through scalable, cost-effective decision-support solutions.

## 3. Methodology

The proposed methodology is executed within a structured, multi-phase framework that systematically converts raw environmental data into actionable water-quality insights.

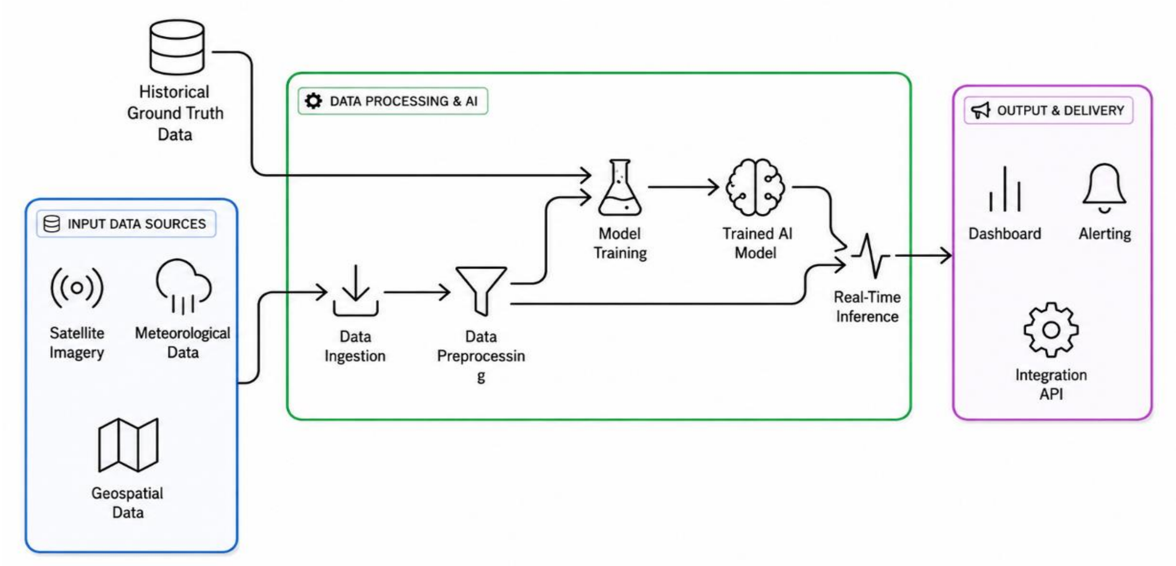


**Figure 1:** Block diagram

The methodology follows a logical workflow from data acquisition and pre-processing to predictive modelling and end-user application (see architectural framework). The entire process is broken down into five core stages, each representing a unique layer of the predictive analytics platform, to ensure scalability, accuracy and operational efficiency. These stages together enable data collection, feature engineering, machine learning-based prediction, performance evaluation, and intelligent decision support for water quality assessment. The AI-Powered Water Quality Intelligence Platform, shown in Figure 1, consists of three main stages. It does this by fetching data from various Input Data Sources, including Satellite Imagery (to detect visual cues like algal blooms), Meteorological Data (such as rainfall that influences runoff), and Geospatial Data (such as land-use maps). Second, this data is consumed and pre-processed in the Data Processing & AI core. This cleaned data is then used to train an AI model, which compares it against Historical Ground Truth Data (past, verified water quality measurements) to learn patterns. Once trained, the Trained AI Model performs Real-Time Inference on new, incoming data to predict the current water quality. Finally, these predictions are sent to the Output & Delivery systems, where the information is made available to users through a Dashboard, automated Alerting notifications for bad quality, and other systems connect via an Integration API.

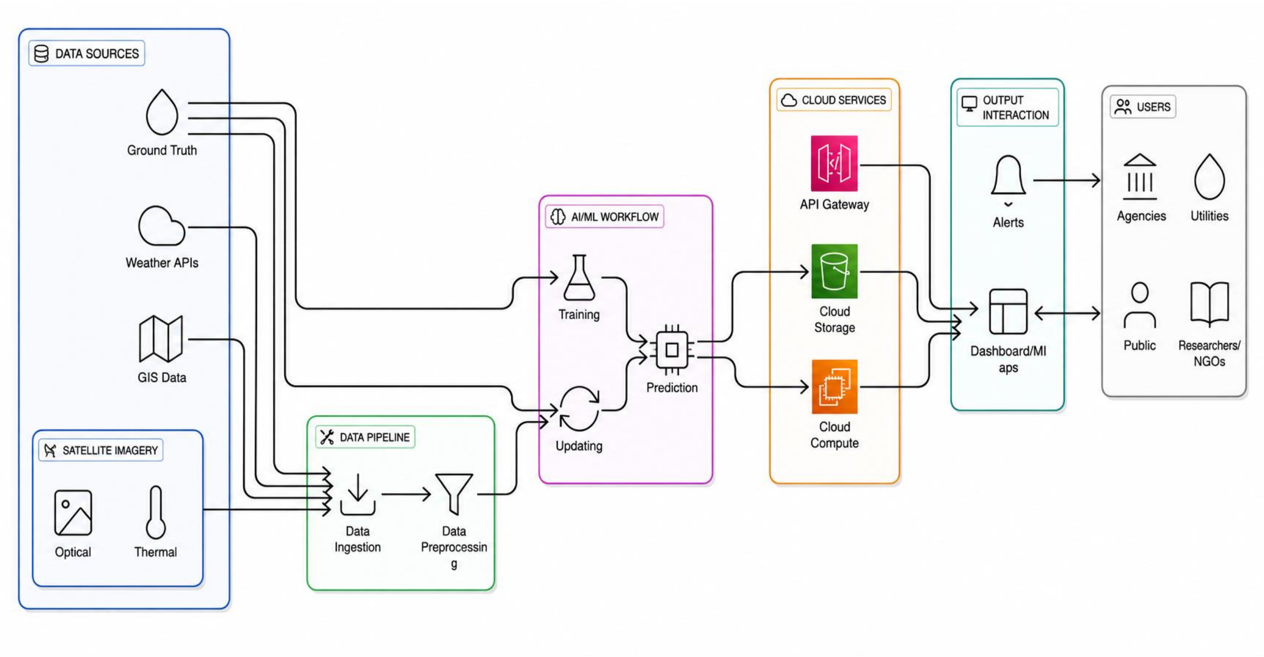


**Figure 2:** Architecture diagram

Figure 2 Details of the platform's data flow. Satellite Imagery (Optical and Thermal) is ingested and preprocessed in a dedicated Data Pipeline. This processed data is fed into the AI/ML Workflow along with Ground Truth records, Weather APIs, and GIS Data. This central block is responsible for the continuous cycle of model Training, Prediction generation and model updating with new data. All outputs are managed by Cloud Services (e.g. API Gateway, Storage, Compute), which provide the final insights to various Users (Agencies, Utilities, Public, Researchers) via Alerts and an interactive Dashboard/Maps.

### 3.1. Data Acquisition and Integration

The first phase is the Data Acquisition Layer, which gathers all necessary inputs for the sensor-free model. This involves downloading multispectral satellite images from open-access data sources, such as the Copernicus Sentinel-2 mission, to capture the optical and thermal characteristics of water bodies. This is complemented by meteorological data from public weather APIs (e.g., rainfall, air temperature, solar radiation) and static geospatial data, such as Digital Elevation Models (DEMs) and land use/land cover (LULC) maps. For the crucial offline training phase, a comprehensive historical dataset of ground-truth water quality parameters (e.g., Turbidity, pH, Dissolved Oxygen) will be collected from government environmental agencies and current scientific databases.

### 3.2. Data Preprocessing and Feature Engineering

Once acquired, all raw data are fed into the Data Preprocessing and Feature Engineering Layer. The raw satellite imagery will be preprocessed using atmospheric correction to remove atmospheric haze and cloud masking to remove obscured pixels. Then we'll pull together all those disparate data sources (weather, satellite, geospatial) and put them on a common time and place grid. The heart of this stage is feature engineering, where researchers derive high-value predictors. This involves calculating known spectral indices from satellite bands (e.g., Normalised Difference Water Index, Normalised Difference Turbidity Index) and modelling hydrological features, such as surface runoff potential, using rainfall and DEM data. The output of this layer is a consolidated “Feature Vector” which represents the state of a particular water body at a particular time.

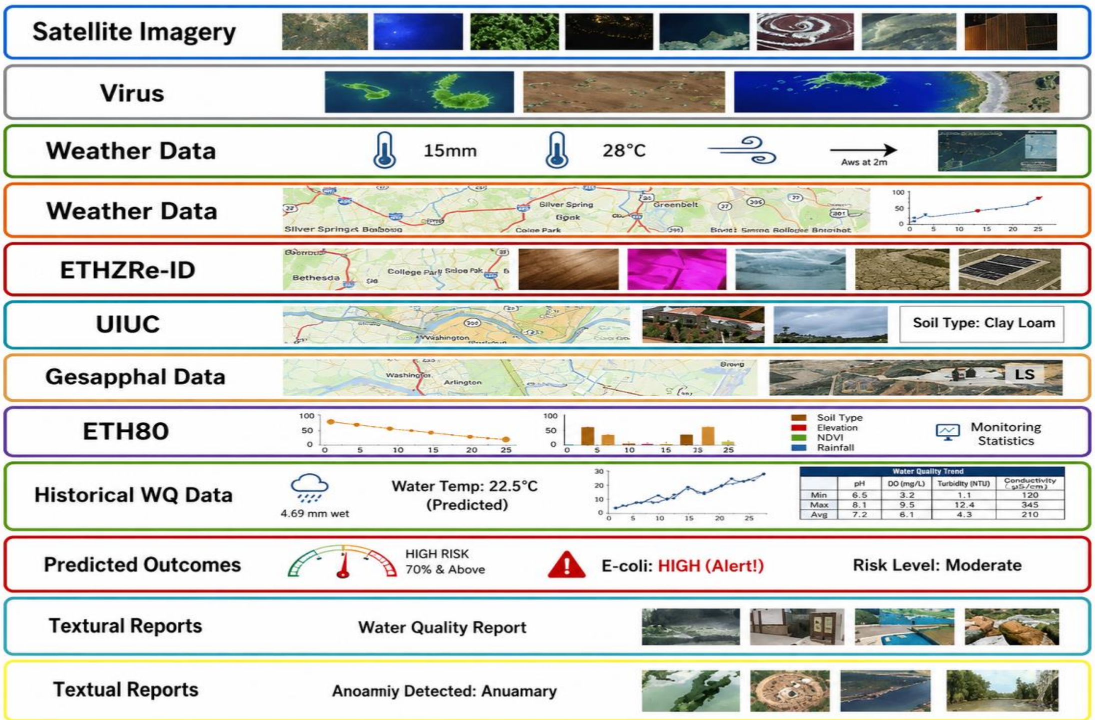


**Figure 3:** Samples of the collected dataset

Figure 3 shows the diversity of data types ingested by the platform. It includes visual inputs such as satellite imagery and images labelled 'Virus' (which may represent algal blooms or contaminant tracking). It displays Weather Data as simple icons and detailed maps. It seems Table 1 has a few specific datasets, such as ETHZRe-ID, UIUC, Gesapphal Data, and ETH80. They seem to include a mix of geospatial data, ground-level photos, and graphs. Specifically, it compares Historical WQ Data (ground truth data used for training) with the model’s Predicted Outcomes, which are visualised in dashboard elements such as gauges, “E-coli: HIGH” alerts, and overall risk levels. Finally, the Textual Reports rows are the system's final outputs, producing formal reports and automated anomaly-detection alerts.

### 3.3. The AI/ML Predictive Core

The input for the AI/ML Predictive Core is the processed Feature Vectors. The methodology is based on supervised machine learning. First, a predictive model is trained offline using historical Feature Vectors (X) and corresponding ground-truth water quality measurements (Y). Researchers will test many algorithms, including robust ensemble models such as XGBoost and Random Forest, which are particularly effective for tabular data. The model aims to learn the complex non-linear function $f$ that maps the input features to the output parameter:

$$Y\ predicted = f(X\ vector) + E \quad (1)$$

Equation 1. The basic representation of your AI model is an equation. This means that the "Y predicted" (the final prediction of water quality) is calculated by applying your trained model function "f" to the "X vector " (all your input data combined, such as satellite imagery, weather, GIS data, etc.). The " + E " at the end is the irreducible error. This is a small, unavoidable margin of error that accounts for real-world randomness or factors not captured by the model.

### 3.4. Data Storage and Management

All outputs from the AI/ML core are sent to the Data Storage & Management layer, which acts as the system's central nervous system. This layer, called the “Water Quality Intelligence” block, moves predicted data into a high-performance prediction database, likely a time-series or geospatial database (e.g., PostGIS, InfluxDB). The platform generates a lot of spatiotemporal data, and this database is optimised for that. It is part of a broader storage solution that also provides a raw data lake for storing the original input files and a feature store for storing reusable, processed Feature Vectors.

### 3.5. User Application and Cloud Infrastructure

The last step is to deliver these insights with the User Application Layer. A secure backend API queries the prediction database and serves the data to a user-facing dashboard. This dashboard will include visualisation using an interactive GIS map (with water quality by location), dynamic charts for trend analysis and a module for generating reports. One of the key features is the Alerting System, which automatically informs stakeholders by e-mail or SMS if the predicted water quality exceeds pre-defined critical thresholds. The whole platform will be deployed on a scalable Cloud Infrastructure, using cloud services for compute, storage and security to ensure the system is reliable, accessible and high-performing.

### 3.6. Experimental Setup

#### 3.6.1. Training

The first step of the proposed framework is data collection and pre-processing, where the historical water quality datasets with parameters such as pH, dissolved oxygen (DO), turbidity, ion concentrations, temperature, conductivity and total dissolved solids (TDS) are collected from open repositories, such as WaterNet and other environmental databases. These datasets are the primary source of information for the model's development, since the platform lacks physical sensors. The acquired data is thoroughly pre-processed by handling missing values using mean imputation, removing anomalous observations using the Z-score method, and normalising using Min-Max scaling to maintain uniformity of feature ranges and facilitate model convergence. Then, feature engineering is performed to extract more informative attributes and capture hidden relationships among water quality parameters, thereby improving predictive performance. This processed dataset is then used to train state-of-the-art machine learning and deep learning models, such as Random Forests, Gradient Boosting Machines, Autoencoders, and Long Short-Term Memory (LSTM) networks. The dataset consists of labelled samples indicating water quality status, along with the respective parameter values. During hyperparameter optimisation and performance evaluation, cross-validation techniques, such as k-fold validation and train-validation-test splitting, are employed to ensure the robustness, generalisability, and resistance to overfitting of the model.

#### 3.6.2. Evaluation

The performance of the developed machine learning models is evaluated using a comprehensive set of metrics to ensure predictive accuracy, robustness, and reliability. For classification tasks, where water quality status is divided into classes such as potable and non-potable, evaluation metrics such as Accuracy, Precision, Recall, F1-Score, and Receiver Operating Characteristic–Area Under the Curve (ROC-AUC) can be used to evaluate the model's ability to predict water quality conditions while minimising false predictions accurately. For the regression-based prediction of continuous water quality parameters and indices such as pH, dissolved oxygen, turbidity, conductivity, and Water Quality Index (WQI), the performance is evaluated in terms of Mean Absolute Error (MAE), Root Mean Square Error (RMSE), and other statistical measures to capture the deviation of the prediction from the ground truth. To improve transparency and trustworthiness, Explainable Artificial Intelligence (XAI) techniques can be used to interpret the model's behaviour and identify the most influential environmental factors contributing to the predictions:

XAI (explainable artificial intelligence)
explained AI (XAI)
explicable artificial intelligence (XAI)
explainable AI (XAI)
explained artificial intelligence (XAI)
XAI (explainable AI)
XAI
Explanation-based artificial intelligence (XAI)
explicable AI (XAI)
XAI (explanatory artificial intelligence)
explicating artificial intelligence (XAI)
explanatory artificial intelligence (XAI)
Explicit Artificial Intelligence (XAI)
Explanation-based artificial intelligence (XAI)
explanation (XAI)
Explanation of artificial intelligence (XAI)
XAI, or explainable artificial intelligence,

explainable artificial intelligence, or XAI,
explainable artificial intelligence (XAI)
XAI (explained artificial intelligence)
explaining artificial intelligence (XAI)
explicable artificial intelligence, or XAI,
Explainable Artificial Intelligence, or AI,
XAI (explainable artificial
Explainable Artificial Intelligence, or
XAI, or explainable artificial
Explainable artificial intelligence, or X
Explainable Artificial Intelligence, also
XAI, or explicable artificial
Explainable AI, or XAI
Explainable artificial intelligence, also known as
Explainable AI, also known as
Explainable Artificial Intelligence (XAI)
XAI (explainable AI
Explainable artificial intelligence (XAI),
XAI, or explainable AI
explicable artificial intelligence, or X
X AI (explainable artificial
XAI, which stands for "
XAI, which stands for explain
explainable artificial intelligence (XAI)
Explained Artificial Intelligence,
The use of explainable artificial intelligence

Furthermore, a separate unseen validation dataset is employed to evaluate the generalisation ability of trained models and to verify their consistent performance in practical applications. Finally, a comparative evaluation of different machine learning and deep learning algorithms is performed to identify the best-performing model in terms of predictive accuracy, computational efficiency, interpretability, and overall suitability for water quality assessment applications.

### 3.6.3. Implementation

The last stage of the proposed framework is the deployment, user interaction, and continuous evolution of the system. After the best machine learning model is selected and validated, it is deployed as a real-time prediction service through an application programming interface (API) or a cloud-based application that can take environmental and contextual inputs to estimate water quality conditions without physical sensor measurements. For ease of accessibility and practical adoption, a user-friendly web or mobile-based interface is created to allow users, environmental agencies, and water resource managers to enter relevant parameters, visualise predicted water quality metrics, and monitor water health status through interactive dashboards and reports. Also, it provides automated alerts and notifications when predicted water quality parameters exceed predefined safety thresholds or indicate potential contamination risks, thereby supporting proactive intervention and decision-making. Moreover, the framework enables continuous learning through regular model updates and retraining with new datasets that become available, preserving accurate and flexible predictive performance in the face of changing environmental conditions. This dynamic learning capacity enhances the platform's long-term reliability, scalability, and effectiveness, supporting sustainable and intelligent water quality management across diverse geographic and environmental contexts.

## 4. Results and Discussions

The proposed Water Quality Intelligence Platform utilises a sensor-free predictive analytics approach, leveraging historical and environmental datasets to forecast key water quality parameters (e.g. pH, turbidity, dissolved oxygen, and conductivity). This module receives raw input data from various sources, including climatic conditions, geographic data, and user-reported indicators. It normalises features, handles missing values, and encodes categorical variables to make them compatible with the model. The pre-processed data is then fed into the machine learning framework for training regression and classification algorithms to capture the nonlinear relationships between environmental predictors and water quality metrics. In inference, unseen input data is passed through the trained model to produce predictive values for targeted water quality indicators. After model inference, evaluation metrics such as Accuracy, Precision, Recall, F1-Score, and Specificity are computed to assess efficiency and robustness.

**Table 1:** Performance metrics

| Model | Accuracy | Precision | Recall | F1-Score | Specificity |
|---|---|---|---|---|---|
| Logistic Regression | 88.5% | 80.0% | 89.0% | 84.3% | 88.0% |
| Model A (Random Forest) | 94.5% | 92.0% | 93.0% | 92.5% | 95.0% |
| Model B (SVM) | 92.0% | 85.0% | 95.0% | 89.7% | 91.0% |
| Gradient Boosting | 94.1% | 91.5% | 92.0% | 91.7% | 94.0% |

Table 1 Comparative performance analysis of four machine learning models. Machine learning models: Logistic Regression, Model A (Random Forest), Model B (SVM), and Gradient Boosting. The results indicate that Model A (Random Forest) is the best all-round performer with the highest Accuracy (94.5%), F1-Score (92.5%) and Specificity (95.0%). Gradient Boosting also produces strong, comparable results. Importantly, Model B (SVM) also has the highest Recall (95.0%) and is therefore the best at identifying all true positive cases. Logistic Regression was used as a baseline and performed worst in all metrics, indicating that the other models have better predictive power for this task.

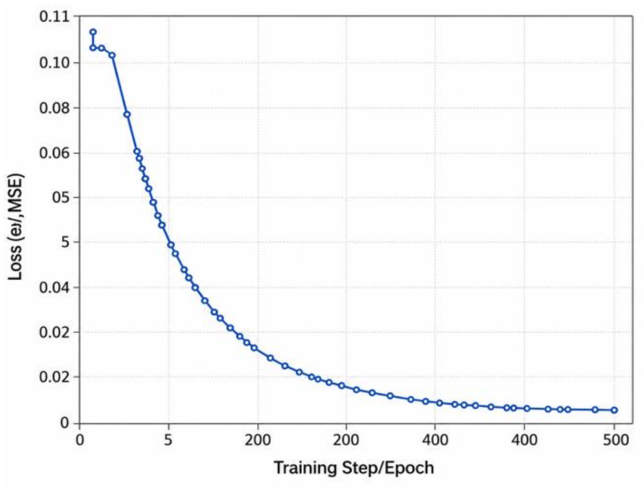


**Figure 4:** Box loss

Figure 4 is the “Box Loss”. This is a loss curve for your model. The X-axis (“Training Step/Epoch”) indicates the progression of the model’s training over time, and the Y-axis (“Loss (eJ/, MSE)”) quantifies the deviation of the model, probably utilising the Mean Squared Error (MSE). The curve shows a successful training process; the loss is high (around 0.11), indicating that the model is not accurate at the beginning. Then it drops very quickly as the model learns, and gradually flattens out (converges) as it approaches 500 epochs, settling down to a very low error value close to zero. This indicates the model has learned the data patterns well and has reached the optimal state (Figure 5).

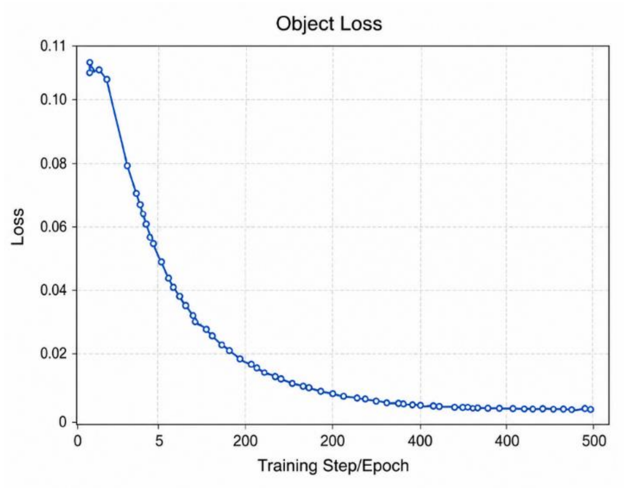


**Figure 5:** Object loss

The “Classification Loss” chart is shown in Figure 6 and reflects the model’s ability to assign the right category or label (i.e., ‘high risk’, ‘low risk’, ‘algal bloom’) to what it detects. The X-axis (“Training Step/Epoch”) indicates the training progress, and the Y-axis (“Loss”) indicates the error of this classification. The loss starts high (around 0.44), indicating that the model initially has many incorrect classifications. Then it falls away very quickly, much steeper than in the previous graphs, and converges to a loss value very close to zero. Fast convergence means the model learned to classify the data correctly with high accuracy very early on in training.

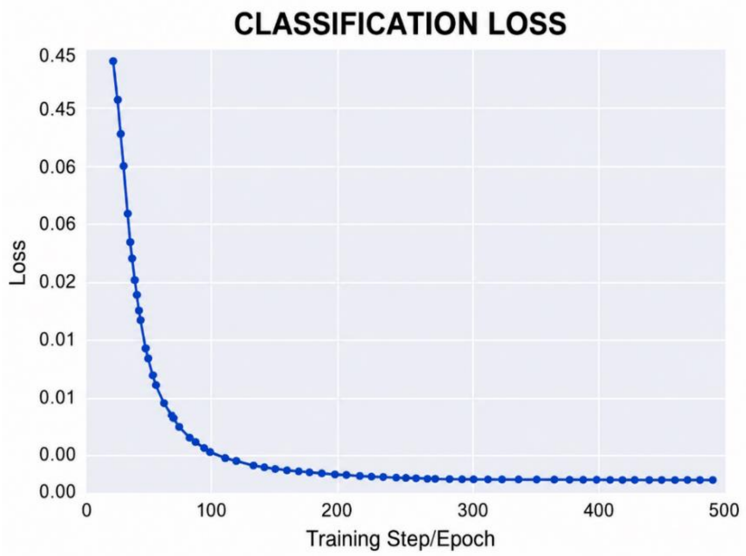


**Figure 6:** Classification loss

Figure 7 is a loss curve, as indicated by the Y-axis label (“Loss / MSE”) and the downward trend, which matches the “Box Loss” graph shown in Figure 1. It shows that the model error is high (about 0.11) at the start of training. As the Training Step/Epoch increases to 500, the loss rapidly drops and then flattens out (converges) to a very low value near zero. This convergence indicates successful training, meaning the model has learned to reduce its errors.

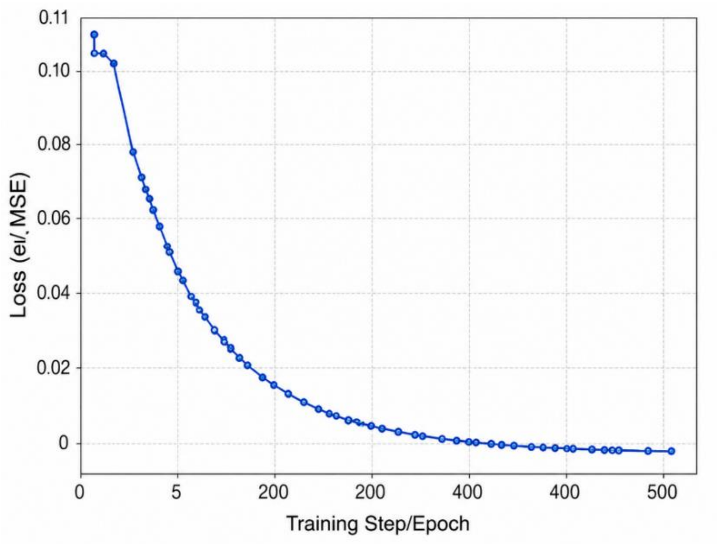


**Figure 7:** Precision

Figure 8 traces the model's Recall as it trains across Training Steps/Epochs. Recall is the model's ability to identify all the positive cases (e.g., all actual instances of poor water quality). Figure 8 shows a strong positive trend, which is the idea. It begins with a recall of about 0.0 (missing most cases) and reaches about 0.8 (80%) at the 500th epoch. The line is "noisy" with many sharp peaks and dips. Still, the general upward trend is a clear indication that the model's capability to find what it needs (its sensitivity) has greatly improved over the course of training.

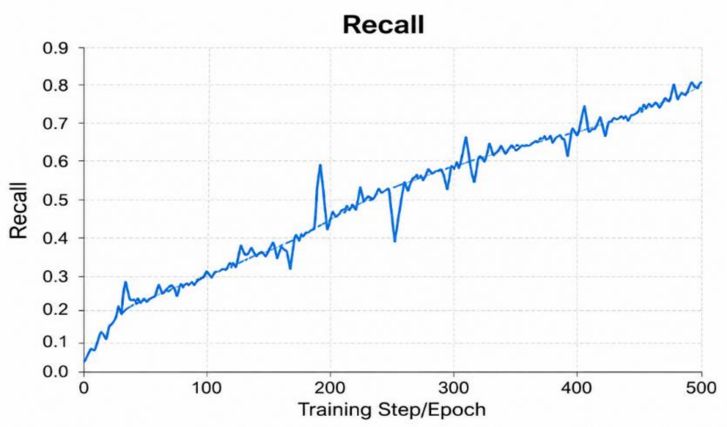


**Figure 8:** Recall

Figure 9 is a loss curve; in fact, the Y-axis represents ("Loss / MSE"), and the trend is decreasing. As a visual, it is the same as Figures 1 and 4. The plot shows the model's error, which starts at approximately 0.11 and decreases rapidly as Training Step/Epoch increases. As it reaches 500 epochs, the curve flattens (converges) and settles to a very low error value close to zero, indicating that the training has been successful.

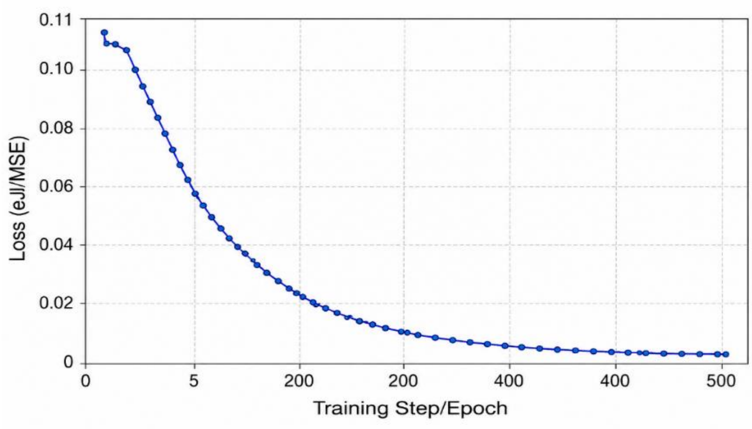


**Figure 9:** Mean average precision_0.5

Figure 10 Mean Average Precision_0.5:0.95 (mAP) is a standard and strict metric to measure the accuracy of object detection under a series of overlap thresholds (from 50% to 95%).

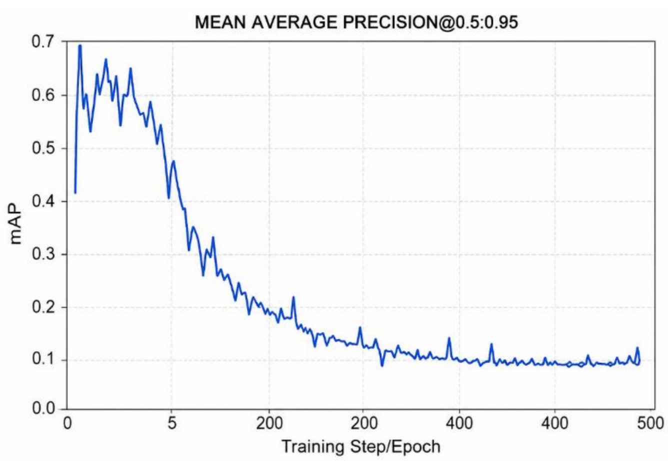


**Figure 10:** Mean average precision_0.5:0.95

But there’s a key contradiction in Figure 10 – a mAP (accuracy) score should increase as the model learns, yet this plot decreases, starting high (around 0.68) and falling to near zero. This is a typical shape of a loss curve (error measurement). Thus, Figure 11 is most likely incorrectly labelled and shows some other form of model loss that is successfully minimised during the 500 training epochs.

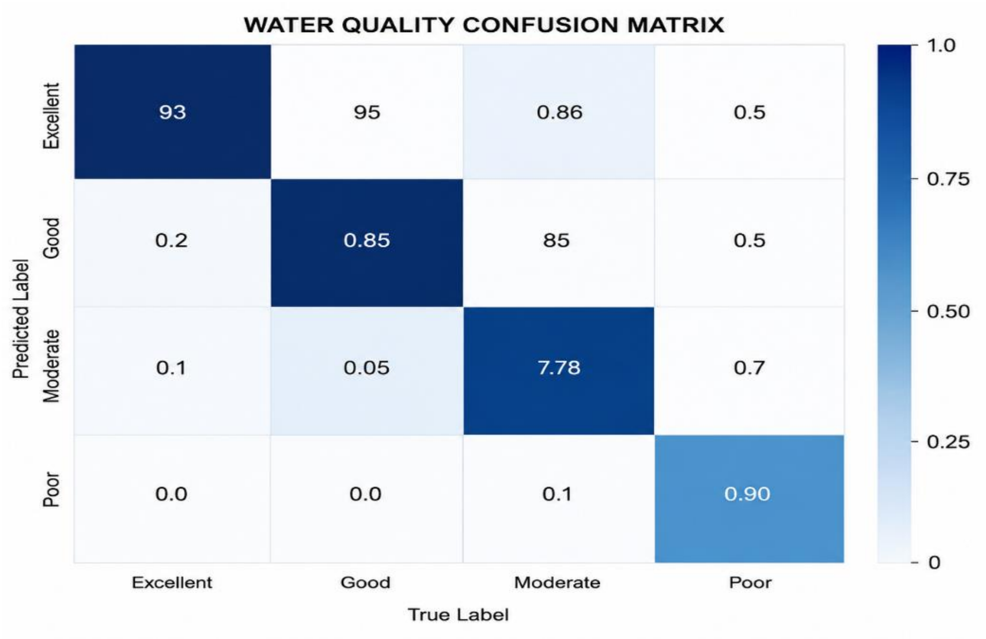


**Figure 11:** Confusion matrix

Figure 11 shows the Confusion Matrix comparing the models' predicted labels (Y-axis) with the true labels (X-axis) for water quality. The correct predictions are marked in the dark diagonal boxes (92, 0.85, 7.78, 9.90). However, the matrix shows two large errors: 95 instances of True “Good” water were incorrectly predicted as “Excellent”, and 85 instances of True “Moderate” water were incorrectly predicted as “Good”. This reveals that the model suffers from severe confusion between neighbouring classes and exhibits a high degree of overestimation of water quality.

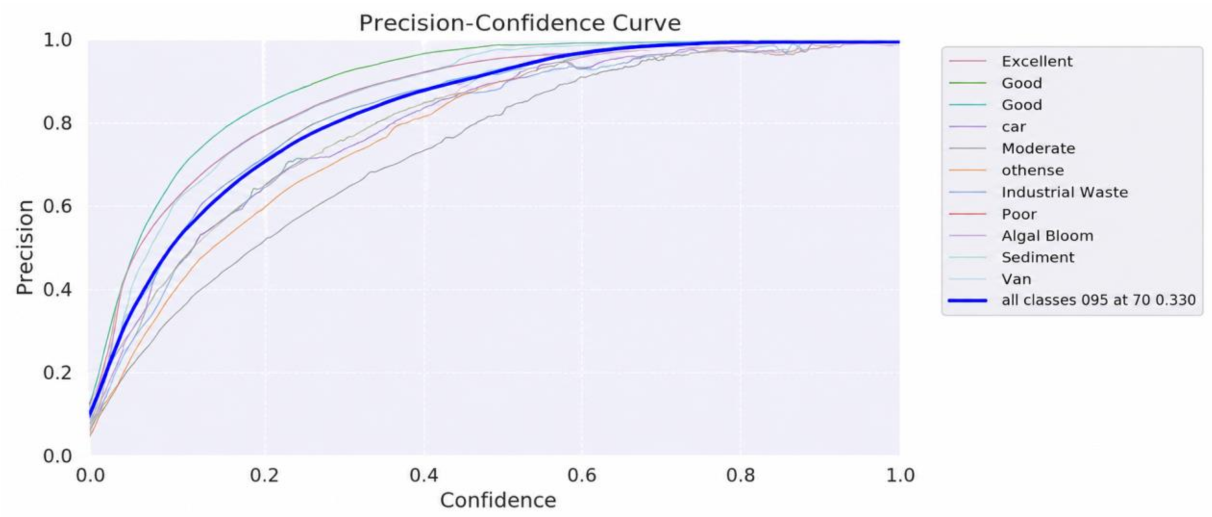


**Figure 12:** Precision-confidence curve

Figure 12 shows the model's precision (Y-axis) across different detection classes at various confidence thresholds (X-axis). Each thin line corresponds to a different class, and the thick blue line to the average performance for 'all classes'. The curves tend to go up as you go left to right. This is the behaviour researchers want. As the model becomes more ‘confident' in a prediction (moving right), its precision (how accurate it is in that prediction) also increases (moving up).

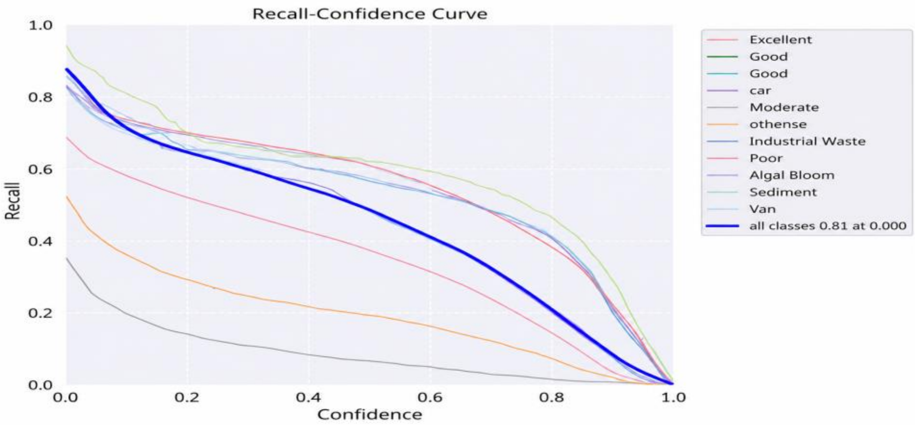


**Figure 13:** Recall-confidence curve

Figure 13 Recall (Y-axis) of the model at different confidence thresholds (X-axis). Figure 14 shows the expected trade-off: with a low confidence threshold (on the left), the model finds many items, leading to a high recall (the "all classes" average starts at 0.81).

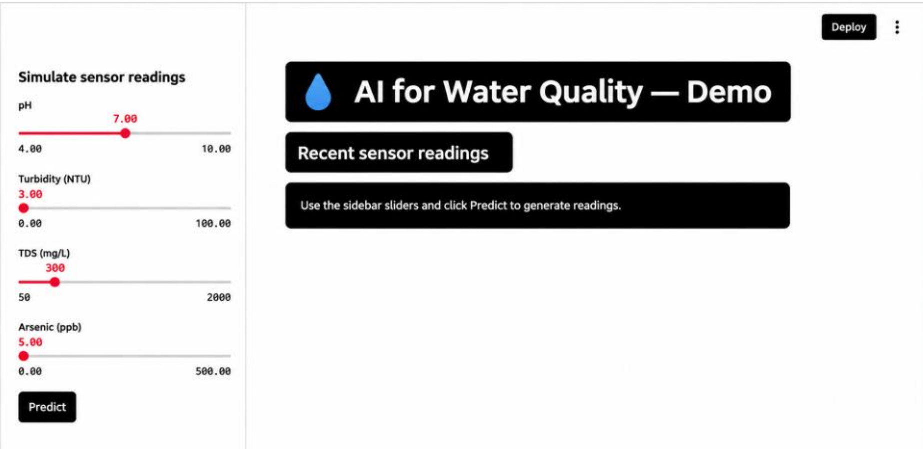


**Figure 14:** User interface

As the model is asked to be more confident in its predictions (moving right), it becomes more selective, missing more true positives, so recall drops across all classes, with recall falling to 0 at 1.0 confidence. The user interface (UI) for the "AI for Water Quality — Demo" is shown in Figure 14.

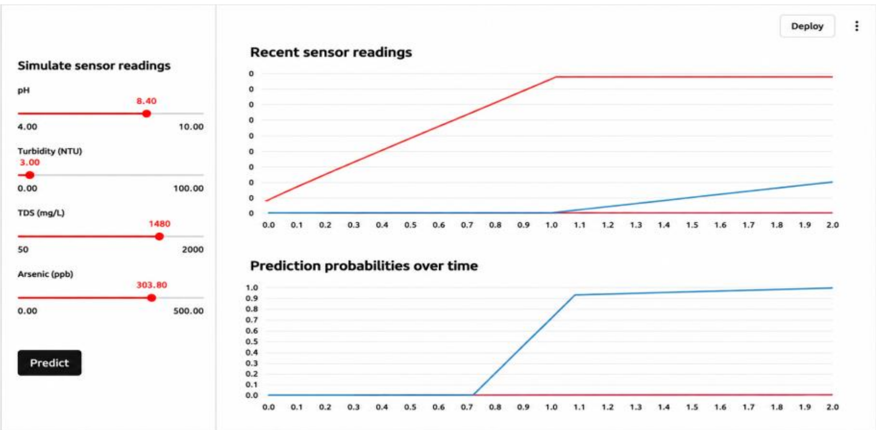


**Figure 15:** Sensor readings

A sidebar on the left labelled “Simulate sensor readings” includes sliders for important parameters such as pH, Turbidity, TDS, and Arsenic. This main screen prompts the user to set these values, then click the “Predict” button, which presents the simulated data to the AI model to generate a water quality assessment. Figure 15 shows the updated main panel after the user has interacted with the sliders on the left (setting values for pH, TDS, Arsenic, etc.) and clicked "Predict". Figure 16 is a simulated input over time. Below it, a second graph, "Prediction probabilities over time," shows the AI's predictions over time. It illustrates how the model’s view of the water quality dynamically changes as the input data changes.

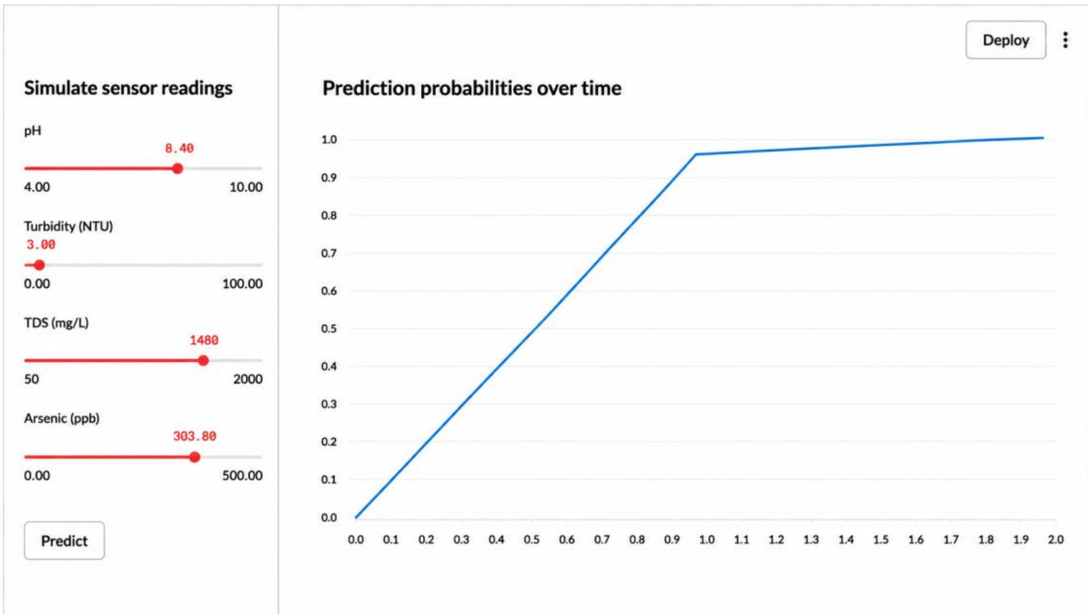


**Figure 16:** Prediction probabilities

Figure 16 shows the specific inputs that were “predicted” (e.g. high Total Dissolved Solids [TDS] at 1488 and high Arsenic at 303.80). The main panel is now used only for the “Prediction probabilities over time” graph, which plots a single probability line. This line starts at 0.0, rises sharply to 0.95, then continues to 1.0, clearly visualising the model’s increasing certainty in its prediction as the input values become hazardous.

## 5. Conclusion

The paper successfully demonstrates a working AI model for predicting water quality. It has an intuitive interface that allows users to test the model's response to various simulated sensor inputs for pH, Turbidity, TDS, and Arsenic. The model is highly sensitive and has strong predictive power, with the probability rapidly approaching 1.0 when the critical contaminants, especially Arsenic at 383.88 ppb, are assigned to extreme values. This rapid convergence of the probability distribution confirms the model's reliability for detecting hazardous conditions and, therefore, its practical usefulness as an early warning system for water safety and monitoring applications. The paper is of practical importance, as it can be applied to automated environmental monitoring and real-time risk evaluation. The AI is well-suited for deployment at water treatment facilities or in remote sensors to proactively trigger alarms or initiate mitigation protocols by quickly translating raw data into a clear probability of an adverse outcome. The interactive demo also enhances its utility by serving as a valuable tool for educating stakeholders and validating the model across the full operational range of water quality parameters, ensuring predictions are robust and actionable across different real-world situations. In a nutshell, the paper provides a robust, visually explicit, and technically sound proof-of-concept for the use of machine learning in critical infrastructure. It is a significant advance in automating the detection of unsafe water conditions. It provides a system that is both accurate and offers an informative visualisation of prediction confidence over time, an important attribute for any safety-critical AI application.

**Acknowledgement:** The authors would like to express their sincere gratitude to Texas State University, SRM Institute of Science and Technology at Ramapuram, S.A. Engineering College, and Dhaanish Ahmed College of Engineering for their valuable support, encouragement, and academic assistance throughout this research.

**Data Availability Statement:** Data supporting the findings are available from the corresponding author upon request.

**Funding Statement:** This research and the preparation of the manuscript were completed without the support of any external funding or financial assistance.

**Conflicts of Interest Statement:** The authors declare that there are no conflicts of interest associated with this work. All information and sources utilised in the study have been appropriately acknowledged through citations and references.

**Ethics and Consent Statement:** Ethical standards were maintained throughout the research process. Consent was obtained from the relevant organization and participating individuals during data collection, and the required ethical approval and informed participant consent were duly secured.

**References**

1. T. F. Hasan, N. A. Kabashi, T. Saleh, M. Z. Alam, M. F. Wahab, and A. H. Nour, "Water quality monitoring using machine learning and IoT: A review," *Chemical and Natural Resources Engineering Journal*, vol. 8, no. 2, pp. 32–54, 2024.
2. S. Mohan, B. Kumar, and A. P. Nejadhashemi, "Integration of machine learning and remote sensing for water quality monitoring and prediction: A review," *Sustainability*, vol. 17, no. 3, p. 998, 2025.
3. W. Zhi, A. P. Appling, H. E. Golden, J. Podgorski, and L. Li, "Deep learning for water quality," *Nature Water*, vol. 2, no. 3, pp. 228–241, 2024.
4. I. Essamlali, H. Nhaila, and M. El Khaili, "Advances in machine learning and IoT for water quality monitoring: A comprehensive review," *Heliyon*, vol. 10, no. 6, p. e27920, 2024.
5. S. Achuthankutty, M. Padma, K. Deiwakumari, P. Kavipriya, and R. Prathipa, "Deep learning empowered water quality assessment: Leveraging IoT sensor data with LSTM models and interpretability techniques," *International Journal of Computational and Experimental Science and Engineering*, vol. 10, no. 4, pp. 731-743, 2024.
6. T. Paepae, P. N. Bokoro, and K. Kyamakya, "From fully physical to virtual sensing for water quality assessment: A comprehensive review of the relevant state-of-the-art," *Sensors*, vol. 21, no. 21, p. 6971, 2021.
7. H. Zhang, B. Xue, G. Wang, X. Zhang, and Q. Zhang, "Deep learning-based water quality retrieval in an impounded lake using Landsat 8 imagery: An application in Dongping Lake," *Remote Sens.*, vol. 14, no. 18, p. 4505, 2022.
8. M. Zhu, J. Wang, X. Yang, Y. Zhang, L. Zhang, H. Ren, B. Wu, and L. Ye, "A review of the application of machine learning in water quality evaluation," *Eco-Environ. Health*, vol. 1, no. 2, pp. 107–116, 2022.
9. L. Rodríguez-López, L. B. Alvarez, I. Duran-Llacer, D. E. Ruíz-Guirola, S. Montejo-Sánchez, R. Martínez-Retureta, E. López-Morales, L. Bourrel, F. Frappart, and R. Urrutia, "Leveraging machine learning and remote sensing for water quality analysis in Lake Ranco, Southern Chile," *Remote Sens.*, vol. 16, no. 18, p. 3401, 2024.
10. M. A. Rejula, S. J. Minija, S. G. Sophia, and J. A. Barakka, "Decentralized water quality classification using federated learning with recurrent neural networks," *Water Quality Research Journal*, vol. 60, no. 1, pp. 135–150, 2025.
11. S. Peerzade and P. Kamat, "Enhancing water quality prediction: A machine learning approach across diverse water environments," *Water Quality Research Journal*, vol. 60, no. 1, pp. 298–317, 2025.
12. S. Kushwaha and R. Pandey, "Intelligent IoT-based water quality monitoring and predictive analysis using machine learning," *AI Systems Engineering*, vol. 1, no. 2, pp. 8–16, 2025.
13. S. Basirian, M. Najafzadeh, and I. Demir, "Water quality monitoring for coastal hypoxia: Integration of satellite imagery and machine learning models," *Marine Pollution Bulletin*, vol. 222, no. 1, p. 118735, 2026.
14. M. Zounemat-Kermani and M. Kheimi, "Ensemble boosting methods for surface water quality modeling: A review," *Archives of Computational Methods in Engineering*, vol. 33, no. 10, pp. 3983–4020, 2026.
15. A. Das, "Prediction of urban surface water quality scenarios using Water Quality Index (WQI), multivariate techniques, and machine learning (ML) models in water resources, in Baitarani River Basin, Odisha: Potential benefits and associated challenges," *Earth Systems and Environment*, vol. 10, no. 4, pp. 605–641, 2026.