



Deep Learning and IoT Enabled Real-Time Water Quality Surveillance and Decision Support Systems

Dr. Yuvaraj. B¹, Dr. D. C. Jullie Josephine², Dr. T. Ganesan³, Dr. Madhankumar⁴, Dr. Raghavendra M Ichangi⁵

Abstract

Clean water is a basic human right and a global emergency – over 2 billion people suffer from a lack of safe drinking water. Current water quality monitoring methods require manual sampling and laboratory analysis from time to time, and involve a lot of time lag and human error in the sampling process, making them ineffective in detecting contamination events quickly. The paper proposes a complete framework which combines the Internet of Things (IoT) sensor network, multi-tier edge-fog-cloud computing structure and the latest deep learning models to provide continuous, real-time, automatic and smart water quality monitoring and intelligence decision support. The system we propose uses eight different physicochemical sensor modules (pH, DO, turbidity, EC, NO₃, ORP, temp, heavy metals) at geographically distributed monitoring nodes. Readings are sent to edge gateways through quantized deep learning models which detect anomalies in less than 20 milliseconds and are transmitted using time series technology over LoRaWAN and 4G/LTE. It utilizes a novel Temporal Fusion Transformer (TFT) along with attention-augmented BiLSTM layers (TAT-BiLSTM), which are trained using more than 820,000 labeled data from five different monitoring stations over a two-year period.

Experimental testing shows that the proposed TAT-BiLSTM model makes a classification accuracy of 97.3%, with a precision of 0.971 and a recall of 0.975 for five water quality index categories, outperforming all the baseline deep learning architectures by 1.7–6.5 percentage points. The system is capable of calculating a real-time composite Water Quality Index and in the event of abnormal conditions, alerts authorities with SMS/email messages via multiple levels within seconds. The decision support layer is completed by integrating with a web dashboard with GIS functionality and mobile application. From sensor reading all the way to the notification to regulatory authorities and utilities, the system's end-to-end latency is less than 25 seconds, allowing both regulatory agencies and utilities to be proactive in addressing water safety issues.

The proposed framework goes beyond monitoring, and focuses on predictive intelligence and adaptive environmental management. The system not only monitors the current water quality but also continually processes the water quality data from the sensors to predict the likelihood of a contamination event occurring before it reaches critical levels. This predictive power can help regulatory bodies, water companies, and environmental organisations to take proactive action and not just a reactive one. Additionally, the framework uses state-of-the-art temporal learning algorithms to identify subtle seasonal trends, periodic pollution patterns, and long-term degradation trends, which are not captured by traditional threshold monitoring systems. This makes the architecture proposed to be appropriate for both environmental surveillance and emergency response applications. The scalable and modular design is another important contribution in this work, and provides a foundation for the deployment of the monitoring network throughout rivers, lakes, reservoirs, groundwater resources, industrial discharge sites and urban drainage systems without extensive changes. As a sensor node is autonomous in its operations, and connected to each other by a shared communication infrastructure, the addition of another monitoring station to the system does not impact the system performance. Moreover, edge computing can drastically cut down on communication overhead because it can filter out anomalies and other information at the edge and only send actionable data to higher layers of processing. This helps to reduce bandwidth usage and aids system reliability in remote locations with weak internet connection. The proposed architecture can be operated with energy-efficient hardware that saves power on the sensors, intelligent duty cycling, and solar-powered deployment, which reduces operational costs and maintenance demands from a sustainability point of view. The decision support platform also makes it easy to use; intuitive dashboards, historical trend visualization, automated report generation, GIS-based spatial mapping and configured alert mechanisms facilitate the easy interpretation by non-technical stakeholders. In summary, the proposed IoT–Deep Learning framework illustrates how the current artificial intelligence tools can revolutionize the traditional approach to environmental monitoring, creating an intelligent, autonomous and proactive environmental monitoring system that can help to support evidence-based policy decision, enhance public health protection and sustainable management of precious freshwater resources.

¹ Professor, Department of Artificial Intelligence and Machine Learning, Kings Engineering College, Sriperumbudur, Chennai – 602117, Tamil Nadu, India, Email: byuvarajb@gmail.com

² Professor, Department of Computer Science and Engineering, Kings Engineering College, Irungattukottai, Sriperumbudur, Chennai – 602117, Tamil Nadu, India, Email: drdcjulliejosephine@gmail.com

³ Professor and Head, Department of Artificial Intelligence and Data Science, Kings Engineering College, Irungattukottai, Sriperumbudur, Chennai, Tamil Nadu, India, Email: ganesan21@gmail.com

⁴ Professor, Department of Computer Science and Engineering, Madras Engineering College, Chennai, Tamil Nadu, India, Email: manomadhanbe@gmail.com

⁵ Assistant Professor, Department of Artificial Intelligence and Machine Learning, School of Computer Science, Central University of Karnataka, Kalaburagi, Karnataka, India, Email: rmichangi@cuk.ac.in

Keywords: Internet of Things (IoT), Water Quality Monitoring, Deep Learning, Edge–Fog–Cloud Computing, Temporal Fusion Transformer (TFT), Attention-Augmented BiLSTM (TAT-BiLSTM), Water Quality Index (WQI).

1. Introduction

Water is the lifeline of life, public health, agricultural productivity and industrial sustainability. According to the World Health Organization (WHO), 485,000+ children under 5 years of age die annually from diarrhoea, which is caused by unsafe water, poor sanitation and the lack of hygiene. The frequency and intensity of freshwater contamination events have been dramatically increased in recent years around the world due to rapid urbanisation, climate variability, industrial proliferation and intensive agriculture. Rivers, lakes, groundwater aquifers and reservoirs, which were once thought to be secure sources of water, are becoming very susceptible to an acute spill of chemical pollutants, agricultural run-off containing nitrates, pesticides and other chemicals, microbial contamination from sewage overflows, and heavy metal contamination from mining activities.

The current conventional water quality monitoring systems have inherent shortcomings. Even if carefully designed and followed every week or month, grab sampling programs will fail to detect transient pollution events that can last for hours, but may have long-lasting deleterious effects. Laboratory-based analysis is extremely accurate, but it takes 24 to 72 hours for results to be provided to decision-makers, allowing contaminated water to get into distribution systems, irrigation systems, and/or natural ecosystems. Bench-top sensors which can make continuous in-situ measurements are available, but are expensive, difficult to calibrate, and require a lot of maintenance, making them impractical to use.

Combined with three important technological shifts: low-power IoT sensing hardware, energy-efficient wireless communications and deep learning, this represents a previously unmatched opportunity to address these challenges. In addition to being able to be manufactured for less than \$80, IoT sensor nodes can be equipped with multiple parameters sensors, such as temperature, pH, turbidity, conductivity, and dissolved oxygen, and with low power wide area network (LPWAN) technologies like LoRaWAN and NB-IoT, allowing for long range, battery powered telemetry from remote locations. Deep learning models trained using multivariate time-series data can classify water quality states and predict future contamination risks with human-expert-level accuracy.

Though progress has been made in all of these areas, there are still many ongoing research challenges in the area of integrating the advances in these domains into end-to-end operational water quality surveillance systems. However, data fusion across sensors in the network, dealing with missing sensor readings, model performance over time (concept drift), a limited number of computational resources at the network edge, response time for immediate alerts, and the ability to translate the model output to decision making by non-technical users like water utility operators and municipal health officials, are all challenges.

Sensing technologies and wireless communications networks have become readily available in recent years, and have revolutionized environmental monitoring over the last decade. Smart monitoring infrastructures are becoming widespread in government, environmental protection agencies and research institutions, as they are able to gather continuous real-time streams of environmental data, collected from geographically distributed locations. The key technological developments have been especially useful in water quality monitoring, as aquatic ecosystems are very dynamic and are affected by rainfall, industry, agriculture, seasonality and hydrological changes due to climate change. Therefore, such monitoring systems need to be sensitive to these fast changes, with a high time resolution, and reliable in harsh conditions. Large-scale environmental datasets are now being used successfully for extracting meaningful information, and one of the powerful analytical tools is now Artificial Intelligence (AI) with the emergence of deep learning. Deep neural networks learn multiple water quality parameters' complex nonlinear relationships without the need for any hand-crafted features or linear assumptions used in traditional statistical models. Water quality applications are particularly useful for these capabilities because the relationship between the various parameters measured – such as dissolved oxygen, turbidity, conductivity, oxidation-reduction potential, and nutrient concentrations – can be extremely interdependent. Deep learning models learn these relationships directly from observations in the past and are very accurate in their predictions whilst being flexible to the changes in environment. In addition, edge computing has emerged as one of the biggest pain points of cloud-based IoT systems in recent times - communication latency. Computational intelligence can now be integrated within edge gateways positioned near the monitoring stations, rather than communicating all the readings of the sensors to central servers. This distributed processing model helps detect anomalies in real time, notify emergency responders in real time, and lessen the need for constant Internet connectivity. These capabilities are especially valuable for systems located on remote rivers, reservoirs and rural water supplies where communications could be problematic. The combination of IoT sensing, edge intelligence, cloud computing, and AI frameworks that rely on deep learning, therefore, are a big step towards creating autonomous environmental surveillance systems that enable real-time decision making. The proposed research adds to this emerging field in the integration of these complementary technologies within a comprehensive framework with a focus on scalability, operational efficiency, predictive analysis and viable implementation for water resources management in sustainable development.

1.1 Objectives

The objectives of this research are to:

- Design and test a distributed sensor network for continuous monitoring of multi-parameter water quality for riverine, lacustrine and groundwater applications.
- Propose a new deep learning model (TAT-BiLSTM) for real time classification of water quality status and prediction of Water Quality Index (WQI) values that achieves state-of-the-art accuracy.

- Develop a multi-stage edge-fog-cloud processing pipeline which incurs minimal end-to-end detection and alerting latency and works within the power and cost limits of remote deployment.
- Develop a decision support system that can be used in conjunction with GIS, with intuitive visualization, automatic reporting and configurable notification processes for water quality authorities.
- Test the performance of the system in the field, considering issues such as sensor drift, network intermittency and seasonal changes in water quality.

1.2 Paper Organization

The rest of this paper will be organized around the following. The relevant literature has been reviewed in Sections 2 that cover the IoT water monitoring, deep learning for environmental sensing and edge computing. The detailed system architecture is given in section 3. The methodology used for the deep learning is outlined in Section 4. Experimental set up and data sets are described in section 5. Reports and discusses results in Section 6. Limitations and future work are discussed in section 7. The paper is concluded with section 8.

2. Literature Review

2.1 The water quality data were collected using the internet of things (IoT) technology.

Initial water quality monitoring systems were based on IoT and deployed single sensor that triggered alerts based on thresholds. Postolache et al. (2014) showed the possibility of real time monitoring of turbidity by acoustic sensor array based on GSM telemetry. Vijayakumar and Ramya (2015) later developed similar Arduino-based sensor-microsystems that were connected to the cloud-based platform, ThingSpeak. These systems demonstrated proof-of-concept viability, but limited flexibility to complex scenarios of contamination with multiple parameters. LPWAN technologies significantly increased the scope of operation and cost of water quality sensor networks. Sanchez et al. (2020) deployed LoRaWAN connected sensor nodes over a 12 km river stretch where they were able to obtain a packet delivery ratio of 97.8% and have node battery life of over 18 months. The NB-IoT deployments by Zhang et al. (2020) showed that NB-IoT could deliver high-quality data transmission in underground infrastructure such as pipes and monitoring wells. These works laid the ground work for the communication infrastructure on top of which intelligent, ML-based analytics can be added.

The recent developments in the field of IoT technologies have greatly improved the efficiency, scalability and cost-effectiveness of environmental monitoring systems. Today's sensor networks can measure many water quality parameters at high frequencies and transmit the data in real-time to centralized or distributed processing platforms. The IoT-based infrastructures can also provide a continuous monitoring capability, which means that authorities can be alerted to contamination events as soon as they happen; whereas traditional monitoring infrastructures require human intervention and periodic checks. Data collection can also be done continuously, which allows a better understanding of the temporal dynamics of aquatic ecosystems, including diurnal variations, seasonal changes and pollution events that might not be noticed with traditional sampling methods. Another great benefit of IoT monitoring is that it can incorporate all kinds of dissimilar sensing devices into a single communication structure. Today, the environmental monitoring stations are usually integrated with electrochemical sensors, optical probes, weather stations, GPS modules and wireless communication interfaces into an energy-efficient, small platform. This kind of integration allows to monitor the physicochemical characteristics simultaneously with meteorological and geographical data, offering a more complete picture of the environmental conditions. Moreover, data from different monitoring stations can be connected to the cloud and then accessible to be centrally analyzed, and comparative large-scale analyses made across various watersheds and geographical areas. The drastic change of low-power communication technology such as LoRaWAN, NB-IoT, ZigBee and LTE-M has extended the transmission distance of monitoring system for IoT applications, while keeping energy consumption low. These technologies can be used to ensure data will be transmitted reliably between distant locations, where traditional communication means may not exist. In addition to energy harvesting from the sun, effective power management algorithms, modern sensor nodes in IoT networks can run for several months or even years without demanding much maintenance. In sum, the Internet of Things (IoT) technology is the backbone of next-generation water quality monitoring systems that can enable sustainable management of water resources, disaster prevention and preparedness, compliance with regulations, and evidence-based water quality management at regional and national levels.

2.2 Machine Learning for Water Quality Prediction

The machine learning methods for water quality prediction have changed significantly in the past few years, from simple statistical learning models to complex deep learning models. The traditional approaches such as support vector machines (SVM), random forest (RF), and gradient-boosted trees (GBT) showed good performance with tabular water quality datasets, where temporal dependencies are relatively simple, but were less effective at handling the more complex temporal relationships found in continuous sensor streams. Support vector machines (SVM), random forest (RF), and gradient-boosted trees (GBT) are traditional methods that have been shown to be effective in tabular water quality datasets, which often have simpler temporal relationships, but were found to be less effective in continuous sensor streams with more complex temporal relationships. Gers et al. (1999) introduced LSTM networks, which has been widely used for prediction of environmental time-series since then. The LSTM models performed significantly better than the random forest and ARIMA baselines, achieving better predictions for dissolved oxygen than other models on 72-hour forecasting horizons, as early as 2019 when recurrent neural networks were first used for predictions (Zhu et al.). The attention mechanisms proposed by Bahdanau et al (2015) and adopted in the Transformer (Vaswani et al. 2017) were also introduced to give further boost to the models' performance by allowing them to choose informative time steps to attend to, which is

especially beneficial for detecting subtle signs before a contamination event.

The convolutional neural network is used in the spatial water quality mapping especially in the remote sensing application where synoptic scale coverage of a large water body is available from satellite images. On multi-site monitoring data, hybrid CNN-LSTM architectures (Li et al., 2021) also achieved the best results, as they were able to capture spatial features from the arrays of sensors and temporal features from the individual time series. A summary of the directly relevant prior works is given in Table 1.

Machine learning is reshaping environmental data analysis by allowing for the automatic identification of complex patterns which are difficult to be identified using traditional statistical methods. It is important to recognize that water quality data are inherently multidimensional, nonlinear and time-correlated, as multiple environmental variables change continuously under varying climatic and anthropogenic conditions. The key lies in the ability of machine learning algorithms to learn the complex relationships directly from past observations, without needing to manually specify mathematical equations. With the amount of environmental data being collected from sensors rapidly growing, machine learning is essential for intelligent water resource management. These capabilities have been further extended by the use of deep learning techniques, which have added hierarchical feature learning mechanisms that can learn meaningful temporal and spatial representations from raw measurements from the sensors. Various recurrent neural networks such as Long Short-Term Memory (LSTM), Gated Recurrent Units (GRU), Bidirectional LSTM, Transformer architectures and hybrid deep learning models have shown impressive results in predicting water quality parameters and categorizing water pollution incidents. Their long-term temporal dynamics allow for accurate forecasting of long-term environmental changes, and attention mechanisms enable them to find the most relevant part of the past to help with the decision making process. However, there are still a number of issues, which affect model performance in real deployments. Many environmental sensor measurements have missing data, noisy data, calibration drift, seasonal fluctuations and significant class imbalance as extreme pollution events are very rare. For this reason, good machine learning algorithms need strong preprocessing methods, normalization of the characteristics, filling in missing values and continuous updates of the models to ensure that their predictions remain accurate over time. More recently, research efforts have been directed towards developing techniques in the area of explainable artificial intelligence, which help to make the systems more transparent by revealing what environmental factors are most relevant to the prediction results. This is especially crucial in regulatory bodies that need to be comprehensible before taking decisions for operations. Thus, it is crucial to incorporate sophisticated deep learning models with robust explainability and preprocessing methods into the development of intelligent systems for water quality monitoring in the future.

2.3 Edge Computing for Environmental IoT

Cloud-based IoT solutions add latency, bandwidth usage and availability issues that are not suitable for environmental monitoring applications with strict time requirements. The term 'edge computing' was coined by Shi et al. (2016) and their work described the benefits of using edge computing with applications of IoT where low latency, efficient bandwidth usage and unreliable connectivity are required. The OpenFog Consortium's definition of fog computing is the addition of an intermediate layer of processing between the edge devices and cloud data centers.

Edge inference can be used for water quality applications for immediate anomaly detection without round trip to remote servers. By employing model compression methods such as pruning, quantization, and knowledge distillation, it is possible to deploy complex deep learning models on resource-limited edge devices. Optimized INT8 quantization support for inference with 4-8X model weight compression with minimal accuracy loss is available through TensorFlow Lite and ONNX Runtime.

2.4 Comparative Analysis of Prior Works

Table 1: Comparative Analysis of Related Works

Author(s)	Year	Method	Accuracy	Parameters	Real-Time	Limitation
Zhang et al.	2020	SVM + IoT	87.3%	pH, DO, Temp	Partial	No deep learning
Kumar & Singh	2021	ANN	89.1%	6 params	No	Offline only
Li et al.	2021	CNN	91.5%	Turbidity, pH	No	Single site
Rao et al.	2022	LSTM	92.8%	8 params	Yes	High latency
Chen et al.	2022	GRU + Edge	93.4%	7 params	Yes	No anomaly
Patel & Das	2023	BiLSTM	94.1%	9 params	Yes	Cloud only
Proposed System	2024	Transformer+Edge	97.3%	8 params + meta	Yes	—

Table 1 shows there is no existing system that can support real-time processing, high multi-parameter accuracy and edge deployment at the same time. From the above table, it is evident that no existing system can support all three stated attributes simultaneously: real-time processing, accuracy of multi-parameters, and edge deployment. All three requirements are satisfied in our proposed system by employing the novel TAT-BiLSTM architecture, as well as multi-tier deployment strategy.

3. System Architecture

The proposed system is designed to be hierarchical and has four layers: physical sensing, edge processing, fog aggregation and cloud analytics with a combination of heterogeneous communication protocols, optimized according to the layer latency, bandwidth and power.

3.1 IoT Sensor Node Design

Each monitoring node is equipped with a microcontroller unit (MCU) of the ESP32-S3 system-on-chip, with a dual Xtensa LX7 240 MHz cores, 512 KB of SRAM and 16 MB of flash. The ESP32-S3 supports the I²C, SPI, UART and SDI-12 interfaces of the different sensors we use, allowing us to connect them straight to the microcontroller. Our sensor suite is heterogeneous, and the ESP32-S3 supports the I²C, SPI, UART and SDI-12 interfaces of those sensors. Main core is always running in deep sleep mode to only 1.8mW at 1Hz sampling with a dedicated co-processor managing always-on sensor polling and analog to digital conversion. Energy autonomy is virtually limitless in sites where the daily average amount of sun is more than 2.5 kWh/m² with solar panels of 5W and 3.7 V/10 Ah LiFePO₄ battery.

Two point and three point sensor calibration is performed in accordance with the manufacturers' instructions with certified buffer solutions and standard references. An on-node Kalman filter is used to smooth raw sensor readings and flags those that are set to be more than 3σ from a rolling 24-hour baseline for second level verification. The specifications for all eight of the sensors measured are summarized in Table 2.

Table 2: IoT Sensor Module Specifications

Parameter	Sensor Module	Range	Accuracy	Protocol	Sampling (Hz)
pH	Atlas Scientific EZO	0–14	±0.002	UART/I ² C	1
Dissolved Oxygen	DO Pro CV	0–20 mg/L	±0.05	SDI-12	0.5
Turbidity	DFRobot SEN0189	0–1000 NTU	±5%	Analog	2
Temperature	DS18B20 Waterproof	-55–125°C	±0.5°C	1-Wire	4
Conductivity	Atlas EZO-EC	0.07–500K μ S	±2%	I ² C	1
Nitrates	ISUS V3 Nitrate	0–2000 μ M	±2 μ M	RS-232	0.25
Heavy Metals	HMTS-1000	ppb level	±3%	SPI	0.1
ORP	Atlas EZO-ORP	-2000–2000mV	±1 mV	UART	1

3.2 Communication Infrastructure

The data packets are sent from the sensor nodes encrypted (AES-128) using LoRaWAN Class A on IN865 band. The LoRaWAN gateway (Dragino LPS8N) receives data from 8 channels and can cover the area of 5 km in suburban area and 15 km in open rural area. Gateway to fog communication via MQTT TLS 1.3 on 4G LTE with automatic failover to Wi-Fi (if available). The MQTT broker (Eclipse Mosquitto) is installed on the fog server, where all the sensor topics are guaranteed at least once delivery (QoS Level 1). Data packets are compactly serialized as Protocol Buffers (protobuf) which are binary rather than JSON, with a payload being 65% smaller than that of JSON.

3.3 Multi-Tier Processing Pipeline

There are 3 tiers of processing to balance the latency and computational complexity. The edge tier detects anomalies in real time with a 94.1% accuracy rate and inference latency of 12.4ms with a quantized (INT8) LSTM model. Anomaly detection done in real time with an accuracy of 94.1% and inference latency of 12.4ms using a quantized (INT8) LSTM model, and local alerting for critical threshold violations. The fog tier utilizes the full-precision TAT-BiLSTM model to classify WQI for nuanced multi-class classification using readings from all nodes in a watershed to identify spatial pattern that would allow for detection of diffuse source contamination. The cloud tier hosts the web dashboards and reporting APIs, as well as a time-series database, InfluxDB, for

historical data storage; it also queues up new models for retraining as labeled data becomes available and runs long-term trend analysis.

Table 3: Edge vs Fog vs Cloud Processing Tier Comparison

Processing Tier	Avg Latency (ms)	Throughput (msg/s)	Power (W)	Cost/Month	Reliability
Sensor Node (MCU)	2.1	50	0.8	\$0.10	99.1%
Edge Gateway (RPI4)	8.4	500	5.2	\$2.40	99.5%
Fog Node (Jetson Nano)	14.6	1200	10.0	\$8.70	99.7%
Regional Cloud	45.2	10000	—	\$35.00	99.9%
Central Cloud (AWS)	82.7	50000	—	\$120.00	99.99%

3.4 Water Quality Index Computation

The Water Quality Index (WQI) is calculated on real time basis by applying weighted arithmetic mean which is standardized by the Central Pollution Control Board (CPCB). A weight (w_i) is assigned to each parameter based on its public health relevance and a sub-index (q_i) is calculated by multiplying the measured value of the parameter by a scaling factor that results in a value within the range from 0 to 100. The composite WQI is actually computed as:

$$WQI = \frac{\sum(w_i \times q_i)}{\sum w_i} \quad \text{where} \quad q_i = \left[\frac{V_i - V_o}{S_i - V_o} \right] \times 100$$

V_i is the concentration of parameter i measured, V_o is the ideal concentration in pure water and S_i is the permissible standard. The thresholds used to classify WQI and their implications for operation are shown in table 4.

Table 4: Water Quality Index (WQI) Classification System

WQI Range	Category	Quality Status	Recommended Use	Alert Level
0–25	Excellent	No pollution	Drinking/All uses	None
26–50	Good	Minor pollution	Drinking (treated)	Green
51–75	Medium	Moderate pollution	Industrial/Irrigation	Yellow
76–100	Bad	Significant pollution	Restricted use	Orange
100+	Unfit	Heavily polluted	No use	Red

4. Deep Learning Methodology

4.1 Problem Formulation

Suppose that $X = \{x_1, x_2, \dots, x_n\}$ is a multivariate time series, such that $x_t \in \mathbb{R}^d$ is an 8-dimensional ($d = 8$) vector of sensor readings at time t , corresponding to the 8 parameters monitored. Water Quality classification problem is formulated as a supervised multi-class sequence classification problem: The task is to classify a sequence of 60 consecutive observations $X_t = \{x_{t+1}, \dots, x_t\}$ into one of the 5 categories of water quality (Excellent, Good, Medium, Bad, Unfit) according to the category label \hat{y}_t . Also, there is a regression head that forecasts the cont. WQI value $y \in [0, 200]$.

4.2 TAT-BiLSTM Architecture

The proposed Temporal Attention Transformer with Bidirectional LSTM (TAT-BiLSTM) architecture has four functional modules:

- The Input Encoding Module is a layer of linear projections that takes raw sensor readings and projects them into a 256-dimensional embedding space. Learnable positional encodings are used to ensure that the information of temporal ordering is maintained in the input window.
- Temporal Fusion Transformer Block: Six layers of Transformer encoders, each having 8 attention heads, are stacked to process the embedded sequence. The model's ability to learn long-range temporal relationships between far-flung timesteps is achieved by using multi-head self-attention. Multi-head self-attention is employed to capture long-range relationships and delayed contamination effects between distant timesteps.
- Bidirectional LSTM Refinement: Two-layer BiLSTM with hidden dimension 256 is used to refine the output representations of the Transformer stack. Bidirectional processing enables the LSTM to use both causal context (forward) and anti-causal context (backward) to enhance feature representation at each time step, when processing complete input windows.
- The Multi-Task Output Heads contain a 3-layer MLP with softmax activation to get class probability estimates,

and a parallel regression head with sigmoid activation to get a normalized WQI value. To ensure a balance between classification and regression accuracy the joint loss function is $L = \alpha \cdot BCE + \beta \cdot MSE$, with $\alpha = 0.7$ and $\beta = 0.3$.

4.3 Training Configuration

All the hyper parameters used to train the model are summarized in table 7. All the experimental configurations resulted in a stable convergence using the AdamW optimizer with cosine decay learning rate scheduling and warm-up.

Table 5: Model Training Hyperparameters

Hyperparameter	Value	Hyperparameter	Value
Learning Rate	0.0001	Batch Size	64
Optimizer	AdamW	Epochs	200
Attention Heads	8	Dropout Rate	0.2
Hidden Layers	6	Embedding Dim	256
Sequence Length	60	Loss Function	MSE + CE
LR Scheduler	CosineDecay	Gradient Clipping	1.0

4.4 Data Preprocessing Pipeline

Before ingesting the raw sensor streams into the training workflow, they are run through a 6-stage preprocessing pipeline. Firstly, duplicate reading, due to network retransmission, is detected and eliminated by using the timestamp. Secondly, missing data values are filled in with a physics-informed interpolation scheme based on cubic spline interpolation combined with domain knowledge constraints such as a dissolved oxygen constraint that depends on temperature and its solubility limits. Third, all the parameters are converted to SI units and all the dissolved concentrations are standardized to mg/L. Fourth, all the parameters are per-sensor normalized using a robust (median/IQR) scaling to minimize the effect of outliers. Fifth, the input sequence is of length 60 (60 minutes of 1-minute reading), which results in overlapping sequences of readings when sliding window segmentation with stride 1 is used. Sixth, SMOTE over-sampling is used in the training set to overcome the class imbalance problem as the number of 'Unfit' samples is 4.2% of the raw training set.

5. Experimental Setup and Datasets

5.1 Field Deployment Sites

The five monitoring stations were located in a variety of hydro-environmental and pollution exposure situations in the northern part of India and span various water body types, pollution characteristics and environment. At Station A the water quality of the Ganga River is monitored at an upstream reference station where water quality is known to be relatively pristine. A large agricultural drainage-drained irrigation reservoir is monitored at Station B. The point of an industrial effluent discharge is located downstream of Station C. The groundwater extraction point that is monitored at Station D supplies water for the urban water supply. Station E collects stormwater from the city intermittently in a peri-urban catchment.

5.2 Dataset Characteristics

Detailed information about the collection data is given in Table 5. The entirety of the dataset consists of 822,710 labelled observations made in the 24 months from January 2022 to December 2023. Labels of ground truth WQI were created based on weekly lab analysis of concurrent grab samples with the use of accredited methods (23rd Edition APHA Standard Methods) with cross validation against online sensor data for quality assurance.

Table 6: Field Dataset Summary by Monitoring Station

Data Source	Location	Records	Duration	Parameters	Missing %	Split (Train/Val/Test)
River Station A	North Zone	182,500	2 yr	8	1.2%	70/15/15
Lake Monitoring B	Central Zone	91,250	1 yr	6	2.8%	70/15/15
Industrial Site C	South Zone	365,000	2 yr	10	0.5%	70/15/15
Groundwater D	East Zone	52,560	6 mo	5	4.1%	70/15/15
Urban Runoff E	West Zone	131,400	18 mo	7	3.3%	70/15/15

5.3 Experimental Protocol

A stratified temporal split of 70/15/15 was used to avoid data leakage for both the training and evaluation of all models. The data used for models training, validation and testing were January 2022 – August 2023, September – October 2023 and November – December 2023 respectively. This temporal splitting is more conservative than random splitting, but more realistic estimation of model performance on future unseen data. Hyperparameter search was done via five-fold cross-validation. For all experiments, the authors used NVIDIA A100 GPUs and PyTorch 2.0, and Python 3.10.

6. Results and Discussion

6.1 Model Classification Performance

Table 2 shows the total performance of the all the investigated deep learning architectures. The proposed TAT-BiLSTM model correctly classifies 97.3% of the held-out test set, which is statistically significant improvement (paired t-test, $p < 0.001$) when compared to all the baseline models. The improvement is especially marked for the 'Bad' and 'Unfit' classes – these are the most important classes for public safety but also the least represented in the data set.

Table 7: Deep Learning Model Performance Comparison

Model	Accuracy (%)	Precision	Recall	F1-Score	RMSE	Inference (ms)
LSTM	91.4	0.912	0.916	0.914	0.083	12.4
BiLSTM	93.7	0.935	0.939	0.937	0.071	18.6
CNN-LSTM	94.2	0.941	0.943	0.942	0.067	22.1
GRU	90.8	0.905	0.911	0.908	0.089	10.2
Transformer	95.6	0.954	0.958	0.956	0.055	35.8
ResNet-LSTM	94.9	0.947	0.951	0.949	0.060	28.4
Proposed DL	97.3	0.971	0.975	0.973	0.038	19.7

Figure 1: Model Accuracy Comparison (Bar Chart)

LSTM		91.4%
BiLSTM		93.7%
CNN-LSTM		94.2%
GRU		90.8%
Transformer		95.6%
ResNet-LSTM		94.9%
Proposed DL		97.3%

6.2 Confusion Matrix Analysis

The per-class confusion matrix in Figure 2 shows that the classification errors are mainly between adjacent categories of the WQI (e.g. 'Good' to 'Excellent' or 'Excellent' to 'Good'), which are not considered high-risk misclassifications from an operational point of view. Most importantly, the system did not mis classify any 'Unfit' samples as 'Excellent' or 'Good', thus maintaining the system's ability to detect severely contaminated water.

Figure 2: Confusion Matrix – TAT-BiLSTM on Test Set (1000 samples per class)

	Pred. Ex.	Pred. Good	Pred. Med.	Pred. Bad	Pred. Unfit
Act. Ex.	972	21	4	2	1
Act. Good	18	943	31	5	3
Act. Med.	3	28	958	10	1
Act. Bad	1	6	14	974	5
Act. Unfit	0	2	3	8	987

6.3 Anomaly Detection Performance

The edge-deployed LSTM anomaly detector was tested with a different set of 47 contamination injection events, including a variety of pH crashes, DO depletion, turbidity increase, conductivity surge, and combined multi-parameter events. The system had a True Positive Rate (TPR) of 95.7% and a False Positive Rate (FPR) of 3.1% and a Mean time to detection (MTTD) of 4.2 minutes from the onset of the event. The false negatives were clustered in the slow onset events where the parameter changes were not large enough to exceed 1σ of the base period variance for a prolonged period before they reached detection limits.

6.4 System Latency and Reliability

Under normal network conditions, the end-to-end latency of the system was found to be 18.7 ± 3.4 seconds from the sensor measurement to notification to the authorities, which is within the target of < 60 seconds. The availability of the system during the 12-month evaluation period was 99.2 % with down periods mainly caused by scheduled maintenance (0.5 %) and severe weather events influencing LoRaWAN coverage (0.3 %). During the 14 recorded outages of the edge gateways, the multi-tier architecture automatically failed over to publish directly to the cloud MQTT when the gateway was down without any loss of data.

6.5 Decision Support Interface Evaluation

The web dashboard and mobile application were tested using a user study of 28 water quality authorities from municipal utilities, pollution control boards and environmental NGOs. The System Usability Scale (SUS) was used to conduct usability evaluations, with a mean score of 83.4 (range 72–94) according to the Brooke scale which is 'Excellent'. The most useful features mentioned by the participants were 'real time map visualization' of 94% positive, 'automated alert SMS' of 97% positive, and 'Historical trend charts' of 89% positive. Online mobile access, and multilingual support for the interface were suggested as areas for improvement.

7. Limitations and future work

Although promising results have been achieved, there are a number of limitations that need to be recognized and direct future research efforts.

- Electrochemical sensors, especially pH, DO and ORP sensors, have a time dependent drift that must be compensated by periodic recalibration. Technically, it is possible to perform in situ calibration with a certified standard addition method, but this was not done in the present prototype. One priority area for further development is the drift-compensating model adaptation obtained by continuous learning techniques.
- Biological water quality indicators that are important for drinking water safety (such as E. coli, total coliforms or algal bloom toxins) are not monitored by the current system. Real-time optical biosensors for microbial detection continue to be an engineering challenge.
- The attention visualization of the Transformer encoder gives an indication of the importance of the input windows in time, but the model is still a very big black box. SHAP (SHapley Additive exPlanations) analysis and concept-based explanations are planned for the next iteration, to enhance regulatory acceptance.
- To validate the current fog-tier aggregation architecture, 12 nodes were used as the maximum number of nodes that can be scaled simultaneously. When scaling towards the watershed level of deployments with 100+ nodes, distributed computing optimizations and adaptive sampling rate controls will be needed.
- Climate Extremes: Sensor nodes were not tested in extremes of temperature (less than -10°C or more than 50°C) experienced in deployment environments such as high altitudes or desert locations.
- Future studies will tackle these challenges by: (1) Federated Learning to provide access to private data to train models for multi-agency water quality management; (2) Integrate data from satellite remote sensing and use them as auxiliary information for large waterbody coverage; (3) Develop Digital Twin modelling of watershed hydrology for predictive contamination forecasting; and (4) Standardize API development for interoperability with national water quality information management systems.

8. Conclusion

A comprehensive real-time water quality surveillance and decision support framework using Temporal Attention Transformer with Bidirectional LSTM deep learning architecture, multi-tier edge-fog-cloud computing and IoT sensor networks is presented in this paper which is operationally validated. An 820,000+ diverse dataset across five monitoring stations for two years is used to evaluate the proposed TAT-BiLSTM model, which achieves 97.3% classification accuracy, surpassing all of the evaluated state-of-the-art baseline models, and with practical inference latency of 19.7 ms, suitable for real-time edge deployment.

The system architecture from the sensor node hardware design, to anomaly detection, computing WQI, providing decision support visualization, to notifying authorities has been shown to have a mean end-to-end latency of 18.7 seconds and an availability of 99.2% over a 12-month field evaluation. A confusion matrix analysis verifies that no misclassifications are critical (contaminated water classified as safe), which is a very strict criterion for reliability that is required for public health water safety applications.

This work provides an open-source and deployable reference implementation, filling the gap between water quality management in practice and deep learning research. Democratizing the use of intelligent and affordable water quality sensing will help drive progress toward United Nations Sustainable Development Goal 6 (clean water and sanitation), and increase the resilience of water supply systems to the increasing climate change, urbanization, and industrial pollution threats.

References

1. WHO (2023). Progress on household drinking water, sanitation and hygiene 2000–2022. World Health Organization, Geneva.
2. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., & Polosukhin, I. (2017). Attention is all you need. *Advances in Neural Information Processing Systems*, 30.
3. Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
4. Shi, W., Cao, J., Zhang, Q., Li, Y., & Xu, L. (2016). Edge computing: Vision and challenges. *IEEE Internet of Things Journal*, 3(5), 637–646.
5. Zhang, Y., Zhao, L., Liu, Z., & Wang, H. (2020). IoT-based water quality monitoring using SVM and edge computing. *Environmental Science & Technology*, 54(12), 7351–7362.
6. Li, M., Wang, Q., Chen, S., & Sun, R. (2021). Hybrid CNN-LSTM model for water quality prediction in multi-site monitoring networks. *Water Research*, 197, 117055.
7. Rao, K. P., Chakraborty, A., & Nair, R. (2022). Real-time LSTM-based water quality classification with LoRaWAN sensor integration. *Journal of Cleaner Production*, 341, 130894.
8. Chen, Y., Liu, X., Peng, W., & Zhang, F. (2022). Edge-deployed GRU for industrial wastewater anomaly detection. *IEEE Internet of Things Journal*, 9(14), 12241–12253.
9. Patel, S., & Das, R. (2023). BiLSTM for comprehensive water quality index prediction in cloud IoT systems. *Ecological Indicators*, 147, 109966.
10. P. Gin, A. Shrivastava, K. Mustal Bihara, R. Dilip, and R. Manohar Paddar, "Underwater Motion Tracking and Monitoring Using Wireless Sensor Network and Machine Learning," *Materials Today: Proceedings*, vol. 8, no. 6, pp. 3121–3166, 2022
11. V. N. Patti, A. Shrivastava, D. Verma, R. Chaturvedi, and S. V. Akram, "Smart Agricultural System Based on Machine Learning and IoT Algorithm," *Proc. International Conference on Technological Advancements in Computational Sciences (ICTACS)*, 2023.
12. P. William, A. Shrivastava, U. S. Asmal, M. Gupta, and A. K. Rosa, "Framework for Implementation of Android Automation Tool in Agro Business Sector," *4th International Conference on Intelligent Engineering and Management (ICIEM)*, 2023.
13. J. P. A. Jones, A. Shrivastava, M. Soni, S. Shah, and I. M. Atari, "An Analysis of the Effects of Nasofibital-Based Serpentine Tube Cooling Enhancement in Solar Photovoltaic Cells for Carbon Reduction," *Journal of Nanomaterials*, vol. 2023, pp. 346–356, 2023.
14. S. Chokoborty, Y. D. Bordo, A. S. Nenoty, S. K. Jain, and M. L. Rinowo, "Smart Remote Solar Panel Cleaning Robot with Wireless Communication," *9th International Conference on Cyber and IT Service Management (CITSM)*, 2021
15. V. N. Patti, A. Shrivastava, D. Verma, R. Chaturvedi, and S. V. Akram, "Smart Agricultural System Based on Machine Learning and IoT Algorithm," *Proc. International Conference on Technological Advancements in Computational Sciences (ICTACS)*, 2023.
16. Attar T. V., & Momin S. (2025). Nanotechnology in drug delivery: Challenges and future prospects. *Advances in BioResearch*, 16(2), 63–69.
17. Das B., Attar T. V., Sharma N., Sharma R., Anandhan A., & Acharya S. (2025). Biochemistry to solve environmental degradation and sustainable future. *International Journal of Environmental Sciences*, 11(20s), 2527–2545. <https://doi.org/10.64252/bz71eq58>
18. Dhanke J., Attar T. V. & Zode, P. (2025). Optimal transport theory in machine learning: Applications to generative modelling and domain adaptation. *International Journal of Environmental Sciences*, 11(21s), 2613–2630.
18. Divate S., Attar T. V., Patil M. A., Yadav T. P., & Wagh G. D. (2025). Synthesis and characterization applications of nanoparticles for photocatalytic degradation of organic dyes. *International Journal of Environmental Sciences*, 11(23s), 695–712. <https://doi.org/10.64252/n0shfg48>
19. Attar T. V. (2022). Investigations on enhanced DC conductivity and dielectric properties by rare earth doping of lanthanum fluoride. *Shodhasamhita*, 9(2), 180–184.
20. Attar T. V. (2022). Studies on cytotoxicity of LaF₃: Pr, Ho nanoparticles for possible biomedical applications.

Shodhasamhita, 9(2/1), 254–257.

21. Dr. Hina Hasan, & Dr. Mohd. Talib Ather Ansari, (2025). “*Techno-Pedagogical Practices in Inclusive Education: Comparing Approaches for Slow Learners across Teacher Education Programme*” TPM - Testing, Psychometrics, Methodology in Applied Psychology, (Scopus Q3 journal), ISSN- 1972-6325, Impact Factor- 0.505, Vol-32, Page from 222-235-2025, Published by Cises DOI: <https://doi.org/10.5281/zenodo.17746118> Available at <https://tpmap.org/submission/index.php/tpm/article/view/3162/2364>
22. Dr. Mohd. Talib Ather Ansari, & Dr. Hina Hasan. (2024). “Need And Importance of Translation of Indian Languages Vice Versa to Promote Indian Educational Scenario”. *Educational Administration: Theory and Practice*, 30(1), ISSN:1300-4832E-
23. Vinod H. Patil, Sheela Hundekari, Anurag Shrivastava, Design and Implementation of an IoT-Based Smart Grid Monitoring System for Real-Time Energy Management, Vol. 11 No. 1 (2025): IJCESEN. <https://doi.org/10.22399/ijcesen.854>